FULL LENGTH TEXT

AUGMENTED COLLECTIVE INTELLIGENCE

HUMAN-AI NETWORKS IN A VIRTUAL FUTURE OF WORK, AND HOW THEY CHANGE OUR WORLD

PRACTITIONER'S GUIDEBOOK

GIANNI GIACOMELLI

SUPERMIND.DESIGN

Version 11.7

Last edit Jan 18th, 2023

First version published: March 2020

Download the latest at www.supermind.design

Contacts:

info@supermind.design

https://www.linkedin.com/in/ggiacomelli/

Cover and chapter pages artwork credits: Bob Holzer, the Noun Project

"Problems can never be solved with the same thinking that originated them."

Albert Einstein

Contents

Acknowledgments

The following people and organization have played an important role in the creation of this document.

MIT Center for Collective Intelligence (CCI): Prof. Thomas Malone, founder of the center and Patrick J. McGovern Professor of Management at the MIT Sloan School of Management - for the many years of work generating research and fostering interest in collective intelligence; Kathleen Kennedy (executive director); Peter Gloor and Rob Laubacher (research scientists) and the rest of the CCI team including Annalyn Bachmann; as well as MIT Media Lab's David Kong (director of the community biotech initiative).

Genpact's CEO NV "Tiger" Tyagarajan and CHRO Piyush Mehta for their sponsorship of the world's first known attempt to use network intelligence to reskill tens of thousands of people and make them a more adaptive workforce. Microsoft Asia's president Ahmed Mazhari for his practical ideas on applying collective intelligence. Others provided input and feedback throughout. Among them: Peter Temes at the Institute for Innovation in Large Organizations (ILO) for insightful research on large enterprise adoption of predictive markets; Eric Alexandre and his thoughts on ecosystems; April Strategy's Flavio Poli; Genpact's Jay Scanlan for his insightful critique of parts of this document, and Sanjay Srivastava for his practical work on AI for the flow of work; Nesta's Aleks Berditchevskaia for the focus on sharing knowledge about collective intelligence; Tejasvita Saran for her help in the application of social media to network analysis; Azeem Azhar's Exponential View community for input and feedback, as well as Azeem's book "Exponential" for a well-rounded perspective on "techno-social" matters; University of Florence's PhD Enrico Imbimbo for his curation of relevant materials; Jacqueline Minor for her curation of network literature; Jason Compton for editing part of the manuscript.

About the author

This report benefits from my exposure to both real-life management of large complex organizations, and academia's unconstrained thinking. I am an innovation professional focused on large organizations, who's spent a career harnessing the power of networks to generate and implement new ideas.

I work as an affiliate of **Massachusetts Institute of Technology's** (MIT) **Collective Intelligence Center's Design Lab** - aimed at exploring and designing people-machine ensembles able to generate radical innovation.

As Chief of Innovation and part of the executive management team, I oversee the innovation efforts of New York City headquartered GE-spinoff **Genpact** (NYSE: "G"), a 100,000 people professional services company involved in the digital transformation of the operations of large, complex enterprises. My focus is to create the environment (including the systems, the processes, and the capabilities) for innovation to flourish across the enterprise.

I am also part of advisory boards of organizations focused on social impact that leverage the power of collective intelligence.

For more, look up the LinkedIn profile.

And most importantly, if you feel this journey is yours too, get in touch.

Prologue

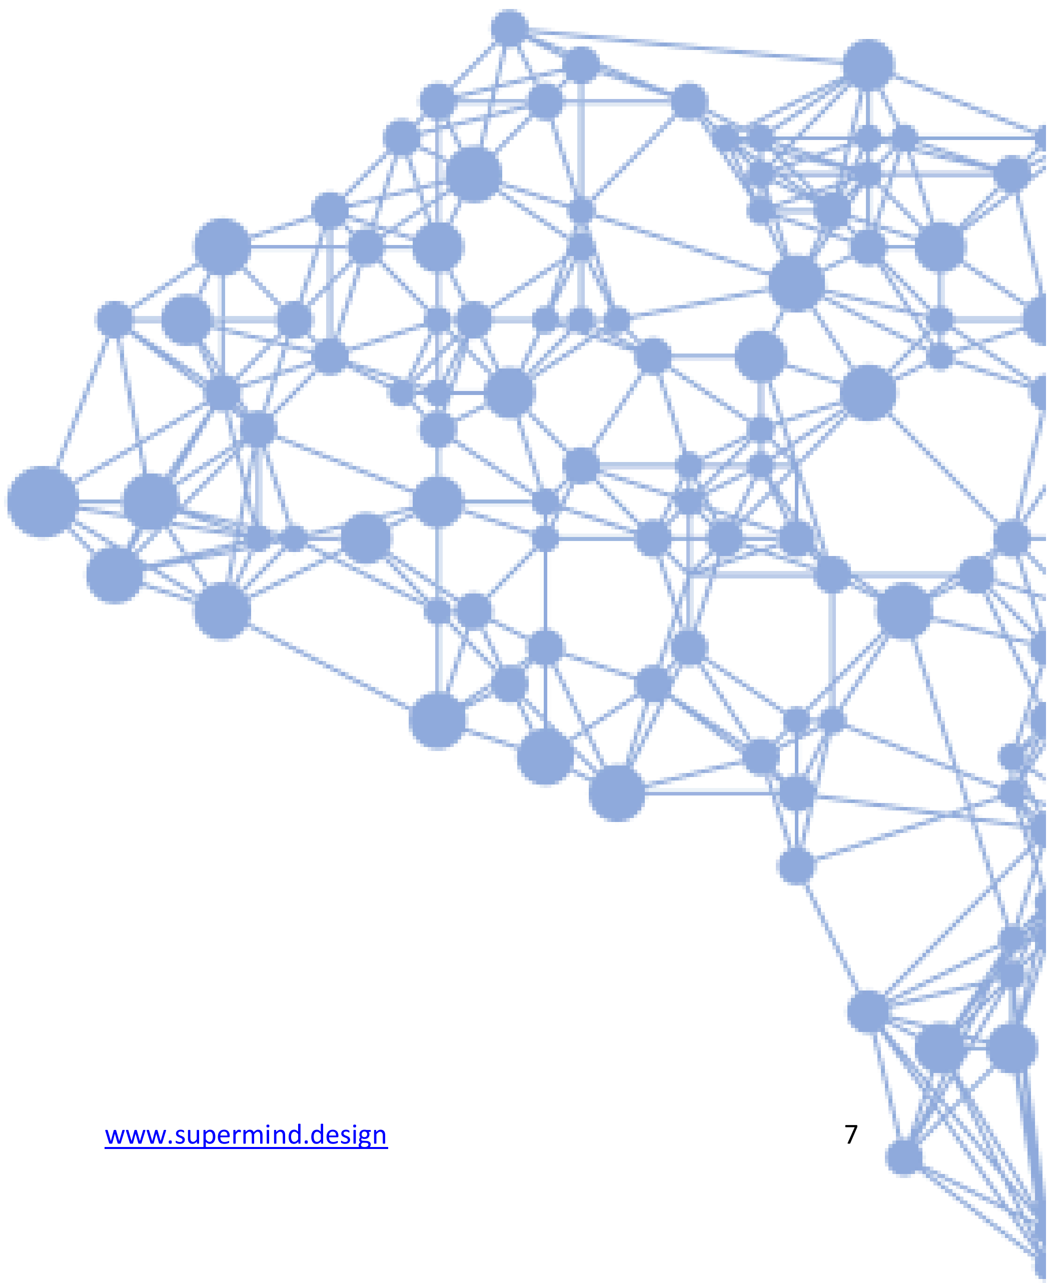

When SARS erupted in 2003, it ripped through a world whose intelligence was far less connected, but whose physical flows of people and goods were comparable to today. The collective response in the first days of COVID-19 tells the story of how much has changed. Within seven days of the official announcement in early January, the virus' gene was sequenced and some of its key traits understood – also by leveraging a "data Commons", an open-access repository of genome sequences run by the American National Institute of Health (GenBank)[1]. By January 13th a German university hospital together with experts in Holland, UK, and Hong Kong, had generated the first reliable test. Within five weeks after the official emergence of the virus, over 500 papers had been written on it, covering topics including vaccines, therapy, diagnosis, and clinical treatment. All of this was an order of magnitude faster than what the world had managed less than twenty years earlier. AI played a role too, as several projects used algorithms to identify relevant data points in patient outcomes and granularly model the virus' propagation. Perhaps most strikingly, AI-powered engines sifted through natural language scientific literature and identified existing drugs whose properties could help treat the virus[2].

While this is heartening and displays how the world's collective intelligence was harnessed, the speed of response could arguably have been even faster. Li Wenliang, a Wuhan ophthalmologist, was alarmed by some unusual symptoms in his patients. He used Weibo's social media platform in late December 2019 to send a warning to his colleagues. Local authorities prevented that signal from spreading and being verified by a broader group of competent parties. Had his message been heard more widely, the world could have gained precious days and mounted a global response before the virus spread exponentially. Instead, China and much of the world were forced into a mass-scale mobilization. By January 23 Wuhan, a city of 11 million people, was quarantined. Anyone who had been in contact with a person at risk was geofenced[3], and an app notified people close to that person in real time. By March, such draconian measures had largely contained the first wave of the pandemic.

The initial reaction of the world's collective intelligence (Wuhan's local authorities, worried about a potential wave of social media hysteria) was hierarchical. Hierarchical responses are fast and decisive, but often cannot sense and confirm systemic patterns quickly, and lack diversity of thinking. Communities (such as specialized social media groups, paramedics, doctors) and ecosystems (governmental bodies, pharma and technology companies, and crisis managers) would have collectively done a better job, if the quality of their output, and their impact on public opinion, could have been controlled (instead, the spread of disinformation created a distraction and made the democratic reaction of citizen less appropriate, especially in groups with lower social capital and trust.)

Communities could have also been harnessed more systematically to crowdsource such data as social interactions or pre-existing personal conditions that led to highly severe infection to discover behaviors leading to contagion. Consider the example of Taiwan, which countered the spread faster than most other countries among other things because of its embrace of social media as a collective intelligence system. There, citizens were able to inform government decisions, and misinformation was quashed by the involvement of citizens – together with very transparent central authorities and social media groups that tagged misinformation early.

Contagion is the spread of a virus in a *network* of people. Curbing it requires isolating the network spreaders. To do so, South Korea implemented network *sensors* (pervasive testing), and by tracing back the networks of the infected was able to contain the spread. Compare that with Italy, which – thanks to the signal from Asia - had a head start in sensing the threat. Italy started testing early, though not as pervasively as South Korea. But the actual *response* of both authorities and the public was too slow and uncoordinated to prevent the virus from spreading in cross-generational households and in the dense social fabric of Italian society. Neither security rules (prohibiting gatherings and tracking spreaders) nor social norms changed fast enough. The collective response was a democratic one without executive powers to coerce the collective. Each person chose how to behave for too long, and then it was too late to stop the spread. After a few weeks, Italy's fatalities overtook China's, a country whose population is fifty times larger. It is worth mentioning that Italy's lesson didn't help the largely hierarchical decision-making in the US, UK, and Brazil, which ended up witnessing devastating consequences.

Just possibly, national governments could have coordinated their responses better, and shared more data from their propagation networks. Enterprises caught in a sudden change of pace, could have avoided slow reactions and costly mistakes.
Even scientific contributors, typically open to international collaboration and ultimately the single best performer in our societies' reaction, narrowed their network in the interest of speed and included fewer international contributors[4]. This trend may now be worryingly propagating

beyond Covid research, due to, at least partly, US-China geopolitical frictions[5] and, possibly, because of scientists' networks adapting to remote work and becoming less reliant on the encounters through conferences, for instance.

This is not to belittle the progress made - researchers widely shared genomics and other data related Covid-19[6]. And when additional variants of the virus started to emerge, genomic surveillance networks across the world proved invaluable in detecting the mutation, and helping clinical trials to adjust– for instance, the Africa Pathogen Genomics Initiative, first established to track Ebola and yellow fever. Public-private partnerships that leverage AI-augmented collective intelligence formed to connect available resources and identify gaps[7]. And the fabric of knowledge and more specifically *office* work changed almost overnight. Millions were forced to work from home and use technologies that had existed for some time but that were still unfamiliar to many. The related digital infrastructure (collaboration and other remote work services and tools) was remarkably resilient as entire companies had to virtualize their knowledge work in a matter of days.

Many struggled not just with the technologies, but more broadly with the new ways of working and their new organizational design. While some rushed to reduce their expensive real estate footprint, the forced lockdown reminded CEOs of why we spend a fortune on office space. The strong (part of work routine) and weak (periodic, based on mutual affinity or chance) ties[8] in our people networks make knowledge-work happen - both business-as-usual and innovation work - and shape culture deeply. The office space literally is a technology to "design the human network" - the physical layout shapes the network structure, with its space between desks, offices, buildings. The issue with remote work is that if left unattended, weak ties dwindle (no more serendipitous encounters), and the effectiveness of strong ties may suffer (less body language, no impromptu whiteboards, sticky notes, etc.). Virtual space is more of a vacuum, making casual connections a daunting free-for-all.

But that's not destiny - we have much technology and human-centered methods to harness those network flows in a virtual world - with the added benefit that the world's brains are now our neighbor. This report provides a detailed view of those practices. Indeed, some companies weathered the storm comparatively better. The more successful were already organized through Agile methods as a team of teams[9] where myriad units aligned by shared goals acted semiautonomously, were familiar with collaboration platforms, and could maintain continuous and granular awareness of the external environment, weathered the storm comparatively better. The healthcare system itself made immediately made more expansive use of telemedicine, especially for triaging patients based on severity, and supporting people remotely. This was all made possible by leveraging distributed workforces that would otherwise not be used. Some companies introduced stronger collaboration technologies, including intelligent bots that automatically match colleagues to strengthen weak ties.

As this document will explain, the world is now hyperconnected, making available to all of us an unprecedented amount of collective intelligence, driven by both computing and people's smarts. But most organizations can't harness it. They're saddled with a limited understanding of the principles of collective intelligence, and now laden with a now deeply entrenched suspicion of social media – which, in many respects, throws the proverbial baby with the bathwater.

The implications are profound. This document was also written because, today, we have levers to craft much more resilient organizations that harness that intelligence, and they should be applied. To understand the potential created by the last twenty years of digital evolution, let's look at another story.

***

In 1998, about two years before the carnage of the dot-com bust pulverized the credibility of over-optimistic "new economy" claims, two Stanford Ph.D. students presented an interesting scientific paper[10]. It contained a big idea that made their company one of the most valuable in the world. Their names are Sergey Brin and Lawrence Page. Their company started as one of many search engines, and quickly gained one of the most lucrative market positions in history. The world has since been entranced with their superior digital prowess, particularly their adoption of Artificial Intelligence (AI). At Google, AI became the mound that protected their market position and valuation. Senior leaders and consultants have studied and attempted to emulate those strategies and practices – with, to date, comparatively scant results. There's another part of the story lost in the mainstream AI hype. That story omits details that can help the rest of us harness that power, especially those whose business isn't built on selling online advertising. To find that detail, let's go back to 1998.

Google didn't invent the search engine. But it did achieve two things that changed the world. First, its queries became more relevant than competitors through a new way of ranking results pages. Google did so by measuring pages' *centrality*[11] in a broad *network* of websites, the web pages created by thousands (and later millions) of people that reference each other through hyperlinks. Second, Google was quick to embrace and transform a new generation of advertising, which used the superior relevance of those queries to target users more accurately and as a result more valuably. The combination of the two achievements was one of the most powerful business model innovations of the twentieth century. What few realize is that Google's platform power comes from masterfully using ("organizing", in Google's parlance) the fruits of the intelligence of billions of people. Google's strengths come from the knowledge we create, and the choices we make when we browse. Google's business model wouldn't and couldn't exist without harnessing the intelligence of knowledge producers, curators, and users. Google does this in ways unthinkable just twenty years ago, and it's still growing (today, almost 2 billion websites exist). It also contributes to that intelligence by enabling the world to easily retrieve knowledge in a collective, comparatively-frictionless *remembering*.

Twitter, Facebook, and LinkedIn exploited a similar opening, as they gave individuals a forum to share pent-up creativity and thoughts. But they did something else that changed the world forever: they explicitly introduced the *people* side of the network. Think about it: the act of following a single person on a social network is mundane, even trivial, but for AI-powered machines, seeing those connections in large volume is a boon. Machines don't predict those connections – they use people's judgment to trace those relationships. They then use that information to curate what the users see – in this case, predicting what you want to see next, based on what your network seems to like. The result is a sticky user experience. That leads to more clicks and superior advertising performance.

The explosion of other social media, from blogging platforms like WordPress to podcasts to Snapchat, fueled the creative fires in millions of people – long before TikTok upped the ante by showing how artificial-intelligence-driven curation can turn young people into truly sticky (some say addicted) customers and arguably opened other forms of expression, not all of them positive. While much of that creation is of dubious intellectual value and can even be downright dangerous, interesting ideas often bloom in these environments – and through platforms such as

Substack or blockchain-enabled non-fungible-tokens (NFT), the number of paid creators is growing exponentially.

Beyond the creation of original content, these technologies are increasingly able to collectively *sense* the environment. From powering the Arab Spring to shedding light on breaking news and crises, a massively decentralized network of sensors (the majority being simply people with smartphones) has emerged. From this ad hoc network flows ever-fresh information.

Network effects underlying collective intelligence have changed the world's competitive dynamics. In late 2021, seven of the ten most valuable companies (Apple, Microsoft, Alphabet / Google, Amazon, Facebook, Tesla, and Tencent) relied on them. It is worth understanding how those networks are built.

Despite the hype, these companies' strengths don't just come from strong AI but from their use of the cognitive power of billions of people who generate information and interact with one another in very specific ways. They're listening to and acting on a sea of information made up of trillions of microevents – all made by human choices - every day. To an individual, those choices aren't worth much. To an advertising machine, they're hugely important, to the tune of trillions of dollars of market value.

Today, we can use the *networks of intelligence* that surround us – both within and outside of our organizations. We can identify the right people (and knowledge-producing machines) to engage and support them. We can measure the energy of those networks and intervene to motivate them at the right time. And while advertising currently skews the access to knowledge that machines give networks of people, new developments in text and language processing enabled by neural network GPT-3 models show remarkable possibilities.

And we can do all of this at massive scale. With technical tools and AI, people and machines can combine relevant knowledge, curate information, collaborate, and do the computational heavy lifting as needed.

Social commerce and supply chains have existed for some time. Everyone knows Uber and how, despite its lowlights, it spurred innovation in the shared-mobility space, but perhaps not everyone knows the now-huge Pinduoduo, the Chinese Consumer-to-Manufacturer infrastructure where hundreds of millions of buyers pool their direct purchases of agricultural products and help streamline the related supply chain.

In an interesting turn of events, Google's use of collective intelligence to power knowledge search was disrupted in 2023 by the introduction of large language models as an interface with knowledge databases. Companies like OpenAI and Perplexity showed that interrogating the world's knowledge can benefit from a conversational interface – not a list of links. The next step is knowledge graphs – databases whose entities (content, people) are connected to each other by relations (edges of the graph) – resulting in much sharper ability to identify causal and distant relationships, for instance. This chapter is being built in real time.

In 1998, most advertisers saw web pages as billboards on a screen. A select few saw deeper, into the network behind them, and recognized the web's ability to upend how we promote things. Those visionaries spawned incredibly valuable businesses powered by collective intelligence. Can yours be next?

Purpose and structure of this document

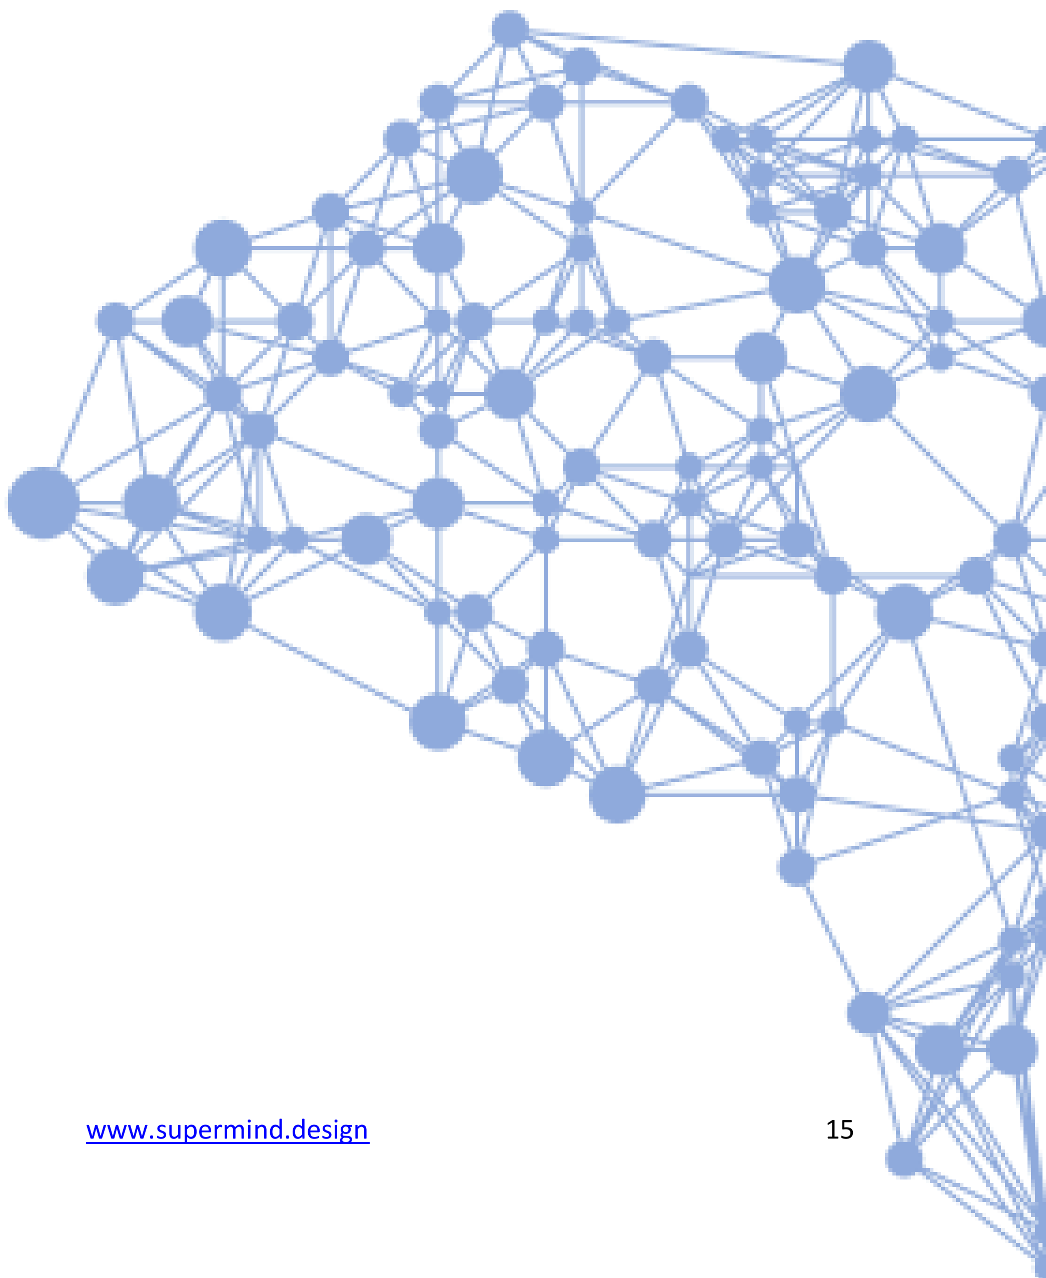

Intelligent networks, made of large numbers of people *and* machines, are a new organizational design ready for widespread adoption. They can help leaders from CEOs to middle managers to movement organizers, harness the full collective cognitive power of their organizations and ecosystems, generate and implement stronger ideas, and adapt more swiftly and effectively to fast-changing conditions.

This compendium provides principles, frameworks, detailed examples and ultimately a playbook for deliberately managing large-scale network activities that today's organizations typically experience in an organic, uncontrolled and poorly understood manner. No matter how good your top management is, they likely have a limited understanding of how problem-solving work gets done – especially when it falls outside of formalized flows, such as product innovation or R&D. Top management cannot see *networks in action*, so they cannot recognize that ideas are created and implemented through large networks of people helped by machines. And without a radical improvement to innovation methods, people who innovate every day, including academia, will continue to struggle with keeping the pace of the "burden of knowledge" – resulting in longer times and more resources needed to develop breakthroughs[12].

The upshot is that at parity of current resources, you can get more from your organization by facilitating the emergence of stronger collective intelligence. Regardless of how smart the individual people and machines are, there is an additional quantum of intelligence generated by the effective *interplay* between them at scale. The network itself, the *connection* of the nodes, generates that additional intelligence. The whole can be made smarter than the sum of the individual smarts.

Anyone who has experienced the frustrations of being part of a large organization knows this doesn't happen naturally. Without a lot of help, the coordination of knowledge flows and collaboration frequently suffer from diseconomies of scale.

To unshackle that power, our organizational and managerial practices must evolve. Our org charts, process maps and people-incentive schemes struggle with today's fluid knowledge flows. And instead of embracing the intelligence of our people-machine networks, our advanced technology is often deployed to reinforce the traditional thinking calcified in existing operating models and processes. Much can be learned at the intersection of four disciplines: management-, computer-, social- and neuro-sciences. Together, they can help us *think in systems*, complementing the traditional solution design that focuses on individual technologies, people, and processes.

The concept of highly networked, collectively intelligent systems is a radical departure from older organizational architectures. Like the "reptile brain" of the ancient dinosaurs, the old ways are very effective in dealing with known conditions but slow to adapt to new circumstances. Adaptation is a trait of the human cortex – but organizational design often fails to address that opportunity, and routinely misses out on the opportunity provided by AI and other digital technology.

Academia, such as Massachusetts Institute of Technology's (MIT) Center for Collective Intelligence[13], has studied the intelligent systems space for years – as have many large internet companies that derive their power from internet data, especially social media. But these benefits shouldn't be limited to researchers and FAANG[14] executives. These pages structure what we know about intelligent systems so more managers can use them. To do so, this document is divided into few sections. In the **overview**, we lay out the key themes so that you can understand the entire concept in about thirty minutes. The overview will help you decide what you're most interested in, and which parts of your organization can benefit.

In the first part, we explain what collective intelligence is, why it matters, and the obstacles to harnessing it in today's organizations. We argue that the current discussion about AI and other advanced technologies largely misses an important point: that the locus of intelligence isn't in the individual, but rather in the emerging collective intelligence that stems from the interoperation of large networks composed of both human and machine nodes. Collectively, Ai-augmented systems can perform the main activities required to **generate intelligence**: sense, remember, create, decide, act, and learn. We introduce a framework for intervention that amplifies the collective intelligence trapped in networks by appropriately connecting the network nodes to create something similar to what MIT's Thomas Malone calls a "supermind"[15]. We discuss what AI can really bring to that organizational vision – *AI's four C's*: Connection, Curation, Collaboration, and all other heavy Computation lift. We then describe the four components of a networked-intelligence supermind: the identification, analysis, classification, and enablement of mutual visibility for the network **nodes; t**he design and implementation of effective **incentives** for their collaboration and their alignment with stated or implicit goals; the enhancement of relevant and ubiquitous **information feeds** that give the network a broader outlook and prevents it from becoming insular; and the creation and management of an appropriate **collaboration platform**.

In the second part, we delve into the individual modules of the framework, and provide additional detail of these concepts. We also discuss in more detail the application of collective intelligence. We describe practical case studies that illustrate real-world application (for instance, massive enterprise-capability building, particularly talent reskilling, through a training infrastructure powered by collective network intelligence), and then in four separate "technical notes" sections we share the detail to help practical application. We also provide a short methodology playbook for design projects, including phases and activities that help teams deliver the blueprint of a networked intelligence.

In the third part, we conclude with the discussion of future-oriented and increasingly important topics. Among others, we discuss the promise of an AI that eliminates the hurdles of foreign language, synthesizes planet-size amounts of human-generated content, and combines it with the machine-generated one to offer fresh input and insight. Clearly, much can go wrong: like all radically innovative powers, these methods can be misused - accidentally or deliberately. We reflect on the importance of ethics and regulation on algorithms that power networks strong enough to hijack society. We conclude that "the genie is already out of the bottle", hence the need to rapidly and proactively intervene, instead of trying to "dumb down" the network's intelligence. And finally, we broach the subject of the need for new leadership skills and in general management practices, to appropriately embrace the power of networked, collective-intelligence superminds.

While we introduce different frameworks, they can be combined to solve practical problems, as shown, for an enterprise environment, in the next chart. That overall framework is applicable across companies and other types of organizations, as well as other types of human-machine ecosystems.

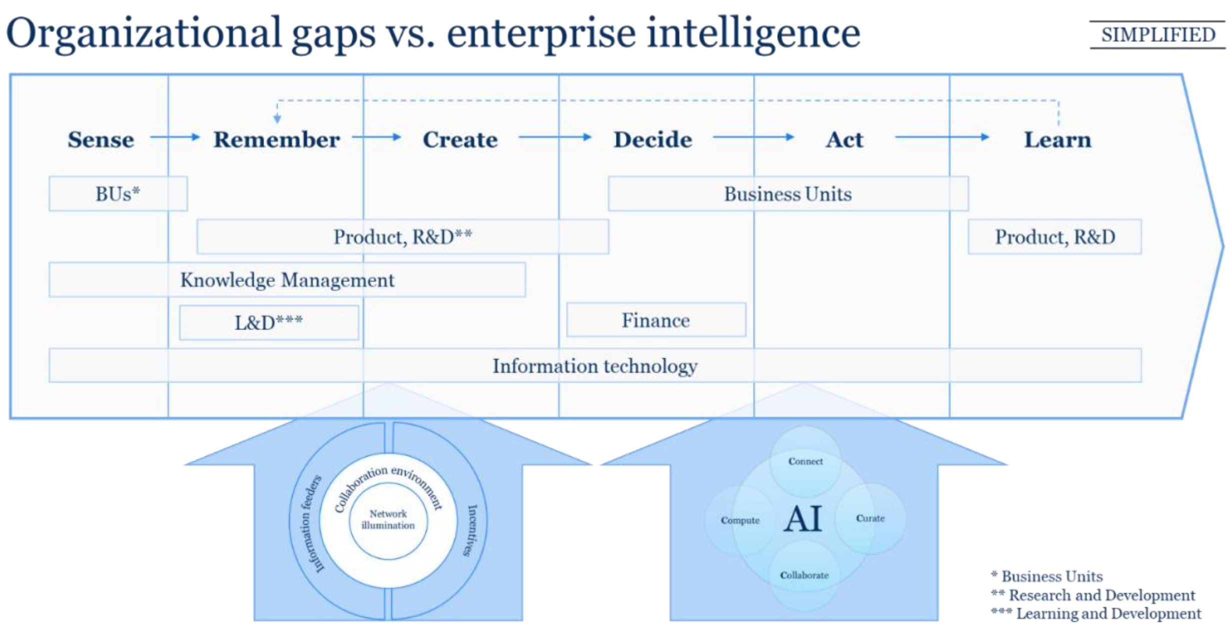

This document is written for two audiences: senior leaders in organizations who want to *understand* how collective, connected intelligence can boost their organizations' ability to innovate, solve hard problems and adapt; and practitioners (consultants, solution architects, developers) who want to *build* those networked systems. This is as an applied-strategy paper, with an appropriate amount of "how-to" guidance - not as an academic or visionary document. Its structure allows for quick browsing. It also provides enough detail for those who need a practical playbook to inform implementation, but don't want to comb through a decade of academic research and management literature.

This is an ambitious effort and this document, an incomplete but hopefully useful compendium of what exists across multiple disciplines, that will only scratch the surface of this space. However, even this initial consolidation and filtering of knowledge that can support application, is likely helpful to leaders.

Once you and your teams have absorbed it, you will hopefully have a better grasp on how to build a networked, connected intelligence - a supermind - of your own.

One more thing. This is an evolving space. This document will become obsolete in parts very quickly – that's "a feature, not a bug". All feedback will be invaluable to keep it a common, shared resource for leaders who want to harness the combined intelligence of people and machines.

Overview

This section provides a birds-eye view of the contents of the entire document. It summarizes why augmented collective intelligence matters, and what steps organizations can take to harness it.

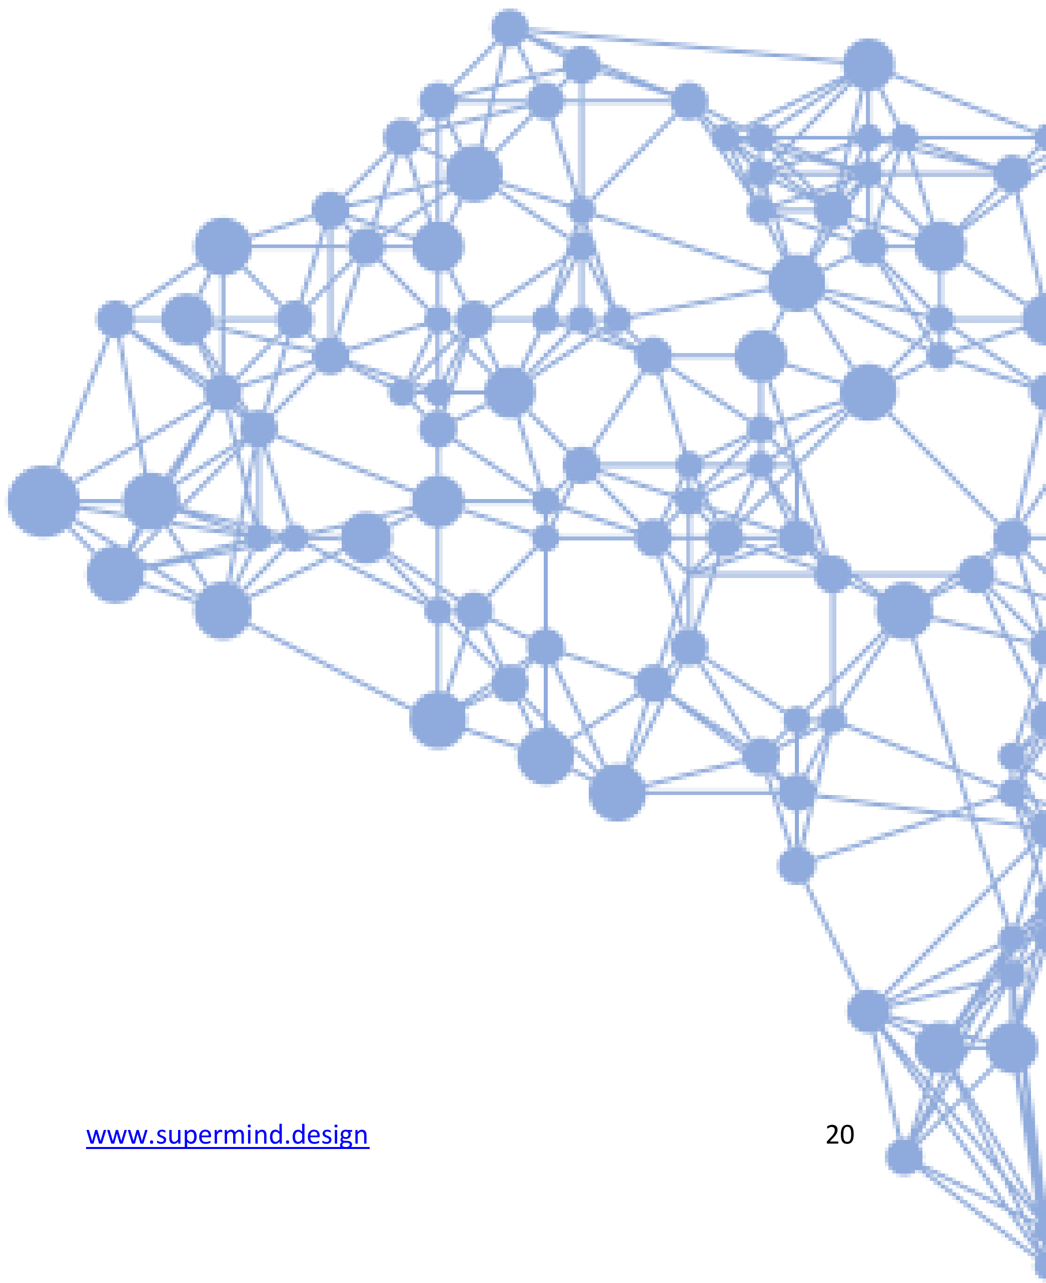

"Human vs. machine": a distracting, wrong perspective

Despite of the hype, the practical, immediate challenge for most of us isn't if machines are better than people. It is that irrespective of how much artificial intelligence (AI) we have today, and how smart our people are, most of us are dissatisfied with how fast we generate and implement great ideas. We struggle with how smart our *organizations* are, how nimble and adaptive they manage to be.

Outside of a small circle of internet-native companies, the average enterprise intelligence arguably hasn't changed much in the last years. Generalized artificial intelligence might one day be able to do complex thinking for us and relieve us from that duty, but it won't do that soon enough to help with the next many years of most nonroutine activities. Specialized AI for the foreseeable future will continue to be about machines that do complicated but narrow jobs – from prediction[16] to automation. We will be able to leverage AI like we use cranes or bulldozers. Which is hugely helpful, but most design, architecture, and building of solutions to complex things will stay squarely with people for some more time. This means, AI won't make our organizations smart by itself.

Machines are excellent – superhuman and world-changing - when the world has been somewhat explained so it makes sense to them. That happens by either feeding them data or placing them in the right part of a process. Interestingly, that doesn't mean that machines can't be creative: the field of generative design, where machines explore huge numbers of combination of e.g., engineering designs once some key parameters and data have been provided by people, proves the opposite albeit so far in a narrow set of cases. Outside of highly quantitative fields, such as some in financial services (e.g., parts of trading activities) or airplane piloting (with sometime mixed results), humans still need to do a lot. Machines are invaluable at industrializing the processing of information, that in turn people use to make non-deterministic decisions. These things make organizations smart, but the human footprint on those processes is still very large, especially when it comes to adapting to change. AI is powerful but still not easy to handle, and our world (including enterprises' convoluted legacy structures), quite messy and rigid. People, ubiquitous and nimbler, are still often needed to make it work.

Yet, this debate misses a big point: "people" does not mean primarily "individuals" anymore. Ours is by now a world where the best ideas and their implementation come from the combination of many inputs, originated in many places, assembled by many hands. For instance, research indicates that the individual genius at the root of yesterday's scientific breakthrough has given way to larger and more multidisciplinary teams that combine a vaster body of knowledge[17]. People have perfected over thousands of years the ability to use language as data transmission protocol. That knowledge can be increasingly processed by intelligent machines, at organizational or if needed planetary scale, and fed back to many people for them to collaborate upon it. That's a better starting point.

Study after study[18] highlight that many **leaders, from the CEO down, struggle with their ability to harness enough of their organization's intelligence to solve complex, nonroutine challenges**. They're frequently left feeling that their organization is unable to fully capture the innovation opportunity offered by new methods and tools. That's especially true when it comes to changing poorly-documented processes, such as complex problem-solving and innovation beyond well-understood R&D and product management. Too often, despite considerable investments in people (both internal and consultants) and technology, the results feel unimaginative and slow to translate into reality. It doesn't help that despite increased company size yields economies of scale for documented, large and repeatable processes, the ability to combine people's intelligence and turn them into results seems to suffer from *diseconomies* of scale. (That handicap is conversely a competitive opportunity for smaller, nimbler players that can make the most of the ideas and creativity of the people they have.)

From R&D to product and service innovation, to digital transformation in large operations, one common theme emerges: much of the struggle revolves around having enough of the right people - employees, consultants, or partners – and ensuring that they work together well, irrespective of where they are. Consider for instance the implication of recent research from MIT and Deloitte[19], which shows that when it comes to embracing digital transformation (arguably the main driver for innovation today), companies that do that well are distinctive in the organization's ability to collaborate and act in a *distributed* manner. These yields, in turn, an unusual ability to be nimble and bold, able to explore and execute - at the same time.

The experiments large organizations naturally generate (insights, prototypes, proofs of concept, pilots, etc.), if scaled fast when successful, can enable enterprise to adapt faster to new conditions. Even more powerfully, they might enable to *create* those new conditions, in ways that competition finds hard to emulate. The large web-native companies today are the clearest example of enterprises whose operating models and culture can generate unprecedented growth and command price/earnings multiples that implicitly assume continued dominance. What can the rest of us do?

Conventional management practices aren't enough

To harness this power for our organizations, we must look beyond current management methods. Consider the following four problems:

Org charts do a decent job at showing decision processes in a world where hierarchies are strong, and most decisions and communications happen that way. But they represent poorly how an "enterprise brain" works on complex challenges and say very little about the dynamics of its nodes and connections - its "social physics."[20] **Documented business processes** present a critical but partial view of enterprise cognition, as they typically only account for the routine, transactional part of the work, and struggle to account for evolution. A lot of "action" doesn't fit within a flow chart. **Machines are often siloed** in rigid workflows or confined to isolated groups of computer scientists or data scientists, whose connective tissue into the company's fabric is still nascent. **Leadership practices** including the prevalent ways of working (think knowledge or project management, or collaboration techniques) haven't kept pace with the ability to harness large, complex, dynamic networks instead of smaller and slow-changing static ones. Consequently, the practices of most managers continue to focus on individual leaders' work, without adequately driving the behavior of the respective networks. All these management practices haven't kept pace with the quantum leap in **collaboration** potential between large groups of people, and machines – including machines helping people to collaborate. A collaboration whose fluidity and scale indeed resembles how organic neural structures work, rather than well-established organizational designs. Our leaders don't have yet the tools to harness the intelligence that emerges from the interplay between people and machines, at scale. Consider the figure below: what happens when we *add* to the traditional management methods (in figure 1, 2, and 3) the ability to *deliberately orchestrate* networks that span across organizational boundaries (in pane 4)?

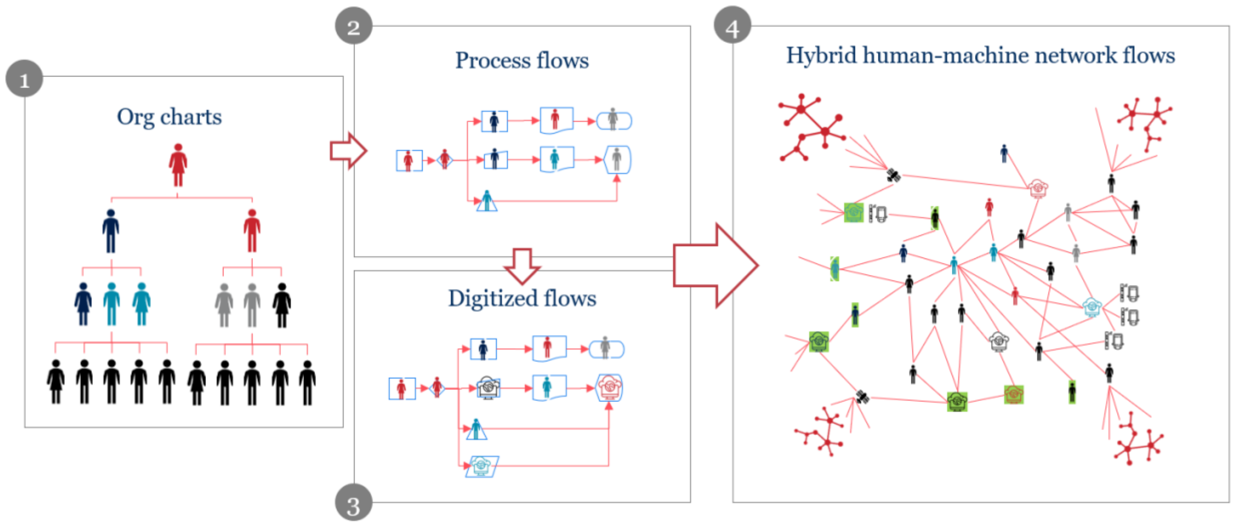

This doesn't mean that org charts, process flows, and their reimagined digital version aren't useful – especially in repetitive processes and in steady conditions. Like the inner parts of our brain, derived from our vertebrate ancestors, those enable organizations to react fast and accurately, instinctively, to known patterns. But true intelligence developed in mammals thanks to the cortex, the outer part of our brain, particularly its prefrontal parts that coalesce input from multiple places to make sense of new conditions – hence enabling faster adaptation. What is the equivalent in organizations? I argue that pane (4) above is a useful analogy, and indeed some new forms of organizational design such as Agile as well as others described in the next picture have started addressing that opportunity.

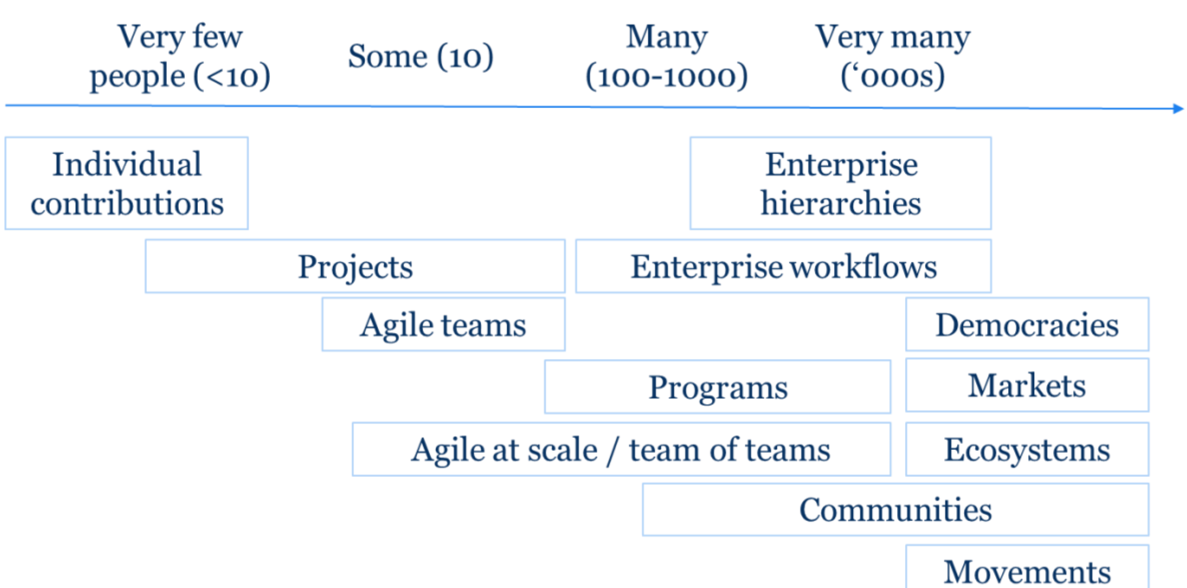

To make useful progress, we must also challenge the current, somewhat uncritical point of view on the impact of AI, which might have sidetracked many of us for a few years. Think of an analogy. The age of space exploration that started in the ‘60s didn’t give us flying cars and moon colonies: it brought us satellites that enabled the GPS, which then changed the world in profound, but also unintuitive and nonlinear, ways. A similar dynamic may be at play here: we have so far been looking at talent, and machines, largely as individuals or elements of technical and process architectures. That’s possibly missing the forest for the tree. Instead, we should think of the power of meshes of human and artificial intelligence as networks, like forest ecosystems or grid computing. Their *combined* intelligence is already a mighty force in the world, and one that we can truly harness for the good of our organizations and societies.

The rise of machine-enabled networks of people, *and* machines

What if we could use millions of machines to curate, consolidate, pre-process input and ideas, and inject them back into close loops with their human-mind counterpart? What would happen to organizational processes that don’t fit in rigid process maps? What would be emerging in the top-right corner of the picture below?

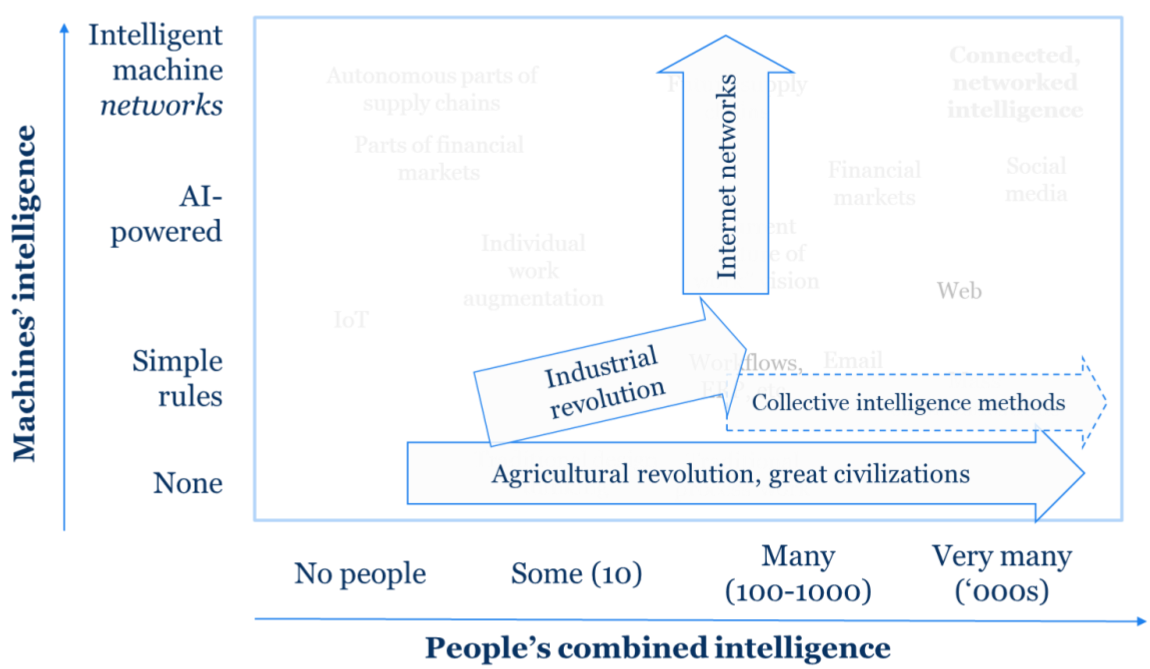

This is not just about technology. Major technological, economic and social trends have *jointly* shaped the state of where we are. Think of the rise of great civilizations, that used laws and common language, as well as currency, as "technology" to coordinate the activities of myriad people. The current situation is the *resultant of the vector* of forces illustrated below – with a late and comparatively underdeveloped entrant: modern collective intelligence methods, from contemporary organizational design to design thinking.

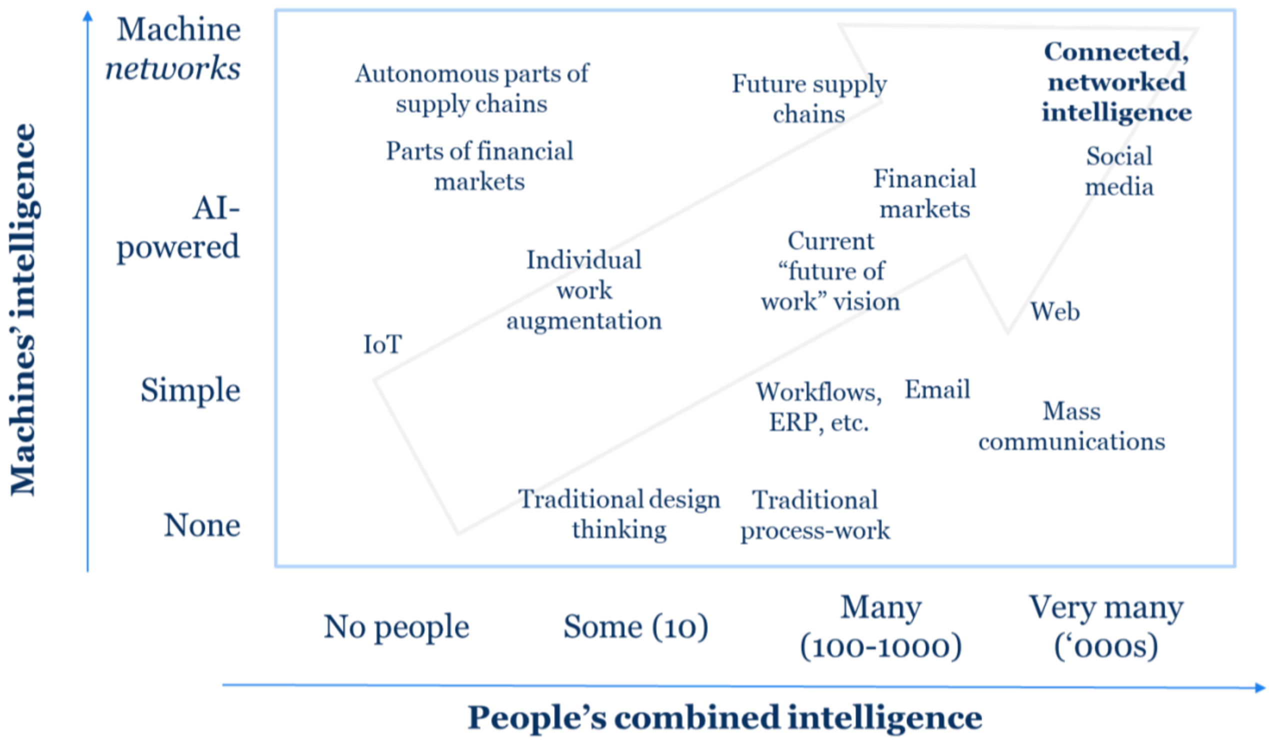

What is really new now is the ability to combine the power of tools and organizations with the power of networks.

The value of networks has been studied for some time. Metcalfe's law is the best known, stating that network's value grows exponentially as more nodes are connected – that's why telephones and email have become ubiquitous. Less well known but arguably equally important is Reed's law[21]*, that scales even faster, as more and more* groups get connected, and people connected with them. In other words, groups generated additional data on top of the individuals. That's why social media communities have grown immensely.

The building blocks, part of well-understood megatrends, exist already. When our societies and organizations moved from physical mail to telephones, the flows of communication skyrocketed, and the structure of organizations, slowly but surely, changed. Then we moved from facsimiled memos (a combination of paper mail + telephone) to email, and the structure – or at least how people interacted across org chart boxes - changed again. Email has now partly given way to truly networked communications through myriad other tools, both real-time (synchronous) and asynchronous. Search has now become central in the cognitive functioning of people – by enabling frictionless "remembering" of our collective memories.

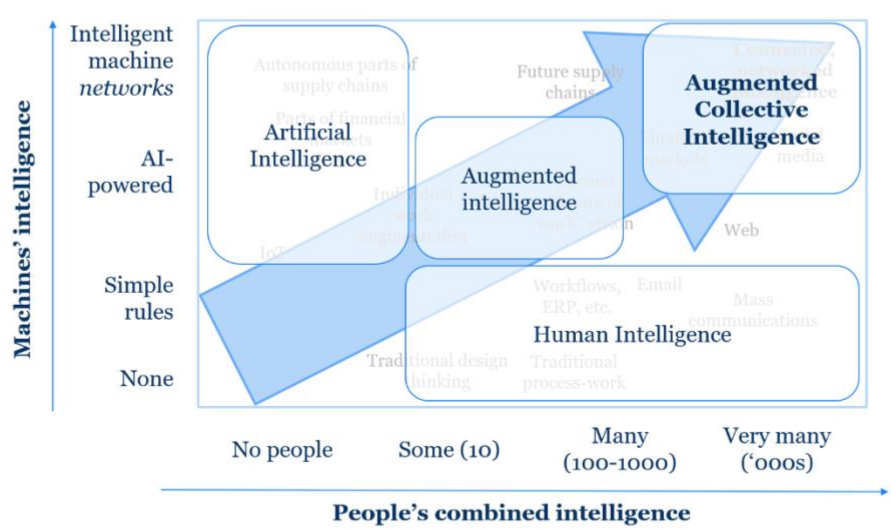

And all of those are enhanced by AI's ability to process visual inputs, natural language, find relevant content and discover connections between people. For example, when a new outbreak of SARS happened in early 2020, an AI algorithm combed through social media data, combined it with airline information data, and predicted (and communicated) the spread faster than official organizations[22]. Without social media posts as sensors, and AI's computing power, that prediction would have been impossible. As suggested in the chart opposite, the result can propel our organization beyond the current, somewhat tired, man vs. machine debate, and even beyond mere "human augmentation".

It is important to explain what dimensions underpin collective intelligence, so that one knows what to augment. Intuitive understanding doesn't do justice to that concept. Ultimately, collective intelligence rests on the power of the individual nodes, amplified by the (somewhat independent) power of the network where those nodes interact at potentially exponential scale. The chart below[23] attempts to separate the main elements, and approximate the position of intelligent entities that we are familiar with.

First, the intelligence of the individual nodes. Autonomy is the ability to explore and obtain data through some form of active inference; and Computational Power is the ability to process data. Then, network's intelligence which is partially independent of individuals', due to nodes' Interaction ability - language, theory of mind, dedicated organs (e.g., neurons, emitters/sensors of signals e.g., vocal cords, ears); and network size e.g., number of people, servers etc.

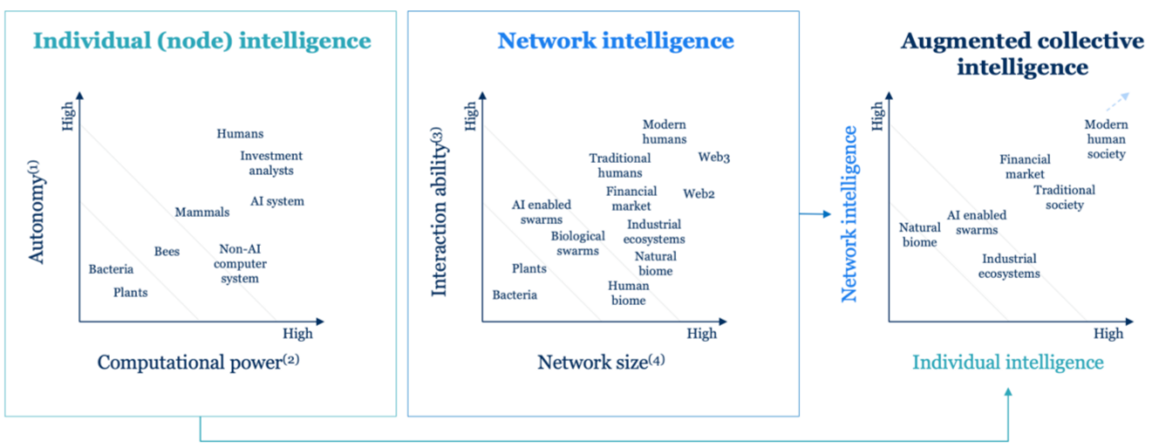

(1) Ability to explore and obtain data through some form of active inference (2) Ability to process data (3) Ability to interact e.g., language, theory of mind, dedicated neural and other organ structures (e.g., emitters and sensors of signals like noise, such as vocal cords and ears) (4) Network connected e.g., number of people, servers etc.

Now look at the picture below. Neural-like networks like those, enabled by AI, now span vast numbers of sources of knowledge, especially people but also machines. They weave them together and spread their ideas thanks to web-enabled hyper-connectivity, generation of sensor-based data (from weather to stock inventory, to citizen's warnings on Twitter), selection of content ("curation"), display of that content to relevant parties through prediction (think advertising and social-media choosing what you'll like to see), and connection (web publishing, synchronous communications, etc.). There, traditional media meets user- (or machine-) generated ideas, propagated by today's communication technologies that have AI at their core.
This is a snapshot of the fabric of our collective brain.

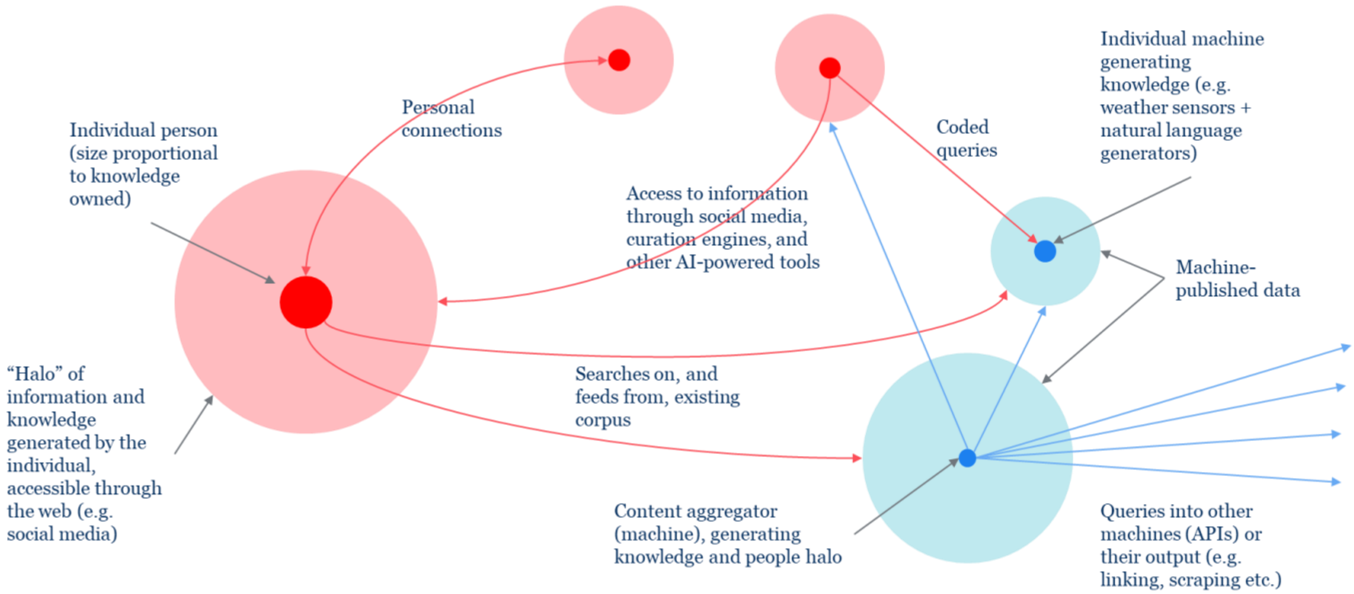

Within individual organizations, that's an explosion of knowledge available, absorbed and filtered by organizational networks now solidly wired through the likes of Outlook and Slack. That knowledge is then amplified and evolved by social networks and is simultaneously immersed in – meshed with - continuous streams of other ideas curated by AI-based algorithms. Not coincidentally, one of the fastest-growing spaces in the field of data is the one of "knowledge graphs", whose biggest advantage is to document and process relationships similar to those displayed above.

Take an example that many executives experience daily: innovation teams working on the next big thing, or simply on business-as-usual continuous improvement. Their network structure looks like something like the following – and will increasingly add "blue" nodes, that is, networked machines that complement and amplify the knowledge of people.

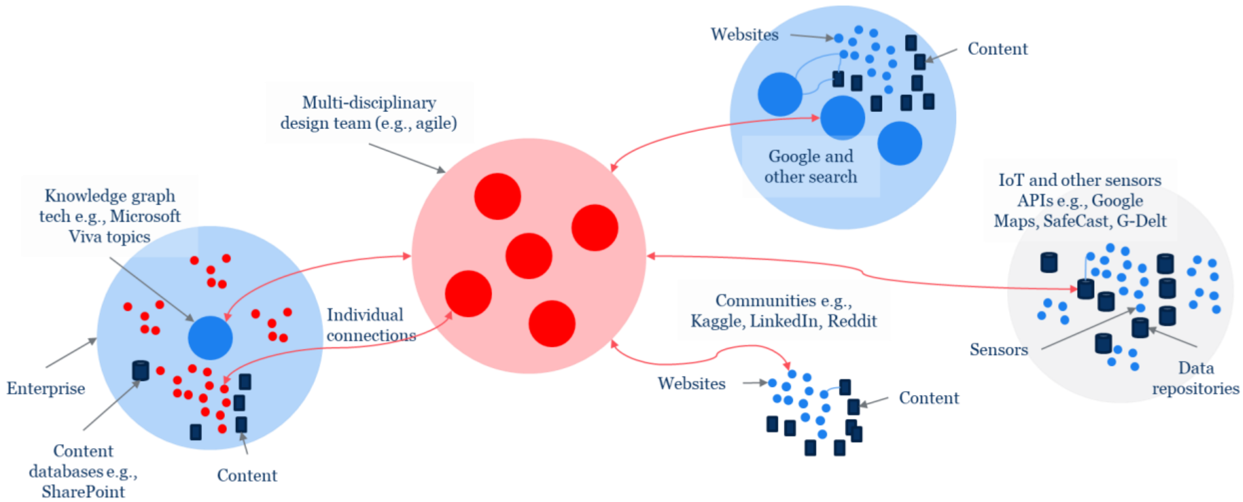

That's here, today.

What could our organizations become by fully using people-machine networks made more intelligent by AI's "four C's", i.e., its ability to exponentially improve the following four things?

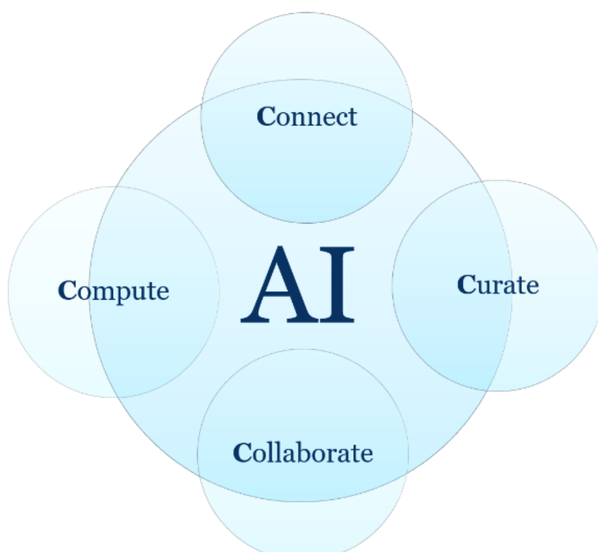

Connect entities (people, and machines – by, for example, helping pinpoint the right nodes in the network and making them discoverable to each other through search, irrespective of how distant those nodes are). Think of the need for experts from different disciplines to find each other quickly during a pandemic.

Curate knowledge. For instance, semantic searches and computer vision that identify the most relevant text and visual content, and cluster it for people to process it more easily; sensor data from network edges (e.g., internet of things "IoT" devices); "automated data discovery" which combs through multidimensional and even unstructured data sets to identify patterns that humans would take too long to evince; combine knowledge like it is done in the field of generative design where human intelligence completes the selection of hybrid ideas generated by machines. Again, think of the importance of machines triaging a vast amount of literature generated in response to the COVID-19 emergency.

Collaborate across the network entities, be they people or machines. For instance, natural language processing that automatically translates content, or machine learning that optimizes video and voice transmission. In the case of the coronavirus emergency, a large part of the economy based on knowledge work leaned on remote collaboration enabled by technologies vastly more reliable thanks to AI's ability to optimize data loads.

…and **Compute** any other prediction[24]. For example: using machine learning to determine which participant is worth rewarding, detect spam or fake inputs; or identify strictures in high-volume process flows) that can either trigger an automated response or be passed on to a human (or a network) for validation; or identify patterns to make prediction, like in the case of an ML algorithm that mines news to make economic forecasts[25]. Machines can also compute feedback loops between parts of the network, for example to identify the effectiveness of natural experiments in a network such as diverse modes of containment of an epidemic, or the detection of actual dangerous side effects – typically surfacing through contact center or doctor's records - in patients as new drugs are introduced quickly.

It is also important to note that machines and people work alongside each other in networked intelligence. To cite just one common example, human crowds often provide input for the machines (e.g., through annotations that train AI in fields ranging from language translation to autonomous vehicles). For more detail on the design of processes that harness both inputs, see the section "A network intelligence's architectural blueprint (or its garden layout) " on page 63.

As noted in the opening section, as of the end of 2021 seven of the ten largest companies in the world by market capitalization owed a good portion of their power to network effects – that is, the emerging strength from creating networks and harnessing their intelligence. Apple, Microsoft, Alphabet (Google's parent), Amazon, Facebook, Tesla, and Tencent all rely on their ability to connect economic nodes in an advantageous, AI-enabled way, curate the knowledge that comes from them, help the respective superminds (communities, markets, ecosystems, and to some extent hierarchies) to collaborate, and then compute to crystallize the power in data-insight-action arcs across their business processes. Economies of scale, scope, Wright's learning and experience curves[26] are classic concepts that take a new life in those companies and their networks, embedded into business platforms. From the App Store to Microsoft's DOS/Windows standards for programming (and developer ecosystems), to Google's search, to Amazon's supply chain standards, to Facebook's and algorithms and Tencent's WeChat communities, to Tesla's use of actual driving to train its algorithms, those competitive dynamics have been fundamentally altered. Not impressed yet? Take this: these seven companies together were valued at about 10 *trillion* dollars as of October 2021.

Now consider the following examples of *existing* collective-intelligence organizational designs that hint at new possibilities (more examples can be found in the endnotes[27])

Combining normal IQs. Collective IQ "engines" like hackathons, Wikipedia, Linux, Kaggle, and Kickstarter have revolutionized the world of collective knowledge creation and innovation.

Combining unique IQs. Sanjay Ghemawat and Jeff Dean, the founding fathers of MapReduce, are the only level-11 coders at Google (the highest level in that expertise hierarchy). Without them, we wouldn't have Hadoop, big data would still be out of reach, and Google wouldn't be what it is. Sitting side by side, they code on the same machine, complementing each other's brain in real-time. (In 2021, developer collaboration platform Github released Copilot, which uses OpenAI GPT-3 based CODEX to auto-suggest lines or entire functions, effectively providing an artificial "pair programmer".)

Combining anyone's creativity. Social media platform TikTok's Duet allows users to build on each other's video by recording alongside the other, triggering an explosion of memes (of admittedly varied quality).

Combining good-enough people with good-enough machines: the best chess player is neither a human (it hasn't been for a long time) nor a machine - it is a mix of them, and neither of them needs to be best in their field.

Internal markets. Some large companies, such as China's white-goods manufacturer Haier[28], have radically re-thought their organizational structures. Haier created a perpetual tension between internal units, reducing their default bias to work with each other. For instance, Haier's individual-market supply chains, marketing, and sales groups are free to buy external services if they don't feel their internal colleagues' offers are good enough.

Natural Language Processing has become very accurate. AI's ability to understand us and talk back is growing by leaps and bounds, combing existing (often crowd-produced) content at an unprecedented scale. Think of how useful Google search has become. Or just ask Siri.

Design thinking has emerged in the last ten years to become the *de facto* standard for innovation. It is an attempt to achieve collective intelligence by bringing together very different people.

Network-driven knowledge and learning. As we will explore in several case studies, there is an opportunity to enlist subject matter experts to curate knowledge and support the enterprise learning. The experience of GE-spinoff, digital business operations provider Genpact offers instructive insights on how that architecture helps people learn and use fast-changing skills in more scalable ways. They also cost a fraction of the traditional learning and development (L&D) and knowledge-management costs.

Network intelligences around, and inside, us

There's a bigger picture and a common thread. Radically innovative breakthroughs often germinate from the combination of ideas from diverse people from different fields as they interact over time. For instance, in the early 20th century, private parties full of wildly different intellectuals, industrialists, artists, and scientists launched a raft of world-changing ideas. Today, Silicon Valley, New York City, London, Berlin, and (on a good day) social media are all spaces where people and their ideas meet. Successful ideas are more likely to be generated through multiple contributions, as opposed to individual heroics. And certain structures of networks of people increase the odds of success[29]. Effectively combining very different ideas is tricky, however. It requires diverse people with strong networks and the ability to cognitively translate what counterparts share and to *enable* the combination of others' skills. Most people can't do that, and most collaboration environments aren't set up to do that - at least not at the necessary scale.

That's why we need to provide management with a new organizational framework. Let's build it in a few steps. For a decade, organizations like the Massachusetts Institute of Technology (MIT) Center for Collective Intelligence (CCI) have studied what organizational structures, collaboration methods and technology allow large groups of people to be collectively more intelligent. Five types of collectively intelligent networks have been identified by MIT's Malone and his teams[30]. These groups are *superminds*,[31] able to sense, remember, create, decide, act, and learn cohesively. In so doing, they also adapt to their environment.

Traditional, **hierarchical** structures are ubiquitous in organizations. They're the default option for that design and have been so for some time. They're fast and efficient. But they struggle to sense fast-changing environments and are poor at harnessing everyone's intelligence. We talked about the Chinese authorities' slow response and lack of transparency, in the crucial inception weeks of the Coronavirus infection in Wuhan. But it is fair to note how many western governments either lacked decisiveness (Italy, Spain, the US) or conversely and sometime simultaneously abused of hierarchical powers and neglected the collective input of scientists and experts (like in the case of the United States, where epidemiologists were sidelined by then-US-President Trump for weeks which contributed to the comparative lack of preparedness of the country, despite being the last major economy to be hit.)

Democracies are also networks, increasingly influenced by other media networks. In more than a few cases governing bodies have been taken by surprise and have arguably even been hijacked. To an extent, democratic decision-making has been applied at scale in online media from Huffington Post to Reddit. Democracies are important for decisions that require societal fairness and consultation, and they can cast the net more widely than any other type of organization, but they're also susceptible to misinformation and bias. At the outset of the coronavirus spread, people "voted with their feet" on what they felt the best course of action was, based on often unreliable information obtained by unqualified online channels. The political system in many countries also struggled to provide technically competent direction, as public opinion weighed heavily on immediate policy.

Large **communities** and movements are collectively intelligent systems. Communities are increasingly used to produce intellectual property (e.g., Wikipedia), or curate knowledge such as PatientsLikeMe. More generally, online communities are one of the largest emerging social phenomena of the last quarter-century. Apple is one of the most prominent nodes in that world, and the algorithms of both Facebook and YouTube exercise significant control on how those communities behave. Communities provide creativity and additional resources that are typically not accessible by normal hierarchical organization. But they're harder to set up and direct. In the COVID-19 situation, many ground level communities sprang into action, from elderly support to healthcare volunteering to donation of medical equipment to data analysis based on public records. Scientific communities of specialized researchers also provided much needed collective processing power. But others, such as the South Korean religious sect responsible for much of the spread in the country, amplified the challenge.

In the COVID-19 crisis, the scientific community (and its individual communities) collected and interpreted ground evidence in an accelerated fashion compared to traditional peer-review publishing, but arguably at times failed to connect cross-disciplinarily or innovate at scale beyond existing disciplines (for instance, by establishing broad collaboration channels, or crowdsourcing data and technology interventions). This slowed down the adoption of simpler, effective practices, such as the use of personal masks whose usefulness was perhaps obvious to more practically oriented designers.

Markets. Markets have for centuries exchanged signals through prices, helping decentralized suppliers and customers continuously adapt their operations. The most powerful example of a market powered and shaped by collective intelligence is so mundane that its rapid emergence has blindsided most of us. It's the internet advertising market. Online media sales leverage algorithms that understand what promotes fast and deep propagation of ideas. This is the same artificial organism that, when left unchecked, fuels sordid disinformation through social media, and invades the privacy sphere. Internet giants and hackers got there before the rest of the world did.

One of the reasons why markets can convey market signals is that prices include collective intelligence. A quote attributed to Elon Musk encapsulates the concept well: “The thing we call money is just an information system for labor allocation. […] We should look at currencies from an information theory standpoint. Whichever has least error & latency will win”.

Markets can be fast and efficient in providing information and deciding on it and tend to stabilize over time – but not as reliably as economists once thought. In the example of the Covid-19 emergency, market mechanisms were at least at first insufficient to signal the real risk, and price assets reliably, and they exposed the broader economy to a dangerous uncontrolled crash – so much so that monetary and fiscal policies had to be deployed to rein in the situation.

In yet another interesting example, investment markets’ short sellers had taken aim at financial services firm Wirecard well before its accounting irregularities became public and brought the company to its knees. And in early 2021 retail trading platform Robinhood showed how a “long tail” of small transactions, galvanized by social media (simple Reddit thread) in a swift collective action, could skew the normal movement of some stocks (e.g., Gamestop), unleash havoc for hedge funds short-selling strategies, and make the markets suddenly unintelligible for traditionally information-rich institutional investors.

The concept of natural **ecosystem** and the accompanying natural selection is particularly evocative: over long periods nature “makes choices” based on alternatives created by mutations of genes, as well as sense the environment and expresses those genes differently through epigenetics - and *remembers* those choices in the incredibly efficient database called DNA (a few years ago it was estimated that using DNA to store all world knowledge would make that data fit in the volume of a pickup truck[32]). In a way, this is the way nature "fights" entropy, and manages to keep parts of chaotic systems in equilibrium. MIT's Cesar Hidalgo has hypothesized that the same mechanisms scale up to the level of economies (where knowledge is stored in the networks of productive relationships between economic and social institutions); and down to the level of atoms, where physical concepts like Boltzmann's constant can be explained through the lens of information stored in effectively-structured networks. For Hidalgo, network-based structures, and the information they store, bring order and local equilibrium in otherwise chaotic systems[33]. Broader market-based ecosystems like large-goods supply chains are also networks, enabling the flow of ideas and matching between their supply and demand. Think about the literal village it takes to design and build the smartphones we all carry in our pockets.

Ecosystems benefit from the natural experiments that happen in loosely connected parts of the network and can provide significant creativity as well as resiliency. They're however harder to align and control.

A high-level synopsis of the strengths and weaknesses of these organizational meta-structures is shown below, highlighting the importance for combination of those structures for specific organizational processes.

	Hierarchies	Markets	Democracies	Communities	Ecosystems
Sense	Efficient, but limited to the immediate network; limited processing power	Broad reach; hard to create efficient and broad ones	Broad reach; hard to create fast-reacting scaled one	Potentially broad reach; hard to create intentional ones	Very broad reach but hard to create intentional communications and curate signals
Remember	Efficient yet limited to the immediate network and formally documented resources	Distributed memory of market actors; but hard to predict in advance what will surface in specific conditions	Distributed memory of participants; but hard to predict in advance what will surface in specific conditions	Distributed memory of participants, and communities' dialogue may be documented (e.g. Wikipedia, technical communities)	Distributed memory of participants; but hard to predict in advance what will surface in specific conditions
Create	Efficient, quick, but may lack creativity needed for very hard problems	Limited to provision of novel incentives and communication channels	Limited to provision of agreed norms and guidelines	Diverse input may generate creative output. Nontrivial management of creation process	Natural experiments may yield very diverse results; harvesting of best ideas is not easy if incentives for sharing aren't explicit
Decide (and act)	Efficient and quick, but may lack legitimacy or inadvertently miss subtler implications	Broad based and often fast; may overreact or overcompensate, especially in highly digitized markets	Participatory process provides legitimacy; risk of bias and interference in information sources or voting mechanisms	Possibly slow and prone to hijacking by vocal minorities	Unlikely to decide in coordinated ways except in special circumstances (e.g. emergencies, and definition of standards)
Learn	Efficient and accurate, but limited to network, and vulnerable to loss of key nodes	Crystallized in the price and availability of resources	Highly dependent on the quality of participants and institutions	Broad based and resilient to loss of some participants	Crystallized in network's structure, and participants' own learning processes

Each of these archetypes can be harnessed for the design of our future organizations. As we will explore in the rest of the document, many flows of activities can use a bit of each of the five archetypes.

