AI, Super Intelligence, and the Fear of Machines In Control

Brian Mullins

INTRODUCTION

The advent of Big Data is decades old, and the citadels built atop its resources have redefined the landscape, shifting the power balance away from governments and into the gray area between the public and private sectors. Regulatory systems have yet to keep pace. Power has come not so much from the collection, ownership, or acquisition of data, but more from the ability to direct them into strategic assets. The combinations of what you know and who knows what will become the next decade's most valuable commodities, with those resting on fractured and ineffective decision-making systems losing the competitive battle.

However, it's important to avoid the superstition of superintelligence, waiting for - or fearing - the day that the machines awaken and take control. The ultimate battle will not be between humans and machines. The battle will be hybrid means and those harnessing the power of true human·machine collaboration will come out on top, thereby achieving true organizational intelligence. This article addresses the foundations of organizational intelligence, and how to navigate the shifting sands and strengthen one's financial and reputational position within global power dynamics.

WHAT DO WE KNOW?

The volume of currently accessible data is unprecedented. The World Economic Forum estimates this will reach 44 zetabytes in 2020.[1] By 2025, data generated globally each day is projected to reach 463 exabytes, or 175 zettabytes in total,[2] and by 2025 there are likely to be 30 billion Internet of Things (IoT) device connections worldwide, equating to nearly four for each person on the planet.[3]

Brian Mullins is a graduate of the U.S. Merchant Marine Academy. He is an entrepreneur and technical leader with over a decade of experience in high growth technology companies as CEO, scaling teams from just a few, to hundreds of employees across 4 countries and raising over $330 MM in investment capital. Serving on the frontiers of technology, he has been awarded over 100 U.S. patents, has testified as an expert before the U.S. Senate, received an Edison award for Industrial Design, been named a CNBC Disruptor 50, and was one of Goldman Sachs "100 Most Intriguing Entrepreneurs."

In 2019, Brian joined Mind Foundry, an Oxford University company as CEO. Mind Foundry specializes in high stakes applications of AI across the sectors of Insurance, Government, Security and Defense.

This trend continues to grow extraordinarily and there is no denying that our obsession with data capture in the quest for insight has changed entire industries. To clarify, data itself is not information. Moreover, information is not intelligence. More data does not translate directly into better decisions, and this is the first myth of control.

The first myth of control: more data equals more knowledge

Historically, lack of information has driven uncertainty. We have moved from a world replete with ignorance, to a world saturated with information but still lacking in evidentiary support for decision-making. We are now blessed with an excess of data, but the quest to capture as much information as possible has now found itself at the feet of intelligence. Navigating this dense data landscape has provided us with architectures, technologies and even entire industries dedicated to the pursuit of insight and intelligence. Assessing what you know, and with what certainty have become key beacons of the information age.

Artificial intelligence itself has risen in prominence with the ever-increasing ability to crunch and then translate large amounts of raw data into information-rich assets. While data is created, analyzed, and sometimes tortured to extract its perceived full potential, machines are deployed to sift through mountains of available data to refine more valuable assets. Data as a commodity has been compared to oil.[4] The more data you own, control, and use, the more power you have. Instead of sitting on top of an oil well, you are sitting on top of an infinite well of data at one's disposal. But the analogy stops here. Oil is finite, yet data is - for all intents and purposes - infinite. Oil is single-use, while data can be created and re-used, and exist in many different forms. You are sitting on, and continuously collecting, an asset whose potential, given the right conditions, can increase in size and value over time.

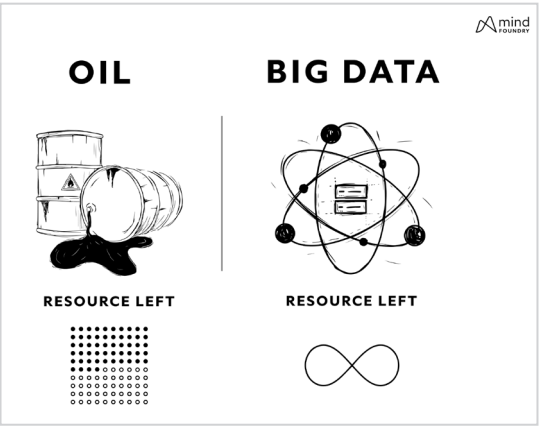

Figure 1. Oil and Big Data Resources. Mind Foundry 2021.

