James Drogan
203-829-3172 • jim.drogan@gmail.com • jmsdrgn.squarespace.com

AI: Questions

James Drogan
8/28/21[1]

[1] This document was first published April 1, 2018.

Overview

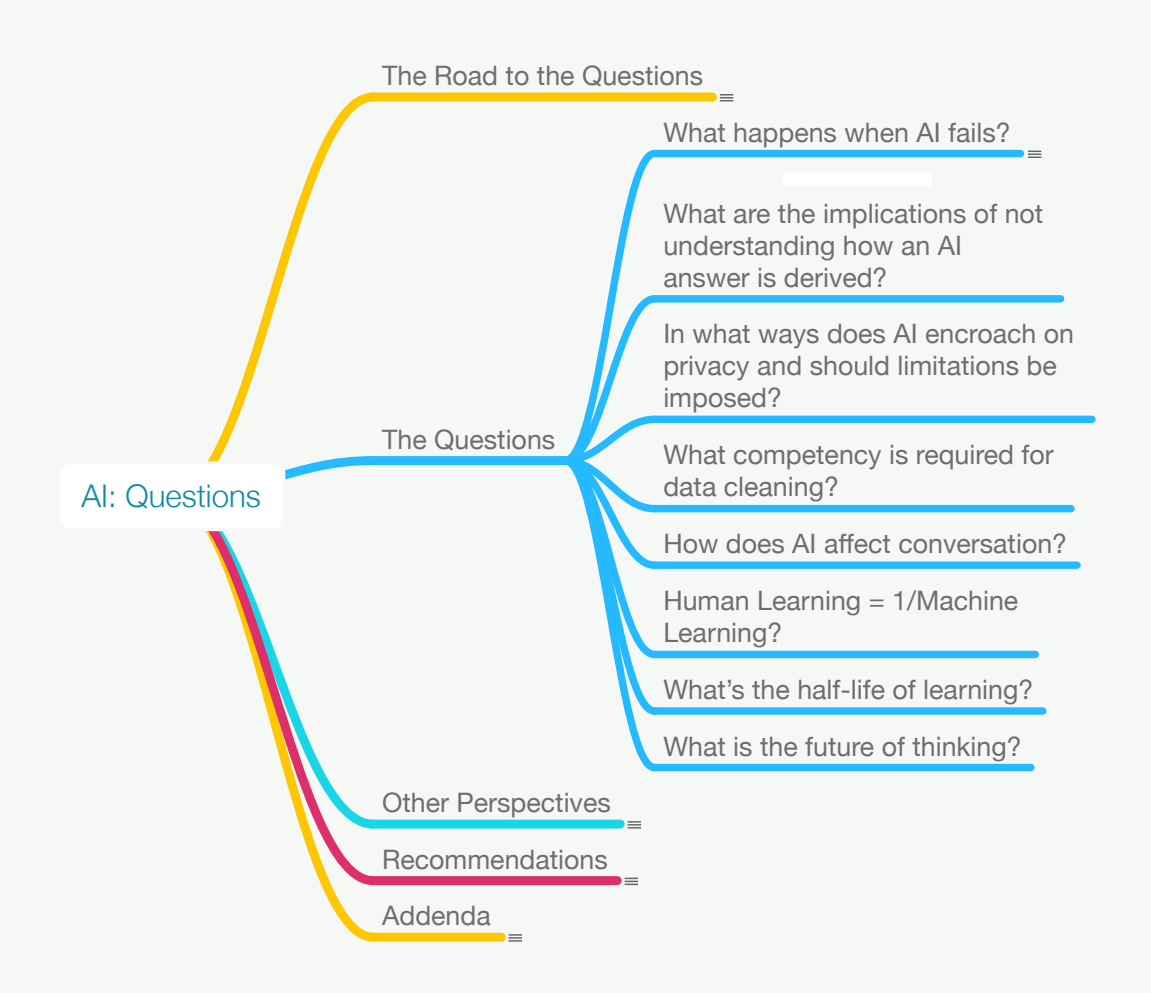

The genesis for this paper is was a meetup on the idea of Building Artificial Intelligence Teams to Transform Business held in Manhattan on September 14, 2017 and sponsored by Bots and Artificial Intelligence.[2]

This provoked my initial interest in writing to explore concerns I had about the subject. I had also been doing a bit of reading on the subject.

The second reason for writing was to provide useful material to students, especially those in my graduate course in information management.

What follows, as the development of the field demands, is work in progress.

[2] https://www.meetup.com/Bots-and-Artificial-Intelligence/

Questions

The questions are more-or-less presented in the order they occurred to me. There is a relationship between these questions that suggests that asking the questions in a different sequence may well lead to arriving at higher quality, perhaps different.

What happens when AI fails?

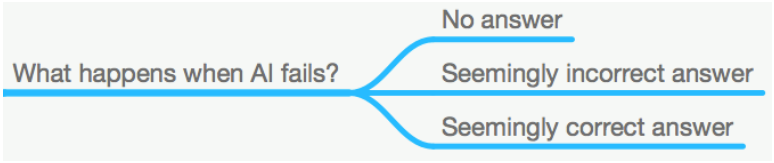

Figure 1

I am prompted towards "Seemingly incorrect answer" by an article on Alpha Go (Wong & Wong, 2016). Here is an excerpt from the article that piqued my thinking.