Now think about the following real-life scenarios:

- Supply-chain planning processes use markets (e.g., spot prices of containers, demand for finished products from retail, the price of carbon for offset) and hierarchies (the decision rules included in the algorithm that planners use).
- Customer support for Apple products leverages a community of users, complemented by some hierarchically organized community managers who propagate the best solutions.
- The prioritization and display of information on a LinkedIn feed combines community-driven signals with an advertising market's dynamic pricing information.
- Or think of product display on Amazon. Democratic user votes guide the visualization of items when searching or filtering by "customer rating," but ultimately the hierarchical Amazon category management can decide to delist the product.
- Think of the early use of so-called *prediction markets* for research and development, used today for political and fashion predictions.

- Or open innovation ecosystems in consumer products, like the now-mature case of Procter & Gamble that has relied for years on a network of partners to inject novel ideas and co-create new products.
- Or betting markets, such as Betfair.

Let's now borrow from neuroscience for a moment. There are some fascinating similarities between how our brain works and the flow of enterprise information and knowledge. For instance, we find specialized areas in both of them, that do nothing but combine inputs from other regions to respond to unprecedented situations. In the brain, that work is done by the prefrontal cortex, and in business that is the role of senior management (or at least one would hope). Many brain areas can play a generalist role and most neural cells can perform any job at their creation, but they eventually specialize because of the connections they create with other cells that fire together in specific situations, like stimulation by sensory input. Many organic systems include very low-IQ (even “no-IQ”) entities have shown that for a long time, from fungi to plants, from insects to animals – over billions of years[34].

Can our management practices emulate some of these combinatorial dynamics? More in the endnotes but consider the following first.[35]

Business process re-engineering has been applied for 40 years, and we know how to create repeatable and efficient organizational processes. We can now design complex workflows that string together people and machines so they can scale. But that's not fully-fledged cognition - at least not yet - despite an increasing amount of robotic or AI-driven automation. It more resembles the reptilian brain, and parts of our own. Its instinctive nature is very effective in known situations, but devoid of a pre-frontal cortex and therefore lacking more complex, learning-based adaptation. To be clear: in specific conditions, that is the exactly the right model. A natural predator, in its environment, is dauntingly effective at using instinct to get what it wants. In much the same way, automation helped by dynamic workflows fed by analytics can generate fast **insight-decision-actions** closed-loops. That can be very effective at enhancing a range of business processes, from catching fraudulent financial transactions to identifying causes for potential failures in a jet engine’s maintenance log. However, those organizational models are still not intentionally harnessing organizations’ deeper functioning as they attempt to find new solutions to tough, fast-changing problems.

Some have observed, based on scientific models of innovation generation and propagation, that combinatorial innovation is the core engine in potentially exponential dynamics of innovation in our future – with the only constraint being the ability to retrieve and process the most promising ideas[36]. What does that networked flow of information look like, both inside and outside our organization? We don't know for sure, but it likely doesn't look like org boxes and workflow swim-lanes. More on this in the endnotes[37].

This is not pie-in-the-sky thinking. Quite a lot of work has happened already. For instance, we now know that the *network signature* of information flows varies, depending on the type of work that is being done. Consider the two extremes displayed below.
The example on the left shows the connectivity between *different* and *diverse* groups. In that

network, ideas are combined and find ways to be pre-tested by people with different backgrounds. Think of the interaction between R&D, marketing, sales, and supply chain, for instance. Or the cacophony of voices on professional social media.
Efficiency-driven organizations have more regimented connections. This example is more typical of G&A functions (finance and HR) or industrialized environments, where business partners orchestrate the communications of tightly managed workflows and corresponding people. This pattern also emergences in times of pressure and when speed is the primary factor, as note in the example of reduced size and international diversity when writing scientific papers during COVID-19.

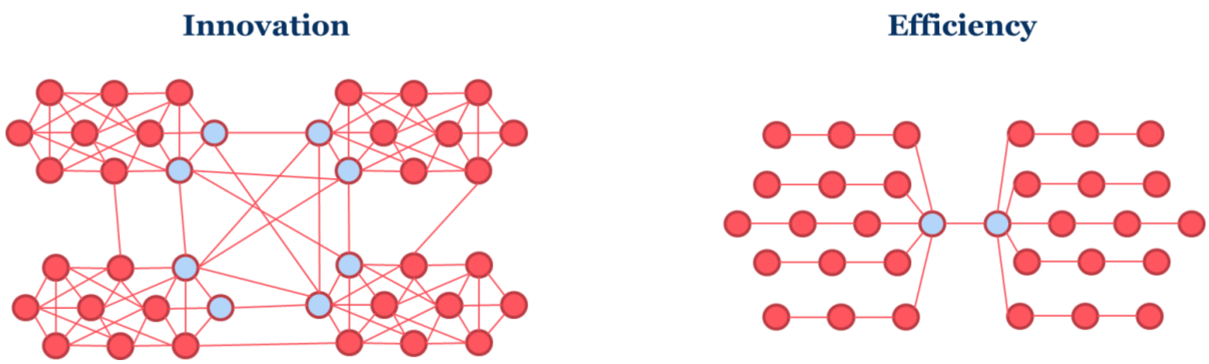

These two simplified structures, once scaled up at an organizational size, create networks that generate very different types of behaviors. As we will see in one of the chapters at the end of this guidebook, the impact of COVID-related lockdowns damaged the "weak ties" that are important to innovation, while the strong ones that underpin efficient network dynamics were largely left unscathed.

Let's now use two relatively well-known perspectives to further illustrate the importance of networks for organizational design.

In "Thinking Fast and Slow", Daniel Kahneman[38] makes a compelling case for the co-existence of instinct (speed, efficient response) and superior intelligence (creation of alternatives, and less deterministic decisions) and highlights the importance of using each for the right things. Indeed, organizations struggle when they mix up their neural structure, and do not apply the most fit-for-purpose structure because of their history and natural tendencies. A second, intuitive example of application of these ideas comes from the world of auto racing, where drivers hone their instinct through data-driven simulations, and then apply it, in split-second increments, during the races (they do think fast). The backbone of that process is provided by data scientists. Armed with an unprecedented number of sensors and able to churn out relevant insights before, during, and after the races, they inform the simulation that the pilots will then go through to reinforce their instinct. The data scientists are applying slower thinking, but those networks of people, augmented by deep data-driven insights, boost the performance of the individual nodes.

The intelligence fabric in both examples is represented in the next chart. Let's simplify for a moment. Individual people owe their capabilities to their experience, which can be increased through learning and development (L&D) activities, and further enabled in the flow of the specific processes they run. (Think of a salesperson using specific collateral documentation for a sales pitch created by a practice leader, or a contact center agent prompted to provide a specific reaction based on a machine-learning algorithm that optimizes the customer interaction).

Both L&D and enablement resources are provided by the organization, and indeed those practices are among some of the best predictors of enterprise effectiveness[39]. But here's the important twist: those capabilities are amplified by the connections that people have, which help them complement their skills and use others as creative soundboards. The combination of individual knowledge (existing, new, and targeted enablement) with network connectivity generates collective intelligence that's superior to the mere sum of people's intelligence and their knowledge[40]. The recently launched Microsoft Viva, which aims at redefining the exploding "employee experience" market space, largely follows this structure and bets on the fact that collective

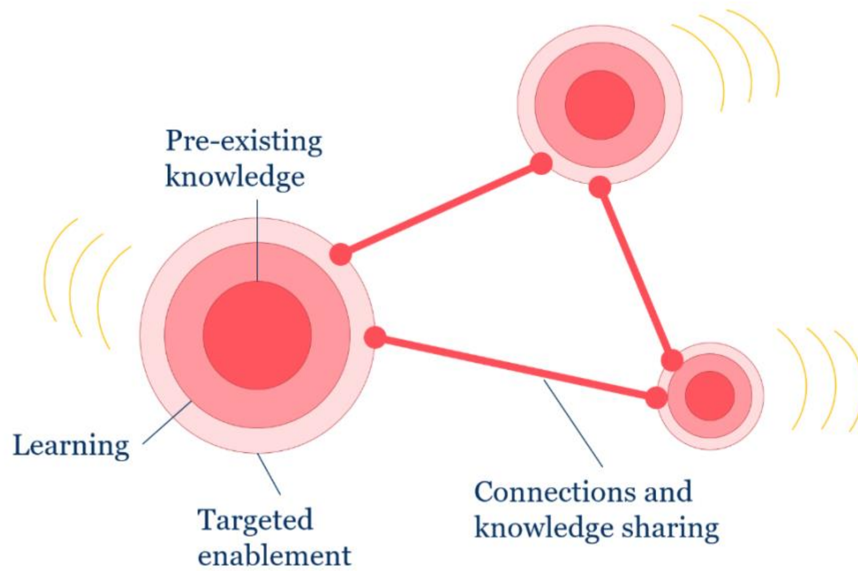

cognition is also linked to employee experience. Neuroscientists would say "there is no cognition without emotions" – an equivalent is likely at play in the enterprise networks world.

But can we create a framework that can be used broadly? Let's borrow again from MIT's Malone, who has abstracted a helpful framework to think about what it takes for a system - *any* system - to behave (think *and* act) intelligently. See the next figure: an intelligent system has processes to sense its environment, to remember what worked in similar conditions, to create options and decide upon them before acting, and then to learn from the result. More on this in the endnotes section[41]. This is a lens to determine what technology, methods, systems, and groups or networks of people are needed to achieve superior intelligence. For instance, a Research & Development organization can get better at sensing promising new trends, retrieving relevant experiments tried in the past, and being able to overcome decision bias. Many other enterprise processes display a combination of some of these cognitive abilities, as roughly depicted below. As noted in the previous section, different archetypes of network intelligences (e.g., hierarchies, communities) have strengths and weaknesses in performing those activities. They require a deliberate design – a combination of them – for the processes to function correctly. (The accuracy of the chart doesn't matter – what matters is that the respective elements exist for those processes to work.)

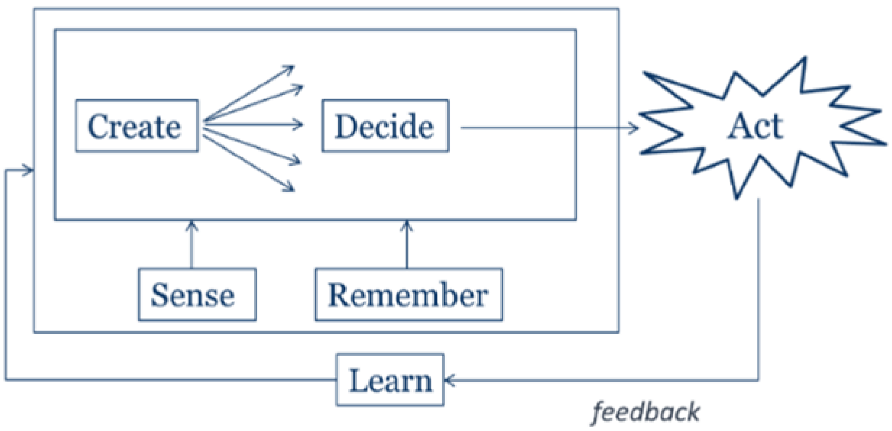

Many of these processes already use technology for each of those cognitive steps. For example, Internet of Things technologies are increasingly deployed to sense the environment. Artificial intelligence is used to remember patterns of what has worked in the past. Collaboration technologies that bring people together to *create* and *decide* are now commonplace. And for many years, enterprise resource planning (ERP) technologies and workflows have scaled to become the backbone of large enterprises. Additional examples can be found in the endnotes[42].

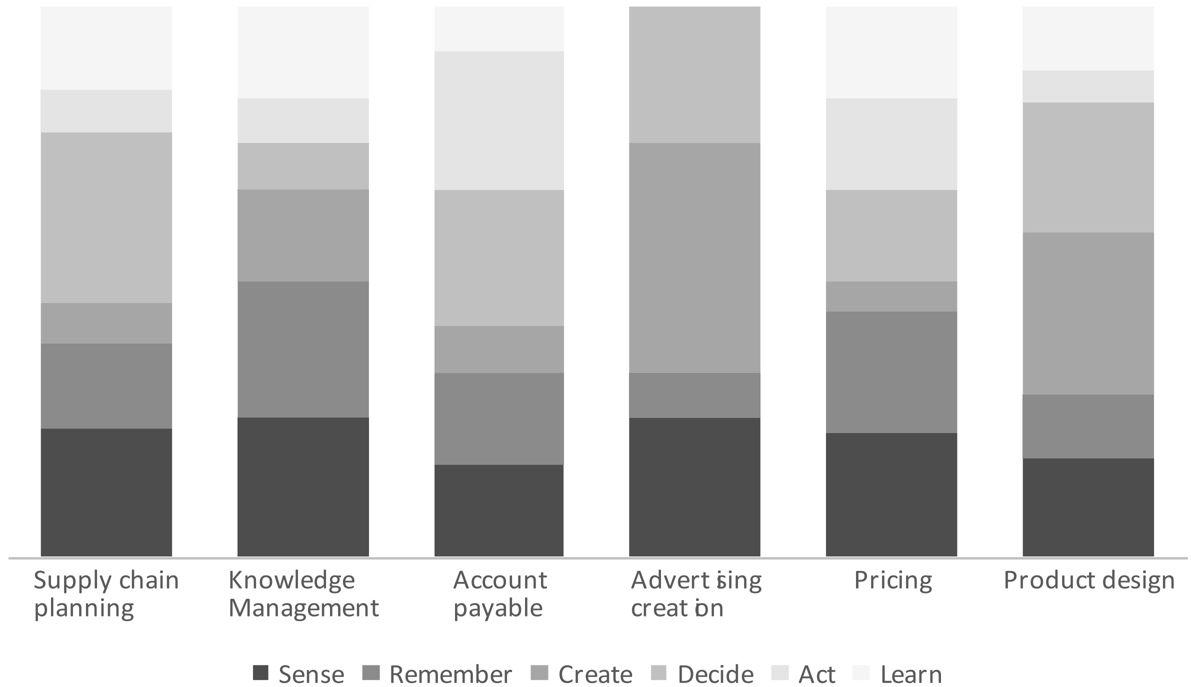

What is very exciting now is the ability of those systems to surface document data and then use artificial intelligence to memorialize the learning so it can be used the next time groups of people need to remember what worked. AI distributes the benefits of the many natural experiments that might have occurred throughout the organization or the broader ecosystem. The other big twist in the story is that in order to fully harness the power of these technologies and methods we need to move away from thinking that organizational charts and deterministic process workflows, and the related practices, are the best ways of achieving these goals. In many respects the organizations of the future will spend a higher proportion of their time exploring what the best solutions are. In the past that exploration and design was reserved for a small minority of people, with the rest of the time and energy dedicated to the exploitation and industrialization of those processes.

Building the four components of a supermind

How do we use everything we discussed so far to design and build an organizational structure able to do things better than today’s combination of org charts, processes, and somewhat siloed machines?

Take the example of a complex crisis like the COVID-19 pandemic, that not only suddenly generated a volatile business environment, but also shattered many of the established ways of working of many organizations. How would a company amplify its supermind to respond to threats, and even move to the offensive in such uncertain time?

Even in normal circumstances, management teams often struggle with detection of threads and opportunities, creation of new ideas and plans, and then concerted yet nimble and intelligent execution. Too often management teams zoom in on collaboration technology as the first point of entry for these challenges. In this document we carefully combine technological advances with organizational design. Because in the end, the challenge tends to be about the change management and empowerment of people, at scale. Technology – even communication technology - in isolation doesn't solve for that.

The four modules listed below help structure the work of invigorating an organization's collective intelligence, augment people networks with machine's – and vice versa. Each of the modules is a lever to steer the system's dynamics towards higher collective cognitive performance. They allow us to take a more *active* management role of that complex system. We will explore each in turn, and in the likely sequence of initial implementation. For more detail on how to implement them, refer to the "How to create a supermind" section.

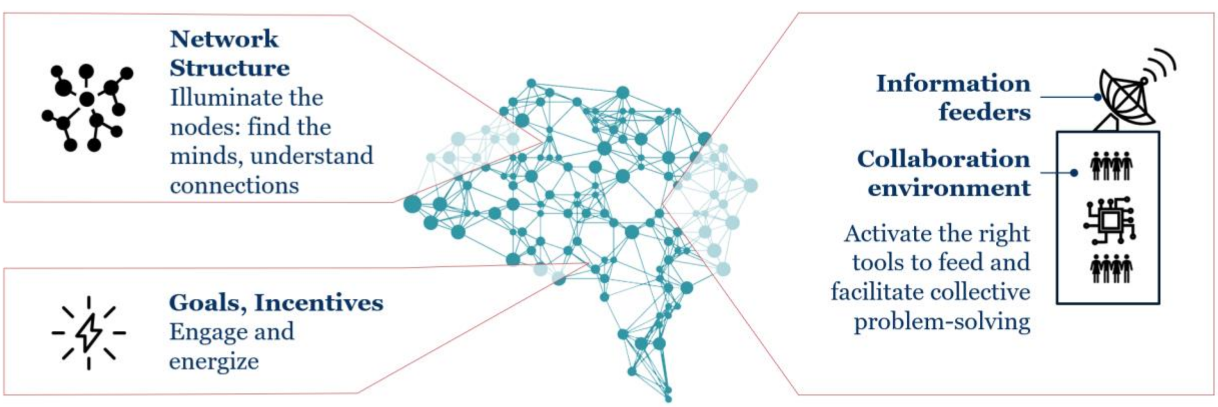

This method can also apply to designing traditional command-and-control business processes, as a complement to established digital transformation and Lean and Six Sigma approaches. But while conventional workflow design already has enough practices (and practitioners), the gap to designing networked intelligences is much wider. So that's what we will focus on.

First, illuminate the network

Find the best **minds** for your challenge - both human and machine - and understand how they're **connected** to each other and to the world around them. While we have good search engines for content, our ability to readily find people based on their skills. And yet, if they're not connected, if the individual nodes struggle to find each other and engage meaningfully, they will not be able to easily build on each other's ideas. The network will look like a foggy cloud, with few and scattered local nodes connecting only to their closest peers. That's where most organizations are at: the immediate reach of individual employees is only a small part of the network, and while that might be sufficient for normal transactional work, it is a handicap when trying to generate new knowledge and implement new plans. And that's where the world is: 1% of the population typically has the IQ of a genius – that equals about 75 million genius in the world…but how many really manage to generate real impact? Irrespective of generic genius traits, many people with strong potential to contribute are systematically overlooked because they're not found, or because collaborating with them with yesterday's means is impractical.

People struggle with *scalability*. We suffer an inability to process large amounts of words because of the time it takes us to understand natural language. We can't hold more than one meaningful conversation at any time and aren't naturally able to maintain connections with very large numbers of people. But a lot of value can be generated by simply making the nodes discoverable to each other in a way that doesn't overwhelm them and helps them connect, even if that's done through basic, conventional means. Think about startup meetups or company yearly kick-off meetings, both of which serve the same function. Help connecting is especially important for people, because we rely on natural language and can flexibly sense our environment and hence connect with interesting (human or artificial) nodes that extend our senses and memories.

The practical applicability of these concepts is enormous. Networks and their analysis have proven significant in adaptive-yet-disciplined environments like enterprise business-to-business sales for instance[43]. And even more obviously, network structure is a key determinant of innovation. Innovations travel in non-linear, iterative ways across the stages of exploration and discovery, to shaping and development, then to diffusion and embedding. Most people only span a few of those stages individually. But their networks may go far beyond that. For instance, an exploration network tends to be broad, made of many (typically unrelated), and (often into multiple parts of each local network, to trigger adoption at scale. Recent research shows[44], that is by implementing, by a tension that leads to change. Find who those people are and help them to understand their roles and to facilitate the right connections. In this phase, AI can help pinpoint the relationships between key ideas, people, and potentially machines that generate information. It can then help curate their knowledge and make that knowledge – and their sources – discoverable to each other, facilitating the emergence of connections at hyperscale.

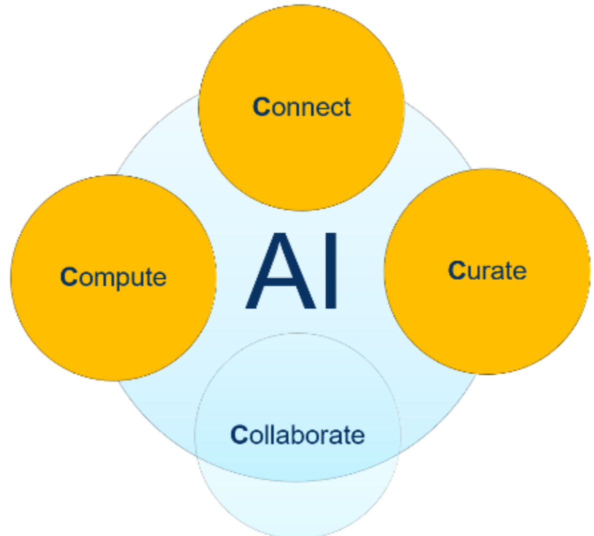

Second, direct and incentivize the network

Now, engage the most important actors, direct their impetus towards a goal, **and** energize it[45]. That's their ability to go above and beyond what they do today, to redirect a bit of their time to what you care about. At the very least, they're open to following the direction of the network you're harnessing. You can do that by designing and enabling the propagation of the right intrinsic (e.g., recognition) and extrinsic (e.g., money) incentives, in a way that's aligned to a meaningful and clear goal. Successful organizations do some of that routinely through cultures that foster *trust* and *alignment* by defining directions and norms that activate people even in the absence of explicit, specific direction.

This is actually very hard to do, and it typically doesn't happen naturally. Many networks fail to ignite because incentives are insufficient to encourage the cognitive effort (and change to personal routine) requested of the nodes. People themselves, if left to their own devices, don't always readily collaborate. They certainly don't naturally collaborate at hyperscale, with some notable exceptions such as highly energized social movements.

Traditional systems of incentives (largely based on individual key performance indicators) tend to do a poor job of galvanizing networks. They focus on organizational structure roles, not network roles which cut across org chart roles. They also tend to disproportionately reward individual leadership. That's probably because they were designed for comparatively low-and-slow-information environments, where it was hard to detect people's contribution outside of a core group with direct influence on the outcome. In those environments pervasive innovation was less of an everyday necessity, and long linear cycles of comparatively command-and-control processes were good enough to adapt to more stably growing markets.

Another route increasingly traveled by many innovators is the design of a seamless human-centered experience for the participants, to minimize the amount of cognitive load required from the contributors. The higher that load, the higher the value of the proposition required to attract people, and the higher the risk of failure. Captcha enlisted billions of people to transcribe millions of books, by showing us a picture and asking us to prove that we aren't a robot, when trying to log in. Waze leverages millions in crowdsourcing data for navigation. These are simple examples of seamless injection of user experience.

An additional good strategy is to go where the nodes are instead of requiring them to come to us and do things they weren't going to do. It is often surprising how much pent-up energy exists in people who are eager to be given a chance to contribute. Think of people who volunteer for special projects on top of their existing workload. Or people who run skunkworks efforts because they're passionate about it.

And we can also use their information exhaust to amplify their contributions. Companies like route-navigation Waze offer a very balanced ratio between what they give (guidance enriched by fellow travelers' actions) and what they ask for (the GPS coordinates of your telephone, and the occasional action to flag an issue on the road). In this area, AI can help detect what drives behaviors, or compute the impact of individual contributions (for instance, to new knowledge) in ways that were impossible until now. And by enabling communications at hyperscale, it helps connect people which, if done right, can itself be motivating.

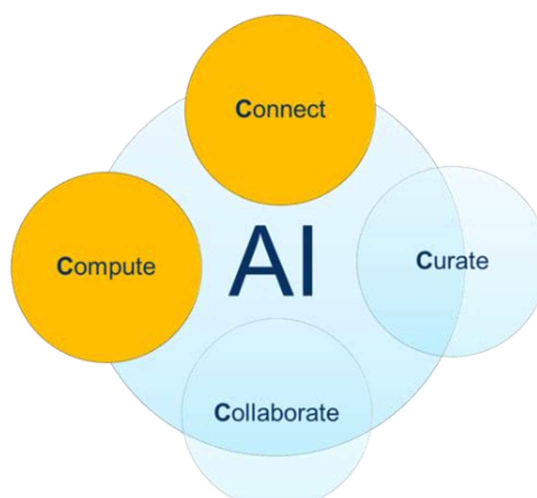

Third, engineer "information feeders"

Third, build a mechanism for **feeding** relevant, usable **information** into that system composed of nodes and incentive energy. Even relatively simple cognitive systems, if fed with enough information, might become capable of expressing intelligent behavior. Examples in the natural realm, from trees networks to ant colonies, abound.

A knowledge feeder is, in a way, a better way to search for relevant ideas without the limitation and bias of standard search engines, and with the additional value of finding adjacent ideas filtered by relevant people. A primitive yet useful knowledge feeder can be obtained simply by a deliberate and systematic sifting through relevant web searches, following relevant people, and feeding the results to the networked intelligence. Quality control can be performed by establishing fact-checking filters at the receiving end – a problem sorely felt during the COVID-19emergency, when most leadership teams struggled with separating accurate, credible information from noise and falsehood. As simple as this may be, it is rarely done intentionally and appropriately in conventional organizational design or project management.

Another basic knowledge feeder is learning and development programs, that train people by providing external input. Training however tends to be slow to evolve, and while valuable, it may prove too rigid for highly volatile situations – such as responding to a pandemic for example.

The first thing a knowledge feeder does is to enable some *sensing* of the environment. It brings signals from the world to the doorstep of the network nodes. For instance, a marketing manager can see competitors' moves. Multiple industry experts in a community can witness the launch of a new product. In Industry 4.0 among other models, internet of things (IoT) sensors become pervasive information feeders.

A knowledge feeder also helps the network *remember* what has worked elsewhere. For instance, medical school and crowdsourced environments like PatientsLikeMe can help medical communities get smarter about what has worked in specific circumstances in the past, sparing individuals the need to form those experiences directly. In many companies, the foundation of a collective knowledge feeder often already exists through the convergence of three things: electronic communications, especially email and more recent collaboration environments like Slack (acquired by Salesforce in 2020) or Microsoft Teams (and related templates that create enterprise wide consistency in people experience) or, in open communities, tools like Circle[46]; organizational structures that include subject matter expert, e.g., practice leaders in consulting firms; and pervasive information feeds fueled by the explosion of internet-based knowledge sharing from social media to specifically curated knowledge feeds.

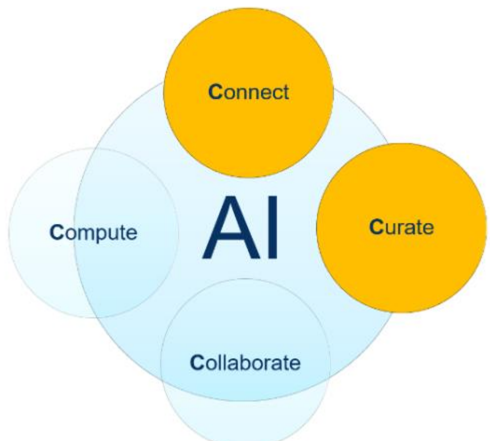

Without AI, we wouldn't be able to curate content at such level of scale and speed while maintaining some level of relevance and accuracy. AI also helps surface the sources of that knowledge, especially people but also machines, and connect them in ways that would have been completely impossible before.

Fourth, provide a platform for collaboration

The next step is to **enable collaboration** at hyperscale by eliminating friction in the interactions and feedback: to make nodes process inputs faster and make them learn. Think of what happened in the aftermath of widespread office shutdowns due to the COVID-19 emergency: the connectivity between the network parts was impaired, at least for some time, resulting in many decision-making processes being slowed down. The reaction, almost uniformly, was a scramble to leverage remote-collaboration tools and practices.

An effective collaboration platform (composed of *both* technology and methods) is also important because *combination* is at the core of innovation, and combination happens best when people interact, compete, and ultimately (especially with hard problems) collaborate with less friction.

Doing that with a finite number of people has been the focus of traditional approaches to creative-group dynamics, including design thinking. However, fostering combinations at scale – *very large* scale – requires a very different infrastructure and ways of working.

Imagine a world where you can access the ideas of the right people. People who have the right skills and are exposed to the right inputs, not just the group of people who works for you or is located close to your team.

Prior to the creation of email, it was hard for groups of people to build on each other's thoughts unless they were in the same place, or on the same telephone call. When they did, for instance through scientific societies, progress was slow. Email was the first tool to enable an asynchronous buildup of ideas. Three decades later, large networks of people now have a plethora of choices to collaborate – from asynchronous threads to synchronous videoconferences. The result has been an explosion of virtual collaboration, and the increasing ease of documenting what's been said – and then extract it.

What does AI do here? Intelligent machines help connect the right nodes in the network, make them discoverable to each other, and enable interaction in the format that's most appropriate for the need at hand, with both synchronous and asynchronous communications. They will become increasingly effective at doing so, as fields like natural language processing (NLP) evolve quickly.

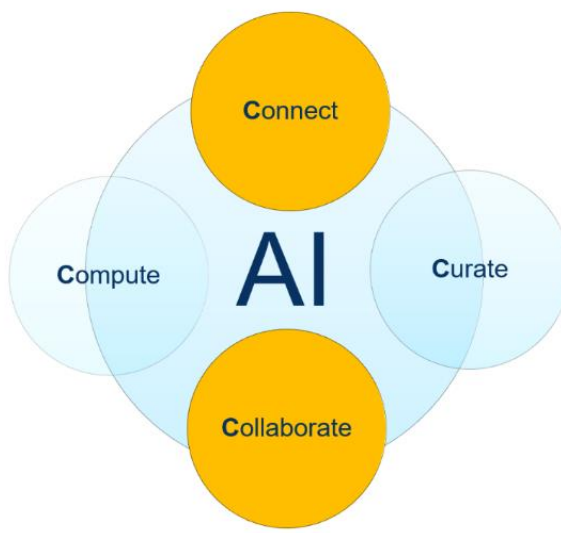

A workplan for a supermind

To summarize, amplifying the collective intelligence of a network of people and machines requires work in four main areas. They aren't typically found in a senior leader's job description, but they should be:

- The identification, analysis, and enablement of mutual visibility for network **nodes (network structure illumination)**
- The design and implementation of appropriate systems of **incentives** for collaboration that align with meaningful **goals**
- The enhancement of relevant and ubiquitous **information feeds**
- The creation and management of an appropriate **collaboration platform**, from new technology to new practice

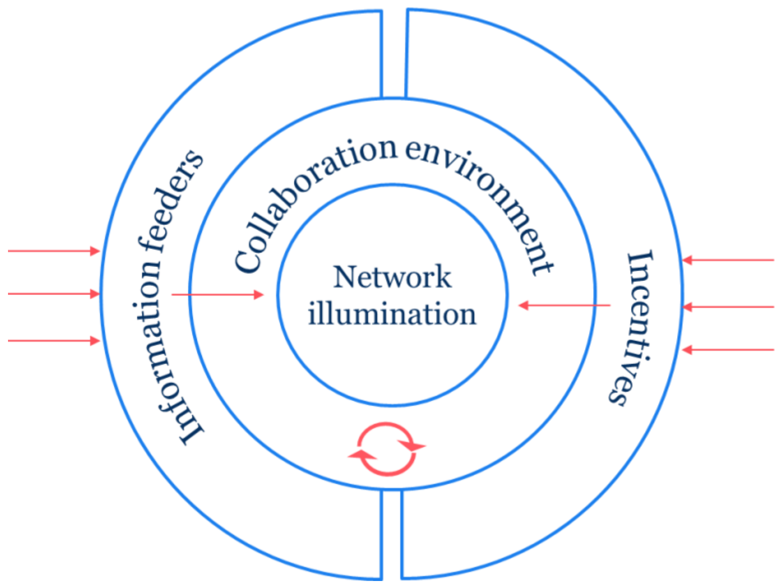

Deliberate management of such an augmented collective intelligence -a new operating mechanism - can lead to superior network intelligence and a more intelligent organization, at parity of current resources.

Combining all the perspectives we explored so far, the following picture emerges of what a network-intelligent enterprise, powered by increasingly intelligent machines but also harnessing its human counterpart, can do. Its ability to sense, remember, decide, act and learn could generate superior ability not to just react instinctively, quickly, and effectively, but also to adapt to entirely new conditions.

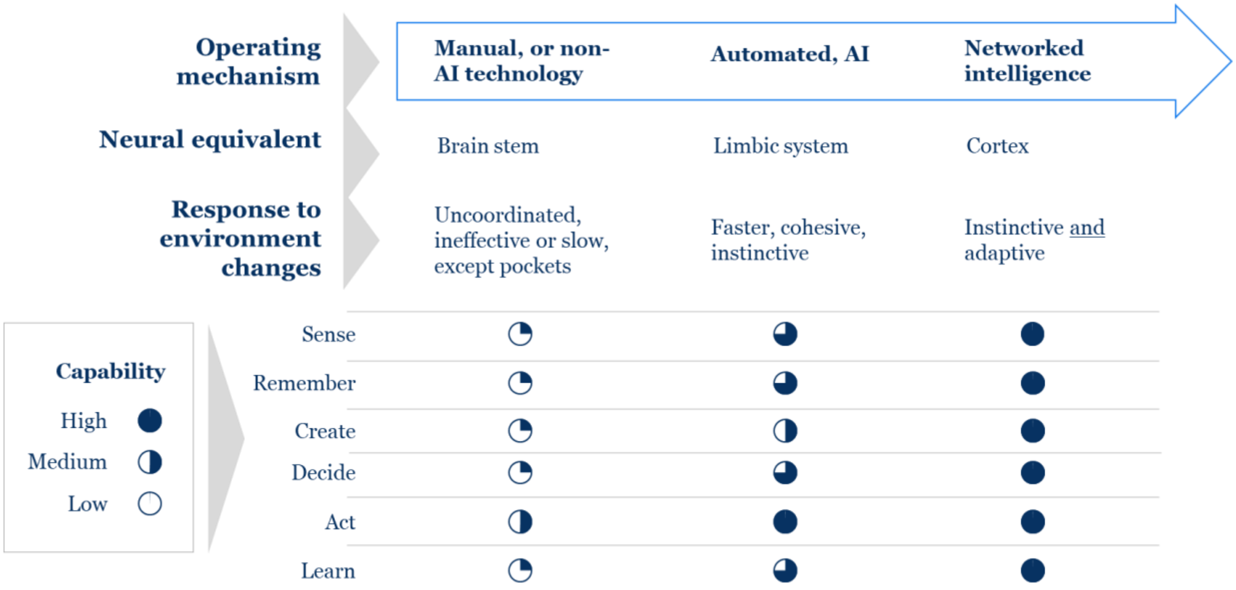

Once more, not all workflows benefit from a networked intelligence of this nature. Just like the human's frontal cortex isn't invoked in everything that we do, some stable processes are perfectly OK if run on little or outdated technology. (Ideally one can still extract data from it that can generate additional insight, such as through process mining technologies increasingly applied to ubiquitous older ERP technologies). Some processes even benefit from being purely instinctive as opposed to fully intelligent, especially those that need to drive fast response to known patterns. In such cases, such as catching spam and detecting typical malware, the insight provided by the machines can be injected into a closed loop enabling fast action.

However, in complex enterprises faced with fast-changing and complex (not just complicated) conditions, a hybrid human-machine networked intelligence is arguably more effective.

What could we use a network intelligence for?

Arguably, organizational design has attempted to recreate the cognitive functions for decades. However, gaps exist. That's primarily because without better technology it was impossible to connect the right entities, provide them with curated knowledge, make them collaborate, and improve those loops at scale. We can achieve the requisite level of connectivity today when we deploy AI to help the four key modules of network intelligence: illuminate the network, provide appropriate information feeders and incentives, and support the creation of a ubiquitous collaboration environment.

A networked intelligence system could power many things - including the examples below. From the launch of new tech products and services to fostering adaptation practices in precision agriculture in the wake of climate change, network intelligences can tackle many large challenges. When pervasive and mature, they can become a source of competitive advantage, one that becomes stronger as more time goes by and more interactions happen.

Innovation	Collaborative invention-to-innovation processes
Consulting & professional services	Effective, efficient, distance-proof collaboration
Enterprise sales	Better engagement of pre-sales and architects
Customer support	B2B; B2C support community (e.g. tech, content moderation)
Customer experience	Flow-of-work optimization across the enterprise
Research and analysis	Thematic research for financial or R&D use
Re-skilling of workforce	Improved learning operating models
Key employee retention	Identification of critical employees, early warning, coaching
Diversity & inclusion	Assessment of inclusion of diverse individuals
Organizational design	Identification of inter-department collaboration gaps
Forecasts and predictions	Superforecasters, collective forecasting, better algorithms
Strategy	Horizon 3 analysis, scenario refinement, validation
Marketing	Thought leadership and research, customer insight
IT organization	Control, orchestration and support of distributed technology landscape
Data science	More pervasive data engineering and science
Cyberthreat response	Better understanding of organizational preparedness
Post merger integration	Acceleration of synergies
CEO, CxO effectiveness	Improved collective problem-solving and adaptability
Drive change initiatives	Propagation and adoption of transformation programs
Better AI	Establish and embed ethical norms in AI algorithms; human-centered AI
Geopolitical diplomacy	Identify critical gaps in connectivity between players
Social investments and ventures	Enhance the synergy in a portfolio of ventures funded by public or private money

Exploring the opportunities in more detail, networked intelligence could be applied to these use cases as described in the table below.

Innovation	Networked intelligence can be applied to private (e.g., new products or services, new business models) or public efforts (say, identification of appropriate practices for precision agriculture in south Asia in the wake of climate change). Identify cross-functional, cross-organizational teams, and couple them with external people who bring very different perspectives – utilize network analysis and NLP based search to identify them. Define incentives for participation upfront (e.g., KPIs, public recognition, special rewards). Consider using external crowdsourced platforms to add and process ideas. Ensure a knowledge feeder with external ideas (e.g., news, blogs) is created using proper semantic search, not just informally managed search engines. Create a dedicated enterprise social media channel (e.g., Teams). Use specialized tools for synchronous virtual sessions e.g., virtual sticky walls to perform scaled-up design thinking sessions. These methods apply to both initial ideation and identification of implementation and related local adjustments.
Enterprise sales	Identify the key salespeople and solution architects. Create a baseline of the type of network connections they have - for instance, how many connections they have in the other group. Incentivize and educate them so that they forge those bonds. Ensure that they understand each other's work, for instance through feeding relevant knowledge as part of mandatory training. Create and foster the adoption of collaboration platforms that minimize the friction in the interactions, for instance through intuitive and powerful synchronous and asynchronous communication tools.
Customer support	In business-to-business services: identify the key responders to customer queries and analyze their network dynamics – both within the responder group, and with the client. Establish which behaviors (e.g., response speed, consistency of process for responding) result in stronger customer

	satisfaction. Evolve existing incentives so they encourage the new behavior. Create a collaboration environment where the behaviors can be displayed as a team and be used for feedback sessions. In business to consumer environments: nurture and incubate communities of expert that complement your in-house customer interaction teams. Use deep semantic analysis to understand trends. Ensure that product management teams are connected to those groups in a bi-directional way so that they can learn from the new queries and help create knowledge artifacts to feed the distributed problem solvers.
Up- and re-skilling of workforce	Identify "knowledge nodes", i.e. both subject matter experts (at different levels), and people who relay that information into the various business units. Incentivize them to spend some time curate knowledge and engage with learners in interactive sessions. Give the system additional external input by curating relevant materials at scale, through professional knowledge managers. Ensure a common collaboration platform exists where knowledge nodes can interact with knowledge managers, and learners can consume content and provide feedback.
Key-employee retention	Identify the population at risk – critical existing employees, new hires. Identify the network patterns that predict attrition, i.e. the "signature" behaviors such as increasingly lower centrality in the network as the person withdraws and focuses on external networks, lower contribution of input to teams around them. Incentivize and educate people around them (e.g., supervisors) so they detect those signatures and start acting, for instance by spending more time with the employees at risk. Specific roles might have particular signatures: for instance, new salespeople who don't grow enough connections with their sales support counterparts are at higher risk. Or consulting staff who don't create enough diversity of connection and strength of links with peers[47].

Diversity & inclusion	Analyze the network of interaction between people, utilizing data that indicates key characteristics such as gender, tenure, etc. Compute the network characteristics for the "diverse" cohorts, to understand if they have suffered from lower network integration (e.g., they're less central, they get responded to less quickly etc.). Incentivize people around them to support them (e.g., by becoming sponsors, coaches and mentors of diverse employees). Inject new practices and thinking about management of diversity. Provide a platform for open exchange and debate.
Forecasts and predictions	Identify nodes across the "data-insight" arc, such as external sources (both machines and people), and internal workflows. Identify additional nodes, such as people in the enterprise network who have a good ability to forecast the actual results (for instance, specific sales, or supply chain people) – irrespective of their formal position in the existing forecast process. Incent them to participate and incent the rest of the network to take their input seriously – for instance, embed that input formally into the actual "insight-to-action" workflows that drive execution at scale. Provide a collaboration platform for the forecasts and feedback loops (i.e., forecast vs actual) to be shared, above and beyond what the existing systems would do (traditionally, ERP systems with very limited collaboration capability).
Strategy and marketing	Identify thought leaders who propagate and relay fresh new insight, especially related to so-called "horizon 3" i.e. areas where the current organization isn't an established incumbent. Assessing the network position of the nodes to ensure diversity of input. Create a knowledge feeder where their signal is machine-curated (NLP, specialized search). Then, filter through knowledge managers – who are enabled and incented to do so. Consider incentivizing some of the nodes in the external network to provide "insider" insights that wouldn't otherwise be surfaced. Establish a collaboration platform through which internal experts (e.g., practice leaders) can collectively work on crystallizing the new insight – for instance through virtual text editors (mindmaps, collaborative writing e.g., Microsoft Office

	365 or Google Docs). For strategy, that collaboration platform is the environment where scaled-up debate happens, irrespective of location.
Data science	Identify the nodes of the "data-insight-action" arc, i.e. the people (and associated machines) who intervene at various stages of those processes in an enterprise network. Ensure the right motivation exists for the adjacent nodes to collaborate with each other. Knowledge feeders may need to be enhanced, to avoid insularity – for instance, by increasing the amount of "ground truth" data received, or external benchmarks. Then, consolidate the collaboration platforms across parties – in some cases, a dynamic but structured workflow would suffice (e.g., for supply chain planning routines), in others there will be a need for more fluid conversation between parties e.g., to design data models. AI is used pervasively here: for instance, dynamic workflows benefit from machine learning e.g., to identify the right escalation path above a certain exception threshold; the models themselves would likely use significant machine learning; the sensors for ground truth may use NLP or computer vision.
Post-merger integration	Analyze the network of connections between the new and pre-existing entities – initially, and ongoing. Identify strictures. For the critical nodes (individuals, but even more frequently groups), ensure that incentives exist for collaboration to become part of their daily priorities. Generate information feeders to those groups so that they start understanding each other's' world, in order to gain mutual respect and forge a shared language. Create collaboration platforms so that individual groups are continuously enabled to interact, and the initial friction is addressed by senior leadership as appropriate, and by other employees who are deliberately acting as connectors (e.g., provide ongoing coaching and guidance).
Change initiatives	Identify the "spine of change", i.e. the networks of people who will act as gatekeepers throughout the organization. Motivate them to participate in the "change movement". Ensure that relevant information is fed into those

and Centers of Excellence	networks through appropriate filtering and communication of news. Create an environment where collaboration between the "new" and the "existing" happens in a frictionless manner – for instance, through succinct FAQs, rapid turnaround of queries, structured and efficient escalation paths.
CEO and CXO effective-ness	Senior executives can benefit from understanding the flow of information and collaboration energy across organizational groups, as well as being able to visualize a sort of "circadian rhythm" that shows the ebb and flow of that energy. They can also benefit from setting up a more effective set of sensors that sift through environmental signals (e.g., competitive moves) and deliver insights on emerging combination.
Better AI	Identify the at-risk parts of the network, for instance advertising markets and news outlets. Also identify supporting resources, such as fact-checking websites, respective AI-enabled automated fact checkers, as well as potential pools of crowdsourced support. Create the right incentives for these parties to collaborate, such as public recognition or regulatory oversight. Formalize the information feeders' flows, so it is efficient for the supporting resources. Provide collaboration platform for continuous learning, and escalation.
Geopolitical diplomacy	Identify the structure of key nodes that influence decisions across and within countries, and ensure that they're properly connected, incentivized, fed with relevant information, and have effective collaboration environments. Irrespective of economic and military heft, network-rich players are likely to punch way above their weight.
Social investments and ventures	Enhance the effective of a portfolio of social ventures funded by private philanthropy catalyst funds, or public investment funds. The components of an augmented intelligence network can be supported centrally by such funds and facilitate the emergence of synergies across venture entities and their ecosystems.

This is no doubt just a partial list. Other applications abound, including those in law enforcement and the military (they've been explored by expert authors[48] and they're outside of the scope of this document.) These use cases fall broadly into three categories, summarized in the next simplified, non-exhaustive chart.

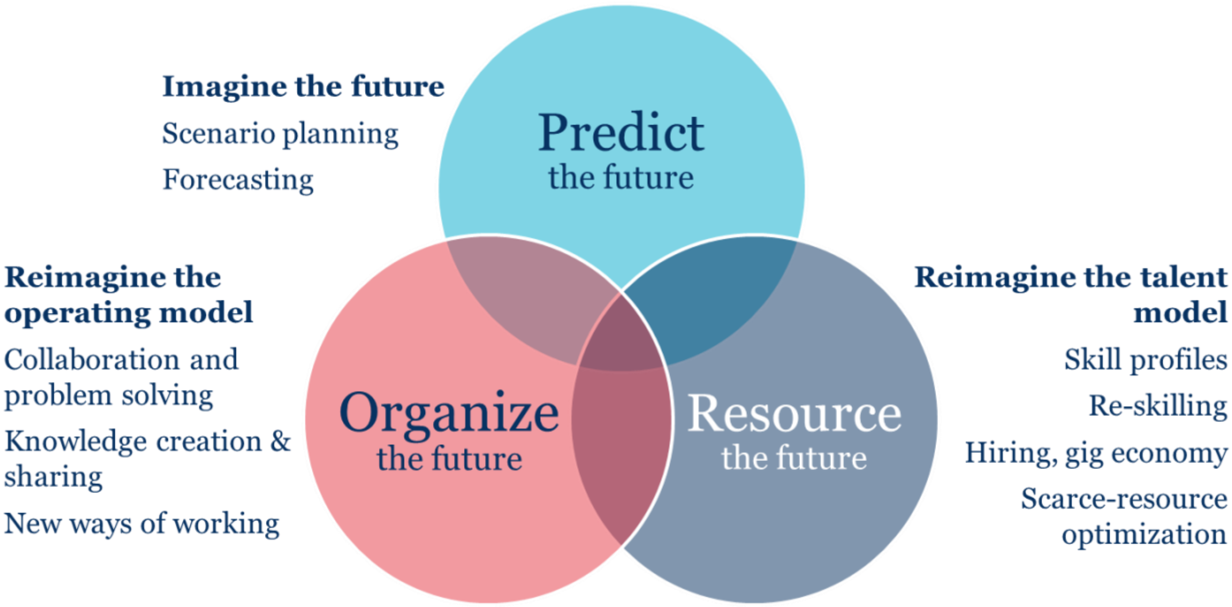

Most management teams don't fully know how to organize through superminds. They will have to learn. Chief Information Officers, Chief Digital Officers, and some Chief Marketing Officers might be inclined to own that agenda, as might some innovation leaders and possibly even some daring Chief HR Officer. But it is unlikely that they will have what it takes, at least at first. CEOs should encourage the emergence of new roles to build capabilities to structure, incentivize, inform and help the collaboration of a network intelligence.

This will take some serious rewriting of job descriptions, from the boardroom down. Leadership's role in a supermind is not that of a traditional architect: superminds aren't inorganic matter that can be 3D printed in a deterministic, tightly controlled way. In the words of former MIT Media Lab's Joi Ito, this is akin to a gardener's job, as seen in a big botanic garden or a precious vineyard.

	From	To
	Extremes of “laissez faire”, rigid organizational design, culture building	Intentional yet “organic” design & management: seed, nurture, prune
Structure	Org charts, workflow, “knowledge work” self-directed poorly-understood collaboration	Clearer understanding of actors, knowledge flow, and related decisions
Information intake	Few centralized structures (e.g. strategy, marketing, R&D) and individual networks (e.g. social media) often siloed or eco-chambers	Broader, increased-relevance, more efficient and systematic input
Collaboration flows	Strong in pockets, few common platforms but unsystematic ongoing improvement	Cohesive strategy
Incentives	Largely extrinsic incentives - individual KPIs that reflect P&Ls rather than behaviors	Mix of intrinsic and extrinsic incentives to energize key individuals

Starting with CEOs, future leaders of complex organizations grappling with hard problems will spend time seeding, nurturing and pruning parts of the enterprise supermind. This will require skills at the intersection of many disciplines. For instance, computer and data sciences, process design, design thinking, social sciences and others, such as social network and online community management. New jobs might emerge in the process. The enterprise cognition opportunity might even reshape the C-suite of companies that redesign themselves not "digital first" but "network intelligence first." Most likely, our children will develop some of these skills as an intrinsic part of their life, just like Millennials have grown up constantly connected with the world via the web.

Future leaders of complex organizations grappling with hard problem will spend time seeding, nurturing and pruning the relevant parts of the enterprise supermind: its structure, its information feeds, its collaboration flows, and its incentives. This will require skills at the intersection of many disciplines including computer and data sciences, process design, design thinking, social sciences and others, such as social network and online community management.

At this point, the question is: can your organization ignore the opportunity to turn into a network intelligence? If you believe the answer is no, then the rest of this document is worth reading.

How to create a supermind

In this section, we delve into the individual modules of the framework. Among others, we describe a practical case study that illustrates a real-world application - massive enterprise-capability building. Then in four separate "technical notes" sections we share the detail to help practical application. We also provide a simple methodology playbook for design projects, including phases and activities that help teams deliver the blueprint of a networked intelligence.

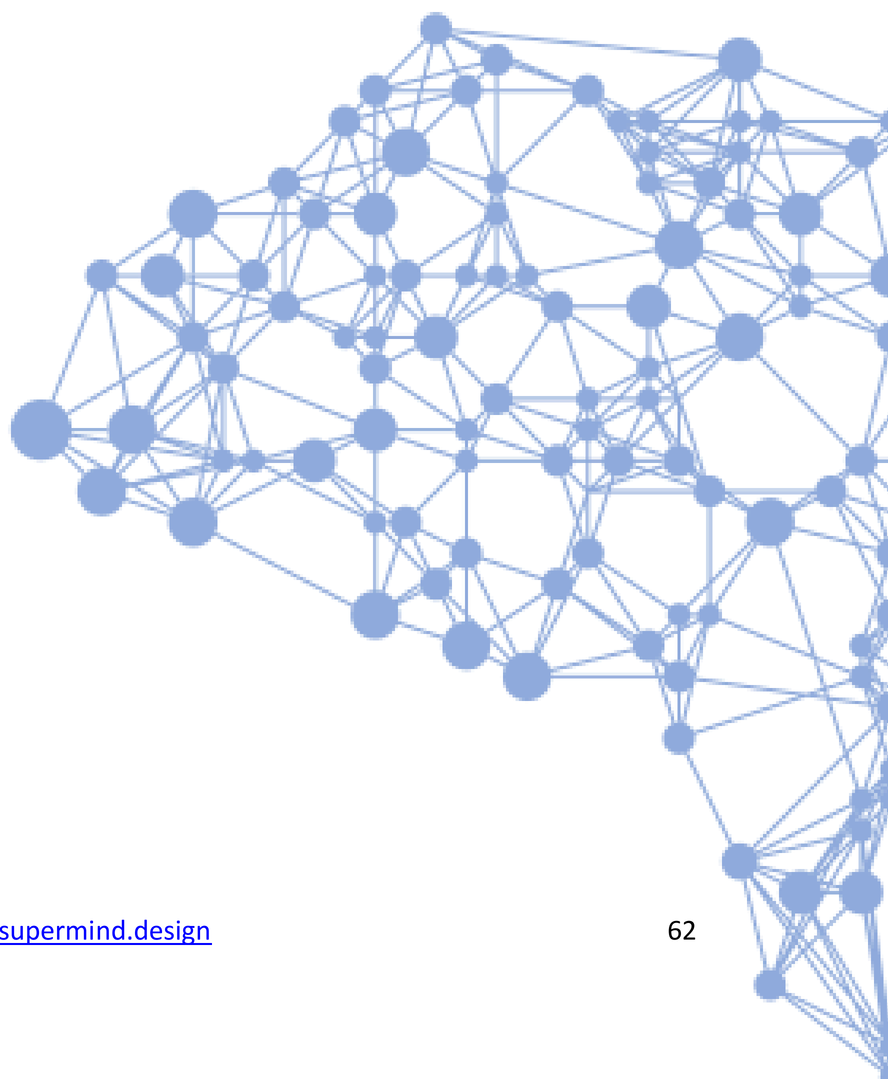

A typical problem

You're a CEO grappling with the fact that you're not getting enough innovation from your organization. You're not obtaining *big I* innovation in the form of great ideas that can reshape your future. Even more pressingly, your enterprise system (your people and your partners) isn't delivering the much-needed *small I* that must permeate everyday deliverables and work, the things that pay the bills. That's where digital innovation should be providing a much-needed lift. Yet despite the grand fanfare of AI and robotics, you're not seeing enough happening at scale.

The quest for new actionable ideas and their execution: most companies struggle with it. They want the ideas to be both fresh and unconventional, but also viable and executable – and they're often frustrated. R&D and innovation teams are typical recipients of that mandate – a mandate that often emanates from the very top of the company. Most companies can't keep ideas flowing fast enough at the ground level, beyond a few centers of excellence and often-insular R&D departments.

This problem isn't new. However, it is now compounded by the speed of change (more technology, more new practices, more funding for disruptors) and the fact that much innovation stems from combining ideas that sit in different parts of the company or even in the external ecosystem, including engineering, digital, customer support, suppliers, and academic partners.

In such a situation, your reaction would likely be to say, "I don't have the right talent." And you may be right. The collective brain that represents your enterprise may not have enough of the right neural nodes. Even so, you might not have maximized the power of your existing talent as a *network*.

A network intelligence's architectural blueprint (or its garden layout)

Let's explore what modules need activation for a connected intelligence to emerge and solve these challenges and others in a new way. Much research, such as MIT's collective intelligence's *genome*[49] has explored the topic for several years and provided good theoretical foundation. Recent research[50], based on extensive analysis of previous work, shows that the collective intelligence in groups depends less on the individual intelligence and skills, and more on the complementarity and cohesiveness of the members, as well as on the process of collaboration. While this research focused on individual groups solving specific and measurable tasks, its implications are crucial for broader-scale networks – including networks composed of people and increasingly-intelligent information technologies.

The framework offered in this document is comparatively narrow and practical. We will aim at the efforts of structuring processes, people, organizations, and technological solutions, and where workstreams are identified as part of projects and programs. The four elements of collective IQ that we introduced earlier (illumination of structure, incentives and goals alignment, information feeders, and collaboration environment) are a good starting point. Each of them is a lever to steer the system towards higher performance. They allow us to take an *active* management role – for instance, by ensuring that more powerful network structure and incentives are designed, or that the right information flows are supported. We will explore each in turn, and in the likely sequence of initial implementation.

It is important to note that this method could also apply to designing more traditional command-and-control business processes, in addition to established digital transformation and Lean and Six Sigma approaches. But while conventional workflow already has enough practices (and practitioners) able to deliver enterprise projects, the gap to designing networked intelligences is much wider. That's what we will focus on.

This implies that a smarter enterprise relies on the intelligence of its individual employees (including expensive new hires in new disciplines). But its ability to display better collective IQ also rests on:

- the aptitude of groups of people to complement each other's deficits when working together, which is a result of the network available
- the incentives to collaboration
- the type of information feeds that nourish the system and the interaction platforms that exist
- the organizational capability to make those individuals learn so that the *quality* element of the collective-intelligence equation increases.

This implies that many of your employees should master the new tools and ways of working that enable *inter-operation* between experts from many areas. The chasm between early inventions and widespread adoption of world-changing innovation originates there. The rank and file is frequently unable to get together and implement ideas at scale. The traditional way to solve the problem is to research best practices, call in consultants, and blueprint the transformation. However, these are *ad hoc*, one-off efforts that in today's world become obsolete in a matter of one or two years. They rely on the (highly variable) quality of those consultants and the employees who interact with them. And they struggle with the fact that it's not easy to change networks of people, existing knowledge, and new knowledge. They aren't easy to rewire to act more like a human brain and less like a workflow.

We need better ways of up-skilling and re-skilling the thousands of people who work in dead-end jobs before automation wipes them out while simultaneously ensuring that we reinforce capacity to implement our innovation plans. That means more than just sending people to training courses. It means activating a networked intelligence that enables people to complement each other's gaps, and to surface the knowledge that many people need.

Organization and information science have attacked this challenge for decades, but the gaps are still visible. Most leaders lack a formal way to design and operate hybrid human-machine networks and their knowledge flows. Marketing leaders, especially those in social media and digital marketing, are the closest to having proper management tools to harness these networks. Even they must put up with comparatively primitive measurements and techniques.

All those *neural nodes* not formally connected through organization charts or workflows are naturally loosely connected at best. They're a disorganized ecosystem.

As hinted at in earlier sections, we can think of a modular design for the architecture that enables it. At the most basic level, this architecture connects the right nodes and ensures feedback networks from the broader ecosystem and the network are established.

First, find the best minds for your challenge, both human and machine. Understand how they're connected to each other, and to the world around them. If they're not connected (if the individual nodes struggle to, with little effort, find each other and engage meaningfully) they will not be able to build on each other's ideas. The network will still look like a foggy cloud, with only few local nodes connected to the closest ones.

Second, engage the most important actors, and energize them so they give you their *cognitive surplus* - their ability to go above and beyond what they do today or more basically, they redirect a bit of their time toward your problem. You can do that by designing and enabling the propagation of incentives that align with meaningful goals. Those goals could be explicit, or implicitly embedded into norms and culture.

Third, build a mechanism for feeding relevant, usable information into that system composed of *nodes* and *energy*.

Fourth, **enable collaboration at hyperscale** by eliminating friction in the interactions. Make nodes process inputs faster and smarter. And make them learn. You can do that by using the right tools and methods to get large groups of people and machines to work together.

These four steps activate the supermind and are the four areas that – if intentionally designed and managed – deliberately harness that network of connected intelligence. Each of them can perform specific collective cognitive tasks. For instance, knowledge capture and dissemination within specific parts of the network enables collective remembering and avoids reinventing the wheel, while external listening posts enable hyperscale sensing.

It is important to note that while these steps are initially somewhat sequential, there are a lot of iterations between them. The identification of new nodes happens continuously, even and especially after the right incentives, knowledge feeders, and collaboration platforms are in place.

The schematics below provide more detail about the necessary modules. Behind these blocks are dozens of methods, practices, and tools that could be used today. These range from social/people network analysis tools, to collaboration technologies (think Slack or Zoom), to behavioral science (think Kahneman's *Thinking Fast and Slow*[51]) to AI's natural language processing. A diagnostic of an organization's readiness in these areas is feasible *today*. And designing systems that leverage superminds is being attempted already, for instance at MIT's Collective Intelligence Design Lab.

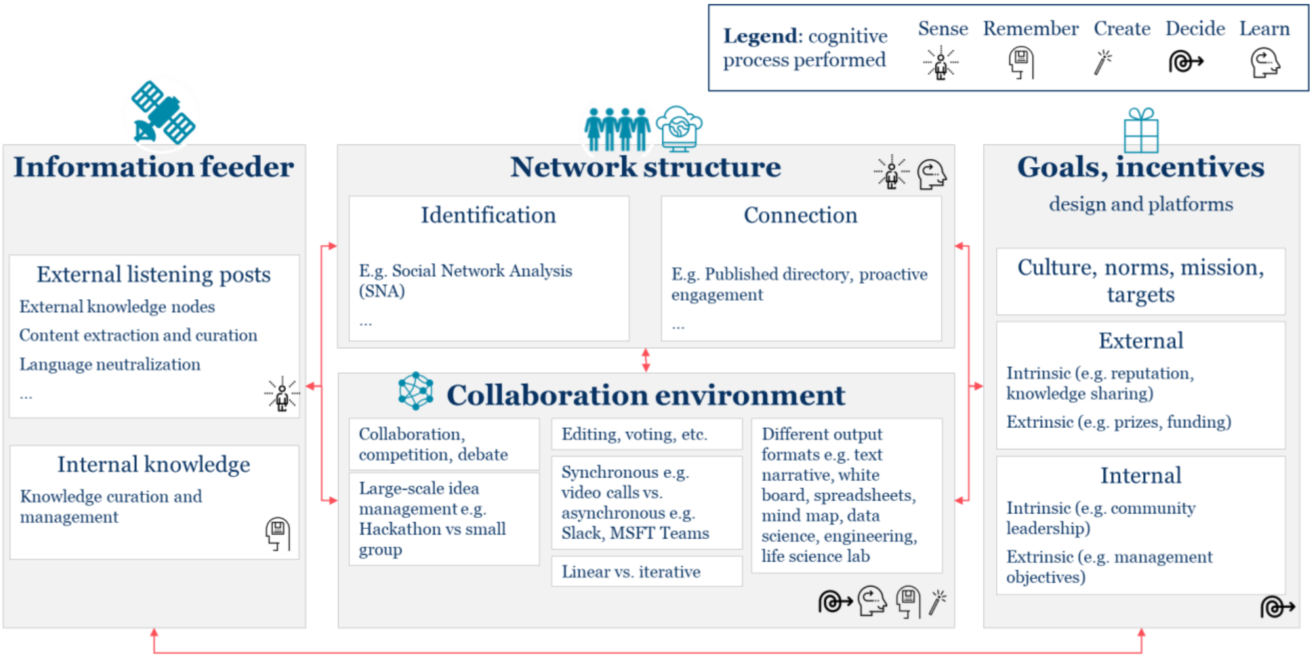

Part of this system is constituted by people. The remainder is made up of machines. Some are simple sensors, others come from the four Cs of AI. Together they support the knowledge flows and exponentially improve how the world does four things:

1. **Connect** entities, both people and machines. *Example:* helping pinpoint the right nodes in the network and making them discoverable across any distance to each other through search.
2. **Curate** knowledge. *Examples:* semantic searches and computer vision that identify the most relevant content (whether text or visual) and cluster it for people to process it more easily; *automated data discovery* which combs through multidimensional data sets to identify patterns that humans would take too long to evince; combine knowledge as in the field of generative design where human intelligence completes the selection of hybrid ideas generated by machines.
3. **Collaborate** across those entities. *Examples:* natural language processing that automatically translates content; machine learning that optimizes video and voice transmission.
4. and **Compute** any other prediction[52] that can either trigger an automated response or be passed on to a human or network for validation. *Examples:* (for instance, using machine learning to determine which participant is worth rewarding; detect spam or fake inputs; or

identify strictures in high-volume process flows). that can either trigger an automated response or be passed on to a human (or a network) for validation.

Machines and people work alongside each other in networked intelligences. The table below visualizes the common types of interactions and some examples, showing that they are two-way streets.

		Initial input…		
		Human	**Machine**	**Both**
…and subsequent processing	**Human**	Traditional collaboration, now at network scale (e.g., WeFarm, enabling automated translation and matching of experts)	People verifying the quality of machine insight, or “finishing the job” (e.g., exception management on items that machines couldn’t process – e.g., Captcha, industrial assets alerts)	Crowdsourced sensor data used by citizen scientists (e.g., GPS OpenStreetMap, Waze)
	Machine	AI training, typically through annotated data sets generated by crowds (e.g., Zooniverse, Amnesty Decoders). Clustering of human language inputs. Automated translation. Automated classification of visual objects (e.g., X-ray footage, Syrian Archive[53])	Traditional automation, now at network scale	Crowdsourced sensor data used by machines (e.g., Google Maps time estimates) Computer vision sifting through mixed-source imagery to enable better agricultural decisions (www.OneSoil.ai) OpenAI GPT-3 predictive text generation
	Both	Inappropriate web content moderation (e.g., Facebook)	Situations where data sets are insufficient for machines to operate independently. Swarm AI predictions (e.g., Unanimous AI), computing likelihood of human predictions. Early Warning Project (Good Judgment Open[54]) to predict mass atrocities.	Generative design software (e.g., Autodesk for engineering design) “Deepfake” content (audio, video)

Most of the quadrants use AI’s 4 Cs. Designing processes that guide the flows above, from machines to people networks and vice versa, is likely to become a new skill in high demand. The chart below highlights the vital roles AI plays in each module of a networked intelligence.

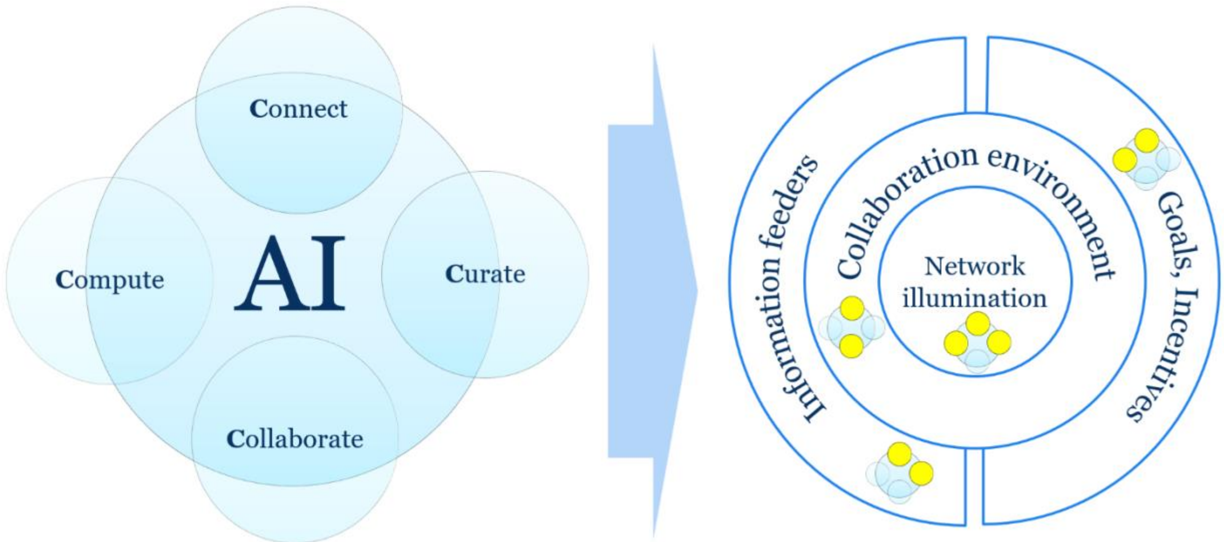

The next four chapters will help you understand each of the modules, their value, and the tools and techniques that can help leaders (enterprise, public, or unconventional) harness the power of a supermind.

The following table describes the more detailed interventions that a supermind project could carry out. Each of the individual points is explored in more detail in the technical notes section following each chapter.

Identification and analysis, and enablement of visibility for **network nodes.** Possible interventions:
1. **Have we identified and analyzed where the key nodes are, and what their interaction seems to be? Does that interaction have gaps that we can bridge?** a. Do relevant managed communities exist, which could help us harness ready-made collective intelligence? b. Do we need to illuminate the network, especially the key nodes and related connections – and then visualize and analyze them? c. Is the network structure ideal? d. Do participants contribute enough? 1. **Do we understand the characteristics of those nodes, especially (but not exclusively) if human? Have we spent resources improving those characteristics, or improving the mix in the most important clusters?** a. For instance, sex, cognitive and interaction styles, experience and skills, location (remote vs close by), engagement level, trust b. Have we deliberately chosen the level of diversity that we want, or are we defaulting to what's in place? 2. **Are roles clear?** 3. **Are "soft rules" that govern the interaction between nodes appropriate?** a. For instance, culture, norms, practices, values, purpose, psychological safety b. Are we helping valuable peripheral nodes getting more embedded in the network? 4. **Have we mobilized the most critical nodes?**

Design and implementation of goals, and appropriate system of **incentives for collaboration.** Possible interventions:

1. **Do we understand the *intrinsic* motivators** that may work for the nodes, so we can channel their cognitive energy?
    a. Among others, **have we explored the value of a supermind in providing a sense of mastery and "flow"**?
2. **Do we use network influencers** and those with authority or clout to energize the rest of the network, for instance by participating in conversations or visibly endorsing people?
3. **Have we optimized task design to avoid unnecessarily taxing cognitive resources,** and instead amplify the flow experience?
4. **How do we give credit to people,** as recognition is important for intrinsic motivation?
5. **Have we understood the implications of** our design of interactions, specifically with regards of the use of **competition or collaboration**?
6. Are we able to muster enough resources to provide **extrinsic motivation** at scale?
    a. Is that the right approach, given your objectives and the culture/values of the network?
    b. With regards to extrinsic motivators, **have we explored the opportunity for crowdfunding**?
    c. Have we explored the use of technology to enable the creation of trust in the network?
7. Do we have the **analytics** engineering and visualization needed to provide feedback to the organization and individuals?
8. Have we created mechanisms that make individual nodes feel like they're going to encounter their counterpart in the **long term**, hence encouraging trustworthy behavior and discouraging freeriding?
9. Have we explored the **business case** that justifies funding for extrinsic motivation?

Enhancement of relevant and ubiquitous **information feeds.**
Possible interventions:

1. **Are we surfacing the output of collective intelligence through conventional or advanced surveys**, taking advantage of web-enabled capabilities?
2. Have we created an **explicit description of the semantic space** of the concepts that we are exploring, so that we can both search for that language, as well as the people who use it – and the people who are connected to them that might extend it?
3. Have we established a comprehensive **net of listening posts** that continuously mine that semantic field and the related people network's input?
   a. Specifically, and when needed, have we created a **network of ground level "sensors"** that generate "ground truth" irrespective of what's already available on the web?
   b. Is our **knowledge management system** fed by, and feeder of, the above?
   c. Are **Business Intelligence (BI) systems** also woven into this flow of knowledge, as another source of ground truth?
4. Have we created a **machine-augmented human *curation* layer on top of the above**, to filter and route the right input at the right time?
5. Have we specifically created feeders **for scientific and patent-related literature** if appropriate?
6. Are our **learning and development** (L&D) systems being fed by (and become feeders of) this system?
7. Are **available resources allocated** so that the flow of knowledge is mined effectively and efficiently?

Creation and management of an appropriate **collaboration platform**, from new technology to new practice. Possible interventions:

1. Have we attempted to reduce the routine, non-value-added repetitive tasks that burden network members, by **automating part of the flow of work**?
2. Are existing "crowds" being harnessed through **crowdsourcing platforms**? Similarly, are **freelancer markets** being utilized appropriately?
3. Are the teams using the right **idea management platforms** and services?
4. If useful, is **citizen science being** used?
5. Are the right **unified communications** technologies (particularly important for synchronous collaboration) established and adopted? Are the appropriate practices understood by participants?
6. Is **enterprise social media** or other community tech used as needed?
7. Are the relevant **techniques, and methods for text and language-based collaboration**, used for the networks?
8. If **debate** and dialectic are needed, are specialized **technologies** used to facilitate the flow of the dialogue?
9. If **predictions** are the focus, are the right methods and tools used for it?
10. Are the appropriate **collaboration** environments available for the specific needs of **data** and **information technology**?
11. Similarly, and where needed, are the right tools used for the needs of **engineering**, or **life science**? If some of the work of the supermind is **scientific research**, are the right tools used for it?
12. Is there a way to use blockchain protocols, potentially even in the form of Decentralized Autonomous Organizations (DAO)? (Section here)
13. Are the right **learning platforms** used to facilitate collaboration at scale around learning materials that help the network evolve?
14. Are we **democratizing** the development of code, and generation of data-driven insight, empowering many people throughout our organization to contribute more to the network?
15. Should we consider **gaming**, **augmented** and **virtual** reality in the design?
16. Is the right **organizational design** in place to make the above possible?
17. If resources permit, is the use of **advanced natural language processing** piloted for any of the above?
18. Is **cybersecurity** important in the design and operation of the platform?
19. Are the **right management practices** being used to facilitate the work of nodes that connect with each other in a many-to-many, two-way conversation? Are teams encouraging the right flow of work, cadence, feedback loops? Are they practicing mindfulness?
20. Are our teams, at different levels, using **data-driven insight** related to the functioning of collaboration platforms, to inform activities and allocate resources?

Module 1: the "who". Illuminate the neural network

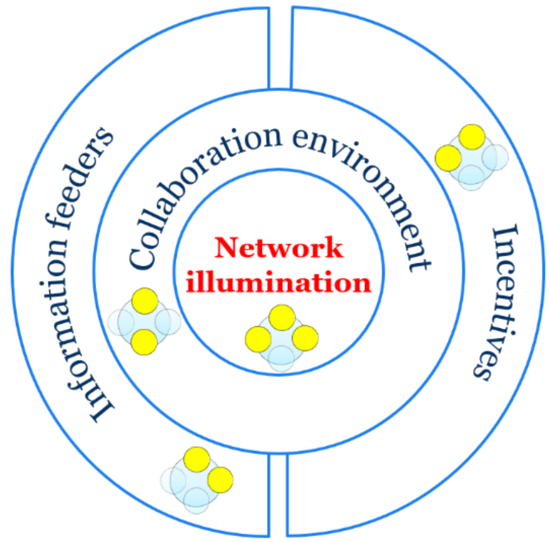

Networks structures drive impact. They contribute to the performance of enterprises. They help people succeed. Some networks, for instance those created during college years, are conducive to entrepreneurship – many startups that went on to change the world, from Microsoft to Facebook and Google, to Snapchat and Reddit and WordPress, were started in college. Ironically, they all share at least one ingredient for success: they harness the power of people's network by enhancing their ability to collaborate.

Bridging holes in networks seems to enhance both business performance, at least when innovation is important, and individual people's performance at least for people in roles where orchestration of new ideas yields value.

Business opportunities or cost optimization can spawn from bridging structural holes in communication networks. Organizations with numerous brokers can harness resources and business insights that can drive performance and help maintain a competitive edge. Research has highlighted the importance for entire companies. For nascent firms, where growth is integral to success, diffused communication networks were associated with 10% higher revenue, firms can better identify growth opportunities or cost reductions[55], better optimize resources (financial or strategic) that are necessary for growth and help to propel performance[56]. Bridging structural holes within communication lines was associated with about a 37% increase in market share, and network brokerage enhances internal firm capabilities which also drives performance[57].

Employees who have better access to new, changing or even contradictory information are given an advantage in staying on top of business trends and can generate valuable ideas. Acting as the pivot point of knowledge flow, brokers themselves can improve their performance and obtain promotional or compensation benefits. Research results show the possible impact[58]. The more concentrated a manager's contacts were within their network, the less positive their annual performance was rated. Network structures drive promotions and compensation among employees: managers brokering connections across fragmented groups had a 68% chance of receiving a promotion or salary increase, compared to 28% for managers with a limited circle of densely connected colleagues. The impact was strongest among senior employees, showing a significant negative relationship between the average salary of senior managers and the degree of diverse connections in their networks. This negative effect became even larger among Vice Presidents and Directors. The relationship between brokerage and an employee's propensity to innovate provides additional insights: managers with more expansive (broad, diverse) networks produced ideas that had a low probability of being discarded (14%) by upper management; managers whose networks were more constrained were less likely to produce valuable ideas and had their ideas were discarded by upper management 43% of the time. (As a sidenote, social networks seem to also have a significant impact on people's health, with significant impacts registered on obesity and happiness, for instance[59].)

This section lays out alternatives to traditional "box and lines" org chart or workflow design and complements existing team-dynamics optimization practices. The objective isn't to do away with the traditional organizational structure scaffolding, but rather to **overlay them with network-enhancing structures and capabilities**. In a way, this is like the way our brain's cortex envelops other parts of the brain connected with the nervous systems' endings, whose behavior is more deterministic but critical to the functioning of the body. The cortex enhances the value of other brain parts thanks to its ability to combine sensorial input and memories, then use its flexible decision algorithms to make decisions and learn.

Individual nodes of a collective intelligence network (whether people or machines) need to sense, remember, and then create, decide, and learn by *engaging other nodes*. For such a system to work well, its structure – how its connections link to each other – should be engineered deliberately. Those connections, those nodes, and how they interact, should change depending on the job that is being done: executing from a fix script should leverage a structure that's different from exploration in chaotic environments, for example.

That's not how most enterprise networks are typically designed today. Org charts and traditional organizational practices deal with connections in mostly hierarchical, deterministic environments. Those methods struggle to identify the *actual* flow of information and knowledge, and the formation of new ideas that happens more fluidly through networks of people. Or consider the limitations of conventional organizational analysis and design in the case of the so-called "*invisibles*[60]", people who don't hold visibly significant positions in an org chart but are nonetheless core to the propagation and formation of ideas. And as we noted earlier, the "weak ties" in personal networks are extremely important in fostering innovation and shaping culture.

Team and organizational design have made significant progress in the last decade, moving in a direction that hints at the importance of harnessing people's groups differently. Take the example of Agile methods that have become standard for innovation execution in technology organizations. In Agile, "*pods*" or "*squads*" (i.e., small productive units) exist, where a product owner collaborates with developers and subject matter experts who complement each other, all orchestrated by a scrum master. They're collectively able to provide output with other pods', and jointly contribute to the final product. The same needs to be done with the most important nodes in a connected-intelligence network. Combined with more hierarchically driven (yet leveraging collective smarts) software code reviews, these methods are foundational to much of the technology progress made in the last years. They have spawned additional organizational ideas like IT analyst Gartner "fusion teams" ("cross-functional teams that use data and technology to achieve business outcomes"). And indeed, open-source movements like Linux routinely apply community-based, collective skills to software code reviews. The combination of these organizational innovations has changed the world.

There's reason to believe that agile-type practices and related "team of teams" are just the beginning, and that the future must borrow from network science a lot more. In one example, a cohort of mid-tier employees was instrumental in the cross-selling of financial products in an investment banking firm, by enabling smaller transactions across different (and otherwise siloed) product groups and effectively serving as integrators across org chart divides[61]. Without dedicated support and recognition, many of these valuable staff members were leaving the firm. Once their role was highlighted, incentives and support structure were put in place and the overall organization benefited.

So, let's deliberately try to harness the network. First, network **"***signatures*", measured by detecting the flows of communications between parts of an organization, have been correlated with the ability to obtain a specific outcome. Think of these concepts in the context of the work of R&D, marketing, sales, and supply chain, for instance. And consider how scientific ideas in radical departure from existing knowledge are often created at the edges of established networks, where people connect with other networks and cross-pollinate[62]. While some of these signatures may be generalizable, some may not, requiring the identification of enterprise-specific "*success*" network signatures. Traditional organization design also struggles to determine what the right amount of networked people, (the right critical mass), is. And as a result, many enterprise-wide efforts end up being over capacity and lacking accountability.

Network analysis has been extensively used to identify the behaviors of the most successful sales reps in large enterprise, business-to-business sales[63]. Connections with senior leadership of their own firm, and *steady* (not just broad) connectivity with sales support resources have been found to be important determinants of success in the firms that were studied.

In general, network analysis conducted in large organizations has shown how traditional organizational and workflow design methods miss out on that large parts of the actual functioning of organizations.

For instance, **innovation** is more likely in networks with much connectivity between different (and diverse) groups, so that trust is established, ideas are combined and "pre-tested" by people with different backgrounds. Importantly, nodes that connect those groups and take the lead in energizing the network will vary over time, with new actors achieving more central status depending on the situation.

