

How Modes of Myth-Making Affect the Particulars of DS/ML Adoption in Industry

EMANUEL MOSS

CUNY Graduate Center / Data & Society

FRIEDERIKE SCHÜÜR

Cloudera Fast Forward Labs

*The successes of technology companies that rely on data to drive their business hints at the potential of data science and machine learning (DS/ML) to reshape the corporate world. However, despite the headway made by a few notable titans (e.g., Google, Amazon, Apple) and upstarts, the advances that are advertised around DS/ML have yet to be realized on a broader basis. The authors examine the tension between the spectacular image of DS/ML and the realities of applying the latest DS/ML techniques to solve industry problems. The authors discern two distinct ways, or modes, of thinking about DS/ML woven into current marketing and hype. One mode focuses on the **spectacular capabilities** of DS/ML. It expresses itself through one-off, easy-to-grasp marketable projects, such as DeepMind's AlphaGo (Zero). The other mode focuses on DS/ML's **potential to transform** industry. Hampered by an emphasis on tremendous but as of yet unrealized potential, it markets itself through comparison, in particular the introduction and adoption of electricity. To the former, data is a mere ingredient, a current, but not a necessary, requirement for the training of smart machines. To the latter, data is a fundamental enabler, a digital, always-giving resource. The authors draw on their own experiences as a data scientist and cultural anthropologist working within industry to study the impact of these modes of thinking on the adoption of DS/ML and the realization of its promise. They discuss one client engagement to highlight the consequences of each mode, and the challenges of communicating across modes.*

INTRODUCTION

In a midtown Manhattan conference room, the audience is nodding along to the presenter's slides. Artificial intelligence seems so accomplished and yet so straightforward, from Google DeepMind's Go-playing AI agent AlphaGo (and successors) and Carnegie Mellon's poker-playing Liberatus AI to Sunspring, a short film based on a movie script, replete with dialogue and stage directions, that was written by a neural network. Let two computers play Go against each other and let them learn from their mistakes until they get better than human Go grandmasters. Feed a neural network with movie scripts until it writes one of its own. The artificial intelligence industry has long been adept at foregrounding the "magic" of AI systems (Elish & boyd 2017). On that day, the audience in the conference room was comprised of employees from an entertainment media company who were identified prior to the event as key stakeholders in how the company collected, analyzed, and utilized data across the many lines of their business. They were interested in using data science and machine learning (DS/ML) for their organization and had sought the help of DS/ML experts to do so. Specifically, they were interested in a "data strategy", a set of project, people, and process recommendations designed to help them harness the potential of

DS/ML. The day began with a recognition of the many magical things that data science and machine learning (DS/ML) is capable of. Over the course of the day, conversations shifted towards practical applications of DS/ML and the conditions that allow DS/ML to succeed within organizations.

This paper grapples with the ways in which the contrasting narratives that surround the development of data science, machine learning, and artificial intelligence present different, at times seemingly opposed, paths forward as enterprises develop strategies and make tactical decisions around these emerging technologies. We identify and examine two contrasting narratives for the emergence and development of these technologies. In analyzing the myths and metaphors that attend to the discursive production of DS/ML, we follow Sturken and Thomas (2004), who observe that “metaphors about computers and the Internet are constitutive; they determine how these technologies are used, how they are understood and imagined, and the impact they have on contemporary society”. So too do these metaphors determine how businesses strategize around DS/ML.

STORIES AND MYTH-MAKING

The history of technological development is populated with spectacular demonstrations designed to hasten the development of and increase public demand for new products. The spectacular demonstrations of electrical lightning at World Fairs, Centennials, and other grand exhibitions of the late 19th century were designed to increase consumer adoption of the light bulb and serve as a (literally) shining proponent of the potential uses of electricity (Nye, 1994). Similarly, prominent players within DS/ML build and promote spectacular demonstration projects. These demonstration projects take a range of forms; they highlight an emerging capability (e.g., the capability of generating new text; natural language generation) or are engaging in ways that generate press coverage (as when a machine defeats a human expert). These demonstrations serve a range of functions; they establish their producers as serious players in the industry, they promote existing products and services offered under the same brand, they burnish the resumes of those who work on them. Primarily, they perpetuate excitement, and investment, in DS/ML.

AlphaGo Zero and the Modular Myth

AlphaGo (Silver et al. 2016) and its successor AlphaGo Zero (Silver et al. 2017) are algorithmic systems built by Google’s DeepMind that spectacularly defeated reigning human world champions of the board game Go. Go is a complex game, with millions of possible moves and billions of possible board configurations. In their release notes of AlphaGo, DeepMind foregrounded the complexity of the game itself and the remarkable achievement of building an agent that can learn to cope with that complexity from human gameplay data. In announcing AlphaGo Zero, DeepMind’s promotional materials foregrounded the ability of the algorithmic system to learn from self-play: AlphaGo Zero learns on its own by playing against itself. In the process, it learns strategies that resemble strategies of human Go players, as well as a few novel others (Silver et al., 2017). The accomplishments of AlphaGo and AlphaGo Zero appear as evidence that AlphaGo must be very intelligent since Go is commonly understood as a complex game that only the most intelligent humans can learn to play well. And while AlphaGo still needed some human help, in form of human game

playing data, AlphaGo Zero freed itself from this requirement, this dependence on human expertise and labor.

The Conditions of Success for AlphaGo (Zero)– Under scrutiny, we discover that, as Andrej Karpathy put it, “AlphaGo is a narrow AI system that can play Go and that’s it” (Karpathy 2017); the success of AlphaGo is grounded in several conditions or “conveniences” of the game Go (see Table 1).