> Lee was doing well, and the two were engaged in a skirmish on the left side of the board. But AlphaGo, playing with the white stones, suddenly attacked deep inside Lee-controlled territory on the right side.
>
> "This was totally in the black area," Li said. "Human players would never think about doing that."
>
> Lee responded, quickly capturing three of AlphaGo's stones. It was a poor move by AlphaGo, or so it seemed.
>
> Twenty moves later, AlphaGo had taken three of Lee's stones in the upper right and occupied about half the area that most human observers had written off as impregnable. Sacrificing three stones turned out to be a key pivot, turning the game in AlphaGo's favor.
>
> "Even in black's area, white got a result. It's unacceptable for black," Li said. "There are huge variations in a Go game, we can't even read 1% of them. We have certain patterns in our minds when we play, so this is the kind of move we would never think about."

In the long run, what seemed to be "a poor move by AlphaGo" resulting in "turning the game in AlphaGo's favor."

Suppose, however, the opposite. A move (answer) that seems good, but in the long run leads to poor consequences.

Go is marked, as I understand it, by a few immutable rules, but many combinations of patterns ("Go and Mathematics," 2017).

International business, on the other hand, is marked by innumerable, constantly changing explicit and implicit regimes and any combinations of patterns. Whether the business combinations are more or less than the combinations associated with Go seems, at least at this point, moot. One is at a disadvantage of one can't imagine workable patterns that others can. Perhaps more importantly, there may not be an understanding of and agreement on the objective.

What are the implications of not understanding how an AI answer is derived?

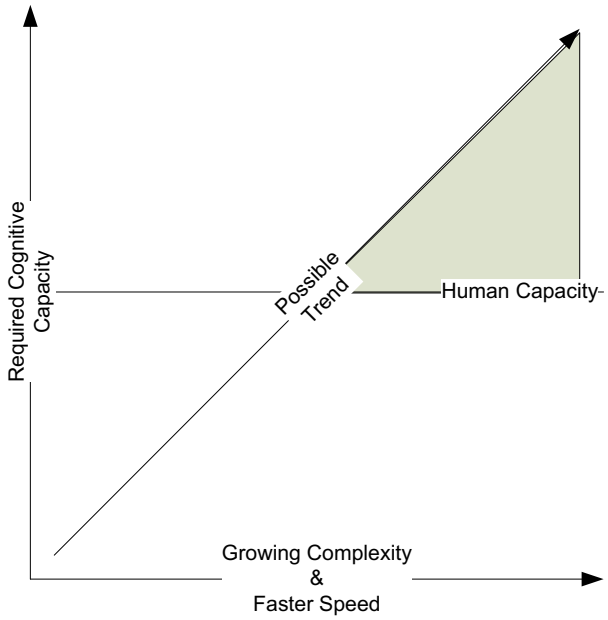

Figure 2

The University of Alberta defines cognitive capacity as, "…the total amount of information the brain is capable of retaining at any particular moment. This amount is finite, so we can say our total capacity is only ever 100%".[3] I interpret *information* to mean what one knows and includes facts and the means of reasoning about these facts. For example, we know numbers and how to they may be manipulated.

While "…total capacity is only ever 100%..." the capacity is assumed to change over time. For example, the capacity of a balloon stays at 100%, but the capacity of the balloon changes depending upon the air pressure inside and out, and the manner in which the balloon is constructed.

The Dana Foundation reports:

> Mental abilities change throughout life, first as a result of brain maturation and later with aging of brain cells and their billions of complex interconnections. As people age, their movements and reflexes slow and their hearing and vision weaken. Until the 1990s, most aging research examined cognitive abilities of adults younger than 80. More recent research includes the fast-growing 80s-and-older population and has advanced our understanding of cognitive changes in the elderly. Scientists in a recent study asked, "When does cognitive functioning peak?" and found evidence for considerable variability in the age at which cognitive abilities peak and decline throughout life.[4]

For our purposes we need to qualify cognitive capacity with the word "relevant."

Figure 2 is meant to suggest a world of growing complexity and speed of change requiring increasing cognitive capability in order to cope with, let alone manage, the change.

[3] https://sites.educ.ualberta.ca/staff/olenka.bilash/Best%20of%20Bilash/cognitive%20capacity.html

[4] http://www.dana.org/Cerebrum/2015/Cognitive_Skills_and_the_Aging_Brain__What_to_Expect/

The hypothesis is that there are limits to human cognitive capacity. We do sleep.[5] AI does not sleep. And, if need be, one simply adds another node. That is, AI can go beyond the limits of human cognition (the triangle at the upper right in Figure 2).[6]

What are the implications of contending with a cognitive capacity greater than one's own? Especially if "No one really knows how the most advanced algorithms do what they do. That could be a problem" (Knight, 2017).

In what ways does AI encroach on privacy and should limitations be imposed?

See:

1. China's Powerful Surveillance System, https://www.facebook.com/video/embed/async/dialog/?url=https%3A%2F%2Fwww.facebook.com%2Fquartznews%2Fvideos%2F1810685592298468%2F

There are two sub-questions here:

1. My observation of our species is that we tend not to welcome limitations, seeing them more as a hurdle to be crossed than something we should respect.
2. Consider Figure 2 again. If the human cannot operate in the triangle then the application of limitations becomes problematical.