Paying attention to network can indeed improve the *process of innovation*. Innovations travel in non-linear, iterative ways across several stages: from exploration and discovery, to shaping and development, to diffusion and embedment into the fabric of the surrounding organization. Most people only span a few of those stages individually, but their network may reach far beyond that. For instance, an exploration network tends to be broad, made of many typically unrelated spaces. Diffusion of prototypes and identification of implementation opportunities instead require the ability to tap into local networks, and often in multiple parts of each local network, to trigger adoption at scale. Embedding into the fabric of work, especially in large enterprises, also requires people who have the network to alter systems of incentives (through, say, sales incentives or budget cycles) and dislocate legacy processes and systems. Most companies don't inventory their people like that. As a result, they lack some of those roles, or neglect priceless people who hold those roles.

In contrast, efficiency-driven organizations have more regimented connections. Think of the G&A (general and administration) functions (such as Finance and HR) that support individual businesses, where business partners orchestrate the communications of tightly managed workflows and respective people. But also remember the example of thinner, less internationally diverse scientific research papers during COVID-19– a pattern prompted by speed.

It is also likely that centralized networks limit the creativity of the ideas they generate, partially because they receive fewer ideas, and partially because of the potential bias of their members which filters out ideas that don't conform with the group's inherent views. Groupthink is real a possibility in those cases, though it may come in exchange of faster decisions and execution – a real tradeoff invoked by many leaders as a justification for close-knit groups.

This said, centralized networks whose peripheral nodes keep strong relations to the outside world may be able to combine speed of execution with reasonable openness to new ideas[64]. And alternating close-knit structures for definition of the objectives and coordinated execution, with decentralized exploration of problems and design of solution prototypes, may also offer speed *and* creativity. Some research indeed suggests the value of alternating centrally-orchestrated collaboration with delegation of activities in peripheral groups. The same research indicates that this strategy works well for both complex as well as simple problems.[65]

Think of the applicability in organizational units that ensure that beyond their decision-taking leaders, an open worldview is developed through other members of the team.

It is also possible that by identifying and analyzing the information flows from the key nodes in the network, senior executives can better understand the ebb and flow of information and collaboration energy across organizational groups. As well, they can better visualize a sort of circadian rhythm that shows the ebb and flow of that energy over time. Identification and monitoring of external network nodes can provide early signals useful in detecting shifts in the market and possible competitive developments.

Network structure and performance isn't just about creating and implementing new ideas. The use of network analysis is also helpful in **understanding the effectiveness and efficiency of working teams** – for example, in collectively generating the right response for their clients[66]. These methods also help in understanding the overall degree of collaboration between separate organizational units which are part of the same end-to-end process (think order-to-cash processes); or in shedding light on the "*rooting*" of new entities into a pre-existing network (useful for post-merger integration), or in the equivalent rooting of new critical hires (such as salespeople in complex business-to-business environments). In the technical notes for this chapter, we will identify the characteristics that individual nodes can have, and how to avoid inadvertent neglect and imbalance. Without going into details here, individual nodes can be differentiated because of their role in enabling specific supermind modules (e.g., such as having a primary role as an information feeder or being a collaboration facilitator of the collaboration) or, not unlike the Agile example above, because of their skill (e.g., as a product innovation owner, subject matter expert, or coordinator) or behaviors.

To illustrate visually, consider the following chart. Network analysis tools, such as those described in the "How to do it" section, can identify how many communities exist in a specific organization (or network) for each quadrant. The objective is to have enough critical mass in the quadrants that are needed for a specific activity. For instance, some companies don't have enough communities in the top-right quadrant, that is they don't have enough groups that connect across large networks e.g., across the enterprise, and where people's role varies over time (for instance, where their leadership network signature "oscillates" between being central and being more peripheral). Those companies may struggle with generating enough serendipity and combination of varied ideas coming from diverse networks.

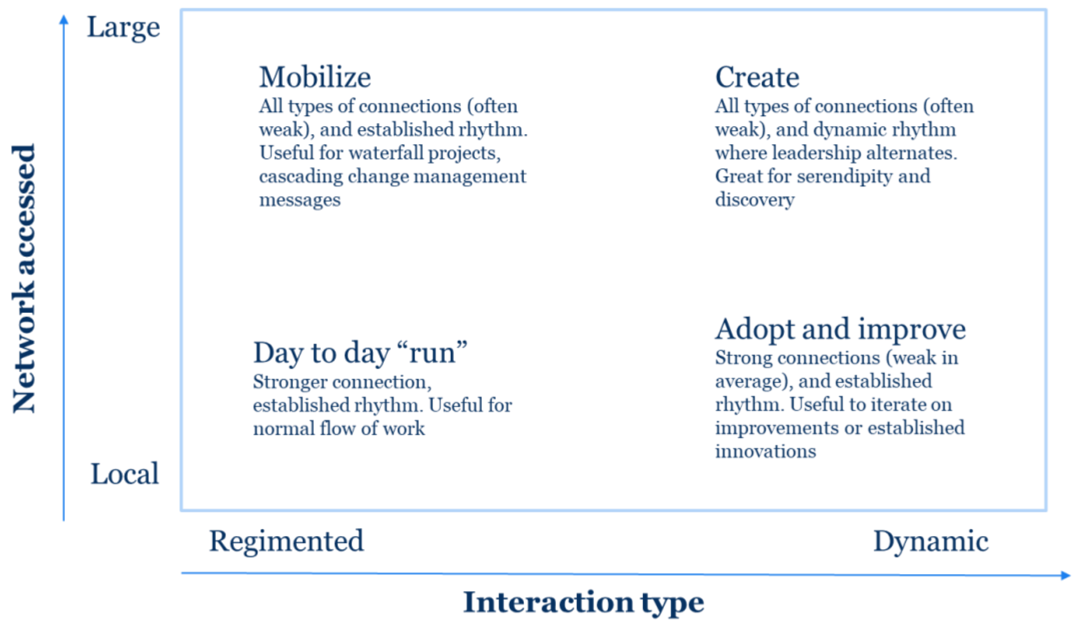

There's strong promise in this approach, but **where does one start**? We will describe these aspects in more detail in the technical notes section, but some of the basics are highlighted below.

The first step in achieving this structure is to identify the network components and start connecting them with each other, in an intentional manner as opposed to the current, largely accidental one. This activity can generate a lot of value with comparatively little effort, and it requires making the nodes discoverable to each other in a way that doesn't overwhelm them and helps them to connect even through basic, conventional means. Think about startups meetups, or company yearly kick-off meetings, which serve the same function. The explosion of professional networks such as LinkedIn shows the importance of identification and discovery of network nodes.

That's especially important for people, who rely on natural language and can flexibly sense their environment, then and hence connect with interesting (human or artificial) nodes to extend their senses and memories. Our problem as humans is *scalability*, that is specifically the inability to process large amounts of data, because of the time it takes people to read natural language. We can't hold more than one meaningful conversation at any time. Also, humans aren't naturally able to maintain connections with large numbers of other people. That observation was first made by Robin Dunbar[67], who estimated at 150 the maximum amount of people an individual can be meaningfully connected to. This limitation has fueled some the criticisms levied at Reed's Law that we mentioned earlier, which states that the value of networks scales exponentially faster than Metcalfe's Law, because groups add connectivity to the peer-to-peer individual connections. That's one of the reasons why (over for several centuries, and despite the development of advances in communication technology,) the size of army units has floated around the number 150: in the imperial Roman army, a unit was about 120 soldiers. Today it is about 180[68].

The flow of ideas has traditionally been traditionally constrained by individuals' networks. That is, if I want to access the ideas of 1,000 people, I need to connect with a smaller group of nodes, and hope that they will process (filter, curate, or relay) the right information to me, after capturing it from their respective, broader networks.

Of course, access to ideas has grown exponentially, and with it the collective intelligence of our societies. From our ancestors' oral traditions to the invention of writing, from the first libraries to the printing press, from radio to smart mobile devices, access to knowledge has changed beyond comprehension and will continue to do so. But the recent COVID-19 emergency shows how, in times of extreme pressure and for the sake of speed, otherwise highly-networked scientists reduced the size of their international networks and the generally the size of the groups involved in doing research together[69]. Is that tradeoff warranted?

Networks **– and the skills embedded** in them as they enable different abilities to come together in individuals exposed to them, regardless of their individual expertise. For instance, data scientists and supply chain experts may need analytics translators able to create common ground between the two fields. Doing this helps identify supply-chain use-cases where data driven insights can be found, even though the translators themselves won't be able to do the actual coding required or to provide specific functional requirements.

In the absence of formal translators, individuals from each discipline will be able to better interoperate with one another if they have a certain level of proficiency in the other discipline. In learning science, that's called a *T-shape skillset*. Its opposite, the *I-shape skillset*, reflects a person with narrow expertise who is very deeply skilled in their domain. A T-shaped individual can bridge across different fields. T-shaped individuals are the connectors and translators in large networks that hold very diverse ideas.

A generalization of the point above hints at the value of **diversity and inclusion** – in all its forms. More detailed discussion follows in the technical notes, but it is intuitive that, as Tim Hartford[70] notes, "if you already have four brilliant statisticians working on a policy problem, even a mediocre sociologist or economist may add more to your team than another brilliant statistician. If you're trying to improve your tennis game, you may do better working with a tennis coach, a nutritionist, and a general fitness trainer rather than three tennis coaches." Much empirical data, from boards of directors to IT projects to company innovation, points at the same pattern.[71]

The **emergence of AI** now offers the potential to further amplify media power by enabling knowledge to be filtered from a very large number of sources, curated to match the need and interest of an individual node, and relayed to that node at the right time.

Think about how much easier the hunt for relevant information has become since the introduction of web search engines. The first mostly computed content relevance based on connections between websites that carried specific words. Others, like Wolfram Alpha, relied on both semantics and curated, structured data derived from internet data sources. The new generation of search engines uses smarter, AI based, semantic analysis to understand language and return the right information.

Additionally, the traces that we all leave when talking to each other give AI the ability to search proactively. Unbeknownst to many, some of this happens already when we access social networks. Our curated news streams reflect not only the content we have been interested in, but also of people we follow and their networks. And additional content recommendations are also based on social media algorithms that identify relevant information based on interest patterns of similar people.

In a sense, social media connects that network already. Except that all of us have many different interests and personal connections, so the (sometime dangerously cost-effective) algorithm cannot create consistently relevant specialized feeds that reflect our complex world. The computationally efficient way to solve that problem has been to amplify information from people networks (dubbed PYMK, or “people you may know”) - that attracts immediate, intuitive reactions driven by the more primitive parts of the brain. As a result, social media ends up providing us with common denominator feeds, and often – because of advertising incentives - artificially homogenous and polarized views.

AI can help calculate the relationships between ideas and people, and potentially with the machines that generate information. It can then help curate their knowledge and make knowledge and its sources discoverable to each other, hence facilitating the emergence of connections at hyperscale. The figure opposite summarizes these aspects that will be explored in the upcoming technical notes section.

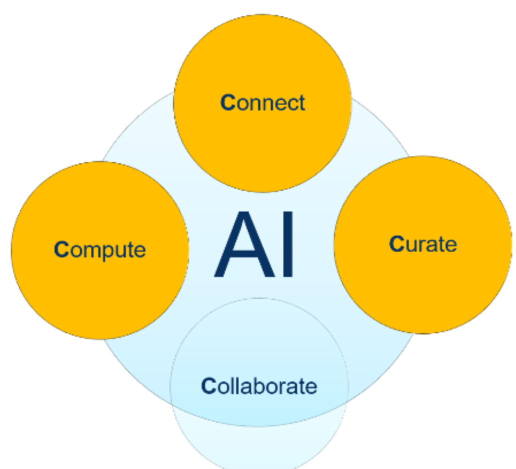

Next, let’s illustrate the art of the possible for this module through real-life case studies.

Case studies – module 1

Many organizations have used related techniques to improve their business or operating models. Social media has long relied on identifying the connections between people to optimize the display of content. As noted earlier, customer satisfaction as well as sales performance in large business -to -business environments, has been shown to correlated with the appropriate people network structures[72]. Crowdsourcing companies, from Amazon Turk to data science Quantopian, harness the power of individual's cognitive abilities at scale. Bank of America, General Motors, and Amazon among many others have pioneered the use of social network analysis to facilitate enterprise innovation. Microsoft's acquisition of social network startup Volometrix helped it develop a suite of technologies now accessible to most Office365 clients.
Using network analysis in a different way helped researcher create a model that could help predicting financial market crashes, by creating a graph that connect stocks which have highly correlated prices and volumes, and then studying the geometric properties of those networks over time[73].

A completely different type of connectivity is used by "Folding at Home", which uses the federated computing power of thousands of home devices to process computationally hard tasks (especially in the field of biology.)

In many respects, risk "crowdpooling" such as peer to peer lending and crowdsurance startup Lemonade, are advanced forms of connection between nodes of the network with the ability to allocate risk-reward through an AI-augmented platform. Peer to peer marketplaces are used elsewhere too – for instance, with the (partially crowdfunded) boat rental community Boatsetter which connects sea lovers, boat owners, and licensed captains.

To rekindle "weak-tie" connections during and after Covid-19 lockdowns, professional services firm Genpact introduced a virtual watercooler bot that automatically matches colleagues based on the strength and diversity of their network connection.

A different way to connect network resources is to create shared databases that draw from the respective communities or ecosystems – data commons that respect privacy and serve the common interest[74]. The UK Biobank uses genomic data from hundreds of thousands of people, and accredited researchers can harness that data. GenBank, run by the American National

Institute of Health is an open-access source of genome sequences that was famously used for the understanding of Covid-19's virus.

More examples are in the "How to Do It" section of the Guidebook, as well as the database on www.supermind.design/database. The next section presents one case study in more detail.

In-depth case study module 1 – upskilling at scale

A professional services firm is up-skilling and re-skilling tens of thousands of employees to make them able to withstand the dislocation engendered by digital and AI technology. The company is Genpact, a digital transformation and operations management services that spun out of General Electric (GE) in the mid-2000s. Genpact adopted a very unusual strategy to overhaul their learning and development (L&D) operations, one that leveraged the concept of network intelligence at least in three main respects as described in the picture below.

First, the aim for capability building should consider teams, not just individuals, and their collective end state. The implication is that individuals should learn not what makes them individually helpful, but rather able to interoperate in teams.

Second, the knowledge needed for the employees sits within the organization and its knowledge *nodes*, not just in external trainings. That requires a deliberate effort to identify and empower those nodes so they can crystallize and share the knowledge with others, above and beyond the job they have today.

Third, AI is at the service of the network's needs, helping curate, connect, collaborate, and otherwise support heavy computation throughout the process of knowledge formation and sharing.

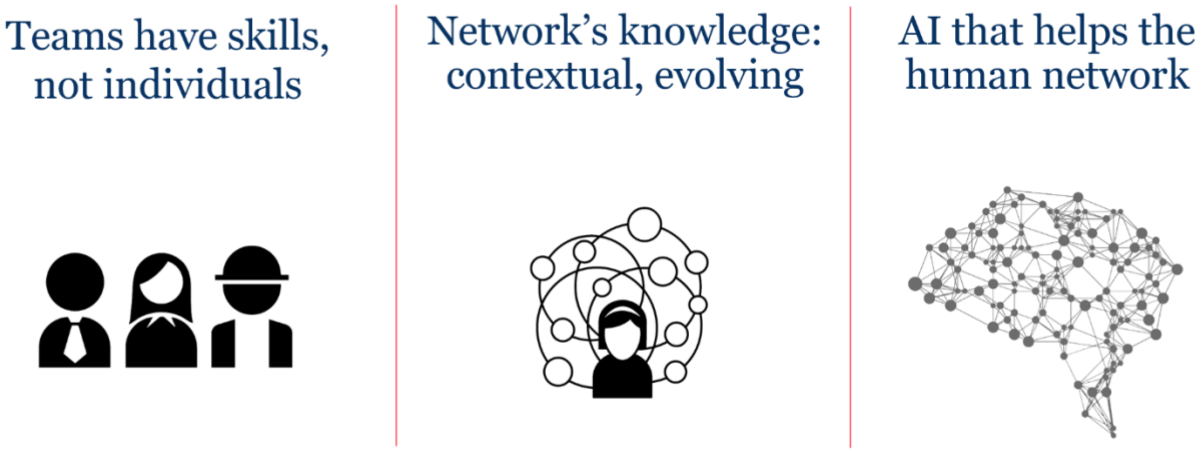

These three principles required the design of an unprecedented learning architecture, whose first pillar is indeed the illumination of the network's key nodes.

The first step deliberately surfaced the network of knowledge-rich nodes and assessed the gaps. Several key competence areas were identified, and two different exercises performed on the back of an all-employee survey. First, many employees rated themselves on their proficiency for every competence area (say, Python coding, process data mining, or insurance underwriting). Then they mentioned the names of colleagues who are go-to-resources for a subset of those skills (such as design thinking or robotics). The latter set of data was further processed through a network analysis that identified the most central employees in the new world, as well as the ones who possessed those individual skills the most.

That data-intensive effort enabled the creation of a large *skill inventory* database later used to allocate learners into *communities of practice* led by those knowledge nodes. The chart below illustrates the process.

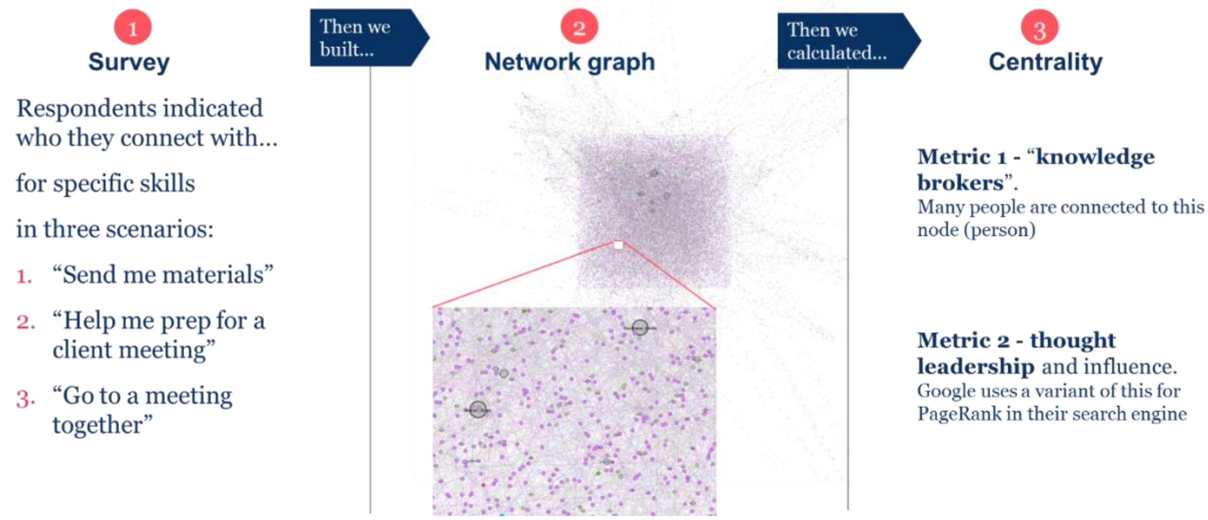

Those central competency nodes also sense and triage new information related to their area of expertise and can consolidate it better than any other employee or even the L&D team. For example, all queries about "robotics" are intercepted by those nodes, giving them the ability to understand the evolution of demands as well as generating new ideas about practical applications. The knowledge nodes for specific skills are also listening posts in broader, external communities: they sense a much broader environment, and act as filter and relay into the firm of that new knowledge. One practical use of this is their role as filters for new learning content that is exposed to learners, from new interesting articles to external courses.

However, some experts were less connected than one would have hoped. This highlighted their lower network-dexterity skills, often because they were less at ease in highly networked environments. Among those, some were chosen to act as subject matter expert gurus and they were provided support to improving their ability to *work the network,* for instance by providing them a support infrastructure for engaging with their followers, as well as data-driven *Honest Signals* diagnostics.

Some of this is normal in the professional services industry. For instance, many consulting firms have communities of practice that form around appointed leaders and serve as watering holes for consultants. However, in this case, the firm uses those nodes as deliberate actors of a tightly knit learning platform enabling the reskilling of tens of thousands of employees. An important element of this infrastructure is the exoskeleton of resources that helps the knowledge nodes to scale. From knowledge management curators, to community managers that enable smooth interactions with the followers, to advanced analytics that constantly feed the nodes with the latest information about learning patterns and learner experience diagnostics, nodes obtain a rich feedback loop from the system around them and redirect the execution accordingly.

One significant decision in implementing the network intelligence for learning and development described here was to amplify the strength of those central nodes – called *Masters* and *Gurus*.

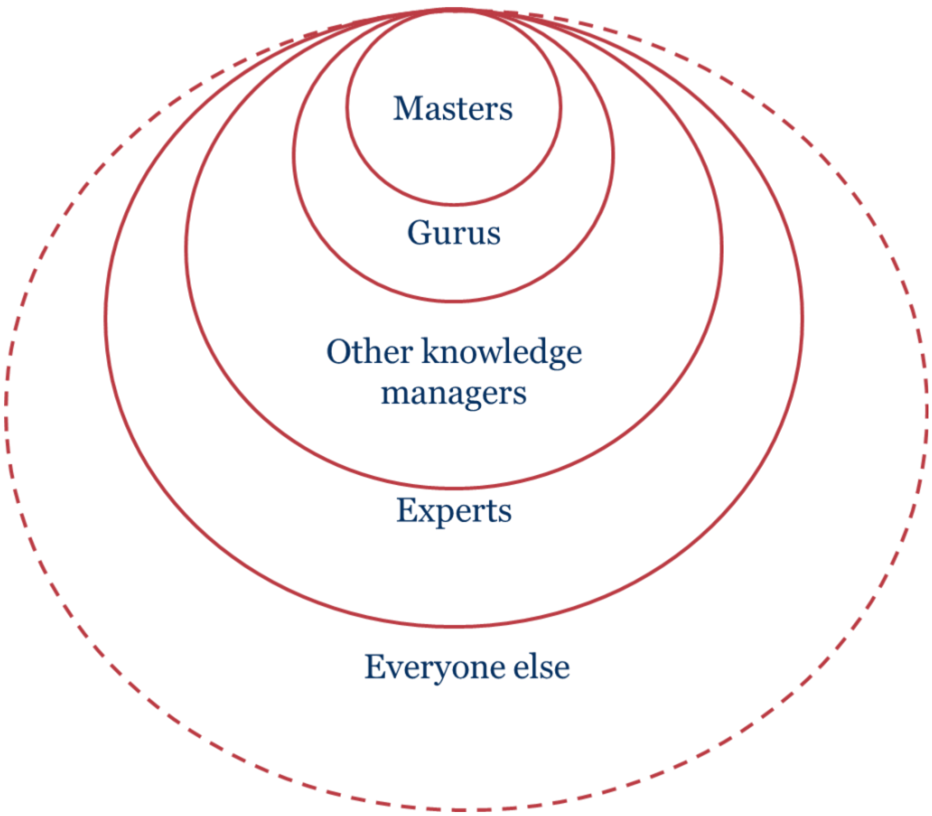

The company realized that many of those individuals couldn't, by themselves, contribute meaningfully to the broader ecosystem, but they could do so if a small group of support resources was made available to them. These knowledge managers and community managers in turn help them crystallize and curate the input of an even larger network of globally distributed experts as well as capture the crowdsourced contributions of thousands of people who unearthed useful knowledge in the flow of their daily work. This organizational design is shown in the picture opposite.

Module 1 - How to do it

[Reminder: the specific examples documented in this technical note section may be used to compile a list of "portable ideas" that enable more creativity during the design process (see the "The process of designing an augmented *network intelligence" chapter)]*

The "how to do it" notes for this module will cover the following areas so they're managed deliberately – and help teams avoid "blind spots". These notes offer practices and inspiration, as well as structure the challenge so it can be addresses.

1. **Have we identified and analyzed where the key nodes are, and what their interaction is? Are there interaction gaps that we can bridge?** (Link to section here)
    a. Do relevant "managed communities" exist, which could help us harness ready-made collective intelligence?
    b. Are the humans in those networks familiar with each other?
    c. Do we need to illuminate the network, especially the key nodes and related connections – and then visualize and analyze them?
    d. Is the network structure ideal?
    e. Do participants contribute enough?
2. **Do we understand the characteristics of those nodes, especially (but not exclusively) if human? Have we spent resources improving those characteristics, or improving the mix in the most important clusters?** (Section here)
    a. For instance, gender, cognitive and interaction styles, experience and skills, location (remote vs close-by), engagement level, trust
    b. Have we deliberately chosen the level of diversity that we want, or are we defaulting to what's in place?
3. **Are roles clear**? For instance, have we designed archetypes of roles, and shared them with the network (or individual groups)? (Section here)
4. **Do we understand the "soft rules" that govern the interaction between nodes? Do we like what we discover? If not, are we spending resources fixing that gap?** (Section here)
    a. For instance, culture, norms, practices, values, purpose, psychological safety
    b. Are new or peripheral nodes helped to tie with others?

   a. Are the rules, possibly embedded in culture and norms, leading to effective and efficient collaboration
5. **Are we monitoring the developments in the network, and constantly supporting its evolution in the right direction?** (Section here)

Much research has gone into assessing the right size of groups, and the right type of people in the group. By and large, there is agreement that complex problems require larger groups of high caliber people. But the correlation is not straightforward, and research has shown that other characteristics like social perceptiveness of individuals may play a role. The combination between people and machines at large scale is also being researched actively. All these aspects should be considered when trying to identify the right composition and structure of a network intelligence.

In this section we will not do justice to all that corpus of academic research, and we will clearly fall short of providing a structured, cohesive framework. Instead, we will focus on the practical steps that can help strengthening the network to make it more effective.

Network nodes' identification and interaction analysis

There are a few tools that can be leveraged for different steps of this module – all have a role in illuminating a network to further connect and nurture it. We will start with the identification of the important nodes, and the understanding of their interaction to surface and address gaps. Over time, we will likely see the emergence of technology based on knowledge graphs that enable searches for "similar people" not just the current "similar content". For now, however, there's no shortcut to some heavy lifting – with significant ROI.

The easiest way to harness a networked, collective intelligence is when there's already one in place and it is being managed in a way that third parties can access it. That is what **managed communities** do.

Some networks are already well mapped. Many years ago, the emergence of marketplaces like Craigslist or eBay showed the potential of giving people the ability to trade outside of more-hierarchical retail environments – and the rest is history. Today marketplace-type communities exist for thousands of niches, including local networks such as NextDoor ("the neighborhood hub for trusted connections and the exchange of helpful information, goods, and services") or LetGo (peer-to-peer circular economy). These environments can be very specialized – think about a Facebook group called "Welt Frauen" (German for "World's Women"), constituted of German-speaking female expatriates who exchange tips and even accommodation listings. Or, in a similar but opposite situation, think of the WhatsApp and similar support groups that form around communities of disenfranchised and abused expatriate workers from poor countries[75]. Another example, at the intersection of academia's invention and funded innovation at scale, is Europe-centered Crowdhelix – "that forges links between an international network of excellent researchers and innovating companies, so that they can plan and deliver pioneering collaborative projects". These platforms enable people to identify other relevant nodes in the network and guarantee a modicum of trustworthiness. Some networks are emerging, for instance the "maker movement" (such as Fab Labs) that foster collaborative experimentation especially with scrappy hardware "hacks". Citimart is a platform that helps cities to get more innovative in their procurement processes, by enlarging the vendor and subject matter experts' pools, and enables the sharing of outcomes and learnings from respective implementations. Restor.eco is a science-based open data platform to support and connect the global restoration movement, accelerating the global restoration movement by connecting everyone, everywhere to local restoration.

What happens when the community and the trading is one of ideas? Think of LinkedIn, Reddit, Quora and other online communities that while not fully managed, benefit from infrastructure and rules that a third party has created. For example, LinkedIn has interest groups that connect people with shared knowledge. Even beyond knowledge work, networks can exist, as shown by JobCase.com. And the members of specialized communities are harnessed surgically for commerce, for instance by online learning marketplace Teachable.

The traditional approach is to identify relevant experts and connect with them, even possibly hire them. Identifying experts is a core tenet of illuminating a network – even if in the end, one doesn't use the network – and related practices have been around for years. Crowdsourcing is the practice of getting micro-work or project work done by contractors. It typically uses competition, as opposed to extensive collaboration. It also bypasses the traditional credentialing mechanisms and encourages contribution from non-traditional people. While crowdsourcing has limits, it has also shown great successes. That's happened both at the low-end of the cognitive range (where work doesn't require strong ideas), as well as the high end (where unprecedented, fresh ideas are generated often from unexpected people prima facie devoid of relevant domain skills). One of the key characteristics of a successful crowdsourcing contribution is the definition of the "challenge" – which requires de-jargonizing the language used and enabling a wide range of people to participate meaningfully. Research shows that complex scientific challenges have been most often successfully solved by "marginal"[76] participants, that is people who either don't possess the expected technological knowledge or come from networks very far from those that originate the challenge.

Some companies have started filling this space – identifying people that help for specific purposes, and enabling them to use open, purpose-built platforms: Amazon Turk is the classic example, and has been used for micro-tasks (say, labeling pictures) for more than a decade. CrowdFlower is the other leading, public provider. Other entities that provide microwork platform are Appen, Lionbridge (specifically for translations), Amara (for subtitling), Clickworker, and iSoftstone to name a few. Great hopes have also been placed on Quantopian, whose contests populated by thousands of "quants" from various fields aim at creating new, and radically diverse, financial markets trading algorithms. Another interesting, specialized organization is Mindsumo, which crowdsources work to university students. With the increased need to train AI, most of the major players have instituted some sort of microwork platform for instance in the case of Microsoft's UHRS (universal human relevance system).

Databases of contractors are the foundation of many other companies, from UpWork (from small technology jobs) to now-defunct UpCounsel (focused on legal support) to TaskRabbit (from handymen to anyone who could help running errands - acquired by IKEA). The common trait here is the credential and trust system, that's based on user ratings in addition to the company's vetting, as well as a transaction-clearing system that takes care of anything from payment to exception management (e.g., refunds).

When a community is already managed, like for example in existing community-based environments like Kaggle (data science), Innocentive, LEO ("LeadingEdgeOnly") and OpenIdeo (general innovation), 99Designs (custom graphic design), Behance.net (owned by Adobe, a collaboration for creative work), many of these activities are already performed by that organization. Those environments however generally don't allow to identify the members, which means that the network will stay hidden to the participants.

We also now observe the emergence of service providers, such as consulting and software-as-a-service firms, whose job is to create and manage communities. One interesting example is www.ubuntoo.com (full disclosure: advised company), focused on environmental sustainability innovation. A second one is platform-enabled consulting boutique shapeable.ai focused on societal challenges.

While collective intelligence depends on the structure of the connection, one element has clear importance: familiarity between members of the network. For instance, research shows that in competitive sports, regardless of the quality of players, players' familiarity with each other tends to be a strong predictor of success[77]. In enterprise teams, the same effect is visible and typically addressed with some "kickoff" or bootcamp rituals, among others.

However, when attempting to create a supermind, we go a step beyond. **Illuminating a network from scratch** requires a different level of effort, and that's what we will explore below.

First, we need to find individual nodes - for instance, experts (e.g., academics, specialized consultants) or sources of information (say industry magazines). We can analyze who follows them for those topics - that already illuminates part of the network.

An example is the mapping of the "maker movement" we referenced earlier, that spans slightly different categories (from Fab Labs to entrepreneurs) and geographies. The picture below shows the conceptual segmentation[78] that could traditionally be performed by market strategists or management consultants, but which benefits from new network analysis tools.

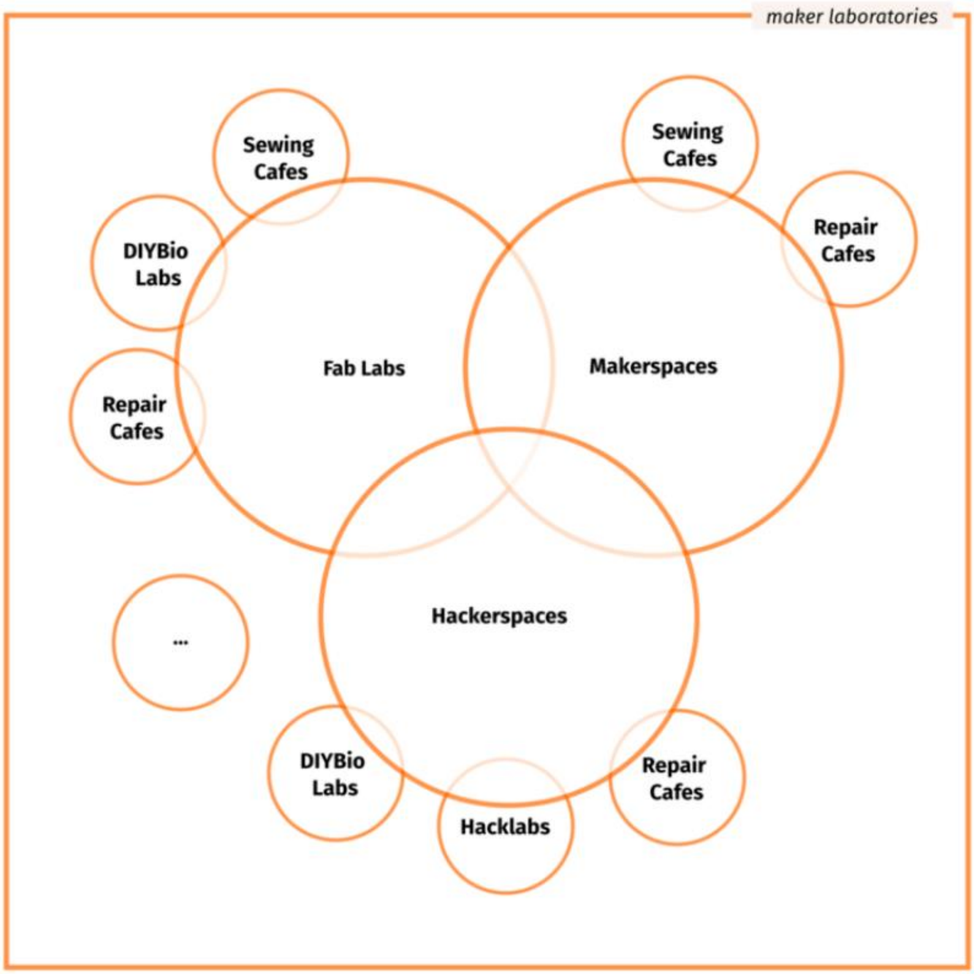

Network analysis, which will be discussed in this section, sheds additional light on the structure of those networks, hence facilitating the connections between groups and the propagation of ideas. A view of that analysis[79] is presented below.

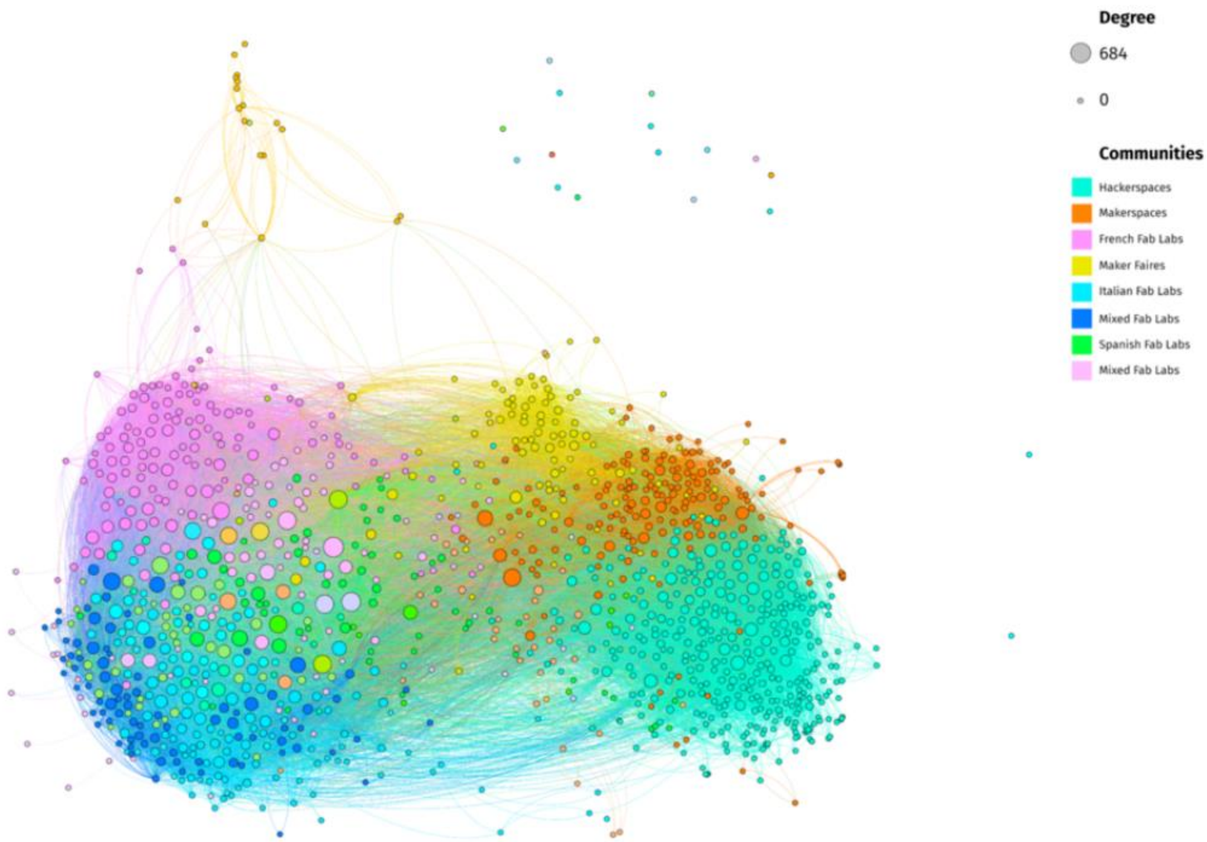

But we can also mine the semantic network, the words that describe their work - and find connections with adjacent environments. For example, reskilling has to do with content curation and generation as well as propagation (hence the connections with knowledge management practices). It has also to do with employee and learner experience, which relate to human centered design for employees, and with cognitive science. By broadening the aperture and mapping the semantic network beyond the obvious words, we also find new people that can help generate novel solutions. These two steps can be performed in the broader world, by mining social media and other available sources, but also within an individual company. When connections between people in physical environments are important, in-person interactions can be mapped through RFID devices such as the one from MIT-spinoff Humanyze.

Then, we need to plan on connecting those neural nodes, on creating the equivalent of synapses. As shown through the discussion of Dunbar's number, that isn't easy unless we reduce the cognitive effort required by the individuals in connecting with each other and give them enough incentives to participate. (For machines, we need effort to connect them).

The first cognitive process that a large and loosely networked group of entities (human/machines) would likely try do is sensing each other - that is, understanding who is at vicinity. That inclination provides some initial energy in the system that we should harvest and use, because the inertia of a static system is possibly the largest inhibitor to the creation of an effectively networked group. In the case of a company, that can be achieved by inventorying people involved in specific activities, and then publishing their names so that the broader company, as well as workers with similar/complementary skills can find them easily. Interesting, one of the most important acts in activating a networked intelligence, a supermind, is to enable people who are engaged in similar or complementary activities to connect with each other. Just the act of enabling that connection can ignite higher-intelligence properties.

This a well-established field and of use in illuminating the collective intelligence networks. **Visualization and analysis** of social networks (**social network analysis**, SNA) is now routinely performed through tools and companies like OrgMapper, Kumu.io, TrustSphere, Cyram, Gephi, Netminer, but also add-ons to established software like CompassHQ (that works on social collaboration software Slack) or Microsoft analytics that draws on the data from Outlook. A significant amount of literature is also available about sociometrics[80].
A small group of consulting firms, ranging from established management consulting companies to specialists (like Valdis Krebs' Orgnet, or Peter Gloor's Galaxy Advisors) have perfected this work over the years. To obtain a better picture of connections between groups of people, additional sources of data for network analysis can be other communication technologies, for instance videoconferencing systems like Zoom, that increasingly provide potentially rich data like turn-taking of individual participants. It is also possible that users of technology like business process mining (Celonis or TimelinePI for instance) will increasingly utilize social network analysis to complement their workflow analysis.

Communities can make use of these technologies to identify "who knows/does what", like in the example of Turkey based design community Atölye.io.

A different type of visualization of nodes, in the space of scientific authors, is the interesting www.connectedpapers.com. Apart from showing connections between papers, it enables the visualization of the connection between authors in ways that would be impossible just with citation maps.

In the reskilling case study that we use throughout this document, network analysis was performed by leveraging survey data from 10,000 employees, who were asked the names of colleagues they ask for help in specific circumstances, and then computing those connections in Gephi. The metrics used were standard (betweenness centrality, which identifies people who are on the shortest path between everyone in the network; and eigenvector, that emphasizes people whose connections are themselves very well connected, conceptually similar to what Google Page Rank algorithm do).

MIT's "honest signals" theory[81] takes this work one step beyond the statistical analysis, attempting to derive conclusions on people's attitudes as demonstrated in their network behavior. While emotional states are important in explaining the creation of connections between people, sentiment analysis is only one of the parameters used, and in fact very little text is used in email-based analysis. For instance, well-functioning, adaptive networks display certain traits:

- the presence of central nodes that are strongly connected to many of the most relevant nodes in the network; they constitute the "management backbone" of the network and have an important role in sharing information from different parts of the network. Many-to-many connections are good for the creation of ideas, but less so for devising and executing action plans. Betweenness centrality, or Eigenvector are two metrics that can be used
- the periodic change of influence of those nodes, so that new nodes emerge, and their ideas and connections become important
- shared context in the form of shared language semantics well understood by the group, as opposed to erratic jargon that might impair the ability of some nodes to collaborate effectively
- ability to respond to each other quickly hence keeping the flow of ideas fresh and "in the moment"
- balanced contribution, which is the amount of input provided by the most important nodes in relation to the input received. When collaboration is important to generate multiple points of view, a balance between the contribution and the receiving is a symptom of a well-functioning network

The application of these methods has spawned interesting possibilities, like in the case of encouraging large groups of people to modify their behavior through "virtual mirroring" – where the data about their social actions (speed of response, breath of network they engage, etc.) is played back to them with an explanation of what those behaviors do to their counterpart. One application is for instance in coaching large, outsourced groups on how to engage their counterparts.[82] Another is the use that Yva.ai does of network analysis among other signals to predict employees' resignation. And some work has also gone into predicting success of network-dependent roles, like salespeople, based on their "network signature". Finally, as we will explore further in the next chapter, the use of social network analysis could support the data analytics needed to provide constant feedback and incentives to network participants.

Other interesting experiences in that field have been made at Bank of America, General Motors and now Amazon – at the hand of Michael Arena, building on previous work with Babson College's Rob Cross. He identifies four key network roles that all organizations need in order to enable "Adaptive Space." Those are: brokers, who have connections across many groups and hence able to discover new things; connectors who can connect local groups to those who have the cross-enterprise networks, and hence able to implement new things; energizers who can provide intrinsic or extrinsic incentives – more on this in the next chapter; and challengers who challenge the status quo. These roles help to facilitate the movement of ideas and information across the firm, and therefore, enable the organization to positively disrupt itself."[83] Proactively identifying those people and enabling their activity (for instance, by protecting the disruptors, or creating spaces that enable local connections) can help increasing the intelligence of the network.

One important aspect of the identification of key nodes in a network is the possibility to understand not only who those people are, but also how they are connected to each other. "Strong ties", that is connections that are strong, typically with few people we know well, are important to shape our opinion, but they may suffer from so-called "small world effects" in that the ideas and beliefs transmitted through those networks may become homogenous. "Weak ties" instead connect people with others who are farther from them and can bring fresher perspectives if they're listened to attentively. In cases where those perspectives are uncontroversial, they can result in behavior change – that's what happens with some fashion trends or electronic gadgets, for instance.

Social media sites like LinkedIn and Twitter likely walk a tightrope between prioritizing content for our viewing from weak-tie and strong-tie connections. It seems plausible that at some point people will have the ability to customize settings, so they can get exposed to more of one or the other.

One particularly interesting field of study (Damon Centola's "broad bridges[84]") has focused on **network engineering to enable the diffusion of more complex behavior**, that is the change of behavior that's more rooted. Think for instance of the spreading of contraceptive practices in villages in developing countries. In those cases, Centola's work indicates that critical mass is crucial: that is, focusing on changing the behavior of and "enlisting" a few relatively unrelated nodes in each group (e.g., a village) is more effective than trying to cover more groups through only one "converted" node in each group. The influencing effect of few role models in each community tends to compound and overcome the resistance typical of established habits.

It is likely that the future will see ubiquitous use of AI-augmented **graph databases[85]** (a specialized field of data science that enables better analysis of connections between entities), like Neo4J, or specialized tools like Semantic-Web.com, Diffbot (which is already used for recruiting, business intelligence, and sales & marketing – and takes a grammatic-rich approach to the problem), golden.com, Kyndi and Primer.ai. DeepMind, now owned by Alphabet, is also doing some work in this space. Think of these as being able to perform a "reverse google knowledge graph", that does to people what google does to knowledge: finding people who are associated with specific topics, and then people who are associated with those people. Within individual companies, a few firms are starting to specialize in the automated identification of experts – WhoKnows, Starmind and Profinda are three such players. Microsoft Teams' "who bot" is an example of artificial intelligence bots deployed to identify the relevant people in a network, so they can be accessed by those who need them. With the famous acquisition of LinkedIn, the less well known one of network-analytics firm Volometrix (at the foundation of Microsoft Delve – now Workplace Analytics), and more recently Project Cortex (application of graph database technology) that turned into its new product Viva, Microsoft is clearly one of the players to watch.

Back to our **search for key nodes of knowledge and influence**. Social network analysis does a good job at helping identifying nodes within the organization. But what about the broader ecosystem? A process to do that, starting with the identification of the right people is illustrated in the picture below, and further described in this section.

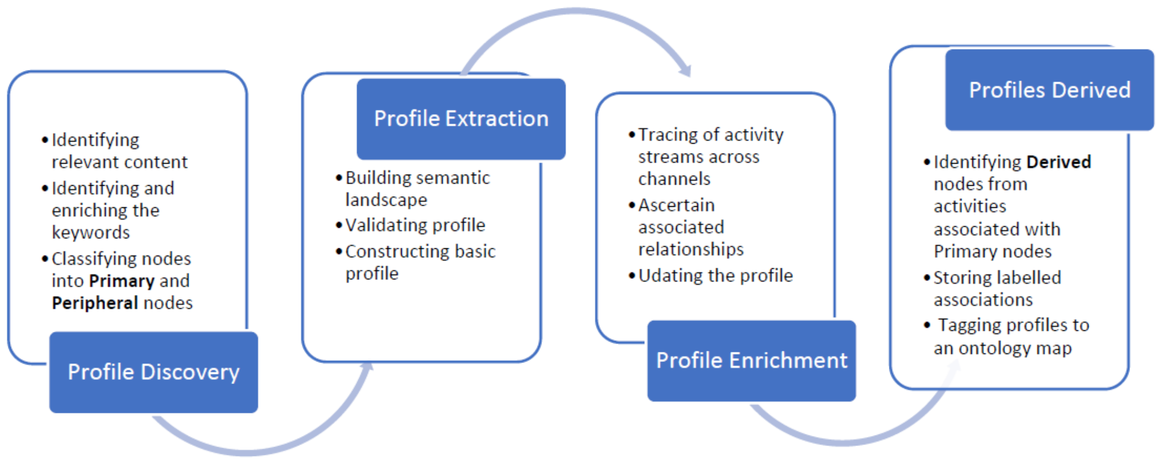

To map relevant people and "illuminate" the external network we first need to precisely define the language for topics of interest which can be disambiguated through what is called "entity extraction", which helps identify the semantic field, i.e. the words, that reflect that theme – for instance "artificial intelligence" and "machine learning" and "neural network" would be branches and separate entities. The next step is some social media scraping, for instance by combing the likes of LinkedIn, Twitter and searching for connections or followers of initial nodal people. For instance, experts in reskilling would likely follow a few dozen core experts in that field. Having a database of key people can help connecting them – as we will discuss in the next chapters. The figure below displays an example of one such semantic network structure.

Tools of use in this area can be social media scrapers (like Phantombuster APIs, Scrapestorm), email finders (like Hunter.io, Clearbit, Voila Norbert), and analytics tools especially for entity extraction (R and Python).

Interestingly, mapping machines instead of people, at least at some basic level, works the same way. Tools like Curata and Feedly can help identify aggregation of relevant content, and their sources.

One last word on physical meetup environments, where individuals can be discovered and engaged in person. Trade conferences, Meetup.com, and the plethora of special-interest local chapters are all examples of places where networks can be identified. If they leave an “internet breadcrumb trail”, they’re typically detected in much the same ways illustrated above.

But **how to tell if the participants are contributing enough**? Contribution levels typically fall into three main groups, and it is important to understand the role each plays. Leaders spend most of their time enabling the network and making key decisions. They’re expected to do over 80% of the effort required to create the network. Collaborators spend some of their time helping to enable the network. Their contribution is much lower. Finally, occasional contributors typically do a much smaller amount of the work.

This is particularly important in communities, that can belong to one of four archetypes resulting from two dimensions: contributor growth (high/low) and user growth (high/low). The two aren’t correlated, and indeed when lacking the right level of modular and scalable architecture, high-user-growth communities collapse under their own weight, as the growth of contributor (and their coordination) is insufficient to sustain the volume of queries and exception management. This is one of the reasons why many communities never take off.

Regardless of the actual amount of "heavy lifting" done, each group is valuable as long as the community infrastructure makes the collaboration seamless (hence the importance of collaboration platforms that we will discuss in a specific chapter.) For instance, occasional contributors are often the source of serendipitously disruptive ideas, as they can connect concepts from very different fields.

Wikipedia and Linux both exhibit these layers, as do other collectively intelligent systems. Think of distributed mobility providers like Uber, where a central group of managers and coders create and industrialize the key processes. The processes are complemented by drivers who not only generate data but also contribute input (often very vocally) on those processes. Finally, end users (riders) contribute data in the form of patterns of behavior, ratings, and comments. Uber's system wouldn't be able to effectively function at all if that layer was removed. Occasional contributors surface data that would otherwise be invisible to the network's central leaders.

Individuals' characteristics

Apart from the network properties, it is useful to understand the *individual* characteristics of the nodes. From that perspective, a people-machine network is different from a human brain, where the individual cells are comparatively more similar to each other.

A core premise of collective intelligence is that there's *emergent* intelligence from the network itself, and that conventional management practices focus excessively on the individual, it is still true that the components of a network matter. For instance, beyond being intuitive, there is scientific evidence[86] that individual sensitivity (or rather, the lack thereof) of any member of a group can result in significant degradation of the collective intelligence of the group. The "no-jerk" norms often enforced in organizations do help addressing that.

A deliberate choice of the components of the network is an important aspect of the creation of a supermind. In some cases, we can't choose what's in the network. That choice may be out of our control, or it may not be legally or ethically appropriate. But a thorough understanding of the characteristics of its individual components gives at least important hints to ensure an appropriate management of its dynamics.

To start with, the correlation between **individual IQ** and collective intelligence isn't as simple as one would think. Higher average IQ isn't necessarily conducive to better group performance. Other factors seem to be more important (like social perceptiveness, and the proportion of women in the group[87]). Additionally, research seems to indicate that high average individual IQ is important when the task is hard, and group size helps increase the collective intelligence. That finding was replicated when evaluating the optimal group size in collective editing, in which condition the ideal group size was found to be about 30 in experimental conditions[88]. However, with more straightforward tasks, the opposite is true: more intelligent, and more numerous people result in limited gains.[89]

For instance, **sex**, **cognitive styles,** and **social perceptiveness** can influence the value of individual nodes in the network. Research has shown[90] for example that the collective intelligence of groups is potentially influenced by the proportion of women, possibly because they tend to have better social perceptiveness and as a result be better at connecting individual minds. "Reading the mind in the eyes" is a well-known test that helps people diagnose their social perceptiveness.

Biological feedback matters. Recent research has also indicated the role of **hormones**[91] (stress-related cortisol, and testosterone) in the ability of groups to produce good output: high testosterone and low cortisol have been associated with strong output, but particularly so when the groups where homogenous. In diverse groups, high testosterone seems to generate adverse dynamics that prevent harnessing the cognitive power of diversity.

Think of the implications: fast results can be obtained by low-diversity groups where testosterone may be enhanced for instance through some competition-type exercises and, if the cortisol is low (i.e. members don't feel stressed), the outcomes may be improved. However, for complex problems, diversity of thinking is important, and the role of competition-induced testosterone can be a drawback – as could be, for instance, the presence of too many "alpha males" whose testosterone is high.

As noted earlier, human brains in stress conditions trigger dopamine-endorphin "feel-good" hormone cycles when they interact with other people (include those well beyond immediate social circles), giving our species a distinct social collaboration incentive when facing complex situations. However, and because of personal biological or behavioral predispositions, people react differently to stress conditions, and some also find collaboration itself to be stressful. Borrowing from emerging coaching practices, stress itself can be reframed as a constructive and natural element part of solution-finding, and generally part of the process of getting better. There seems to be evidence that after reframing, individuals affected are able to lower the perceived levels of stress to a more functional "eustress" threshold, which is just enough to generate adequate arousal to perform cognitive tasks effectively.

Many management practices implicitly try to coach employees that way, but there's reason to believe that people in networks would benefit from a deliberate reframing of collaboration-related stress.

Humility as a shared norm could also prove beneficial in collective intelligence, because humility enables better collaboration and openness, as well as limits the aggressive behaviors that may lead to the expression of stress-related hormones which can marginalize the prefrontal cortex.

Experience and **skills** matter in networks too. New research shows that skill congruence (the complementarity and synergy between group members' skills), as well as individual skills, matter – although they do differently, depending on the type of task. For example, brainstorming requires team level skills congruence, whereas completing a Sudoku game benefits from individual skills[92]. In the same research, remote work seems to require stronger individual skills for the team to be performance, possibly implying that collaboration between remote individuals is stifled by distance and matters less than personal, independent abilities. Much of the study was conducted before the Covid-19 pandemic forced a significant improvement of remote collaboration skills, but it is noticeable how without a learning curve, the quality of the collective output certainly degrades. Interesting, the same research points out that high age diversity may act as a hindrance to group-level effectiveness, hinting at the need to enable cross-generational working teams with the tools to collaborate more effectively.

Another example: people with experience in the *transformation* of processes, and design of new solutions, will have different characteristics from people who are more experienced in managing *established* processes. For instance, while both groups are important in creating new ways of working (think of a new business process, a new technology, or a new service), the former will be crucial in thinking out of the box and ensuring new solutions are designed fast and thoroughly, as well as prototyped and experimented in a disciplined fashion; the latter will be important to "sanity check" the prototype, and will become essential when rolling out the new solution – as they will be the backbone of the respective operations and scale.

Most companies have relatively primitive skill inventories (i.e., know which people has what skills), and even fewer apply that knowledge in decision-making processes. Credibility is conveyed informally, either through rank or executive presence – and both can be flawed inputs to decision-making. While the complexity of that task is understandable, it seems intuitive that is a gap. One company that has put much attention into this is the top hedge fund in the world, Ray Dalio's Bridgewater, and their dot system[93], where people rate each other nonstop to weigh people's opinions and reach the best decisions. They do so through algorithmic analysis of their collective cognitive process and allocate people accordingly. (It is remarkable and has been translated into a commercially available software called PriOS, though in ten years we will likely see the current version as a quaint and incomplete prototype.)

There's been much discussion about "skills of the future", with one relevant insight for us: professionals involved in innovative efforts tend to be "T-shaped", which means that they can go beyond the deep but narrow expertise of "I-shaped" individuals and are able to help connect dots between different ideas. That's particularly useful when working in groups, whether small or very large. Hence, a well-functioning group will likely need to contain people with those skills, for instance able to straddle at least superficially between functional or industry expertise, digital, experience design, process design etc. Deliberately design the composition of groups with that in mind can help achieve the best results.

This said, specialized expertise is unlikely to go out of fashion. Experts' input is very valuable when combined with the work of others. Crowdsourced innovation platform Innocentive has found that retired workers, for example, bring a wealth of expertise that allows group to deliver the best ideas[94]. One good way to detect the characteristics of future experts, and ensure that they are reflected in the supermind, is to mine the evolution of job descriptions, as Burning Glass Technologies do.

Distance creates natural limits to what networks of people can do, but less so than in the past. That's why **remote workers** are becoming increasingly important in the structure of organizations. Notably, research indicates that remote work seems to require higher individual skills to generate strong output, likely to palliate dilution in the collaboration's contribution, possibly pointing at the need to get better at collaborating remotely to unlock quality output[95]. Apart from technical challenges (e.g., video and voice quality) that are being increasingly solved with today's solutions (more on this in the collaboration platform chapter), the presence of these network participants influences ways of working and managing, as well as has a direct bearing on the type of work that can be performed. While this is not the place to discuss in detail about these workers, it is important to realize that there are big advantages in a remote workforce: the ability to accommodate women and seniors in certain phases of their career; and the availability of accessing talent at the right time at the right price. However, the sense of belonging and connection that an office environment provides is often lacking in those groups, and that problem requires deliberate addressing – for instance, by ensuring that remote workers have constant touchpoints with some of their groups; or that the remote workers own tasks that are best done in solitude; or that those individuals are able to sustain long term focus without the need of colleagues nearby; and clearly, ensuring an intentional design of the technology and work practices so that these people feel part of a group.

Network dexterity of the individual node is another factor to consider. For instance, some people or machines may be well versed in using networking to solve problems, while others despite their other individual competencies may not know the tools and practices of working in a networked environment. Those nodes may need some help to wire with others. While the obvious way to determine the level of network dexterity is through survey, that might not be possible or practical at scale. However, network centrality measures may be a good predictor of that, as in certain cases it may be safe to assume that little-connected nodes with otherwise good knowledge would score low on this measure.

Human networks' functioning also depends on **time zones**. While work can increasingly be done remotely[96], and remote employees are increasingly able to feel part of the group, time zones influence the functioning of a supermind in multiple ways.

First, the amount of time-overlap for groups operating at the extremes (say, West Coast USA and

India) impacts their ability to do synchronous work. That’s especially true when members of the network require redirection and constant iterations to work efficiently – for instance when doing highly iterative, feedback rich, Agile work, or with comparatively unskilled workers. These considerations often prompt the creation of local “hub and spoke” network structures: there, far-away groups can operate independently when they can’t connect with each other, thanks to individuals in each location that can act as a “proxy” for the global community.

		Time zones overlap for synchronous communications	
		Low (e.g., US / India)	High (intra-region)
Need for constant iteration across network	High (creative work)	Hard, slow	Easy, assuming synchronous comms available
	Low (e.g., standard operating procedures)	Easy, if process is well designed	Easy, but often doesn't optimize for cost-effectiveness of operation

The flipside is that these groups will tend to prefer asynchronous communications, which may be ineffective for certain types of work - for instance when immediate responses are part of the cognitive process between individuals. Typically, routine, or individual-driven work, for instance some analytical work, doesn't require constant, synchronous communications. Creative work may not be as effectively performed by decentralized groups, unless the tasks can be broken down to minimize the need to correspond synchronously with others. An additional possibility is to create groups situated midway between the extreme time zones, who literally orchestrate the extremes, process much of the interactions bilaterally (e.g., Europe-US and Europe-India) and ensure that the scarce, overlapping times of the day are utilized most effectively for true cross-time-zone discussions. More on this is the "Things we must get smarter about" chapter later in the report.

One last point connected to the effect of time zones overlap, but that goes far beyond it: individuals who struggle to be candid or don't embrace candid and explicit **feedback** may find it hard to work across locations, and especially locations with low-time-zone overlap. That is, unless their work is clearly described in tight operating procedures. The problem here is that latent confrontation may result in pent-up distrust and productivity loss, disengagement and generally friction between groups. Specific coaching and training may help mitigate some of this problem, but there's still a lot of debate about "optimal feedback" – with a deep divide between supporters of "positive feedback"[97] and those who militate for "direct feedback" (think of management consulting firms, or investment outfits – including extreme examples like hedge-fund major Bridgewater). While that debate is far from being settled, it seems clear that large virtual networks, with their disembodied nature and their need for quick propagation of input, would benefit from clear signals – which is not to everyone's liking. That's a struggle requiring constant attention and deliberate management, not only through training, but also through the formation of clear and socially accepted norms for communication.

According to recent research, **mindfulness** is conducive to better team collective intelligence, because it helps participants become more focused and balanced, and supports leaders in becoming more self-reflected, among other things[98].

Networks depend on **individual engagement** too. MIT research shows that an organization with top-quartile employee experience typically achieve twice the innovation, and customer satisfaction, and significantly higher profits than those with a bottom-quartile one[99]. In neuroscience, it is well understood that "there is no cognition without affection". Apart from taking specific steps to engage the network (discussed elsewhere in this document), it is important to get an initial reading of the level of engagement of the network, recruit people with a high level of engagement, identify who's most engaged, and support the transmission of that engagement across the network. And even for machines, the level of responsiveness and capacity allocation, often dependent on contractual arrangements and related service levels, has an impact on network dynamics – think about the latency of a remote server response in high-velocity financial trading.

Finally, **diversity** is at the very core of collective intelligence: without it, adding nodes to a network may just add capacity but not the ability to think differently. Research has shown the impact of diversity, writ large, on creation of quality knowledge[100], technology projects[101], innovation[102] and enterprise growth[103]. Diverse individuals, whether people or machines, can complement each other's shortcomings. It is well understood that very diverse group take longer to "storm and norm" and are better at seeking novel solutions for a broader range of areas when given enough psychological safety, whereas more homogenous groups will typically generate solid results faster in their area of expertise but will likely shift paradigms. Another potential implication is that persistency of collaboration between diverse people is key to effectiveness of those groups, arguably more so than for more-cohesive groups[104].

Let's now consolidate the characteristics in the list below. They can provide guidance for interesting diagnostics and ultimately highlight areas that need work. Ideally, the humans in the network (or at least segments thereof) should be assessed with these parameters. Some of them are easy to measure at scale, while others will only be of help when working with smaller groups:

- "**Quality" metrics**: like IQ, performance
- **Personal competencies** such as skills, or abilities as defined in labor taxonomies like US Department of Labor O*NET. Experience in different phases of the business cycle, i.e. design, build and run of new solutions, as well as general management ("run") of established processes
- **Age**
- **Generalist or specialist** in a specific domain
- **Workload**, cognitive surplus available
- **Collaborative** ability
- **Styles of communication** such as those identified by natural language processing algorithms used by several research companies such as Crystal Reports
- **Immediately creative** or requiring pause and analysis to produce output. Those who need time to reflect will find it hard to contribute well in synchronous exercises, and that is especially problematic when time overlap between people is limited – because the iterations will take a long time to converge
- **Social perceptiveness** measured for instance by "Reading mind in the eyes (RME)"[105]
- **Cognitive / emotional profile** of individuals: Deductive vs inductive, better with Semantic (language) vs semiotics (visuals, symbols) driven, geospatial reasoning, intuitive vs analytical, level of curiosity
- **Myers / Briggs** traits and derivatives (although their validity has been questioned in recent times)
- Practice of **mindfulness**

- Perceived levels of **psychological safety**[106] – are one of the most important predictors of innovation and creativity in teams. Measuring them at individual level is useful. Google's project Aristotle[107], meant to discover the traits of successful teams, showed that this is a key metric.
- **Favorite roles** of individuals[108], e.g., leaders vs active collaborators vs passive supporters; Sensors, People Connectors, Promoters, Idea connectors; Ideators; Energizers
- **Formal roles of individuals.** Google's project Aristotle demonstrated that highly effective teams define individual roles well.
- **Employee engagement levels**, as measured through surveys (like TINYpulse or Pol.is, or LinkedIn's Glint) or SNA (like Keencorp)
- **Trust level** (overall, and for specific people in specific situations)
- **Physical location** (including office-based vs remote), and time zone. These factors have a bearing on the ability to do in-person or at least synchronous work during enough of people's waking hours, which is critical in some creative work.
- **Organizational structure** – and the position of the individual
- **Organizational complexity** of the group where the individual sits
- **Network dexterity** of the individual node (assessed through surveys or network centrality measures)
- **Other** proxies for quality and individual characteristics, like those from software companies that specialize in behavioral trait profiling (like Talocity or Knack, currently used in the recruiting space)

Roles in networks

Google's "Aristotle" project showed that one of the most accurate predictors of teams' effectiveness was clarity of goals and roles.

Unfortunately, very often role definitions are too high level and delegates the task of specifying rolls in actionable detail to members of the team who may not be competence enough to ensure an effective role definition.

Additionally, the communication of roles to a network is not a simple as was writing a charter page on a program management PowerPoint presentation. Instead it resembles more of a change management exercise where multiple communications layers need to be harnessed to increase repetition and use network effects that will finally make sure that individuals no what needs to be done by them and how.

Also, as recent research shows[109], each network node has one of at least four functions: they can connect people between very distant groups; they can "*localize ideas*", that is implement concepts that come from distant groups in their own local one; they might be really good at energizing the entire network by generating or sharing incentives; and they might continuously challenge and destabilize the status quo, creating a tension that leads to change. Other methods can also be applied to the identification of the roles attributed to the most important nodes. Do our people know what their roles are? And are we evaluating the potential machines who have nodal positions with this lens?

Analysis of culture, norms, practices

Groups of people behave in a certain way also because of cultures, norms, and practices. Connected, collective intelligence benefits from understanding what those are in the respective group, and possibly shaping them to achieve the appropriate level of diversity, or uniformity (the same principles as described in the previous section are likely to apply here).

Understanding the common **purpose**, or the lack thereof, in a group is an important first step, as is the creation of a purpose that rallies individuals together and creates energy and incentives to participation (more on this in a subsequent chapter). One interesting framework to create a common purpose is Marshall Ganz's, focusing on shared stories, relational commitments between members, clear structure, creative strategies, and effective actions.

Many companies have achieved superior levels of collaboration by creating and enforcing shared **principles**. One famous example is Ray Dalio's principles[110] at Bridgewater, the well-known hedge fund that helped create one of the most unique company cultures in the world.

Analyzing **values** is another way to peer into the set of shared assumptions in a group. One model for the analysis is Gunther Weil's Minessence[111] and another one is Shalom Schwartz[112] where values are defined as: beliefs inseparably linked to affect; refer to desirable goals that motivate action; transcend specific situations; guide the selection of actions, policies, people; can be ranked; and their relative importance guides behavior. Through the analysis of natural language, it is possible to understand the values of parts of the network. Interestingly, values framed that way can also explain motivation in a similar way that Maslow's and other motivational theories do, and that can be helpful when designing the incentive system for the supermind that we will cover in "module 3" chapter:

a. Conservation rests on values such as humility, conformity, rules, traditions, societal and personal security
b. Self enhancement rests on reputation, capital, control, and achievement
c. Openness to change relies on creative freedom, autonomy, and challenges
d. Self-transcendence rests on nature, justice, tolerance, loyalty, kindness.

Practices move the focus from mental and cultural assumptions to organizational and process "ways of doing things". Adoption of appropriate ways of working is quite important in knowledge organizations – hence the focus that for instance professional services firms put into establishing standardized foundations helping teams interoperate well when quickly forming and then disbanding. Communities of practice are formed for that reason, with regional and global leaders enforcing the standards and acting as nodes of knowledge. In the recent past, practices related to Agile management and Design Thinking have become the standard to harness the power of multifunctional teams.

A particularly important consideration relates to the *efficiency* and *effectiveness* of interactions between nodes. As we noted, effective collaboration for solving hard problems requires diversity, and enough people with different views, which in turn necessitates the existence of a culture where people collaborate openly with others, beyond strict hierarchical lines, and irrespective of similarities of views and ways of working. Conservative, hierarchical organizations tend to be poor performers in those situations. There's however a flipside: the collaboration of too many people may reduce speed, which in turn hinders agile processes and decreases the incentives for participants who don't see progress happening. Amazon's "two-pizza principle", whereby agile teams' meetings can't be bigger than what two pizzas can feed, is an attempt to communicate through a simple meme some very important norms that have bearings on a system's effectiveness and efficiency. Similar considerations apply to companies whose meetings participations are bloated, especially when virtual meetings make large meetings very easy to set up.

“The way we do things here” is also a key organizational trait, very visible to every new joiner in a company. It can promote or inhibit the ability of large groups of people to collaborate and achieve successful adaptation. Among others, McKinsey has developed an “agile organization assessment” that can act as a framework to diagnose specific situations.[113]

A deliberate analysis of the cohesiveness of the above, and activities to form common ground in case of divergence, is foundational to well-functioning networks. As in the case of skills, homogeneity or diversity should be intentionally designed – the former for fast and efficient execution, the latter for broadening the world view and exploring new dynamics of collaboration. Teamscope.io is an example of technology intended to provide diagnostic tools to initiate that process, enabling a better understanding of where team members (though not large network’s) stand compared to each other.

Helping peripheral nodes connecting deeper into the network is often nobody’s job. Many organizations create a “100-day” plan for new recruits, but they fail to maintain focus on that. Otherwise, supporting the formation of connections should be a primary responsibility of people managers, but the job is very granular and often falls by the wayside in the day-to-day work.
An app called Bridge simplifies network introductions. The typical approach of "email forward" is clunky and often inappropriate for busy executives whose network is large. Bridge (brdg.app) makes the experience of connection more frictionless and enables the connected parties to exchange accurate information.
Some software companies that specialize in collaboration technology, like Slack, facilitate some of those connections. For instance, a Slack add-on called Donut proactively forces serendipitous or programmatic connection between people, and a Slack functionality (Sparkly) attempts to connect newcomers to others with similar interests. A Slack app called “random coffee” pairs up people in a randomized fashion. Microsoft Office365 version of that is a bot called Icebreaker. In mid-2020, MIT’s Center for Collective intelligence released an open-source code base and concept called Minglr[114], that focuses on facilitating impromptu synchronous interactions between people who want to talk to each other – using some of the same principles that matchmaking apps (like Tinder) have perfected and applying them to replicating “watercooler” conversations. The future likely holds much innovation in this space, possibly through smarter algorithms whose match making is based on stronger knowledge graphs interpreting affinity

signals (like TikTok would, by doing away with people network as a driver for curation, and cutting straight to interpreting interest in the content itself), or drawing from network analytics strategy – for instance, how about encouraging connections between units whose relationships aren't sufficient, such as newly integrated entities in a post-merger scenario?

Ongoing monitoring of the network's development

Many of the tools discusses so far lend themselves to the creation of dashboards. Those can be used by central units to understand the development of the network, including parameters like: Evolution of the key nodes; size of the network; engagement of the individual nodes; other "honest signals" monitoring.

Monitoring a supermind isn't the same as typical, prescriptive "mission control" operations. It is more akin to taking stock of the evolution of an organic system, in order to mobilize resources to the relevant "high leverage points" (e.g., incentive systems.)

Some key metrics must track required and intolerable network behaviors and should be monitored. At the very least, related behaviors should be consistently encouraged.

- General engagement metrics related to simple activities, not unlike what is used for online audience to number of visits and other key actions. Examples include content, comments and ratings provided.
- Network-specific metrics, such as *Honest Signals* developed at MIT. These are described in more detail in the technical notes. They typically correlate with innovation and creativity in networked-intelligence systems. For instance, they identify the presence of a quite a few central nodes strongly connected to many of the most relevant ones in the network; the periodic change of influence of those nodes, so that new ones emerge and their ideas and structures become important; shared context in the form of shared language semantics well understood by the group; quick responses keeping the flow of ideas fresh and in the moment; and balanced contribution, which is the amount of input provided by the most important nodes in relation to the input received.

Module 2: the "why". Create and disseminate incentives aligned with meaningful goals

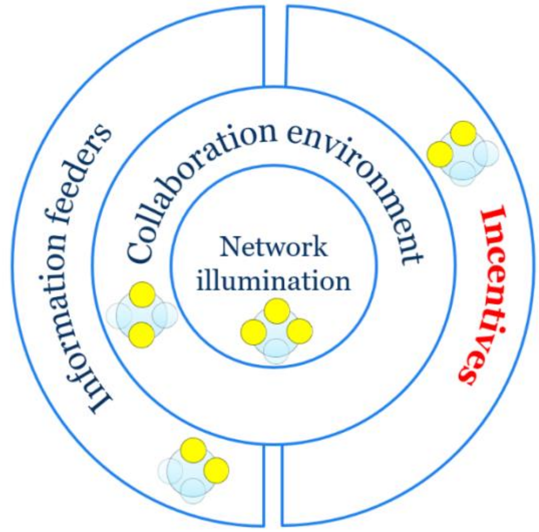

Most network structures are too weak to generate impact. That's often because enough energy is needed from participants, so they stay connected to the network. The issue is not people's tendency to connect – the human brain is wired to be connected, so much so that the loss of meaningful connections excites the same parts of the brain involved in registering physical pain, and painkillers have been shown to soothe such distress. The issue is that human connectivity is organic and hard to direct. People can come and go without forming stable connections, because the prioritize the interactions in the "here and now", and typically can only deal with a small number of "strong ties" – a concept akin to Dunbar number[115] of roughly 150 people that one can be meaningfully connected to on a sustained basis.

And yet, without *liquidity*, the critical mass of people which makes networks vibrant and attractive to more people, collective intelligence may not take off. Network effects through which the network increases its value to participants as it grows has been noted as a foundation of competitive advantage for companies[116].

A system of incentives is necessary to maintain that energy and encourage nodes to connect. At the same time, those incentives need to align efforts towards a commonly agreed and meaningful goal. That goal must be either explicitly stated or embedded into the culture or norms of the organization. In the presence of trust between participants in the network, incentives make the supermind viable and desirable, and ultimately effective as an organizational construct.

That intersection is not trivial. One of the reasons why many employees feel disengaged at work is because the goals are not inherently aligned with their emotions, their belief system, or their values. The feeling of working in teams, with people whose mission and value is shared, has proven to be a significant driver of employee engagement[117] – in other words, the network itself *is* the reward, and one that is often not managed intentionally. Similarly, many intelligent networks fail to "ignite" because incentives are insufficient to encourage the cognitive effort (and changes to personal routine) requested of the nodes. That value needs to be enough, for instance, to encourage the participants to trade off some of their spare time (including machine-time) and redirect their cognitive surplus[118] to a specific challenge.

This isn't a theoretical system dynamics concept. Think of Google's project Aristotle[119], which studied quantitatively the determinants of performance of highly interconnected teams: a key parameter identified was clarity of goals, as well as "meaning", a sense of purpose. Also, think about the role of senior leaders in large organizations and their incessant meetings. Their schedule is full of connections that aren't just strictly problem-solving and functional. They spend much of their time energizing nodes of the network to influence the pace and direction of their organization. This is a form of incentives distribution. Unfortunately, most of that effort is based on intuition and while many executives are intuitively good network leaders, they could still benefit from proactively optimizing the time through structured analysis of networks and audits of their own activities.

Second, systemic incentive-design has been widely applied in ways that have changed the world, sometimes with a double edge. For instance, standardized tests in education send powerful signals to school systems. Individual schools and teachers have, on average, modified the ways they teach to optimize standardized test scores. In quite a few cases, the impact on the quality of education has been questionable. Similarly, environmental incentives aimed at decentralized improving land use (such as planting trees to offset carbon emissions – a form of market collective intelligence) may backfire if they inadvertently promote indiscriminate afforestation of marginal land or encourage monoculture substituting previous, diverse flora.

Similarly, one the of the most ubiquitous systems of incentives for large networks, the rating system (a trust mechanism), has facilitated the emergence of the gig economy. As web3 technologies such as blockchain become more prevalent, the credentialization of people online may become "deplatformed", that is, independent of individual website and instead attached to the individual as they travel through the web – opening significant new possibilities. Also, tens of thousands of people across the world receive advertising-driven micropayments from Google and others, and millions more receive some element of value in exchange for content they produce. For instance, the free-to-use tools provided by sites from Blogger to YouTube compensate the creative producers who use them. Part of the explosion in content available on the web is due to this dynamic. Although much of the content is useless, much that is meaningful and insightful has been surfaced or created that way. In summary, incentives to networked systems do have an impact, and the direction of that impact isn't easy to control with today's organizational methods.

Traditional systems of goal setting and incentives also generally struggle to galvanize networks. First, they focus on organizational structure roles, not *network* roles. Network roles cut across org chart roles, and their incentive structure draws from a variety of practices beyond traditional organizational design and human resources, like digital marketing and social network analysis among others.

Standard systems of incentives also tend to disproportionately reward individual leadership. We live in a winner-take-all society where simple stories of heroic leaders are glorified, while we pay comparatively little attention to the web of contributors. The hidden figures often seem less newsworthy and even distracting compared to a clean, hero-centered narrative. But collective intelligence's power rests on hidden figures. That is a common challenge for networked intelligences. Every big collective effort, from Wikipedia to Linux, has always walked a tightrope around allocating credit to many contributors while recognizing the leadership of central nodes. Successful networks ultimately avoid concentrating rewards or glorifying one person.

Another problem with traditional incentive systems is that they were designed for comparatively low-and-slow information enterprises. In those environments it was hard to detect contributions from people outside a core group with direct influence on the outcome. In that world, innovation at scale was less of an everyday necessity, and long linear cycles of comparatively command-and-control processes were good enough to adapt to more stably growing markets.

We see significant impact from these conditions on innovativeness on and adaptation is likely significant, at least judging by the amount of profit redirection to share buyback, and general dissatisfaction from most CEOs dissatisfaction with their company's innovation performance. Recent work[120] summarizes the all-too-common dynamics that might hinder collaboration: for instance, determining individual contribution in large companies becomes hard if one only uses standard key performance indicators. Many individuals spend time cultivating their relationships because company politics are an effective substitute for precise effectiveness metrics. But the value of naked company politics in determining success tends to be an obstacle to more free-flowing collaboration. Also, when organizations are large and hierarchical, promotions are on everyone's mind (especially if they come with a hefty pay increase) but they create an incentive to succeed individually, not collectively.

The problem is not just organizational incentives. Left to their own devices, people don't always readily collaborate. They certainly don't naturally collaborate at hyperscale, with notable exceptions such as highly energized social movements where the alignments of goals and/or alignment of the culture of the group is often very strong. Other research points out that people tend to collaborate more readily in large groups when confronted with complex problems, but they seem to underestimate (to their detriment) the value of collaboration for simpler things. That bias might deprive a networked intelligence of some cognitive energy and may be worth proactively countering.

Interestingly, the same issues may affect science[121]. There, individuals should cooperate in a frictionless way for the sake of society, reflected in the public money that funds research in the first place. However, and especially as public funds dwindle, even scientists experience a tension between the real need for collaboration and the need to prove one's individual worth in a visible way – that is, by *visibly* authoring peer-reviewed papers. Peers review itself is a form of collective intelligence, where a diverse body of people contribute to fact-checking and improvement of initial ideas. But the authoring of papers, while more diverse, is still a few-people affair. Collective authoring is being experimented with, and in the future, we will likely be able to recognize people whose early ideas were recombined into breakthroughs by mining the spread of related language. Papers that explicitly focus on refuting the validity of prior work, and papers that focus solely on meta-analysis of pre-existing literature (for instance, review of COVID-19 studies during the emergency) exist, but they are a small fraction of the total despite their usefulness. And generally new research is surfaced so that authors are recognized for it in a controlled way, which also limits incentives to broadcast preliminary thoughts to widespread audiences through informal blogs (such as LinkedIn articles). But when strong incentives are provided, things can change: during the COVID-19 emergency, pre-print papers (not formally peer-reviewed, but often scrutinized by open communities) played an important role in providing options to policymakers and their quality proved to be close to standard peer-reviewed ones, while their speed was an order of magnitude faster[122].