Table 1. Conveniences of Go (adapted from Andrej Karpathy [2017])

Deterministic	The rules of Go describe possible game states without any randomness or noise.
Fully Observed	Each participant knows everything about the current state of the game simply by looking at the board.
Allows only discrete actions	There are a quantifiable number of different moves that are possible without gradations between these moves.
Is simulatable	It is easy to simulate a game of Go and this simulation will be identical to the game itself.
Is short	Each game of Go lasts approximately 200 moves.
Has a clear outcome	There is a clear definition of what constitutes a ‘win’ or ‘loss’.
Is well-documented	There are hundreds of examples of human gameplay to supercharge the initial knowledge that AlphaGo begins learning from (AlphaGo Zero, of course, freed itself from this condition).

Few ‘real-world’ problems, problems that one is likely to encounter in industry, share these conveniences with the game Go. Real-world problems are full of imperfect information, vaguely defined in terms of a success metric, rare enough that trainable examples are hard to come by, or they involve continuous phenomena rather than discrete moves that allow for gradations between possible states. Arguably, most real-world problems are more complex than the game of Go (see also, Elish & boyd 2017); in DeepMind’s promotional material and the paper detailing the algorithms that power AlphaGo (Silver et al., 2016) and AlphaGo Zero (Silver et al., 2017) complexity is defined in terms of combinatorics, the number of possible board configurations, a narrow definition of complexity. Furthermore, most real-world problems are only simulatable through deliberate decisions about what is and is not part of a system that do not come close enough to approximating reality to be good representations of the problem at hand. The weather does not affect the outcome of a game of Go, yet it is likely to be relevant for algorithms that steer self-driving cars; most real-world problems tie into dynamics part of the world that require us to make decisions about what is relevant and what is not when we model the system. As Karpathy concludes his analysis of AlphaGo (and AlphaGo Zero), it demonstrates not so much the power of DS/ML, but rather shrewdness in choosing a tractable yet impressive problem, as well as the power of Google to devote its resources to create a system that can tackle such a difficult, if singular, problem (Karpathy 2017).

AlphaGo and AlphaGo Zero are only two of the more recent spectacular projects, of course, that demonstrate the dramatic promise of DS/ML. Carnegie Mellon's poker-playing Liberatus AI beat humans in games of Texas Hold 'Em, for example.

The Modular, Bolt-On View of DS/ML – Spectacular demonstrations of emerging DS/ML capabilities solve isolated and isolatable challenges without recognition of the conditions critical to the success of these demonstration. They encourage a modular, bolt-on view of DS/ML. This view encourages us to see DS/ML as an add-on with high interoperability: conditions do not matter. The view suggests that DS/ML can be added to existing software without reconfiguration; existing processes can be complemented, or augmented, by DS/ML without transformation. The modular, bolt-on view of DS/ML suggests that DS/ML can be deployed as a layer that sits atop or replaces existing products and processes, as neatly as desktop word processors seemed to replace typewriters.

But, even in the transition from typewriter to word processor, problems of translation required halting and stepwise adjustments from one to the other. The dot-matrix printer, TrueType font libraries, and skeuomorphic user interfaces (Laurel 2013) all filled in the gaps between how people had designed their work processes around the typewriter and the new, unique affordances of desktop publishing. The modular bolt-on view of DS/ML draws attention away from the very particular conditions that enable DS/ML successes as well as the expertise and labor that is required to conceive of, develop, and refine these technologies over time and often over many rounds of trial and error; this is at the heart of what we call the modular, bolt-on view of DS/ML. Spectacular demonstrations have broad appeal because of the human tendency to misunderstand what constitutes a computationally difficult problem, to see proof of technological capability as proof of pragmatic capability. Arguably, this human tendency is exploited in the choice of demonstration project to achieve reach and effect. AlphaGo Zero, for example, is a dramatic proof of concept of reinforcement learning (Silver et al. 2017). But, it is not a persuasive proof that reinforcement learning can accomplish tasks we as humans understand to be on the same order of complexity as the game of Go. Indeed, human intuitions for what are easy or difficult problems to solve do not map on to computational difficulty. This problem has long bedeviled AI researchers who have struggled to explain, for example, just how hard comprehension tasks are for computers, when they seem so 'easy' to humans. This human tendency to misunderstand what constitutes a computationally difficult problem allows spectacular demonstrations, like AlphaGo Zero, to recast a range of problems that seem 'easier' from the perspective of human intelligence as within grasp of being solved by the technologies that are showcased by the demonstration project.

DS/ML as the New Electricity

Because of the prominence of DS/ML as modular and bolt-on, exceptions to this narrative are worth examining. One such narrative that stands in exception to the modular addition of capabilities story is told by Andrew Ng, formerly of Baidu, now Co-Chairman of Coursera and an Adjunct Professor at Stanford University. He describes the challenges of adopting DS/ML as an emergent technology in terms of the challenges that faced industry around the turn of the 20th Century as the emergent technology of electricity began to replace steam power. At that time, electricity was far from the omnipresent and almost invisible

commodity that it is today. Few aspects of the technology had been standardized, from voltages to outlet plug shapes, and ensuring that a new investment in electrification would pay off was far from certain. In Ng's telling, *"a hundred years ago, electricity was really complicated. You had to choose between AC and DC power, different voltages, different levels of reliability, pricing, and so on. And it was hard to figure out how to use electricity: should you focus on building electric lights? Or replace your gas turbine with an electric motor?"* According to Andrew Ng, *"thus many companies hired a VP of Electricity"* (Ng 2016).