I appreciate Asimov's Three Laws of Robotics:

1. A **robot** may not injure a human being or, through inaction, allow a human being to come to harm.
2. A **robot** must obey orders given it by human beings except where such orders would conflict with the First **Law**.
3. A **robot** must protect its own existence as long as such protection does not conflict with the First or Second **Law**.

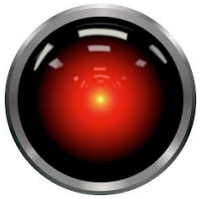

Figure 3

[5] However, I'm pretty sure that in my case there are "background processes" at work. I often awake with a clearer appreciation for a problem I'm working.

[6] This prompts me to wonder where in the AI life cycle (to be defined) is cognition (requires a more detailed definition) most effectively applied.

What competency is required for data cleaning?

The Facebook group *Applied Artificial Intelligence & Machine Learning For Business* brought to my attention[7] *Data Needs Cleaning Before Machine Learning Can Find Meaning* ("Data Needs Cleaning Before Machine Learning Can Find Meaning," 2018).

This item is the basis for some comments and questions. These are aligned with statements from the Visionnet Systems web page.

> *Enterprises large and small need to use all means at their disposal to gain a strategic advantage for their business. To remain competitive, these organizations must make use of contemporary machine learning tools to unlock value and meaning from the wealth of data they have about their customers, products, employees, work processes, and even competitors.*

Context is important and is critical to selecting the data relevant to producing the desired outcome (Drogan, 2009).

> *As is so often the case, the answer lies in numbers. Machine learning algorithms crunch numbers, so the transactional and analytical data that enterprises have in their databases and data warehouses needs to be selected and prepared to be fed into these algorithms. There are several data preprocessing steps, from data imputation that addresses data sparseness to techniques for normalizing and standardizing data, allowing the appropriate machine learning model to be applied.*

Einstein reminds us that, "Not everything that counts can be counted; and not everything that can be counted counts."

> *Common techniques for data imputation provide default values, such as 0 where none exist, or correct for erroneous and outlier data, thereby reducing "data noise".*

There are essential preparatory steps to the data imputation process (Drogan, 2009).

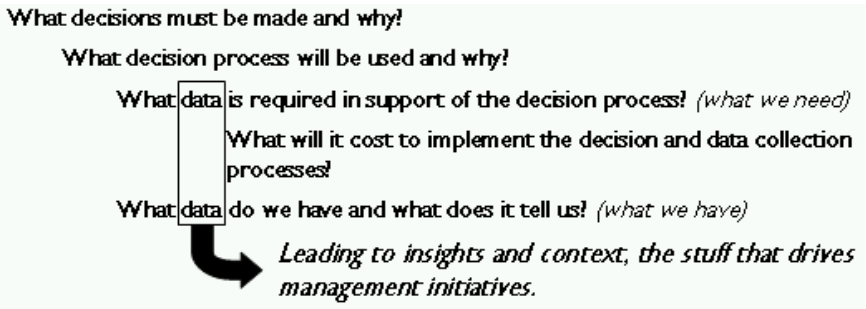

Figure 4

This notion of data cleaning must be continually checked against the goals and objectives of AI.[8]

> *Other data normalization techniques involve reducing the amount of data that needs to be presented to the machine learning model so that the models can efficiently process the data and, more importantly, only consider the data that falls within the business domain of the question or problem being solved.*

Visionet seems to recognize the importance of context. Yet, we have this next-to-last sentence on the web page.

---

[7] July 5, 2018

[8] We might, of course, simply want to go explore data to see what's there. We might not want to clean that data for fear of grubbing away something essential. Another paper, *What Can I Do with This Data?*, is intended to explore this question.

All of this work requires a trained data scientist with a keen understanding of statistical and mathematical data manipulation techniques, technical skills in data processing tools like Python, R, sklearn, and pandas libraries, and in case of big data tools like Apache Spark, the ability to quickly prepare the data set and select relevant machine learning models for enterprise data.

What happened to the recognition of the importance of understanding the context? The data scientist needs to have or employ more "keen understanding" than is mentioned in this list.

This item also raises another interesting question. Consider the following.