Hopefully, the above makes the case for a deliberate design and implementation of a system of incentives that energizes the network. The way a network absorbs, generates and propagates energy in the form of motivations for its constituent parts is a fascinating topic and the object of much research[123]. The technical notes section of this chapter will describe the necessary areas of focus, but intrinsic (e.g., love, glory) and extrinsic (money) motivations clearly play a big role.

Another reason why harnessing collective intelligence is so hard and requires strong incentives is that people tend to be fickle. They are also unwilling to change their routines unless powerfully motivated to do so, or if their contribution is made *really* simple – that is, minimizing the cognitive expenditure required. Especially after the World Wide Web was created, collaboration become often free of charge. But the process of collaboration may still require people "paying" with their time, which counters other collaboration incentives. A route increasingly explored by many innovators is to **design a seamless human-centered experience** for participants to minimize the amount of cognitive load from the users. The higher the cognitive load, the higher the value required to attract people, and as a result the higher the risk of failure.

What happens when collaboration is *not* free of charge? That situation clearly generates an additional impediment, as it is clear when thinking about the cost of communications before the Web was created. Arguably, telecommunication costs have been an obstacle to cross border collaboration for a long time, and conversely the historically lower cost of long-distance communications in places like the United States may have contributed to its growth and innovativeness. Similarly, the European Union's move to unify cross-country telecommunication tariffs will likely play in favor of more seamless collaboration within the EU bloc.

We have so far not talked explicitly about **trust**. Lack of trusts creates disincentives to collaboration. It dampens the strength of connection between nodes.

Social capital is crucial in networks, and the nodes that enjoy it tend to produce stronger network effects. Individual's trust in other individuals, and even more importantly in the broader group, is a determinant of people engagement and collaboration. That trust and drives the amount of reciprocity that individuals may expect before contributing further. Trust has been at the core of the effective functioning of networks well before technology made it easy to do.

This is not a "soft" concept, as a simple example can illustrate. Most people, and virtually all adults in the United States, are familiar with personal credit scoring systems, where a market is used to determine financial trustworthiness. Markets routinely pull together information from myriad touchpoints through aggregators and then act especially in de-averaging pricing (based on risk adjusted premium). While imperfect, such collective intelligences are being extensively used to allocate financial resources. When those aggregators don't exist (e.g., legally imposed restriction on healthcare disclosure) the load is shared more evenly but there may be freeriding and limited incentives for good behavior (e.g., proactive healthy behaviors).

Many large companies depend on the trust between employees who work together over long periods of time –, and spend significant resources in creating trust, for instance through travel and facilities that foster familiarity with others. Some of the principles researched by Francis Fukuyama[124] in the 1990s apply well to a supermind. His work sheds light on the ability of large advanced, large societies to collaborate more effectively when trust is plentiful. Strong social movements, where the energy of the individuals is funneled into the collective, feature high levels of trust within the members of the movement. In more recent times, trust has been shown to generate positive stimulation of human brains (including the release of mood-enhancing oxytocin.) Organizational trust is correlated with engagement, productivity, lower burnout, and a higher sense of accomplishment[125].

All these conditions indicate that trusting individuals will engage more intensely in the work of the network. Unfortunately, the trust levels in many of our institutions, from government to individual companies, has significantly degraded in recent years. This can be easily seen in disengagement from communities and even from the democratic process. The result is arguably a much lower aggregate intelligence expressed by those organizational structures.

While there's no direct equivalent of social capital for machines, it is evident how trusted information is treated differently by network machines. Websites have *trust scores* used by search engines, which in turn make trusted sites more central and powerful than others. Blockchain-based information, combined with the capability of Web3 is also expected to increase the trust level of nodes who provide it, as well as administer granular incentives without needing existing banking systems to allocate that value. As a result, being able to leverage high-trust machines in a network might become a key design element in future network intelligences. Blockchain is ultimate a collective-trust creation mechanism that operates by virtue of a collectively distributed system (a collective intelligence in itself). The use case of bitcoins shows the potential for the creation and distribution of incentives through, for instance, Initial Coin Offerings (ICOs). Similarly, the emergence of non-fungible tokens (NFTs), will strengthen the ability of creators to show proof of their work, and the ability of buyers or funders to believe it – both resulting in the possibility of opening up fractional ownership of real or virtual assets to communities and ecosystems.

But trust can also create herding effects. Collective intelligence emerges from respectful and fact-based critique, which needs encouraging both through norms and incentives[126].

A separate discussion is needed for the creation of a **business case for collaboration** - part of the system of incentives for *budget owners*. Networked intelligence's technology and tools will be detailed in the next chapter, while the technical notes in this section deals with the incentives for organizations to embrace them. Suffice to say, business cases for the formation of collaboration platforms are imperfect art forms, as a large part of the value is hard to quantify. The results may be a more engaged workforce or the organizational ability to respond faster to future volatility.

The web3 space, leveraging blockchain infrastructure, is fast evolving so few established examples exist. It is however worth keeping an eye on, because of its potential.
Quadratic funding uses new algorithms to crowdfund projects without inadvertently falling prey to special interests. Funding public goods through this method is intuitively attractive.
Retroactive funding's vision is to fund foundational work that typically wouldn't attract funding because of its risk but, because of the traceability of the lineage of successful solutions, blockchain could retroactively be used to allocate incentives to the work that was at the inception of those solutions. Funding science through this method also sounds intuitively appealing. Time will tell if these approaches are viable, but at the time of writing, this space is buzzing with experimentation (also see the database for the latest web3 examples).

Before moving to the detail section, it is worth remembering how AI enables the creation and propagation of incentives in unprecedented ways. For instance, AI can help detect what drives behaviors, or compute individual contributions to the growth of the network and its knowledge, in ways impossible until now. And by enabling communications at hyperscale, AI helps connect people in ways that, if done right, can motivate. The picture opposite illustrates some of those elements, some of which we will explore more closely in the technical notes section.

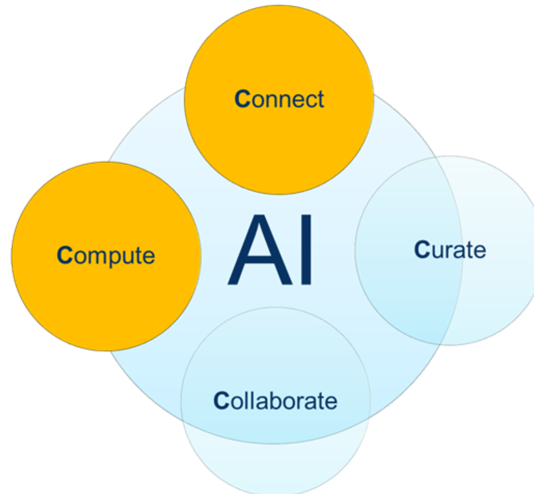

Next, let's bring some of these aspects to life through real-life case studies.

Examples – module 2

For years, incentives have steered the behavior of distributed resources, in organizations from social movements to enterprises.

Google AdWords has incentivized people monetarily to post novel (or simply attractive) content for many years with cash payments.

The Chinese consumer electronics company Haier[127] has constructed a system of incentives for the companies in its portfolio to continuously seek to form new collaboration linkages between them and with ecosystem around them, thereby forcing them to incessantly explore.

Hedge fund Bridgewater uses ruthlessly clear incentives to drive people collaboration in complex problem-solving, and foster a distinctive, self-perpetuating culture.

Wikipedia and Linux have shown how intrinsic incentives related to public recognition and a sense of self-worth can unlock the cognitive surplus of thousands of people operating in a distributed manner.

The Ice Bucket Challenge enabled the formation of a viral campaign of fundraising for Lou Gehrig's disease, relying on intrinsic motivation and leveraging then-novel social media video sharing.

Open-source music-sharing application Napster, before shutting its doors crushed by the legal implications of its own success, showed the power of providing simple incentives ("listen to more music for free") with a strong web-based collaboration platform.

As noted in the previous section, risk "crowdpooling" such as peer to peer lending and crowdsurance startup Lemonade, are advanced forms of allocation risk-reward through an AI-augmented platform.

In China, Pinduoduo enables groups of consumers pool their purchases of farm products and connect directly with farmers and distributors to obtain better price and service, and ultimately better efficiency of the agricultural value chain. Hundreds of millions of people use it today. During the Coronavirus outbreak, which affected traditional wholesale distribution networks, Pinduoduo helped farmers sell perishable stock.

Miners of blockchain receive bitcoins as compensation for their decentral processing capacity use, in support of the blockchain architecture; so-called "yield farming" is a similar example of receiving incentives in exchange for lending capacity to crypto architectures; Helium.com enables a crowdsourced WiFi network, where users add software to their router and use Helium's cryptocurrency as an incentive and governance mechanism.

In mid-2020, in the wake of the controversy about TikTok, Facebook moved swiftly to offer financial incentives to TikTok creators with millions of followers, to encourage them to move to Facebook's Reels; TikTok countered this move with a US $300 million "creative fund" aimed at highly popular creators. TikTok is also an interesting example on how to motivate a community of creators, by making future success on one's creation less dependent on past successes (the "old money effect" in social media) and more dependent on the actual interest of the community members in the current creation. In 2021 social media network Reddit started working on enabling its community of creators to monetize through NFTs, combining extrinsic incentives with its traditional intrinsic incentives.

A more detailed description of some of the above and many other examples are included in the "How to Do It" section of the Guidebook, as well as the database on www.supermind.design/database. Let's explore one specific case study in more detail.

In-depth case study module 2 – upskilling at scale

In the example of the professional services' reskilling architecture, several incentives were designed and implemented for different levels of the company, as summarized in a non-exhaustive manner in the table below.

	Intrinsic	Extrinsic
Top leadership	Genuine interest in making 100,000 employees more future-ready	CEO mandate reflecting indirectly in personal performance
Knowledge nodes	Leadership's attention Mastery Praise from top leadership Expansion of professional network	Bonus payout related to percentage attainment of performance indicators linked to success in propagating knowledge
Learners	Mastery Praise from top leadership	Promotions linked to learning ability

At the center of the incentives' infrastructure sits a gamified *learning score* that tracks learning activities: hours spent learning independently or in groups; time spent learning-by-doing when completing specific learning-oriented tasks; or to a lesser extent time spent updating one's skill inventory in the personal profile or providing feedback on learning content. The top learners are celebrated in multiple ways. The ability to track oneself over time, and the gamified nature of score, ensure that the score isn't perceived as a normal key performance indicator.

Team leaders all the way to the CEO track their own achievement as individual learners, but also through a manager score reflective of the average learning scores of their direct reports. That nudges them to continuously discuss learning progress with their immediate teams. And the model cascades down to the rest of the organization, acting as a force-multiplier for the incentives propagated at the top of the organization.

For subject matter experts acting as *masters* and *gurus*, a balanced scorecard was created to monitor performance of their learning channels. For example, the scorecard reports the number of learners who have completed their learning journeys for a specific skill and indicates the level of satisfaction learners have expressed with the content available. Given the low proportion of users that typically rate online content, the satisfaction measure used is a derivation of Lean and Six Sigma practices related to defect measurement. In this case, the percentage of completed learning cards flagged as “poor quality” by the users is closely watched.

A human-centered design approach was also used to ensure that the journeys for specific user groups (or *personas*) are frictionless, reduce the cognitive load, and create a sense of mastery and belonging. For instance, the process flows involving knowledge nodes were explicitly streamlined and supported by knowledge and community managers, ensuring that the experts’ scarce resources are used efficiently.

In such a system, becoming part of a culture that promotes learning is a large part of the motivation for all parties. To cement these aspirations, a specific visual and language tone was deliberately applied across touch points. This starts with a purpose-built introduction website whose soothing appearance conveys intimacy and personal support, instead of a traditional, corporate, neutral identity. This tone is a radical departure from the company’s standard intranet-based internal communications.

Module 2 – How to do it

*[Reminder: the specific examples documented in the technical note section may be used to compile a list of "portable ideas" that enable more creativity during the design process (see the "*The process of designing an augmented network intelligence*" chapter)]*

The "how to do it" notes for this module will cover possible solutions and inspiration in the following key areas:

1. **Do we understand the *intrinsic* motivators** that may work for the nodes, so we can channel their cognitive energy? Among others, **have we explored the value of providing a sense of mastery and "flow"**? And have we made the network nodes **visible** as useful, knowledgeable resources? (Section here)
2. **Do we use network influencers** and those with authority or clout to energize the rest of the network, for instance by participating in conversations or visibly endorsing people? (Section here)
3. **Have we optimized task design to avoid unnecessarily taxing cognitive resources,** and instead amplify the flow experience? (Section here)
4. **How do we give formal credit to people,** as recognition is important for some intrinsic motivation? (Section here)
5. **Have we understood the implications of** our design of interactions, specifically with regards of the use of **competition or collaboration**? (Section here)
6. Are we able to provide **extrinsic motivation** at scale aligned with visible goals? (Section here)
    a. Is that the right approach, given your objectives and the culture/values of the network?
    b. With regards to extrinsic motivators, **have we explored the opportunity for crowdfunding and potentially web3 technology**?
    c. Have we explored the use of technology to enable the creation of **trust** in the network?
    d. Are we encouraging disagreement and detection of errors, to avoid herding effects?

7. Do we have the right **data engineering and visualization** to provide incentive related feedback at different levels of the organization? (Section here)
8. Have we created mechanisms that make individual nodes feel like they're going to encounter their counterpart in the **long term**, hence encouraging **trustworthy behavior** and discouraging freeriding? (Section here)
9. Have we explored the **business case** that justifies funding for extrinsic motivation? (Section here)
10. If we are trying to incentivize a market, are we aware of the specifics of that use case? (Section here)

More detail about intrinsic and extrinsic incentives is in this section. These are building blocks in architecting well-functioning networked organizations, and their deployment requires intentional process, technology, and data work.

Before we describe avenues to generate and propagate energy in a network, it is worth remembering that a slightly opportunistic approach may be helpful. Energetic nodes, as discussed in module 1, often already exist in a network. "Following the energy"[128], i.e. identifying the sources of existing network efforts and amplifying them, can be a good way to accelerate connected intelligence without needing to expend much effort.

The previous chapter touched on culture, norms, and people characteristics. Understanding these aspects is important when crafting the right incentives. For instance, homophily (being with people of similar traits, whichever they may be) is important in creating a sense of belonging for some individuals. Given the risk of creating groups that are too homogenous (they tend to be less creative, and less inclusive), it is important to use the initial phases of the engagement with the network to create common ground, such as purpose, values, and norms, to stress similarities in otherwise diverse group.

Interestingly, well-functioning groups don't seem to simply need "happiness". In fact, strong output seems to be uncorrelated with the amount of happiness in the group. Other factors are more important and should be considered in the design of a supermind.

One intuitive classification is Tom Malone's "love, glory, money" – explicitly intended to cover collective intelligence systems. At high level - Love can be of others, of nature, but also of self. Glory can be reputation, fame and the like. And money is any material comfort that can be had - irrespective if actual currency exchange is obtained - and it is indeed a large part of workplace incentives.

Maslow's theories can be of some use here as well, as described in the points below representing his pyramid[129] and related examples for the participants in a networked supermind.

- **Self-actualization:** Ultimate feeling of accomplishment through connection to broad network of ideas, people
- **Esteem**: Sense of accomplishment, fame from becoming a well-known entity in a community
- **Love and belonging:** positive feedback from others
- **Safety**: physical and professional safety through access to networks and information
- **Physiological needs**: safety through access to networks and information (or machine equivalent, like access to resources such as energy and capacity)

Along similar lines, management consulting McKinsey[130] has used for some time a framework for **creating meaning through superior narratives**, and hence motivating individuals to participate in organizational transformation that includes:

- Beating competition, impact on investors and market - which is the classic one, and often overused
- Impact on society: for example, making the world a better place
- Impact on customers: for instance, serving clients
- Impact on working teams and co-workers: e.g., improving the work conditions
- Impact on the individual, such as better opportunities for career improvement, pay etc.

Sometimes a strong network incentive may just be the ability for the participants to **find useful resources**, in the form of useful information and people. Finding "experts" creates and exchanges an implicit currency (expertise) that benefits from network dynamics. In the example of reskilling, it can be quite helpful to find others who are looking at those problems, so that everyone has a place to exchange ideas, a very specialized watering hole that efficiently channels what everyone senses, or remembers, or has learnt. If the respective UI is designed well, the cognitive expenditure for the participants in this scenario is low - mostly akin to normal social media share.

In other cases, a central curation and combination of information and ideas can provide enough value to generate solid incentives - for instance, a central group of people can use machines to eliminate irrelevant information from the network to facilitate the absorption of the rest. It can also be useful for centrally located people to connect dots between different disciplines - for instance, showing how network analysis can help reskilling by enabling the identification of subject matter experts.

Mastery is a powerful incentive. For instance, learning methods leverage it to keep people going. Credentials (like Lean Six Sigma, and obviously Karate, Black Belt) In Wikipedia, mastery is a strong driver for community members' efforts. In consulting firms, Practice Leaders and in general SMEs derive part of their satisfaction from that mastery – and by sharing it with the people around them, as well as obtaining powerful feedback from those networks. Mastery is also a social construct, as "masters" often enjoy from the perception of people around them. Klout[131], a measure of online influence now discontinued, was one of the first online scores deliberately built on algorithms that measure the network effects of someone's opinions. Leaderboards of all sorts are a common management practice, and they typically mean something only within a specific group: few people care about who the employee of the month is, outside of that specific organization.

Flow, the concept developed by psychologist Mihaly Csikszentmihalyi [132] refers to the positive experience felt by people who are absorbed in an activity they master. Top athletes, businesspeople, and in general top performers achieve it during what is called "peak experience". Flow is a strong motivator but is not just limited to top performers in individual efforts. "Group flow" is a derivation of that – good examples are dance troupes and other performing arts, religious rituals, sports teams and clearly military troops collective exercises such as marching together. A deliberately engineered networked group, a supermind, can enable people to experience flow, hence further reinforcing the incentives to participate in that network. The deliberate management of this flow is something we will explore in the chapter about collaboration platforms, where we will observe the value of pacing communications, for instance.

Task design is an example of how to engineer work and other activities to facilitate people's achievement of a certain level of mastery and flow. Educational pathways use that in instructional design, so that for instance beginners get positive feedback that helps them progress. The restaurant McDonald's famously redesigned its kitchen layout and standard operating procedures to enable employees to achieve a certain level of collective flow. In that respect, much workflow in large enterprises, starting from early industrial revolution restructuring of assembly line, had some "group flow and mastery" at its root.

These examples belong to a broader "love", writ large, theme. Mastery and flow are a form of love of oneself, as they generate similar emotional and physiological responses. Love, in its broadest meaning, is obviously a powerful motivator. Think of receiving feedback, guidance and attention, or even being communicated to, as potentially making people feeling that are loved, cared for.

Supermind type organizations typically create a very strong sense of common values and **belonging** (a close relative of trust), as discussed in the Module 1 chapter. They do so by often participatory design of those values, and the processes that make the group function the way it does. They also foster that through communal ceremonies, like parties or other functions that so far tend to be in-person, possibly through local chapters of a global network. Wikipedia and Linux have operating procedures that make them cohesive, and so does arguably the most famous hedge fund, Bridgewater. It would be easy – and wrong – to dismiss these aspects as generic "culture" traits: culture is formed by doing, day after the day, what that specific group has decided that it was worth doing, in the way that's right. Getting to that "what" and "how" can be done deliberately, and certainly requires a participatory design element.

Incentives affect the behavior of people in often perverse ways – related to short- and longer-term **trust** - as shown many years ago in the "*prisoner's dilemma*"[133], the best-known game theory problem in game theory which shows why two completely rational individuals might not cooperate, even if it appears that it is in their best interests to do so in the scenario, the actors have (the reason is largely due to imperfect information in a one-off situations). It has since been demonstrated that in situations where individuals are going to face their counterpart later, a more pro-social behavior emerges.

Interestingly, the systems of incentives created by societies have tended to create mechanisms that substitute meeting again with one individual with meeting *members of a community*. – For instance, silk route merchants and traders may not have met with one another again, but they would very likely meet with members of the respective groups. So, incentives for cooperation can be established even in low-information networks, and the individual "tit -for -tat" translates to the network level, beyond individual retribution. The resulting creation of social capital and trust makes large networks more effective, as shown in the work of Francis Fukuyama[134]. These examples are important primitives for architecting modern high-connectivity networks of much larger size – just think of the trust mechanisms that are embedded in many modern peer-to-peer ecommerce platforms.

Even historically, this issue has created the conditions for clever organizational models. Think of the ingenious credit system that merchants along the silk route used for centuries. The system enabled people who didn't know each other to engage in commerce, because they always had very trusted nodes between them. In more recent times, Grameen Bank, has shown that social networks can enable trust mechanisms that allow provision of credit to previously "unbankable" communities. Some inspiration can be drawn from those low-information environments.

Designing incentives in a supermind requires clear and granular choices about **collaboration** and **competition**. The same task, performed as a competition, can be off-putting for large groups of people who would prefer it to be a collaboration, and destroy the motivation that's supposed to build. For instance, design thinking sessions and in general brainstorming practices deliberately reduce the amount of competition in the group, for people to be able to build on each other's contribution. Conversely, competitions like those performed on Kaggle or MIT Climate CoLab's include an element of competition to stimulate participation, clearly signaling to participants that their reputational value would increase were they to win.

As a result, intrinsic, non-monetary (love and "glory" related) incentives are particularly powerful, in groups – including very large groups that include machine-curated information.

For instance, as noted earlier, **influencers and people with significant clout** could inject themselves at crucial junctures in the threads and conversations, and visibly endorse or praise people – or simple show with their presence that they matter. Senior executives do that systematically, by participating in departmental townhalls and in general trying to meet as many of the relevant network participants as possible (for instance, newly appointed leaders routinely spend the first 100 days in meeting people both in person and virtually). The use of network analytics and technologies like Microsoft Workplace Analytics can help network influencers to pinpoint the intervention points and be more efficient.

Another good strategy emanating from this observation is to go where the nodes are, as opposed to requiring them to come to us and do things they weren't going to do. It is often surprising how much pent-up energy exists in people who are eager to be given a chance to contribute.

For example, rich source of energy is the *information exhaust* which can amplify an individual's contribution. Companies like Waze understand that well. Waze offers a very balanced ratio between what it gives (guidance enriched by fellow travelers' actions) and what it asks for (the GPS coordinates of your telephone, and the occasional action to flag an issue on the road). More generally, as we will discover in module 3 and 4, information feeders and collaboration platforms can minimize the actual amount of work asked of the nodes – for instance by mining what they say and what they do.

Alternatively, or additionally, one can use collaboration platforms' social features like "thumbs up" and similar, to crowdsource peer-based endorsements. This age-old practice started with Facebook's "likes" can provide both positive and negative reinforcement through for instance emojis, that discretely nudge people into collectively-acceptable behavior.

There are other ways to incentivize networks of people and machines, some of which are just embryonic as of today. Energized, connected systems already exists - think of the success of things like Kaggle.com, TopCoder or Kickstarter that combine both intrinsic (mastery, fame) with extrinsic (monetary prizes, often however, mostly a reinforcement of the intrinsic motivation). Numer.ai, a crowdsourced hedge fund, uses cryptocurrency payouts to incentivize coding of models done by thousands of data scientists.

Gamification has also been increasingly used as part of the user experience to incentivize participation. More examples are presented in the "Module 4" section.

A similar set of considerations applies to academic research papers, where there are very clear rules about who gets **credit** for work. While the academic research system is mature, and its main output – academic papers – very well understood, there is reason to believe that they may be some opportunity with regards to harnessing the power of larger networks of people. Some experiments are currently happening at MIT's MediaLab[135] regarding synthetic biology, where the operating model is more the one of a movement than a traditional academic collaboration. Additionally, it is not unrealistic that at some point we could benefit from the advances of natural language processing. Imagine if one day whatever breakthrough ideas could be disassembled, and its individual semantic-network components traced back to raw ideas that other people had originated in some form - and those people credited with those inceptions. The hidden figures' contours could be brought out automatically and fluidly without needing complex forensics, and science historians. If this sounds far-fetched, let's just think of the propagation of memes and how they can be traced, how their lineage can be made explicit.

But what about **pure money**? How are money and other **extrinsic incentives** created and used in large groups – apart from the obvious market-based economy and personal use? Clearly, financial markets are large-group economic tools intended to pull risk and reward. The invention of paper money itself, was a social construction that required shared norms and assumptions to function. Organizational incentives such as key performance indicators associated with bonus payout are frequently deployed on groups of people, not just individuals, to secure collaboration. No wonder that employee performance and employee-engagement software is a fast-developing space, with companies like Kudos, Workhuman (Globoforce), Workstars and others vying for the respective enterprise budgets. But what's next, what can fuel the incentivization of activities of large groups of people, and the deployment of other resources (think of cloud computing time in large networks)?
A different example comes from the DARPA "red balloon challenge", where participating teams competed to find the location of ten red balloons released randomly over a very large territory across the United States. The winning team, an MIT group, successfully used incentives to

stimulate participation of thousands of people, not unlike methods used in multi-level marketing[136].

An open question is the effectiveness of concentrated, "winner takes all" prizes, as opposed to more pervasive rewards of lower individual amount. Initial research[137], under somewhat simplified experimental conditions, seems to indicate that especially in groups, "winner takes all" rewards tend to incentivize more dramatically creative ideas, while a less polarized reward system encourages more incremental contributions.

To an extent, and despite the visible significant challenges, arguably in a very successful way, online advertising has brought collective intelligence, supermind mechanisms into the internet. Through a market structure, ad space and placement are dynamically allocated to bidders, based on a complex set of rules informed by the audience's profile. This in turn results in monetizable activities for creators that publish on the internet.

In response to the significant market control of established online companies such as Google, new intermediaries (such as Substack, or its not-for-profit counterpart ghost.org, for newsletter creators) are emerging to provide better incentives to creators. Apart from Blockchain architectures described below, there are new solutions for the generation of granular payment streams from the users to the providers, e.g., those proposed by open-source Webmonetization.org and its web-monetization wallets. Grantfortheweb.org is notable in its mission as "a fund for the development of open, fair and inclusive standards and innovation in web monetization". Coil.com, supports both organizations and offers end customers the ability to browse all web-monetized content through an extension in the user's browser.

The largest **crowdfunding market** in the world, Kickstarter, is a good example. Its structure allows people with specific interests (e.g., art, fashion, consumer electronics, but also others) to provide monetary support for what they believe are worthy causes. In return, those funders receive some extrinsic reward (the finished product) but mostly a strong intrinsic reward, such as "insider" information about the progress of the project. Kickstarter has made the matching of supply and demand for monetary incentive relatively frictionless, so much so that funders aren't just the traditional wealthy patrons of the arts anymore. And because the funders receive strong intrinsic incentives, it enables much more risk taking on the part of the entrepreneur – as well as the funding of endeavors that don't have strong monetary output.

This model has inspired several other market-like mechanisms for funding. From ad-hoc bitcoin offerings to Crowdjustice (legal action crowdfunding – which echoes funding by citizens in Singapore, for instance), Patreon.com (artist crowdfunding), GoFundMe.com (American crowdfunding platform, covering everything from celebration to expensive cures), WeMakeit.ch (Swiss funding for inventors), to Omaze.com (online sweepstakes). And in the meantime, many VCs and angel investors put some early financial seeds into crowdfunding with the objective of following new trends at their inception.

For enterprises, there has been long a need for giving employees a little bit of time and money to start something. Often, determining “who” and “how much”, is hard. CultivateLabs enables companies to give people some time and budget to allocate, in order to crowdfund promising ideas. In some respect this is similar to early prediction markets that were dispensing fictional money to observe which bets employees would take – except that in this case, the time and/or money are real and will be effectively allocated once the employee has made that decision, and the candidate project has achieved a minimum funding level to be approved and funded.

The model is followed by Companisto in Germany[138]. Founded in 2012, Companisto is an equity-based crowdfunding platform in Europe. The firm collaborates with business angels, corporate finance specialists and venture capital companies to raise capital. Companisto offers a sectorial approach focused on information technology, software, consumer products, e-commerce, social media and energy sectors. Similar mechanisms are at play at Republic.io, Ourcrowd, Seedatthetable.

This space is worth watching, as it is still in its infancy and the combination of mass affluent class with social media networks (and ability to segment individuals more finely based on their interest) will likely generate an increase in the number of transactions these markets clear, as well as an extension of the fields they cover. As a sign of the increasing maturity of this market, in late 2020 the US Security and Exchange Commission (SEC) increased the maximum amount for crowdsourcing equity from just over 1 million US$ to $5 million, which will spur activity – for instance through intermediaries such as Crowdfund Capital Advisors and crowdfunding platform Wefunder. For now, however, we are clearly missing a more sophisticated marketplace that includes micro-incentives, or a broader reach than Kickstarter. There’s no “amazon of incentives” out there – with its ability to provide so broad a choice that funders could find exactly the right topic (say, sustainable urban architecture made of bamboo) at the right level of risk (e.g., certified international architects, publicly funded university, or local outfits). Think of micro-incentives (cents of a dollar) that Venmo or PayPal could disburse in exchange for some useful micro work like answering a question from a website or identifying a bug in a browser-based system.

Also think of the ability to transact more precisely and securely, in a radically decentralized manner, thanks to **blockchain** based systems. One example is Ocean Protocol that enables those transaction to happen closer to sensitive data, behind protection wall, or 0x.org for the exchange of tokens enabling the secure allocation of reward to distributed workforces (both human and machines). An interesting new player, built on the blockchain architecture, is commonstack.org, intent in "building commons-based microeconomies to sustain public goods through incentive alignment, continuous funding and community governance [through a] library of open-source, interoperable web3 components [that] put effective new tools in the hands of communities, enabling them to raise and allocate shared funds, make transparent decisions, and monitor their progress in supporting the Commons."

In early 2021, non-fungible tokens (**NFT**) based on blockchain have started to emerge as a potential architecture to fractionalize ownership. NFT provides proof of ownership rights to a broader network, so that members of network can acquire that ownership even partially. This can help communities to own and fund new ideas, while obtaining some ownership rights in exchange – above and beyond the general concept of funding that crowdfunding has promoted.

The concept can potentially also be applied to the scientific and R&D spaces: for instance, the University of California auctioned off an NFT for the documents related to the work of a Nobel prizewinner, and World Wide Web inventor Tim Berners-Lee did the same for source code of the original web browser[139]. NFTs are now used to store individual experiences, through POAP.xyz, a DAO.

The largest NFT marketplaces are about art and collectibles[140] and use Ethereum. The most prominent are OpenSea, Axie Infinity (headquartered in Vietnam, and deals with online games), NBA Top Shot (US basketball). Decentraland, a virtual world referenced elsewhere in this document, also has an NFT marketplace to support transactions. NFTs are now traded on Coinbase, Binance and other exchanges.

Another example is Crucible's Emergence SDK, that simplifies the access to Web3 Digital Trust Layer, with a specific focus on games.
This space is bound to develop fast in the near future and will generate additional viable use cases.

Also, while there are risks and unresolved questions, Initial Coin Offerings (ICOs) also show the potential to raise large sums and bypassing traditional financial systems – and doing so with a large reach, thanks to a collective trust-enabling mechanism due to bitcoins and their blockchain underpinning.

As noted elsewhere, **the web3 space**, leveraging blockchain infrastructure, is fast evolving so few established examples exist. It is however worth keeping an eye on, because of its potential. Quadratic funding uses new algorithms to crowdfund projects without inadvertently falling prey to special interests. Quadratic Funding[141] is the mathematically optimal way to fund public goods in a democratic community, because in the process of allocation of funds, it weighs the number of people who support a specific use-of-funds as well as their individual funding. Funding public goods through this method is intuitively attractive.
Retroactive funding's[142] vision is to fund foundational work that typically wouldn't attract funding because of its risk but, because of the traceability of the lineage of successful solutions, blockchain could retroactively be used to allocate incentives to the work that was at the inception of those solutions. Funding science through this method also sounds intuitively appealing. Time will tell if these approaches are viable, but at the time of writing, this space is buzzing with experimentation (also see the database for the latest web3 examples).

As discussed, trust can also create herding effects. Can we incentivize people to disagree when appropriate? Research suggests that fostering diversity of opinions is possible by encouraging detection of errors and rewarding that more than maintaining agreement with previous observations, if appropriate. Clearly, examples like peer reviews and quality-control exist, but these concepts are generalizable[143].

The Module 4 "how to do it" section on collaboration also lists numerous blockchain-based decentralized autonomous organizations (DAO) examples, some of which support the deployment of incentives.

No matter what incentive strategy one follows, at some point the scalability of the model hinges on **dashboards and analytics**. For instance, in large organizations it is often complicated to understand who really contributed to what. In the case studies presented in this section, this problem was addressed, and resources deployed against data engineering and visualization and facilitate the provision of formal feedback to the system, for instance as part of individual and departmental KPIs. With that, significant data engineering (particularly master data) challenges may emerge and need to be addressed, typically through dedicated teams and using some of the many specialized tools, like Tamr and Informatica. A more innovative way of solving the problem is the analysis of the network and the identification of the most collaborative people. That analysis could be performed through the SNA tools that we discovered in the previous chapter.

A final topic that we will not dive into in detail is **incentives for markets**. Apart from the general economic theories about market behavior and related incentives, what we are interested in to generate a market-driven supermind is the ability of the market to be rewarded for providing insight, not just transactions. For instance, Gnosis.pm prediction markets conditional tokens (blockchain based) – a decentralized autonomous organization (DAO) - enable more precise risk-adjusted pricing and matching demand with supply of predictions, including the ability to predict, trade, and hold funds.

Much of what we said so far lends itself to creating a dashboard that monitors the “incentives supply chain”, that is the flow of creation of incentives, and their distribution. That applies to both extrinsic and intrinsic ones. For instance, in the case of the reskilling supermind, some key metrics being monitored are:

a. Feedback from the users to the key nodes
b. Key nodes’ progress toward success, both in absolute terms and compared to others – and related intrinsic benefits derived from being exposed to senior management
c. Learning score for the learners, especially those who perform particularly well who obtain rewards for that behavior

A separate set of considerations applies to evaluating the benefits of collaboration platforms and making the respective **business case** that incentivizes senior management to allocate financial resources for building and maintaining them. The costs of these systems have decreased radically but are still a factor. Technologies like Microsoft Office 365 are a typical example, where the benefits claimed by their vendors, and dissected by technology analyst firms, need to be evaluated when building the financial case and can be broken down across: the value of holding real-time meetings instead of waiting for scheduling to happen; document and enterprise collaboration, enabling teams to co-edit prose, slide or spreadsheet type documents among others; team messaging that substitutes the more formal and cumbersome email flow; similarly real-time customer interaction; more efficient use of collaboration time and resulting time savings, including more efficient meetings; reduced time switching between applications; reduced employee turnover thanks to better employee engagement and feeling of being integrated; more efficient resource onboarding; time savings working with outside companies; frontline worker collaboration; generally faster time to decision; reduction of overnight trips; and reduced IT compliance[144] thanks to a more explicit infrastructure that can be audited by information security specialists.

Module 3: the "what". A knowledge and information feeder

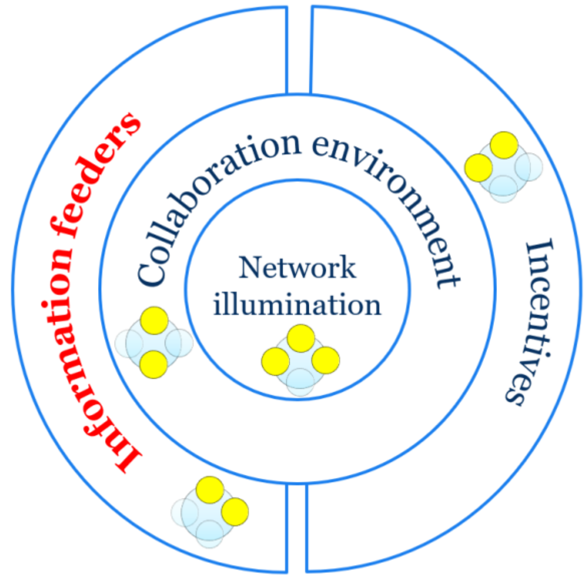

Let's start with a simple example: anyone involved in an innovation project typically needs to know more than they already do. They need to access the latest developments for inspiration and to avoid reinventing the wheel. This is a form of sensing and remembering, traditionally done by exploring knowledge from a variety of sources: school, executive education, trade magazines, and events. For instance, marketing departments have collected press clippings for decades, curating news about competitors. Scanning the environment is the standard when forming strategy. One of the best examples is a disciplined scan of multiple ecosystems and their nodes (education, infrastructure, government, geopolitics, economy, public health, demographics, environment, telecommunications, wealth distribution, and technology)[145].

How does a collective intelligence enhance those (and related) processes, preventing our collective intelligence from being blindsided or reverting to groupthink?

Let's take a step back. The development of cognition in any sentient life form requires external stimulation. Neuroscience indicates that our senses play a large role in the development of our intelligence. Language, in particular, allows us to interact more deeply with the world and enables us to greatly expand our cognitive capabilities. Indeed, language changes the very physical shape of our brain to accommodate higher levels of cognition.

Similarly, machine learning and most AI methods require a significant volume of external stimuli to make (narrow, for now) sense of the world and shape the neural network so that it can learn how to predict unknown situations more accurately. Remarkably, even comparatively simple cognitive structures like AI's neural network perceptron layers (basically arrays of probabilistic numbers) can perform complicated cognitive tasks like image recognition with somewhat simple instructions. That is, *if* they have the scale of multiple layers and they're fed enough data.

Bottom line, a relatively simple cognitive system, if supplied enough information, can exhibit intelligent behavior. What's the implication for the design of a supermind? That a purpose-built knowledge feeder enriches the data that a group of human and artificial minds, a supermind, can process - and arguably ignites a higher level of cognition. This isn't an abstract philosophical concept or merely a belief. Consider the following:

A knowledge feeder is a better way to surface relevant ideas without the limitation and bias of standard search engines, with the additional value of finding adjacent ideas filtered by relevant people or machines. A primitive yet useful knowledge feeder could simply be obtained by deliberately and systematically sifting through relevant web searches, constantly following relevant people, and feeding the results to the networked intelligence. As simple as this may be, it is rarely done intentionally or given appropriate capacity in conventional organizational design or project management.

The first thing a knowledge feeder does is enable a certain amount of **sensing** of the environment. It brings signals from the world to the doorstep of the network nodes. For instance, a marketing manager can see competitors' moves, or multiple industry experts in a community can witness the launch of a new product. Or natural resource knowledge, for instance in threatened ecosystems, can be surfaced by harnessing the input of people on the ground[146].

Intriguingly, these sensing knowledge feeders don't require a significant amount of cognition to deliver effective results. Ants emit pheromones (chemical traces) that enhance their collective behavior by making each individual ant a sensor of the environment around them. This optimizes routes used for the scouting of food with remarkable accuracy. Interestingly, ants don't just optimize the search direction. They also optimize the amount of "ant resources" deployed in the search, to ensure their collective energy is used correctly at the right time. Some do so by simply modulating the speed of exit from the anthill, achieved by sensing the speed of incoming ants. Returning ants bump into those who are ready to exit. The faster the rate of return (signaling more food available, because it took less time to find it) the more ants are sent out. A bee's "dance" serves a similar purpose.

These processes aren't complicated, nor do they require particularly intelligent *individual* parts. A simple computer program written by a high school student could replicate their basic functioning. What makes these algorithms superior is the scale of the collective, and its ability to generate enough natural experiments to identify what works across many different situations.

It is plausible that the same applies to human-machine groups. They're able to efficiently sift through mountains of data generated in their environment, from Internet of Things (IoT) devices to local news. For the last 20 years, knowledge has become easier to retrieve. And with the addition of new sensors, new knowledge is created in previously dark areas.

A knowledge feeder also helps the collective mind "remember" what has worked elsewhere. For instance, medical- school instruction, or crowdsourced environments like PatientsLikeMe can help medical communities get smarter about what has worked in specific circumstances in the past, sparing the individuals the need to form those experiences directly[147].

In many companies, a sort of ubiquitous collective knowledge feeder already exists through the convergence of three things: electronic communications (email and collaboration environments like Slack or Microsoft Teams); organizational structures that include subject matter experts (think *practice leaders* in consulting firms); and pervasive information feeds fueled by the explosion of internet-based knowledge sharing, from social media to specifically curated knowledge feeds.

Looking at this list you may wonder: aren't Twitter or Google enabling all this already? The answer: not fully. For all its might, and with some notable exceptions, many of Google's development efforts have focused on improving search algorithms to make advertising more relevant to users, and on the retrievability of the data (Google's BigQuery), not the creation of knowledge feeders. An issue with traditional search algorithms is that they don't focus much on *network edges* - that is, the connections between people. For example, the connectivity between person A and B (measured from social media and citations) isn't the primary input an algorithm uses to return results. If I search "Crohn's Disease", the algorithm won't surface people who relate to individuals who talk about that search term more consistently, and whose webpages are visited and connect to others extensively. That's OK when it comes to identifying topics to sell contextual advertising, but it might miss people with additional and related knowledge who don't talk about it or are not well-referenced on the web. That's why we don't rely on information feeders to be built solely by giant search curation engines.

This said, the data surfaced by Google is being increasingly mined in ways that could create solid knowledge feeders. For instance, Google search queries for instance have been used to gain a new perspective on social dynamics. They can identify to the gap between declared and actual behavior (many lie to pollsters, but few lie to a search engine text- input field)[148]. Another obvious application is healthcare: Google Health was discontinued in 2011 but examination of search queries haves, for instance, been used to shed light on health needs in the African continent by exploring which demographic segments are interested in certain types of cures[149].

In the case of Twitter, the original model rested on the generation of "knowledge graphs" based on people and semantics connection, and surfacing the respective tweets. Easy enough to do with the AI and computational power of the 2010s. The problem though is that to keep most people engaged, the content would need to "scream" novelty, as opposed to provide deep insight. Bubbles and divisiveness ensued. Fueled by the incubation of "hypercuration" ideas made by companies like Substack and Medium, in late 2020 and early 2021, a variety of alternative models started emerging. Forums like Clubhouse and Twitter Super Follows showed that smaller audiences, relying on a more intimate connection with the creators, could also be financially attractive. The result is, hopefully, more thoughtful, insightful, and civil information feeders.

Beyond these examples looms the opportunity to harness the varied contributions of a network of *people/machine nodes*, whose serendipity is at the root of innovation. As discussed earlier, good ideas often come from the combination of points of views of diverse networks of people[150]. But the sharing of information sharing doesn't always require active human intervention. The flow of social media curated for each of us relies on algorithms that examine our Twitter network, or the communities on Reddit or LinkedIn. Or consider the globally- sourced news repository G-DELT that disgorges machine-translated news from thousands of publicly accessible sources worldwide; or the repository of free internet page structures offered by Commoncrawl.org.

And in one of the most significant recent advances, OpenAI's GPT-3, a language model trained on one of the largest knowledge datasets in the world, is increasingly able to summarize natural language and other text-based knowledge. Think of it as a "auto-complete" on steroids, but with an unprecedented ability to synthesize information, and even doing so in fascinating ways, for instance by using a specific corpus of data that would, say, enable it to simulate the answers that a person (say, Alan Touring), would provide to a specific question. As of early 2021, the software was still unreliable and the outcomes' quality heavily depended on the underlying dataset's biases, but the promise is enormous especially when combined with human judgment at scale. So much so, that this strand of AI could become the fastest to reach some degree of general collective intelligence (AGI). And remarkably, it would do so because of the reliance on the underlying human knowledge. Initially engineered by Google, Transformers are a new architecture that enables computers to understand language with radically improved accuracy, and they're architecturally the fourth archetypes of neural networks (after convolutional and recurrent neural networks, and multilayer perceptrons). Facebook launched in 2021 a *speech*-generation model (called Generative Spoken Language Model or GSLM) based on similar technology.

While we are not reliably there yet, imagine the possibilities: uncover ideas for climate change resilience that have been tested elsewhere (a form of scalable remembering) or measure the local public's reaction to the introduction of climate change related regulations (a form of granular and scalable sensing); enable faster reactions by authorities and the international community to water scarcity violence by sensing initial signals from news originated in upstream places, like the Himalayan Pakistani-Indian water basins.

Before moving to the more detailed section related to information feeders and their implementation, let's once more reflect on the role of AI. Without it, we wouldn't be able to curate content at scale and speed while maintaining relevance and accuracy. AI also helps surface the sources of that knowledge, both machines and people, and connect them in ways that would have been completely impossible only ten years ago. The diagram opposite illustrates these aspects.

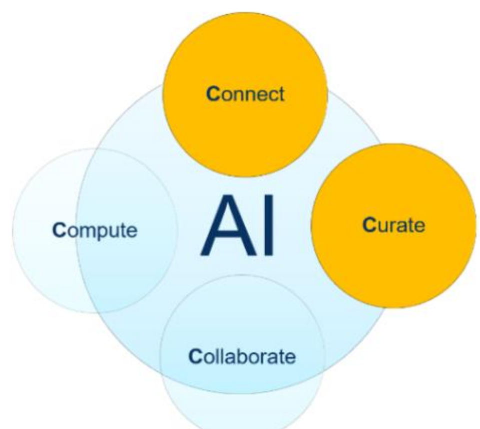

In the technical notes we will explore even more advanced methods and technologies that comb human knowledge wherever it is publicly logged, from the web to patent filings to news. And we will see how AI does so in increasingly effective ways, well beyond simple semantic searches.

Next, let's illustrate the application of this module with through real-life case studies and examples.

Examples - Module 3

Distributed information feeders, often organized around common standards by ecosystems and even communities, are a significant part of global sensory systems.

Some early examples have been around for decades, such as Xerox's Eureka, one of the first community-based technical help forums, where technicians could exchange tips – and saved the company millions of dollars.

During the Covid-19 emergency, information sources became the lifelines of many large organizations, from governments to enterprises. Misinformation and just plain mistakes were rampant, with social media often amplifying the most outlandish memes. The accuracy of preprint scientific papers and repositories such as Biorxiv and Medrxiv, in an effort to speed up the sharing of information, became a more concern, so much so that MIT Press launched "Rapid Reviews", a peer-review journal that uses artificial intelligence to curate the most promising papers, and then commission peer reviews in collaboration with University of California Berkeley. The journal is hosted on PubPub, referenced in the Module 4 chapter of this paper in this report as an interesting example of infrastructure for facilitating collaborative editing effort. Arxiv, the grandfather of pre-prints' repositories started in 1994, has seen tens of millions of downloads, and has spurred the creation of specialized-domain pre-prints, such as PsyArxiv and SocArXiv – before the healthcare related ones mentioned above. On a related note, many call for even freer access to scientific literature, like the creators and users of Sci-hub, a (currently illegal) repository which, according to Wikipedia[151] "is a shadow library website that provides free access to millions of research papers and books, without regard to copyright, by bypassing publishers' paywalls in various ways".

The recent effort to contain Covid-19's spread has also heavily relied on distributed sensors, tracing the movements of people at risk through their cellphone GPS signal. The same data is mined to understand how fast and how easily the disease spreads, which gives important cues to epidemiologists and public security officers. In the wake of the pandemic Apple and Google started cooperating to create robust solutions that help collecting data while maintaining privacy controls, so that democratic governments can use them.

These efforts use distributed sensors harnessed by hierarchical powers (the public security branch of governments). Many other examples of collective intelligence rely on more organically organized structures. During an earthquake hitting Nepal, Facebook's Safety Check located millions of people automatically and notified 150 million about their safety almost in real time[152]. The famous Safecast[153] is a classic citizen science effort. Safecast is the collective, citizen-led response to Fukushima's nuclear reactor failure, when hundreds of DIY radioactivity sensors were deployed by laypeople, enabling a fast response to the crisis. Similarly, Takingspace.org's Airbeam enables crowd-based "*aircasting*", i.e. the monitoring and analysis of pollution exposure across large groups of people. A competitor, Awair, is doing similar things. Co-timed with Earth Day 2020, Microsoft launched its "Planetary Computer" project, aimed at collecting, storing and processing (through AI) data sourced across the planet, and make it available to researchers, including citizen scientists[154].

This effort is reminiscent of the Air View[155] project where Google Street View cars measure air quality as they drive through cities. Google's launch in mid-2019 of the Environmental Insight Explorer (EIE)[156] is another interesting development, as it offers the ability to append environmental data, measured or estimated, to geospatial layers enabling further insight for both organizations and end users. One of the first use cases is to help cities manage pollution.

In another example, a network of sensors tracks biodiversity degradation (iRecord in the UK[157]). Expanding on GEOSS (Global Earth Observation System of Systems), the Citizen Observatories projects are funded by the European Union. One of them is GROW (soil-moisture citizen observatory) aimed at climate change and including scientists, NGOs, communities and individuals. Computer vision sifts through mixed-source (people, machines) imagery from both people and machines to enable better agricultural decisions at OneSoil.ai.
The iNaturalist community (an effort run by the California Academy of Science and the National Geographic Society) uses specialized computer vision specialized classifiers to allow the mapping and sharing of observations of biodiversity across the globe using a free mobile app[158] and has contributed more than 60 million observations of more than 300,00 species[159]. Similar concepts are applied elsewhere. A similar example comes from Cornell University's eBird, aimed at obtaining pictures (over a million logged so far) and bird songs – which are then run through an AI application called Merlin. The results are then used, among others, by scientists,

like in the case of State of India's Birds.

Mercicorps' Agrifin provides a platform for African countries to use multiple sources, including crowdsourced input, to detect the spread of locusts[160].

Platforms to help users act are also appearing, such as in the case of Oregon-based Wild Me whose software provides image recognition to cameras positioned in the wild – so that individuals can be tracked as they move in their habitat. The Grevy's Zebra Trust, for instance, has used the platform to support volunteer and scientists rallies in a wildlife conservancy area. The emerging risk that these platforms need to manage is that they may inadvertently give poachers a way to find animals. Authentication and identity control can be a solution, not unlike in other parts of the web.

The Mexican state of Tabasco has been mapping human behavior in response to adverse weather events, combining mobile and other sensor data, to pinpoint needs and interventions[161]. A comparable effort is at play in Dar es Salaam, Tanzania, where the Ramani Huria project, (based on OpenStreetMap – an open-sourced mapping effort), helps local communities map relevant infrastructure and land features especially in informal human settlements. Europe's BlackShore.eu uses gamified crowdsourcing to create maps.

Another example comes from the current effort to map the ocean's floor. The Nippon Foundation and GEBCO Seabed project, launched in 2017, has by now mapped over 10% of the Earth's oceanic floors. The next phase will include crowdsourcing of data by appending echo-sounding equipment to all sorts of vessels, based on their routes. Along similar lines, swarms of bots could possibly one day act as sensor networks canvassing larger spaces in a dynamic, adaptable manner[162].

Similar strategies are applied by HealthMap.org that in their words "deliver real-time intelligence on a broad range of emerging infectious diseases for a diverse audience including libraries, local health departments, governments, and international travelers. HealthMap brings together disparate data sources, including online news aggregators, eyewitness reports, expert-curated discussions and validated official reports, to achieve a unified and comprehensive view of the current global state of infectious diseases and their effect on human and animal health. Through an automated process, updating 24/7/365, the system monitors, organizes, integrates, filters, visualizes and disseminates online information about emerging diseases in nine languages, facilitating early detection of global public health threats." That's an information feeder based on collective intelligence, and arguably makes a range of healthcare stakeholders smarter already.

At the other end of the spectrum, an array of companies, from Mechanical Turk, to Appen, to Accenture and Genpact are using crowds to train algorithms. AI crowdsourcing is also being used to avoid bias, like in the case of Mozilla's Common Voice initiative where crowdsourced voices train speech recognition software.

HumanDX.org is an effort to improve patients' therapy by combining human practices through machine learning. It is "a worldwide effort created with and led by the global medical community to build an online system that maps the steps to help any patient. By combining collective intelligence with machine learning, Human Dx intends to enable more accurate, affordable, and accessible care for all". "Let's end unequal access to medical knowledge" is their motto.

OpenStreetMap is a similar effort. For over a decade thousands of volunteers have toiled to provide an open-source alternative to commercial geo-mapping, further annotating and correcting computer-vision-detected terrain. Additional uses of this networked intelligence include rapid re-mapping of areas impacted by natural disasters. BlackShore.eu operates in a similar space.

Intelligence also has another meaning, and collective intelligence methods apply there as well – as demonstrated by the emergence of open-source intelligence methods (OSINT)[163]. The Syrian Archive[164] is an initiative using tools that scrapes social media feeds to detect and classify imagery pertaining to war equipment, with the intent of documenting and making searchable evidence of war crimes. The neural net AI curates and classifies information automatically. Bellingcat uses crowdsourced information for investigative journalism – resulting for instance in clarifying responsibilities in the MH17 flight downing by a military missile over Ukraine, and a similar situation in Iran. It also runs workshops to train people on using web resources for research investigation. Europol, the European Union's policing agency, harnessed the resources of amateur detectives to sift through photos and identifying clues in child sexual abuse cases. More recently, readily available satellite pictures helped students identify ballistic missiles silos being built in China, by leveraging Geo4Nonpro's community infrastructure. Routes of airplanes (e.g., through FlightRadar24) and ships are available publicly, which adds data to understand, among others, CO_2 emissions sources.

A different set of possibilities is opened by AI helping people to write, such as what Hyperwrite.ai (powered by Open AI's GPT-3) does. Additional possibilities of democratizing content production could be unlocked when combining AI-generation capabilities with storyline tools such as Lynit.app.

The ubiquity of sensors, combined with the availability of machine learning and other computationally sophisticated methods and infrastructure, is also bound to change the field of economics. GDP measurement is a hierarchical effort, lacks granularity and speed, and is subject to significant manual intervention. Information like restaurant booking OpenTable can help surface early signals to sense changing conditions and inform trend estimates in ways that were previously impossible[165].

The explosion of data sources will no doubt increase the extent of these efforts, and open new possibilities. More examples are in the "How To Do It" section of the Guidebook, as well as the database at www.supermind.design/database, ranging from environmental to infrastructural and even cultural heritage protection. But first, let's explore one in more detail.

In-depth case study module 3 – upskilling at scale

In the case of reskilling of a large workforce, information feeders are applied at multiple levels.

First, all knowledge nodes, (the experts at the center of learning communities who curate content) are supported by knowledge managers who help them filter and repurpose external information. Those knowledge managers are often the same who work on knowledge management across the company, above and beyond their role in the reskilling program.

The experts have their sources of information – from trade contacts to client engagements – and can tap into their pre-existing networks. The networked intelligence model applied here enables them to expand that network both externally (as knowledge managers help comb external sources more thoroughly) and internally. The latter is quite important. Thousands of people now provide the experts feedback from the natural experiments that they carry out daily. They also re-direct to the experts any interesting tidbits of knowledge that they might come across. As a backbone to all of this, multiple Microsoft Teams channels facilitate the interaction between experts and knowledge managers.

In a learning supermind the content curated acts as an information feeder for the learners. Specialized curation and learning portals identify, sort and allocate the content to the right places, such as a learning platform like EdCast, or a central knowledge repository built on a common Microsoft SharePoint. The social nature of these platforms makes it easy for users to provide feedback and share with others, further improving both relevance and salience of the content.

Module 3 – How to do it

*[Reminder: the specific examples documented in the technical note section may be used to compile a list of "portable ideas" that enable more creativity during the design process (see the "*The process of designing an augmented supermind*" chapter)]*

The technical notes for this module will cover the following areas, which make a good checklist for teams intent in creating a supermind. These notes offer practices and inspiration, as well as structure the challenge so it can be addresses.

1. **Are we surfacing the output of collective intelligence through conventional or advanced surveys** (and related research firms), taking advantage of web-enabled capabilities? (Section here)
2. Have we created an **explicit description of the semantic space** of the concepts that we are exploring, so that we can both search for that language, as well as the people who use it – and the people who are connected to them that might extend it? (Section here)
3. Have we established a comprehensive **net of listening posts** that continuously mine that semantic field and the related people network's input? (Section here)
    a. Specifically, and when needed, have we created a **network of ground level "sensors"** that generate "ground truth" irrespective of what's already available on the web, ideally by using the new semantic structure? (Section here)
    b. Have we used existing knowledge management infrastructure to feed the network? (Section here)
    c. Are **Business Intelligence (BI) systems** also woven into this flow of knowledge, as another source of ground truth? (Section here)
4. Have we created a **machine-augmented human *curation* layer on top of the above**, to filter and route the right input at the right time? (Section here)
5. Have we specifically created feeders **for scientific and patent-related literature** if appropriate? (Section here)
6. Are our **learning and development** (L&D) systems being fed by (and become feeders of) this system? (Section here)
7. Are **available resources allocated** so that the flow of knowledge is mined effectively and efficiently? (Section here)

Information feeder's building blocks

A knowledge feeder is, at its core, a curation engine. That is, an environment that makes the timely retrieval and dissemination of relevant and salient content frictionless. Most teams and all organizations have some sort of process to do this – the problem is, they're very often amateurish at it.

Traditionally, **surveys** have served that purpose – and a whole industry was born around it, including recent developments like frequent engagement enabled by the internet. Companies like SurveyMonkey have made the administration of simple questionnaires online (asynchronous) extremely easy. Semi-synchronous polling, for instance on sensitive topics during employees townhalls, is possible through tools like Feedy.

Synchronous, real-time polling is possible through a variety of collaboration platforms, like videoconferencing technologies e.g., Zoom, Skype for Business, WebEx etc. (more on this in the next chapter). For more insight on the real-time views of large groups, tools like remesh.ai can be useful. For more sophisticated visualization of asynchronous opinions, other tools like Crowdscope.com exist. Allourideas.org is a variant of surveys that can be administered through wiki, with an in-built statistical analysis that lends itself to deciding between options during a debate. Ureport.in enables polling through telephone messaging, thus making it accessible in developing countries where internet connectivity is hard to achieve. Apptivism.org uses custom-built bots to engage constituencies and collect their feedback. The US department of state used a combination of traditional survey and AI-curated collaboration with its citizens to improve its passport services in 2016, through a participatory platform called insights.us[166] - the platform is now available to cities and government agencies.

And sophisticated text analysis, able to connect themes in large sets of data, is increasingly possible even for inexperienced analysts, through tools like Infranodus.

At the most sophisticated end of the spectrum sit consulting and research companies, with very mature yet often somewhat traditional practices (large knowledge management organizations, and "practice leaders" surrounded by subject matter experts). In recent years, we have seen the emergence of nontraditional players that claim a different position in the curation ecosystem by utilizing social media (like HfS Research, focused on the back-office operations market) and natural language processing engines associated with advanced analytics (like CB Insights, that monitors the venture-capital backed space).

An effective knowledge / information feeder identifies, selects and propagates the right information in the right channels, and solves the problem of proliferation and duplication of efforts that most companies incur – illustrated in the picture below. Multiple people search the web and filter with raw google and other searches, follow individuals on twitter, LinkedIn or ResearchGate (for scientific researchers), subscribe to newsletters, and in rare cases through RSS feeds or specific curation applications. That's inefficient and doesn't scale well, as knowledge will tend to stay siloed with the recipient and their immediate network.

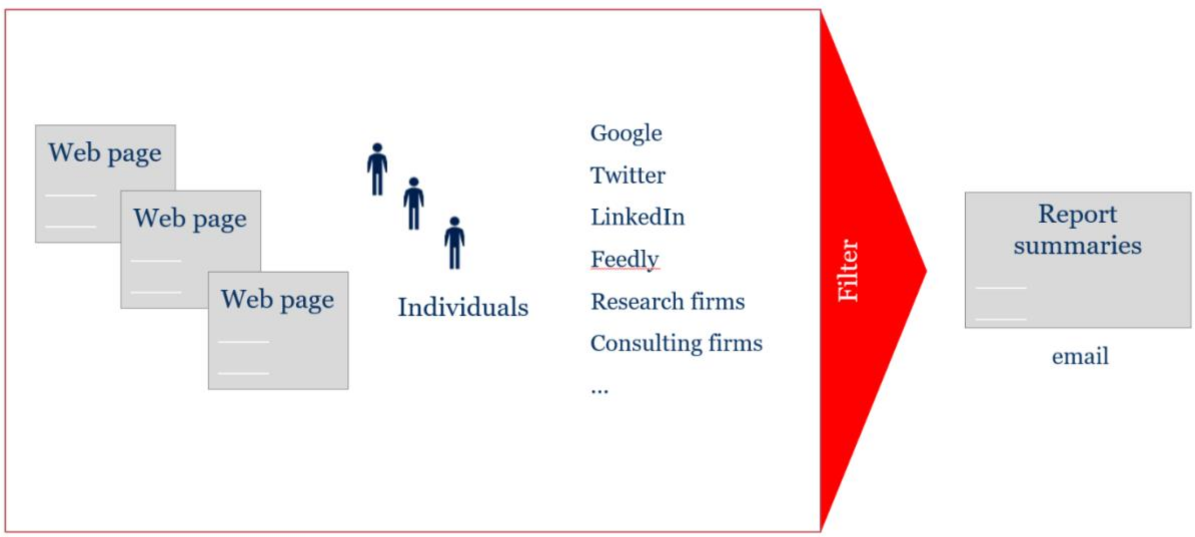

A more effective model instead is the one illustrated below and outlined in the following bullet points. Throughout the process, the effort is a combined machine-human one, ideally facilitated by the construction of a graph database.

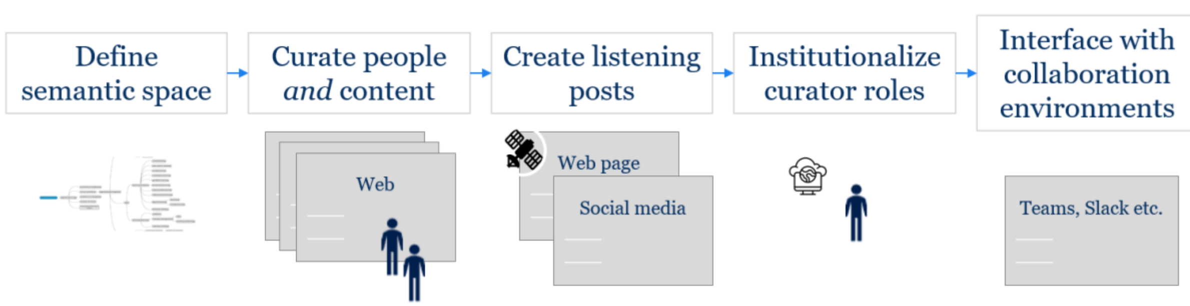

It starts with the **identification of the words** that describe the space being searched (semantic network that delineates ontologies e.g., words that are ideally MECE - mutually exclusive and hopefully collectively exhaustive). That can be done either heuristically by asking experts to map out the words and clustering them in groups, or even analytically by extracting so-called "entities" by using R or Python NLTK (or more specialized entity visualizers like DisplaCy) to identify words recurrence in typical content.

Then, the **words are used to find people** – not just content. A simplified schema of the process is represented below. Also remember that Google search can often be configured in a precise way by using the standard Programmable Search Engine[167].

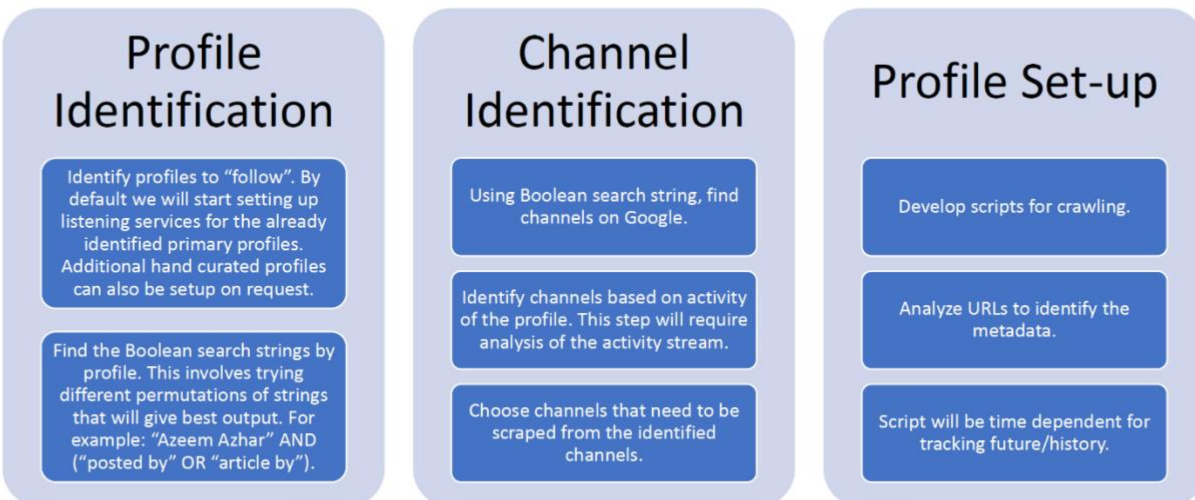

As discussed in the Network section in module 1, people are nodes of knowledge, and identifying them is the first step to sourcing relevant content. One example of the process looks like the following. Keep in mind that a rudimentary version of it could be achieved through brute force curation of web searches, and filtering of information coming from relevant nodes by following them for instance on Twitter, LinkedIn, ResearchGate and the like (when identified as explained above). At small scale, and with enough manpower, this crude method can be effective too. The description below applies to a more scalable method. It will no doubt be superseded by better tools over time. More detail on the process followed (e.g., for extraction of URLs etc.) is in the note appendix[168].

Listening posts are then created, typically through some social media tools (like Feedly, Hootsuite, Tweetdeck), but also by giving internal users, typically subject matter experts, the ability to act as “level-1 curators” who, by virtue of their own network, can act as initial filters that complement the listening posts. Curation engines that aggregate content through a combination of human and AI intervention also exist (e.g., Curata[169] but also competitive intelligence tools like XiQ.io or technologies like Discover.ai). The individual person’s ability to efficiently collect and distribute curated knowledge is also the object of recent developments, for example through Factr.

Those listening posts can also obviously point at specialized information sources to facilitate searches that would be very noisy and time consuming with normal search engines. A range of data capture (including specific Boolean queries coded to systematically mine Google with the right filter), text analytics, web scraping tools[170] can be used for these purposes. Kaggle, the data science contest community, was used during the COVID-19 pandemic to write natural language processing algorithms helping scientists interpret insights emerging from research done around the world. Some of these tools, like ModuleQ, can look at both external and internal (documents, customer data, meetings) data, and then convenient inject them into the flow of work, for example through Microsoft Teams. Microsoft’s own Project Cortex (now Viva Topics) is going in the same direction, but for now with more of a focus on internal knowledge.

To tap into news trends, G-DELT[171], a global database of published sources, auto-translated from dozens of languages and available as a public resource or through Google Big Query is a great resource.

Some search engines are indeed customizable and can be turned into usefully curated sources - like Google Scholar, Semanticscholar.org, Microsoft Academic Graph or iris.ai that only focus on academic papers, as well as Google Dataset Search[172], DataPlanet or WRI.org (consolidated sources for environmental and resources) for data. Companies like Dataminr.com then combines multiple sources of text and image data, as well as audio streams, both crowd and machine sourced – in the case of “First Alert” system it uses them to highlight potential crises early. A similar space is the object of the work of Qatar’s Crisis Computing group[173].

To see how graph theory and natural language help understanding complex topics, it is worth viewing the visualizations that www.connectedpapers.com provides. The enable the mapping of the corpus of scientific knowledge and connect papers well beyond citation mapping.

Zencity.io is a solution for city-specific collection of feedback and input, enabling local-level analysis and even support to crisis management. Citibeats.net have specialized in mining people's opinion to help cities and public organizations extract insights from their constituencies. As an alternative, cHistoronnections between philosophical concepts (manually curated!) is at https://www.denizcemonduygu.com/philo.

Sometime the data sought is not readily available on the web but the individuals participating in the supermind can generate some of it and **local network presence to generate "ground truth**". The classic example is the use of mobile phone GPS data by various map applications, from Google Maps to Waze. Many other applications have for years used some form of metadata generated by network participants to generate insights – from LinkedIn to Google. But a better example of grounds-up, community-based activation of a sensing network is Japan's Safecast that deployed IoT devices to detect radiation in the aftermath of Japan's Fukushima disaster. Apart from the explosion of IoT devices that provide additional data at often marginal cost (which can be consolidated through technologies like DescartesLabs to facilitate research and analysis), there is a need for deployment of those sensors, whether IoT or others. Companies like Premise.com offer "ground truth services" in spaces as different as logistics and movement monitoring and risk and they're used for private, government and international development purposes (for instance, to obtain fresh data about Zika's virus spread). The platform enables on-ground operators to perform data surfacing tasks.

Some **human curator**, with the appropriate level of business domain knowledge, can add value to the machine by filtering knowledge and allowing it to be loaded on an appropriate collaboration or knowledge management[174] platform (discussed in more detail in the next chapter) where end users can decide what feed to use. A simplified schema of the process is represented below.

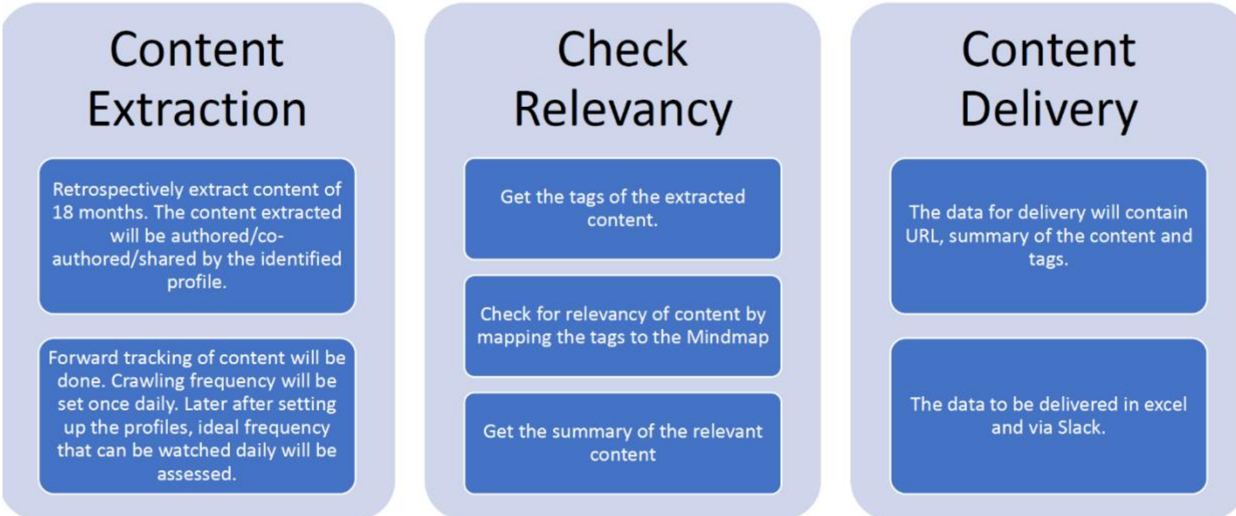

Big data analytics languages like Python are already equipped with extensions to mine social media data. For instance, for Python, Tweepy enables it to connect relatively simply to Twitter's API and analyze its natural language.

Human curators can also use web text annotations tools, like getMemex.com, effectively appending metadata to text parcels of browsing history, in a way that's somewhat like what Amazon Kindle allows to do on kindle-format books. A more collaborative approach to web bookmarking is taken by hypothes.is, that enables communities (for instance, academics, teachers, publishers or journalists) to tag and annotate documents, and discuss them online – providing not only an information feeder but the beginning of a collaboration platform (more on these in the Module 4 – How to do it at page 178). Bridgit.io (in alpha version as of January 2020) enables people to annotate web objects to manually create knowledge graphs. An additional possible use of some of these technologies and the related network intelligence is fact checking (more on this in the "What can go wrong?" section at page 222.), with particularly important implications for multi-lingual situations[175].

With the advances in Natural Language Processing, machines are getting better at helping humans make sense of connections between elements of large content sets – as in the example of Quid.com or Forestreet that can map technology market landscapes and draw connections between market participants and related technologies; or golden.com that provides accurate semantic search on companies or topics of interest, for use in consulting, investment or financial services for example. The company Golden in particular has emphasized the synergy between machine and human curators, by enabling workflows allowing editors to improve content quality. Diffbot's media monitoring, using their knowledge graph's technology, is another example.

SparkBeyond launched in 2019 an automated data exploration technology that sifts through multiple data sources (web, patents, news etc.) and visualizes ideas that are connected to each other (co-occurring entities), hence facilitating the mining of potential solutions to existing problems: think about being able to quickly query that global corpus of knowledge, largely a curation of human-generated ideas, to find solutions to pressing problems – from disposal of plastics to reducing fossil-fuel energy consumptions for residential cooling. Quid.com, iProva and Forestreet work in similar spaces.

Iris.ai focuses on facilitating scientific literature review and helps the discovery by offering the user a structuring of the semantic fields related to the search query, hence enabling the combination of papers derived from various disciplines.

US Department of Energy's Berkeley National Lab used NLP to create word-based bodies of text vectors in its word2vec (machine learning) algorithm, which proved to be able to predict material science's advances[176].

In 2020, BenevolentAI came into the limelight for using machine learning and other AI techniques to identify patterns in existing literature through knowledge graphs. As stated on their website[177] "in the face of the growing global health crisis, BenevolentAI set up a specialist scientific team and launched an investigation using its drug discovery platform to identify approved drugs which could potentially stop the progression of COVID-19, inhibit the cytokine storm and reduce the inflammatory damage associated with this disease. Benevolent's research findings were published in [scientific journals] as a potential treatment with both anti-viral and anti-inflammatory properties, for COVID-19 patients admitted to hospital prior to the development of critical lung damage."

A different approach is used by Omnity.io which challenges' the premise of Google's traditional search (based on citations and connections) to enhance the search of relevant information. Interestingly, Omnity doesn't need keywords, as it is able to scan entire documents (for example, research notes) and find relevant materials on the web, irrespective of the amount of links to other websites. While Google has enhanced its search engines through knowledge graphs that enable semantic search, its algorithms still significantly relies on web links because of the advertising value of well-connected web real estate.

Carnegie Mellon researchers have also devised methods that use crowdsourced groups and AI to identify **analogies** that can be fed to innovation groups to enhance their ability to think in unconstrained ways and can be useful in design thinking exercises like the so-called "alternate worlds".[178] For instance, biomimicry (the use of natural analogies to solve complex engineering design challenges) is a type of analogy-seeking exercise. Another famous example is the design of foldable solar panels for space vehicles that was helped by origami experts. Carnegie's approach leverages collective intelligence of crowdsourced groups able to find analogies in distant fields (something that humans can do well), helped by AI that winnows large quantities of potential analogies before turning them over to humans. This method is interesting also in that it attempts to formalize specific variables that describe the problem to enable nodes with less context about the problem (such as a crowdsourced worker or an algorithm) to build on that input. For instance, the foldability challenge mentioned above would have been deconstructed into variables such as weight, strength, flexibility and then fed to the external nodes to identify viable analogies.

Knowledge feeders should be synergistic with existing **knowledge management** (KM) systems. In fact, many of the internal contributors to KM will be curators of external information, if given the opportunity through an appropriate user experience (UX). In the example of reskilling at scale, the company utilized the KM group to support the subject matter experts, enabling them to use the resources already stored in the KM systems, and further use the new materials to enhance the existing KM system. In other words, the learning and development flow of knowledge was made synergistic with the KM work. Apart from the pre-existing KM system (largely SharePoint based, though many other specialized systems exist[179]), the team utilized hubs built as channels in Microsoft Teams for the subject matter experts and KM support team to interact with each other.

Largely thanks to new cloud and AI technology, documentation and knowledge management technology is a dynamic space, with a variety of best-of-breed players also complementing and competing with the traditional enterprise software vendors – see examples such as Notion, Bloomfire, Slapdash, Slab, Aodocs, Fireflies.ai, Humap, and Airtable.

Knowledge management systems are increasingly complemented by what “insight engines”, that effectively enhance the search and combination of concepts excerpted from the available documents. Companies like Lucidworks, Attivio, Coveo and Sinequa among others operate in this space.

It is important to note the knowledge management role that some nontraditional platforms play. YouTube is arguably today’s leading video-based knowledge management platform in the world. From home improvement DYI to drone photography to convolutional neural network presentations, YouTube has it all and guides users through much of the same search algorithms as the rest of Google. The limitation is also the same – in that semantic search accuracy is balanced with ability to monetize by marketing to homogenous consumer segments. Similarly, managed expert communities are the backbone of user support for many highly-engaging advanced technologies, for instance Apple. There, aficionados post solutions to problems at a level of granularity that even Apple wouldn’t be able to do – and the manufacturer’s community managers play largely the role of orchestrators and quality controllers.

Part of the knowledge search and curation happens with **scientific literature**. While Google and SemanticScholar offer specialized search engines[180], companies like Reed Elsevier have historically consolidated research and made it accessible through their portals, behind a paid firewall, or in databases like Clarivate’s Web of Science. Access to scientific knowledge happens increasingly through open source portals. Project Aiur’s objective is to create an open access to quality-controlled content that isn’t restricted through the traditional publishing model and has launched a bitcoin offering to incentivize scientists to collaborate in peer-reviews, hence disrupting the access to scientific information. One obvious current limitation of the publishing of academic papers is their size and specialized, often convoluted and certainly jargon-laden language, which impairs their consumption in larger audiences. While better search capabilities will simplify some of the access, there's still much to do to make those materials faster to scan and use at scale. Publishing-grade summarization and language streamlining can certainly help there, but at present there doesn't seem to be an obvious economic model for that to happen, except in organizations (such as specialized consulting firms) who can afford such curation services.

Specialized search in specialized data sets can be however quite helpful. Epistemonikos[181], used by the government of Chile to set new health guidelines based on the latest evidence, uses machine learning to identify specific types of reviews in the respective scientific literature.

An additional source for insight is **patent** literature. Google Patents for instance enables the search and reading of full text of patents around the world, as well as identify prior art in non-patent literature. SparkBeyond among others also uses the semantic field search to identify co-presence of words in patent databases, hence facilitating the exploration of connected concepts at the edge of today's inventions.

Learning and development (L&D) environments are, from this standpoint, a special case of knowledge intake. They're specifically made to enable people to learn, and as a result their content is calibrated not for novelty and depth, but rather for clarity and engagement. And instructor driven sessions, the workhorse of traditional L&D, are in themselves a form of knowledge intake. They indeed create a form of collective intelligence in that they pass pattern recognition capabilities onto much broader groups. In the example of the professional services firm, the instructor-led training is largely performed by subject matter experts (SMEs), which enables them to obtain useful feedback from a large network of individuals, including people who would typically not get in direct contact with the SMEs – hence enabling additional "sensing" cognitive pathways in the collective brain.

Interestingly, many collaboration platforms are also knowledge feeders, and are routinely used for learning. While Slack, Microsoft Teams and the like may not be top of mind as learning platforms, they enable individuals to significantly cut the number of network "hops" required to get to people who have the right answers, which in itself is a form of knowledge retrieval.

With the emergence of AI, we will likely see L&D resources woven more naturally into the daily flow of work. Just like Gmail gives Google's advertising search algorithms hints about what you are interested in, the new generation of L&D courses will routinely proactively surface relevant nuggets of knowledge based on what people work on. Some of these features are already in production in some of most advanced learning experience platforms (LXPs).

Traditional **Business Intelligence** (BI) and related analytics dashboard environments are also a form of information feeders. They've been used for decades to surface so-called "ground truth" for management to act.