One's ability to process, refine, curate and sustain assets helps define success in this new age. Maintaining evidentiary chains of custody for these assets and capturing their use in decision-making are foundational tools needed by all. This in turn has led us from ignorance being understood as 'not knowing', to now being defined as 'not knowing enough.'

The Myth of Superintelligence

In the pursuit of knowledge, we've fallen prey to various theories of superintelligence. A common thread among these theories is that more data coupled with ever more powerful machines will produce superintelligent machines. The idea that a machine or series of machines will suddenly be endowed with sufficient data or computational power to qualify as 'superintelligent' is worse than unrealistic – it is superstitious.

The commoditization, fragmentation, and complexity of the data landscape alluded to above requires increasingly more sophistication to manage. There is reason to believe that Moore's law – the observation that computational capacity doubles roughly every two years as the size of transistors gets smaller – is becoming more inaccurate.[5],[6] What happens when the amount of computation power plateaus, dependent not solely on power, but also mineral resources for chips?[7]

Worse, we still quake so much in the shadow of impending technological singularity that we foist misdirected control systems on semi-autonomous systems in the hope of keeping them in line and buying ourselves time to stop an 'intelligence explosion.' Our focus should be on how we can rise above this fear and build a world where we can interoperate with trust.

The second myth of control: a human will always have the final say

We need to more fully understand what governance and oversight look like in this space. Outdated methods, including human-in-the-loop, are increasingly becoming inadequate fail-safes. In the face of automation, these human-speed inefficiencies will become primary automation targets, taking with them incumbent security mechanisms. In a truly hybrid domain, the

efficiencies of human and machine agents operating on collaborative tasks require more sophisticated mechanisms of oversight. The best path towards a governance framework that will optimally fit an AI-enabled workforce begins with a basic understanding of the hybrid systems – everything from their benefits and known failure modes to their interactions with each other.

SO, WHO KNOWS WHAT?

We have been too obsessed with the pursuit of data-driven knowledge – so much so that subsequent action is sometimes an afterthought. The true performance indicators of how well an organization is harnessing its incumbent knowledge are often poorly architected or understood.

The intelligence-processing domain provides a good example. As in many industries, the tsunami of data at the intelligence community's disposal exceeds human processing capability. There are established methods for data collection and connection, for disseminating intelligence reports, and for creating actionable products. There is an entire domain dedicated to the processing of intelligence data, which involves multiple layers of sanitization. This is no mean feat but generally, the field-to-field delay time is too long.

A piece of information collected yesterday that is not actionable until tomorrow – or three months from now – represents an intelligence chain that at best is sub-optimal, and at worst, broken, and undermines the effectiveness of those whose job depends on decisions and action that can exploit that information. Given the stakes, many industries simply can't afford such delays. Even where incumbent capabilities exist, efficiencies break down when the systems cannot achieve the necessary performance. From an organizational perspective, this can be the ability to act as desired within a particular time frame, or simply to act in general.

The third myth of control: once something is known, everyone knows it

That someone somewhere knows something is insufficient. Living in a connected world, we assume that once something is known that it is immediately disseminated to everyone that needs to know it, which simply is not the case, even within tightly knit organizations. Mechanisms to enable effective organizational decision-making require explicit architecture and thought. Also required is the ability to adapt rapidly as the environments around us change. Relying solely on broadcast mechanisms simply amounts to turning up the volume and letting the noise get louder. It is better and smarter to build systems that effectively get the right information to the right person at the right time.