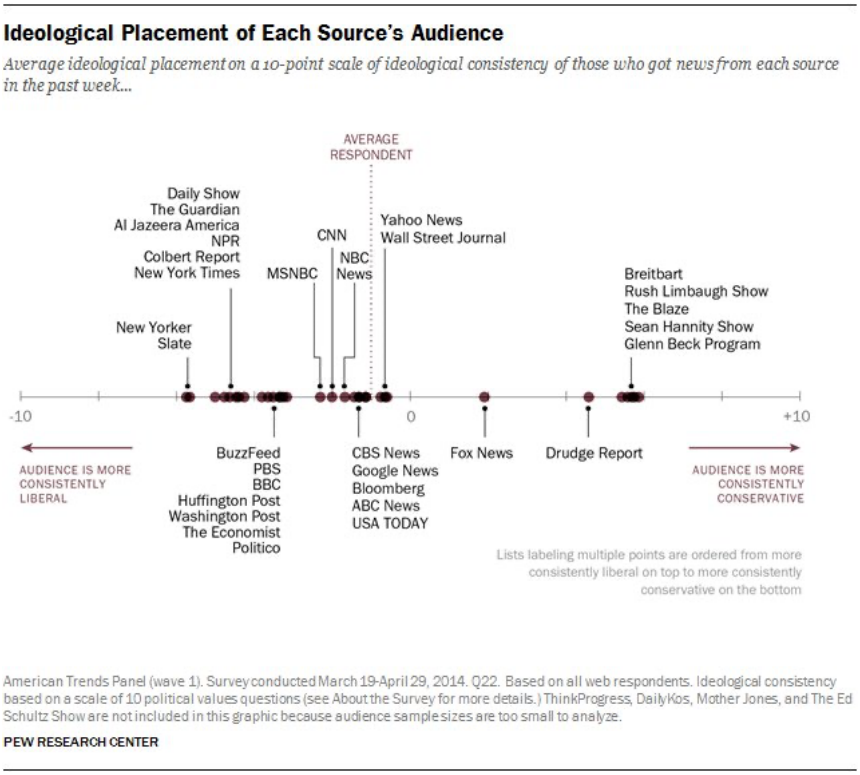

Figure 5

Where on this spectrum of values one samples makes a difference to the outcome. We need to be careful on this point.

How does AI affect conversation?

Conversations with a human differ from conversations with an AI machine.[9] Conversations with a human, especially if face-to-face, can be rich beyond the conversations with the AI machine. For example, Siri cannot, at present, see my eyes roll at her answer and interpret that as part of the conversation. This may not last forever.

I will also argue that the opportunity for learning is maximized when engaging in rich and meaningful conversation. This naturally leads to the question of how the level of education is related to, say, the quality of life.

[9] An AI Machine is considered herein as a combination of algorithms, data, processing capability, and connection to the external environment. Siri is an example of such a combination.

It would be interesting to compare the Social Progress Index ("2017 Social Progress Index," 2017) measures for countries against the level of education in these countries. I suspect the results of this comparison will be positive, but maybe this represents correlation and not causation.

In this regard I came across a journal article, *The Social Progress Index in International Business Site Selection: Three Case Studies* (Pate & Sweo, 2016), that may provide guidance for further research.

So, if conversation is essential for learning, and quality of life as measured by social progress is positively impacted by level of learning, does this suggest that the rise of AI may dampen the rise of quality of life?

The next three questions are provoked by the following:

> "The human operator loses the opportunity to develop through experience, the deeper skills and talents, such as assessment and judgment. The problem with machines that think is that they give rise to people who don't need to think. This presents a significant challenge in the dynamic and complex maritime world, where assessment and judgment, based on experience and total situational awareness, are fundamental to making the 'right' decisions – often under tight time constraints – that can spell the difference between safe passage and disaster." Captain George Quick, IMO Safety Committee (Hansen, Tikka, & Docherty, 2018).

Human Learning = 1/Machine Learning?

I've represented Captain Quick's comment in this fashion.

It's not that people will not want to or cannot think, but that AI will relieve them of the need to think as much about as much. Their stock of knowledge and its relevance thus becomes less over time. As a consequence they are at an increasing disadvantage as the machine learns more.

I think there are examples of this occurring around us.

> "Siri, how do I get to Chillicothe, Illinois?"

We need to know little more than how to drive and obey instructions. Even this minimum amount is moot in the age of autonomous vehicles.

What's the half-life of learning?

As a consequence of the above the half-life of learning in the individual continues to decrease. In a harsh way, we lead an increasingly irrelevant life.

What is the future of thinking?

In 2016 the “elephant chart” made its appearance (Milanovic, 2016).

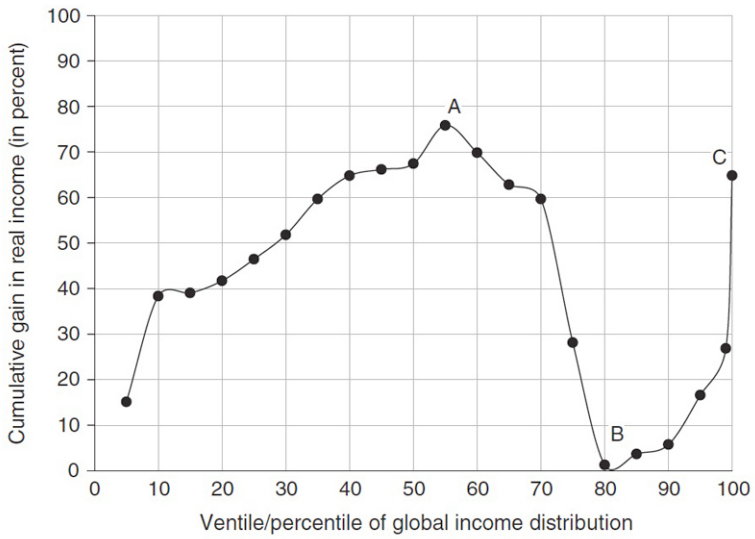

RELATIVE GAIN IN REAL PER CAPITA INCOME
BY GLOBAL INCOME LEVEL, 1988–2008

Figure 6

Some years ago I had a similar thought that was based on little more then intuition based upon experience.