However, their sensing input has typically been machine-based data emanating from enterprise software systems. In the future there will be increasing sources of data (thanks to internet of things devices) that identify patterns from human behavior (for instance, movements of staff in a warehouse).

Additionally, many lightweight software tools (called "systems of engagement") can be created to consolidate the input of large groups of people. Think of the example of a large professional services firm that runs operations on behalf of clients (e.g., finance, risk etc.), and that can use "alerts" coming from the operations "shop floor" to direct attention and intentionally escalate situations that may become problematic. While this is a simple derivation of classic Lean Management principles, the ubiquity of software will make it possible to collect and filter input from more touch points than in the past, and granularly inform their decisions. Even here, the concepts discussed so far apply: it is important to identify the nodes to be monitored, engage those nodes with the right communications (e.g., a CEO's message that emphasizes the importance of alerts), and then share the insight with the right community of people.

Knowledge flows may be overwhelming, and yet it is necessary to be strategic about **allocation of resources** towards the most useful ones. Wikipedia and other crowdsourcing websites (like Zooniverse) have done that for years. One useful analysis is to identify where the knowledge supply / demand gap is biggest – in the case study of reskilling at scale one can see it from the stats from knowledge management and learning systems, where it is possible to track the number of people accessing content of a specific type, and the respective amount of content available. At the other end of the spectrum, individual human users will increasingly receive insight and nudges on their allocation of time spent interacting with the network, so that they don't inadvertently disengage or burn out due to over-collaboration. Microsoft Insights, with important new analytical insights features launched in late 2019, will stay for some time the most ubiquitous tool available, as it consolidates and processes data from Outlook. Apart from creating awareness on gaps, incentive design is also helpful in ensuring the supermind allocates resources as needed.

Like in the previous modules, there is often value in carving out capacity to monitor the key metrics related to the sourcing and usage of information and knowledge. The following metrics are some examples of what can be measured in order to understand what works and allocate resources accordingly

a. Sources of information that are most used
b. Sources of information whose content is most liked
c. Strongest users of determined information types
d. Other factors (and combination of factors) that explain the variance in use, that can be modeled as features in a data model

In the example of the reskilling supermind, the most liked and used content was an important part of the feedback loop to the curators (both subject matter experts and respective knowledge management support).

In the **future**, it is likely that sophisticated and increasingly automated knowledge feeders will evolve more independently from their human users, based on the feedback that they receive. Machines will also possibly route knowledge to be curated to the right people at the right time, after an initial pre-processing (for instance, tagging and prioritization) hence making the human counterparts much more efficient. This already happens with traditional workflows. It is also possible that machines will enable much deeper learning, by acting as interpreters of "natural experiments" that happen in a broad ecosystem. Think again of G-DELT news database, and its ability to identify newsworthy actions taken by billions of people: from adaptation to water scarcity, to early detection of potential conflicts due to environmental stress, to surfacing of the perceived popularity and adoption of public policies, the world becomes a petri dish that teams elsewhere can read, avoiding the proverbial reinvention of the wheel.

Module 4: the "how". Enable the network's nodes to interoperate

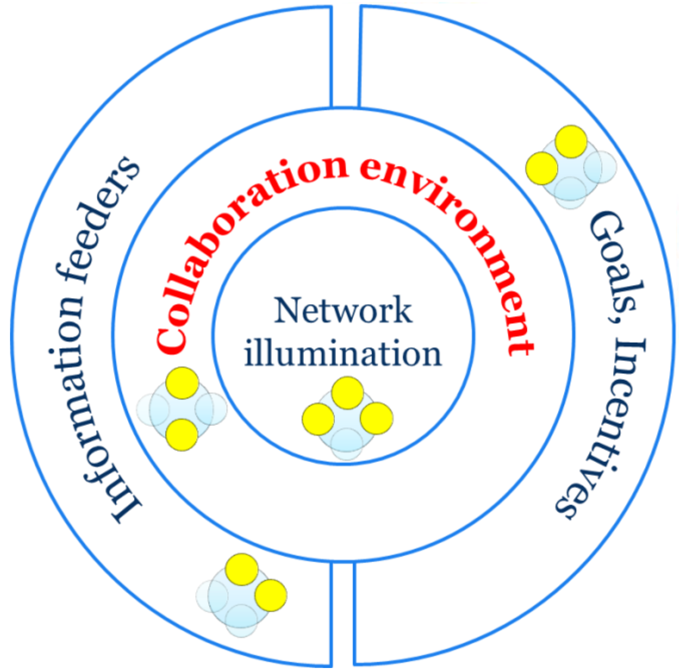

The next stage of connected intelligence growth will come from the use of an appropriate interaction platform. The COVID-19 emergency highlights the importance of collaboration platforms and the related new ways of working – and how they can evolve, adapt, and arguably succeed in delivering outcomes in the face of disruption. The takeover of Afghanistan in 2021 by the Taliban is a somber reminder of how effective the pervasive use of comparatively simple collaboration technology (in this case, WhatsApp groups) can be.

Let's start with the evidence of scientific research[182], which reviewed a large corpus of existing work and revealed the most critical elements in effective collective intelligence in groups: the process of collaboration (in particular, the congruence - cohesiveness and complementarity – of the members' skill, as well as the collaboration steps) is the dominant factor, way more important than the individual skills and other characteristics, such as IQ. These findings need to be extrapolated at network level, because that's how organizations and their ecosystems do the most meaningful work.

Why is a low-friction network collaboration platform so important? Because **combination is at the core of innovation**, and combination happens best when people interact – ideally, with some form of intelligent machines to complement them. Bringing together different perspectives generates value by creating hybrids between diverse perspectives, and those hybrids generate breakthroughs. Doing that with a finite number of people has been the focus of traditional approaches to creative-group dynamics, including design thinking. However, fostering combinations at scale – *very large* scale – requires a different infrastructure.

The current state can be quite problematic at times. In the absence of low-friction collaboration between distributed nodes, **global organizations tend to become bloated and slow**. Think for example about the regionalization of operating structures, where similar resources need to be replicated in each region in order to facilitate local collaboration. This organizational design consumes resources in each region, resulting in loss of economies of scale. Consequently, people in each region are forced to make do with the resources that will be available to them at a specific location at a specific point in time. This might saddle them with a suboptimal match between needs and available capabilities, and potentially force them to delay important interactions until the requisite resources become available. Another consequence of this situation is that senior citizens and many working mothers struggle to meaningfully insert themselves into the flow of work because of logistical constraints.

Now imagine a world where you can access the ideas of the right people - people who have the right skills and are exposed to the right inputs – not just the people who work for you or are located close to your team. Imagine the creation of a meta-city and its exponential innovation power, where each citizen's address is just a time zone.

Imagine an intelligent entity using those networks. First, being able to *sense* the external conditions, for instance the places where the potential demand for a new product is stronger and where potential consumers are. Or collectively *remembering* things that might have worked (or failed) in the past. Then having the capacity to harness all of that to *create* ideas and then *decide* what's best. Harnessing that cognitive power irrespective of distance could even open the door to a new way of living, where talent can reside geographically close to their emotional ties and remain functionally close to the professional ones. While we are far from being able to dispense with in-person interaction and the motivation it sparks, we are already able to harness the ideas and creativity of people and machines irrespective of distance. We can, that is, if we use the right platform and the right methods.

Don't lose sight of the goal of all this: creating new ideas. For centuries, capturing and **combining analogies and adjacencies** has been the spark behind innovation's fire. Traditionally, the people best able to harvest that power have had a combination of cognitive skills and contact networks. They have intuitively leveraged their personal network to amplify their own power. They knew how to broaden their horizons and seek new connections with new ideas. They knew the right ways to "share the sunshine", that is to **give credit and share benefits** with others when those ideas yielded results. They could spot early feedback and iterate systematically. And they were keenly aware of their role in the network, be it a connector across disparate groups, an expert at adopting and adjusting faraway ideas for a local group, or an energizer that galvanizes network nodes. But how do we teach those behaviors? And how can we ensure that new disciples can connect at hyperscale, now that distance matters less than before?

Interaction at hyperscale isn't trivial. To start with, an interaction platform for connected-intelligence networks needs to **address structural tensions**[183]: competition is a strong motivator, but hard problems require collaboration; evolution of great ideas takes time, but individuals' attention is fleeting; honest dialectics and strong points of view are necessary, but people struggle to establish enough the confidence in each other to challenge others openly and objectively.

Another problem pertains to the **establishment of a common language**, especially among people from different backgrounds. Here, the best strategy is to *de-jargonize*: that is, counter the tendency of domain experts to use accurate but obscure technical terminology that prevents others from participating productively in the conversation.

An additional challenge relates to the **natural tendency towards specialization** that emerges in environments with frictionless connectivity. Specialization is useful because it increases the competence in the network, and it happens naturally in complex networks like our own brains. But the flip side is that unfounded and irrational ideas can attract enough people into specialized echo-chambers that legitimize the validity of those beliefs, such as in the well-publicized cases of "flat-earthers", anti-vaxxers and QAnon.

A different but related **problem in frictionless environments (virtual ones tend to have less friction) is the difficulty that those systems have finding equilibrium** in the absence of well-designed checks and balances. In other words, slowing down the *data-decision-action* flow may be beneficial when network intelligence may be faulty – for instance, when human-bias incentives fuel the propagation of misleading political advertising[184].

And finally, ideas don't naturally combine unless they're **modular** enough to do so. Often, people are required to enable that combination, especially when ideas are "packaged" in the form of natural language. But as the evolution of open source software shows, deliberate modularization (structure, and process for changes of existing modules) is critical to the scalability of the distributed-creation model.

And obviously, hyperscale collective intelligence needs to contend with the challenges of distributed work. We will talk about this more in the "What We Need to Get Smarter About" section, but for now let's just keep in mind that remote work groups have often counterintuitive dynamics of their own – for instance, they seem to succeed faster, but fail slower, than their in-person counterparts[185].

So how do we intentionally design better hyperscale collaboration?

While the solution needs to be human-centered, let's start with technology – as it has historically been a limiting factor. There are many types of collaboration platforms, some of which are explored in the technical notes. Each type of supermind (hierarchy, community, democracy, market, ecosystem) requires different functionality. Also, the best group dynamic requires different functionality: contests where collaboration is secondary and instead the value resides in surfacing many different perspectives (markets and ecosystems tend to do this well); collaboration where the results largely depend on the ability to work together with others (communities tend to do this well). One size doesn't fit all.

An additional challenge that very often goes unmanaged is that **preferences in collaboration technology varies by generation of workers**, as shown in the next chart.

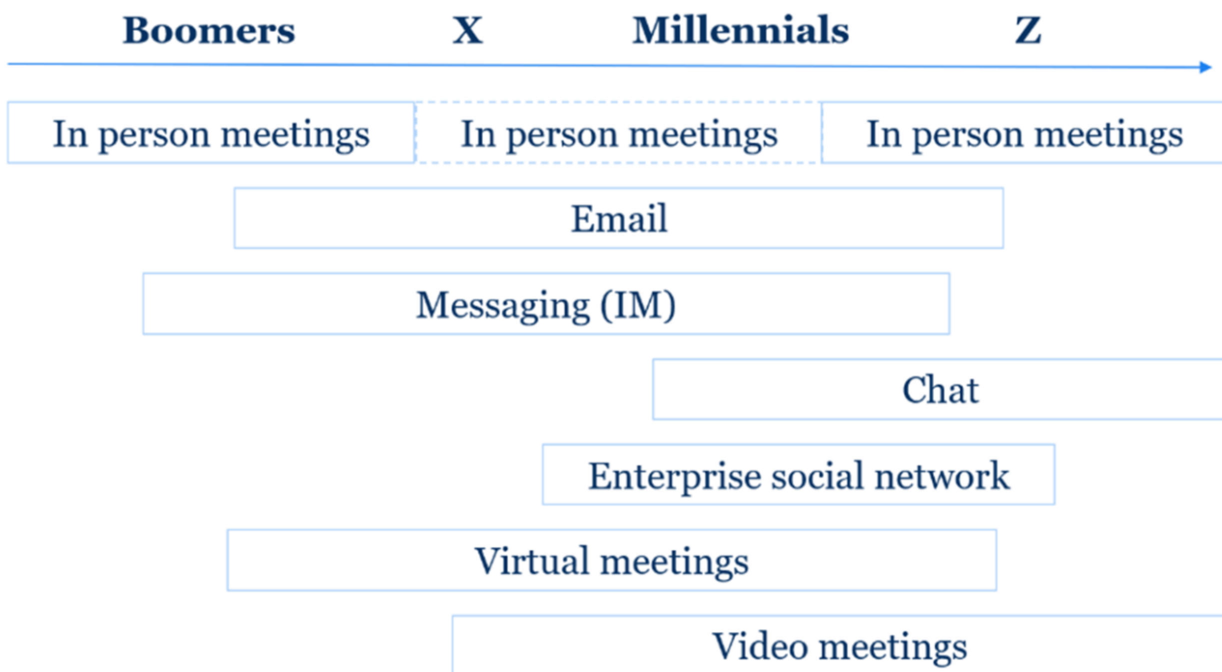

Regional preferences also complicate matters. For instance, because of the disruptive growth of WeChat in China, the use of chat among workers there is stronger than in the West, where chat is mostly used by younger generations. If left unattended, these preferences may segregate parts of the network and prevent groups from working seamlessly together.

The **design of the process of collaboration** (for instance, through self-contained stages where roles and time commitments are clear), and the actual platform (for example, through gamification and generally engaging user experience) **mitigate some of these challenges**. Let's explore how that can happen.

Basics first: a lot of collective-cognition work, especially regarding formation of new ideas, shares common principles. Those tenets have little to do with platforms like Facebook, and much to do with the practice of innovation in the last decade, where the value of combination of ideas has been exposed in all its power. That's where anyone interested in harnessing collective intelligence for generating new solutions should start.

Quite a few methods for innovation have been tried over the years. Design thinking has been particularly influential in creating environments where people come together through a structured (yet not overbearing) process, to meld minds. Many contemporary collaboration platforms can work both on the identification of relevant ideas at scale (as explored in the previous chapter) as well as providing the virtual space for the combination to happen. Yet current approaches may just scratch the surface of what is possible.

Before we dig deeper into interaction platforms, remember the distance we have traveled, and the **trajectory we are on**.

Before email, it was virtually impossible for groups to build on each other's thoughts unless they were in the same place, or at least on a telephone call. Email was the first tool to enable an asynchronous buildup of ideas. Despite its limitations (the dreaded *reply all*) it is still the workhorse of remote, asynchronous collaboration. Email's limitations also explain why we need better practices and tools. And email has evolved in response, with the ability to show conversations in threads, the auto-classification of email types, and Gmail's automatic text suggestions. Despite its limitations (think of the dreaded "reply all") it is still the workhorse of remote, asynchronous collaboration. An estimated 230 billion emails are exchanged every day[186]. And email is still evolving, through bolt-on tools such as Superhuman and the weaving of email with workflow such as Sedna.

These are all examples of machines stepping in to simplify the way people work via email.

Group collaboration, especially the hyperscale part of it, has also evolved through social media. Think of the changes introduced in recent years. Segmented communities from Facebook to Pinterest, from Reddit to LinkedIn help members find more relevant information from others with similar traits and interests. Enhanced user interfaces enable seamless sharing of content. The limitations of social media in enabling collective intelligence interaction are also more visible, particularly the infamous *echo chambers* that, while profitable for an advertising sale, have distorted public discourse.

Enterprise social media is comparatively new. Yet it is already a mature concept at the core of enterprise information technology strategies. Most innovative companies use social media tools very heavily, reducing their reliance on email. Slack's IPO was one of the largest in 2019 In that same year Microsoft's CEO heavily emphasized its Teams product. Major software vendors like Salesforce and SAP have their own social media platforms and strategy. A dozen or so tools have sprung up over the last years to organize communities outside of enterprises. This is a big space to watch.

Less well known but equally valuable are the inner workings of big communities like Wikipedia or Linux, where people use purpose-built yet simple tools and rules to build on each other's work. The cultural aspect of these environments is as important as their technology platforms.

An additional major trend is the **democratization of information technology** that will boost the power of individual nodes. Just like text editors and later blogs multiplied the creative power of millions of people, and spreadsheets made calculations more frictionless, the increasingly intuitive technology-application user interfaces will enable millions to become software developers and data scientists. That means that millions of network nodes will start processing information (turning data sets into insights) and creating flows of knowledge that scale. The upshot could be an increased automation of workflows enabling people to focus on problem solving. But it could also be an explosion of the number of nodes able to make inferences based on the world around them, as more people assemble sophisticated prediction machines, forecasting everything from biodiversity loss to weather-related disruption to supply chains.

These environments allow individuals to *sense* what the rest of the participants sense, *remember* what they have experienced, and open conversations that help *create* and *decide*. And the resulting *learnings* made are somewhat crystallized in the network, both in the form of documents, and through the establishment of new connections between people. In a way, that's similar to the way neural cells learn to *fire together* and *wire together* when exposed to new conditions, when able to provide an adequate response, "fire together", and subsequently "wire together". This is happening in organizational networks, right now. Virtualization of ed interactions adds a big advantage: the digital breadcrumbs that those activities leave behind are a rich environment for analysis, and – not just for network analysis (as described at page 74 in "Module 1: the "who". Illuminate the neural network"). They are also rich in opportunities for by AI technologies like Microsoft Viva Topics (formerly known as "project Cortex") to make sense of the network of people and the content they share.

Apart from the technology and methods themselves, collaboration platforms require active **management of the norms and rules** that enable nodes to work together. As we saw in module 1, intentional diversity of nodes is a great asset when seeking novel solutions, but it also makes interplay between nodes more complex and potentially less efficient. Successful online communities, and generally movements of all types, have shown the importance of the establishment of common ground to ensure the right level of motivation (e.g., belonging to a group of like-minded people) and the right process for collaboration (common processes understood by all). Some research has proposed that a few common aspects are essential to the effective collaboration of large groups[187]: trust, transparency, listening, fairness, honesty, and forgiveness. Their absence generates friction that saps the network of its momentum. Finally, the ethical aspect is increasingly important to collaboration. The power of AI can be used to amplify the darkest human tendencies, and creates new, disturbing possibilities[188].

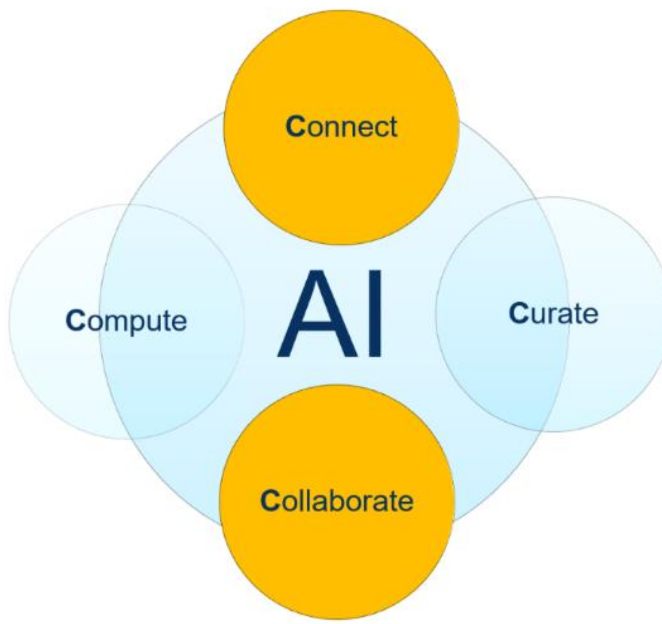

Let's recap **what AI brings** to collaboration at hyperscale in the diagram opposite. Intelligent machines help connect the right nodes in the network, make them discoverable to each other and able to interact in the format that's most appropriate for the need at hand – in both synchronous and asynchronous communications. They will become increasingly effective at doing so, as fields like natural language processing (NLP) evolve quickly. This is just the beginning, and we will explore more technology-enabled collaboration opportunities in the next section.

Much to use already, but also much to discover: this space is in constant evolution, so be sure that your teams check out the latest in each of the areas that we will discussed.

But first, let's illustrate the value of this module with real-life case studies.

Examples - Module 4

For years, new businesses and organizational structures have been sparked by technology-enabled collaboration platforms, and software companies have been ready to embrace the related opportunities. From Microsoft Teams and Skype for Business to Slack and Zoom and even Facebook's Workplace, collaborative functionality, made more intelligent by AI, has become an increasingly significant part of their development efforts. These followed the trend of building increasingly mature application interfaces (APIs) and development environments, enabling decentralized power users to build interoperable software code and containerized microservices that can be recombined like Lego blocks, hence enabling novel applications. New crowdsourcing businesses, from software community TopCoder to innovation platform Innocentive have shown the value of specialized workflows and in orchestrating collaboration. In a different example, American retailer now Lowes uses peer-to-peer countrywide cashier chat to drive collaboration, where frontline employees share best practices and early alerts as like in the case of the early withdrawal of a pet treat that was later the object of a lawsuit.

At another end of the spectrum, some examples will clarify how collaboration platforms can also work in more informal, less tightly managed groups. The BP oil spill on the US Gulf Coast in 2010 triggered a networked intelligence response, where groups of citizens armed with simple tools (such as aerial images captured by balloons to capture aerial imagery, combined with other information sources) initiated a process of monitoring and decision support processes that enhanced public response to that crisis. The resulting platform, PublicLab.org, has been adapted and used elsewhere in the world since. In 2016, Mexico City "open- sourced" the input to their city's constitution[189], where contributions from hundreds of thousands of people were translated into legal language by a central committee. Madrid is experimenting along the same lines[190] as is Reykjavik with its Better Reykjavik program. And in Africa, Wefarm.co helps millions of small-scale farmers to connect with each other – both through web and SMS - and exchange solutions to their questions.

In yet another context, collaboration can be created to convey emotions into art to foster coping. One of the best examples is the design of the New York City's 9/11 Memorial, intended to memorialize the trauma the city suffered in 2001. The initial proposals from architects and experts around the world had been deemed unable to reflect the true feelings and hence not fit for purpose. The final incarnation of the Memorial is heavily influenced by the voices of communities impacted, in a stunning work that sourced contributions such as soundscapes from ordinary people.

Another example of collaboration platform to deliver a vision for the future, centered around art as a catalyst, is the yearly Burning Man festival[191], organized by Black Rock City Limited Liability Company which uses similar models for other activities. The festival today congregates 70,000 people and has been running since 1986. People build and then completely dismantle a temporary city in the desert, including its infrastructure, security, food and shelter – respecting a strict "leave no trace" credo inspired by sustainable-development ideals. It is a great demonstration of how agilely coordinated bottoms-up activities can deliver complex systems at scale. Burning Man's experience is portable for instance to urban planners, refugee camps, and generally city planning.

In the space of political and social action, and in an effort to generate both new solutions and alignment, the UN, in collaboration with the Brookings Institute and the Rockefeller Foundation created the "17 Rooms" method. Its intent is to find solutions for each of the 17 sustainable development goals (SDGs). Each room, initially physical and increasingly virtual, is a breakout space staffed by a cross-section of skills capabilities, and authorities[192].

In a completely different realm, bitcoin miners de facto use blockchain protocols as the enabler of a massive collaboration environment, where decentralized computed can be "trusted" and its delivery made certain. New blockchain architectures will take this concept to the next level, enabling other use cases and hence permitting the use of decentral capacity to "reliably deliver runtime code". Some examples from decentralized finance (DeFi) can be found in stablecoins, exchanges, credit, derivatives, insurance and asset management (more in the detailed technical note section).

In the space of financial services, Numer.ai is a Hedge fund that harnesses the models of thousands of data scientists and serves up obfuscated data (homomorphic encryption) to over 7,000 data scientists. They give feedback to AI models to predict the stock market and get paid in digital currencies when the trades are successful.

And the even newer field of decentralized science (DeSci) provides some additional examples in the form of purpose-built decentralized autonomous organizations (DAOs) intended to, among others, transparently allocate funding decisions for specific research – like in the case of Ants Review, VitaDAO, or Molecule.

Decentralized Autonomous Organizations (DAO) based on the blockchain protocol are possibly one of the most exciting prospects for a future where technology will enable "business-as-a-software" architectures. A decentralized autonomous organization (DAO) is "a group organized around a mission that coordinates through a shared set of rules enforced on a blockchain. A decentralized autonomous organization (DAO) is a group organized around a mission that coordinates through a shared set of rules enforced on a blockchain.

One of the main benefits of a DAO is that they are more transparent than traditional companies since all actions and funding in the DAO are viewable by anyone. This significantly reduces the risk of corruption and censorship. Publicly traded companies must provide independently audited financial statements, but shareholders only get to see the financial health of the organization at a snapshot in time. Since a DAO's balance sheet exists on a public blockchain, it is completely transparent at all times, down to every single transaction."[193]

A few archetypes of DAO exist today[194]: Investment Syndicates (an evolution on the traditional hedge fund or venture fund model, these DAOs pool resources of members to invest in projects the community believes in); Protocol DAOs (decentralized projects issue governance tokens, enabling communities to weigh in on key issues and make their preferences felt through voting.); Service DAOs (like a new age LLC, these DAOs serve as the focal point for the development of new crypto projects in a decentralized manner); social DAOs (whether it's providing access to digital spaces or physical events, membership based DAOs are a new way to meet and collaborate with internet strangers.); NFT Curators (Collectors band together to pool both their capital and their collections, and with the resulting liquidity, they can support artists in new ways).

More examples are in the "How to Do It" section of the Guidebook, as well as the database on www.supermind.design/database. Let's explore one in more detail.

In-depth case study module 4 – upskilling at scale

In the case of the professional services firm engaged in up-skilling and re-skilling its large workforce at scale, the role of collaboration platforms was paramount.

To start with, the design, development and build phases of the new learning platform were executed by a completely globalized team. Most team members worked from separate offices or from home. External contractors flexibly came and went. Being effective in these conditions hinged on several technology platforms described in the next technical notes section, which facilitated the use of common ways of working that drew mostly from Agile and Design Thinking methods. For instance, all synchronous parts of *sprints* (short project bursts) and *scrums* (specific forms of team meetings) took place through Zoom videoconferencing, with widespread use of virtual whiteboards and Microsoft Teams.

The *run* part of the effort needed advanced uses of different technologies to enable communities of learning. For instance, meetings led by subject matter experts use the same videoconferencing technology with a strong emphasis on real-time chat discussions, virtual room breakouts, and instant polls that give SMEs to rich input from the learners. The SMEs themselves leverage community sites built on Microsoft SharePoint and Teams.

A marketplace of learning-by-doing opportunities was set up. It allows learners to consolidate memorization of the learning curriculum by executing small projects in mostly virtual settings. They leverage Microsoft Office 365 including Teams, as well as the core functionality of a Learning Experience Platform (Edcast). The latter also facilitates collaboration by enabling fluid feedback from the learners on specific parts of the curriculum.

Feedback loops are extremely important. They provide daily direction for the activities of both the central learning group and the communities. SMEs get continuous input on their choices and sequencing of materials from listening to the “buzz” in each learning channel. At the same time, data about learning patterns generated by users – both by compiling the feedback received on the e-learning platform as well as mining the digital traces left by learning journeys – is fed back to central teams, SMEs, learners, and their managers. That’s a departure from conventional instructional design which relies on infrequent oversight by learning experts largely disconnected from the organization-specific flow of learning.

The learning supermind system doesn't fully self-regulate. The overall flow of interaction requires tight orchestration. All input and issues surfaced from multiple touchpoints (sensors such as troubleshooting tickets, content ratings, and other feedback) are consolidated by a shared service group which makes insights accessible to the appropriate parts of the collective intelligence. This ensures that feedback is used quickly to fix mistakes, add functionality, or simply communicate more clearly and uniformly with the learners.

Module 4 – How to do it

*[Reminder: the specific examples documented in the technical note section may be used to compile a list of "portable ideas" that enable more creativity during the design process (see the "*The process of designing an augmented network intelligence*" chapter)]*

The technical notes for this module cover the following elements that can inform specific work streams:

1. Have we attempted to reduce the routine, non-value-added repetitive tasks that burden the members of the network, by **automating part of the flow of work**? (Section here)
2. Are existing "crowds" being harnessed through **crowdsourcing platforms**? Similarly, are **freelancer managed-markets** being utilized appropriately? (Section here)
3. Are the teams using the right **idea management platforms** and services? (Section here)
4. If useful, is **citizen science being** used? (Section here)
5. Are the right **unified communications** technologies (particularly important for synchronous collaboration) established and adopted? Are the appropriate practices understood by participants? (Section here)
6. Is **enterprise social media** or other community technology used as needed? (Section here)
7. Are the relevant **techniques, and methods for text and language-based collaboration**, used for the networks? (Section here)
8. If **debate** and dialectic are needed, are specialized **technologies** used to facilitate the flow of the dialogue? (Section here)
9. If **predictions** are the focus, are the right methods and tools used for it? (Section here)
10. Are the appropriate **collaboration** environments available for the specific needs of **data** and **information technology**? (Section here)
11. Similarly, and where needed, are the right tools used for the needs of **engineering**, or **life science**? If some of the work of the supermind is **scientific research**, are the right tools used for it? (Section here)
12. Is there a way to use **blockchain** protocols, potentially even in the form of Decentralized Autonomous Organizations (DAO)? (Section here)

13. Are we **democratizing** the development of code, and generation of data-driven insight, empowering many people throughout our organization to contribute more to the network? (Section here)
14. Are the right **learning platforms** used to facilitate collaboration at scale around learning materials that help the network evolve? (Section here)
15. Should we consider gamification, augmented reality (**AR**) or virtual reality (**VR**) for some of the collaboration? (Section here)
16. More generally, is the right **organizational design** in place to make the above possible? (Section here)
17. Are we **modularizing** the knowledge units, so they can be modified and recombined in a disciplined, scalable way? (Section here)
18. If resources permit, is the use of **advanced natural language processing** piloted for any of the above? (Section here)
19. Is cybersecurity considered in the design and use of the collaboration platform? (Section here)
20. Apart from the tools, are the **right management practices** being used to facilitate the work of nodes that connect with each other in a many-to-many, two-way conversation? Are teams encouraging the right flow of work, cadence, feedback loops? Is feedback sought and given extensively, and in a civil manner? Are iterations scrupulously performed? Are network nodes practicing mindfulness? (Section here)
21. Specifically, for the management of the platform, are our teams, at different levels, using **data-driven insight** related to the functioning of collaboration platforms, to inform activities and allocate resources? (Section here)

Tools and platforms

This section dives deeper into the methods and tools available today to those who want to harness the combinatorial power of a supermind.

Most of the analysis will refer to at least partially remote (virtual) collaboration, though other methods like in-person hackathons are well understood and widely practiced.

While possibly counterintuitive, the first technology class that should be investigated is not a collective-intelligence one: **automation of routine flow of work**. The reason is that much cognitive load is generated by duplication and repetition of efforts, which in turn reduces the amount of time that people can dedicate to higher value add activities. Low-code (such as Appian, or Genpact Cora Sequence) and even no-code software (such as KnomAI[195]) are two examples. Additionally, software in the "collaborative workspace management" attempt to streamline somewhat repetitive and even project-based work that doesn't fit squarely in a traditional workflow – examples in that space include Wrike, Workfront, Asana, Hive, Unito, Taskade, and Basecamp, as well as other enterprise social media that we will discuss later.

The first phase of collective intelligence platforms was ignited by the emergence of **crowdsourcing** (see also the respective section in the "module 1" chapter). Captcha, the now-famous method to enable authentication for user sign in with the objective of limiting automated login and related misuse of internet assets, was initially conceived as a clever mechanism to use people to help machines tag visual content from millions of scanned books that computerized character recognition couldn't decipher. It is in a way the largest crowdsourcing effort in the world. Amazon used it for tagging millions of pictures with humans – which spawned Amazon Mechanical Turk, the grandfather of crowdsourcing platforms. Google for many years has used humans to label data that then is served up to artificial intelligence algorithms. Companies like Crowdflower or Pybossa have roughly followed the same direction. Crowdsourcing however depends on generally strict workflows, that lend themselves less well to exploration. In many respects, they mostly do a sense and remember, and some simple "decide" collective cognitive job.

Freelancing is not new, but hyperconnectivity and the ability to inventory and quality-control millions of people at minimal marginal cost has changed its dynamics radically, effectively surfacing the cognitive surplus of millions of people and reshaping many jobs. In many respects, Uber was one of the first to enable this for a specific occupation – driving. But the "uberization of work" has many facets. The identification and enablement of a supply network of people at different levels of ability to meet demand has Companies like TopCoder (software development), UpWork (crowdsourced workers for many different jobs), UpCounsel (legal support, now discontinued), TaskRabbit (basic help for the home), ThinkSprint (help to improve innovation sprints based on Agile methodologies). These platforms have in many respects created the current "gig economy" movement, and they increasingly show specialization of roles where very specific expertise is harnessed in a way that wasn't possible before – an idea first publicized by MIT's Malone and colleagues in 2011[196]. Kobo360, a sort of Uber for freight, has helped Nigerian truck drivers to improve a notoriously inefficient, opaque and prone-to-corruption market[197].

The idea of hyperspecialization is also at the core of service providers like GLG (Gerson Lehrman Group) and Alphasights, whose mission is to connect people who have unusually specialized insight with seekers of that insight. For instance, professional services and consulting firms, as well as financial services ones, utilize their services to connect with experts for a few hours, to ask pointed questions that can help validate or disprove their hypotheses. Another example is the learning content marketplace Teachable, where creators generate and sell to an audience of millions of learners.

But like in the case of the crowdsourcing technology tools a la Amazon Turk, the freelancing platforms leave little space for exploration, that is for the "create" part of the collective cognition.

A very different example comes from Squadbox, a proof of concept developed in 2018 at MIT's CSAIL, where small communities of users come together to defend a member who has been targeted by online harassment. The tool enables the group to support the person through an intelligent workflow.

Crowd based, and more generally gig workforce, is a relatively new phenomenon with significant social justice implications. Collective organization of labor workforces in that space is only starting to emerge. Some examples: Opolis.co is a non-union tool for solidarity and economic security; a Danish trade union employs solopreneurs, offers tax records and other back-office support, provides sick leave, in exchange for 8% of their revenue (Hk.dk Serviceureau).

Idea management platforms are built to streamline and optimize the otherwise messy flow of idea input and refinement. Spigit, Wazoku, Hunchbuzz, Ideascale, Brightidea, Ideadrop, Qmarkets, Mindpool, LEO ("LeadingEdgeOnly") are some of the better-known players[198]. Spigit for instance has been behind one of the most successful idea funnels known – AT&T's TIP (The

Innovation Program). They have slight differences, especially in terms of additional services that they offer, but all of them have some ability to reach out to potential participants, display the challenges for which solutions are sought, and shape the workflow of input from the participants. Some of these tools are specialized around a use case, like in the case of citizenlab.co, which focused on consultation, engagement and co-creation with citizens to invigorate participation in democratic and public action.

Some specialized communities have created their own workflows. MIT Futures CoLab, used for instance for the ClimateColab community, is an example – as well as MIT Solve, a competition seeking for community-based innovation enabling inclusion and shared prosperity. Another example is the US Defense Advanced Research Projects Agency (DARPA, also known as being at the inception of the original internet protocol) Polyplexus which combines open-innovation collaboration with Delphi-style expert-community predictions. Innocentive (one of the most established, and famously behind much of NASA's open innovation efforts) and OpenIdeo and even some consulting firms (like Deloitte with their "Pixel" UK efforts) and others (like General Electric with their Geniuslink) share some of the same tenets. In those cases, the technology platform (often outsourced to the technology providers listed earlier in this section) is of somewhat secondary value to the provision of a pre-existing community, and the management of the community that those entities do to achieve the best results (some of the idea management software vendors also double up as managed-solution providers, helping their clients with community management).

At the other end of the development spectrum, peer to peer communities like Wefarm.co help Africa's small farmers share information with each other. More than a million people use it in Kenya and Uganda, with over 40,000 questions answered every day. Using the technology platform is very simple and is also possible through SMS for those who don't have access to internet. Interestingly, natural language processing is utilized to facilitate the exchanges in three regional languages in addition to English.

Beyond idea management, **citizen science** has been used for years to get humans to "finish the job" that machines couldn't. Zooniverse.org is one the most accomplished of these environments. Scistarter.org added a list of COVID-19 specific challenges during the emergency. Among many others, its GalaxyZoo project (using volunteers to classify galaxies based on deep-space telescope imagery) has inspired the very birth of crowdsourcing. It now covers projects in fields as diverse as arts, biology, climate, history, language, literature, medicine, physics, social science and, indeed, space. In these environments, the collaboration is mostly structured, and text based. Forms follow workflows and allow for attachments and links to add flexibility. A slightly different organizational construct is used by Amnesty International when using its volunteers in Amnesty Decoders[199], for instance in Decode Darfur which asked participants to identify destroyed areas in the settlements of Darfur using satellite images.

In the environmental space, Restor is a science-based open data platform to support and connect the global restoration movement, accelerating the global restoration movement by connecting everyone, everywhere to local restoration. Restor connects people to scientific data, supply chains, funding, and each other to increase the impact, scale, and sustainability of restoration efforts. Participants can analyze restoration potential for any area anywhere, Monitor and manage all their restoration projects in one place, share projects to show progress and support the movement, and Connect with the movement (e.g., explore sites locally and globally, find potential partners doing similar work, search using a range of scientific and social filters, share contact information to connect directly with other users.)

Citizen science has long used some sort of gamification to enable people as well as motivate them. For instance, the classic online FoldIt game where players contribute to solving protein-folding challenges that predict protein structures - and in some scenarios, the respective reactions to potential drug treatments. Built During the COVID-19 challenge, it created specific challenges for the identification of coronavirus related solutions.

A certain amount of gamification was also applied by Folding@home, arguably one of the largest distributed computing infrastructures in the world. Aimed at helping scientists develop new therapeutics by simulating protein dynamics, it heavily relied on giving points to the 30,000 owners of individual computing powers.

In 2021, the protein folding game (no pun intended) was upped by DeepMind's AlphaFold 2 and RoseTTaFold, both open sourced and enabling scientists to leverage this powerful AI machinery. Or consider the QuantumMoves game from the University of Aarhus in Denmark, through which path-choices by game players create potential paths for laser beams moving sub-atomic particles in experimental environments. Another experiment called BlockByBlock[200] is being conducted by UN-Habitat, Microsoft's Mojang (maker of online game Minecraft), where thousands of people in thirty countries brainstorm on the layout of their neighborhoods. Along the same lines, Neureka is an app intended to engage laypeople in helping scientists better understand neurological conditions such as dementia by having them play a game. And finally, the Phylo DNA puzzle (from Canada' McGill and some Canadian governmental agencies), through which large numbers of people identify alignments in clusters of genetic data. The recent "Planetary Computer" initiative at Microsoft is part of the same trend, this time with Microsoft Azur's significant resources behind it.

Much citizen science focuses on somewhat technical challenges, not organizational ones – and clearly not freeform, unstructured ones. But what if the collaboration is not easily guided by a workflow or enabled by a simulation environment and its rules? That is the case in most work that requires an element of freeform improvisation. Whether asynchronous or synchronous, large groups of people need appropriate technology to get that work done at scale – and typically virtually.

Enabling communications is the next big building block. Connecting people, so that they're discoverable to each other and they can communicate - this is one of the simplest yet most important elements of platforms supporting collective intelligence. Language is the key vehicle of transmission of information. That's why the foundation of group collaboration is still language-based, and the respective technology focuses on it. And one of the key values of network-enabling applications is that they enable members to discover each other (LinkedIn became what it is by doing just that). As noted elsewhere, this is one of the single most important triggers for collective intelligence. And this factor prompts the expansion of the network, where satisfied members reach out to their own networks to make them use this environment. (The sales strategy of all enterprise social media companies uses this effect, by planting small groups of adopters and enabling them to become evangelists – to a point in which the network has become so powerful that the choice for the CIO is sealed).

Unified communications are an obvious collaboration environment, specifically catering for **synchronous** or near-real-time (e.g., chat) interactions. These interactions are important because many people tend to be more productive and creative when they can bounce off thoughts with each other synchronously (mirror neurons may have something to do with it). Additionally, live video feed from the other participants, especially when the face of the counterparts is big on the screen, enhances the level of engagement and tends to help creating an *esprit de corps* that would otherwise lack especially in virtual teams. In many respects and within reason, real time interactions are a form of incentives, as participants in the network derive energy from personal interaction, both with peers and with senior executives, celebrities, etc. (refer to Module 2, “Create and disseminate incentives” for more detail).

This technology space has seen significant advances in the last decade, starting with the disruption that Skype introduced into the telecom industry, but also with entrants competing successfully against the hegemony of Cisco Webex and to a lesser extent Adobe Connect. Citrix GoToMeeting, Google Hangouts, BlueJeans and possibly most notably Zoom have made personal videoconferencing a reality. The vitality of best-of-breed players also shows that this space’s technology continues to evolve both through software and hardware innovation– see the examples of Around, Fuze, Lifesize, Loom, Owl, Starleaf, TechSmith, Hopin.

The recent coronavirus emergency boosted the adoption of these tools, especially as the cloud-based version of them was already extensively used in organizations, enabling scalability while maintaining an appropriate level of security (in most cases, though the infamous “zoombombing” cases demonstrate that user behavior is often unpredictable.) There are even open-source, specialized communication platforms for gaming communities, such as Discord – increasingly able to spawn bots that moderate automatically some of the conversation.

Their audio and video have improved so much that traditional videoconferencing equipment has come under siege, and the network of people able to access these tools has increased exponentially. The lag between image frames and voice transmission between parties has become so low, and the picture definition so high even on large screens, that the participant’s brain doesn’t perceive it as “remote” and responds with neurological (and likely hormonal) signatures akin to in-person ones. That makes conversations more natural. Additional capabilities being deployed now are simultaneous captioning where the speaker’s language is transcribed in real time to facilitate the understanding for non-native speakers; and AI-based noise canceling (such as in Krisp.ai) that eliminates any frequency but the speaker’s voice, greatly improving the experience for those who work at close quarters with others.

One current and significant problem of video-based synchronous communications is the misalignment between gaze and camera, an effect called parallax that makes people's brain interpret the other person's behavior as somewhat less engaged compared to in-person meetings. The loss of personal engagement affects motivation and possibly reduces the effect of so-called "mirror neurons" that are thought to play a role in enhancing problem solving in groups. Solutions will be brought to market in the future, for example through Apple introducing eye correction software for FaceTime. For now, the best solution is to encourage speakers to look into the camera, or to have an external camera placed on the line of sight through for instance flexible arm on which the webcam can be mounted.

Since 2018, Microsoft has been actively promoting its real time translation of spoken word during Skype sessions, has added automated subtitling to PowerPoint presentations (Presentation Translator - which uses the language corpus from the presentation itself to enrich its vocabulary) and Google launched simultaneous earbuds intended to automatically translate language between two or more people in an in-person meeting – and while the current results are not effective enough for professional use, they will soon be close enough, especially for groups with mixed proficiency of language. Meanwhile, Zoom videoconferencing has enabled simultaneous translation manned by human translators (professionally certified contractors working for a third party translation company) to enable real time translation between critical languages such as English, Chinese, Russian, and German. In the zoom example, human translators are automatically connected to upcoming meetings in real time, as part of an additional service that Zoom subscribers can purchase.

At the end of 2020, the introduction of Clubhouse created a new hybrid: real-time audio-only conversations, in private invite-only rooms, where (typically) influencers host discussion in which participants can raise hand and intervene. This one-to-many model is de-facto a hybrid between social media (podcasting) and unified communications. While it is too early to tell if this model would be effective, in early 2021 Clubhouse was already valued as a Unicorn.

Other examples of confluence of methods are journalist and entrepreneur Azeem Azhar's Exponential View media source, which includes a podcast in collaboration with Harvard Business Review synergistic with a newsletter as well as a community of thought leaders run on Slack and Circle.io. The three elements reinforce each other.

Indeed, the increasingly lightweight data transmissions and the improvement of bandwidth access and quality have made these tools usable by many, in many conditions (including companies where the IT infrastructure is poor or complicated). There's little doubt that the introduction of 5G will bring out some additional ubiquity of use for workers on the go. With that, the network effect of many more people being able to connect is becoming very prominent. What's missing at this point is a few years of use for the less tech-savvy users so to build the habit.

An existing gap is that use of synchronous collaboration technologies requires a degree of formality, through for instance calendar meeting requests, that creates barriers to serendipitous encounters. Tools like Microsoft Teams have integrated "**presence**" indicators that show individual people's current status (e.g., busy in a meeting), so that others can reach out at the appropriate time. Other tools (Jamm, Embryo, Pragli, Remotion, Sococo, Sidekick, Tandem, Teemly, gather.town. workadventu.re among others) have even more directly focused on the design of that spontaneous encounter experience, trying to provide additional ways that people can "locate" each other in a digital/physical space to trigger interactions. Often integrated with tools such as Microsoft Office and Slack, the ultimate objective is to replicate a "virtual office", or as analysts Gartner calls them, "ambient virtual meetings".

While synchronous collaboration enabled by these platforms is a powerful engine for combination of ideas, synchronous work has two main drawbacks that warrant a cautious and efficient use of synchronous *time*.
The first problem is that large groups of people struggle with being able to communicate meaningfully at the same time. Groups above 7-10 participants exhibit a quick decrease of engagement, because some members get to rarely speak and their attention flags. And some participants tend to monopolize the use of time at the expense of often thoughtful but quiet others. While some of the new platforms have capabilities like breakout rooms (where smaller groups can congregate) and "turn taking" indicators (showing the participants respect of meeting etiquette), much still needs to be done to enable a smoother and more scalable flow of synchronous interaction.
The second is that due to calendaring conflicts (and the complexity of managing individual calendars) planning even for small scale synchronous meetings is hard and slows things down.

Companies like doodle.ai and X.ai have started attacking the problem to make the user experience a lot more intuitive (for instance using email's natural language). Other applications that integrate with Microsoft Outlook, and Microsoft themselves are actively working at the problem, to eliminate the labor-intensive work of finding common slots. Limitations due to information security restrictions will likely soon be resolved, and the only constraint will be user awareness. It is probable that one of the most useful applications of AI will then be the work of finding common slots that minimize the inconvenience for all participants.

A last set of considerations pertains to the business practices related to synchronous, remote collaboration. Given that most today's professionals have only had a comparatively short exposure to these tools (compared to, say, email), the proficiency in using these methods varies significantly. It is a good idea to attempt to reduce such variance by deliberately training users on a common set of basic standards, such as etiquette (background noise, avoidance of multi-tasking), use of equipment (microphones, cameras, lighting). While the tools themselves are increasingly able to provide some input to the user (for example in the case of poor audio quality), large networked teams general practices should be the object of an intentional effort. Failing that, it is likely that some groups will constantly undershoot their network potentials, as they may not get exposed to the best practices that others possess.

Community-building software, whether enterprise social media or outside of companies, has become very important. Apart from the often-cited Slack and Microsoft Teams, many other platforms exist: Mattermost, SAP Jam, Jive, Basecamp, Khoros ("community and social media management software that makes it easy for marketing & support teams to deliver the best customer experiences"[201]), Telligent, Vanilla Forums, LumApps, Unily, HigherLogic, Tribe, Circle.so and many others – each with their own slight specialization, depending on the type of community, for use within the organization, outside of it, or in a hybrid model. Interestingly, Facebook has also made some steps into space with their Workplace product. Many of these solutions started with asynchronous communications, but increasingly added synchronous capabilities such as video and voice.

These solutions often focus on facilitating collaboration *within* organizations and provide adequate cybersecurity to do so – though Microsoft Teams has made steps to make federation of organizations within the same environment, and Slack as of 2020 had been very vocal about its strategy to facilitate the emergence of cross-organization ecosystems. This said, as of today, connecting people *across* organizations leverages many other tools. To an extent, Reddit and LinkedIn Groups perform some of the functions of these tools: get groups of people "discoverable" to each other and ensure that they can post content visible to all, to enable frictionless debate. The upshot is that creating templates of some of these environments (Microsoft Teams for instance), when done through experience design practices, can help drive new behaviors throughout a company. Some experience-design firms, such as RightPoint (a Genpact company) and Accenture perform that kind of work.

The following types of technologies can be combined to cater for varied collaboration needs revolving around the enhancement of the use of **text and visual communications** to increase their impact.

- **Messaging and chat.** While these represent one of the older categories, there's a rich and vital supply of best-of-breed applications, with distinctive feature including security and ease of use, as well as the ability to share voice messages seamlessly. Discord, Dialpad, Quill, Spike, Threads, Yac are some examples.
- **Repositories** like Box and Google drive were the first tools to enable the sharing of documentation.
- **Blog** editors that allow basic commenting like WordPress
- **Professional-looking editing to enhance user-generated communications e.g., presentations, multimedia, and web design** from the likes of Canva or Squarespace. Include collaboration for co-creation.
- Co-editable **long form text** editors: from Google Docs to Microsoft Office 365
- Co-editable **spreadsheets** and slide editors, including those from Google and Microsoft, or Airtable (which for instance was used to catalog COVID-19 community based and crowdsourced solutions[202] and has also been used to create small open source reusable utilities, such as appointment setting for business-to-consumer services[203]).
- More **flexible editors** like Coda.io, that combine text with basic database, workflow and other types of editing to create flexible process guidance for teams working on the same sheet

- More specialized tools exist for **scientific long form writing**, namely papers where the citation burden needs to be managed. Overleaf, offering LaTex support, is one such tool. MIT PubPub is an attempt to create a text environment allowing for large scale annotation by a broader, typically scientific, community.
- **Virtual whiteboards** – for synchronous screen share, like in the case of the Zoom whiteboard, conveniently available on its IOS application, hence allowing the use of the intuitive Apple Pencil on an iPad. This is a marked improvement from previous options, such as the use of screen sharing of documents (like PowerPoint) while using optic pens such as Wacom's that while precise, have a steeper learning curve for the casual user
- **Virtual walls**, where virtual post-it notes (the workhorse of creativity sessions) can be generated, modified and moved by groups of roughly up to 50 people. Padlet.io, Mural.ly, Miro, Jamboard, Klaxoon and Stormboard are some options here. These cater for mostly text type entries but do allow for attachments and links, but they also accommodate other content. Dynamic, real-time input and visualization of text by groups of people is supported by tools like Mentimeter, Slido, or ThinkTank.net.
- **Mindmaps** help creating outlines of documents. Companies like iThoughts, Kumu.io and Mindmeister allow collaboration in this, largely text-based format. Also, these cater for mostly text type entries, but do allow for attachments and links. TheBrain.com offers a more expansive vision that leverages mindmap type relationships for more than just text – the result is possibly less intuitive than the typical mindmap software.
- **System dynamics maps.** While hyperscale collaboration may not be fully possible on these tools, they are quite useful to capture the interplay between components of a supermind, hence enabling a better understanding of the combined effects of different elements of a network. Mental Modeler[204] is one such tools.
- **Dictation.** It is likely that in the future we will rely less on keyboards, because of AI's significantly improved accuracy of voice-to-text, which in turn will dramatically increase the quantity of text (if not necessarily the quality) that humans can produce. Microsoft Office already allows voice dictation, for instance.
- **Design** co-editable environments like Figma.

Collective cognition and decision-making are being increasingly supported by specialized providers, like in the case of debate and prediction platforms.

Debate platforms, building and expanding on early attempts like MIT Deliberatorium and DebateGraph[205], Kialo[206], atstakegame.org (small-group turn-taking and debate), D-agree.com (Japan), and Changeaview.com. There, opposing point of views are discussed in civil and constructive manners to enable the formation of better, more cohesive positions. A different model is used by Pol.is (also used in Taiwan's extensive public consultation platform), where the points of view of the various constituencies in an organization are plotted against each other on a map, to show signs of convergence or significantly spread out opinions. A somewhat similar approach is adopted by Crowdscope.com. US government's Crowdlaw is an attempt to engage broader constituencies in the legislative process. Survey tool allourideas.org, mentioned in the previous chapter, is a variant of surveys that can be administered through wiki, with an in-built statistical analysis that lends itself to deciding between options during a debate. Decision making after debate is the focus of other applications, like Ethelo.com, democracyos.org, partiunion.org, bangthetable.com, decisiontree.io, choiclaweb, Airesis, Occupy.here and Loomio.org. which can be used for engaging stakeholders and teams; or liquidfeedback.org that enables feedback and voting on motions. Tapvote.org, Delib.net and Common Ground for Action[207] are another platform for public deliberation online. Horizonstate.com is a "secure ballot system" based on blockchain distributed-ledger technology. Decidim.org includes participator planning, budgeting, assemblies support as well as consultations of citizens and networked communications. Specifically conceived for use for local government or smaller countries, consensus.ai harnesses "verified signals gathered from citizens in tight feedback cycles to help governments interpret sentiment as it evolves".

Prediction platforms that combine the opinion of people have been around since the early 2000s, for instance in the form of "prediction markets"[208] where individuals can "bet" on specific outcomes (like the future trend in a market or the success of a new product) and the platform's algorithms combine and weigh the results. Companies like Eli Lilly, IBM and Google have all used them extensively in the past to obtain "signals" to improve their development of new solutions. IBM's alphaWorks used to track downloads of "alpha" (i.e. very early release) solution to predict future demand. Eli Lilly identified many MDs and PhDs who proved extremely far sighted in predicting future drug's commercial performance.

Predictions based on employee crowd judgment has been used extensively in companies like Ford, Tchibo, Best Buy, Siemens and others – on a range of topics from sales forecasts to project overrun early warning.

Interestingly, the use of prediction markets hasn't grown as one would have expected given those early successes. Prediction markets' challenge is not in the process or technology – it is in embedding it into the fabric of strategic decision-making so that it doesn't become anecdotal. In fact, it is hard to maintain an active group of contributors if the model is seen as a one-off, and the strategy, marketing or R&D departments don't use it consistently.

Publicly available providers that leverage that model now exist - see Predictit.org, MindPool and Metaculus for instance.

Betting applications have leveraged some of the prediction markets ideas, and some of them have been quite successful, like in the case of Betfair or Smarkets. Examples are Metaculus.com and Hypermind.com (open forecasts) or Tinycast[209] which can be deployed by enterprises and other organizations to research specific topics. Crowdsmart.io applies a combination of expert member community with AI to product startup funding success. In 2018 Augur.net, the first decentralized prediction market based on the Ethereum platform was launched, hinting at possible use of blockchain technology to allow for a broader range of queries, unlike the fixed ones that normal prediction websites like BetFair allow. Polymarket is another, even more recent Ethereum based player. While it is unclear how successful this direction has been to date, given the difficulty for a market to create enough liquidity to support itself, the idea of making predictions a market holds the potential of creating a paradigm shift, in that market structures can be disaggregated into smaller units and made much more scalable. In this regard, think for instance of the attempt by JPMorgan (ROAR data) to create the infrastructure a marketplace for data science that supports predictions at a granular level. While some of the ideas are similar to the ones used in contest platforms like Kaggle, the concept of utilizing data science for the creation of building blocks for a market is novel. Similar concepts are being deployed by Microprediction.org, where modularity of data science algorithms are envisioned to meet an increasing flow of data sets.

The concept of "superforecasters" has emerged from the crowdsourcing methods, indicating the ability to harness the superior prediction capabilities of specific individuals. GoodJudgment[210] initially driven by University of Pennsylvania to support IARPA (US intelligence community's equivalent to DARPA), including its GJOpen, as well as the UK's government Cosmic Bazaar focus on that.

Predictions models for healthcare are the focus of Evidencio, which enables users to get all relevant medical prediction models from a single online library.

On the incentive side of prediction markets, Gnosis.pm has built "conditional tokens" with richer semantic characteristics that allow more granular enablement of prediction market pricing. Another example is Unanimous.ai, that attempt to harness people's ability to predict discrete alternative outcomes (for instance, sports and political) by monitoring not just their opinion, but also their behavior as they do so. They do so by using among others AI to detect patterns in the movements of the cursor as people decide between alternatives.

Data science and **information technology** leverage virtual collaboration in possibly even more straightforward ways. After all, Linux showed to the world how a coding environment can be created with rules that enable both solid coding and participatory work – and even spawned an ecosystems of successful private enterprises like RedHat (acquired by IBM) that use and extend that code and add additional services on top of it (like maintenance.) GitHub, recently acquired by Microsoft, is a thriving example of collaboration collective code enhancement and reusability through a clever form of collective revision control based on its now famous forking and committing (indeed, Linus Torvald worked on the early code for that process). StackOverflow is an independent, open community for anyone that codes, with the joint objective to get answers to the toughest coding questions, share knowledge with coworkers in private (and find a job, which acts as a powerful motivator). Other software vendors have been collaboration communities for many years, such as SAP Community Network.

All these network environments have intentional and codified methods for making changes (e.g., "pulling") to modules of knowledge. Those processes are at the core of their scalability[211].

Collaborative software development technology has hence become a significant market segment, with examples such as GitLab, Repl.it, Glitch.com, Codestream, Clubhouse, Linear, and

JetBrains. Specifically, for data scientists, Jupyter, Dominodatalab, Algorithmia, CrewSpark and in a different way Kaggle (data science contests) or Numer.ai (crowdsourced hedge fund) all enable collaboration at scale.

Even **engineering**, **life sciences, and neuroscience** and more generally “hacker- and maker-spaces” (such as Fab Labs, where participants collectively hack practical problems with scrappy, limited resources, as part of the so-called “maker movement) have their participatory, collective intelligence networks. Arduino’s community (open-source electronic prototyping platform enabling users to create interactive electronic objects), Hackster.io (“a community dedicated to learning hardware, from beginner to pro, where participants can share their projects and learn from other developers”), Wikifactory and OpenBCI.com (dedicated to brain computer interfaces) and are three examples. The Open Source Pharma Foundation[212], in India, attempts to use open source methods to share data and processes to generate breakthroughs in the field of life sciences.

In **life sciences**, MIT-linked Community Bio is aiming at creating a global “movement” that can reimagine how science can be done. Starting from a core of synthetic biology, its modus operandi is “anti-disciplinary”, in that it not just brings together hundreds of people from very different disciplines but also encourages thinking of solutions that could redefine the future of those disciplines, on the premise that solutions will require a radical rethinking of the boundaries of those methods. It does so sometime in person, but more often virtually through a mix of synchronous and asynchronous tools. The community concept and some of the participants was leveraged during the Covid-19 pandemic to generate new ideas as part of a MIT CoLab in collaboration with healthcare equipment manufacturer MilliporeSigma.
The Openinsulin.org is “working to develop the first practical, small-scale, community-centered model for insulin production to make insulin accessible to all”.
Another attempt to give large, decentralized life sciences and possibly other scientific communities new methods for doing research, JOGL (Just One Giant Lab) has created a technology platform that enables sharing of resources and attempts to provide the foundational infrastructure needed for participatory and decentralized innovation.
One more example is Eyewire, a science project where volunteers help neuroscientists to map the neural connections of the brain, taken through an electronic microscope. As AI-enabled

computer vision and related machine learning will improve, the volunteers' input will likely be able to train the next generation of machines so that they can "finish the job", at scale, and ask human input only for exceptions.
And finally, collective responses to crises like Covid-19 have been harnessed by the "maker movement" community, such in the case of viralresponse.io where "a community of makers, designers, engineers, experts and creative problem solvers who innovate globally, and produce locally [to overcome] existing supply chains and intellectual property restrictions, to get the right products in the right hands, wherever they're needed."

A generalized **scientist research** platform is synapse.org, that helps researches organize their digital research assets, get credit for their research, and collaborate with others, through a platform that enables structured sharing.

Blockchain is a form of collective intelligence, typically able to ensure "trustless" verifiability of transactions. The future of blockchain might also bring significant change to collaboration – potentially even collaboration between machines, through decentralized autonomous organizations (DAOs). Several examples of blockchain-based collective intelligences are emerging, specifically helping various parts of the financial services system to collaborate without the need for intermediary by leveraging so-called "smart contracts". Take the following as examples for financial services[213]

1. ***Stablecoins** seek to maintain a constant value of a token relative to some asset, most commonly the U.S. dollar or other major fiat currency. Non-custodial stablecoins function as DeFi services themselves. Custodial stablecoins are centralized but may be incorporated into DeFi services.* An example is DAI (governed by MakerDAO) based on Ethereum.
2. ***Exchanges** allow users to trade one digital asset for another. DeFi exchanges avoid taking custody of user assets, either through a decentralized order book or by matching orders and setting prices algorithmically.* Uniswap and Sushiswap are examples.
3. ***Credit** involves the creation of time-limited interest-bearing instruments, which must be repaid at maturity, and the matching of lenders and borrowers to issue those instruments.* Flash loans are an example, enabling borrowing from decentralized protocols such as Aave and dYdX. Another example is Compound *"a money market protocol that lets users instantly lend to or borrow from a pool of assets in a smart contract"*
4. ***Derivatives** are synthetic financial instruments whose value is based on a function of an underlying asset or group of assets. Common examples are futures and options, which*

reference the value of an asset at some time in the future. An example is Synthetix, a synthetic asset issuance protocol built on Ethereum.

5. ***Insurance** provides protection against risks by trading the payment of a guaranteed small premium for the possibility of collecting a large payout in the event of a covered scenario. Asset management seeks to maximize the value of an asset portfolio based on risk preferences, time horizons, diversification, or other conditions.* An example *is* NexusMutual which offers smart contract insurance cover, for protection against a smart contract bug, in the UK.
6. ***Asset management** seeks to maximize the value of an asset portfolio based on risk preferences, time horizons, diversification, or other conditions.* An example is Set Protocol, *a decentralized portfolio management protocol enabling users to create a token that represents a fully collateralized portfolio of other digital assets, including Bitcoin, Ethereum, and stablecoins.*

As mentioned in Module 2's case studies section, DAOs are mushrooming, and their granularity enables the composition of aggregates which significant scope. A DAO's structure typically includes some of the following functions, with specialized players that support them: Formation; Communication; Community; Governance; Compensation; Treasury. Some examples here: POAP uses NFTs to "bookmark life event"; Snapshot helps coordinating resources through distributed voting, to help create proposals, support varied voting strategies; Gitcoin supports the "building (and funding) or Open Web" with a focus on public goods, through a token (GTC) that helps making decisions on allocation of resources; Boardroom helps governance e.g., voting; withTally is a voting dashboard to aggregate information from disparate Defi governance structures; llama helps manage treasury, as does Parcel which also supports payouts; DAOHaus is a no-code platform for creation and orchestration of DAOs; CollabLand is a community management system; PleasrDAO organized people who outbid others in buying Edward Snowden's NFT; PartyDAO organized people to build PartyBid which simplifies the setup of DAOs to bid on auctions; Krause House is a DAO aimed at buying an NBA team; SyndicateDAO wants to challenge the venture capital business model; Seedclub.xyz is a DAO focused on tokenized communities; and others are springing up in multiple areas, such as the media industry. In mid-2021, real-world biopharma research and IP was transferred and funded on-chain as an NFT, and now owned by a DAO (Vitadao.com), with the first objective of funding a research project at the University of Copenhagen.

As of 2021, DAO's user experience remained clunky, hence limiting the actual democratization of its use, but developments are underway, and libraries will emerge so that code objects could eventually be reutilized. Similarly, DAO-to-DAO (D2D, currently funded by PrimeDAO) collaboration and interoperability mechanisms will emerge, enabling the formation of larger DAO ecosystems.

Democratization of access to means of **coding** will likely create an explosion of the number of people able to develop applications, and as a result, the speed of information through their respective network. Platform-as-a-service tools, including AI modules such as computer vision and natural language processing; low- or no-code process development will put workflow design and execution into the hands of business analysts; and automated process discovery (Celonis or TimelinePI for instance) will enable many people to better understand the actual flows in their formal organizational processes, so that they can focus on the suboptimal part of it, and they can even simulate future conditions. If GPT-3 lives up to the hype it has generated, it could also radically alter how people write code, by enabling millions to just provide plain-English instruction and let the neural network do the rest. The upshot is that networks will be able to connect with applications, as opposed to with people, hence greatly increasing efficiency and throughput, and freeing humans to spend time on problem solving, as opposed to transacting.

The same democratization process will have a major impact the fields of data and analytics. There, citizen data scientists will use tools to build AI-powered models and other solutions that so far were the precinct of highly trained, scarce, and costly data scientists. These advances will also unshackle the elite data scientists from the need to painstakingly prepare data and basic models, so that they can focus on more complex tasks. And finally, the reduction of the friction in processing data will enable both business and data scientists to spend more time having business discussions, informing hypothesis and analyzing insights, as opposed to spending most of their time trying to communicate with each other. The result of these new conditions will be an increased breadth and depth of data processing, with much more – and fresher – insights being injected into the networks. Google AutoML, which helps developers to train sophisticated models and can even provide them with (often crowdsourced) annotation of data, is an example of things to come.
And again, GPT-3 has started demonstrating promise in no-code, plain-English writing of SQL

or Python code OpenAI launched GPT-3-based Codex in mid-2021. The accuracy of AI's code, and certainly its creativity, are still well below the standard required for mission-critical work, and the actual writing of code is only a part of the developer's work – so we shouldn't expect AI to take over developers' jobs anytime soon. However, another option is possible, as exemplified by Github's Copilot, which uses Codex to support developers through machine-generated suggestions for lines or entire functions.

All these platforms have some common tenets, and in particular their ability to leverage a **common technical environment** for building the solutions. Well beyond standard APIs and a common technical language, many platforms provide extensive software development kits (SDKs) that enable programmers to leverage preexisting code, as well as put into production theirs with comparatively minimal efforts. Most large software companies, from pioneering Salesforce to SAP, have created these environments. More recently, Android and iOS enabled millions of people to create mobile technology code. Gaming communities are among the most interesting new developments in this space. Importantly, these environments exercise some element of quality control – most famously through the comparatively slow process that Apple follows before allowing apps on its AppStore.

Learning platforms are, as noted in the "reskilling at scale" case study, a form of collaboration platform, although they're less two-ways than typical collaboration environments. There, feedback from the users enables curators of learning to fine-tune their choice of materials, and even learn from patterns due to user needs or the change external conditions. Just like in a retail store, the learning platform acts as a "knowledge superstore" that shows the type of demand for knowledge at any point in time, reflecting what is needed and enabling a better, data driven decision about what knowledge investments are needed.
A slightly different twist on network-based learning are networks of mentors, for instance those focused on startups or underserved segments of the working population – for example the mentorship program of The Five Network which "connects talented 9th and 10th graders to a team of mentors and motivational learning content via a mobile app for 5 years".

While these are early days, using **videogames, virtual and augmented reality** for collaboration between people is worth mentioning. Virtual spaces like Second Life have for many years hinted at the possibilities. Nintendo's Animal Crossing, a life-simulation game with its own currency

and central bank, is quite active also owing to the COVID-19 lockdown. The Wild[214] is possibly the first current example of fully collaborative 3D models where people can work together in a virtual space to gain insight on anything from architecture to product design. Real time, immersive collaboration holds the promise of providing a three-dimensional canvas, and an emotional immediacy that traditional virtual spaces don't have. No wonder that engineering and other design spaces are watching this space, in the attempt to build immersive "digital twins" that combine the ability of users to experience virtual spaces without the limitations of physical location, together with the possibility of data-driven simulations. Nvidia, one of the leading manufacturers of high-power processing units (and the inventor of the GPU), has recently launched its Omniverse environment to cater to these possibilities.
Another example is the (re)use of Pokemon Go to help people socialize in partly physical, partly digital spaces, for instance in the case of US military troops stationed in Afghanistan's Bagram base[215].

In mid-2021 Facebook launched Horizon Workrooms, virtual reality environments that help people generate their own 3D avatars and meet in three-dimensional spaces.

Collaboration platforms are the foundation of new, collective-intelligence based **organizational design**. At the time of writing, there are some experimental examples of platforms intended to create radically new organizational structures, whereby people work on the tasks that they like and they're competent for, and they get fairly rewarded for it. They borrow some of the concepts of holocracy (used for instance at retailer Zappos, software firm Satalia, or to an extent videogames producer Valve). Stanford University's "Flash Teams" management software is an attempt to provide a software that enables management at run time for such groups[216]. Colony.io attempts to create a protocol for interoperability between an emerging set of technologies based on Ethereum, for instance, and is focused on the management of the human resources – in contrast to Augur (predictions), golem.network (distributed computing power) and 0x.org (exchange of tokens). McKinsey's OrgLab enables many people to collaboratively design the organizational structure during a redesign effort. The ideas expressed by Aaron Dignan[217] in his OS Canvas also focus on the collective and its functioning, and pivot around areas consistent with the four pillars of collective intelligence: purpose, authority, structure, strategy, resources, innovation, workflow, meetings, information, membership, mastery, and compensation.

All these tools enable the recombination of ideas. However, proper exponential scalability for that combinatory act depends on how **modular** the knowledge is. Search, rework and eventually potentially generative design rely on the ability of people and machines to find the right nuggets, compare them, and build on them. Content tagging in a database, annotation on free text content, and generally metadata associated with content elements are key architectural enablers. With the advent of advanced knowledge graphs aided by AI, such as in the case of GPT-3 or Diffbot, we may be able to obtain properly structured natural content, that humans (or machines, in the case of generative design) could build on. While much of this is out of reach for most companies, working more diligently on natural language is not – hence the importance of the next few paragraphs.

The tools reviewed so far are mostly enablers of hyperconnectivity, but they hardly use artificial intelligence on **natural language** for anything more than basic search, grammar and spellcheck. This is a gap, given the important role that language plays in enabling large groups of people to innovate and learn together (more on this in a subsequent chapter). Natural language processing (NLP) however has improved the ability for people to exchange large knowledge sets, by enabling for instance better search functions. Tools like EyeOnText simplify data science on text analytics. Text "x-ray" like what's done by Amazon Kindle allow the automatic creation of themes that humans can peruse and recombine.

The promise of a semantic web, where knowledge is assembled and mapped, tagged or otherwise structured enough to search for content in a more effective way than using traditional search engines, is still somewhat unfulfilled, but there are interesting developments. Semantic MediaWiki[218] "is a free, open-source extension to MediaWiki – the wiki software that powers Wikipedia – that lets you store and query data within the wiki's pages". MIT-backed Underlay[219] is another attempt to move in that direction. In their words "The Underlay is a global, distributed graph of public knowledge. […] This is an attempt to replicate the richness of private knowledge graphs in a public, decentralized manner.

It is likely that collective intelligence platforms will in the future use artificial intelligence to sift through language-based content much more effectively, for instance enabling the clustering of similar ideas at scale, hence allowing people with similar ideas to connect and cross pollinate much more fluidly with each other. Driven by easy-to-monetize market dynamics, some of these technologies are now being used for social media to push content to homogenous groups and hence maximize advertising revenue. But there's much more that can be done. Just think of the power of an NLP engine that helps cluster sticky notes in a hyperscale design-thinking session where thousands of participants surface their ideas with short snippets of text. Or think of the ability to follow the propagation and evolution of idea strands over time, showing how they organically mix their respective semantic fields with each other and generate new, cohesive ontologies.

Another quick look at what could go wrong (and as repeatedly said, a lot could). **Cybersecurity** considerations intervene for some organizations. For instance, corporate environments that deal with sensitive information tend to adopt Microsoft technologies that promise a stronger control. Some of the tools discussed so far allow a single-sign-on authentication that permits access to some of these environments (including idea management platforms) only to authorized users. The reality is that as of today some of these collaboration environments lack multiple-factor security that more traditional tools guarantee. Some of that security will become available at a cost – namely, at the cost of privacy.

Ensuring appropriate data management is a key problem in collective intelligence. Data integrity and anonymization are significant challenges for IT teams involved. At this point, the trust that individuals and companies put into data aggregation services is generally low, or even very low as in the case of sensitive data like healthcare. Some solutions may be able to scale thanks to blockchain based systems, such as the Ocean Protocol that enables those transaction to happen closer to sensitive data, behind protection wall, or Singapore's announced APEX[220] common application programming interface (API) for allowing public agencies to shared data with public and private organizations.

Network engagement and management - turning ecosystems into deliberate superminds

A network isn't a collection of entities. It needs to connect and wire, which requires a certain amount of energy at first. For instance, any community, from Linux to Wikipedia to creative design e-commerce marketplace Etsy, needed an initial amount of energy through founders, founding and in general labor, to get to a critical mass that makes it self-sustaining. That first phase of engagement and management of the network takes deliberate effort around several activities. Some of them are listed below.

As mentioned, much collective-cognition work, especially regarding formation of new ideas, shares a few **common principles** – best embodied in today's design thinking practices. Lots of good resources around these topics are already widely available. For instance, Stanford design school's "bootleg[221]", and Google Venture's Sprint[222]. For this reason, we won't look at these methods in much detail here, but it is worth going over some of the key principles – they will inform our use of new techniques and tools, especially the ones described in the previous section of technical notes. They are elusively intuitive, but beware: their implementation at scale is just the opposite of trivial – especially with large groups of people who aren't used to embracing them:

> **Think about "why", "what", and especially "who", first**. Focus on the journey of the "human in the mix" first. Use process flows if people and their feelings aren't as important as the behavior of other resources, for instance in a conveyor-belt type operation (but think about that twice, as it often it isn't). Never, or at least very, very rarely, think about the "how", the tools and technologies needed, first. As we saw in the chapter around incentives, reducing the cognitive load while maximizing the incentives, the value perceived by the human *users* is crucial. In the absence of those conditions, an interaction platform fails to achieve critical mass and potentially fail – especially if a large network is needed for that job.
>
> **Fall in love with the problem, not with the solution**. Many people fail to spend time in dissecting the problem and empathizing with those who struggle with it. Many professionals are wired to do quick pattern recognition and move fast to conclusions. When problems are really hard, that shortcut is a surefire way to end up with boring, trite solutions. Kahneman's talked about "thinking fast, thinking slow". There is a collective-intelligence version of that. Fast group-thinking is dominated by a few people who "want to move on", but especially in the initial phases of innovation processes, the opposite is needed: the group must slow down, and painfully get to a real deep understanding of the problem at hand. This sounds easier in theory than it is in practice. Most innovation workshops are too short, and most participants too impatient, to get to that deeper level of understanding - unless a group of people tightly orchestrates the flow and injects appropriate incentives to keep the group humming despite some may feel that they're meandering.
>
> **Be strict with the process but be flexible with everything else.** Large groups tend to struggle with keeping pace and staying engaged with poorly designed collaboration processes, or processes that aren't fit-for-purpose once the conditions have changed. While collective intelligence requires that

nobody is strong enough to unilaterally impose solutions, there is a need for someone (or a group) to define the process and adapt it based on conditions' changes, while avoiding knee-jerk reactions.

Ensure that "safe space" exists. People often take their own inhibitions, cultural reticence and emotional bias into collaboration processes. That's poison to a creation process. Leader must enable psychological safety and reinforce the need for improvisation, going out on a limb, polarized thinking, and crazy ideas. Those are often where the germs for real breakthroughs are, either directly or through recombination and evolution.

Make sure that reliability is monitored. Google's project Aristotle[223] demonstrated that highly effective teams were composed of people who could be trusted to deliver the right quality of output at the right time. At the most fundamental level reliability of team interaction is simply a pro social behavior that encourages individuals to respond to their counterparts in a timely and satisfactory way - network analysis shows not quite often some team members do not treat everyone equally when it comes to responding quickly get the messages they receive. Beyond this daily tactical behavior, reliability tends to be part of the culture, a norm that reinforces itself the more the network uses it. the considerations that we made in the "module 1" section when we talked about network characteristics are helpful here, because the creation and propagation of culture is a crucial driver of reliability.

Pay attention to interaction dynamics. Smart groups are often diverse (see the discussion in the "module 1" section), and certain types of people will naturally take the lead or monopolize the conversation. Collective intelligence is about giving all the right people a voice – not just the loud ones, or the ones who look and sound right.

Selection and filtering of people, and initial allocation of resources: Network analysis and construction of a database of potentially relevant entities is the first step that we discussed earlier, followed by the identification the first entities to bring together, and the decision about their role. For complex challenges, one needs a larger group of stronger-quality nodes. However, large groups tend to suffer from individual lack of accountability, and free-riding. Creating the right motivation (for instance, through transparency on contributions) is crucial. An additional solution can be to break down the problem into subcomponents that can be addressed by smaller groups (Agile methods advocate that for some type of work, for instance).

Are those people "contributor" nodes, or should they be **community managers**, helping engaging members, connecting nodes, solving queries and re-posting or re-sharing relevant content?

Do they have **equal rights**, or would they need to respect sophisticated rules to ensure that everyone's given a voice? This can be done like in Wikipedia are able to express themselves within well-detailed, community-enforced operating procedures[224]. Or it can be done by spreading rules and etiquette, and acquainting participants with diversity of communication style. For instance, mastery of language tends to vary in large groups, with the native speakers of the language used having a natural upper hand in the dialogue – that playing field must be leveled somehow.

And, where appropriate and possible, **is work allocated deliberately**, for instance with a sort of ticketing system whose algorithm matches resources (people, machine time) with tasks, or is it up to the contributors to decide where to work out of the open tasks? In the examples Wikipedia and Linux, for instance, the latter option is adopted, but contributions are subjugated to contributors having the right credentials. Or through some data and information feeds provided to the community, so that members can orient themselves to the right tasks even in the absence of a formal workflow? A supermind is often however non-hierarchical, which makes it impossible to force participants and cognitive resources to work on specific areas. Linux famously let its community work on whatever the individuals liked, which in the end turned out to be the most useful parts. It is unclear how much of that was due to "founding father" Torvald's charisma and presence, as well as the cohesiveness of the community (all coders), or the implicit incentives that led people to focus on useful things.

More generally, collaboration overload is a real problem, especially with central nodes don't know how to manage the load that their network position generates[225]. One solution is indeed a more deliberate allocation of work, with less-naturally-central but specifically trained people (e.g., community managers) taking some of that strain. Also important is coaching the central nodes (e.g., prioritization, use of technology, etiquette, access to support resources), so that they know how to deal with their network position.

Initial energization of the group: We will discuss incentives in more detail later but suffice to say that there needs to be a strong **value proposition** for individual entities to join a new network. To simplify, that can be anything that helps individuals fulfilling any level of Maslow's scale of needs – from basic food and shelter, to self-actualization. A new network's most important energy source is its ability to provide information, knowledge, connection that caters to the above functional and emotional needs of its participant.

In that respect, this step isn't different from any creation of an appealing product or service – for instance, at parity of resources, the venture-capital startup rule "go narrow" is applicable here as well: that is, it is easier to create a vibrant group around a specialized and underserved yet active segment.

Like any successful product or service, no group is interested per se in what the founders want. A new place is also attractive for individuals if it is consistent with some of their characteristics such as values and purpose – revisit the earlier section on mapping individual entities' traits.

At the same time, a strong "**marketing**" muscle can help. Outreach tools, starting with social media (e.g., Dux-soup, Followers DM, or Sprinklr.) They can also start celebrating punctuations of the lifecycle of the new network (think of what Change.org does), and in general create an amount of drumbeat that makes the network's pulse felt. These interactions tend to be virtual, but in-person contacts have positive roles where applicable, such as where there's enough critical mass in a specific locale to justify the creation of a "local chapter" of sorts – Meetup.com or tools like Eventbrite are obvious candidates.

Large groups of people can and increasingly will be initially engaged by **bots**, that can execute a simplified yet human-like conversation with the targets, taking an example from recruitment campaigns that have started utilizing such tools (e.g., Kriya.ai, Talkpush).

The **management of a collective intelligence platform** is not trivial, as the dynamics of such a system require the maintenance of a significant level of energy and excitement for the human nodes of the network, while providing enough guidance to avoid chaos. The discrete but present hand that govern Wikipedia and Linux, for instance, have shown how that balance is not trivial to achieve[226].