WHERE WE GO FROM HERE

Whilst writing this article, I was lucky enough to get the thoughts of Professor Stephen Roberts, Director of the Oxford University Center for Doctoral Training in Autonomous Intelligent Machines and Systems. Our conversation touched upon how we must begin from a systemic vantage point to design systems that interface with the complexity of data and a range of

human and machine stakeholders. He commented that, “we live in an era of hyper-abundant data. However, solving the challenges we face requires more than data alone. It requires a deep understanding of the dynamic relationships found not only in the data, but also between agents and stakeholders associated with it; as generators, consumers or actors within the data-sphere. Understanding such complexity hinges upon our abilities to create robust models able not just of dissecting complex data, but capable of managing and orchestrating its flow and engagement across stakeholders, be they software, hardware or human agents.” (Professor Stephen Roberts, Personal Interview, December 2021). Thus, not only must we be able to harness and operate on the increasingly vast amounts of available data; we also need to anticipate and direct our data-collecting activities towards those most valuable. Finally, we need a way to imbue our values and principles into the systems we design and use, and to better understand what true human-machine collaboration looks like, with reliable performance indicators that will ensure immediate and long-term success.

This is not so much about learning how to harness AI as a tameable beast, but more about how to bring it into your team as a trusted and responsible member. This is best achieved by ensuring an unwavering commitment to continuous organizational learning that enables both human and synthetic agents to learn and improve as they collaborate towards defined goals.

This was touched upon further when I recently connected with Professor Mike Osborne, one of the world’s leading experts in collaborative AI technologies, from Oxford University’s Machine Learning Research Group, we discussed the necessity to embed human context in the design of artificial intelligence systems. Mike explained that, “It is misleading to think of AI today as being a like-for-like replacement for a human worker. AI today is powerful, but severely limited. Even with today’s data volumes, AI without deep human collaboration is useless at best and harmful at worst. The best data that exists is embedded in the heads of those stakeholders who best understand the problems to be solved – only if an organization designs its AI solutions with and around those stakeholders will it truly deliver value.” (Mike Osborne, Personal Interview, December 2021).

Figure 2. Organizational Intelligence. Mind Foundry 2021.

WHERE DO WE SEE THIS OPERATING AT SCALE TODAY?

Organizational intelligence is not superintelligence as defined by philosopher Nick Bostrom[9] and others. Today, the hallmarks of superintelligence are seen not in our machines, but rather, in the organizational effectiveness of some governments and major corporations. This requires a large orchestration of humans and technologies, the dissemination of information, and decision-making against common goals. This interconnected network of agents working together is the seed from which true organizational intelligence will sprout.

Taking a subset of these organizations – those with the desire and the commitment to harness the true potential of the data age – and providing them with the tools to collaborate with their AI counterparts, you will see a new kind of superintelligence evolving a truly scalable hybrid intelligence.

BARRIERS TO ORGANIZATIONAL INTELLIGENCE

Organizational intelligence is difficult to achieve, and goes well beyond simply adding AI capabilities to your toolbox or hiring lots of good data analysts and hoping for the best. Of the many impediments to the true realization of organizational intelligence, these are the top four killers:

Rigidity of organizational structures

For centuries certain outfits have enforced control by relying on archaic and rigid chain-of-command structures. While clarity in autonomy for decision-making has its benefits, examples where rigid, inflexible structures are unsuitable abound.

Retired U.S. Navy Captain and best-selling author, L. David Marquet, sees bringing decision making closer to key information as pivotally important, enabling a distributed (or federated) decision-making environment that allows for fast, informed decisions untethered, or at least less tethered, to a central chain of command. Immediate access to relevant and evidentiated intelligence should enable strong, even prescient, decisions, provided the organization can make decisions as quickly as that information is made available.

There will always be decisions that cannot and should not be outsourced, and any organizational intelligence framework must be vigilant regarding autonomy given to agents, both human and AI. Humans can adapt well to broad context and new situations. Hybrid organizations must harness this adaptability.