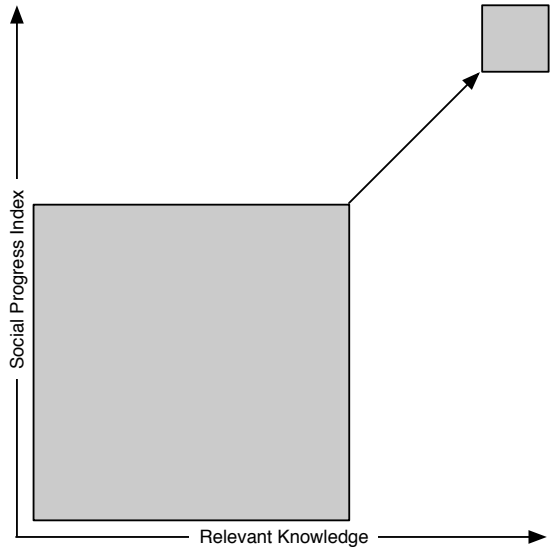

Figure 7

My initial expressions were cruder than what you see here, but the point was the same. The level of relevant knowledge has, in my opinion, much to do with one’s position in the social progress index as well as income.

I suspect that by most measures knowledge is accumulating and changing faster than before. If the perusal, analysis, and subsequent decisions are left to the machine, what happens to us?

What interpretation of "Greensill Bank in hands of German watchdog as parent seeks rescue" (FT 3-3-21) is AI likely to make?

If we can get by with vague questions, how can we judge the value of any answers?

Re Google Outlines Future of Its Search Engine[10]

In the print version, the subtitle is "AI technology named MUM promises to answer the vaguest of questions but critics warn of backlash from publishers." Emphasis is mine. If we can get by with vague questions, how can judge the value of any answers?

If you read the article online (see link below) be sure to scroll down and read the comments.

Thoughts of Others

In an article by Sherry Turkle (MIT) on *Technology and Human Vulnerability* in the September 2003 issue of the Harvard Business Review she opines:

> "We know that technology changes our lives -- but could it be changing our selves as well?

I suspect she was correct then and most surely correct today. I offer in support the relationships we have with our smart phones.

Coterminous with Turkle, Andy Borowitz in ' Danger, Stupid Human! Danger! ':

> "Our world is becoming a smarter place. We have smart phones, smart cards and smart bombs… An unprecedented surge in smartness, however, applies only to inanimate objects and has avoided humans altogether. In fact, as our high-tech devices have grown smarter, we've become much more dim-witted.
>
> …
>
> A car's global positioning system does, as advertised, help position us on the globe. It also strips us of our ability to read maps, plan itineraries and, ultimately, find our way from the bedroom to the bathroom."[11]

What does it take to accommodate and not be inundated by this change? Relevant knowledge seems to me to be a pretty good answer.

The third of Arthur Clarke's Laws is:

> "Any sufficiently advanced technology is indistinguishable from magic."

If the advanced technology can be distinguished by those that have the relevant knowledge, how is this gained in sufficient numbers such as to avoid the dangers inherent in believing in magic?

Michael Scott-Morton wrote in The MIT Report of December/January 2000:

> "Success takes a balance of forces, but technology is not always necessary and is almost never sufficient."

Indeed, in my experience, sometimes the most important technology decision one can make is to invest in paper and pencil. Successful adoption and adaptation of technology is fundamentally dependent upon human acceptance.

---

[10] https://giftarticle.ft.com/giftarticle/actions/redeem/a9813a76-bece-4a13-b9cf-a487e193bdc8

[11] This was originally found at http://www.smartmobs.com/archives/001798.html, but no longer appears.

Which leads us to John Gall, author of *Systemantics: How Systems Work and Especially How They Fail* (Gall, 1975), and his contention that:

> "Systems aligned with human motivational factors will sometimes work. Systems opposing such vectors will work poorly or not at all."

Gall did not, as I recall, distinguish between factors in support from those in opposition. Both need to be considered. Machiavelli wrote about this issue some 500 years ago in *The Prince* (Machiavelli, 1995):

> "There is nothing more difficult to take in hand, more perilous to conduct, or more uncertain in its success, than to take the lead in the introduction of a new order of things. Because the innovator has for enemies all those who have done well under the old conditions, and lukewarm defenders in those who may do well under the new."

I've been reading Brian Arthur's book, *The Nature of Technology: What It Is and How It Evolves* (Arthur, 2011). On pp 159-160 he discusses knowings:

> "Real advanced technology — on-the-edge sophisticated technology — issues not from knowledge but from something I will call *deep craft*. Deep craft is more than knowledge. It is a set of knowings. Knowing what is likely to work and what not to work. Knowing what methods to use, what principles are likely to succeed, what parameter values to use in a given technique. Knowing whom to talk to down the corridor to get things working, how to fix things that go wrong, what to ignore, what theories to look to. This sort of craft-knowing takes science for granted and mere knowledge for granted. And it derives collectively from a shared culture of beliefs, and unspoken culture of common experience.
>
> It also means knowing how to manipulate newly discovered and poorly understood phenomena, a type of knowing that comes from practical experimentation and research built up on local universities and industrial labs. A knowing that again becomes part of a shared culture."