Flow and cadence of communication and engagement, for instance, play an important role. A simple example is the daily check-in through any instant messaging platform that many leaders of large, remote groups do. The point is akin to the ceremonies and rituals in methods like Agile at scale, for instance through scrum of scrums and the individual agile squad (production unit). Autonomy and alignment periods need careful design to ensure that individual resources express their potential but also drive the collective goals. For fact finding and creating situational awareness ("sense" mechanism), and always-on connection is good. But for creation of novel solutions and decisions, there is evidence that punctuated periods of connection and separation lead to better outcomes[227]. The choice of frequency (intermittent as opposed to continuous) of communications also impact group effectiveness[228] - whereby periods of "disconnection" from the rest of the group helps individuals explore, and punctuated reconnection helps maintain norms and avoid duplication of work. The use of synchronous and asynchronous tools can support flow and cadence, but only when done deliberately.

A related and interesting topic relates to the use of mindfulness practices to drive effective collaboration between the human nodes in the network. A recent study shows that mindfulness improves the collective intelligence of groups because it enables participants to be more present, balanced, and encourages leaders to become more self-reflected[229].

Providing **feedback** to participants is also a complex matter. While some of the general human capital management rules apply here – especially those that deal with the management and engagement of remote, virtual workforces – dealing with thousands or more people on a platform leaves little room for improvisation. For example, common norms must be created for ensuring that incivility is not tolerated, and only the appropriate behaviors are encouraged – for example, that people take turns when speaking to each other synchronously (something that Ambit.ai has addressed through its software), or that ideas from more junior or minority participants are taken seriously. Some of the concepts that we explored in the "network illumination and engagement" section of module 1 apply here too.

Iterations are particularly important especially when working with large groups of novices, laypeople and non-experts. The risk of failure in the solutions designed by those groups is higher than when only experts are used[230], though the upside is that ideas could be fresher and more radically new. It is no coincidence that lean startup and design thinking methods have made the process of iteration a core fixture of the process.

Working in large networks requires participants to monitor their language and optimize it even more than if they were in the office. That's because others will not be able to see body language as well as if the exchange were in person, especially when relying on non-video communications, and particularly when asynchronous (e.g., Teams). Participants should consider prefacing some of messages with "assume positive intent" language, and use (but not abuse) emoticons when appropriate to convey emotions.

One last word about the use of **data-driven insights** in supporting a collaboration platform. Most critically, collaboration platforms are very different from normal organizations in that it is very hard for individuals to "feel the mood of the team" by walking corridors and meeting people in an office setting. At hyperscale, "management by walking around" becomes extremely hard to do. Apart from being extremely gifted at operating in this environment, the organizers of a collective intelligence network can use data and multilevel communications.
For instance, in the reskilling supermind example that we have used as a case study so far, the provision of insight through advanced analytics was a crucial part of the work. In that case the central team ensured a periodic and deliberate flow of insight to the subject matter experts in charge of curating and sharing the knowledge, in a way that was both exciting (e.g., "see how you've performed this month compared to your peers") and functionally helpful (e.g., "see what type of content is most in demand"). At the same time, the learners, who else needed to be supported and motivated to expend cognitive resources in using the platform, were given a periodic score based on their learning behavior. The score showed their progress, compared it to others, and broke it down into areas so that the individual could understand how much effort they put into learning specific parts of the curriculum

We will explore some more the type of people who are needed to build these alternative organizational networks in one of the final chapters.

Monitoring of the collaboration metrics

As in the case of the previous modules, and especially in the case of such a complex infrastructure, it is useful to carve out system resources and capacity to monitor key metrics in order to understand what works and allocate capacity appropriately. Unlike other mechanisms that provide some quantitative feedback in a decentralized manner, what we are talking about here is a somewhat centralized function that can step back and consolidate data sources.

Some examples below:

a. Contribution of nodes to specific parts of the infrastructure

b. Consolidation of metrics related to network structure, incentives and information-processing from the previous two modules, in order to obtain a comprehensive view of who has done what, and what the outcome has been.

In the example of the reskilling supermind, a central team constantly monitors the activities that generate content for the end users and uses that to allocate senior leaders and troubleshooting capacity. In the absence of such visibility, directing attention in a timely way to what (and who) matters would be akin to finding a needle in a very big haystack.

The process of designing an augmented collective intelligence

If in the early 2000's I had asked the best innovation consultancy in the world to reimagine the Encyclopedia Britannica, the answer – even through the best design thinking methods – would have sounded like better user interface, better technology for engagement with the users, more interactivity etc. But thinking about an encyclopedia as a dynamic collection of knowledge enabled by an internet-based *community* was what ultimately changed the world. By "forcing" the innovators to think about the properties of a community, or other forms of collective intelligence, can propel creativity beyond our traditional, largely individual-focused and hierarchically driven mindset.

Designing an effective supermind takes more than just a combination of the modules already discussed. Just like any other design process, it requires a set of methods that enable a team to collaborate and deliver a concept, a prototype, or a final blueprint. That team itself can be a supermind, comprised of a small group of core members who engage systematically and appropriately with a much larger group. By designing a system (organization, processes, and technology) composed of a connected-intelligence network, one supermind that designs another.

The people involved in that solution design will benefit from the four-module architecture and other concepts illustrated in this document. However, the actual process of designing a solution benefits from a series of steps enabling the participants to truly *think in systems*, as opposed to following traditional solution design that focuses on technology, people, and related processes.

An underlying consideration, especially as the use of machine learning and AI increase, is the one of participatory design. Bias and exclusion issues have been documented in algorithms for many years, and it is critical to keep them in mind as we design solutions. The use of the methods discussed in this section and throughout the document helps contain the worst problems: human centered and inclusive design, designing with humans in the loop, participatory modeling (where stakeholders groups are well represented). Literature exists that addresses this topic specifically[231].

In this section, we share a collective-intelligence network design playbook. At its core, design thinking provides a useful methodological scaffolding to the design of a networked, connected supermind. There are also a few new methods that MIT's CCI has experimented with which can help design such a system, and we will look at them briefly in this section. They can intuitively overlay design thinking's "double diamond" structures – from the exploration and definition of the problem, to the exploration and synthesis of solution (and related iterations) – to complement existing design practices.

For example, in the exploration of the challenge, a customer journey exercise could be "augmented" by thinking about what individual personas experience in terms of a different type of supermind – for instance, by asking "is this behavior or feeling influenced by the fact that this user belongs to a community of people with specific norms". Similarly, problem-tree or even a traditional lean-six-sigma value-stream mapping analyses could apply the lens of alternative cognitive frameworks, for example by asking "is this part of the process trying to sense the environmental conditions or creating new solutions". In a later stage, such as the during the exploration of possible solutions, the collective intelligence framework could be invoked again, for instance by asking which technology solutions (from module 4 in this guidebook) can enable the incentive mechanism (from module 2) needed for a community (from module 2) to sense (module 3) and create.

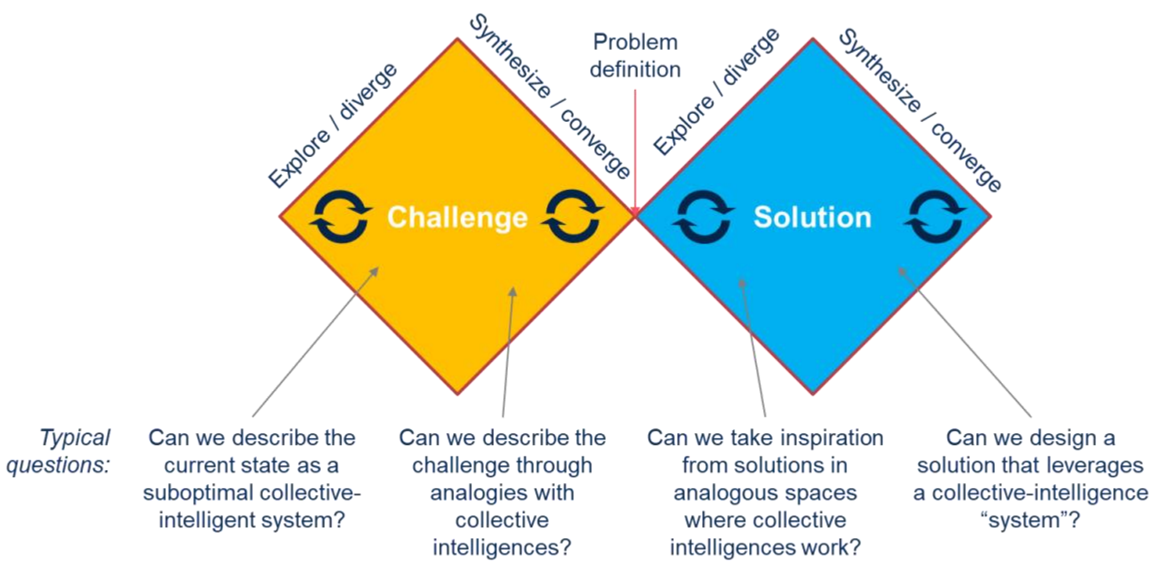

Apart from the thoughts in this guidebook, you may be interested in additional blueprints for the design of collective intelligence, replete with step-by-step artifacts, that are emerging: two examples are Nesta's toolkit repository[232], which includes for instance a specific Civic AI Toolkit[233].

The highly iterative process is summarized in the diagram below – which, once more, can be overlaid to each of the two sections of a double-diamond structure. Iterations are necessary to ensure fresh ideas from novices and laypeople are filtered and improved, hence reducing the risk of failure of the final solution. The chevrons at the top ("define goals" through to "design the system"), as well as the broad "diverge" (explore) and "converge/sustain" phases are relatively common. The way they come together in a supermind's design process is not. A design process enabled by a supermind starts by sensing (defining the problem and identifying the network available) and remembering what exists that can be used (an information feeder), then proceeds through creation and decision phases that require collaboration platforms and methods.

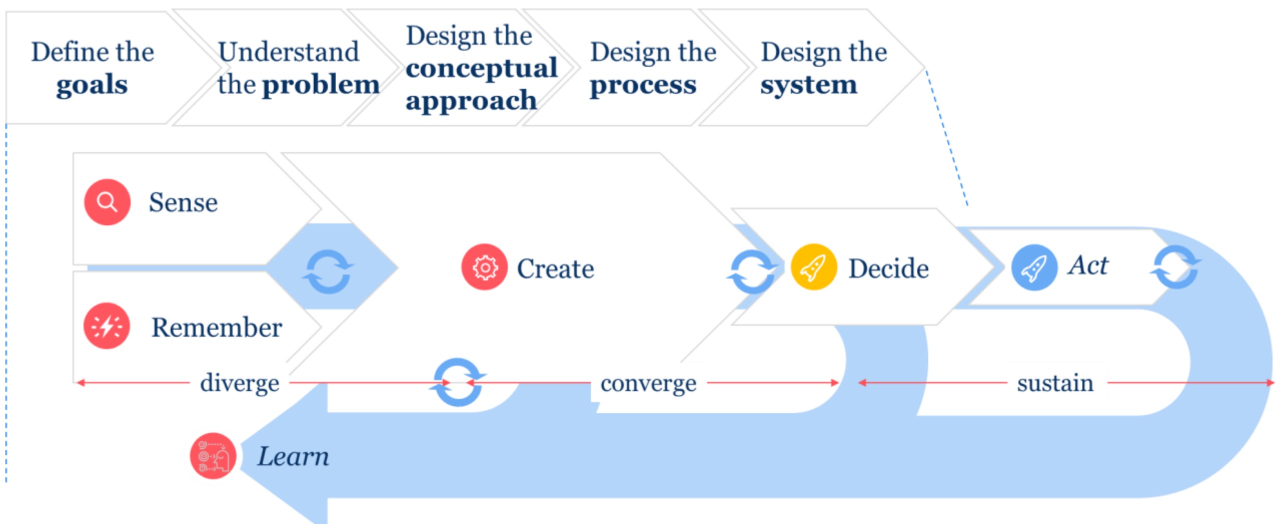

The overall process can last a few months, with important milestones and punctuations (workshops, standard meetings, steering committee meetings as represented by dots on the chevrons) and interim deliverables that enable the client organization to obtain visibility on the progress. There are three major steps in the design process: diagnose and understand; create; and decide.

All these steps can of course be performed through traditional in-person meetings. But it is more powerful to adopt collective-intelligence methods and harness larger groups with no location restrictions. This helps cast a wider net and accesses a more varied and vibrant group of people who (if led well) can generate more creative and unusual ideas as well as validate their viability.

Keep in mind that teams will need to oscillate between engaging a small number of people and large groups, as shown in the chart below. Smaller groups typically find it easier to collaborate synchronously, while larger groups are easier to manage asynchronously (e.g., through surveys).

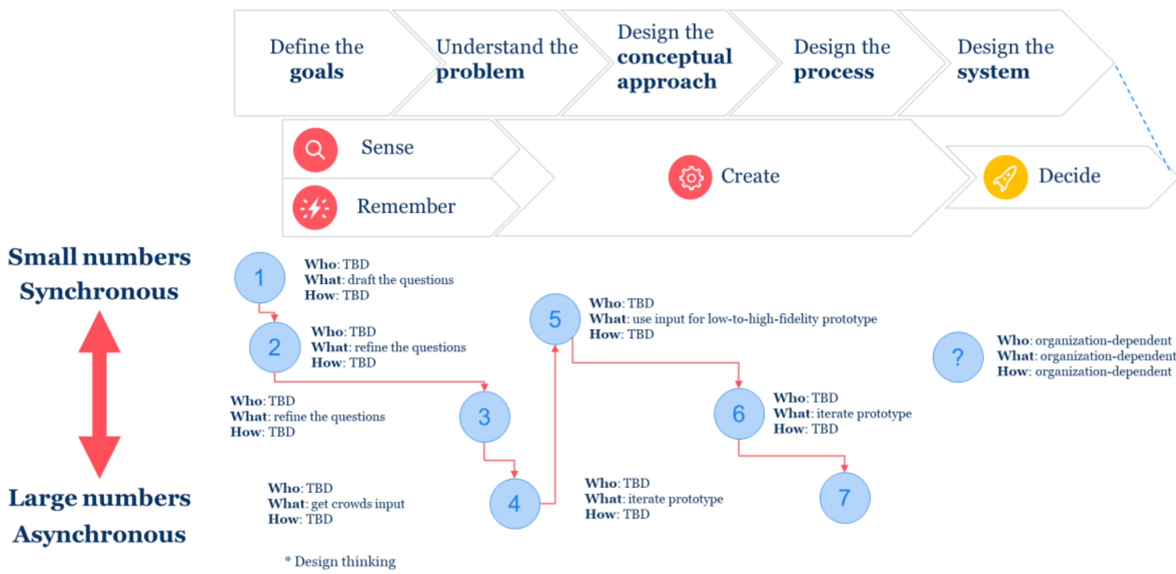

That oscillation isn't fully new: "World Café" methods have promoted breaking down large groups into smaller ones as part of workshops for many years. What can be novel is to do that at hyperscale, possibly going beyond synchronous workshops, for instance through asynchronous threads in private social collaboration groups (e.g., Microsoft Teams or Slack.) And as we noted in the "collaboration platform" chapter there may be some synchronous collaboration in large groups, for instance using today's virtual canvas (e.g., Miro or Mural) to collect ideas.

The below table summarizes (without being exhaustive or overly prescriptive) which type of groups can be leveraged when designing and building a supermind.

	Explore / diverge	Converge – concept to low-fidelity prototype	High-Fidelity prototype to MVP
Crowd (very large number e.g., 50+)	Provide perspectives. Sense.	Critique ideas. Provide different conceptual angles for the concepts. Remember what worked / didn't.	Critique / support prototype (decide). Tests MVPs.
Expert groups (large number e.g., 10+)	Critique the initial framing of the problem. Provide initial remembering of what worked.	Synthesize concepts from previous phase. Remember what worked / didn't.	Critique prototypes and MVPs.
Individual experts (<5)	Frame the problem (create).	Synthesize ideas and create concepts.	Critique and rework prototypes and MVPs.
Other facilitators (e.g., design experts, scrum masters, knowledge curators)	Knowledge management support. Help clarify intent and language. Engage communities and nudge people into action.	Synthesize ideas and create themes / building blocks.	Synthesize and clarify language, visuals etc.

In a design thinking process, the choice of design exercises is critical. Many exercises are adapted to small-group design and particularly lend themselves to specific phases of the process. Some can be portable to large-group, supermind-type creation processes. And all of them can be useful in designing connected-intelligence networks.

One important point to note: many of the techniques below are derived from design thinking or similar frameworks, so it should be relatively simple to find people who know how to use them. There are a few distinctive activities though that can't be fully undertaken without fully understanding the four-module framework laid out earlier. In particular, the technical notes of those modules include many real-life examples that can be brought into the "alternate world" and "advanced technology matrix" exercises.

The simplified table below shows a schematic example of the combination of design thinking techniques complemented by *supermind innovation* activities. It elaborates on a real case example that will be further described in this section – the *ideation* phase to generate possible ways to treat depression by enlisting the support of networks of people and machines. The efforts are targeted at disassembling the current processes that tackle depression and creating desirable options. As with standard design thinking techniques, significant emphasis is placed on empathizing with the people who are part of these processes.

	Existing	Exercises / activities	Raw ideas for future use
WHAT	Diagnosis and interventions	Abstraction: e.g., depression to happiness, reduction of stigma	Reframe the challenge to address happiness and social stigma, not just clinical depression" to remove friction in the early identification of symptoms, and enable patients to seek help through communities
HOW	Psychological, pharmaceutical and to a limited extent, other interventions.	Alternate worlds (e.g., "finding" happiness" Pokémon Go, Google Maps' data layers). Combination with new technologies (e.g., virtual assistants, social media networks, sentiment analysis, gig economy platform)	Virtual assistant helps patient find "happy places", "happy things to do"; as well as subtly support friends and family as caregivers
WHO	Doctors, specialists, and to a lower extent, lay people	Persona and journey analysis	Friends, colleagues, family, volunteers, community centers as broader network whose performance is augmented by AI-powered machines that pinpoint needs and opportunities

A list of activities is in the table below: all entries are classified based on their ability to enable specific phases of the design process. For more detail and templates on each of the most typical techniques, a quick web search will provide the latest details on how they're currently practiced.

Collective Intelligence Design Methods		Sense	Remember	Create	Decide
Problem trees	Logical breakdown of the components of the problem - can be sequential/directional (think fishbone diagrams in lean management) or amorphous (just mapping connections between elements, and then clustering elements that have affinity with each other).				
Persona mapping	Participants detail the characteristics of the personas (peoples in specific roles) involved in the challenge we are solving. E.g., passengers in an elevator we want to make faster; or building engineer. Where appropriate, one can map the characteristics of an entire network as one supermind, with its own aggregate cognitive and behavioral traits.				
Journey mapping	Drawing the journey of the personas involved in the challenge and identifying the emotional connotation of every step of the journey - with the objective of both mapping the sequence of events as well as developing empathy for the subjects. Where appropriate, one can map the journey of the entire network as one supermind, with its own aggregate cognitive and behavioral traits.				
Abstraction laddering	Identification of the right framing of the problem.				
Extreme users	Extreme users, typically super-users who use existing services and products to an extreme level, generate interesting signals for a design process. Groups of them often create unique dynamics, and even individual ones generate network dynamics worth analyzing. Prototyping with and for extreme users provides additional insight.				
"Which supermind" assessment	Analysis of current or future state, by utilizing the supermind lens (hierarchy, community, market, democracy, ecosystem). This activity can be applied independently, or overlaid on top of other activities (e.g., stakeholder mapping)				
Cognitive processes analysis	Analysis of current or future state, by utilizing the supermind cognitive process type lens (sense, remember, create, decide, act, learn). This activity can be applied independently, or overlaid on top of other activities (e.g., journey mapping)				
Stakeholder mapping	Two-dimensional network diagram of the stakeholders in the challenge. Where appropriate, one can map the interrelations of entire networks as superminds, with their own aggregate cognitive and behavioral traits.				

Collective Intelligence Design Methods		Sense	Remember	Create	Decide
Journaling / interview / observation	Any ethnographic methods aimed at developing a fact base to empathize with the human subject of the study. Potentially complemented with quantitative surveys to understand prevalence of attitudes etc.				
Affinity clustering	General method that surfaces ideas (often on simple sticky notes) and then clusters them together logically, developing implicit taxonomies in the process.				
Alternate worlds	Utilize examples from other fields, including somewhat unrelated ones, to foster creative ideas. Can be done in a scalable, collective way through mindmaps, google docs, Mural.ly, Padlet, etc. Human-machine approaches could eventually be used too[234]. Refer to the technical notes section to compile a list of examples and ideas.				
Statement starters	Engaging participants to fill blanks of a specific problem statement intended to solve a broader challenge "how might we ____". E.g., when needing faster speed for elevators "how might we decrease the time passengers perceive".				
Advanced technology creative matrix	Table that forces collision between two conceptual groups: available technologies and practices (e.g., NLP, predictive markets) vs. type of augmentation for the humans involved (ingestion of data, consolidation of data, prioritization of work, recommendation of decision, action/decision automation, collaboration). In the future, smart semantic search (e.g., SparkBeyond) could augment this step and dynamically surface unconventional options based on "naturally-occurring experiments". Refer to the technical notes' sections of each of the modules to derive a list of technology ideas.				
Strategy recombination permutations	Individuals or small groups identify the key dimensions (e.g., price, cost, distribution, tech) for a problem (e.g., a strategy, a product). Then generate all permutations of the values from those dimensions assessed by individuals or small groups, to identify combinations that might have merit. Then ask group to further refine.				
Polarize ideas	Generate additional ideas by removing some conventional mental filters (e.g., "what if we had all resources in the world? And if we had very limited resources?" or "if we could change moral conventions to make the idea acceptable"?) etc.				

Collective Intelligence Design Methods		Sense	Remember	Create	Decide
Concept posters / low/high-fidelity prototype	Increasingly precise mapping of the constituents of the solution - from general concepts to eventually UI. Needed for obtaining actionable feedback.				
Write your press release	Participants write a fictional press release imagining they were publicizing the success of their solution. In so doing, they highlight WHY/WHAT/HOW of the solution.				
A/B testing	Voting dyads of alternative solutions.				
Voting (e.g., rose, bud, thorn)	Participants vote (typically with round colored stickies) on concepts or constituents of concepts. Three positions "great idea" "hard to do / less valuable" "has potential but needs development".				
Pre-mortem and post-mortem	Storytelling done by individuals and then consolidated. Topic "in 3 years into the future, if this initiative has miserably failed, what would have happened for it to end so?". A form of scenario planning. Similar to the classic “How might this fail”				
Shared values (community)	Individuals surface their set of values, debate the emerging values from the community, and create a set of joint values narrative that highlights the common ground.				
Shared purpose (community)	Individuals surface their purpose in the community, debate the emerging purpose of the community (story of me, us, now), and create a set of joint narrative. Part of this can live online (further elaboration, voting).				
Community summit / jam	Large scale event to generate cognition + affection for large community formation. In person.				
System usability scales	Derivation of UI scorecards.				
Feature trade-off	Participants force tradeoffs between components at the solution, to arrive at a ranking of importance.				

Collective Intelligence Design Methods		Sense	Remember	Create	Decide
Network intelligence structure modules	Participants identify gaps and opportunities in their process/system understanding of the supermind by utilizing the 4-module supermind architecture				
Importance / difficulty matrix	Participants map the components of a solution on a two-dimensional space (importance/risk of the component, vs difficulty/cost/risk)				

It is worth taking a deeper look how some of these play out during the innovation process. It is likely that in the future semantic search and purpose-built libraries of alternate worlds and technology trends will do a more scalable job of mining the collective intelligence (and respective artifacts) that exist on the web and other knowledge repositories. The next chart highlights some examples of that information feeder. Large companies and consulting firms will be able to build such libraries and search mechanisms, and there may be demand for socially shared open-source utilities for specific uses, such as health care or education.

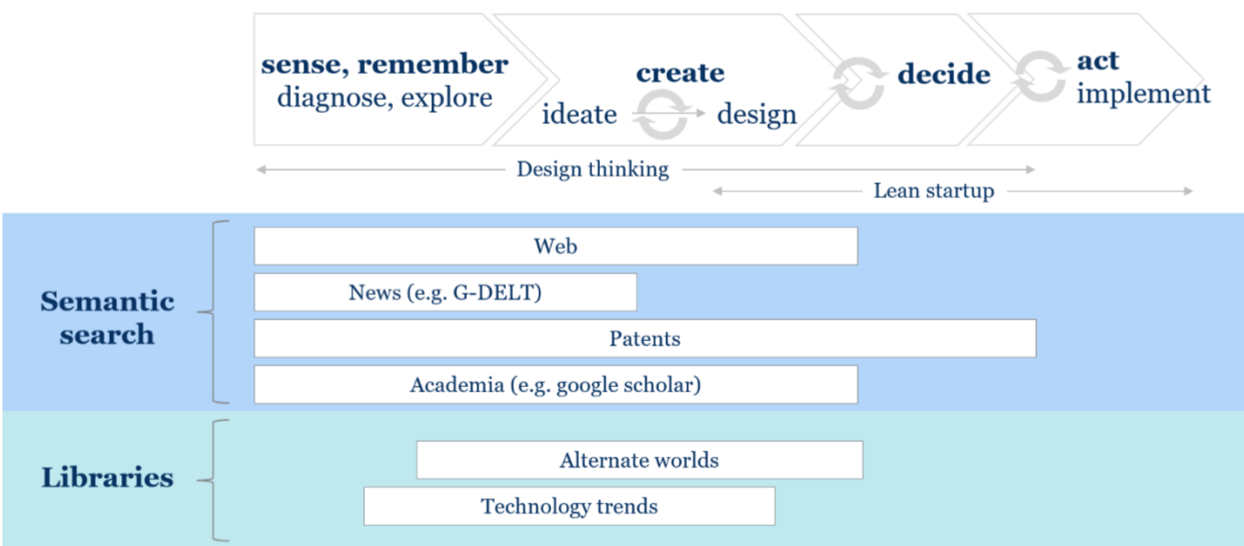

Let's close this chapter with a case study that describes the application of similar methods in a real situation.

Case study - Building superminds, and innovation capacity, at Takeda pharmaceuticals

Even in the absence of accelerants like the COVID-19 pandemic, the world of life sciences is undergoing a profound transformation. Big breakthroughs seem harder and more expensive to achieve. Genomics have yet to yield scalable impact. AI's promise is still somewhat unfulfilled. Digital transformation, improved prevention, and community-based healthcare continue to impact the industry. And many pharma executives, brought up in a different paradigm based in biomolecular science and R&D, struggle to embrace the new possibilities.

All the while serious conditions persist. One is depression, which severely impacts tens of millions of people of all ages worldwide, costs societies hundreds of billions of dollars annually and in many countries was only recently considered a real disease.

It is against this backdrop that Japan-headquartered Takeda in early 2019 launched a highly unusual initiative, spearheaded by Magda Schoeneich, leader of the company's Center for Scientific Leadership and Innovation. Inspired by a decade of research at MIT, Takeda identified a specific and particularly difficult healthcare challenge – depression in Japan – to test-drive new ways to improve healthcare and build its own leadership capabilities. It chose a cadre of executives to experiment with new ways of designing patient solutions based on superminds. Takeda's big question: how can superminds be part of healthcare's future?

The collaboration between Takeda and MIT aimed at developing about 50 Takeda executives, gathered from many functional areas and operating geographies. This process helped innovators *think in systems*, and effectively acted as an extension of design-thinking innovation techniques. The approach included virtual and in-person sessions, and leveraged a diverse ecosystem of people, including MIT faculty, students and affiliates that added new perspectives to the solution.

The outcome was the conceptualization of a technology-enabled platform, a *healing infrastructure* that harnesses the ability of networks of people and increasingly AI-powered machines. Takeda's concept detects depression's signals as soon as possible and provides the right support to the individual who deals with it. This including patients as well as those who surround them, including caregivers, doctors. The significant difference in this approach to depression is that ecosystems, communities, and specialized markets (e.g., health and wellness workers) made of people and their networks, and enabled by increasingly intelligent machines, are an active part of addressing the patients' condition. This is a departure from the traditional approach centered on hierarchically-organized individual doctors and medicines, and from top-down, technology-centric (not people-centric) solutions.

The network's power is amplified by surfaced data (such as interaction data with early-stage patients that wouldn't otherwise be collected) and ultimately brings the power of people to the forefront, with machines as a force multiplier. The resulting ensemble of people and machines networks is likely to adapt better to uncertain future conditions than humans and machines individually. The concept is provided to the public under a Creative Commons license, so that anyone can build upon it.

And Takeda gained a cadre of executives experienced in these methods, able to propagate them within the company through their own networks. For more detailed reports on this case study, visit the MIT Collective Intelligence site[235].

Things we must get smarter about

In this section, we conclude with discussion of future-oriented and increasingly important topics, such as the impact on leadership roles, distributed-workforce collaboration, the evolution of language and communications, and the risks to society of increasingly intelligent networks.

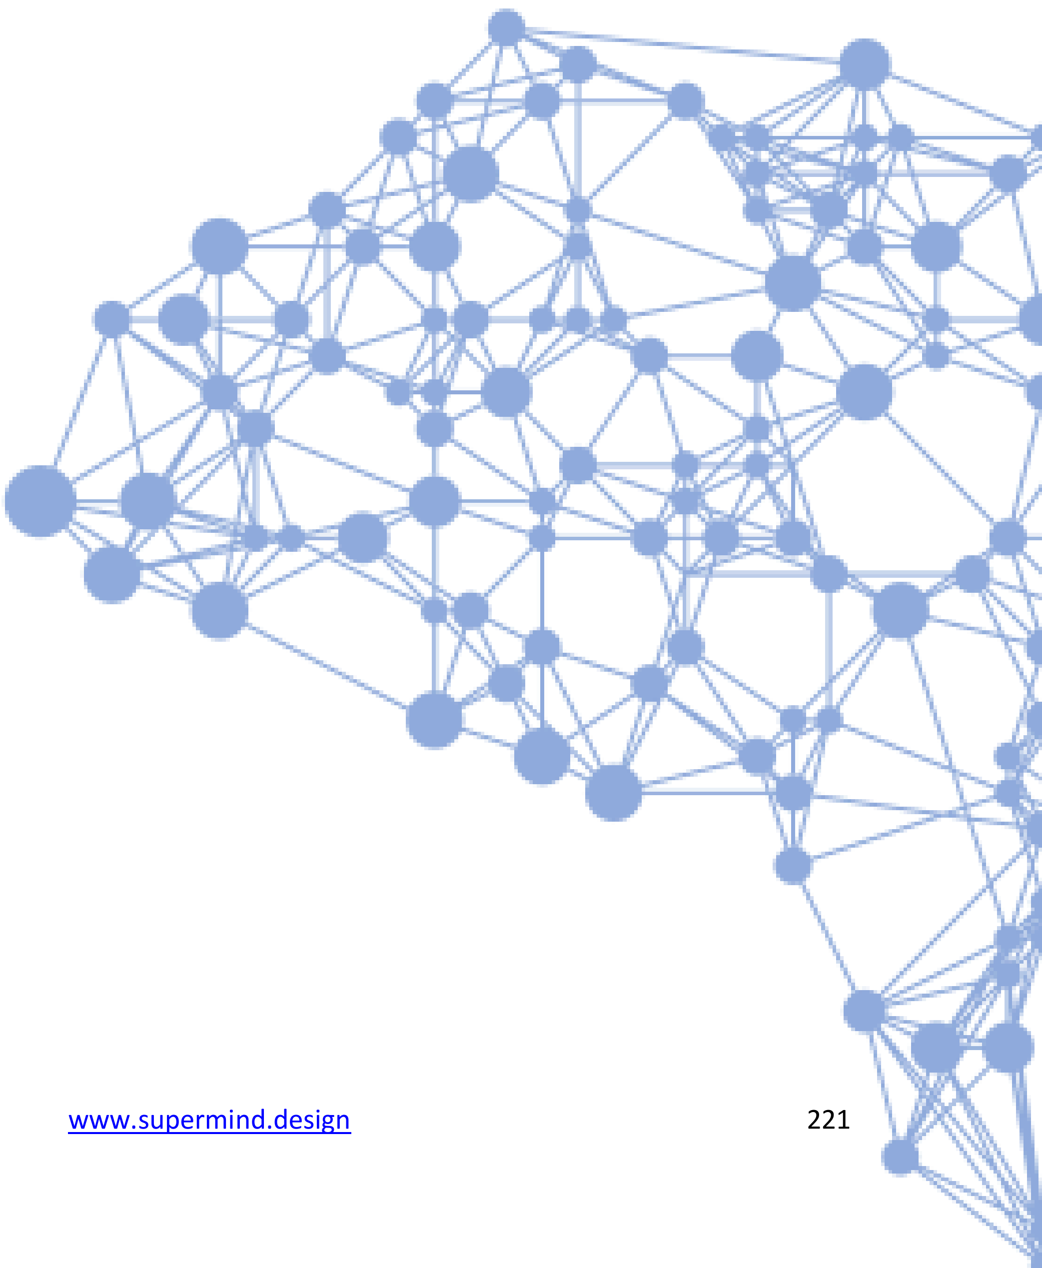

What can go wrong?

A lot, actually.

In the year 2002, during a trip to Myanmar's capital, Yangon, I witnessed a unique spectacle: in the city's main exhibition center, three public-use computers were, for the first time in the country's history, officially connected to the internet. Swarms of teenagers flooded the scene. Their intent? Mostly downloading pictures of celebrities, especially famous soccer players. Access to the internet was still officially prohibited. Invited by a local NGO, I spent an evening discussing what the internet is with a group of saffron-robed monks. It was hard to do, as they had only heard about it. A decade and a half later, social media networks leveraging the same technology and hyperconnectivity were used to fuel ethnic hatred, contributing to acts of genocide. It was the same country. It was possibly some of the same kids and monks. And it was for sure the same web.

During the COVID-19 epidemic, researchers estimated that about half of twitter accounts spreading news about the epidemic might have bots – generating content from somewhat innocent feel-good to possibly state-sponsored misinformation[236]. All generated confusion and jammed communication channels meant for humans.

Political polarization, especially but not only in the United States, has been extensively blamed on social media reinforcement of homogenous communities.

A big, global debate is raging on how to monitor the rationale of AI engines' decisions and ensure they're ethical and accurate. From social-media bubbles that polarize public opinion – including in sensitive election times - to algorithmic bias that denies credit to certain segments of the population, one thing is clear: threats don't look like what they used to.

We know that in nature, network intelligences can be hijacked. Parasites' can take control of a host's mind (this has been so far studied in insects[237]). Intestinal bacteria can sway animal and human behavior. Can analogous dynamics happen in superminds? For instance, homophily (the tendency of people to follow other people like them) and general bias prevent some nodes from contributing meaningfully. These behaviors also encourage the network's intelligence to listen too much to wrong information. The women leadership gap in large or technology-oriented enterprises is an example of important parts of the network whose impact is routinely dampened by retrograde social norms.

Because individual brains are biased towards exaggerated or outlandish claims and amplified by algorithms which maximize web-browsing stickiness, social networks have the tendency to propagate negative and false information[238]. This effect is also rooted in other social dynamics, like the ability of vocal, extreme minorities to make themselves look much bigger than they are: for instance, climate change skeptics have a much higher share of voice in online threads than their actual number would warrant.

As a result, misinformation is easier to spread than ever, and that will get worse before it gets better. The viral nature of inaccurate or deliberately misleading information will grow, due to increasingly sophisticated rich media: think of "deepfake" videos where digitally retouched pictures and voice impersonates influential personalities. By one estimate, there were 15,000 fake videos in 2019 – most of them so-called "revenge porn" aimed at women (including some high-profile cases like Rana Ayyub, an Indian investigative journalist). Synthetic media is an exponential trend, especially now that the tools for building reasonably credible audio and video are being democratized. Social media echo-chambers can sway superminds' information feeders too. And sophisticated tools inadvertently help networks turn insular and self-referential, showing patterns that network analysts call "small world", where reinforcement of established biases can shut out necessary self-doubt. AI itself is notoriously prone to bias because of data or algorithm limitations.

Blockchain itself may be misused, for instance by fueling unpredictable swings in financial markets (in addition to cryptocurrencies). The impact on individual investors, as well on the trust for the financial system, could range from significant to enormous.

AI is being deployed to catch some of these issues, especially on large social networks. But the ability of AI alone to automatically eliminate inappropriate content has its limits, especially because social norms and free-speech laws are hard for machines to fully apprehend. Even the exception management of less than one percentage point of social media content requiring human intervention amounts to hundreds of millions of content pieces that need human cleanup each year.

There's more. Machines alone are good at catching some things better than others. Spam is relatively easy to identify, as are bots impersonating people accounts, and nudity. Hatred is much less so. As a result, legions of content moderators are deployed worldwide, literally augmenting machines, not the other way around. To address this, some solutions are being devised through both private and public bodies, and include things like:

Rework the algorithms to prevent or mitigate these problems, for instance by introducing rules that limit the amount of unproven information they feed to large groups, or force serendipity into the choices of content served (an example of this is Spotify Discover Weekly functionality that can encourage exploration beyond one's comfort zone; a second one is the "diversity" setting in Genpact's virtual watercooler app, that encourages encounters between colleagues who share very varied enterprise networks).

Slow algorithms down so that their impact can be seen and countered more easily (not unlike some of the discussion about high-frequency trading in financial markets). Facebook did that in 2019 when it limited the number of people who can receive a message forwarded by another user.

Visibly label and possibly ban bots that amplify despicable human actions. Social media networks routinely eliminate bots from their user base; and cyber security companies monitor new malware strands to stop them before they proliferate. The same concepts could prove useful in securing collectively-intelligent networks.

Improve, and deploy at scale, **natural-language-processing and computer vision filters** that block not only inappropriate but also inaccurate content, or at least proactively queue it for review. This is not unlike what happens, after the fact, through content moderation organizations in social media companies.

Bolster fact-checking platforms in a partly automated[239] fashion using natural language analysis and network-credibility analytics such as Factmata.com or fullfact.org, the European Union's Disinformation Lab[240], Amnesty International's Citizen Evidence Lab which collects and compares additional metadata (such as location, weather), Snopes.com, or partly crowdsourced[241] for instance through some of the technologies discussed in Module 2 – How to do it on page 127. DARPA's MediFor project focused on algorithmic detection of image and video manipulation. Scientific research continues apace in this field (see for example MIT's PhD students "Vitamin C"[242] efforts). At the right level of scale, they could routinely sift through what's being said, detect possible malfeasance, and escalate to human experts for further analysis. As an example, Singapore's government in 2019 enforced regulations that require labeling of social network posts that spread falsehoods, for example. During the Covid-19 pandemic, social media tagged misinformation, partially with the help of users who surfaced suspicious information early. In general, research has shown that relatively small groups of laypeople can identify factual bias from small amounts of text, such as titles and subtitles, which bodes well for machine-enhanced crowdsourcing of fact checking[243].

Prevent algorithmic bias. AI only knows what it is taught by the data it is fed and how it is trained on that data. Bias can creep into both. For example, AI learns bias if the sample data inadvertently excludes certain ethnicities, or if a small AI team passes on its own bias. AI deployment for network information should have processes and governance that minimize mistakes leading to bias. Facebook, for instance, is starting down this path with its content moderation oversight body.

Make algorithms explain their reasoning path. The problem with much AI is that we can't trace the exact logic that led to a choice. But AI explainability is progressing fast across many fields. We should get it to work here too, so that we get to determine why a certain post was shown to a segment of people. That's particularly important as this space becomes the object of legal rulings.

Encourage citizen engagement, so that mainstream views are represented more frequently in the public discourse and online platforms. This may counterbalance vocal minorities' share of voice**.**

Proactively bring pro-social voices into the community forums. An example of this is Google Jigsaw, which for instance was used to build anti-terrorism narratives leveraging AdWords and YouTube to disrupt ISIS' recruitment networks[244].

Mandate the inclusion of alternate, fact-based points of view in newsfeeds that are at risk of becoming dangerously self-referential. The best newspapers have for decades allowed "dissenting voices" to be published on their real estate. Alternate points of view, and their authors, are fairly easy to identify with today's natural language and network analysis. It is not impossible to think that a quota for alternate content could be enforced by law.

Outlaw and publicly ostracize "hate-for-profit" posts, or posts that are likely to lead to personal damage. Natural language processing type AI can identify that indefatigably, if fed with non-biased information**.**

Legally protect individuals' rights to their personal imagery (photos and videos) and restrict social sites rights to the same materials to reduce the risk of deep fakes.

Formalize the role of third-party auditors that have the ultimate right of decision, and possibly institutionalize the role of public, trusted auditors. This envisions a sort of *Supreme Court for facts*. Ensure that powerful social media networks generate, **spread, and adhere to "ethical free-speech" guidelines for user-generated content**. Those guidelines should be based on legally approved frameworks and comply with testing algorithms (not unlike financial regulators do to systemic financial institutions).[245]

And finally, thoroughly **regulate the field** and introduce stiffer legal deterrence against willful disinformation. And possibly, follow Estonia in formalizing a "digital ID" that accompanies our most sensitive activities on the web. A version of that (stricter requirements for identity verification) was introduced in early 2020 by Google, in an effort to prevent political misinformation on its advertising platforms[246].

This is a very primitive, even naïve, list. Aral Sinan and others have analyzed extensively how social media changes the behavior of networks and of individuals, and identified a range of countermeasures to adapt[247]. Some of it will encounter fierce opposition on grounds of free speech but also (and more prosaically) because in the short-term it possibly impairs advertising revenues. This space is in the eye of the storm and will likely evolve fast.

What else can go wrong? AI's trajectory could take us to places we are utterly unprepared for. Are these networks conscious, or could they be? Can they express independent volition – especially the wrong type? Could they decide to modify their environment, as opposed to playing within the boundaries we give them?

Several books have looked at this eventuality, and they make for both intriguing and disturbing reading[248]. Even if these systems have no agency of their own, many scientists and philosophers have discussed the possibility of independent consciousness arising from the combination of their parts – a manifestation called "*phi*"[249] which represents the additional intelligence a network generates above and beyond the sum of its parts' intelligence.

A completely diverse source of risk relates to the propagation of the wrong type of behaviors and tendencies – (not just ideas -) in more tightly connected networks. Many of the complex ideas that travel in a network require "*complex contagion*" mechanisms for spreading, as discussed earlier[250]. What if those "broad bridges" were enough to enable the spread of, say, depressive or other unhealthy tendencies, on the back of rational dialogue.? We know that today's social media networks do carry that risk, and there's no certainty that more complex network intelligence wouldn't become a vehicle for those.

An additional risk relates to **individual productivity**, as people who participate in large online networks may end up being distracted constantly. As noted earlier in this document, these distractions are known to reduce productivity and creativity, as well as negatively impacting the emotional well-being and motivation of people. Individual and organizational practices and culture will need evolve to temper these unintended consequences.

Another clear strand of challenges relates to the treatment of **privacy** for all parties and especially private citizens who lack sophisticated means for defending themselves that corporates and bad actors often have.

Some research[251] related to social network analysis and its risks of infringing privacy laws has been done, especially in the European Union. In those cases, the law and this interpretation are progressing and will certainly become more consistent over time. In the short term, the immediate line of defense for any user is to ensure legal departments in organizations are conscious of the threats and have the right understanding of the application of GDPR and related implications.

One of the founding fathers of the web, Tim Berners-Lee, has spearheaded the creation of a set of methods and tools, call Solid, for ensuring individual data ownership. Other initiative, such as blockchain-based Rainfall.one, attempt to address the same challenge, with the added value to promote monetization for the end users in exchange for the data they create. Self-regulation is also a possibility – if internet players move advertising revenue away from selling advertising based on tracked individual tracking history.

Google started down this path in early 2021. If this is a broader trend, then relevant advertising will rely increasingly on natural language processing of the internet page's content, and less on the internet breadcrumbs based on browsing history.

Data trusts, where the rules are decided based on public good and guidelines shared, are another step in this direction. Open Data Institute and various national governments as well as the European Commission, are moving in that direction[252].

And finally, there's the obvious risk of **cyber threats**. With so much exposed surface in the form of easily accessible devices, including billions of IoT clients, we literally don't know what we don't know. Lots of so called "technical debt" has been incurred by bringing to market quickly devices that were never made for a thoroughly networked world, especially a world where cyberattacks are perpetrated by increasingly sophisticated actors. In the absence of regulations, easy "backdoors" exist in millions of devices – from baby monitors to printers. And now, criminals can use cybercurrencies to easily move relatively small amounts of money in and out of the blockchain world. As Zeynep Tufekci notes[253] "[…] *there will be some moves on the financial side (making it harder to get large sums out) and state-sector side (you can disincentivize another government from hacking your infrastructure, mutually, but it's much harder to do that to independent players). There may also be efforts "to make an example" of a few high-profile attempts: tracking down the people and handing down massive sentences. This isn't as difficult as it sounds, but it requires resources. If the ransomware attempts proliferate, the punishment will not be as effective a deterrent, because most people will not be caught, since so many are making attempts. This is essentially setting up a catastrophe lottery for the ransomware folks: most of them probably will not get caught, but the few that are will be crushed.*

But as long as it is this easy to cause digital havoc and hope to profit from it, and as long as getting small sums out without getting caught or punished is plausible, it will be tried again and again. Despite their potentially terrible consequences, it's difficult to address decentralized threats even if they are not very profitable."

Given this backdrop, one would intuit that such a threat should be address through augmented collective intelligence – whereby the four pillars (identification of the nodes, incentives, information feeders and collaboration structures) are used against such thread. In some respects, that model is already partially used by software security providers, including those who are part of large software firms – who use existing network devices to sense potential threats and release proactive patches to stop them from propagating. The open source software world also constantly produces cyberthreat resistant code. But indeed, many more devices have software these days, and many of the respective software coders don't have the means of the larger software firms – and in the absence of common standards and larger "commons of knowledge" .

The point is: the genie is clearly out of the bottle. Whether we like it or not, networks are developing at speed already. As seen on the Dark Web, where complete anonymity encourages illicit transactions, not all of them are in the hands of socially responsible actors. Like all powerful technologies and methods, our job is to harness this power for good and prevent its exploitation for malicious intent. The time to prepare is now.

If the world knew what the world knows: the future of search and knowledge management

Genius and stupidity seem to coexist at an unprecedented scale in our world. As Edward Olsen said, the interplay between our paleolithic brain, medieval institutions, and advanced tech is at the root of many of our struggles. The collective intelligence emerging from those three elements is constantly tested and often fails — from populism to social media gone awry, to pandemic unpreparedness and climate change. It often feels like we are fighting tomorrow's challenges with yesterday's intelligence.

But there's one significant reason for optimism, as one *very* large resource is largely untapped. *Our world routinely throws away or ignores the knowledge we create*. You can see it in your own daily work, and the work of your organizations: every day, we reinvent wheels, and we don't access the right people (or organizations) at the right time to find (or remember) solutions. Our collective brain isn't functioning as well as it could.

Thanks to the web, in the last twenty years, we have wired our collective brain in unimaginable ways. The world creates an astonishing amount of data — and knowledge — and makes it available online. It connects people in incredible ways that would have felt like sci-fi at the turn of the millennium. Yet, when it comes to harnessing planetary knowledge, we haven't seen anything yet. Today there's immense and untapped potential because of the convergence of a few powerful vectors. Consider these examples.

One of the most important innovations of the last twenty years has been the **search engine**, intended to "organize the world's knowledge", in Google's words. Even video and audio content are easy to search today.

AI's natural-language models have enormously progressed in the last years, leading to astonishing tools such as GPT-3 which have made language understanding and generation a lot easier. Beyond written language, image processing and generation (think of Dall-E2) have also evolved by leaps and bounds.

Knowledge-graph technologies that establish relationships between concepts, people, and organizations ("entities"), make the world's knowledge even easier to mine. New tools use that to enable richer search and, when combined with natural language understanding (for instance, in science, this, this, and this), hold promise for the exploration of specific topics. And the ongoing re-mix of everything made through social media makes connections between ideas, people, and organizations explicit — some of which can be mined through publicly accessible APIs.

Data science, including its crowdsourced citizen data-science form, enables the use of new and existing sources of data, including the increasing amount produced by the Internet of Things (IoT) both public (e.g. heat measurement), private (e.g. Google's land development tracker), and crowd-based (e.g. Arduino based sensors).

People simply share more: thanks to self-publishing tools, and because of the importance of enterprise and personal thought leadership, the web is awash with publicly available content from companies that would have been considered trade secrets only a couple of decades ago. Scientific knowledge is increasingly retrievable, through specialized search engines (e.g. Google Scholar), portals, and networks (e.g. Researchgate), and because of the mounting pressure to make it freely accessible.

And it is not just "asynchronous" knowledge access. Modern cloud technology and sophisticated data compression algorithms make **video and voice** connectivity ubiquitous, at increasingly low levels of data speed, which makes synchronous knowledge retrieval and generation more frictionless than ever.

As a result, augmented collective intelligence (MIT's Malone calls them superminds) is emerging, as the image below illustrates for the example of enterprise innovation teams, who harness the internal and external ecosystem as an extension of their own brain.

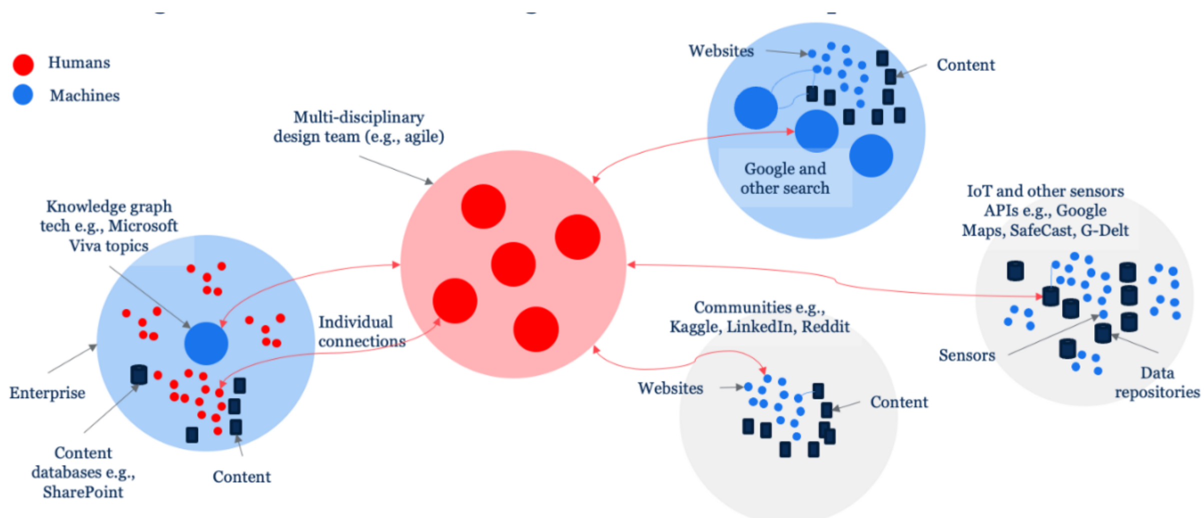

But our collective technology and methods pale in comparison with what happens elsewhere. The world does countless "natural" experiments (both in our society as well as in nature) that aren't harvested— unlike the "active inference" that our brain and in a way the natural world do. Take the following examples:

Search engines' algorithms, and their use, are still largely driven by advertising markets, not knowledge industries. Search engines and commonly-used methods do not make truly advanced search available to most people. For instance, they don't explicitly allow the exhaustive visualization of knowledge graphs, so that one could identify both content and people (and organizations) — as well as explore adjacent fields. Not all meaningful websites and content are inventoried. Much of the "new" action currently remains *within* enterprises through machine-learning-based knowledge management (such as Microsoft Viva Topics), but the overall knowledge ecosystem is many orders of magnitude larger.

Social media algorithms' recommendations optimize for predicted engagement (e.g., likes, or shares), not problem-solving or creativity. And try to follow the right people and the right topics isn't effortless: one can't easily find people to follow based on the field they're competent in. Similarly, **professional social networks** such as LinkedIn are not optimized for skill-based search ("which people work in my field?") and do not facilitate field exploration ("which subfields exist, and who works there?") or validation of ideas (e.g., assessing people's claims credibility by checking their — or their network's — skills).

Natural language models could proactively propose novel combinations of concepts for humans to refine — but they're not used for that purpose yet.

Data science, and the translation of science into respective models, is still very much an elite job. Data crowdsourcing (e.g. through citizen science) and increasingly easy-to-use tooling (e.g. XGBoost) show that the floor can be further lowered so that more people can come on board.

Surprisingly, **language barriers** are still significant and end up siloing up the world's knowledge. Think about it: web searches only show results for same-language sites: if you are in the US and look for "heat pump installation methods", you typically won't see content from (machine-translated) German, Japanese, or Chinese sources. And while translation engines like Google Translate have improved remarkably, they're still not used pervasively yet in a range of potential knowledge-sharing applications.

As a result, *we collectively don't learn enough from the experiments made elsewhere.* Think of "Global South" practitioners quickly learning from cost-effective climate adaptation projects in other countries, irrespective of whether they are documented in Indonesian Bhasa, Spanish, Urdu, Swahili, Hindi, or Chinese. Conversely, developed-countries practitioners fail to access sources of "reverse innovation" — lower-cost ideas developed under significant budget constraints. And, generally, knowledge "backwaters" exist: users in many (non-English speaking) countries prolong the use of old knowledge because they don't have access to the right networks in real-time (think of old schoolbooks and non-English language internet pages for technical topics).

Sadly, our **organizational design** practices reflect the issue: strategic knowledge creation and management isn't a C-suite role, and that job is often fragmented across departments — domain practice groups, the CIO, sales support, etc. which weakens the much-needed enterprise transformation. Across even broader ecosystems, incentive systems are still broken, as attested by academia's struggles to give appropriate credit and encourage more creative exploration.

And finally, and ironically, the respective digital product ecosystem doesn't attract as much attention and investment as others (venture capital anyone?).

In order to amplify and accelerate innovation cycles we need to build "supermind utilities" — possibly as public, or partially open-sourced, goods so that the global community can access them. They could be financed by governments, private individuals, or corporations. Over time, the return on such investments will attract more private capital, crowdsource contributions, and help develop business models that eschew advertising and make money by stimulating our pre-frontal cortex, not our amygdala. (The potential promise of some web3 technology could help, as and when it gets out of its current hype and greed cycle.)

To be clear — it is very likely there's a solid business case to build commercially viable digital products that cater to a type of "knowledge super users". The current challenge is to show them (and their C-suite), an easy and exactly quantifiable return on investment. As is often the case, the most sophisticated users, and the companies with the most foresight, will end up leading the pack.

Over 2,500 years ago the library of Alexandria ignited innovation across a chunk of the ancient world, and innumerable efforts have helped build repositories of knowledge over the centuries. The word "university" originally meant "community", and universities received funding to strengthen those (analog, organic) superminds — helping the respective networks and their knowledge converge. In the 21st century, augmenting the world's collective intelligence by building such knowledge utilities sounds like a reasonable thing to do.

These superminds will generate a superior intelligence, emerging from the network of knowledge and skills that exists below today's comparably superficial web-based interactions. They will help us fight tomorrow's challenges with tomorrow's intelligence, across:

- **known-knowns:** problems whose solution exists elsewhere, so collectively remembering and learning what works,
- **known-*un*knowns**, by creating and deciding on solutions that we struggle with, and
- ***un*known-*un*knowns**, by sensing low-signal but high-momentum trends that could quickly turn into major opportunities or threats.

Of course, lots can go wrong. To start with, we will need ways to mitigate our collective tendency to fall for unsubstantiated claims, counter rogue actors, and generally reduce trolling and abuse. But with the right incentives, methods, and capability, it sounds plausible that we will be able to emulate, for instance, Wikipedia and its collectively-enforced quality control.

Every single hour, the Earth receives from the Sun the amount of energy that the entire human civilization consumes in a year. We are getting better at harvesting that power. There is reason to believe that we are "leaving knowledge on the table" in similar proportions, and by harnessing our collective knowledge, we could harvest our collective *cognitive* power.

Building on today's technologies and methods, there's much that we can do about it. Let's solve tomorrow's problems with tomorrow's intelligence. Let's go build superminds.

How large should the supermind be?

Possibly the hardest practical question when harnessing collective intelligence is - "what's fit for purpose"? Is a small team enough, or do we need to orchestrate large internet-based communities with all their complexity? Most people will be familiar with the difficulty in getting real creativity out of small teams who know each other well, and the slog of sparking the momentum in larger groups of people who get together infrequently. The tradeoffs are real.

That is, depending on the challenge at hand, what group size is needed? And vice-versa – given the size of the group that can be reasonably be assembled, what can it do?

The general and not particularly useful answer is "do as small as you can, but not smaller". And most people typically end up doing much smaller than they would need to, trapped in hierarchical organizational designs born in low-tech eras.

The table below provides some succinct guidance. First, what size of the group tends to be better at what type of task. Second, what levers harness that collective cognitive power, specifically

- how to connect "network nodes" so they're discoverable to each other
- how to incentivize their contribution
- how to feed fresh ideas into them to fuel their inspiration, and
- how to make them collaborate to collectively produce – not just surface – better ideas, and execute them.

These are not guidelines, since the exact combinations of tools and methods are very situation-dependent and can't be summarized easily. But the indications in the table provide some guidance and at the very least some inspiration.

The lists of tools are not exhaustive – they're just meant to indicate directionally what tends to work at that level of scale. What will possibly surprise some is the variety of "how-to" which include digital and physical technologies and span a large spectrum from the traditional (e.g., office spaces) to the tried-and-true (e.g., enterprise communication tools) to the emerging (e.g., Web3).

	Very small networks	Small network	Mid-size (enterprise / web) network	Large enterprise community	Large web community	Very large web community / ecosystem
Examples	Strong-ties network, team under one manager or direct colleagues	Purpose-built community e.g., project specific. Strong ties	Purpose-oriented community (interest group like Subreddit, practice). All ties types	General communication channels (corporate Yammer). Mostly weak ties	Information broadcasting groups (LinkedIn). Mostly no or weak ties	Broadcast with feedback (Medium), newsletters, social media (Twitter) Mostly no ties
What it does best	Decide, learn, (create)	Create, decide, learn	Sense, create, (learn)	Sense, (create)	Sense	Sense
Connect it	In-person offices; regular governance cadence	Enterprise facilities (e.g., canteen, meeting rooms). List of virtual group (e.g., Slack channel participants)	People finder databases, directories / inventories. Enterprise facilities e.g., canteen, auditorium	Yammer, virtual watercooler type technology. Enterprise facilities e.g., canteen, atria, corridors	LinkedIn, Reddit, Discord, Clubhouse. Curated newsletters. Digital marketing	Online publishing, newsletters, social media (Twitter). Digital marketing outreaches (inbound / outbound)
Incentivize it	Hierarchical, extrinsic incentives (KPIs / OKRs), peers-related extrinsic incentives	Hierarchical, extrinsic incentives (KPIs / OKRs), peers-related intrinsic incentives	Extrinsic incentives: (KPIs / OKRs), external e.g., client, competition; peers-related intrinsic incentives	Extrinsic incentives (KPIs / OKRs, bonus plans). Intrinsic (enterprise purpose, culture, reputation)	Intrinsic (reputation, followership). Potentially extrinsic (web monetization)	Intrinsic (reputation, followership). Potentially extrinsic (web monetization), especially with Web3
Feed it	Team meetings	Meetings e.g., Agile scrums, learning tools, knowledge mgmt.	Posts feeds, learning and knowledge management	Posts feeds, enterprise social media, learning and knowledge mgmt.	Interest-based news. Search tools powered by knowledge graphs. Digital marketing	Search tools (knowledge graph), matching algorithm newsfeeds
Make it collaborate	Enterprise communication tools (Microsoft, Google, Slack, Zoom)	Slack / Teams channels, Zoom. WhatsApp. Google apps. Project management tools.	Slack, Teams, Yammer. Zoom. Subreddit, Discord, WhatsApp, Google apps. Facilities. Program mgmt. tools	Hackathons; townhalls; corporate social responsibility projects	Likes. Comments. Polling tools (Pol.is). Decentralized Autonomous Organizations (DAOs)	Likes and other social actions (share, etc.). Comments. DAOs

Whatever you do, it is critical to recognize the importance of being human-centered in the design of this architecture – because that's the key to tapping into human networks' cognitive surplus. But digital entities also have a role as "intelligent nodes" – think of knowledge management bots that help serve relevant information.

Preventing machines' input from biasing human judgment

Does AI impair human knowledge? Does it worsen individual and group decisions? What is the impact in traditional workflows, but also in broader decision-making processes?

These are broader topics than we can do justice to here, but it is worth providing some indications based on new data. There's a strand of research that investigates that, and it is key to understanding the value of digitally augmented collective intelligence. Just as a primer, a new paper ("Will humans-in-the-loop-become-Borgs"[254]) attempts to provide guidance on how to support individual and group decision-making through AI suggestions. It also presents a useful review of the existing literature, if you're interested in a backgrounder.

The experimental setting has limitations - e.g., asking people to classify images with or without the help of AI is hardly generalizable to all decision processes - but some of the insights are food for general thought. Some of them reconfirm previous studies' results, but taken together they hint at new possibilities:

> **[Human+machine] is better than either alone.** AI does improve individual and group accuracy, and pairing people with AI does result in better judgment compared to both AI-only and human-only decision making
>
> **Accuracy improves with the increased size of a human group**, and simple tasks require only small group size in order to be consistently performed accurately
>
> **Some humans disregard cues.** Telling humans how confident the AI is in its prediction doesn't automatically result in better accuracy (many people tend to follow AI's advice irrespective of the level of AI confidence)
>
> **Diversity is precious.** The level of accuracy improves when the diversity of human decisions is allowed and encouraged, especially when the people who are better judges (i.e., have better capabilities) are given the ability to influence both the algorithm as well as the rest of the group. To help people more generally, it seems to be important to encourage humans to express their critical judgment and resist the urge of somewhat uncritically validating the machine's prediction - in order to preserve the diversity of thought and leverage the unique knowledge that some people have. This is, for instance,

important when social media drives the formation of opinions but has clear implications on a range of other processes, including workflows, planning, etc.

A final, practical implication of these observations, inspired by some old (classic?) lean and six-sigma thinking: variance, not just averages, is a useful data point when optimizing most decisional processes. In this specific case, “variance/confidence” should be proactively presented to decision-makers in addition to the prediction, i.e., the UI/UX should encourage people to consider that variance seriously. In particular, the “confidence” level should be both based on AI’s own confidence, and also on the variability of the judgment expressed by typically-competent humans.

The net - don't trust machines and humans in isolation, and "out-of-the-box" AI-powered decision-support user interfaces. Better decisions can be made when they work together, in a smartly designed way.

Managing language in network intelligence

This topic could have been part of the "collaboration platform" module, but it is so important that it serves its own space.

Language is what makes our brains different[255]. It was the original killer application for our species. The adaptation and innovation in our societies and organizations rode, and still ride, on it. The pre-frontal cortex, the part of the brain responsible for much of our complex thinking especially in unknown situations, and hence for our superior adaptation and intelligence, heavily depends on language. It likely is also what makes our *collective* brain different. Language is the code that makes human nodes interoperate with others. It is the language of the instructions that we give to each other in large groups so that we sense, remember, create, decide, act and learn as societies, not just individual brains.

Human language evolves more organically and flexibly than any program code. Machines require more regimented semantic structures and syntax. The organic nature of language has made it resilient and evolutive and has led to the structure of modern languages like English that are so foundational in the functioning of our societies.

But we have hit the limits of the paradigm. Humans struggle with volume. Our cognitive processing power is limited, and our sensorial instrumentation can't handle much more than one input from a person at any time. From social media to IoT, today's deluge of information makes it difficult for organizations to absorb new ideas. Businesspeople in particular don't have time for verbose descriptions. A firehose of knowledge is indeed too much to drink, to the point that it is sometime more effective to reduce the intake, either by introducing curators and filters, and in the worst (and common) case just by cutting down the number of sources.

This limits our ability to expand beyond the original band of primates to harness intelligence on a global scale. Connected-intelligence systems draw strength from feedback loops, where nodes (people and machines) give feedback to each other's input. That input is, often, language.

Let's borrow from system dynamics to understand what happens. If we slow down feedback loops across the system (and they do slow down when the system is large and the language is complex) and if the amount of "knowledge stock" is high (which it increasingly is), there is strong *inertia* in the system. That means by the time feedback is fully understood and visible, it is also late. Late feedback requires effort and luck to recover, meaning we may make significant mistakes or miss out on important developments.

That's why we need faster feedback loops for the world's knowledge, and the feedback loops won't operate with people alone. Scalability of interaction is key, but our original one-to-one voice-based synchronous communication doesn't scale because no one can have multiple conversations at the same time. One-to-many broadcast scales to deliver information, but it doesn't provide scalable feedback loops. That is, I can talk *to* hundreds of people at the same time, but I can't talk *with* all of them. I still need to process their individual feedback linearly – if they're even patient enough to give me feedback without an individual conversation. Asynchronous, structured communications like those on enterprise social media scale a little better, for instance with idea roll-up in crowdsourced ideation, voting and the like, but are not widespread.

The most immediate hurdle however is not our finite biological processing power. It is our culture and behaviors. Mark Twain said, "If I had time I would write more concisely." He was aggrieved about the waste generated by unnecessary words, in an era when readers had more time. Today, our society is time starved. Unfortunately, most communicators (both in business and especially in academia) take time detailing their points without don't taking much time to streamline their language. For decades, American English usage has encouraged more concise prose. This may be because generations of migrants and non-native speakers had little interest in flourish. Just possibly they were more obsessed with economic practicality (and efficient interactions) than grand-sounding statements. That's likely a good thing. Now, what more can we do?

Remember that the value proposition of two of the most important innovations in this century – internet search (think Google) and social media (think of Twitter's character limits) – enabled humans to comb through text faster.

Below are a few thoughts aimed at new, more efficient communication practices that enable connected intelligence to flourish. Some demand research and innovation and can possibly be done by machines with increasingly good natural language processing ability. But if you can't wait, some can be turned into practical ways of working that most organizations should be able to implement today.

Content layers. Can we write books in 30 pages instead of 300, but better than the often-superficial summary services (GetAbstract, Blinkist) that exist today? And can those books help readers the ability to dive into concepts progressively, in a "pyramidal way" that combines management consulting "*Minto principles*"[256] with the internet-type feel of "double clicking" on hypertext? Can the note section become 70% of the word count, "side bars" another 20%, and the main narrative the remaining 10%? Can that 10% serve should be exposed to all readers and work as a path through that knowledge, a path that allows for exploration of smaller trails, as needed?

Proper abstracts. Can scientific papers provide a more practically useful and standardized abstract, instead of a content-description text followed by dozens of pages of disquisitions, so that users (yes, users of science) can peruse the main aspects in minutes? Can we facilitate diagonal reading, instead of forcing readers to improvise it themselves? Could we customize the text based on the amount of time that we can spend reading it?

Force concise prose and arguments. Today's innovation methods rely on stories for testing concepts. From Agile to design thinking, a lot of effort goes into compressing the time it takes to articulate an idea, so that feedback sessions can be held with more (and more diverse) people. There are even art forms that force minimalist expression – for instance, what's the Haiku version of everything we write? Amazon's founder Jeff Bezos famously enforces a concise format for meeting memos, which are religiously read by participants prior to the start of the meeting. Many companies have adopted it. Good editors are skilled at shrinking text, and they're increasingly available as gig workers. Can more of us use those practices? Can word processors do more editing, and can we use purpose-built apps to help us shrink our writing? For instance, Textio is used to craft language that avoids gender bias. Grammarly and Hemingway App support good writing. Can their algorithms do more to streamline our verbiage?

Algorithmic syntax for language. Can we tag and sequence our arguments in a logical way, proceeding from WHY this topic matters to WHAT the scope is to HOW? Experts start with this mindset, but it makes boring reading for everyone else. Can algorithmic language become master data, so that machines can combine text from any number of writers?

More visuals and hybrid narrative. Most books, articles, even PowerPoint decks are mostly words. But a picture is worth a thousand words, especially if accompanied by some textual context. That's why PowerPoint caught on in the first place. That's why so many got enthusiastic about Prezi.

Better visualization. Writing developed because our brains and hands can only tell a linear narrative, and our eyes and brain can only follow one while reading. Gutenberg had only paper for his printing press. The way we write today is still mired in those initial limitations, meaning the written word is the equivalent of a road with no intersections, underpasses, branches, ramps, or shortcuts. Can we navigate the semantic maze instead of relying on written lines of text to get us to do that? HTML links changed the world, but a true semantic web that weaves language together irrespective of the actual pages doesn't exist yet. Can we take more steps in that direction? We could structure visualization based on some of the concepts above, letting syntax drive more dynamic text outlines. Can we use mind maps such as iThoughts or MindMeister more, to look at content differently, or to complement increasingly rich outlines? Can related concepts be visualized in a more intuitive tree structure as overlays to the text we want to dive into, instead of requiring a manual scan of additional pages?

Metanarratives instead of everyone's individual story. From academia to thought leadership, individual authors are revered, and their contributions highlighted. Very little praise goes to the creators of compendia who summarize and weave concepts together. Can metanarratives, the creation of a story woven from many other stories, become heroes of science and innovation? Can the role of knowledge curator and weaver become as important as the role of author? Can we change incentives so that happens? Academic tenure depends on authorship when doing research. Should it depend more on making research useful by knitting it together in more digestible forms?

Dimensionality reduction. A simple example: if you want to know what 1+1 is, all you need is "2", i.e., you don't need the "1+1" data. The formality of algebra doesn't have a precise counterpart in grammar, unfortunately. But text analysis is an established discipline and can be brought to bear – for instance some language editors, like Grammarly, do help with scoring for conciseness, and AI-powered copywriting is becoming a thing. Written language often packs everything together: the "why" of a concept, together with the "what", "how" of it. Each element can be useful in specific circumstances to specific audiences, but they are never all useful to everyone. Can we write prose that only shows the reader what they need at that time, and allow them to dig deeper if they want? Can we apply to language the equivalent of algorithmic dimensionality reduction, simplifying features and data? What is the equivalent of a data science auto-encoder for language?

Eliminate accent-related unconscious bias to level the playing field. The ability to persuade another is linked to inspiring confidence. As a result, people with an accent aren't always perceived as competent. Foreigners sound less sharp because their vocabulary is less precise, and part of their cognitive processing power is used up processing language, slowing them down. A native-speaking audience may unconsciously form a bias about those things. The upshot? The power of some nodes is lost in translation, so to speak. The best way to prevent this is to develop and cultivate patience with non-native speakers. This is possible. In many parts of the United States and Canada with highly heterogenous speakers, there seems to be more leniency with stuttering foreigners. How does that compare with other advanced economies like France and Germany?

Train people to communicate better. The focus on STEM can't make us forget that technically minded people need language to work with others. Anybody involved in innovation should learn, and *keep* learning, how to communicate. Similarly, the way we learn foreign languages should change now that natural language processing and crowdsourcing are mature. Apps like Duolingo have paved new paths, but more can be done. Think of the ability to identify the words that a certain person uses most in a specific job and create a customized learning plan that reinforces their use.

De-jargonize. Domain experts tend to use words that, while accurate, are obscure for most other readers. They also tend not to spend time making their arguments explicit enough to be understood without specialized vocabulary. As noted in the practice of successful crowdsourcing communities, that's a real problem, because the best ideas are often the result of cross-pollination between different people with experience in different fields. Spending time making language accessible to more people and limiting reliance on obscure jargon can help ideas flourish in diverse networks.

Simplify languages. French, German, English, Japanese, Mandarin, Bahasa: every language has its own quirks, but the ones that have caught on for global communications between very different people have been simplified[257]. Some say they have they lost their ability to be precise, but it is unclear if that's fully accurate. Fewer words for the same meaning lower the barriers to entry for non-native speakers, and that makes more network nodes interoperable.

If we were able to do some of the above, we would gain the following:

- **Better feedback loops** within the connected-intelligence network
- **Better initial ideas**, as authors can spend more time structuring ideas and discovering gaps, instead of putting linguistic ribbons around them
- **Better ability for machines to make sense of language** beyond search enabled by word co-occurrence methods available today. Instead, they could further summarize ideas, and start connecting them for us by applying machine learning**.**
- **More machine curation.** Google changed what search meant to the world. Before it, search was an individual librarian's job, and dramatically less efficient. Can we develop more technology expressly dedicated to conciseness and curation? Many marketing organizations already use the Curata.com curation engine. G-DELT is arguably the largest assembly of news data machine-translated into English. Can machines do more of that, and humans move to more value-add curator roles? The promise of OpenAI GPT-3 is, at the time of writing, a possible avenue to an exponentially scalable future of knowledge curation.

Equally interesting is the opportunity presented by the impact of today's *language neutrality* technologies. They will eventually enable accurate translations, even of real-time conversation. The growth will start in engineering, science, and management disciplines where nuances are relatively easy for machines to understand and the corpus of existing knowledge that machines can use is larger. Google Translate has evolved over the last five years, growing from very inaccurate and grammatically primitive, to become accurate enough to convey meaning in many situations. In many respects, its evolution rests on an exercise of mining collective intelligence, as noted by Richard Baldwin[258] "the EU Joint Research Center posted a dataset with human-translated sentences in twenty-two languages (it has over a billion words). Not to be outdone, the EU Parliament released a dataset with 1.3 billion paragraphs that had been translated into twenty-three EU languages. Another massive database, uploaded by the Canadian Parliament, has millions of paired, human-translated sentences from the parliamentarian debates. With data and the computer power to process it, Google translations improved more in a month than they had in the previous four years.” Today, it is good enough to power asynchronous workflow for back-office and other applications. Combine that with spoken-word natural language processing that already provides remarkably good real time transcriptions. Then possibly add intonation analysis for additional cues. Couple that with ubiquitous mobile connectivity and the fact that sounds are comparatively much lighter to transmit than video, and you have a good combination of technologies on the cusp of becoming production ready. The result? Synchronous communication between people who speak different languages isn’t a far-fetched, futuristic dream. Google, Timkettle and Waverly earbuds, as well as Skype real-time translation have debuted and shown potential despite their current limits. The difference between 90% and 98% accuracy isn’t 8%, it feels more like *multiple times* better. It is just a matter of time to get there, especially for standard, non-extravagant language. And for now, as noted elsewhere, some videoconferencing providers such as Zoom offer real-time translation services as an additional service.

And even if we won’t get fully usable synchronous language translation in the immediate term, language neutralization technologies can become a good-enough foundation for translation at scale, where content gets automatically translated and then human translators “finish the job”. Most translation agencies do that today with an army of contractors, but it is not impossible to think that a proper marketplace of gig worker translators could be created, helped by machines, and possibly including volunteers for translation of educational or socially useful work.

Language is only one of the elements that makes collaboration more effective, but the implications of the above could be very significant. Speakers of languages like Chinese and English, where the corpus of data is large and the language translation more accurate, will become more interoperable. Speakers of language with a vast corpus of language data for management, science and engineering applications (like Japanese, German, French, and Russian) will also benefit. With companies like Appen and Lionbridge continuing to use crowdsourcing to enhance translation algorithms for specialized uses, other languages may quickly become available for machine-assisted translations. Soon, Spanish, Italian, or Indonesian speakers could see their accented English recognized more accurately by translation apps, limiting the need for machine translation of their native languages.

All of these could generate very strong network effects. Larger and very specialized pools of talent might become quickly available irrespective of location. A more expansive knowledge base would be accessible to all. A new level playing field would benefit the Chinese who can now access more English-language resources. But the same would be true of English speakers (native or not) who would be able to access Chinese, Japanese, German, French and Russian knowledge more directly.

There may be some losers in the process. People who don’t speak any of those core languages well enough will be dramatically disadvantaged, trapped in an old world of delayed, selective, and manual translation that limits their access to fresh information and knowledge. It could prevent them from productively participating in broader collective intelligence networks.

Also, what would be the impact on publishing of these scenarios? Clearly, some publishers would find themselves disrupted, more so than when Google digitized the content of millions of books. That matters. A large amount of scientific knowledge is currently hidden by private publishing houses who often only allow fee-based access to the text. They would not take disruption lightly. But, just possibly, the upshot of infinite searchability and summarization could be that publishing moves closer to management consulting and advisory, and dynamic publishing could step much closer to research-reports production.

In any case, this sounds like a small price to pay in exchange for a world where the biblical curse of Babel is mitigated, and our collective brains can function exponentially better together.

Managing thoroughly remote people networks

Behind the fog of COVID-19's war, a new world beckons for knowledge workers. Despite all the personal and economic trauma, a big silver lining is emerging. Hundreds of millions of people are now experiencing some of it. I am not talking about confinement. I'm talking about the world where you and I can be with people we work with, without physically needing to be there.

It is a *phygital* world, a hybrid of physical and digital, where the time zone is my address and no other physical constraint matters. It is not just a better world of work. It is a better place. Let me describe it to you.

The daily commute doesn't exist anymore. It doesn't take weeks or months to get the right people to work on a problem together effectively. You can say hello to your children and spouse as they have breakfast. Most days, you can sit with them for dinner. I can walk around my suburban neighborhood near New York City and the air smells of the countryside. I can hear birds sing instead of being drowned by the noise of cars and airplanes overhead. People in Paris, London, Milan, Delhi, Shanghai, Tokyo enjoy the same. In that world, not only we are more efficient and effective, but we are more resilient and creative. And we are happier.

It is not too good to be true.

The way we move will change. One simple constraint has shaped our lives: for more than a century, our engineers have attempted to reduce the time it takes for our bodies to get somewhere. That was the design principle. It is this simple: from cars to airplanes to elevators, faster has always been better, almost no matter the cost, as our environmental footprint and energy geopolitics attest. These crazy times gives us a chance to think differently.

One classic article[259] tells a beautifully insightful story. Imagine you're building a particularly tall skyscraper — tall ones. Your clients complain that it takes too long to get to their expensive top floors. Challenge an engineer to solve it, and they would respond by tweaking the tech, and likely making the machinery a lot more complicated, expensive and energy-intensive. They save you 30 seconds on each for the ride to the top. Now ask a great designer to solve it. They would show up with several big, cheap mirrors bought at the nearby IKEA store, place them on the elevator's walls, and gleefully watch people lose track of time. For just $100, we get the same results. Beyond a point, physical speed doesn't matter. Only our mind's perception of it does.

We will continue to travel and commute for work. But we will do so to establish a human, personal connection with our colleagues and clients. We will not need to commute every day to do that. Instead, we will reserve time together in physical space to create personal bonds. That will happen through primarily social, not functional, activities: no slouching on chairs around a boardroom table, no whiteboards, no PowerPoints. Instead: a lunch, a coffee, a walk. Looking at someone in the eyes, genuinely shaking hands (yes, we will go back to shaking hands). For everything else, we can summon everyone we need in a virtual space in a matter of seconds.

Our vehicles will probably change in that world. Laws of physics don't overly collaborate, as elevator engineers know. Let cars and buses and airplanes slow down, be quieter, consume a fraction of the energy, and emit a fraction of the noise. But they will be fully connected, and we will not notice time go by because in them we can work comfortably or be entertained. And autonomous vehicles? In many cases, we won't even need them. To most of us, a bus is an autonomous vehicle (neither you or I likely drive it). We just need combinations of public transport to feel effortless, and there shall be apps for that. Indeed, the big idea isn't to make cars drive themselves, but not to need cars that often.