Innovation in high-stakes environments

There is always a complex interplay between evidence-based innovation and the new evidence that arises from that innovation, and as stakes increase, so too do barriers. To test novel technology and can afford a million failures before performance attains success, fine.

One can afford to lose a million chess games, but it is obviously unacceptable when human lives are on the line. I simply note here in passing those valid fears often dampen innovation in high-stakes environments, which often results in the continued use of archaic strategies.

Failure to adapt

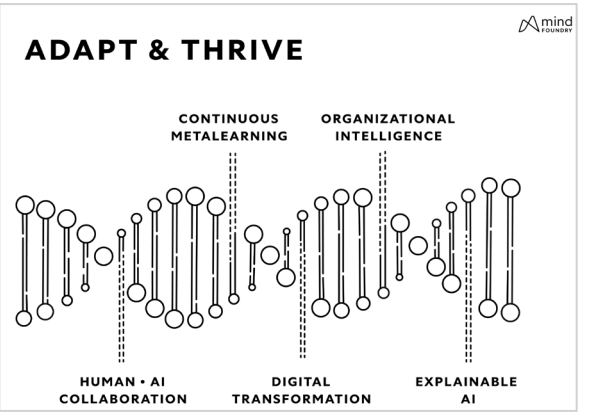

Figure 3. Adapt & Thrive. Mind Foundry 2021.

Current industry standards are insufficient to tackle the demands of the big data age head-on. Solutions designed to mitigate today's issues, standing alone, will not help us surmount future hurdles. With change as the one constant we all must live with, adaptation must always be the cornerstone for any successful organizational strategy. Data also is in a constant state of flux, as are our technological capabilities. Thus, our response must be continually dynamic.

Even the most mature AI technologies shouldn't lull us to assume that the future will resemble the past, which includes the data we now collect and hold. Ability to handle a future that looks nothing like the past grows ever more important. While one approach might be to stay at the cutting edge of new technologies, this is only a partial solution. Nimble adaptability is needed at the organizational and human level to optimize the modern-day hybrid workforce. Again, change is our only constant, so while investment in the future is crucial, adaptation needs to be part of a business's everyday operations.

THE QUESTION IS ONE OF TRUST

Each pitfall flagged above can kill organizational intelligence in its crib. So, how is this problem averted?

Trust and accountability are central to a hybrid workforce: they enable autonomous and semi-autonomous AI agents to execute their remits effectively with proper collaborative and auditable functions. Accordingly, both the human and AI agents must:

- understand their remits and their allowed space of operation,
- definitively explain their reasoning for making certain decisions,
- indelibly explain any actions they have taken based on those decisions, and
- collaborate effectively with other agents against combined problems.

THE NEW AGE

These tenets comprise the backbone of trust and accountability. Increasing amounts of autonomy should not be granted to systems that fall short of these imperatives. This is a pivotal moment: our next steps will reshape the workforce. The sheer amount of available data will continue to expand geometrically. Those who prioritize quantity and the status quo over quality, understanding, and adaptation will lose out. The new age will not find machines ruling us, or humans working within archaic organizational structures. The new age should (and will, if we do this right) find humans and machines working complementary, with the lag between data and action significantly reduced, and human-AI collaboration much better serving the interests of mankind.

NOTES

1. How much data is generated each day? WEF, April 2019. (1 zettabyte = one trillion gigabytes) (1 exabyte = one billion gigabytes)
2. The Digitisation of the World IDC, November 2018.
3. IoT Statistics, Finances Online, 2020.
4. The world's most valuable resource is no longer oil, but data, *The Economist*, May 2017.
5. We're not prepared for the end of Moore's Law, *MIT Technology Review*, February 2020.
6. Compute trends across three eras of Machine Learning, 2022.
7. Kate Crawford, Atlas of AI, 2021.
8. Expanding AI's impact with organizational learning, *MIT Sloan Management Review*, October 2020.
9. Nick Bostrom, How long before superintelligence? updated 2008.