It seems to me these "knowings" are the stuff of intuition gained through the skillful application of knowledge over time on dealing with the realities of the world. These knowings constitute a portion of the relevant knowledge of which I wrote earlier.

The Spring 2021 Capstone course in the SUNY Maritime MS in International Transportation Management program asked two questions.

1. The possibility for two AIs to engage in alternative solutions to a common issue, say container positioning, such that no solution emerges.
2. The investment required to apply AI is such that less capitalized organizations are unable to compete.

The final point I will make here comes from Mark Twain:

> "It ain't what we don't know that hurts us, it's what we know that ain't so."

Conclusions

Power Thesaurus[12] list 27 synonyms for "artificial intelligence." Some (e.g., neural network) we would likely agree with. Others (e.g., data processing) we would likely balk at considering.

Here is chart showing the growth in AI Startups in the US (Shoham, Perrault, Brynjolfsson, Clark, & LeGassick, 2017, p. 16).

[12] https://www.powerthesaurus.org/artificial_intelligence

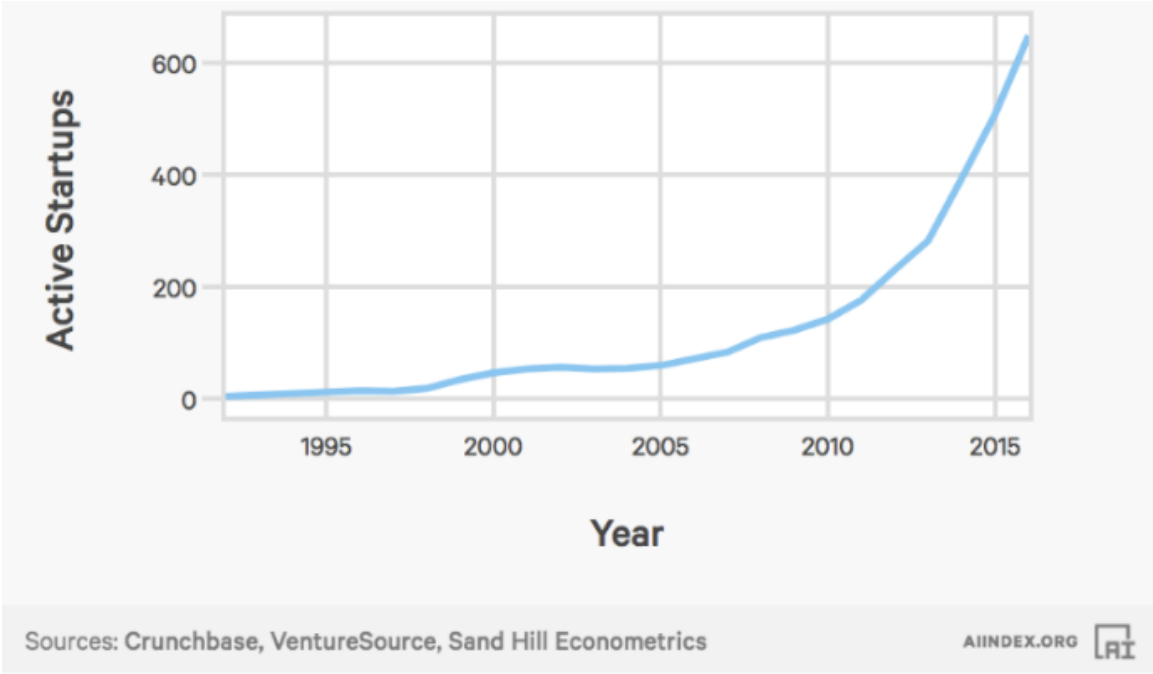

14x *The number of active US startups developing AI systems has increased 14x since 2000.*