Our real estate will change, as our address is our time zone. Our commute will change. From every day, it will turn into maybe a couple of times a week. Teams will timeshare a central office, and a better-looking space. For those who don't have space at home, they will walk to a local shared office. We might be in the office for a few hours and we might still meet all together in person once a month, or once a year. But that will primarily be to *be together*, not to *do together*.

All of that will change where we can live: that is, where we decide to, not where we must. The main factor will not be the commute to the office but your kids' distance to school (it is still good that they meet face to face!). The cost of living will decrease as we back away from overpopulated places, with a reduction of congestion and infrastructure complexity of megacities. It will be a good impact.

Our equipment will change. At home, our lonely laptop screen will be substituted by three oversized screens. One of them shows people's heads, almost life-size, in the ubiquitous video conference. We will want the best bandwidth money can get. My webcam and my microphone will be between me and my screen so that everyone feels that I'm indeed looking at them in the eyes. (The effect is called "parallax," and it is a good engineering challenge. Apple promised that they would address it in an upcoming FaceTime release, and others will follow suit.) I will have a dedicated iPad and pencil as my virtual whiteboard. And if I want to, I could use one of those big screens to permanently show a "gallery view" of my entire direct team, as if I had a window on the entire office. I will have infinite virtual walls on which I can append an infinite number of sticky notes (the workhorse of problem-solving) together with any colleague anywhere, anytime.

Expensive? Today this might cost an extra $1,000 a year. Expect that to go down 20% next year. That's the cost of one 3-day day trip between New York and Boston, Paris and London, Munich and Berlin, Delhi and Bangalore, or Sydney and Melbourne! And the equipment will last a few years. That's an ROI of one month or less for many of us.

Our approaches to problem-solving will change. The way we think together will also evolve radically. We've known for a while that to create radical new ideas we need diverse people. For centuries that diversity was limited to the people close by, typically in large cities. What happens if we live in a *meta-city* where I can work with anyone from Japan to the US West Coast throughout my day? I wager that as soon as we master the ability to identify and collaborate with the right people at hyper-scale, connectivity will ignite a collective intelligence the like of which we've never seen before.

Our daily flow of work will change. People need people, and they need structure. We will have check-in routines every morning for a little virtual chit-chat. We'll be on a video half of our time just like we are in meetings today, except that our office background will look like an AI-generated Caribbean beach or some other expression of individual creativity. Office365, or its equivalent, will tell us which people in our network we have not talked to recently. Some apps will take care of making sure that we mingle with people outside our typical network, just like we would at the office canteen. We will use specialized apps to crowdsource ideas and debate at hyper-scale. We will have virtual happy hours, and virtual wine-and-cheese office parties. And those managers who worry that "out of sight is out of mind" should feel reassured: research shows that most home office workers actually work harder and are more productive, as long as they can collaborate well.

And the few who misbehave are very easy to identify. Virtual collaboration leaves behind digital breadcrumbs that are easy to detect. By the way, risk management of decentralized work will benefit from the same tools and techniques. In large groups, there's always the statistical possibility that someone makes mistakes or even deliberately goes rogue, and a remote environment may feel a little more unguarded. But cybersecurity specialists tell us that for a long time there has been no physical perimeter to guard.

Artificial intelligence's hype will come to fruition in this world, as AI allows us to quickly discover people we should connect with (watch out for ubiquitous applications of knowledge graph technology) and make those interactions seamless with whatever bandwidth I have (more here). Yes, AI server farms use a lot of energy, but it's a pittance compared to the energy required to shuffle around millions of tons of metal every day. And AI farms are amenable to be powered by renewable sources.

And finally, quality of life will change. We have known for some time that a happier and longer life doesn't depend on how much money you make. Research says that beyond about $75,000 per year, money no longer matters. What matters is how many meaningful people are in your life, and how often you can be with them. Older people typically live better and longer when they are close to their offspring, especially their grandchildren. The opposite is also true: young children benefit from the care and knowledge transfer that grandparents provide. And multigenerational neighborhoods are good for people in their prime years as well — as someone once said, if you want to learn something new talk every day to someone above 70 and below 10 years of age. Just one of many opportunities for diversity in the new world.

Managing teams, and managing self, in a more remote and distributed work environment

There are lots that technology can do for us in the new normal of work, in which more people work from anywhere, more of their time. But in the end, it is all about humans, and the trust they build with each other.

What “remote bosses” need to do (managing teams and organizations):

- Get better at goal setting. Manage people by objectives and outputs more, not watching inputs and breathing on people's neck.
- Build more meaning in the work you give. In-person physical presence promotes gregarious behaviors and helps create momentum. In the absence of in-person gatherings, it is more important that you spell out the “why” of your team’s job, what impact it generates - from the benefit to your employee’s families, to that of their communities, their teams, the company, their clients, or the environment. Connect those dots.
- Help people pace themselves. Create routines and rituals that help people pace their day, now that the physical office routine subsides.
- Build trust deliberately. To be trusted, you need to be authentic and have the good of your teams at heart. Manipulative bosses will likely suffer.
- Learn new methods. Learn how to manage creativity and complex problem-solving meetings remotely. Remote meetings are NOT just talking and slides - there are lots of tools (e.g., virtual whiteboards, use of polls, etc.) and methods (design thinking basics) that can make people very effective at creating new solutions.
- Evolve work politics. This is hard to fix, and it might be a good thing. The politics of sitting down for lunch and dinner is important but doesn’t translate well into Zoom calls. Try to catch up and being helpful to your network. And indeed, Machiavellian plotting may just become harder.

In the end, managers are the orchestrators of the collective intelligence of their teams. That requires connecting people well, giving them goals and incentives that work, ensuring they get the external input they need to avoid insularity, and create the right collaboration environment.

The employees' part (managing self):

- Overall, work on yourself. Remote work requires more self-direction, and possibly personal maturity. It is not for everyone, at least not naturally. But it can be learned if you focus on it. Avoiding multitasking, ability to connect with others to connect and recharge, taking rests without being sucked into distractions, etc.
- Commit to the plan. Show that you have understood what's needed. Minute meetings actions, and memorialize them on Slack, Teams, etc. - and stick to those and their timelines. There's nothing more frustrating than remote employees who make project planning hard.
- Be transparent. Be extra transparent with the use of your time. If you think that being far away means that you can hide, think again: people may end up not trusting you and that is bad, bad news. For instance, do what you've promised to do (previous point) or make the professional part of your calendar publicly visible to your team, to incontrovertibly signal that you're confident about how your time management and commitment.
- Be direct. Don't fudge stuff during synchronous meetings. Bosses don't have much time to ferret the truth out of conversations, and they need to walk out of the meeting with a sense of clarity. This includes - answer direct questions directly.
- Learn how to do it. Don't assume that "what got you here will get you there". Do learn the techniques to help to problem-solve remotely. If all you can do is slides and talk, your professional life may suffer.
- Be flexible yet protective of your rhythm. Yes, there will be odd hours here and there. Take advantage of the flexibility and lack of commute. But also, don't always allow your work to occupy your mind. Ring-fence time and consider mindfulness and meditation.

Failing this, careers will likely be in peril.

In a nutshell: employees shouldn't try to "hide". In a virtual world, where you can't just "show face", your reputation is all you have, and that depends on the trust you build. And indeed, it is also your responsibility to figure out how to be easy to work with when you're remote.

A blueprint for hybrid-remote work

One of the reasons that make senior executives fret about partially remote (hybrid) work, is that many of the established ways of working (leading, collaborating, etc.) won't cut it anymore. We aren't just talking about Zoom screen-shares and occasional Slack or Teams chat. It is about getting people, at scale, to embrace interaction possibilities that didn't exist in 2019. And most importantly, it is about using the new ways of working to enhance the employee experience which fosters innovation, culture, and ultimately strengthens engagement.

There is no established blueprint yet, and without the right approach, we may end up with a lot of unwieldy, complex change management. Ignoring the problem won't work either. A new equilibrium may not be reached organically: people with disproportionate in-person presence and hierarchical (or another convening) power will tend to naturally dominate the interactions and polarize their location. For others, the result may be the weakening of interpersonal ties and lower engagement, resulting in talent loss through attrition and reduced productivity, and less-diverse thinking. That's why a significant number of CEOs are ambivalent about the future of partially remote work.

How do we structure the challenge so we can direct resources to solve it? The following framework is based on what we know about the design of a supermind.

While perhaps daunting at first, the "hybrid-work cube" below is actually quite intuitive. It identifies eight distinct "cells" where new ways of working (e.g., norms and tools) need to evolve (read: project teams need to be deployed to storm and form.) There are three dimensions to the solution, and their combination is where intentional design can happen.

1- Time: Is a particular interaction (e.g., team meeting, townhall, or brainstorming) between people synchronous or asynchronous? For instance, in pre-virtual times, most interaction was synchronous - with exceptions such as memos and physical libraries. We now spend much more time in asynchronous interactions with a corpus of knowledge created (crowdsourced) through the organic work of large groups of people - think of Slack threads or SharePoint repositories, and Microsoft Viva in the future.

2- Place: Does the interaction happen in a physical facility or a virtual one? Or maybe partially in both, which I call "phygital" - for instance, when some people are in-person, and others remote; or when the same person constantly moves from a physical space to a virtual one?

3- Type of connection: This is something that most people don't think about and yet interactions depend on it. Strong ties are often those between people connected in day-to-day, functional work; weak ties relate to more serendipitous encounters across a broader network (e.g., water-cooler) and are crucial for innovation and culture formation. Organizational network analysis (ONA) reveals that both are essential. More recent evidence points out that strong ties tend to be resilient to shocks like a lockdown and sudden change of location, but weak ties may shrink. That's another reason why many CEOs fear remote work.

On any given day, people will cross many of the cells across these three dimensions. Designing the right experience within each cell, and ideally also in the handoffs between them, is the job at hand.

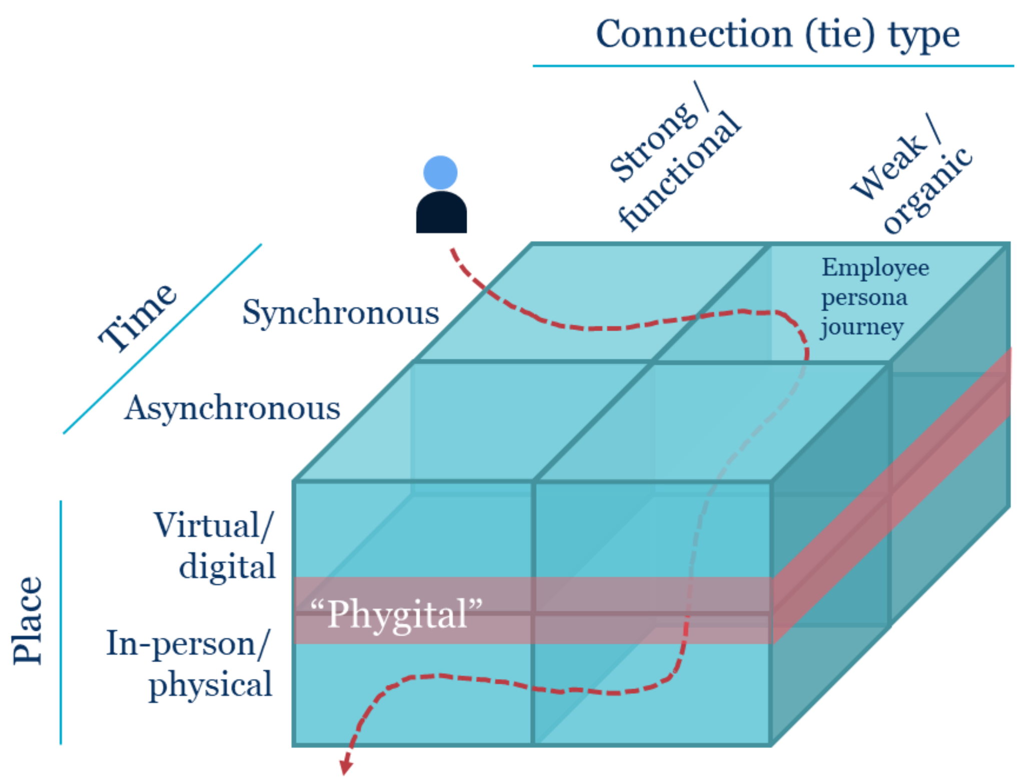

DEFINITIONS: Strong ties: day to day, functional work; Weak ties: serendipitous encounters across a broader network (e.g., water-cooler); Synchronous: video calls; Asynchronous: Teams / Slack threads, Knowledge Management; Virtual: interaction happening virtually, even if the person is in an office; Physical: interaction in an office (e.g., canteen, meeting rooms); Phygital: hybrid (e.g., some members virtual)

What does this mean in practice? At the very least, this framework should serve as a checklist to identify possible blind spots. But ideally, design teams across employee experience, IT, and lines of business should use it to blueprint and then drive the adoption of new solutions: that is the new tools, and the social norms that will propagate their use. Here are some ideas of design interventions based on the framework:

- Strong-tie, in-person: better traditional offices meeting spaces, based on new office frequentation patterns (likely more space for those, as more of the individual work can be done remotely)
- Weak-tie, in-person: traditional recreational space (canteen, water-cooler, corridors, atria – synchronous; bulletin board - asynchronous)
- Strong-tie, virtual: standing Zoom video calls (synchronous), Microsoft Teams channel thread (asynchronous)
- Weak-tie, virtual: intelligent matchmaking (synchronous), Slack affinity groups (asynchronous)
- Strong-tie, phygital: (synchronous) meetings where some members are remote, e.g., "one dials-in, all dial-in"; or standards for documenting in-person meetings for those who couldn’t participate (asynchronous) e.g., Zoom / Teams video recording and minutes of conference-room meeting
- Weak-tie, phygital: employee townhalls for big announcements; multi-location conferences; “virtual windows” into other offices' shared facilities. All of them should avoid the old-style one-directional download. This is mostly uncharted territory. In a world where distance is no longer a physical constraint, can one bounce into a remote colleague on the way to the corporate-office canteen (perhaps at the hand of an algorithm)?

If some of this sounds unlikely, let's remember that in March 2020 most people would have said that remote work at scale is functionally impossible. What is unlikely is that we "stumble into" a scaled-up solution to this challenge. It is time to structure the work and get on with it.

Does remote work damage innovation?

There's a heated debate on the impact of "work from anywhere" on innovation. CEOs are piling on it with their own (varied) views, and lots have been written about it – including materials referenced in the four pillars sections of this guide.

One side states that innovation doesn't work well with distributed workforces. To varying degrees, many CEOs, from Microsoft to Netflix to Goldman Sachs, fret over the long-term implications of remote work for the vibrancy of their workplace. Remote workshops are often clunky, and people end up spending time figuring out the tools instead of focusing on their fleeting genius. The face-to-face connection, the physical, helps to get people "in the zone" which promotes positive ideation dynamics. Serendipity of encounters that lead to new ideas, the fleeting watercooler or canteen interaction, are hard to replicate virtually. A recent Microsoft Research260 paper clearly states that remote work, without deliberate interventions, can damage innovation because the weak ties between people dwindle and people don’t use the right tools to collaborate.

On the other hand, there is reason to believe the opposite could also be true. Work-from-anywhere employees self-report261 higher levels of innovativeness compared to those who can't access remote work. And working more flexibly, without being bound by physical location, helps to connect more readily with the right people - not just the people who are in the vicinity or those who can get to that location in time. A good analysis of academic research262, for instance, shows that productive academic collaborations require(d) the first meeting in person, but their success was not dependent on co-location later. Big-I innovation (and invention) requires the collective mastery of a larger, cross-disciplinary body of knowledge, which means enough people with diverse backgrounds. For most companies, there's a real tradeoff between physical location and the availability of the right talent - the right people often work elsewhere...even if they work for you. And, by the way, just the fact that we *can* productively work virtually is in itself one of the most significant innovations of the last decades.

How do we reconcile those positions before they turn into dogma? A few thoughts:

Being in person comes more naturally to most people. The brain likes full-size, three-dimensional people, being able to use the sense of smell, and the oxytocin that comes from being in person.

Some things are best done in person. For instance, manipulation of physical matter - from life sciences labs to apparel - still requires co-location. Also, building trust and team energy benefits from (quality) real-life person contact.

If you have the right people nearby, especially for mission-critical punctuations, do strong workshops. The best ones can be truly cathartic experiences for people and organizations, both cognitively and emotionally. If you do fly people in, make it count.

Use digital technology to curate the right input and connect, collaborate and energize the network- irrespective of location. Novel ideas are typically the germination of thoughts through the combination of pre-existing ones and are triggered by people literally bumping into each other's ideas - not just bodies. And in today's world, the amount of new knowledge is forcing people to specialize, and teams to become bigger - and more interdisciplinary263. We badly need this broad collective intelligence. Let's use today's tech well, and let's realize it will get better (and our use of it will do so, too). Ideas are memes and should be helped travel through networks of people. Those networks may not be, and may not need to be, in person. For instance, networks of people and ideas are captured in knowledge graphs and can be made to connect. An example is below - office space creates strong ties and facilitates serendipitous connections, but also marginalizes others (those who are not physically there). Increasingly, AI-powered knowledge management (e.g., Microsoft Viva Topics) will create solid knowledge feeds that can fuel thoughts over time. And virtual watercoolers can help.

Innovation is a process[264], not a workshop. And asynchronous collaboration should be used better. We must harness the flow of ideas through *asynchronous* means, such as the creation of ongoing communities that document their discoveries in digital environments (e.g., Teams, Slack, or idea-generation workflows, for instance) which means idea-provoking conversations can continuously be overheard - upstream and downstream of workshops. The move to asynchronous collaboration is a megatrend that with the help of AI will radically change how we work. Clearly, asynchronous work can be a scourge too - for instance, too many emails don't generate creativity, and often delay the convergence to a common understanding that could be achieved in a quick synchronous meeting. The challenge is to get people to use the right method depending on the circumstance.

Watch the time(zone). Irrespective of location, you may need continuous synchronous interactions, especially in a tight-timeframe sprint, and you need people to work in somewhat compatible time zones.

Hybrid is "a thing" in innovation. Today, most interactions where some people are co-located and others distributed tend to be frustrating for the latter. That doesn't need to be destiny, and there's value in deliberately orchestrating effective workshops with smaller in-group people and others (e.g., subject matter experts) virtually there. Generally, an innovative hybrid workplace needs to be designed to be effective, or it won't happen naturally.

Train **people on updated innovation methods.** Machines are merely helpers there. Virtual sticky notes and whiteboards can be transformative, and so are other forms of collective input, especially as you can get large and truly diverse groups to contribute - but most people still struggle to master them. The learning curve will continue to flatten, as Zoom, MS Teams, Mural, Miro, Slido, Menti and a plethora of others cross-fertilize and compete. Yet, people won't wake up one fine day in their virtual-innovation-master avatar.

Using network analysis to understand the drivers of employee attrition and disengagement

The engagement gap. The new implicit contract. YOLO. All have a bearing on the “Great Resignation”. But here's a different perspective that leaders should urgently consider - with some very practical implications for the retention of senior professionals.

For some time, some companies have been using enterprise social network analysis to identify the reason for employee attrition. The foundational social-network science on this is established (MIT's Honest Signals and Social Physics, for instance). We know that if we want, we can often anticipate (and possibly prevent) people's resignations by looking at their network signature.

But is there a more proactive picture? Is there a "prevention-medicine" equivalent, instead of just curing the individual problem when it happens? There's ground to think so.

Two reasons why good people go

There are many reasons why people leave but many of them relate to these two questions:

Do they **feel they’re part of a (good) group**? Do their strong (e.g. manager) and weak (colleagues met at the watercooler) network ties give them a sense of emotional support, group "flow", as well as functional help?

Are they able to **make an impact** - a recognized and rewarded one - where they are? People's impact depends on which depends on being associated with the right things to do, getting access to knowledge and learning, doing work frictionlessly, and getting access to the right people to influence actions. The last one is especially important for senior people who typically work across organizational silos, and outside of fixed, rule-based workflows.

Interestingly, there's evidence that the answer to both questions depends on network structure, but let's focus on the "ability to make an impact" because that's where many of the surprises are.

Having **impact through a network** means that, at parity of individual capabilities, (a) you can access hard-to-find (e.g., undocumented or unindexed) information and you can do that because your direct and indirect networks have it and (b) armed with the resulting recommendations, you can influence how things are done by nudging directly and indirectly the right stakeholders at the right time. Specifically, getting (b) done in the case of unpopular or counterintuitive decisions is what makes companies strong, and managers successful.

(The spreading of useful but against-the-grain knowledge has been researched thoroughly: some call it "complex contagion": it is information that, unlike easy "memes", isn't believed and amplified by the network unless the recipient hears it from multiple, unrelated sources.
See here for instance).

The tipping point

Good people who fail to create a strong network in their current company often can't make enough impact, and consequently, they may have less to lose and could end up leaving. For them, *the skills that they have - not the network - are the biggest asset*. Those skills are marketable externally, and the loss of a company network doesn't penalize them much.

Think for instance about people whose networks are particularly strong outside of the company (e.g., some salespeople, senior consultants): they would also have less to lose when changing employers. That's especially true when remote work complements physical-office work, which makes moving elsewhere easier.

Contrast this situation with that of good people who have been in a company for long: their network’s ability to generate impact represents a higher proportion of their professional assets. And it’s not as portable: that is, they would lose a lot of that if they moved elsewhere.

That's what the simplified chart below illustrates. The hypothetical at-risk employee is the blue, larger icon in the left quadrant.

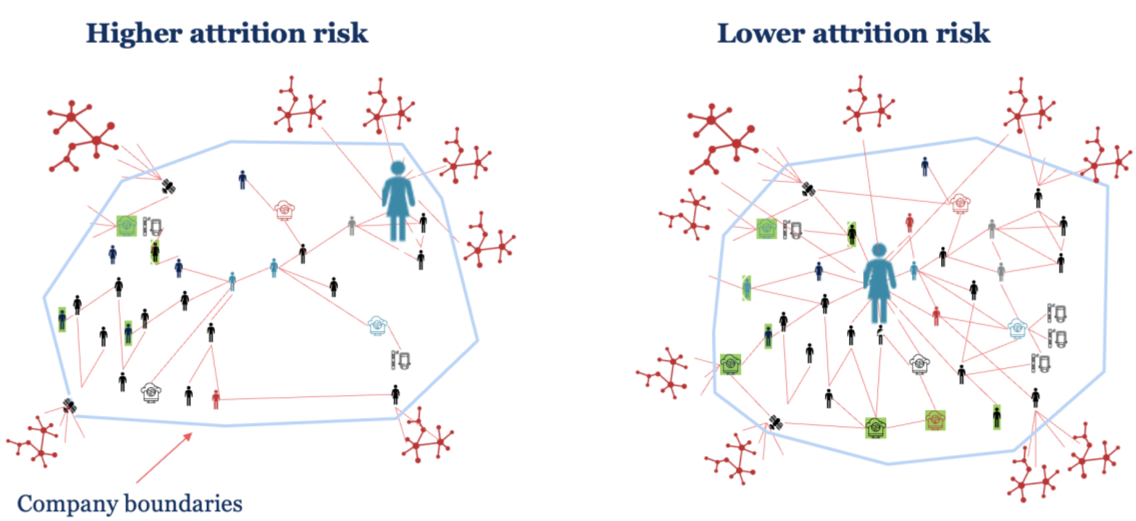

The real brain in the brain drain

Step back for a second, because there's an even bigger picture. What we are saying is that people use a "collective brain" to sense, remember, create, decide and learn. That's what generates the real impact.

It is a form of what some call a supermind: a cognitive engine constituted partially by that individual person's brain, but also by the "neural" network structure in their organization, or in their ecosystem. That brain is marketable internally, or externally - depending on where the network is comparatively strongest.

What can be done

A clear implication: in companies where knowledge is easy to retrieve, and where networks are easy to access, relatively new people typically become impactful faster, which should reduce the probability of their attrition.

Conversely, companies where the network is all-important yet tribal and knowledge isn't easy to retrieve will have an easier time retaining people who have been there for a long time: these employees' success is predicated on painstakingly building that network, which has now become the primary asset, and possibly a competitive moat against other employees.

So, while the creation of internal networks is important for long-term retention, it can become a hurdle for others, especially but not exclusively new people, and that will make them attrite. But if companies **invest in enabling people's "network-based impact"** and improve **knowledge access** (including learning and knowledge management), the **collaboration tools** (e.g. virtual whiteboards, asynchronous conversations), and the ability to **network effectively** (including serendipitous encounters, affinity groups), they might be able to retain more of their best employees. Individual managers are responsible to make some of this happen, and they may need specific sensitization and training. But CEOs need to invest in augmenting the digital infrastructure that caters to the broader system and making the related change management happen.

Harnessing the world supermind's "circadian rhythm" for innovation processes

For knowledge workers, distance now suddenly matters less than it used to. COVID-19 tore through years of timidity, and technology infrastructure held up remarkably well.

This is a brave new world where geography matters only if it draws time zones because the time zone might replace your office's address on your email signature and LinkedIn profile (forget the paper business cards). Time zones span from the North to the South Pole and can be thousands of miles wide (think of Central European Time that stretches from Spain to Poland, GMT+9 which covers the entirety of China, or US Eastern Time, that goes from the Great Lakes to the tip of Maine.) People in adjacent ones are only marginally farther away.

This development will profoundly transform complex knowledge work, where expertise and diversity of thinking are essential. We will get used to harness that talent everywhere - from East Asia to Europe, all the way to the US West Coast...if they're awake. And that's the twist in the story.

Where are the world's cognitive resources? The next chart simplifies the point - 24 time zones, 12 broad economic areas, their GDP, and their population - their talent. The world is flat. Or is it?

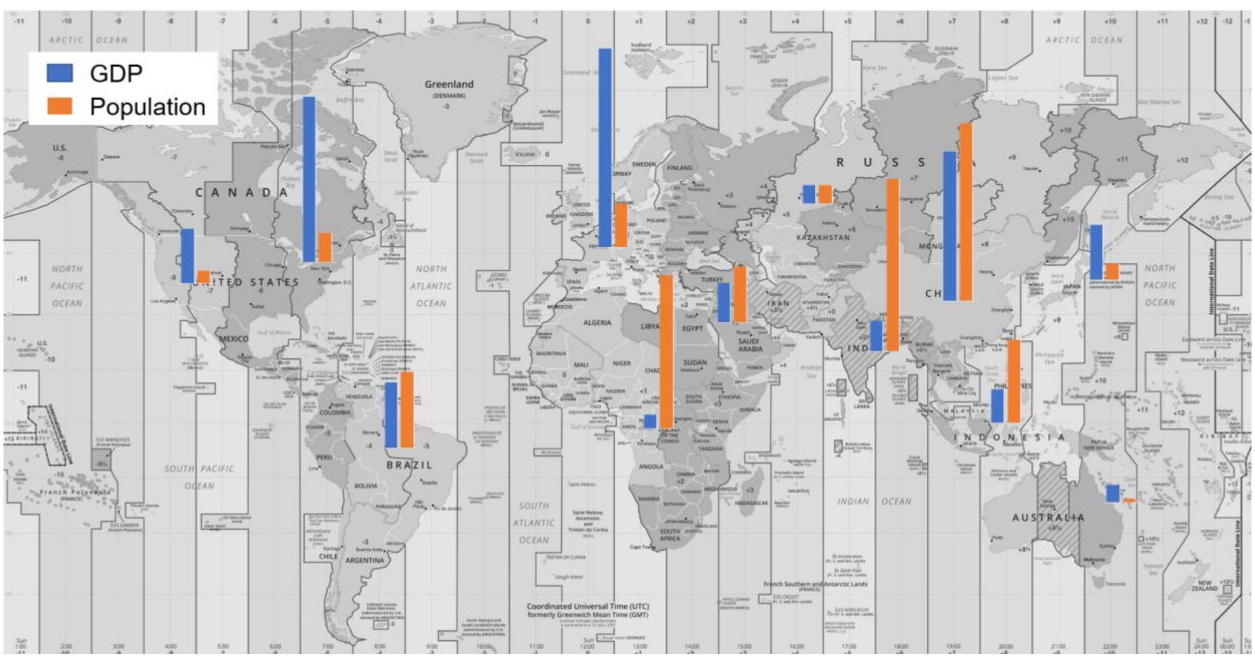

Much routine activity can be digitized, or happen asynchronously, through electronic messages and workflows. But for new, nontrivial solutions, we need the right people to engage each other through real-time, synchronous communications.

Anyone who runs global organizations knows that you can't dispense of the talent (and money) from Asia, Europe, and America, but their collaboration needs orchestrating across time zones. East and West relay information to each other, and co-create it, through calls, workshops, scrums, and the like. This is not trivial to do. In theory, orchestration work could get done anywhere. But there is one part of the world where it is easier. So much so, that management practices need to account for it.

So - when is the world's brain fully awake? Does a supermind have a circadian rhythm?

Fly from San Francisco to Tokyo, and endless water will parade below you. Continue from Tokyo to London, and you will only see land. Each hop crosses roughly the same amount of time zones - but the former doesn't fly over any major talent pool, or human activity, for over 5,000 miles: hardly a place for workshops.

What that means is that there is a time during a world's day when most people are offline - and that matters for processes that require synchronous communication. Let's do the math. I took population and GDP data; approximately broke them down by time zone; and estimated when they are available (I assumed a 12-hour window when people work during their day.) Then I plotted the available "global brain's processing power" (represented by how many people are up, and their respective GDP) against the cycle of a day. I used Greenwich Mean Time (GMT) for simplicity, but it turned out to be more than just a convenient choice.

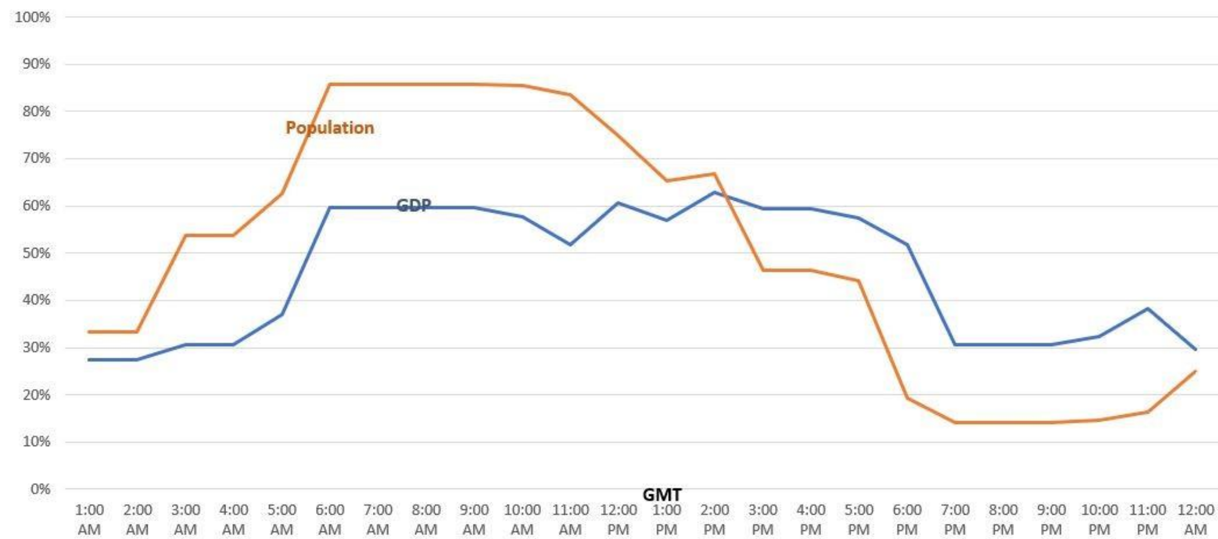

The peak of concurrent resources being "online" (that is, awake) falls between 7 a.m. and 6 p.m. GMT if we use GDP as a proxy for highly skilled knowledge workers, and a little earlier if we use general population. The former is more realistic for today’s advanced work. The latter is where the world is heading if Asia's workforces continue to develop as they are. Africa and Latin America might also contribute hundreds of millions more if they play their demographics smartly.

In a way, this isn't surprising. London investment bankers have long benefited from being able to trade during working hours with both Asian and American exchanges - harnessing the intelligence (on a good day) of all of those financial markets. But while trading has been virtual for a long time, the rest of the knowledge work is just becoming so.

The implications for organizational design, and evolutions of ways of working, are significant.

First, Europeans must get better at orchestrating between East and West. Europe is home to tens of millions of highly educated workers that can collaborate with the East in their morning and the West in their evening - and in doing so, increase the meaningful time overlap between them. It is a huge opportunity.

Second, American and Asian knowledge workers must become more adept at collaborating - and their leaders must make that possible. Today, a lot of that collaboration is at arm's length and often done through rigidly defined specifications. That's not what complex problem-solving, at scale, should be. The challenges we face, from epidemics to climate change, can't be addressed if Asia and America do not work together. But it is difficult to build trust between peoples from different cultures who can talk to each other only for a few hours every day.

New ways of working, and possibly that European brokerage and relay, can help generate trust and results. I described some of those methods earlier in this document.

Of course, the world might go in the opposite direction: regional ecosystems could end up mirroring fragmented spheres of political influence, where people mostly work with those in nearby time zones. Even virtual citadels. That world would lose part of the cognitive diversity and coordination necessary to solve our most complex problems. The fragmented response to COVID-19is a testament to what an East-West rift can do. Our brain uses all its parts for its most complex jobs. Our planet's collective intelligence can do the same, if we relay the signal across it.

Cutting climate tech Invention-to-Innovation time through collective intelligence

The cycle of invention (idea successfully prototyped) to innovation-at-scale (widespread implementation of the new practices) typically takes decades. For climate change, we just don't have that time. Yet, we can intentionally compress it with existing technologies and organizational design.

Every science- and ultimately technology-based revolution typically takes decades to percolate deeply into the world, because of the slow process of “learning by doing” - often starting with academia, some R&D pilots being successfully executed, then the most innovative managers adopting the new practices, and finally most others following - many years later.

This is innovation's death by thousands of small cuts. One of the most striking examples of that problem is the slow progress in evolving healthcare systems across the world, with its immense amount of variance in the deployment of tried and true practices (e.g., India's Aravind eye care process for cataract treatment, with its order-of-magnitude costs improvement at comparable quality, which a decade after scale, is still not adopted extensively worldwide).

We can't rely on established knowledge-transmission mechanisms

For climate change, the "typical cycle" is not nearly good enough. We can’t wait for five years until heat-pump installation capacity ramps up; we can’t wait for established regenerative agriculture practices for specific microclimates to spread to enough farmers who don't speak English well; we can’t wait for enough municipal utilities to learn how to incentivize and enable citizens for efficient energy usage; we can’t wait for a serendipitous uptake in the long tail of cities, regions, and countries that are not exposed to the most recent technology, methods, and applications.

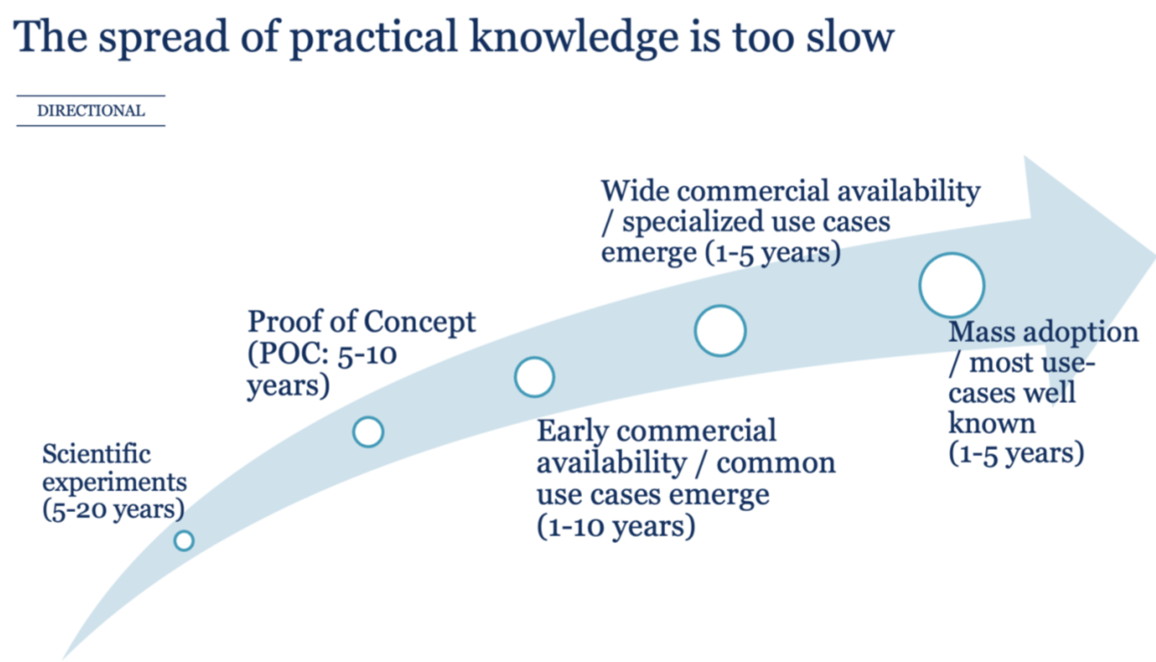

Many strictures exist: from upstream scientific to downstream practitioners access to knowledge repositories is not always as easy as it should be (academia and media); many professionals don't know how to thoroughly harness social media where new ideas surface; language barriers make it hard for the "global South", among others, to access and share new things. And the natural tendency of experts to silo their knowledge and try to find the next big thing in their field, whereas we know that innovation comes from the combination of existing ideas.

Industry and generally internet media are also not doing that job well enough. Thanks to algorithms tuned to maximize advertising and stickiness, meme-able noise often obfuscates the signal, and finding what's relevant is still too hard or expensive (e.g., paywalled content).

So despite the excitement about climate startup funding and corporate net-zero commitments, at least one aspect remains seriously neglected: **the intentional crystallization and sharing of practical, specialized, knowledge so it productively "touches the ground" and can be recombined with existing ideas, processes, operations etc**. That's a clear multiplier of impact but unsexy for many entrepreneurs and investors, and often left to either individual firms' marketing, or to well-intentioned but under-resourced NGOs and other public institutions -

including educational ones - that struggle with both granularity of information and speed of change. The outcome is a frequent reinvention of wheels.

We can do better today

This is not just about media or training. Both help, but in isolation, and when executed in a traditional manner, they have significant limitations. What works is a new organization for the *knowledge of networks of people, augmented by intelligent technology*: **Augmented Collective Intelligence**.

Today we have access to methods for knowledge formalization, retrieval, and sharing, vastly superior compared to the past. Google, Wikipedia, and the Web2 revolution (from WordPress blogs to Reddit, LinkedIn, Substack, Medium, etc.) have shown potential; yet they're not yet "finishing the job" of making relevant and practical climate-change information efficiently available to most relevant people. A minority of experts and practitioners know many information sources and can monitor them efficiently, but most others can't. That's significant leakage in the invention-to-innovation cycle. We can do better.

The table below summarizes the main idea. Hyperspecialized collective-intelligence "utilities" could accelerate the spread of high-momentum/low-signal *content* (both practical enablement and broader learning), and support the identification and engagement of relevant *people* (experts and practitioners). These infrastructures can use new natural language capabilities, and build *knowledge graphs* that facilitate two crucial processes: first, finding and combining granular information, i.e. the "what" (e.g., new ways of implementing heat pumps cost-effectively in areas where energy is expensive and unreliable); and second, pinpointing experts, i.e. the "who" (e.g., people or organizations who have codified the respective processes and can help on the ground).

The spread of practical knowledge can greatly accelerate

DIRECTIONAL

What	Who		Knowledge tools		Speed	
	Supply	Demand	From	To	From	To
Scientific experiments	Science labs / R&D	R&D	Scientific papers, some research portals	Addition of explicit knowledge graphs enables combinatorial innovation, identification of players	5-20 years	TBD
POC	Translational R&D, pioneers	R&D	R&D, science papers, whitepapers		5-10 years	TBD
Early commercial availability / common use cases emerge	Startups, specialized corporates	Innovation dept. of pioneering firms	Enterprise thought leadership; some training from suppliers; some press; specialized social networks; academic courses	Hyperspecialized "collective-intelligence utilities" accelerate the spread of high-momentum (and low-signal) knowledge (both enablement and learning), and more effective allocation of relevant experts and practitioners	1-10 years	50% less
Wide commercial availability / specialized use cases emerge	Scaled up new players, corporates	Early majority's operations dept.	Marketing; training (incl. enterprise and early vocational training); social networks; press		1-5 years	90% less
Mass adoption / most use-cases well known	Large new players, corporates, disruptors	Mainstream operations dept.	Marketing; commoditized training; vocational learning; social media	Explicit knowledge graphs facilitate identification and allocation of practitioners	1-5 years	TBD

Time

CORE IMPACT HERE

The uptake would be that the new granular, practically implementable knowledge could now reach not just the pioneers or the "hackers", but also mainstream professionals open to new ideas. That is the *early majority* of users.

There isn't a clearly defined category for this type of work. Its sits between social media, professional networks, education, training, open-source solutions, and even thought-leadership marketing. But the **building blocks already exist**. For instance, Microsoft has Viva, Linkedin, and Bing, which - combined - potentially have the full solution both within and *outside* of organizations. Others could use off-the-shelf tools that combine content and social media scrapers, perhaps using additional sources such as Google's open-source science and data repositories, or the amazing G-DELT machine-translated world news, or interesting new tools like Diffbot. Climate solutions startups like Ubuntoo (disclosure: I am an advisor there) already curate knowledge for innovation. Content providers, from scientific journals to Twitter, Reddit, and Quora, could make it easier to access rich APIs for this.

The sharing and combination of the world's relevant collective knowledge can be intentionally engineered thanks to new digital technology and practices. We could soon live in a world where detailed, specialized *"how-to"* knowledge for climate mitigation and adaptation is available on a browser that millions of people can readily access. Then, a broader base of people will have a fighting chance to tackle the largest challenge humanity has ever faced.

Applying collective intelligence to environmental innovation

A common misconception is that most innovation is invention and new technology. That's the bias in many fields, including climate-change technology solutions. And it leads to blind spots and insufficient investment in critical resources.

The world runs on an operational stack

Entrepreneurs and their net-new tech are essential. But impactful innovation is typically not just about that: it is also - and crucially - about their **embedment into the fabric of work, at scale**. That requires the evolution of the "legacy stack", i.e., all layers of activities related to tech infrastructure, data, applications, processes, policies, and people (and their ways of working). In the absence of that change, the power of new technology will not translate into real-world impact.

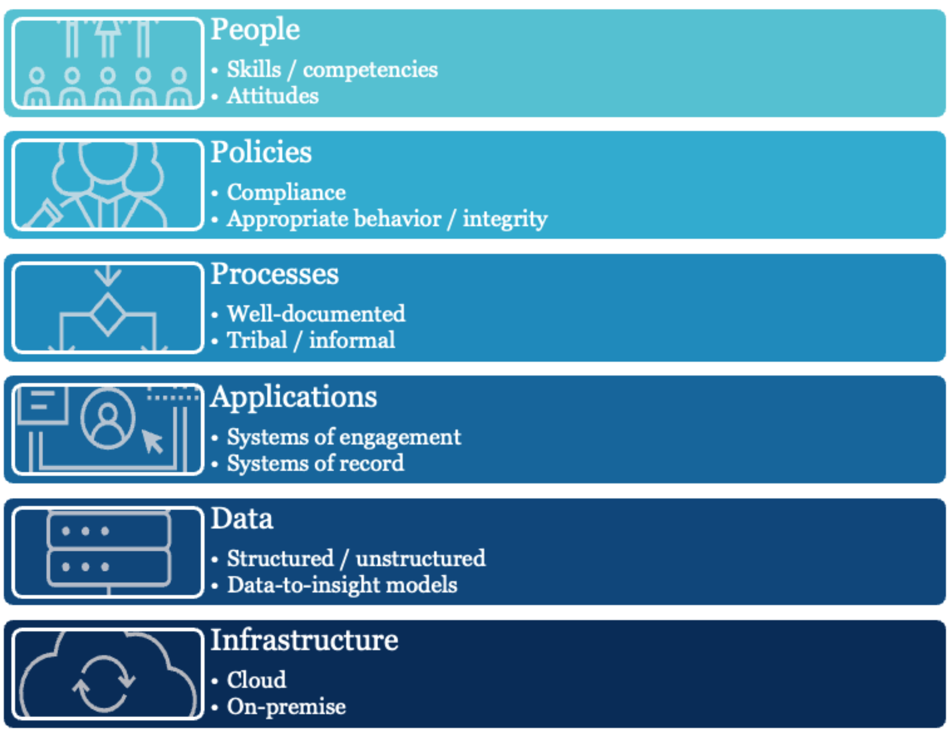

Resources need to be invested in addressing the strictures *across* the stack, not just at the technology level, and in doing so at scale to shorten the typically lengthy cycle between inception and mature adoption.

Learning the lessons of digital transformation

The slower-than-expected impact of AI, modern data science, and in general "digital" on enterprise effectiveness show that excitement, money, and technical genius struggle to overcome the massive inertia of how work is *currently* done. In another example, the global rollout of the Hepatitis B vaccine took more than 17 years - from regulatory hurdles to, quite simply, change management through healthcare systems. McKinsey estimates that in healthcare 30-40% of the disease burden, especially in terms of reduced quality of health in the final years of life, could be eliminated by looking beyond the healthcare technology itself, and instead focusing more on other factors such as mental, spiritual, and social health as well as prevention. These levers require processes, and people, more than just technology.

While we tend to ascribe hero status to scientists and especially entrepreneurs, we often neglect the critical role of operations in the innovation cycle.

Let’s take an extreme and well-known yet widely misunderstood case: we wish that there were more “iPhone moments" where entire categories are reimagined through a superior innovative *product*. Transportation, energy, food, you name it - the next Steve Jobs might emerge there. That should be certainly possible at least for a telephone, which is a product, one would think: you buy it, you update it, you trade it back or throw it away. That kind of product would be the most obvious place for new technology (or novel recombination of existing technologies, with a world-changing design) to be totally transformative.

But even the iPhone needed a lot of non-tech-product efforts: from the coordination with telecom carriers and their byzantine requirements, all the way to the game-changing AppStore which is a marketplace with all its (clever) back-office processing before it is a technology. Steve Jobs was initially lukewarm on the new device because he wanted to retain control of everything and rely on his startup methods - he wasn’t fond of engaging ecosystem players like the telecom companies, and wasn’t keen on the AppStore's operations and its community of developers. The messiness of those worlds disturbed his inspired design. In the end, the world changed because he let go of that qualm (also because of pressure from his teams) and intuited that the combination of those practices, and skills, was the future.

And - the world of climate change isn't remotely as product-centric. Rarely does the equivalent of an AppStore exist, and so does the option to plug and play your new tech effortlessly. The mosaic of economic and organizational structures is an unruly entanglement of legacy stacks, sedimented over decades.

The way to evolve them is to work on operations. One of the most successful climate stories, the emergence of electric vehicles largely at the hand of Tesla, was facilitated by a radical reimagination of sales and support operations - and not just an obviously superior product. Tesla did away with dealerships, whose incentives and competencies biased them in favor of combustion engines, and invested in a seamless experience - including an online purchase experience supported by showrooms, and delivery to the buyer's home.

Innovations for climate are large, tech-enabled *operations*

One can intuit the portability of this example. An analysis based on data from the seminal Drawdown.org work on climate innovation will illustrate the point further. To simplify, let's categorize the quantitative impact (tonnes of emissions abated) of each climate innovation very roughly, into three broad categories.

- Category 1 (blue): **engineering jobs**. E.g., wind turbines, carbon sequestration, and nuclear fusion - and products derived from new product design (e.g., new lighting, electric cars)
- Category 2 (yellow): **industry-based, or public-policy-based** change of practices. E.g., landmass use (agriculture, forestry, etc.), public transport
- Category 3 (green): **social change** e.g., girls' education, family planning

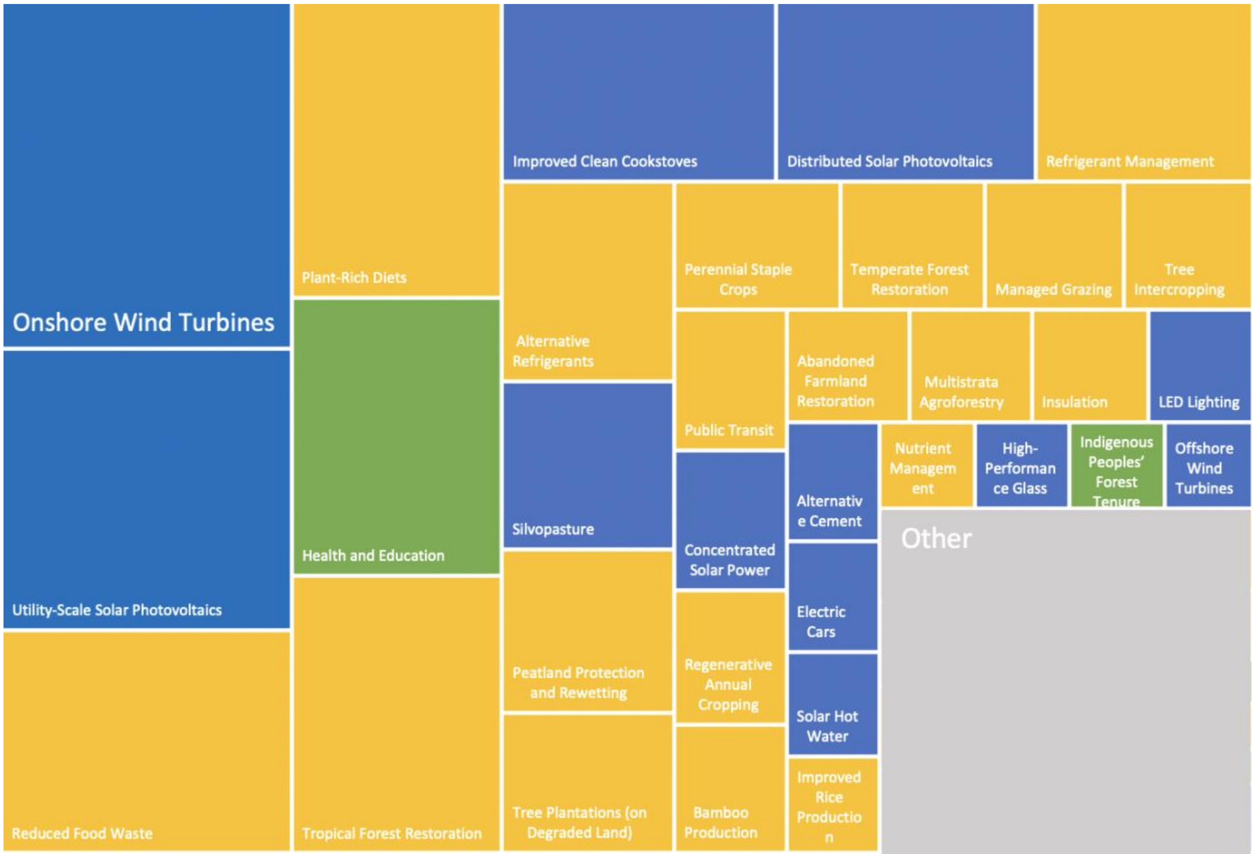

In *all* of them, the processes and the ways of working - not just the technology - will need to change. Even in category 1, where arguably a lot of the value is embedded in the technology asset, innovation requires a substantial element of process operation - from supply chain to maintenance, from financing to compliance, and from customer support to training.

- Environmental regulations' **compliance** is needed for value chains to change, including new areas like carbon offsets. They will look a lot like finance processes, especially finance processes of the future that use blockchain and advanced data capabilities (see Digital Gaia for example). And other critical areas, such as obtaining permissions, will look like large-infrastructure approval processes.
- **Supply chain optimization** for the environment, including for instance logistics and procurement, means adding explicit environmental variables (say, CO_2 intensity and cap-and-trade offset management), but they still rely on supply chain optimization processes.
- **Industrial asset maintenance** and optimization will need to enable the reduction of use, the decommissioning or mothballing of old equipment, as well as the conversion and

installation of new capabilities, and the engineering for maintainability. A lot of those processes are reasonably portable from old assets to the new (for instance, Terra.do helps oil & gas professionals retrain to cater to green economy assets).

- Most new environment-focused **financial services** such as banking (e.g. project finance, funds) and insurance (e.g., specialty insurance line that insures coastline properties) products will continue to look, at least to an extent, like financial products. Certainly, they will be powered by new data sources, modeled in different ways, possibly sold, underwritten, and their claims managed in different ways - but banking and insurance methods will be at their core.

At least in the early decades of the transition, **the legacy of financial and industrial structures will be overlaid and complemented by environmentally-sound practices.** We aren't going to "do over" extremely large and complex structures overnight, as they have sedimented over decades and trillions of dollars have cemented them where they are.

One more thing: People's behaviors, particularly those reflected in our ways of working, are steeped in our neural endowment and **are possibly the deepest legacy** of the entire economic system. There needs to be an inspired **design of experience for today's users** - from farmers to school teachers, from underwriters to mechanics, making the new ways of working desirable will be a crucial part of adoption. That's where better digital user interfaces, use of new visualization tools (AR, VR), and user of collaboration platforms will be crucially designed as *part* of modern processes - and not just plugged in.

Develop and combine capabilities, and embed them into the operational fabric of work

Another necessary yet vastly neglected condition for an invention to become innovation at scale is **skills formation**. The next chart illustrates the required combination of skills needed to (a) **design**, (b) **build** and, crucially (c) *run* the new environmental solutions. Climate innovation doesn't fully happen until the last step is done, at scale, everywhere on the planet. What capabilities are needed, where, and how to get that done at scale?

Domain expertise is foundational - say, understanding how a wind turbine works and its energy grid implication, or an intimate comprehension of the drivers of deforestation. It is essential that **technologists** are able to work with domain experts to design those assets (say, materials science, or synthetic biology), and support them with digital tools (e.g., cloud architecture and product engineering) that make it easier to operate and smarter - especially when combined with machine learning that harnesses the enormous amount of data that new sensors provide (say, weather stations, or satellite imagery). And finally, it is also important that those technology solutions are designed with the help of **experience** designers, design thinkers, and **process** designers - to ensure that those solutions work around humans, for instance, to limit NIMBY behavior or enable communities to counter the encroachment of unscrupulous developers.

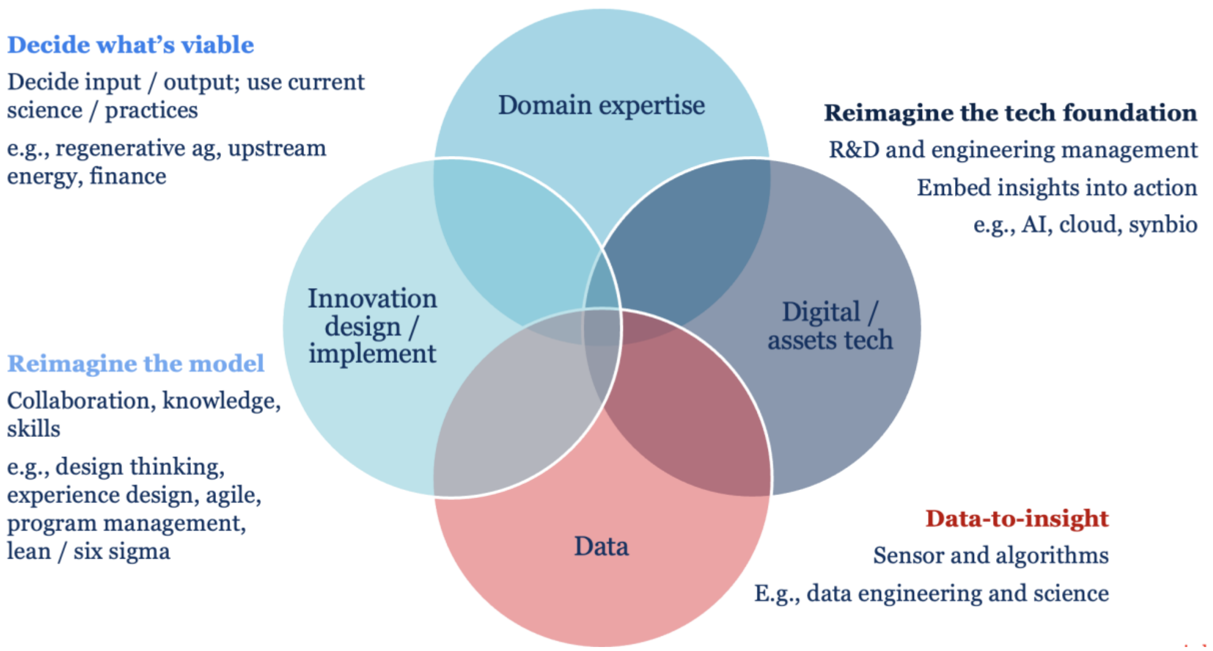

Climate and environment innovators will need to "get the village together" across four areas: domain, tech, data, human-centered experience, and attendant process design. Else they will become insular and ineffective, and delay the propagation of the impact.

Buttress learning operations

Many of the best and most novel solutions are typically born at the intersection of those domains. Environmental innovation will emerge from the entire system's collective intelligence, from the interplay in the networks of people that connect different areas of expertise. Those dynamics can yield novel and practical combinations that enable the adjustment and implementation of the net-new technology (more on this here).

To make that happen, people's skillsets need to increasingly “T-shape”. That is, they need to possess the narrow-and-deep stem of the T to master one area of expertise. But they also need the "range" afforded by the flat side of the T to go inch-deep into other fields in order to collaborate effectively with the respective experts. (More on reskilling for tech-led transformation here, and the implications on the climate space here).

None of this will be possible if we don’t **train tens of millions of people,** which is in itself massive-scale operations and won't happen fast enough if don't engineer the respective processes methodically.

Make no mistake: our educational and training systems, whether schools or other academies, will find it hard to address this challenge: they are not built for this level of granularity, interoperability, and for this speed of change. The lines between training, media, and knowledge management will need to blur if we are to shorten the cycle between invention and scaled-up innovation (more on this here and here). We saw the extent of this problem in the recent digital-tech waves.

World change = new tech + shared knowledge

None of the above is to belittle the value of radical tech shifts, and startups will do the world a great service by continuing to challenge incrementalism and foster the sense of urgency that large companies sometimes lack. However, glossing over the world's messiness might be fine only for some tech startups and their VCs when they focus on greenfield, narrow product-market-fit of ideally plug-and-play solutions, and the early part of the new innovation waves (that's where climate venture money goes). That’s when some of them may exit. But that falls short of solving messy environmental challenges, and their scale. And, like it or not, scale is often what incumbent large companies, established players - and their management, tend to do well. (More on this here).

Therein lies a conundrum: the two worlds of new-tech and scaled operations overlap somewhat, but they're not one and the same. Dangerously, more than a few people from each consider the others as “not fully fit for the job”, which is an act of self-preservation and personal positioning and significantly dilutes the impact we can create.

In conclusion, if done well, operations transformation is a core component of innovation success (more on this here). For that to happen, we need the collective intelligence of people *across* the operational stack. And to harness that collective intelligence, we will need infrastructures so that the right people, with the right knowledge, spurred by the right incentives, engage on the right collaboration platforms - and get the work done. That's what will drive scale *and* speed in climate tech adoption.

Society's collective intelligence is emotional: leaders, pay attention

Long-range scenario planning for enterprises' global operations is part of strategy efforts. And yet, many seem to miss something that matters, and that could directly or indirectly be influenced—the analysis drivers of society's stability.

The recent US election shows how important collective mood swings are in one of our largest decision-making mechanisms: democracy. Democracy is a collective-intelligence organizational model, a "supermind" as MIT Center for Collective Intelligence's Tom Malone calls it. And like other intelligent systems, it self-regulates, except when it doesn't. Just like the brain's cognition (thinking, decision-making) is directly linked to emotions, socially-shared emotions can likely sway society's choices. They can hijack our collective intelligence.

Among other things, this matters hugely for leaders exposed to the vagaries of policies driven by those mood swings, and also to those leaders whose organizations influence those swings.

To try and (very roughly) predict potential future scenarios, I analyzed some data series from the World Bank. The question was framed as "*Which economic and social indicators predict where things can go berserk*?" That is—when do social conditions empower parts of individual (and paleolithic) brains (e.g., amygdala) to take over; these conditions are poorly-controlled by our old institutions designed in times where information was scarcer and slower to propagate; and these conditions are triggered by "godlike" tech, such as potential social media diffusion of misinformation.

The hypothesis is this: all things being equal, societies' collective emotion is more stable when the following three parameters are kept in check.

GDP per head (not just total GDP) is sufficient and grows unconstrained by inflation- in absolute money terms, not as a percentage. And that's purchase-parity-adjusted, not the nominal GDP figure that's much more susceptible to inflationary distortion. Growth in how much goes into people's pockets, year over year, matters for obvious reasons. A shrinking population, if it is not accompanied by lower GDP per head, may not be destabilizing. And bursts of high inflation create fear, which likely makes people discount the value of their increased wealth.

Therefore inequality, as measured by the Gini coefficient for simplicity, should stay within limits. The growth of GDP doesn't mean that everyone benefits. I use this as a proxy for many other things, in particular access to opportunity and individual agency.

The last factor is particularly important when the above factors are in the wrong range. A society's expressed mood is more stable when the proportion of young males to adult males is lower. There are strata of society where the lack of progress can fuel violent revolt, and potentially even harbor the seed for terrorist organizations (link to a relevant blog). Young males are disproportionately more likely to feel the pinch of lack of agency, economic marginalization—and are more likely to revolt idealistically and violently. And thanks to today's technologies, they're able to find each other, amplify their moods and coordinate their actions more than ever before.

The above bear similarity to drivers of collective intelligence—identify nodes in the network, incentivize them, curate information, and provide a platform for collaboration—that I described in another article (here). This time, they double up as network-emotion drivers.

The model below has plenty of imperfections. I used GDP growth averages over ten years, because they may be a proxy for social infrastructure stock, but that may not be true everywhere, and I didn't look at family disposable income—and in some countries, GDP doesn't trickle into families' pockets. The momentum of growth may be as important as the average of 10 years. GDP PPP (purchase-parity prices) and Gini aren't awfully precise. I clearly do not account for idealism (including religious extremism) and the endemic presence of armed conflict which polarizes people even more and for a longer period. And in some countries, the data may just be wrong.

And yet—the results are worth a look.

I broke down the list of countries into three roughly equivalent groups, ordered based on their GDP per head growth over the last 10 years. But the heatmap is more important than the sequence itself.

First, 50 countries with most indicators in the right place, which should keep most of them clear of general unrest. Note that economic turmoil such as high inflation, or high inequality levels, especially when coupled with young demographics or with sustained fear of economic disruption, may still leave segments of the population at risk.

Country Name	2000-19 infl. Avg	2010-19 growth per head (avg) PPP	GDP per head PPP avg 10 ys	GDP/head/growth/$ 2010-19 PPP	Gini 2018	Young vs adult male ratio
Ireland	0.6%	8.2%	62,521	5,142	33%	34%
Singapore	1.7%	3.4%	87,757	2,949	40%	16%
Iceland	3.1%	4.7%	48,742	2,308	28%	31%
Panama	2.4%	8.9%	24,261	2,170	53%	42%
Lithuania	1.8%	7.4%	28,998	2,154	36%	25%
Switzerland	0.0%	3.3%	62,630	2,052	33%	23%
Malta	1.4%	5.7%	35,587	2,012	29%	22%
Estonia	2.3%	6.6%	29,917	1,979	33%	27%
Germany	1.3%	4.1%	47,865	1,972	31%	22%
Austria	1.9%	3.9%	50,280	1,941	30%	22%
United States	1.8%	3.3%	56,080	1,871	41%	29%
Denmark	1.2%	3.7%	50,001	1,865	27%	26%
Latvia	1.5%	6.9%	24,779	1,721	36%	27%
Romania	2.8%	7.4%	22,874	1,694	36%	24%
Czech Republic	1.7%	4.9%	33,930	1,662	26%	25%
Belgium	1.8%	3.5%	46,404	1,645	29%	27%
Netherlands	1.6%	3.2%	51,261	1,625	29%	25%
Australia	2.1%	3.4%	46,376	1,599	34%	30%
Sweden	1.1%	3.1%	48,763	1,532	28%	29%
Israel	1.1%	4.3%	35,412	1,532	41%	48%
France	1.1%	3.6%	41,536	1,499	32%	30%
Hong Kong SAR, China	3.3%	2.7%	56,054	1,499	47%	21%
Poland	1.6%	5.5%	26,942	1,492	33%	23%
New Zealand	1.6%	3.9%	37,629	1,452	32%	32%
Slovenia	1.2%	4.3%	32,823	1,410	25%	23%
United Kingdom	2.1%	3.3%	42,140	1,391	34%	28%
Hungary	2.5%	5.2%	26,719	1,381	30%	23%
Finland	1.3%	3.1%	43,964	1,364	28%	26%
Turkey	9.8%	5.3%	23,985	1,283	41%	37%
Canada	1.7%	2.8%	45,514	1,278	34%	24%
Korea, Rep.	1.7%	3.4%	37,105	1,276	32%	18%
Seychelles	2.4%	5.1%	24,504	1,243	47%	34%
Spain	1.2%	3.2%	35,572	1,150	35%	23%
Croatia	1.1%	4.7%	23,783	1,125	32%	23%
Bulgaria	1.6%	5.7%	18,767	1,076	37%	23%
Malaysia	2.1%	4.1%	24,799	1,025	44%	34%
Montenegro	1.8%	6.0%	17,050	1,017	40%	28%
Norway	2.1%	1.6%	63,498	1,013	27%	27%
Slovak Republic	1.5%	3.5%	29,235	1,011	27%	23%
Russian Federation	6.9%	4.0%	25,159	1,008	39%	29%
Italy	1.2%	2.6%	38,641	992	35%	21%
Portugal	1.2%	3.3%	30,234	991	36%	22%
Georgia	3.7%	8.4%	11,671	981	38%	33%
Kazakhstan	7.2%	4.0%	23,676	955	30%	48%
Japan	0.5%	2.4%	39,190	933	33%	21%
Mauritius	3.0%	4.8%	19,369	930	37%	25%
Costa Rica	3.2%	5.4%	16,449	889	49%	31%
China	2.6%	6.8%	12,851	879	41%	26%
Dominican Republic	3.7%	5.9%	14,604	868	48%	44%
Moldova	6.0%	8.7%	9,564	834	32%	23%

Second, countries where stability is still a work in progress, and where leaders (and their global partners) must maintain a close eye on some of these indicators. Note how Chile's recent violence eruptions might be explained by a high disparity of wealth distribution, despite the country being in aggregate the most affluent country in Latin America. Also look at India, whose economic indicators show that the path is still a work in progress—but also highlights how, on balance, less inequality and the creation of opportunity for young people thanks to a thriving services economy, have played a stabilizing role in the last thirty years.

Country Name	2000-19 infl. Avg	2010-19 growth per head (avg) PPP	GDP per head PPP avg 10 ys	GDP/head/growth/$ 2010-19 PPP	Gini 2018	Young vs adult male ratio
Chile	3.0%	3.7%	22,422	826	47%	29%
Cyprus	0.5%	2.4%	34,488	814	33%	25%
Maldives	0.9%	5.0%	16,152	805	37%	20%
Fiji	3.2%	6.9%	11,293	779	38%	45%
Armenia	3.4%	7.4%	10,479	771	32%	35%
Bosnia and Herzegovina	0.7%	6.0%	12,160	728	33%	22%
North Macedonia	1.4%	5.2%	13,943	726	37%	24%
Serbia	4.4%	4.5%	15,352	687	39%	24%
Thailand	1.6%	4.3%	16,008	684	39%	25%
Sri Lanka	5.2%	5.8%	11,278	655	39%	39%
Uruguay	8.0%	3.3%	19,900	650	43%	33%
Mongolia	8.0%	6.0%	10,617	638	33%	49%
Mexico	4.0%	3.3%	18,257	599	49%	42%
Ukraine	12.6%	5.5%	10,709	588	26%	25%
Bhutan	6.0%	6.3%	9,288	586	39%	35%
Botswana	4.8%	3.7%	15,882	586	60%	58%
Colombia	3.7%	4.3%	13,253	567	54%	34%
United Arab Emirates	1.5%	0.8%	68,017	565	33%	12%
Argentina	43.9%	2.7%	20,705	558	46%	40%
Albania	2.1%	4.6%	11,723	544	32%	26%
Belarus	19.8%	3.0%	18,119	535	28%	27%
Lao PDR	3.7%	8.8%	6,017	528	35%	52%
Vietnam	6.1%	7.9%	6,085	483	36%	35%
Bolivia	4.3%	6.5%	7,123	461	51%	50%
Indonesia	4.8%	4.2%	10,328	432	36%	39%
Peru	2.8%	3.6%	11,634	419	47%	38%
Philippines	3.0%	5.7%	7,237	414	46%	48%
St. Lucia	1.5%	2.9%	14,050	411	51%	26%
Paraguay	4.2%	3.1%	11,548	364	51%	46%
Ghana	11.6%	7.4%	4,681	345	43%	63%
Cote d'Ivoire	1.2%	8.9%	3,832	343	42%	75%
El Salvador	1.2%	4.5%	7,529	342	45%	46%
Greece	0.7%	1.2%	27,810	337	35%	22%
Ecuador	2.6%	3.0%	10,952	327	49%	44%
India	7.3%	5.8%	5,506	319	35%	40%
Egypt, Arab Rep.	12.9%	2.8%	11,002	312	31%	57%
Bangladesh	6.9%	8.2%	3,557	291	33%	41%
Kyrgyz Republic	5.4%	6.6%	4,299	282	30%	54%
West Bank and Gaza	1.6%	5.5%	5,020	275	35%	67%
Guatemala	4.0%	3.3%	7,866	258	52%	58%
Kenya	7.3%	7.7%	3,292	254	44%	68%
Jamaica	6.3%	2.8%	8,895	253	47%	36%
Honduras	4.6%	5.2%	4,764	248	54%	50%
Myanmar	6.1%	5.7%	4,285	246	34%	40%
Nicaragua	5.6%	3.9%	5,068	199	48%	48%
Namibia	5.2%	2.0%	9,786	198	61%	64%
Cabo Verde	1.2%	3.1%	6,247	191	47%	42%
Nepal	7.5%	7.2%	2,622	190	38%	53%
Papua New Guinea	5.2%	4.6%	3,795	176	42%	59%
Tonga	2.9%	3.1%	5,735	176	38%	63%
Morocco	1.2%	2.4%	7,016	167	40%	43%

And third, countries that clearly need all attention and help they can get. Many of them have systemic importance, be it because of their army (including nuclear capability in the case of

Pakistan) or because of their migratory flows, or simply because they may harbor extremism. In all of them, the threat of populism and its ability to amplify destructive tendencies is present.

Country Name	2000-19 infl. Avg	2010-19 growth per head (avg) PPP	GDP per head PPP avg 10 ys	GDP/head/growth/$ 2010-19 PPP	Gini 2018	Young vs adult male ratio
Samoa	1.7%	2.7%	5,950	158	40%	66%
Timor-Leste	5.1%	5.8%	2,720	157	31%	64%
Ethiopia	13.4%	9.4%	1,620	153	33%	73%
Sao Tome and Principe	8.6%	4.1%	3,577	146	40%	78%
Eswatini	5.7%	1.7%	8,556	146	53%	67%
Mauritania	3.7%	3.2%	4,591	145	37%	71%
Tajikistan	6.8%	4.7%	2,992	141	33%	64%
Cameroon	1.9%	4.0%	3,221	128	44%	78%
Sudan	32.2%	3.0%	4,208	125	35%	73%
Benin	1.3%	4.2%	2,838	119	43%	79%
Tunisia	4.9%	1.1%	10,451	114	37%	38%
Guinea	11.8%	5.5%	2,036	111	39%	86%
Pakistan	7.5%	2.6%	4,334	111	32%	59%
Comoros	1.8%	4.0%	2,741	110	51%	69%
Rwanda	4.2%	6.1%	1,771	109	47%	71%
Brazil	5.8%	0.7%	14,941	108	55%	31%
Jordan	2.8%	1.1%	9,493	100	34%	54%
Senegal	1.0%	3.2%	2,987	96	40%	84%
Venezuela, RB	79.3%	0.5%	17,449	89	50%	44%
Burkina Faso	0.7%	4.7%	1,821	86	39%	87%
Lesotho	4.9%	3.2%	2,699	86	48%	52%
Sierra Leone	9.7%	5.2%	1,560	81	37%	72%
Guinea-Bissau	1.9%	5.0%	1,621	81	43%	80%
Tanzania	7.2%	3.2%	2,363	75	39%	83%
Nigeria	11.8%	1.4%	5,182	75	39%	83%
Kiribati	-0.7%	3.6%	2,039	74	37%	63%
Solomon Islands	3.0%	3.3%	2,217	74	42%	73%
Gabon	2.2%	0.5%	15,089	72	40%	61%
Mali	1.0%	3.4%	2,034	70	37%	96%
Congo, Dem. Rep.	5.4%	6.8%	876	59	42%	92%
Togo	1.4%	4.2%	1,380	58	44%	74%
Liberia	10.3%	4.3%	1,322	56	35%	74%
Mozambique	7.6%	3.7%	1,183	44	49%	88%
Vanuatu	1.9%	1.4%	3,024	42	38%	70%
Lebanon	3.1%	0.3%	15,231	41	32%	39%
Angola	5.9%	0.6%	7,235	41	49%	93%
Micronesia, Fed. Sts.	1.9%	0.9%	3,390	32	41%	49%
Haiti	8.4%	1.8%	1,697	31	41%	55%
Niger	1.0%	2.1%	1,146	24	37%	107%
Uganda	6.2%	1.0%	2,109	22	43%	94%
Madagascar	7.3%	1.2%	1,581	20	43%	72%
Burundi	7.0%	1.9%	753	14	36%	89%
Malawi	16.4%	0.9%	1,056	9	43%	83%
Azerbaijan	5.3%	0.1%	15,330	8	28%	36%
Gambia, The	6.1%	-0.6%	2,170	(14)	42%	85%
Chad	1.5%	-1.1%	1,692	(19)	42%	93%
Yemen, Rep.	11.9%	-1.3%	3,708	(49)	36%	69%
South Sudan	89.5%	-7.5%	1,846	(139)	46%	77%
Algeria	4.6%	-1.1%	12,472	(141)	28%	49%
Congo, Rep.	2.3%	-3.8%	4,136	(158)	48%	75%
Iraq	2.4%	-1.3%	12,470	(159)	29%	66%
Iran, Islamic Rep.	20.3%	-2.0%	15,306	(302)	41%	37%

Can even the most powerful of us do anything about it? If I had just one wish—give young folks a real job. Apart from public institutions, that's something that CEOs and their workgroups can do, and many are already doing so. Think of the role of industry bodies like India's IT-services NASSCOM in promoting employability.

Second, and just possibly, see if you should apply the same model to your organization. Likely – swap GDP growth for salary growth, distribution gap (Gini), etc. While it may not be a democracy, it is possible that your employee engagement, attrition rates, and overall company vitality depend on similar dynamics. And if you're a financial investor, you may want to monitor that in your portfolio companies.

Collectively intelligent networks – a principle for all types of intelligence?

Having come so far, it may be appropriate to attempt another leap – admittedly more speculative and less grounded on observation, but perhaps able to inspire more and more-radical thinking.

I would like to explore a potential theory, maybe just a framework, for now, that explains the emergence of intelligence in any system with certain characteristics – and explains the intelligent behavior of many more things than just organizations. Highly speculative, but possibly interesting.

[This is not a scientific review of literature, but I will borrow liberally from theoretical work conducted by scientists and researchers across disciplines e.g., MIT Malone's work on collective intelligence[265], UCL's Friston's[266] Active Inference based on the Free Energy Principle (FEP), Anil Seth's work on the brain as a prediction machine[267], former MIT's Cesar Hidalgo on collective learning[268], Gupta, Taylor & Kaufmann[269] on modeling agent-based system optimization based on FEP, Jessica Flack[270] on collective computation in natural systems, Hawkins' "thousand brains" theory[271], etc. I make no promise that the below is an accurate representation of any of that work, and I will rather use it to inspire a radically new idea.]

So, here's the key question.

What if the most efficient and effective way to generate intelligence is to establish connections between nodes in a network, enable them to feed each other with simple but not random input and feedback, and store the results of what has worked best? Can that generate planet-scale intelligence, starting from its smallest constituents, without requiring any complex intentional design?

If that is true, it might just take enough natural experiments until the chemical connections lead to biological and then possibly neural structures of any type. All one would need is nodes that can respond to the impulses of other nodes in a network. There's no reason such a structure would be unique to the human brain. The nodes can be of varied nature – from neurons to individual brains to computer servers to entire companies and industry value chains to internet servers. And possibly the other way around, down to cells, DNA, other molecules, possibly even atoms. We will get back to this at the end, hinting at the scale of collective intelligence, from the smallest to the largest.

These structures are quite simple at their root-nodes, their individual predictions stored in some sort of memory, modulated feedback from increasingly sophisticated and deliberate actions, and potential change of behavior as a response to that experience.

Their permutation brought about by large-scale network interactions generate practically infinite outputs to practically infinite inputs. But these are not just random feedback loops. The better networks, those that manage to develop more and longer, do so because they develop a property of prediction.

It is no coincidence that some neuroscientists see the brain as a prediction machine, embodied in a sensory environment which also includes the feeling of one's one body as part of a prediction network (e.g., the guts' - of the bladder, skin, etc. reaction to a thought). Neuroscientist Friston's framework borrows from thermodynamics the concept of Free Energy Principle (FEP) and sees intelligent systems as able to encode predictions. Any system of that nature, including but not only the brain, has a prediction algorithm that creates a model of what's going to happen, and performs so-called active inferences to continuously refine the model based on actions and resulting experience (even simulated ones).

A (neural) system with these characteristics uses less energy and is more effective, hence surviving more, than alternatives. It doesn't require any more magic than machine learning (for a comparison, see here) – yes it is complex, and it needs lots of data, but it is far from unthinkable. This mechanism would also work in the absence of large brains, through natural selection and natural experiments (especially when living beings can move) once the right data storage exists (in nature's case, DNA either through genetic mutations or through epigenetics).

There is no reason why that would only apply to brains. Connections with feedback and reinforcement exist everywhere. Yes, humans do it especially well in some ways: we connect with other people through sensorial feedback, which is explained by the "theory of the mind" whereby we interpret each other's signals especially through looking at people in the face (particularly the eyes).

Let's look at the underlying architecture - to see if such a structure could emerge somewhat spontaneously from unexpected environments. All it might take is sensors outside of an entity, connecting with other nodes (people), and then nodes inside the entity with their ability to do active inference. Of course, this works better with language, which people excel at. But the basic architecture (and certainly scale) is hardly unique to humans.

To function, these systems must be able to do five things[272]: sense the environment, remember previous conditions and successful solutions, create new solutions, decide on the course of action, and learn (possibly, using a combination of the previous four). It doesn't need to be done the way humans do it, and most importantly it doesn’t need to be at an individual level. Hence the importance of collective intelligence - and any technology, including biological ones.

That architecture just needs some components, broadly akin to the four modules discussed in the Augmented Collective Intelligence work: connections, incentives, information feeders, and collaboration environments.

First, instructions – the software, and its algorithm – to create based starting from basic IF/THEN. That can also be evolutionary, and from the emergence of modern AI we now understand that complex programming of algorithms is by far not the most important characteristic of an intelligent system. So, for instance a DNA strand is code (and compact too).

Second, processing power (to create) and storage (to remember), e.g., the size/speed of the CPU. Its speed matters because in highly variable environments, or where new experiences can be made proactively (e.g., by trying new things deliberately) more data is more important than better algorithms. Or you just need time - of which Nature has plenty.