Similarly, DS/ML is, today, "really complicated". Data can be local or distributed in the cloud. It is difficult to know whether and why to use a random forest algorithm or a neural network, or how to evaluate the success of any particular implementation. Furthermore, it is difficult to anticipate the costs of a project; the reliability and cost of machines, data storage, and engineering talent vary widely. And it is difficult to know where to focus one's efforts; should one build an audience segmentation model first or a churn model?

While most commentators gloss Ng's story under breathless headlines like "Artificial Intelligence Is the New Electricity?" (Eckert 2016), the story that Ng tells is more nuanced than one of simple metaphor-making when one focuses on the importance of the "VP of Electricity" to Ng's narrative. It is also more nuanced than the modular addition of capabilities. Through this lens, his story is one of complexity in emerging technologies that requires dedicated expertise to construct new interfaces that mediate between the different needs of the different parts of an organization. For DS/ML, this means preparing data in a way that is easily ingestible, and constructing tools that simplify the underlying complexity but offer affordances for making use of tools that had been previously beyond the reach of non-experts. AlphaGo Zero is presented by its creators as a persuasive proof of reinforcement learning and its capacity to solve 'complex' problems. However, reinforcement learning works best on problems that have been adequately abstracted to sufficiently resemble the kind of closed problems that reinforcement learning can solve. That is, real-world problems must be made sufficiently deterministic, observable, discrete, simulatable, short, evaluable, and well-documented before they can be addressed by the emergent technologies embedded in AlphaGo Zero, as Karpathy points out above. Furthermore, these emergent technologies must also be reshaped to accommodate real-world problems, even in their abstracted conditions, as inputs. This work of abstraction and accommodation is drastically different than the work of software development attuned to understanding DS/ML as the modular addition of capabilities. And yet, most promotional materials for DS/ML tend to further this narrative, evoking the sense of magic that Elish and Boyd identify in their work (2017). In business settings, these narratives fulfil specific functions; modular capabilities are easier to sell as products, and they are easier to explain to customers as discrete technologies. Furthermore, they lend themselves to the very same spectacular demonstrations that we have discussed above. These spectacular demonstrations are countered by Ng's metaphor, which argues that until DS/ML can be utilized as easily as a lamp can be plugged into a standard wall outlet, a dedicated form of expertise will be required to make it have any value for an organization.

CONSEQUENCES FOR PATTERNS OF ADOPTION

Metaphors matter, they guide adoption of emerging technology (Sturken & Thomas, 2004). And, they shaped how the audience of stakeholders communicated with the DS/ML expert

consultants that had gathered in that Midtown Manhattan conference room. Those stakeholders and expert consultants were gathered to develop and implement a data strategy (see above) for the entertainment media client. The goal of a data strategy is to help companies realize the potential, and potential value, of their data for their organization. Recently, over the past couple years or so, companies have started offering consulting services to help craft such data strategies, responding to a need in the market thereby acknowledging the difficulty of translating the spectacular successes of DS/ML into industry applications, from Amazon's ML Solutions Lab¹ and Google's ML Advanced Solutions Lab² to the startup Element AI³ (to mention the more prominent players).

As a producer of original content, from written text to short-form video, for a variety of different audiences, the client was interested in natural language generation, from de novo generation of content, from text to video, to the automatic tailoring of existing content to appeal to different sets audiences. In addition, they were interested in internal-facing chat bots to increase operational efficiency (e.g., a bot that suggests to re-publish existing content). This set of projects, while feasible, suggests a modular view of DS/ML. Furthermore, in preliminary meetings prior to the workshop, there was little to no concern for the conditions that allow DS/ML to succeed within organizations, yet another hallmark of the modular view of DS ML.

Over the course of the day, the consultants met with business stakeholders, content creators, software and data engineers, and data analysts. They started the onsite with a presentation on data science, machine learning, and artificial intelligence (AI) designed, one the one hand, to define a common language, and one the other, to set realistic expectations for what can be accomplished with DS/ML and the work it takes to achieve these possible accomplishments. In particular, the presentation was designed to create awareness for the conditions that need to be created for a long-term, successful, in-house DS/ML practice that could develop text generation algorithms and smart bots for internal efficiency.

Throughout this paper, we use the following definitions of data analytics, data science, machine learning, and artificial intelligence. We do not claim that our definitions are better than their alternatives, there are many competing definitions, in part because definitions suggest and drive the particulars of the adoption of new technology. Our definition of AI, for example, sidesteps a thorny issue (the definition of "intelligence", which is highly political). Here, we merely define our use of terms for the purposes of this paper to avoid confusion.

The Destruction of Modular Myths

The presentation defined, first, terms such as data analytics, data science, machine learning, and AI. There is a lot of confusion about these terms, in part, because their definition is shifting. Looking to attract talent, companies have started rebranding their data analyst positions as data science positions, for example. Artificial Intelligence is a particularly confusing term; founded as an academic discipline in the 1950s, it has been rebranded several times over the past decades with an emphasis on goals (mimicking intelligence

Table 2. Definition of Terms

Data analytics	Data analytics is the craft of counting. Data analysts count “daily active users”, for example, to inform the business about its performance. In doing so, they make use of descriptive statistics (medians, means, variation, etc.).
Data science	Data science is the craft of making predictions using and surfacing patterns from data. Data scientists use machine learning, from supervised (e.g., classification) to unsupervised (e.g., clustering) techniques in addition to descriptive statistics.
Machine learning	Machine learning is a set of tools, from supervised (e.g., classification) to unsupervised (e.g., clustering) techniques including techniques such as deep learning and reinforcement learning.
Artificial intelligence	Artificial intelligence denotes a set of capabilities or behaviors, from object recognition to goal-oriented decision making to (natural, human) language understanding and generation, that appear, to an observer, to demonstrate some kind of intelligence. Generally, these capabilities are displayed, and behaviors performed, by systems that take a set of inputs and produce outputs guided by internal states, a kind of memory.

behavior), tools (machine learning, logic, etc), or as what is just outside the grasp of current technological capability. The consultants introduced data analytics as “the craft of counting”, data science as “the craft of making predictions using data and surfacing patterns from data”, and machine learning as “a set of tools used by data scientists” to yield insights and to contribute to products that may display (seemingly) smart behavior. To the client, they suggested to leave artificial intelligence out of the day’s conversations, to focus on data analytics, data science, and machine learning, to limit potential for confusion (see Table 2).