Figure 8

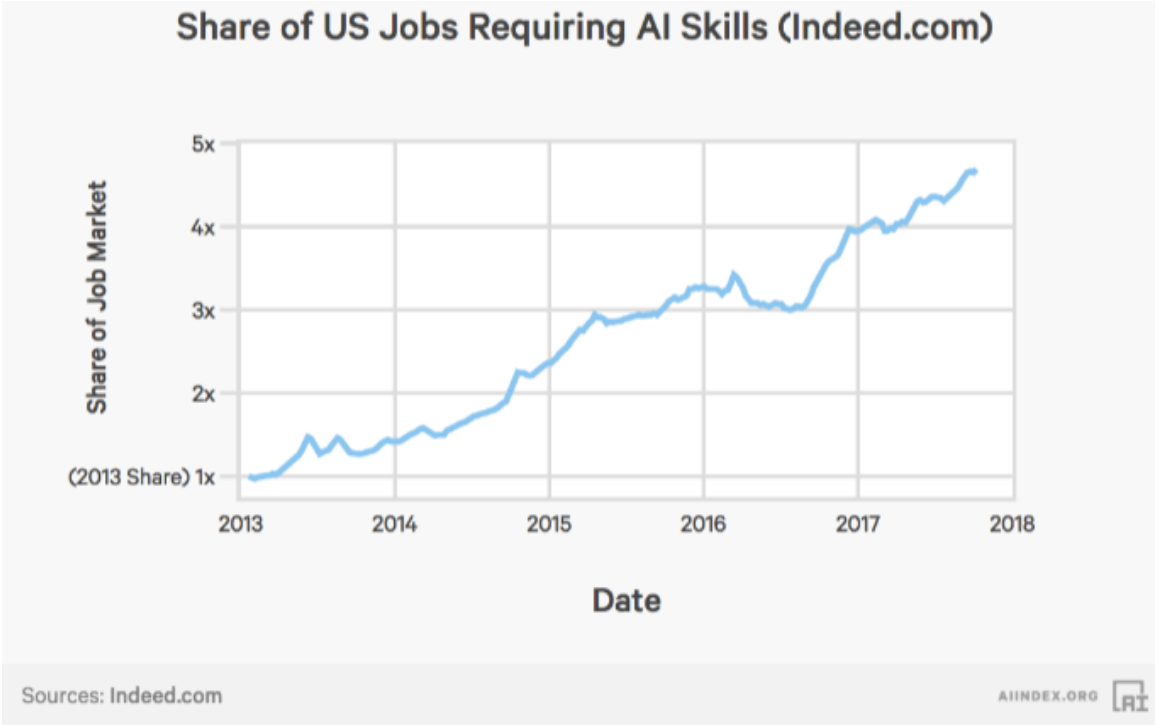

Figure 9

Here is another (p 18 of the same report showing an estimated of job growth related to AI.

Forbes (Columbus, 2017) reports "Vanson Bourne selected 260 IT and business decision-makers VP-level or higher from organizations with a global revenue of more than $50M a year were interviewed in July 2017." These interviews revealed that:

- 80% of enterprises already have some form of AI (machine learning, deep learning) in production today.
- 30% of enterprises are planning on expanding their AI investments over the next 36 months.
- 62% expect to hire a Chief AI Officer in the future.
- 91% see barriers ahead with lack of IT infrastructure (40%) and lack of talent (34%) as the most significant.

This brief paper ends a bit too abruptly for my taste, but my thinking is where my thinking is at the moment.

I have a better appreciation for the benefits and risks of AI than I had before. More troubling, to me at any rate, is the associated uncertainty.

■

Appendix: AI Bibliography

Following is a bibliography of AI-relayed material from my collection.

Agrawal, A., Gans, J. H., & Goldfarb, A. (2017, Spring). What to Expect From Artificial Intelligence. Retrieved July 11, 2017, from (Shoham, Perrault, Brynjolfsson, Clark, & LeGassick, 2017, p. 16)

Article 19. (2018). *Privacy and Freedom of Expression in the Age of Artificial Intelligence*. Retrieved from https://privacyinternational.org/sites/default/files/2018-04/Privacy%20and%20Freedom%20of%20Expression%20%20In%20the%20Age%20of%20Artificial%20Intelligence.pdf

Artificial Intelligence Will Make Forging Anything Entirely Too Easy. (2017, July 1). Retrieved July 31, 2017, from https://www.wired.com/story/ai-will-make-forging-anything-entirely-too-easy/

Banavar, G. (2016, November 29). What It Will Take for Us to Trust AI. Retrieved July 22, 2017, from https://hbr.org/2016/11/what-it-will-take-for-us-to-trust-ai

Bughin, J., Hazan, E., Ramaswamy, S., Chui, M., Allas, T., Dahlström, P., … Trench, M. (2017, June). How Artificial Intelligence can Deliver Real Value to Companies. Retrieved September 30, 2017, from http://www.mckinsey.com/business-functions/mckinsey-analytics/our-insights/how-artificial-intelligence-can-deliver-real-value-to-companies

Bughin, J., Seong, J., Manyika, J., Chui, M., & Joshi, R. (2018, September). Modeling the Global Economic Impact of AI. Retrieved September 9, 2018, from https://www.mckinsey.com/featured-insights/artificial-intelligence/notes-from-the-frontier-modeling-the-impact-of-ai-on-the-world-economy

Bughin, J., & van Zeebroeck, N. (2018, September 6). The Promise and Pitfalls of AI | by Jacques Bughin & Nicolas van Zeebroeck. Retrieved September 7, 2018, from https://www.project-syndicate.org/commentary/artificial-intelligence-digital-divide-widens-inequality-by-jacques-bughin-and-nicolas-van-zeebroeck-2018-09

China May Match or Beat America in AM. (2017, July 15). Retrieved July 16, 2017, from https://www.economist.com/news/business/21725018-its-deep-pool-data-may-let-it-lead-artificial-intelligence-china-may-match-or-beat-america

Chui, M., Manyika, J., Miremadi, M., Henke, N., Chung, R., Nel, P., & Malhotra, S. (2018, April). Sizing the Potential Value of AI and Advanced Analytics. Retrieved April 28, 2018, from https://www.mckinsey.com/global-themes/artificial-intelligence/notes-from-the-ai-frontier-applications-and-value-of-deep-learning

Confrete Problems in AI Safety. (n.d.). Retrieved from https://arxiv.org/pdf/1606.06565v1.pdf

Conn, A. (2017, February 3). How Do We Align Artificial Intelligence with Human Values? Retrieved September 13, 2018, from https://futureoflife.org/2017/02/03/align-artificial-intelligence-with-human-values/