Third, crucially, the ability to connect with other nodes, and “get their take” on the five cognitive processes - in particular, when a common goal is defined. Connectivity helps sense what other nodes sense, as they connect with the external world, for instance through the embodiment of neurons into the peripheral nervous system connected to our senses – vision, hearing, skin; the internet, with its http protocol does that too, connecting servers; but also, in a society, communication tools and the rule of law regulating interactions and contracts, serve that purpose. Connectivity also helps tap into the memory of the network – synapses biochemically do that, and in many ways collective memory is a form of connectivity (for instance, neural pathways activation, but also industrial value chains and possibly others.

Why would something so simple also be powerful? Because of its ability to scale, that is the network’s ability to permute connections between nodes through a quadratic law – x nodes, can combine in x-squared ways one-to-one, but a lot more when you're considering all the other permutations between groups of nodes – that's factorial (n!), which grows even faster than an exponential e(x). That's a lot of computing power emanating from a simple architecture, when scaled.

So - where are these structures? A visual example first[273].

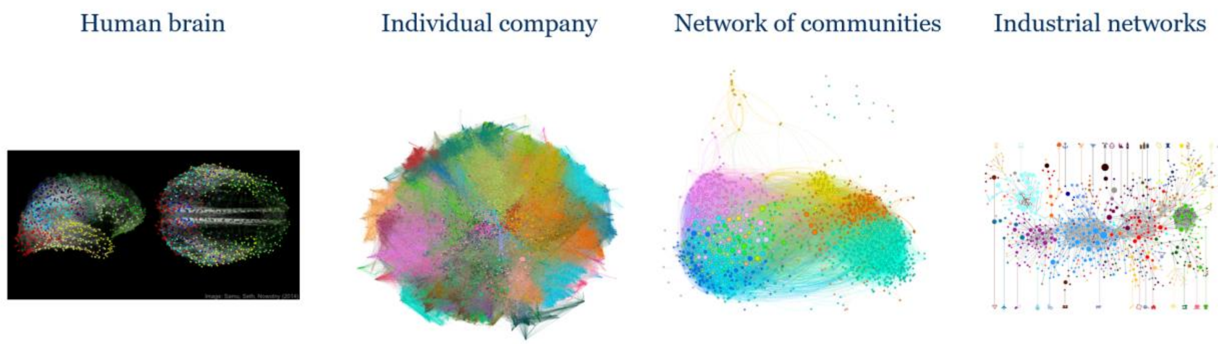

Whether or not each of them is accurate, I think the following list should at least provoke new thinking and debate.

- Molecular ties – some of them are stable and collectively "remember" sustainable structures. And they can evolve: chemical autopoiesis is a likely precursor to biology. In particular, mRNAs, microtubules, and particularly DNA is a precious example of that – with its ability to store designs (remember), but also constantly change them, creating new combinations whose best ones the ecosystem collectively learns through natural selection.
- Cell's interplay is the next level up – starting with proteins acting as information links between adjacent cells, then fungi, whose ability to create sophisticated networks in response to environmental conditions is all but random. Slime molds for instance can efficiently shape their growth to search for food, and to adapt their behavior to time[274]. So-called basal cognition has been studied for decades, and rich data is available about the interactions that help it emerge[275].
- Then plants – whose network connectivity is established through pollen and roots. Think of the resilience of Pando, the 106-acre aspen colony considered to be one of the oldest and largest living organisms on earth – and effectively one tree. Or the fact that flowers

seem to increase the amount of sugar in their nectar when they perceive the buzzing of pollinating insects (and not others)[276].

- Then animals, whose recombination of DNA through sexual reproduction enables natural experiments; and whose collective behavior helps their colonies: from bees and ants, from colonies of jellyfish to schools of fish, to mammals including heterogeneous ones like zebras and wildebeest migrating together.
- Then entire ecosystems, highly interdependent from each other and thought by some (Lovelace, Gaia theory) as becoming a planet-level mechanism that not only regulates life, but also the environment where life happens.
- Next up, the brain's neural networks – with the clever structure and connections of their regions, from the reptilian brain to the subcortex to the neocortex that's totally primed for network interaction. But also, the interplay within the neural network in Hawkins' "thousand brains" theory. And possibly, the fact that through dreams and memory we can connect and interact with our other or former selves. Additionally, there is mounting evidence that our gut microbiome (dependent on the environment we live in) influences our brain, including possibly steering it toward more social behavior[277], possibly adding another collaboration mechanism with external nodes.

But what about digitally-, often AI-augmented collective intelligence?

- Organizations are the aggregate of not just their people, intellectual property, and processes, but also their connections. The power of individual people can often be explained in those terms too, well beyond individuals' IQ.
- Think of media, starting with the invention of the printing press, that helped many more nodes (authors and readers) connect in a more organic and less hierarchical way – especially between people who shared a common transmission protocol i.e., a language.
- The internet, search mechanisms like google, social media networks, blockchain through DAOs enable interactions between network nodes in ways that are both organic and synthetic.

- Social media networks evolve properties of their own, often quite different from the individual contributors', shaping collective sense-making, political systems and markets, among others.
- Economies are in themselves network structures enshrined in business, social, and legal processes – where the interplay between nodes is now facilitated by web technologies. Think logistics chains, financial markets, for instance.
- Political structures are also power-oriented networks, and they're increasingly under pressure because of their inherent "narrowband" nature due to their often-hierarchical mechanisms.

(Note that neuroscientist Tononi and other researchers have even used network properties, expressed as Phi, to estimate the *consciousness* of systems[278] through so-called "integrated information theory", which is not directly covered here; the concept was also used by collective intelligence researchers to assess group interactions[279]).

Many of these systems are under pressure. In an ecosystem, stress is a key part of evolution, as it forces natural experiments. Humans may not have existed without an environmental upheaval that wiped out the earth's previous overlords. But adaptation is a balancing act and depends on speed and impact.

What happens to biological systems when parts of the ecosystem collapse due to unusually fast climatic change? What is the impact on our individual brains when we are being overstimulated by being connected nonstop to virtual networks? What happens to economic networks' intelligence when globalization goes into reverse? What happens to trust in societies when social media becomes "truth" for millions of people? What is the effect of cyberthreats in highly networked economic and military systems? What are government bodies and even nation-states going to become when business and ideas can truly happen globally and virtually?

That's where the intentional design of a supermind, an augmented collective intelligence, becomes important. We can build better systems – systems where collective intelligence emerges, if we deliberately focus on their network design.

Conclusion

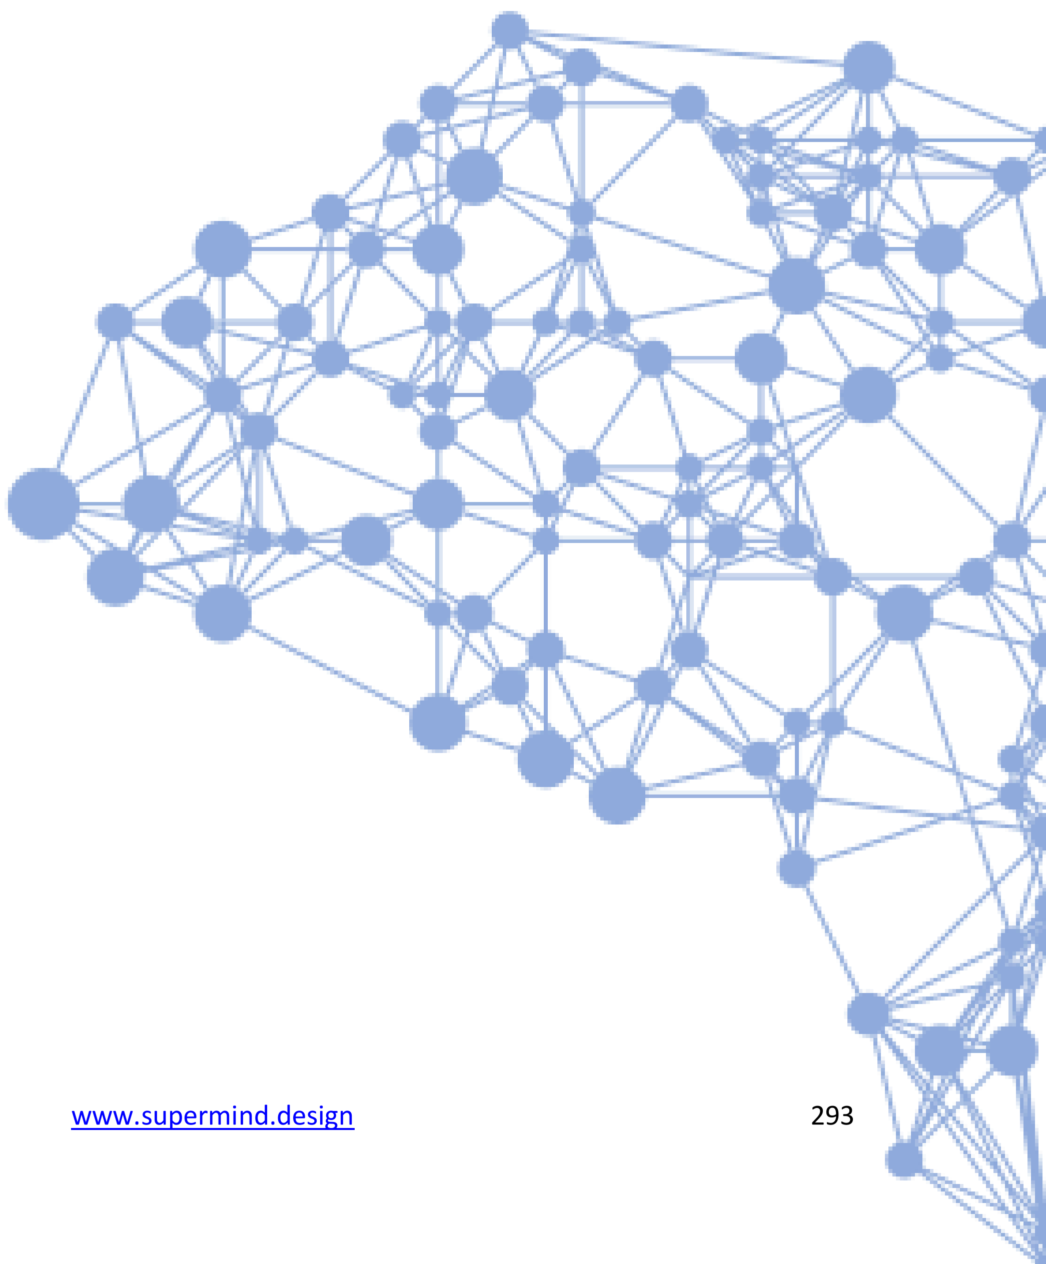

E Pluribus Unum

The intelligence of groups of people has powered some of the most spectacular economic successes in the history of the world, and more than successfully countered the reduction of individual brain's size that happened roughly at the time when modern civilizations emerged[280].

The Roman empire assimilated conquered peoples and harnessed their best minds through a system of administration based on better communication flows, facilitated by the world's best road network in the world and Latin, as then the new standard of communication for a large territory.

Ancient China's kingdom leveraged similar methods, until the Ming dynasty retrenched in an effort to contain the Mongols and severed most ties to the rest of the contemporary world[281].

The networks created by Venetian traders brought together East and West and gave birth to collective-intelligence financial innovations enabling allocation of resources across networks (including accounting, insurance, and risk management), all amazingly functional even in times of slow communications.

The expansion of the British empire, and its ability to invent and harness of world-changing steam power[282] was based on commerce, and related networks of people who could speak a common language.

The rise of the United States of America as the dominant economic power over the last hundred years was enabled by millions of diverse, comparatively inexpensive, and educated immigrants. These "huddled masses" deployed across a large and increasingly well-connected territory thanks to railroads and telegraph. That momentum was accelerated by New York City becoming the most important financial center in the world, where the information from market superminds was used by a very diverse workforce able to communicate seamlessly in English.

The rise of modern China has also been facilitated by large networks of people ruthlessly competing in the internal market. Notoriously open to adopting new things, having seen a 30-fold increase of standards of living over little more than a generation, the Chinese have been collaborating in Mandarin (and to a lesser extent in Cantonese) on multiple platforms and reaped the value of the country's critical mass. And they have simplified their language – especially in written form - to make it more accessible to more people.

Israel, born out of international treaties and grown through the repatriation of very diverse people with a common underlying background and a unifying language, has turned into an innovation powerhouse for possibly the same reasons.

What's most striking is that all of them adapted to constantly new circumstances and hence managed to last for centuries – a manifestation of intelligence.

All of them rested on significant alignment of incentives as well as the collaboration enabled by legal and trade platforms, with common language and information channels that facilitated the exchanges. These exchanges happened through networks enabled by organizational rules, commerce, and conquest. Collective intelligence has likely been a foundation for world power for a long time.

But most of these value-creation epics predate the full use of networks made more intelligent by AI-powered connection, curation, collaboration, and computation – AI's four Cs. What happens now that we can achieve critical mass *and* specialization *and* speed of connection at the same time, largely irrespective of physical location? Could that give us the much-needed smarts to help us deal with our challenges, and sustainably create a better world?

And conversely, what happens if we throw these trends into reverse? If we partition our collective brain for fear of too much connection? Or if we don't develop deliberately a management discipline out of the current muddled, incoherent sets of practices we use to harness the power of human-machine networks? Would we deprive ourselves of the ability to address Albert Einstein's concern, that no problem can be solved from the same level of intelligence that created it? Much of the current backlash against globalization could end up fragmenting our collective intelligence, at a time in which existential challenges such as climate change require our *best* collective thinking.

One last point, and a request.

This document is a starting point for what I hope will be a vast collaboration between like-minded yet diverse people. All feedback is invaluable to support the maturation of the discipline of managing connected-intelligence networks, superminds. There's arguably been too little "collective intelligence *on* collective intelligence" so far. This is an attempt to break that cycle. Our own network is incomplete, and we will need to create better information feeders and collaboration environments. None of this seems out of reach if each of us resist the urge to look like the smartest neuron in the brain.

Reach out, and let's wire up a better world.

SUPERMIND.DESIGN

References

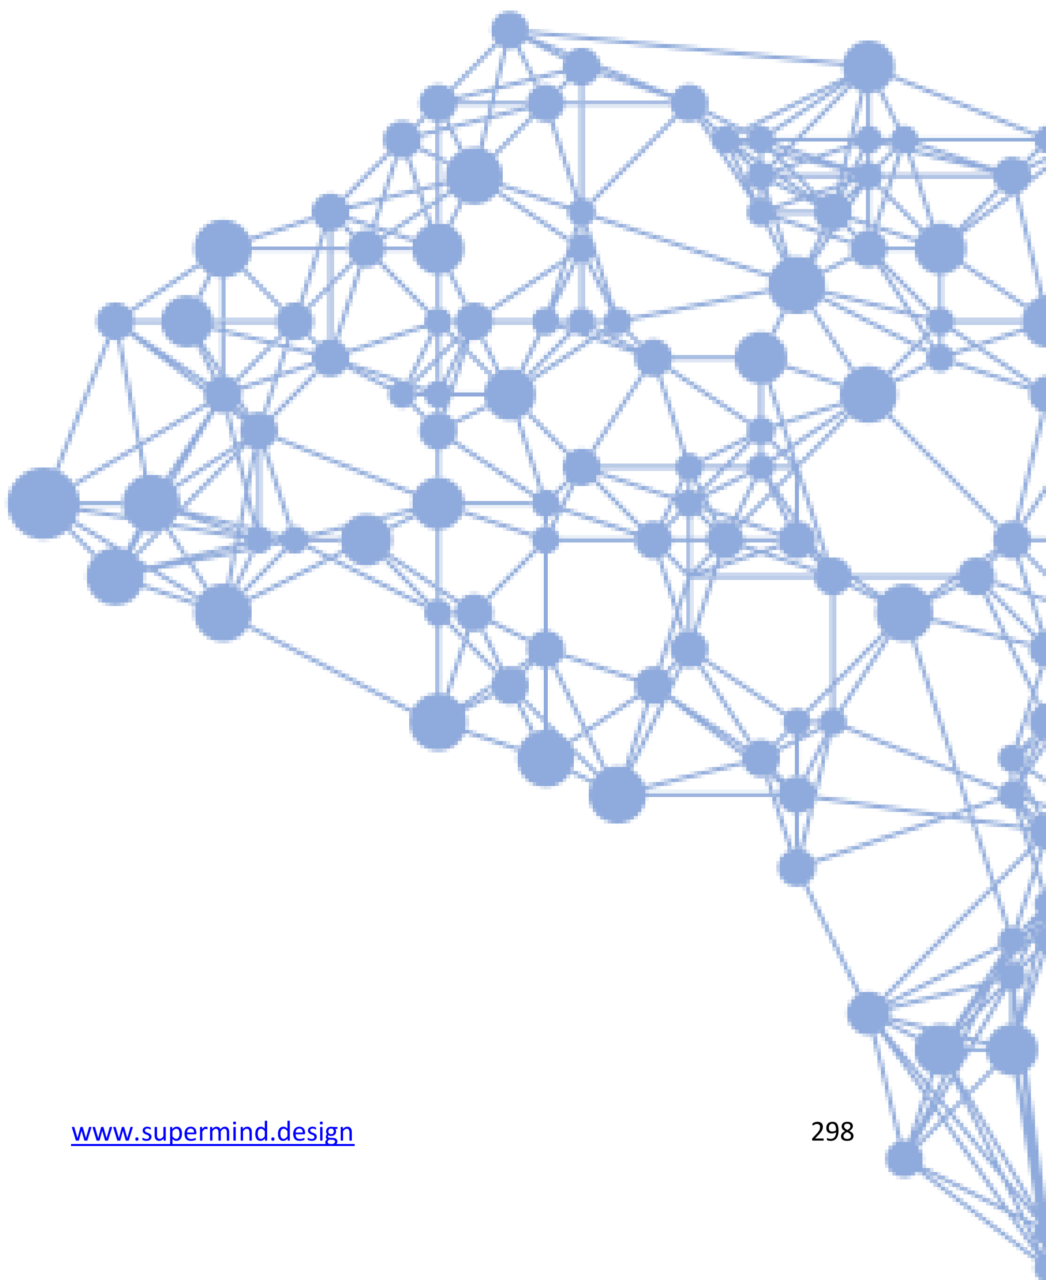

---

[1] "Exponential: How Accelerating Technology Is Leaving Us Behind and What to Do About", Azeem Azhar, 2021
[2] https://benevolent.ai/news/potential-treatment-for-COVID-19-identified-by-benevolentai-using-artificial-intelligence-enters-clinical-testing
[3] Geofencing is the practice of identifying the physical position of an individual or asset through digital means (for instance mobile phone GPS), and monitoring its position to detect trespassing and trigger additional action (such as notifying people at vicinity, or security forces).
[4] Fry, Cai, Zhang, Wagner "Consolidation in a Crisis: Patterns of International Collaboration in COVID-19 Research", 2020 https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3595455
[5] "How the COVID pandemic is changing global science collaborations", Nature, June 16th, 2021
[6] For instance, the Global Alliance for Genomics and Health
[7] www.Caiac.org. Quote from ZDNet https://www.zdnet.com/article/new-global-coalition-aims-to-help-policymakers-leverage-ai-against-covid-19/ "The Future Society is establishing CAIAC along with the Stanford Institute for Human-Centered Artificial Intelligence, with support from UNESCO (an agency within the UN) and the Patrick J. McGovern Foundation. The alliance will establish an advisory group that includes UNESCO and other UN entities, including UN Global Pulse (an initiative of the United Nations that attempts to "bring real-time monitoring and prediction to development and aid programs"). The group is working closely with private sector partners to get the platform off the ground, including C3.ai, stability.ai, Element AI, Axis, GLG, and Planet. Initially, the alliance is focusing on three uses cases for the platform, zeroing in on use cases that are especially relevant for multilateral organizations: Tracking and tracing of contagion chains via mobility data and artificial intelligence; Identifying and addressing inaccurate information on COVID-19; Finding marginalized areas most affected by second and third-order pandemic impacts to deploy the appropriate interventions needed."
[8] Granovetter, "The strength of weak ties", 1973
[9] McChrystal "Team of Teams", 2018; McKinsey "An operating model for the next normal: Lessons from agile organizations in the crisis", June 25, 2020
[10] http://snap.stanford.edu/class/cs224w-readings/Brin98Anatomy.pdf
[11] While the exact functioning of Google's search is a trade secret, Google has officially confirmed that it uses the number of so-called "backlinks" to weigh the value of a page in its results. As Google's founders explained in the original Stamford paper ("The Anatomy of a Large-Scale Hypertextual Web Search Engine"), a backlink is a link that another website has toward that page. The more numerous the backlinks, the more valuable the page toward which the backlinks point. But also, the more backlinks point to the website that generates those backlinks, the more valuable that page becomes. In other words, links from websites that in turn are the destination for many other backlinks, are more important.
[12] Matt Clancy, Are ideas getting harder to find because of the burden of knowledge?, https://www.newthingsunderthesun.com/pub/zsc23qxz/release/5; Innovation (mostly) Gets Harder https://www.newthingsunderthesun.com/pub/bvmu4ol2/release/3?readingCollection=9f57d356
[13] https://cci.mit.edu/
[14] Facebook, Amazon, Apple, Netflix and Alphabet
[15] Malone, Superminds, 2018
[16] Agrawal, Prediction Machines, 2018
[17] Uzzi, "The Increasing Dominance of Teams in Production of Knowledge", 2007
[18] Several research over the years have pointed at the same broad trend. For one of the latest, see McKinsey "The key to unlocking a successful digital transformation" (2018) https://www.mckinsey.com/~/media/McKinsey/Business%20Functions/Organization/Our%20Insights/U

nlocking%20success%20in%20digital%20transformations/Unlocking-success-in-digital-transformations.ashx

[19] Quoted in Kane, Phillips, Copulski, Andrus "The Technology Fallacy: How People are the Real Key to Digital Transformation"

[20] Pentland, Social Physics – How social networks can make us smarter,

[21] David P. Reed. https://en.wikipedia.org/wiki/Reed%27s_law

[22] https://www.wired.com/story/ai-epidemiologist-wuhan-public-health-warnings/

[23] Inspired by: Sole, R.; Seoane, L. Evolution of Brains and Computers: The Roads not Taken. Preprints 2022, 2022050121 (doi: 10.20944/preprints202205.0121.v1)

[24] Agrawal "Prediction machines" ibid.

[25] https://bankunderground.co.uk/2020/10/20/machine-learning-the-news-for-better-macroeconomic-forecasting/

[26] https://en.wikipedia.org/wiki/Experience_curve_effects

[27] **Combining normal IQs**. Throw more average brains at some hard problems, and those problems might crack. If organized appropriately, a few average-IQ people who collectively take an IQ test can achieve an above-genius score. In software testing, the common wisdom is that "given enough eyeballs, any bug is shallow". These simple realizations fueled the emergence of crowdsourcing. Collective-IQ "engines" like hackathons, Wikipedia, Linux, Kaggle, Kickstarter, Innocentive have revolutionized the world of collective knowledge-creation and innovation. They're surprisingly good at competing with - and often beating - the traditional experts.

Combining unique IQ. Sanjay Ghemawat and Jeff Dean are the only level-11 coders at Google, and the founding fathers of MapReduce. We wouldn't have Hadoop without it, and big data would still be out of reach. Google wouldn't be what it is without them. What's their secret? They code, sitting side by side, on the same machine - complementing each other's' brain in real-time.

Combining good-enough people with good-enough machines: the best chess player is neither a human (it hasn't been for a long time) or a machine - it is a mix of them, and neither of them needs to be best in their field.

Corner office's secret. The most innovative CEOs harness their people network, not only their own intellect, for ideas and execution. Their genius is not only due to their own big brain. Some will even admit it.

Tenured employees' value. In companies where informal flows of knowledge is important, tenured employees tend to form valuable network connections, which makes them more effective as time goes by. Sometimes these informal flows become so effective that the formal ones recede. The downside is such an environment tends to create artificial hurdles for new joiners – especially the precious ones who bring new skills.

Attempts to avoid enterprise amnesia. "If my company knew what my company knows...", solutions would often emerge faster and more inexpensively. You've heard this before. Many companies are trying to make knowledge management tools better - to varying degrees of success.

"Trim teams". Amazon, with their famous "two-pizza principle", favors continuous experimentation through independent and scrappy teams. This allows for the work of many, loosely coordinated teams to stay focused and edgy and experiment enough so that some of them will hit big results at some point – cost-effectively.

Crowdsourced innovation. Among many other interesting examples, toy manufacturer LEGO famously embraced its client base's reverse-engineering of LEGO's robotic kit Mindstorm. The result was an explosion of innovative ideas, effectively turning a product into a crowdsourcing platform, and significantly contributing to LEGO's return to the commercial success.

Internal markets. Some large companies, such as China's white-goods manufacturer Haier[27], have radically re-thought their organizational structures to create a perpetual tension between internal units, and reduce their default bias to working with each other. For instance, Haier's individual-market supply chains, marketing and sales groups are free to buy external services if they don't feel their counterparts offers are good enough. **Mass-scale orchestration**. Hong Kong supply chain company Li & Fung, a $20bn enterprise specialized in apparel, personal products and furniture, acts as the enabler of massive flows of information and decision. It sources raw materials, manufactures and enables the logistics of finished products, without owning a single factory and keeping its own asset base to a minimum.

Natural experiments at macro-scale. More generally, in China the use of "natural experiments" and iterative innovation is comparatively very pronounced. Organizations there are unencumbered by decades of organizational orthodoxy, and the vast market is full of risk-taking consumers that typically tolerate the imperfections of fast product iterations. In addition, China has little qualms about using artificial intelligence extensively on consumer data of those experiments, which adds to the ability of local players – such as Alibaba – to improve their algorithms.

Natural Language Processing. Natural language processing is becoming very accurate. AI's ability to understand us and talk back is growing by leaps and bounds, combing existing (and often crowd-produced) content at an unprecedented scale. Think of how useful Google search has become. Or just ask Siri.

Known key organizational traits. Recent McKinsey research[27] sheds some light on the behaviors of companies best able to adopt digital transformation practices and succeed: they use multiple sources of customer data; their business leaders dedicate time to learn about digital technology; their teams share test-and-learn findings systematically; they reallocate digital talent aggressively; they evaluate opportunities as a portfolio, and divest proactively; they reallocate capital expenditure intentionally, and defund underperforming initiatives with no remorse. All of these deal with the ability of improving the arc of data-to-insight-to-action and related processes (the brain is very good at it), by connecting dots between ideas and reallocating resources.

Design thinking has emerged in the last ten years to become the de-facto standard for much innovation. It is an attempt to achieve collective intelligence, but it relies mainly on the human mind of a few people gathering in a shared physical space. Its ability to sense is limited to the information those people have. Ditto for the memories available to that group. So, while we love what it does, design thinking is leverages the cognitive equivalent of the brain of an insect. And one that doesn't even communicate with its swarm. Lots of headroom there.

Network-driven knowledge and learning. As we will explore in several case studies, there is an opportunity for enlisting subject matter experts to curate the knowledge and engage the networks of people who need it. The experience of GE-spinoff, digital-operations services provider Genpact offers instructive insights on how that architecture helps people learn and use fast-changing skills in more scalable ways, for likely a fraction of the traditional learning and development (L&D) and knowledge-management (KM) costs.

[28] https://hbr.org/2018/11/the-end-of-bureaucracy
[29] Arena, Cross, Sims, Uhl-Bien "How to catalyze innovation in your organization", MIT SMR, 2017

[30] Malone, Superminds, 2018. "a *democracy* [is] a supermind where decisions for the group are made by a vote of its members." "a *market* [is] a supermind where decisions are made by individuals mutually agreeing to trade resources with one another." "a *community* [is] a supermind where decisions are made by informal consensus or according to shared norms, both of which are enforced through reputations and access to resources." "an *ecosystem* [is] a supermind where individuals interact without any overall framework for cooperation. In the short term, decisions are made by the law of the jungle: the individuals with the most power get what they want. In the long term, decisions are made by survival of the fittest: the individuals that survive, grow, and reproduce most successfully control the most resources."

[31] Malone, Superminds
[32] "DNA could store all of the world's data in one room", Science Magazine, March 2017
[33] For more on this topic, see the work on knowledge stored in networks by Cesar Hidalgo ("Why Information Grows: The Evolution of Order, from Atoms to Economies").
[34] A slime mold called Physaurum Polycephalic, which looks like a fungus displays some of the characteristics of animal instinct, like optimizing the search of food, and reacting to setbacks. As Robert Wright wrote in "Nonzero: The Logic of Human Destiny" ""The slime […] sits at the boundary between society and organism, never making a firm commitment. Its cells spend lots of time on their own, scooting along the forest floor looking for nutrients, occasionally reproducing by splitting into two. But when food grows scarce, the first cells to feel the shortage emit a kind of alarm call in the form of a chemical called acrasin. Other cells respond to the call, and a transformation ensues. The cells bunch up together, form a tiny slug, and start crawling as one." Trees communicate and even share resources with each other to fight off threats. We now know that the roots exchange chemicals across trees in ways that help them alert each other of threats and defend (jokingly called "wood wide web"). Pando, possibly the oldest living organism on earth, is one large group of interconnected tree roots that have successfully adapted to varying conditions for over 80,000 years in a way that is far from random. Apart from the use of fungi, plants species communicate with other species through volatile organic compounds (VOC). They use VOC to attract or repel insects, which is another form of network connection. Ants optimize their exploration in search of food by tracking each other's movement through pheromones (chemical traces) or can form a raft in case of a sudden flood. African wildlife migrations leverage the complementary skills of thousands of animals, including different species (e.g., zebras and wildebeests). In these examples, the underlying dynamics are simple, but their combined cognitive power is remarkable. What is also remarkable, is that network density also seems to have an impact on the ability of animals to remember their communication language – as observed in the Regent Honeyeater's impaired song production ability as their population dwindles.

[35] Let's illustrate the concept with a simple example – what's the value of collaboration between entities? An entity can be anything that has some level of processing capacity: a simple organic one, like

bacteria and plants; more complex but still organic ones, like people; or an inorganic one, from a simple sensor (say, a humidity meter) to a sophisticated computer. Each of these will have a finite set of abilities: for instance, some people will be good at math and statistics, and others at reading people's emotions. Some machines will be good at sourcing data, others at applying advanced statistical models, others at moving "work packages" (called "case" in process-design jargon) through large workflows that weave together decision-algorithms and people. Those different capabilities are shown as A, B, C, and D in the chart for an illustrative case of two entities. When successfully combined, the resulting entity could benefit from the ability level of both. Agile methods, design thinking, and brainstorming to mention just a few, have ultimately attempted to get people with diverse abilities to complement each other. The upshot of the combination of skills is not just the elimination of each party's gap: it can be the building on and hence multiplying each other's capability, to achieve what individually wouldn't have been possible.

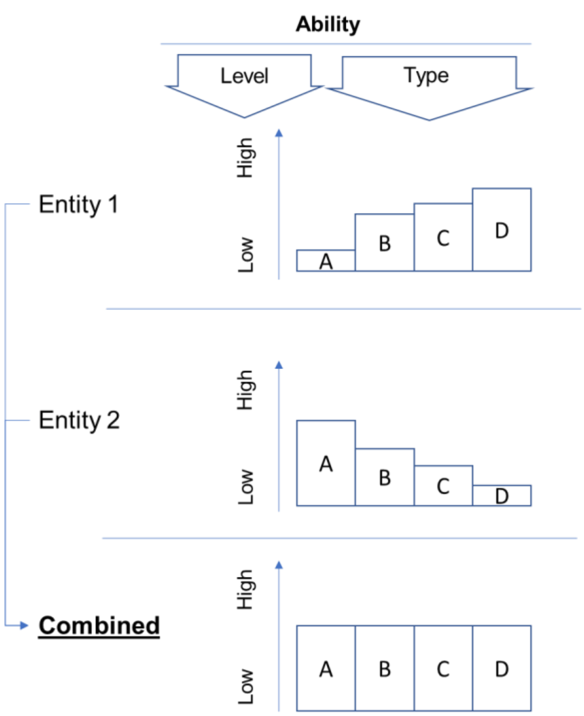

Such simple concept is at the root of innovation practices since time immemorial, and clearly at the foundation of methods like design thinking. After leaving Pixar, Steve Jobs famously became a connector between Silicon Valley entrepreneurs and ideas, which helped him to leverage the individual groups working at the components of the iPhone. As it turns out, experts in each of the subgroups didn't get to complement each other's capabilities – they got to learn how to collectively create one of the most transformative devices in history.

The idea is intuitive at small scale, but what happens if we exponentially multiply the interactions? To start with, Internet of Things (IoT) devices like sensors, with some computational power at the network edge, can similarly interoperate with cloud-based computers able to run massively complicated computation at scale. The emergence of architectures like service-oriented architecture (SOA) and more recently advanced programming interfaces (APIs), especially when combined with standard web transmission protocols, facilitate that interoperability. Of late, companies like Headspin.io have started making those programming jobs easier. The result has been an explosion of data-driven applications that help a variety of things, from telecom networks to agricultural equipment and commodity markets. This trend will not abate anytime soon: the emergence of 5G cellular data will only fuel the adoption of more edge computing, now that small hardware (e.g., increasingly small drones) can be placed at the edge of a network and perform significant processing with no time lag.

Let's add humans to the equation. We are witnessing - and arguably we can accelerate - the emergence of a system composed of millions of specialized and interconnected parts. For instance, there are human-machine subsystems that:

- optimize the trading of cocoa, through traders that use information and other technology, and the signals that markets provide – ultimately matching price with quantities that, barring market distortion, constitute a good equilibrium between incentives for all parties;
- choose songs at the Eurovision contest and spread memes on the internet, through social media and related "democratic" up or downvoting further amplified by algorithms;
- find cures for cancer by involving hundreds of thousands of health experts, but also laypeople such in healthcare community website PatientsLikeMe;
- weather-forecast models, enhanced by ubiquitous weather sensors, enable proactive redirection of physical resources in supply chains, including those that transport humans (like Uber's algorithms that by modulating monetary incentives optimize the deployment of cars, also based on weather predictions).

This is a planet-sized mesh of information streams, partially generated by people, partly by machines – and partially *curated* by humans, and partly by machines. These systems are ostensibly imperfect, and they do lend themselves to new and sometime disturbing forms of hijacking, but they're often and increasingly much more effective than their older, centralized-decision counterparts.

Social networks have been investigated for a long time, and in the last decades the advances in data science, combined with the increased connectivity between people, have started to make their presence felt. Some researchers state that based on people network properties, each of us can influence directly or indirectly up to *three* degrees of adjacency away (I know someone who knows someone who knows someone who knows someone). If that doesn't sound like much, consider that it has been estimated that the whole world's population falls within *six* degrees of separation. Which means that each of us can reach halfway through the entire world. In other words, in a hyperconnected environment, there's a theoretical possibility that we influence, and are influenced by, the ideas of half of the world – at least. Now that that technology infrastructure has become ubiquitous, it has started wiring a neural-like network composed of billions of people and machines. We don't see it, and it isn't exactly the equivalent of a biological brain, but it does something that a brain does: processes information as an entire brain, that is as a network and not just individual neurons. It also does so at much higher level of scale: able to ingest more stimuli, and process information in exponentially combinatorial ways.

Today, all these systems are fueled by a new set of incentives including very sophisticated online advertising markets and innovation-hungry capital markets and are enabled by more powerful technology.

Meanwhile, millions of experts optimize their own organizational and process and technology design, but most of them are working on the equivalent of individual neural cells or synapses - taking just a few bits of that collective, connected brain at the time. Some large internet-native firms like Google are already harnessing that new, ubiquitous processing capacity.

[36] Matt Clancy, "Combinatorial innovation and technological progress in the very long run", November 2021
[37] The analogies with neuroscience are intuitive and fascinating. There are differences for sure: neural cells differ from one another in that they have different connections, but they're otherwise comparatively similar. They're certainly more similar to each other than people are – as people are the result of the *permutations* of those cells within their brains, and their intelligence is the emergent quality of those permutation. And biology is a very different technology compared to what is used in people-machine networks. But some of core properties are similar.

For instance, we find specialized areas in both that do nothing else but combine inputs from other regions to respond to unprecedented situations: the prefrontal cortex, or senior management (on a good day). Many brain areas can be generalist and most neural cells can perform any job at their outset, but they end specializing themselves because of the connections that they create with other cells that "fire together" in specific situations – like when stimulated by sensorial input.

Brain networks react in particular ways when stimulated by hormones (e.g. serotonin for wellbeing, and cortisol for stress, among others), inhibiting or emphasizing pathways that lead to specific reactions and behaviors. Think of craving for sex or sport, or the implications of anxiety on how the brain directs the body. Interestingly, human brains in stress conditions trigger dopamine-endorphin "feel-good" hormone cycles when interacting with other people (including those well beyond our immediate social circles), giving our species a distinct social collaboration incentive when facing complex situations.

Similarly, human networks shape themselves based on emotional reactions, as when they generate eco-chambers of like-minded yet dangerously siloed views. Think of those as "intrinsic motivators" (behavior driven by internal rewards) for both brains and people-machine networks. Now think about extrinsic motivations. Brain circuitry requires energy, and it will create specific behaviors to get it - food, sugar, and some say, even benefited from the invention of technology (cooking) that made its assimilation more efficient and allowed our species to spend more time in activities other than foraging[37]. Similarly, people-machine networks need money to buy time, whether man- or CPU-hours. Markets allocate that cognitive capacity. In that respect, the emergence of cloud computing and the related market-based competition has generated a lot more fluid allocation of technology resources. Similarly, the inception of crowdsourcing and more recently, the surge of the gig economy have restructured the allocation of people's cognitive capacity.

One more similarity. According to modern neuroscience, one of the main contributors to human intelligence is active memory, that is the ability to keep multiple ideas in mind at the same time so that they get recombined by the pre-frontal cortex. What would be possible if machines could help large groups of people to "keep fresh in mind" dozens of relatable ideas?

It is not implausible that some of these similarities have the potential to spawn entirely new disciplines to deal with extremely large, hybrid human-computer networked systems. Something that shifts the focus from individual people and machines, towards a focus on the networks they form, where hundreds, thousands or even millions of them mesh into an ensemble, where each node's intelligence – whether human or machine – is augmented.

[38] Daniel Kahneman, "Thinking fast and slow", 2014
[39] Mckinsey Quarterly, "Organizational Health: a fast track to performance improvement", 2017
[40] A few papers have explored the importance of being exposed to adjacent knowledge. Matthew S. Clancy, Paul Heisey, Yongjie Ji & GianCarlo Moschini "The Roots of Agricultural Innovation: Patent Evidence of Knowledge

Spillovers", 2020; and Philipp B. Cornelius , Bilal Gokpinar , Fabian J. Sting "Sparking Manufacturing Innovation: How Temporary Interplant Assignments Increase Employee Idea Values", 2020

[41] **Sense** – Machines can be invaluable in scaling up human's ability to sense, both by accessing people who might have new information faster, or by connecting to artificial sensors such as internet of things devices (IoT). For instance: What's the snowfall in a specific zip code? What's the best trending ad? What's the mood of the electorate? What color of lipstick is trending for a particular age group? What is the engagement level of my staff? Is there toxic behavior in any of my teams?

Remember – Machines can help people dramatically increase the active memory of what happened, what worked or didn't in specific conditions – both by enabling access to people who may remember, or to repositories where that knowledge is stored. For instance: What ideas are relevant to my problem now? What worked best in previous conditions? What existing financial model predicts best my cash flow for the next quarter? What is the price sensitivity of that customer segment? Where are the people with that skill in my organization? Is there a consultant with that exact expertise in the market? Does every production unit have the right standard operating procedure available? And here are two examples of the emergent response from the network: YouTube is full of "how-to" videos, for instance, crowdsourced from millions of people; and Wikipedia has leveraged millions of sources through tens of thousands of people.

Create - new options. For instance, new fashion or product designs on Etsy.com, or a new algorithm on Kaggle to predict your Netflix movie preferences. Or use generative design where machine learning-based algorithms help create and filter dozens of options that humans, either individually or in groups, enhance and decide upon.

Decide - on which options to take. Diversity of opinion and input, independence of thinking, decentralization to draw on local knowledge, and appropriate mechanisms to aggregate input to arrive at the final decision are key attributes for collectively intelligent decision networks[41]. From social media "likes" that are mined by web advertising firms, to Amazon's distributed product ratings that drive individual-consumer and supply chain decisions. Or think of investment banking, where routinely analysis is performed by machines trained by humans, and humans inform decision protocols after validating the result of algorithms. Machines can help simulate event outcomes at scale, or simply enable the connection of many people and guide the decision process (from mundane workflows, to more complicated and judgment-based ones).

Act – while this is not a cognitive process, it is inextricably connected to decisions made by large, coordinated groups (think workflows in an enterprise to process a healthcare claim, for instance). And in recent years it has become evident how actions are useful to make machines learn: from humans labeling road situations to train self-driving cars, to adversarial networks where machines play against each other literally by building strategies to respond to the actions of the other algorithm (for instance these techniques are used in helping computers make sense of fake artifacts like videos).

Learn – the results of the process of creation of options, as well as the experience of the action. Principles of machine learning can be applied to superminds too - as they describe what worked best and enable the remembering next time. For instance, the threshold of outstanding credit is the limit beyond which my clients may default their payments. The snowfall level beyond which supply-chain delivery time gets elongated by 10%. The rate of clickthrough to beer ads that feature smiling faces, for target audiences in the 21-25 age range. Learning is preparing for future retrieval of memories.

[42] Let's use the example of a complex enterprise process – for instance a consumer products' supply chain planning. There, planners use input from a software system connected to internal production and logistics technology, and external demand sensing. Planning is about predicting demand and ability to supply, optimizing the fixed cost of factories, the variable cost of raw materials and distribution, and balancing them with the risk of foregoing sales in the short but also in the longer term. And planners do that on thousands of different products, multiplied by dozens of factories and hundreds of retail distribution centers, and under dozens of different promotional conditions.

That process is typically hierarchical, with supply chain planners providing input that is finally decided upon by management, typically a group of senior people across supply chain and sales. The recommendation is validated by a software algorithm to avoid overcommitting products to a specific channel, or inadvertently disregarding the constraints in some factory. Software for this work has already become quite sophisticated and will continue to weave together all these cognitive processes in hybrid human-machine systems - superminds connected not by physical proximity but rather because of the availability of cost-effective computing and human resources. Such a system would likely display a collective IQ higher than the individual IQs of its parts. That is, there a collective intelligence that emerges from the combination of individual intelligences as it can:

- *sense* accurately the ground-level demand through improved sensors at multiple levels, starting from the end retailer's shelf; as well as the potential vagaries of weather and traffic conditions. And the prices of the mode of transportation;
- *remember* past conditions and thanks to machine learning, it enables the recognition of patterns, whereby new conditions are compared with similar ones from the past;
- *create* new options, partially automatically and in part through human judgment – both from individual planners, and increasingly from groups of planners working together and triangulating each other's views. In particular, the emergence of outsourcing of certain types of data work now enable the centralization of heavy data work – with teams increasingly located across the world;
- *decide* – typically through a human decision heavily augmented by the machine, and in simpler cases, totally driven by the machines;
- *act* – in this case the system sends the orders to the various production and logistics systems, and locks out any other request;
- *learn* – by comparing what has worked after the fact (for instance, mismatch of final demand with the supply, and consequent loss of revenue) the system is now able to alter its future recommendations.

Some of it exists today, especially in large and sophisticated firms that employ state-of-the-art technology. However, innovation may come from reimagining the process with a more thorough disassembly and network-level reassembly. For instance:

- **Better sensing** - Many more sensors, both machine and human (e.g. gig workers monitoring promotional activities in stores) throughout the supply chain, providing more accurate data for both the supply and demand side
- **Better remembering and creating** - Increased accuracy of the algorithm as cloud computing accesses specialized "prediction utilities" that improve the accuracy in specific regions, weather conditions, product segments etc.
- **Better decisions** – for instance including additional people in the process, based on the quality of their previous judgment calls in similar conditions, but also with algorithms able to correct the systematic bias that those people might have
- **Better learning** – for instance by giving more parts of the supply chain the ability to become aware of execution issues that caused problems, such as: mistakes in preparing the custom promotional materials for a limited-quantity batch which resulted in a disproportionate loss of demand in a critical channel; or the revision of order-management algorithms to prevent them from delaying particular orders (e.g. because of customer credit lines temporarily exhausted), adjusted dynamically in conjunction with the upstream planning systems.

[43] Fuller, "What makes great salespeople", Harvard Business Review, July 8, 2015
[44] Michael Arena, Rob Cross "How to catalyze innovation in your organization", MIT Sloan Management Review, 2017
[45] https://en.wikipedia.org/wiki/Cognitive_Surplus
[46] www.circle.so
[47] Cross, Davenport, Gray "Collaborate smarter not harder", in "A manager's guide to the new world of work", MIT Sloan Management Review, 2019
[48] Using Teams to Solve Hard Problems by J. Richard Hackman, 2011
[49] https://cci.mit.edu/genome-of-collective-intelligence/
[50] Christoph Riedl, Young Ji Kim, Pranav Guptag, Thomas W. Malone, and Anita Williams Woolley "Quantifying collective intelligence in human groups", 2021
[51] Daniel Kahneman, "Thinking fast and slow", 2014
[52] Agrawal "Prediction machines" ibid.
[53] The Syrian Archive (www.sirianarchive.org) is an initiative using tools that scrapes social media feeds to detect and classify imagery pertaining to war equipment, with the intent of documenting and making searchable evidence of war crimes. The neural net AI curates and classifies information automatically.
[54] https://www.nesta.org.uk/feature/ai-and-collective-intelligence-case-studies/early-warning-project/
[55] Batjargal, Bat. 2013. Institutional Polycentrism, Entrepreneurs' Social Networks, and New Venture Growth
[56] Hite, J. M., & Hesterly, W. S. 2001. The evolution of firm networks: From emergence to early growth of the firm
[57] Zaheer, A., & Bell, G. G. (2005). Benefiting from network position: Firm capabilities, structural holes, and performance.
[58] Source: Burt, Robert. 2004. Structural Holes and Good Ideas
[59] https://blogs.cornell.edu/info2040/2011/10/03/social-networks-and-public-health-the-hidden-influence-of-social-networks/
[60] https://blogs.cornell.edu/info2040/2011/10/03/social-networks-and-public-health-the-hidden-influence-of-social-networks/
[61] Cross, Davenport, Gray "Collaborate smarter not harder", in "A manager's guide to the new world of work", MIT Sloan Management Review, 2019
[62] Painter, Daniels, Laubicher, "Innovations are Disproportionately Likely in the Periphery of a Scientific Network", 2020
[63] Fuller, "What makes great salespeople", Harvard Business Review, July 8, 2015

[64] Shore, Bernstein, Jang "Network centralization and collective intelligence: a randomized experiment", 2020
[65] Xi'an Jiaotong (2016),"Effects of centralized, decentralized and reintegration group structures on groupperformance",Operation Research and Management, Vol. 25 No. 5, pp. 6-14.
(PDF) Literature review on collective intelligence: a crowd science perspective.
[66] https://hbr.org/2019/03/a-novel-way-to-boost-client-satisfaction
[67] https://en.wikipedia.org/wiki/Dunbar%27s_number
[68] Kristhakis, Fowler "Connected: The surprising power of our social networks", ibid
[69] Fry, Cai, Zhang, Wagner "Consolidation in a Crisis: Patterns of International Collaboration in COVID-19 Research", 2020 https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3595455 which includes the chart below showing the network structure of coronavirus related scientific research.

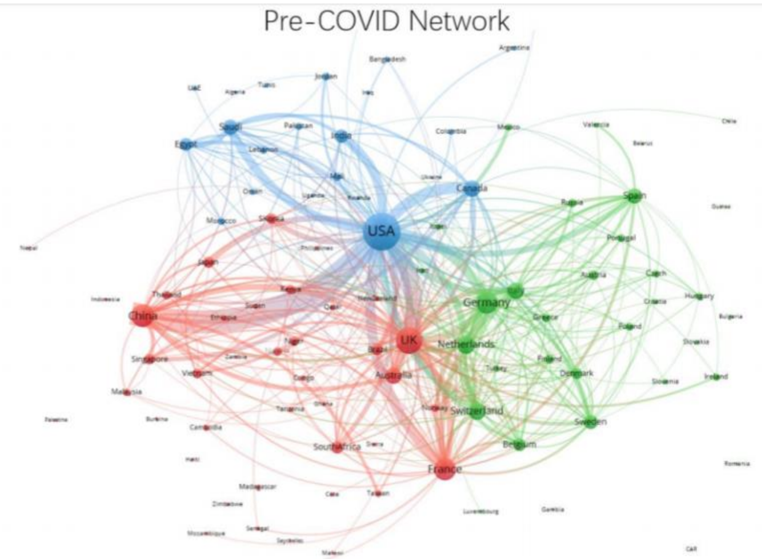

Figure 1. Network of International Collaborative Relationships in pre-COVID-19 Period, January 1, 2018 to December 31, 2019. Edges lower than 4 are removed.

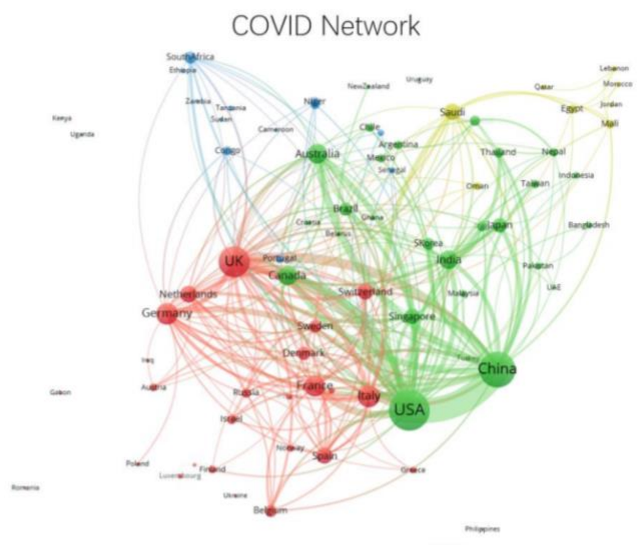

Figure 2. Network of International Collaborative Relationships in COVID-19 Period, January 1, 2020 to April 8, 2020. Edges lower than 2 are removed.

[70] Tim Hartford, “Messy: The power of disorder to transform our lives”, 2019
[71] Scott Page “The Difference”, 2007; For instance, analysis of millions of medical papers shows mixed gender teams do more impactful and cited work; randomized teams with different gender breakdowns ran a startup for a year - teams close to 50/50 male-female did best. Yang, Woordruff, Tian, Jones, Uzzi “Gender diverse teams produce more innovative and influential ideas in medical research”; Hoogendron, Oosterbeek, van Praag “The impact of gender diversity on the performance of business teams: evidence from a field experiment”
[72] https://hbr.org/2019/03/a-novel-way-to-boost-client-satisfaction

[73] Areejit Samal et al, Network geometry and market instability, *Royal Society Open Science* (2021)
[74] “Exponential: How Accelerating Technology Is Leaving Us Behind and What to Do About It”, Azeem Azhar, 2021
[75] https://restofworld.org/2020/migrant-workers-virtual-solidarity/
[76] Lakhani and Jeppesen, “Marginality and Problem Solving Effectiveness in Broadcast Search”, 2010, mentioned by McAfee, Brynjolfsson, “Machine, platform, crowd”, 2017.
[77] “Prior shared success predicts victory in team compeittions”, Mukherjee, Huang, Neidhardt, Uzzi, Contractor, In Nature Human Behaviour 3, 2018
[78] “First Exploratory Geographical and Social Maps of the Maker Movement” Massimo Menichinelli and Alessandra Gerson Saltiel Schmidt, 2020-06-22
[79] “First Exploratory Geographical and Social Maps of the Maker Movement” Massimo Menichinelli and Alessandra Gerson Saltiel Schmidt, 2020-06-22
[80] Peter Gloor “Sociometrics and human relationships”, 2017
[81] https://mitpress.mit.edu/books/honest-signals
[82] https://hbr.org/2019/03/a-novel-way-to-boost-client-satisfaction
[83] Michael Arena, “Adaptive Space”, 2018
[84] Damon Centola “How behavior spreads”, 2018, Princeton University Press
[85] https://www.technologyreview.com/2020/09/04/1008156/knowledge-graph-ai-reads-web-machine-learning-natural-language-processing/
[86] Meslec, Aggarwal, Courseu “The Insensitive Ruins It All: Compositional and Compilational Influences of Social Sensitivity on Collective Intelligence in Groups”, Frontiers in Psychology, 2016 https://www.frontiersin.org/articles/10.3389/fpsyg.2016.00676/full
[87] Wooley, Chabris, Pentland, Hashmi, Malone 2010 “Evidence for a collective intelligence factor in the performance of human groups”
[88] Hashmi, Malone “Effects on Collective Intelligence of Online Collaboration Tools and Group Size”, 2018
[89] Bates Gupta 2017, Crede Howardson 2017. Capraro coccioni 2016, Moore Tenbrunsel 2014, Heylighen 1995, Imbimbo Guazzini 2019
[90] Malone – women proportion in groups with high collective intelligence
[91] Modupe Akinola, Columbia University
[92] Christoph Riedl, Young Ji Kim, Pranav Guptag, Thomas W. Malone, and Anita Williams Woolley “Quantifying collective intelligence in human groups”, 2021
[93] Ray Dalio, Ted Talk, “How to build a company where the best ideas win”
[94] Quoted in Kane, Phillips, Copulski, Andrus “The Technology Fallacy: How People are the Real Key to Digital Transformation”
[95] Christoph Riedl, Young Ji Kim, Pranav Guptag, Thomas W. Malone, and Anita Williams Woolley “Quantifying collective intelligence in human groups”, 2021
[96] Gartner researchers found that even prior to the Covid-19 emergency, the effectiveness of teams wasn’t degraded if at least 50% of the team was co-located, and that the impact of reduced diversity would counter a higher location concentration. Gartner “Fusion Teams by the Numbers: An Empirical Analysis of Digital Business Teams”, 2020. It is likely that many teams, forced to change their ways of working during the Covid-19 pandemic, are now more effective in working remotely.
[97] https://hbr.org/2019/03/the-feedback-fallacy
[98] https://www.bcg.com/en-ar/publications/2020/tap-your-company-collective-intelligence-with-mindfulness.aspx
[99] Dery, Sebastian, “Building business value with employee experience,” MIT CISR Research Briefing 17, no. 6 (2017).

[100] Harvard's Misha Teplitskiy, and coauthors showed that diversity of political views in Wikipedia articles produced higher quality content (quoted in The Age of AI: Strategy and Leadership When Algorithms and Networks Run the World by Marco Iansiti, Karim R. Lakhani, 2019)
[101] IT analyst firm Gartner reported that diverse "fusion teams" achieved faster speed and higher value capture in their projects." Fusion Teams by the Numbers: An Empirical Analysis of Digital Business Teams", January 2020.
[102] https://www.bcg.com/publications/2018/how-diverse-leadership-teams-boost-innovation.aspx
[103] https://www.bcg.com/publications/2018/global-landscape-of-corporate-vitality.aspx
[104] This hypothesis is consistent with the results of Yi Bu, Ying Ding, Xingkun Liang and Dakota Murray presented at the SIG Panel for Scientometrics in 2020.
[105] well.blogs.nytimes.com/2013/10/03/well-quiz-the-mind-behind-the-eyes/ ; www.labinthewild.org/

[106] https://en.wikipedia.org/wiki/Psychological_safety
[107] https://rework.withgoogle.com/print/guides/5721312655835136/
[108] Markova, McArthur "Collaborative Intelligence", 2015
[109] Michael Arena, Rob Cross "How to catalyze innovation in your organization", MIT Sloan Management Review, 2017
[110] Dalio "Principles: Life and Work", 2018
[111] www.Minessence.net
[112] Shalom Schwartz, 2012
[113] https://www.mckinsey.com/business-functions/organization/our-insights/how-to-create-an-agile-organization
[114] https://www.technologyreview.com/2020/08/17/1006604/remote-workers-want-to-recreate-those-watercooler-moments-virtually/
[115] https://en.wikipedia.org/wiki/Dunbar%27s_number
[116] www.nfx.com
[117] Hayes, Chumney, Wright, Buckingham "The global study of engagement", ADP, 2018
[118] https://en.wikipedia.org/wiki/Cognitive_Surplus
[119] https://rework.withgoogle.com/print/guides/5721312655835136/
[120] Safi Bahcall, Loonshots, 2019
[121] There's significant literature on the subject. Some interesting materials here Edwards, Roy "Academic Research in the 21st Century: Maintaining Scientific Integrity in a Climate of Perverse Incentives and Hypercompetition", 2016; Anderson, Ronning, De Vries, Martinson, "The Perverse Effects of Competition on Scientists' Work and Relationships", 2006; https://svpow.com/2017/03/22/why-do-we-manage-academia-so-badly/ https://svpow.com/2017/03/17/every-attempt-to-manage-academia-makes-it-worse/
[122] The Economist, May 7, 2020 https://www.economist.com/science-and-technology/2020/05/07/scientific-research-on-the-coronavirus-is-being-released-in-a-torrent
[123] See for instance Nicholas Christakis and James Fowler "Connected", 2009
[124] Fukuyama "Trust: the social virtues and the creation of prosperity", 1995
[125] HBR, The Neuroscience of Trust, 2019
[126] Optimal incentives for collective intelligence; Richard P. Manna and Dirk Helbing https://www.researchgate.net/publication/316558876_Optimal_incentives_for_collective_intelligence
[127] https://hbr.org/2018/11/the-end-of-bureaucracy
[128] Arena "Adaptive space"
[129] Maslow "A theory of human motivation", 1943
[130] Referenced by Keller, Schaninger, "Beyond Performance 2.0", 2019
[131] https://en.wikipedia.org/wiki/Klout
[132] https://en.wikipedia.org/wiki/Mihaly_Csikszentmihalyi

[133] Originally framed by Merrill Flood and Melvin Dresher https://en.wikipedia.org/wiki/Prisoner%27s_dilemma
[134] Francis Fukuyama, “Trust: The Social Virtues and the Creation of Prosperity”, 1995
[135] David Sun Kong, Community Bio
[136] https://en.wikipedia.org/wiki/DARPA_Network_Challenge
[137] Joshua Graff Zivin, Elizabeth Lyons, “The effects of prize structures on innovative performance”, NBER Working Paper Series, 2020
[138] Source: Pitchbook https://pitchbook.com/profiles/advisor/107473-06#overview
[139] “How scientists are embracing NFT”, Nature, June 18th 2021
[140] https://www.blockdata.tech/blog/general/what-is-the-biggest-nft-marketplace-a-top-10-comparison
[141] https://vitalik.ca/general/2019/12/07/quadratic.html
[142] https://medium.com/ethereum-optimism/retroactive-public-goods-funding-33c9b7d00f0c
[143] Optimal incentives for collective intelligence; Richard P. Manna and Dirk Helbing https://www.researchgate.net/publication/316558876_Optimal_incentives_for_collective_intelligence
[144] Forrester Research, The Total Economic Impact of Microsoft Teams, 2019
[145] Amy Webb, in “The 11 Sources of Disruption Every Company Must Monitor.” MIT SMR, March 17, 2020 https://sloanreview.mit.edu/audio/how-to-spot-disruption-before-it-strikes/
[146] Aminpour, Gray, Jetter, Introne, Scyphers, Baltaxe, Arlinghaus “Harnessing the collective intelligence of natural resource users for conversation”, 2020.
[147] PatientsLikeMe is a particularly interesting example, because of their collaboration started in 2017 with Chinese AI company iCarbonX whose stated intent is to “digitize the components of life”.
[148] Seth Stephens-Davidowitz 2017 “Everybody Lies”
[149] MIT Tech Review July 2019, highlighting Rediet Abebe’s work as a reason for recognizing her as a top under-35 innovator.
[150] Steven Johnson’s “Where Good Ideas Come From” is a good primer on the subject
[151] https://en.wikipedia.org/wiki/Sci-Hub
[152] Sinan Aral, ibid.
[153] https://blog.safecast.org/en/
[154] https://innovation.microsoft.com/en-us/planetary-computer
[155] https://sustainability.google/projects/airview/
[156] https://insights.sustainability.google/
[157] Nbn.org.uk
[158] https://www.nesta.org.uk/feature/ai-and-collective-intelligence-case-studies/inaturalist/
[159] “How volunteer observers can help protect biodiversity”, The Economist, Jun 15th 2021
[160] https://www.mercycorpsagrifin.org/2020/12/08/desert-locust-citizen-reporting/
[161] The effort is part of the “UN Global Pulse”, harnessing support from the respective UN agency, the World Food Program, and local constituencies.
[162] Carrillo-Zapata, Milner, Hird, Tzoumas, Vardanega, Sooriyabandara, Giuliani, Wifield, Hauert, “Mutual shaping in swarm robotics”, 2020
[163] https://www.economist.com/briefing/2021/08/07/open-source-intelligence-challenges-state-monopolies-on-information
[164] www.syrianarchive.org
[165] https://www.economist.com/briefing/2021/10/23/enter-third-wave-economics
[166] As documented by NYU GovLab http://www.thegovlab.org/static/files/publications/CI-AI_oct2019.pdf
[167] https://programmablesearchengine.google.com/about/

[168] **Extraction of URLs from Google.** The construction of Boolean string is the basis of Google search. It consists of keywords sourced from the Mindmap. This step ensures getting relevant URLs from the source.

Steps	Tasks	Tools	Output
Construction of Boolean String(s)	1. Pick a set of words or phrases for constructing the Booleans. 2. Create a set of Booleans using a combination of words or phrases. It will include a time frame (to get 18 months old data), profile name and keywords from the Mindmap. 3. Test the Booleans in the Google search engine and identify the Boolean string(s).	Google	The Boolean string is a combination of profile name and other keywords.
Extraction of URLs	4. Use the Boolean search string to get the list of search results. **Note:** The result from Google includes content authored, shared and cited.	Google API, Python	The dataset contains a title, truncated description of the resource and URL

Note: The output of Google search are the articles in which the profile may have cited, mentioned, or authored. It may include all file types such as PDFs, docs, etc. The time limit can be set for a range of date like after 2018 or before 2019. It, however, does not allow data extraction for the desired time as a parameter in the query.

Extraction of URLs from Twitter. The Twitter account of a profile has the following types of tweets:

- Tweets that they liked
- Tweets that they re-tweeted
- Original tweets

Steps	Tasks	Tools	Output
Extraction of tweets	5. Use Twitter API to fetch the tweets and related data of a given profile (within a period). The API will extract tweets from “likes” and “tweets/re-tweets”. 6. Extract hashtags, username and URLs from the tweets.	Python, Twitter API, Twitter	The dataset will contain tweets, username, hashtags, URLs (tweets or original source of the content).
Conversion of URLs to readable URLs	7. Collate the URLs. 8. Convert the URLs to a readable format. 9. Identify links that can be processed further. These are usually articles/video links. Some links need to convert to a readable format. **Disclaimer:** conversion of URLs to readable format is limited to 2-cycle.	Python	The dataset will contain readable URLs.

Extraction of URLs	10. Collect the URLs for further processing	Python	Dataset containing URL

Some of the accounts may be in restricted mode. Such profiles need to be followed manually.

Extraction of URLs from LinkedIn. The LinkedIn activity of a profile can be followed from:

- Recent activity (contains both articles and shared posts)
- Articles (All articles written by the profile)
- Posts (Articles shared by the profile)

Steps	**Tasks**	**Tools**	**Output**
Extraction of URLs from recent activity (for daily following)	11. Use a scraping tool to extract the recent activity. 12. Extract URLs from the post content field. **Note:** the dataset extracted does not show the date and time	LinkedIn, Phantombuster, Python, R	The dataset will contain, URLs, post content, # of likes etc.
Conversion of URLs to readable URLs	13. Collate the URLs. 14. Identify readable URLs from the list. 15. Convert non-readable URLs to a readable format.	Python	The dataset will contain list of all readable URLs.
Extraction of URLs from posts (for 18-month retrospective)	16. Use a scraping tool to extract the shared articles. 17. Understand the fields. 18. Extract URLs from the data for further processing. **Note:** the data extracted does not show the date and time	ScrapeStorm, LinkedIn, Python, R	The dataset will contain title, URLs, # of comments, # of likes, etc.
Extraction of URLs from articles	19. Create Boolean string that will extract articles relevant to the Mindmap and for the given time. 20. Test the string to identify the best search string. 21. Use the search string on Google API to get all the relevant results.	Google API, Google, LinkedIn, Python	It will contain all the relevant URLs of the articles authored by the primary profile.

Extraction of PDFs and the text

Steps	Tasks	Tools	Output
Extraction of URLs having PDFs	22. Identify and filter PDFs from the URLs. 23. Download the PDFs from the URLs.	Python, R	All PDF files

	24. Map the name of the PDF file with the URL.		
Extraction of text from the PDF files	25. Extract text from the PDF files. 26. Add the text to the master data. 27. Get metadata from the PDF files and append to the master data. 28. Get a text summary from the text and append to the master data.	Python, R	The dataset will contain URLs, metadata, text summary and the associated text.

- Collect all URLs sourced from Google search, Twitter and LinkedIn. This will be collated by profile and by source.
- Extract domain names from the URLs.
- Run the URLs through the deduplication process to get unique links.
- Get the keywords, metadata, text, and summary from the URLs using an entity extraction tool and text summarization tool (for URLs other than PDFs).
- Tag the text, title, and summary of the URLs by mapping it with the keywords from Mindmap.
- Filter the relevant URLs that map to the Mindmap.
- Review the URLs and run summary statistics on the data.

Boolean strings used in this process look like the following:

after:2017 ++"Thomas Malone" AND ("collective intelligence" | crowdsourc* | "social network analysis" | "wisdom of the crowd" | "swarm intelligence" | “citizen science”
site:linkedin.com/pulse "valdis-krebs" AND ("SNA" | "ONA" | "citizen science" | "swarms" OR "collective intelligence" OR "social network analysis" | crowdsourc* | "wisdom of the crowd") AND posts AND (comments | shares | likes)
site:slideshare.net/ intext:"Thomas Malone" AND ("collective intelligence" | "social network analysis" | crowdsourc* | "swarm intelligence" | "wisdom of the crowd" | "citizen science") AND posts AND (comments | shares | likes)
site:reddit.com intext:"Thomas Malone" AND ("collective intelligence" | "social network analysis" OR crowdsourc* | "swarm intelligence" | "wisdom of the crowd" | "citizen science") AND posts AND (comments | shares | likes)
after:2017 site:medium.com ("thomas malone" | "@thomasmalone") AND ("collective intelligence" | crowdsourc* | "social network analysis" | "wisdom of the crowd" | "swarm intelligence" | "#SNA" | "#ONA") AND lang:en

For example, this process was followed to analyze the field of collective intelligence. About 1,000 relevant publicly available web articles published over 18 months were identified. The word cloud of the respective keyword showed not only the expected semantic field, but also highly specialized branches.

[169] A longer list is here www.contentmarketinginstitute.com/2018/06/tools-content-curation/
[170] http://blogs.cim.warwick.ac.uk/issuemapping/tools/
[171] https://www.gdeltproject.org/ and https://en.wikipedia.org/wiki/Global_Database_of_Events,_Language,_and_Tone
[172] https://toolbox.google.com/datasetsearch

[173] https://crisiscomputing.qcri.org/
[174] www.g2.com/categories/knowledge-management
[175] Kazemi, Gaffney, Garimella, Hale "Claim Matching Beyond English to Scale Global Fact-Checking", 2021
[176] https://thenewstack.io/ai-makes-new-scientific-discoveries-by-analyzing-old-research-papers/
[177] https://benevolent.ai/news/potential-treatment-for-COVID-19-identified-by-benevolentai-using-artificial-intelligence-enters-clinical-testing
[178] https://www.pnas.org/content/116/6/1870
[179] www.g2.com/categories/knowledge-management
[180] For a comprehensive list of academic databases and search engines, visit https://en.wikipedia.org/wiki/List_of_academic_databases_and_search_engines
[181] https://www.epistemonikos.org/en/about_us/methods
[182] Christoph Riedl, Young Ji Kim, Pranav Guptag, Thomas W. Malone, and Anita Williams Woolley "Quantifying collective intelligence in human groups", 2021
[183] Majchrzak and Malhotra, "Towards an Information Systems Perspective and Research Agenda on Crowdsourcing for Innovation" 2013
[184] For more on this, see Joseph B. Bak-Coleman, Mark Alfano, Wolfram Barfuss, Carl T. Bergstrom, Miguel A. Centeno, Iain D. Couzin, Jonathan F. Donges, Mirta Galesic, Andrew S. Gersick, Jennifer JacquetAlbert B. Kao, Rachel E. Moran, Pawel Romanczuk, Daniel I. Rubenstein, Kaia J. Tombak, Jay J. Van Bavel, and Elke U. Weber "Stewardship of global collective behavior", PNAS, 2021
[185] Mors, Waguespack "Dispersed teams succeed fast, fail slow", Harvard Business Review, Mar 26, 2021
[186] Source: Citi, "The Content Sunday" newsletter, October 2019
[187] Gloor, Swarm Leadership, 2017
[188] Bostrom, Superintelligence, 2014
[189] www.citiesofservice.org
[190] https://madrid-colab.xcolab.org/
[191] https://innovate.burningman.org/
[192] "A new approach to problem solving across sustainable-development goals". Brookings, 2021 https://www.brookings.edu/blog/up-front/2021/03/12/a-new-approach-to-problem-solving-across-the-sustainable-development-goals/
[193] Linda Xie https://linda.mirror.xyz/Vh8K4leCGEO06_qSGx-vS5lvgUqhqkCz9ut81WwCP2o
[194] www.discoverdaos.com
[195] Full disclosure: Gianni Giacomelli sits on the advisory board of KnomAI
[196] https://hbr.org/2011/07/the-big-idea-the-age-of-hyperspecialization
[197] "Exponential" Azeem Azhar, Penguin, 2021
[198] www.g2.com/categories/idea-management
[199] https://decoders.amnesty.org/
[200] www.blockbyblock.org
[201] www.khoros.com
[202] https://airtable.com/shrPm5L5I76Djdu9B/tbl6pY6HtSZvSE6rJ/viwbIjyehBIoKYYt1?blocks=bipjdZOhKwkQnH1tV
[203] https://digitalnative.substack.com/p/roblox-airtable-and-the-building
[204] www.mentalmodeler.com
[205] https://debategraph.org/Stream.aspx?nid=61932&vt=ngraph&dc=focus
[206] For an example of its functioning, view the discussion forum around liquid democracy
[207] https://www.nifi.org/en/common-ground-action
[208] https://en.wikipedia.org/wiki/Prediction_market
[209] www.tinycast.cultivatelabs.com
[210] https://goodjudgment.com/

[211] Eghbal “Working in public: the making and maintenance of open-source software”
[212] www.opsfound.org
[213] Wharton School Blockchain Project, in collaboration with the World Economic Forum, DeFi Beyond the Hype, 2015 https://wifpr.wharton.upenn.edu/wp-content/uploads/2021/05/DeFi-Beyond-the-Hype.pdf
[214] https://Thewild.com
[215] https://www.stripes.com/theaters/middle_east/2021-07-02/pokemon-go-bagram-troop-withdrawal-2027917.html
[216] http://stanfordhci.github.io/flash-teams/
[217] Aaron Dignan, “Brave New Work”, 2019
[218] https://www.semantic-mediawiki.org/wiki/Semantic_MediaWiki
[219] https://underlay.mit.edu/
[220] https://govinsider.asia/digital-gov/singapore-builds-government-wide-api-platform/
[221] https://dschool.stanford.edu/resources/the-bootcamp-bootleg
[222] https://www.gv.com/sprint/
[223] https://rework.withgoogle.com/print/guides/5721312655835136/
[224] https://en.wikipedia.org/wiki/Common_Knowledge%3F:_An_Ethnography_of_Wikipedia
[225] For more on this, read Rob Cross, Reb Rebele, and Adam Grant, “Collaborative Overload,” Harvard Business Review, January–February 2016.
[226] Jemielniak “Common Knowledge? An ethnography of Wikipedia” 2013
[227] Bernstein, Shore, Lazer “Improving the Rhythm of Your Collaboration”, Sloan Management Review, 2019
[228] https://www.pnas.org/content/115/35/8734
[229] https://www.bcg.com/en-ar/publications/2020/tap-your-company-collective-intelligence-with-mindfulness.aspx
[230] Carlson, Lewis, Maliakal, Easterday, Gerber “Supporting Knowledge of Design Risks in Collective Intelligence Efforts”, 2020
[231] Nesta UK “Participatory AI for humanitarian innovation”
[232] https://www.nesta.org.uk/toolkit/
[233] https://www.nesta.org.uk/toolkit/civicai/
[234] https://www.pnas.org/content/116/6/1870
[235] https://cci.mit.edu/takeda-superminds-program/
[236] Carnegie Mellon University researchers, quoted on Vice https://www.vice.com/en_us/article/dygnwz/if-youre-talking-about-coronavirus-on-twitter-youre-probably-a-bot
[237] Ed Yong “Zombie roaches” Ted2014
https://www.ted.com/talks/ed_yong_zombie_roaches_and_other_parasite_tales
[238] As an example, The Economist created a Twitter “clone” bot of US President Trump in 2020, and watched how the algorithm favored polarized content when recommending content. The Economist | Relevant content, 30th July 2020
https://www.economist.com/graphic-detail/2020/08/01/twitters-algorithm-does-not-seem-to-silence-conservatives?frsc=dg%7Ce
[239] See for instance https://hoaxy.iuni.iu.edu
[240] www.disinfo.eu
[241] https://en.wikipedia.org/wiki/Category:Fact-checking_websites
[242] Schuster, Fisch, Barzilay “Robust Fact Verification with Contrastive Evidence” March 2021, https://arxiv.org/pdf/2103.08541.pdf
[243] Allen, Arechar, Rand, Pennycook “Crowdsourced Fact-Checking: A Scalable Way to Fight Misinformation on Social Media”, 2020

[244] Sinan Aral, ibid.
[245] See also Peter Temes "We the People", 2019
[246] https://techcrunch.com/2020/04/23/google-is-extending-identity-verification-requirements-to-all-advertisers/
[247] Aral Sinan, The Hype Machine: How Social Media Disrupts Our Elections, Our Economy and Our Health – and How We Must Adapt, 2020
[248] For instance, Nick Bostrom "Superintelligence", 2014
[249] https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0205335
[250] Damon Centola, ibid.
[251]Kotsios, Magnani, Rossi, Shklovski, Vega "An Analysis of the Consequences of the General Data Protection Regulation on Social Network Research", 2019 https://arxiv.org/pdf/1903.03196.pdf
[252] See for example "Data trusts: A new tool for data governance", ElementAI and Nesta. https://hello.elementai.com/rs/024-OAQ-547/images/Data_Trusts_EN_201914.pdf
[253] Zeynep Tufekci, "Battlestar Galactica Lessons from Ransomware to the Pandemic", Insight newsletter, May 13 2021
[254] Fügener, Andreas and Grahl, Jörn and Gupta, Alok and Ketter, Wolfgang, Will Humans-in-The-Loop Become Borgs? Merits and Pitfalls of Working with AI (July 4, 2021). Management Information Systems Quarterly (MISQ) - Vol. 45 No. 3 , Available at SSRN: https://ssrn.com/abstract=3879937
[255] Yuval Harari "Sapiens", 2014
[256] Barbara Minto, "The pyramid principle" 2008
[257] https://www.economist.com/books-and-arts/2019/08/08/why-widely-spoken-languages-have-simpler-grammar
[258] Richard Baldwin "The Globotics Upheaval: Globalisation, Robotics and the Future of Work", 2019
[259] https://hbr.org/2012/09/are-you-solving-the-right-problem
[260] Yang, L., Holtz, D., Jaffe, S. et al. The effects of remote work on collaboration among information workers. Nat Hum Behav (2021). https://doi.org/10.1038/s41562-021-01196-4
[261] Remote-Work Options Can Boost Productivity and Curb Burnout, Catalyst, 2021 https://www.catalyst.org/reports/remote-work-burnout-productivity/
[262] Clancy, M. (2021). An example of successful innovation by distributed teams: academia. New Things Under the Sun. Retrieved from https://www.newthingsunderthesun.com/pub/5pua5ge3
[263] Clancy, M. (2021). Are ideas getting harder to find because of the burden of knowledge? *New Things Under the Sun*. Retrieved from https://www.newthingsunderthesun.com/pub/zsc23qxz
[264] Mathias Benedek, Martin Karstendiek, Simon M. Ceh, Roland H. Grabner, Georg Krammer, Izabela Lebuda, Paul J. Silvia, Katherine N. Cotter, Yangping Li, Weiping Hu, Khatuna Martskvishvili, James C. Kaufman, Creativity myths: Prevalence and correlates of misconceptions on creativity, Personality and Individual Differences, Volume 182, 2021 https://www.sciencedirect.com/science/article/pii/S0191886921004451
[265] https://en.wikipedia.org/wiki/Thomas_W._Malone
[266] https://en.wikipedia.org/wiki/Karl_J._Friston
[267] https://www.anilseth.com
[268] https://centerforcollectivelearning.org/
[269] Rafael Kaufmann, Pranav Gupta, Jacob Taylor, "An active inference model of collective intelligence", 2021
[270] https://www.santafe.edu/people/profile/jessica-flack
[271] Jeff Hawkins "A Thousand Brains: A New Theory of Intelligence", 2021
[272] Thomas Malone, ibid.
[273] Sources: human brain, Anil Seth; enterprise network, Microsoft Viva Insights; communities networks, Menichinelli and Schmidt; industry network, Hidalgo.
[274] https://www.harvardmagazine.com/2021/11/right-now-can-slime-molds-think
[275] https://aeon.co/essays/the-study-of-the-mind-needs-a-copernican-shift-in-perspective
[276] Veits, Marine & Khait, Itzhak & Obolski, Uri & Zinger, Eyal & Boonman, Arjan & Goldshtein, Aya & Saban, Kfir & Seltzer, Rya & Ben-Dor, Udi & Estlein, Paz & Kabat, Areej & Peretz, Dor & Ratzersdorfer, Ittai & Krylov, Slava &

Chamovitz, Daniel & Sapir, Yuval & Yovel, Yossi & Hadany, Lilach. (2019). Flowers respond to pollinator sound within minutes by increasing nectar sugar concentration. Ecology Letters. 22. 10.1111/ele.13331.

[277] Klein, S. L. Parasite manipulation of the proximate mechanisms that mediate social behaviour in vertebrates. Physiol. Behav. 79, 441–449 (2003), cited in Lewin-Epstein, O., Aharonov, R. & Hadany, L. Microbes can help explain the evolution of host altruism. Nat Commun 8, 14040 (2017). https://doi.org/10.1038/ncomms14040

[278] https://en.wikipedia.org/wiki/Integrated_information_theory

[279] Integrated information as a metric for group interaction; David Engel,Thomas W. Malone; 2018; https://doi.org/10.1371/journal.pone.0205335

[280] When and Why Did Human Brains Decrease in Size? A New Change-Point Analysis and Insights From Brain Evolution in Ants; Jeremy M. DeSilva, James F. A. Traniello, Alexander G. Claxton and Luke D. Fannin; Frontiers in Ecology and Evolution

[281] Robert Wright, “Nonzero: the logic of human destiny”, 1999

[282] For more on this, see Rosen, “The most powerful idea in the world”, 2018