The consultants proceeded with a review of popular, celebrated accomplishments in the field of machine learning including AlphaGo, AlphaGo Zero, Sunspring, etc.. The presentation was designed to refer to accomplishments in the field that some audience members may have heard about to first introduce the reasons why there is much excitement in the field of DS/ML. Second, they introduced these examples to then explain, at the high level, the technologies that enabled these feats, the limitations of these technologies, and the conditions they need to work (seamlessly). In Sunspring, for example, many of the protagonists express lack of knowing: “*I don’t know.*”, “*I don’t know what you’re talking about.*”, “*What do you mean?*”. The consultants explained how the algorithms that underlie the Sunspring movie script, written by a computer trained on movie scripts, led to these kind of patterns. The key intention of this was to highlight the conditions and circumstances that allow these algorithms to succeed, and consequently, the limits within which they can successfully operate: the consultants used the spectacular feats of DS/ML, and respectfully deconstructed them, to guide the client towards a view of DS/ML as an emergent capability that requires expertise to being into new business contexts. While the presentation was well received, it did not have the intended effect, as we found out later and discuss below.

Empowering Data Teams

The consultants talked to the clients in-house data analytics and data engineering team, two data analysts and one data engineer. They were joined by their current, interim manager (who did not have a data science background). In conversation, it became apparent that the data analytics team was overwhelmed by creating reports requested by the business or creatives on the performance of the business or content. Requests were fulfilled in an ad hoc manner, each one custom based on the specifics of the request. The data engineer worked on making data accessible where needed to satisfy requests. The data analysts were eager to develop self-serve approaches, dashboards that could communicate to the business performance metrics on demand, however, ad-hoc requests took priority and occupied the majority of their time: there was little to no time to build this functionality.

This situation is not uncommon. Data analytics teams tend to struggle to handle their workloads often due to the very specific nature of the requests they are asked to handle and short timelines. To remedy the situation, data analytics teams need to log and monitor incoming requests to identify common themes. They then can build self-serve dashboard for on-demand delivery of data insights around those common themes that will cover a range of frequently asked questions. In doing so, data analytics teams, tasked, due to their function, “to count” need to define “what to count”? They need to answer questions such as “What is a daily active user?” or “For how long does someone need to visit a website, watch a video, or interact with content to qualify as a content consumer?”

Within organizations, there tends to be a variety of definitions of terms such as daily active user or content consumer. Often, differences go unrecognized and unacknowledged. They surface when the data analytics team is tasked to count: they need to translate daily active user into a set of instructions (e.g., a SQL query) that demands specificity. Lacking specificity, data analysts tend to borrow details from their own, sometimes idiosyncratic, definitions of these terms. This practice has several consequences. First, asked to count daily active users, different data analysts tend to produce different answers. What is more, even the same data analyst may give different answers depending on the definition, if available, of daily active user passed on by the stakeholder. This discrepancy in answers is, at many companies, gradually eroding trust in data. Second, idiosyncratic definitions prevent data analysts to build on-demand, self-serve dashboards and other tools. At the extreme, every request becomes custom because every request demands a different way of counting a similar, often seemingly same, concept.

To remedy the situation, data analysts need to be empowered, in collaboration with the business, to define what to count. As they receive requests, they are in the best position to record definitions in current use and to consolidate definitions. To do so, they need to set aside time to work on recording requests and consolidating terms. Working with the client, the consultants received significant pushback to these suggestions, despite a clear opportunity to consolidate terms (it was suggested and requested by members of editorial and creative). According to the client, the core function of the data analytics team was to respond to ad hoc requests first, not to define or redefine them, and then, as time permits, to build on-demand, self-serve tools. There was lack of recognitions of the impossibility of accomplishing the latter task without a say in the consolidation of terms. Disempowered to establish the conditions for their own long-term success, the data analytics team was seen as a mere service function to the detriment of the organization.

Understanding DS/ML as modular, bolt-on solutions de-emphasize the importance of data readiness and interferes with deriving value from data for increased efficiency or novel products. By contrast, understanding DS/ML as an emergent capability emphasizes that a robust in-house data analytics capability is the foundation for successful in-house DS/ML projects and products; it portrays data as a resource. Like any resource, data needs to be harvested and managed. Data analysts interact with the data, build an understanding of the data, in counting they establish concepts, such as daily active user, that find use often as labels in the predictive algorithms of data scientists and machine learning engineers: e.g., the success of a piece of content may be measured in how many content consumers it attracted. In the DS/ML as modular mode, the client, surprised by our suggestions, rejected them. Coming from the DS/ML as emergent technology mode, we were surprised by the client reaction. Each mode leads to a different set of expectations, suggestions, and ultimately strategy.

Successful Data Science Requires Data Analytics –The client company had let go of their only data scientist a couple month prior to our engagement after a tenure of less than one year; the data scientist had failed to make an impact. The client’s failures in data science is grounded in the their approach to data analytics. Without a robust analytics function, data science cannot succeed. Data science depends on definitions of what to count as well as data quality and access. Lacking data analytics, data science roles tend to morph into either data analytics roles, the data scientist helps fulfill ad hoc stakeholder requests or does data quality assessment, or helps build the pipelines for better data quality and access. These patterns can be exacerbated by lack of a clear distinction between a data science and a data analytics team, as was the case at the client company. Without a robust data analytics function, data science cannot succeed. In such situations, data scientists tend to leave or, as the more expensive members of the data team, are asked to leave, as happened in this case. Understanding DS/ML as modular deemphasizes data readiness, prevents data science from having an impact within organizations; it does not highlight, much less create, the conditions for successful data science within organizations.

Protection of Editorial and Creatives

In the run-up to the engagement, the consultants were advised by the technology and product side of the business to “tread lightly” so as not to upset editorial and creatives who may fear about changes in or losing their job. During conversations, they found editorial and creatives to be eager to hear about our work, solutions, and possible externally or internally facing data products; they freely talked about their work. They encountered healthy skepticism, not fear. In many ways, editorial and creatives were more receptive to our suggestions and eager for adoption than the technology and product side of the business.

Viewed as spectacle, DS/ML offers modular, bolt-on solutions to add new products or business functions or to replace existing ones; it promotes self-sufficient, mostly autonomous systems. It de-emphasizes the importance of conditions and context. It de-emphasizes the importance of data readiness and the contributions to data and data readiness by people across the organization, from the data analytics team to editorial and content creators. It paints a picture of users as collaborators with machines but on the machine’s terms. Workers are to assist the machines, to be tasked with the edge cases that

machines can't handle, providing the glue between the complex work environment and its simplified version that allows machines to succeed. This view fosters fear of replacement by machines; AlphaGo pitted the machine against the human. AlphaGo Zero excluded humans from training machines.

Viewed as an emergent capability akin to the emergence of electricity, DS/ML is a potential, fueled by its resource: data. It not only emphasizes the need to harvest and manage this resource, it encourages us to think of applications not in terms of add-or-replace model but in terms of an open-horizon model: electricity enabled humankind to build entirely new kinds of products; it gave us superpowers, in many ways. We tamed electricity, and it has enabled us to build products that to many were unimaginable prior to their invention. Our lives changed alongside these inventions, we adapted. The view of DS/ML as emergent technology emphasises the potential of DS/ML without giving it concrete form. It emphasizes that we can change, as technology changes around us by our own actions; the adoption of DS/ML becomes less of a zero sum game with winners (machines, technologists, STEM) and losers (humans, humanities).

In the DS/ML as modular mode, the technology and product side of the business were concerned about editorial and creatives and their reaction to the arrival of consultants at the company and their suggestions. There was a big difference between the expected and the actual situation. Viewed as modular, DS/ML devalues conditions and context and with its interaction with teams across companies, especially outside the technology teams, a potential explanation for this discrepancy. It devalues the importance of knowledge of teams outside technology groups, it can lead to a kind of "benevolent paternalism". Editorial and creatives, on the other hand, were aware of inefficiencies in their work and were keenly aware of what questions they would like to have answered by data. With confidence in their work, they were looking to DS/ML as an enabler, potential partner, more in line with seeing DS/ML as the new electricity; different lines of business may be more susceptible for one way of thinking about DS/ML with consequences for communication across business lines.

Vendor Strategy

For some of their DS/ML solutions, the client relied on vendors; they paid a company for a data product or DS/ML service. In one case, the client shared their data with a vendor company for product/service delivery. It was considered to be a "good deal" since the client company was not charged by the vendor (the payment, of course, is in the form of data).

The DS/ML as modular view sees data as a mere requirement for data products and DS/ML services (as AlphaGo Zero showed us, not even a necessary one). As long as you get a product in return for your data, it is a "good deal". The DS/ML as emergent technology view promotes the idea of data as a resource, an enabler. Data enables an entire suite of data products; sharing your data in exchange for one data product becomes a "bad deal" especially if data sharing enables your competition. Most vendors work with multiple organizations often in the same line of business. Data sharing, via such vendor, can remove competitive advantage that increasingly lies in data, as the DS/ML as emergent technology view emphasizes. The DS/ML as modular view deemphasizes data as a resource, a valuable asset that is best protected; it can lead to decisions with negative consequences for the competitiveness of the business in the long term.

ETHNOGRAPHIC LESSONS LEARNED

The Role of Expertise

Taking as a starting point the “simple premise that expertise is something people do rather than something they have” (Carr 2010), it becomes possible to see this case study as revealing the tensions and misunderstandings that arise from the differing sets of practices that are called upon in the shift towards DS/ML within business enterprises. The two motivating myths presented above constitute DS/ML as two different kinds of capabilities. One myth presents a modular capability that can be added instrumentally to existing practices, the other presents a transformative capability that requires the reshaping of existing business processes to new, sometimes custom, interfaces of the emerging and still unstable technology. Each of these two kinds of capabilities, then, entails a different set of interactions between data, personnel, products, and tools, and therefore a different set of practices through which expertise functions. By understanding these two myths as motivating different forms of expertise, the positioning of various actors in the case study presented above becomes more legible.

What expert practices are motivated by the modular myth? The modular myth lends itself to picking and choosing amongst instruments to be deployed, and expertise in this context would be constituted by performing knowledge about these available tools. Such performances might include conducting cost-benefit analyses on available vendor solutions, performing knowledge about the available packages and implementations, and situating DS/ML development as the stepwise incorporation of such modular tools into existing architectures. Indeed, we observed a reliance on such performances of expertise in the reaction of some in the case study presented above. And in particularly power-laden ways, this exercise of expertise was able to repress challenges posed by alternate forms of expertise (see below) by leveraging existing control of economic resources to prioritize one set of priorities (vendor solutions) over others (reorganizing the DS/ML team).

The expert practices mobilized by the modular myth also draw strength from a particular conception of objectivity mobilized by the modular myth. As a historically- and socially-constituted value, objectivity (Daston and Galison 2007) can take many forms. The modular myth contributes to a form of “mechanical objectivity” that sees human judgement as failable, whereas algorithmic systems can stand in for human actors who may introduce “bias, inefficiency, and discrimination” (Christin 2016). Trusting algorithmic systems over human actors allows those who exert control over the use of such systems to participate in this form of objectivity as a further practice of their own expertise. However, sources of mechanical objectivity, whether they be crime scene photographs or brain scans (Dumit 2004) tend towards a situation in which the products of these tools themselves require further expertise in order to be translated for lay audiences or integrated into other sociotechnical systems. The ability to do so constitutes a form of objectivity Christin identifies as “trained judgment”. It was precisely this form of objectivity that we highlighted by introducing the story of DS/ML as the “new electricity” through the transformative myth presented above. By demonstrating the ways in which trained judgment could form a “hybrid entanglement of human and machine expertise” (Christin 2016), we show that there was a great deal of human ingenuity still required to craft DS/ML solutions for the particular problems the client was facing. How to make those problems legible to the machine were

very human questions, and their outcomes were uncertain, so our best recommendations centered on empowering that form of expertise.

What expert practices are motivated by the transformative myth? The transformative myth lends itself to precisely those practices of expertise that constitute trained judgement, but a trained judgement that extends beyond that which might evaluate between several similar products offered by a vendor. These practices include engaging in forms of collaboration and experimentation that treat DS/ML not as a stable product, but as a set of open, unresolved questions from which meaningful solutions might emerge. Specifically, the expertise of a data science lead (or a VP of electricity, for that matter), is entailed by fostering different lines of communication between disciplinary silos, for example by enacting a process in which data analysts work with data scientists to craft key performance indicators that are useful for machine learning experiments. This form of expertise is also entailed by wielding economic resources to engage in experiments that may be fruitless, but also may produce useful insights or products for further development.

What other expert practices are at stake? The creative team, who in the planning stages of the on-site workshop were to be insulated from any hints that their roles could be automated, was revealed during the workshop to have their own expert practices that actually positioned them to be promising collaborators for the DS/ML team. Indeed, they were central to the business offering at the company, but they also were able to position their work as primarily valuable because they were the ones who ‘crafted’ new content for the media company. By foregrounding this aspect of their work and downplaying the routinized labor they performed, they could have pragmatic conversations about how to automate the routine work without their central expert practices being compromised. The DS/ML team could potentially be given broad latitude in building systems for the curation of past content, summarization of aggregated content, and the monitoring of dashboards without threatening the practices that constituted creative expertise.

In reflecting on “the pervasive sense that technologies transform us in irrevocable ways means that idealistic concepts of technology are always accompanied by the anxiety that they will also promote some kind of loss - loss of connectivity, of intimacy, of desire, of authenticity in some way.” (Sturken & Thomas 2004) we were surprised to realize that this was a far more active concern for those whose expertise depended on control over the technologies that were the subject of the workshop, and not those who were most central to the production of the content offered by the company. This points towards two key findings from the engagement. The first is that where there is resistance to recommendations for a move away from modular solutions and towards transformative capabilities, sensitivity to different enactments of expertise are key. Unless existing expert practices can be reshaped or otherwise adapted to the kinds of practices entailed by a focus on transformative capabilities, a defensive, dismissive, or destructive reaction is possible from those like the CTO, whose existing expertise will be subsumed by such a shift.

The second key insight is not all that different from the lessons Latour drew from examining the history of the pasteurization of France (Latour 1993). While singular inventions and modular capabilities may sometimes be identified as transformative in their own right, they are not enacted or brought to bear on the world without a broad accommodation of the social sphere to the technological apparatus, and of the technical apparatus to the existing practices within the social sphere. As little as Louis Pasteur could accomplish in France on his own, just as little could be done by any one person in the

offices of the client in our case study. Rather, ground must first be laid across the organization to accommodate the kinds of changes that any particular form of DS/ML might take. This groundwork can be done purposefully, but requires the active participation of the entire range of actors likely to be impacted by such changes. It also requires working against the emotional grain produced by spectacular demonstrations of DS/ML. The ways in which such spectacles mobilize the sublime are quite persistent, and effectively immunize against alternate understandings of the technologies as anything but modular.

The Modular Myth

The myth of the modular addition of capabilities contributes in concrete ways to the emergence of “technology” as a “hazardous concept” (Marx 1997) that refuses interaction with anyone besides experts (conditions to not matter). The hazardous concept, in Leo Marx’s analysis, is that of technology as an “singular noun” capable of acting as an agent in history. In his telling, it is the technology that affects people’s lives and reshapes the possibilities for human existence, not the field of individual actors who comprise the sociotechnical system in which technology is embedded. AlphaGo (Zero), and other spectacular demonstrations, mark the unfolding of DS/ML as a succession of particular inventions, recapitulating in miniature the sweeping narratives of human progress that are marked by key inventions — stone tools, fire, the wheel, gunpowder, semiconductors, perceptrons, reinforcement learning — that have played active roles in human history. Marx goes on to point out that narrow conceptions of technology as constituted by discrete objects like the steel plow or the steam engine are “merely one part of a complex social and institutional matrix”, that is entailed by large scale technosocial system. This understanding informs an understanding of technology as a constitutive force that shapes society as a whole, but particularly reshapes the institutions, including corporations, that are intimately bound up with developing, employing, and deploying new technologies. In the context of this paper, technology may sometimes be seen as an active force in the constitution of the corporation.

The Modular Myth and the Technological Sublime – In *American Technological Sublime*, David Nye (1996) discusses spectacular presentations of technology as participating in an experience of the sublime. The sublime, in this context, is not a “self-conscious aesthetic theory” but rather a “cultural practice of certain historical subjects” that continually produces “new sources of popular wonder and amazement”, in Nye’s analysis. The modular myths of DS/ML development gain their mythological status from the technological sublime, and considering these spectacles as such through the lens Nye provides is instructive. The sublime, in it’s Kantian, Enlightenment-era sense has both a ‘mathematical’ and ‘dynamic’ aspect. The mathematical sublime pertains to an experience of scale that produces wonder in a human subject. The Grand Canyon, the vastness of space, and the Great Wall of China all exist at scales dwarfing the normal realm of human experience, and produce, according to Kant, a sense of the mathematical sublime. The dynamic sublime is more closely associated with a sense of terror, as when a crowd gathers for a skyscraper demolition or to watch a passing storm from a safe distance.

The spectacular, modular myths of DS/ML participate in both these forms of the sublime, and indeed are key to understanding these cultural objects as spectacles. The scales

at which an algorithm may run are constantly foregrounded in promotional materials, as in a documentary about the defeat of Go champion Lee Sedol by AlphaGo, which tells us that “a game of Go has more possible configurations than there are atoms in the universe”. The number of petaflops a computer is capable of, the number of cores and GPUs brought to bear on a computational problem, the nearly infinite permutations of possible outcomes, are all made clear to an audience in order to produce a sense of the mathematical sublime. There is terror in these spectacles, too. Even setting aside the many terrifying scenarios of an “AI Apocalypse” in which machines actually attack humanity (see Dowd 2017), in many ways DS/ML mythmaking points towards a world that doesn’t need human subjects at all, self-driving cars, efficiently optimized factories, and flawless recommendation systems sketch out a world in which the human is mostly incidental. Like Niagra Falls, it will keep churning, oblivious to our existence, and that such a world is possible induces a sense of the dynamic sublime.

While these Kantian forms of the sublime are certainly at play in the modular myths of DS/ML, they are also legible as an iteration of the electrical sublime that Nye presents, in which spectacles moved beyond the realm of the natural world, and were developed specifically for celebrations of industry, nationalism, and amusement. These modular myths, as spectacle, make invisible technologies visible. Electricity was made visible through lighting displays, just as AlphaGo (Zero) makes algorithms and data streams visible, as events that pit a human master of a game against a computer: AlphaGo defeated Lee Sedol in front of a human audience.

The Myth of the “New Electricity”

A crucial point Andrew Ng’s “VP of Electricity” metaphor makes is about the complexity of emerging technologies and the necessity of expertise to adequately grapple with that complexity. Because of the siloed nature of divisions in modern corporations (Rumelt 1974), expertise is not easily distributed across an organization. Supporting DS/ML expertise in any one part of a company will not necessarily translate to other parts of that company, unless they are empowered to make changes beyond their own division. And as the DS/ML experts will not be able to influence business practices outside their own division, it becomes difficult if not impossible to transform those practices in ways that integrate well with the DS/ML projects they work on. By placing a DS/ML expert at the executive level, or by explicitly designing processes for distributing that form of expertise across existing divisions, the complexities of the emerging technology can be addressed in a coordinated, rather than piecemeal, fashion. An expert in DS/ML can approach these capabilities as resource-driven, capable of using data to transform existing products and processes in ways that a modular, bolt-on approach cannot.

The tendency of the discourse around DS/ML towards narrating the emerging technology as a modular addition of capabilities rather than as resource-driven is highlighted by the way Ng’s story was bent towards a metaphor of AI as “the new electricity” (Eckert 2016). Portraying it as such is a subtle rhetorical move that foregrounds the power of the new technology eliding the challenges that remain in building practices around it whilst pointing to future, as of yet unrealized potential. The power of electricity is readily visible to any audience that hears that “AI is the new electricity”, even if not all listeners connect AI, machine learning, and the data that drives it with the role electricity has played as a public

utility (as opposed to the private commodity data currently is). Indeed, the challenges that were present in the early days of electrification, however, have receded to the background. It has become infrastructural, visible only when it fails (Star 1999). According to Ng, DS/ML share its eventual invisibility and great power.

CONCLUSION

The algorithms that are powered by data participate in their own set of metaphors. Some are ‘intelligent’, while others are merely ‘smart’. The use of games like Chess or Go as demonstrations of DS/ML perpetuate this metaphor. Such games have long been proxies for human intelligence (Ensmenger 2012), foregrounding certain human skills like foresight, planning and concentration over others like sensitivity, compromise, or even deception. But the use of these games in AI research remain an abstraction of human cognition that fails to capture the entire gamut of human intelligence. These are distinctly human capabilities that set algorithms on an even playing field with people who may feel more threatened than enhanced by their presence in the workplace. This tension between human and machine becomes more acute when DS/ML is described as superhuman, either in terms of being hyper-rational, hyper-vigilant, or omniscient. In some cases, DS/ML is imbued with capabilities bordering on the clairvoyant, as in breathless headlines like, “Google’s AI Can Predict When A Patient Will Die” (Tangerman 2018). Framing the capabilities of DS/ML as on par with, or even as surpassing, human capabilities places it in competition with the humans who must be full participants in any integration of DS/ML into a company. However, this participation is frequently fraught due to an inadequate consideration of the “affective relationship to the product or system, that is, how someone feels about the technology at stake” (Elish & Hwang 2016).

The metaphors of big data tend to treat data as a resource from which value can be extracted. The metaphors of DS/ML tend to treat machines as somehow more than human, which is to say they have many of the strengths of humans (intelligence, anticipation) but few of the weaknesses (inattention, exhaustion). Both of these sets of metaphors elide the uncertainties inherent in the metaphors they employ. Resource extraction is not a linear processes, it involves the failure of exploratory wells, infrastructural costs to move minerals to markets, and shifting price and demand curves relative the costs of extraction. Neither is human intelligence a completely predictable process, particularly where the development of science and technology are concerned.

In this paper, we have discussed the how the prevailing stories that highlight the emergence of data science and machine learning tend towards an understanding of DS/ML as a modular capability. These stories fail to promote transformative practices that might reshape existing business problems into ones that the emerging capabilities of DS/ML can currently address. To do so would require an attention towards data not as an ingredient, but instead as a means through which other things become possible, but also requires a different set of expert practices than those that are currently incentivized by many technical teams, which was particularly true in the case study laid out above. By understanding expertise as sets of practices that can be encouraged and rewarded, rather than as an object that can be possessed by individuals (Carr 2010), we point the way towards undertaking broad shifts in overall business practices by seeking transformative changes that are not siloed within individual departments, but rather have the opportunity to reshape existing practices broadly

in pursuit of interfaces that match the underlying technical capacities of DS/ML with the specific, measurable business needs of an organization.

NOTES

Acknowledgements - The authors would like to thank Cloudera Fast Forward Labs; without their support this work would have been impossible. All conclusions represent the work of the authors, and should not be interpreted to represent the position of Cloudera Fast Forward Labs or any of its employees or officers. The authors would also like to thank Dawn Nafus for her generous notes and Jan Philipp Balthasar Müller for reading a late draft. Emanuel Moss is grateful to the CUNY Graduate Center, Data & Society Research Institute, and the Wenner Gren Foundation for their support.

1. <https://aws.amazon.com/ml-solutions-lab/>

2. <https://cloud.google.com/ai/>

3. <https://www.elementai.com/>

REFERENCES CITED

Carr, E. Summerson

2010 "Enactments of Expertise." *Annual Review of Anthropology* 39 (1): 17–32.

Christin, Angele

2016 "From Daguerreotypes to Algorithms: Machines, Expertise, and Three Forms of Objectivity." *ACM Computers & Society* 46 (1): 27–32.

Collobert, Ronan, Jason Weston, Léon Bottou, Michael Karlen, Koray Kavukcuoglu, and Pavel Kuksa

2011 "Natural Language Processing (Almost) from Scratch." *Journal of Machine Learning Research* 12: 2493–2537.

Daston, Lorraine, and Peter Galison

2007 *Objectivity*. Brooklyn, NY: Zone Books.

Dowd, M

2017 "Elon Musk's Billion-Dollar Crusade to Stop the A.I. Apocalypse." *Vanity Fair*. April.

Dumit, Joseph

2004 *Picturing Personhood: Brain Scans and Biomedical Identity*. Princeton, NJ: Princeton University Press.

Eckert, Daniel

2016 "Artificial Intelligence Is the New Electricity?" *Vox Creative*.

Elish, M. C., and danah boyd

2017 "Situating Methods in the Magic of Big Data and AI." *Communication Monographs*, 1–24.

Elish, M.C., and Tim Hwang

2016 "An AI Pattern Language." *Data and Society*. <http://autonomy.datasociety.net/patternlanguage/>.

Ensmenger, Nathan

2012 "Is Chess the *Drosophila* of Artificial Intelligence? A Social History of an Algorithm." *Social Studies of Science* 42 (1): 5–30.

- Fei-Fei, Li, and Pietro Perona
2005 "A Bayesian Hierarchical Model for Learning Natural Scene Categories." 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05) 2: 524–31.
- Ford, Heather
2014 "Big Data and Small: Collaborations between Ethnographers and Data Scientists." *Big Data & Society* 1 (2).
- Karpathy, Andrej
2017 "AlphaGo, in Context." Medium.
- Koch, Gregory
2015 "Siamese Neural Networks for One-Shot Image Recognition." University of Toronto.
- Latour, Bruno
1993 *The Pasteurization of France*. Cambridge, MA: Harvard University Press.
- Laurel, Brenda
2013 *Computers as Theatre*. New York: Addison-Wesley Publishing Company.
- Marx, Leo
1997 "Technology: The Emergence of a Hazardous Concept." *Social Research* 64 (3): 965–88.
- Mearian, Lucas
2016 "Google's DeepMind A.I. Can Slash Data Center Power Use 40%." *Computerworld*.
- Moravčík, Matej, Martin Schmid, Neil Burch, Viliam Lisý, Dustin Morrill, Nolan Bard, Trevor Davis, Kevin Waugh, Michael Johanson, and Michael Bowling
2017 "DeepStack: Expert-Level Artificial Intelligence in Heads-up No-Limit Poker." *Science* 356 (6337): 508–13.
- Ng, Andrew
2016 "Hiring Your First Chief AI Officer." *Harvard Business Review*. <https://hbr.org/2016/11/hiring-your-first-chief-ai-officer>.
- Nye, David E.
1996 *American Technological Sublime*. Cambridge, MA: MIT Press.
- Popper, Karl
2005 *The Logic of Scientific Discovery*. New York: Routledge.
- Rumelt, Richard P.
1974 *Strategy, Structure and Economic Performance*. Ann Arbor: University of Michigan Press.
- Silver, David, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, et al.
2017 "Mastering the Game of Go without Human Knowledge." *Nature* 550 (7676). Nature Publishing Group: 354–59.
- Star, S. L.
1999 "The Ethnography of Infrastructure." *American Behavioral Scientist* 43 (3): 377–91.