Data Needs Cleaning Before Machine Learning Can Find Meaning. (2018, July 4). Retrieved July 6, 2018, from /blog/data-needs-cleaning-machine-learning-can-find-meaning

Davenport, T. H., & Kirby, J. (2016, Spring). Just How Smart Are Smart Machines? Retrieved July 11, 2017, from http://sloanreview.mit.edu/article/just-how-smart-are-smart-machines/

Davey, T. (2017, July 31). Towards a Code of Ethics in Artificial Intelligence with Paula Boddington. Retrieved August 10, 2017, from https://futureoflife.org/2017/07/31/towards-a-code-of-ethics-in-artificial-intelligence/

de Keyser, L. (2016, June 22). The Art of Mimicking Human Behavior. Retrieved June 22, 2016, from https://www.stratfor.com/weekly/art-mimicking-human-behavior

Dietvorst, B. J. (2017, July 5). When People Don't Trust Algorithms. Retrieved from http://sloanreview.mit.edu/article/when-people-dont-trust-algorithms/

Drogan, J. (2017, July 25). Thinking Prompted by "How Can We Stop Algorithms Telling Lies?" GBAT Department, SUNY Maritime College. Retrieved from http://jmsdrgn.squarespace.com/storage/Thinking%20Prompted%20by%20How%20Can%20We%20Stop%20Algorithms%20Telling%20Lies_.pdf

Drogan, J. (2018, April 1). AI: Questions. Unpublished.

Gans, J. (2017a, October 3). How AI Could Change Amazon: A Thought Experiment. Retrieved October 6, 2017, from https://digitopoly.org/2017/10/03/how-ai-could-change-amazon-a-thought-experiment/

Gans, J. (2017b, November 15). Can a superintelligence self-regulate and not destroy us? Retrieved November 18, 2017, from https://digitopoly.org/2017/11/15/can-a-superintelligence-self-regulate-and-not-destroy-us/

Gray, S. (2017, July 16). How Can We Stop Algorithms Telling Lies? by Cathy O'Neil.

Hanson, R. (2018, April 20). Prediction Machines. Retrieved April 21, 2018, from https://www.overcomingbias.com/2018/04/prediction-machines.html

Hanson, R. (n.d.). *What would happen if we upload our brains to computers?* Retrieved from https://www.ted.com/talks/robin_hanson_what_would_happen_if_we_upload_our_brains_to_computers

Hazan, J. B. and E. (n.d.). Five Management Strategies for Getting the Most From AI. Retrieved April 13, 2018, from https://sloanreview.mit.edu/article/five-management-strategies-for-getting-the-most-from-ai/

Khatchadourian, R. (2015, November 23). The Doomsday Invention. *The New Yorker*. Retrieved from http://www.newyorker.com/magazine/2015/11/23/doomsday-invention-artificial-intelligence-nick-bostrom

Kinni, T. (n.d.). Ethics Should Precede Action in Machine Intelligence. Retrieved July 25, 2017, from http://sloanreview.mit.edu/article/ethics-should-proceed-action-in-machine-intelligence/

Knight, W. (2017, April 11). The Dark Secret at the Heart of AI. Retrieved April 1, 2018, from https://www.technologyreview.com/s/604087/the-dark-secret-at-the-heart-of-ai/

Knight, W. (n.d.). If a Driverless Car Goes Bad We May Never Know Why. Retrieved July 8, 2016, from https://www.technologyreview.com/s/601860/if-a-driverless-car-goes-bad-we-may-never-know-why/

Kurzweil, R. (2017, March 14). Ray Kurzweil on How We'll End Up Merging With Our Technology. *The New York Times*. Retrieved from https://www.nytimes.com/2017/03/14/books/review/thinking-machines-luke-dormehl.html?smprod=nytcore-ipad&smid=nytcore-ipad-share

Kuziemski, M. (2018, May 1). A Precautionary Approach to Artificial Intelligence. Retrieved May 13, 2018, from https://www.project-syndicate.org/commentary/precautionary-principle-for-artificial-intelligence-by-maciej-kuziemski-2018-05

Lee, K.-F. (2018). *How AI Can Save Our Humanity*. TED. Retrieved from https://www.ted.com/talks/kai_fu_lee_how_ai_can_save_our_humanity

London, S. (2017, November). Artificial intelligence in business: Separating the real from the hype | McKinsey & Company. Retrieved November 3, 2017, from https://www.mckinsey.com/business-functions/mckinsey-analytics/our-insights/artificial-intelligence-in-business-separating-the-real-from-the-hype

Mahroum, S. (2018, February 16). The AI Debate We Need by Sami Mahroum. Retrieved February 16, 2018, from https://www.project-syndicate.org/onpoint/the-ai-debate-we-need-by-sami-mahroum-2018-02

Makar, A. B., McMartin, K. E., Palese, M., & Tephly, T. R. (1975). Formate assay in body fluids: application in methanol poisoning. *Biochemical Medicine*, *13*(2), 117–126.

Malone, T. W. (2018, May 21). How Human-Computer 'Superminds' Are Redefining the Future of Work. Retrieved June 20, 2018, from https://sloanreview.mit.edu/article/how-human-computer-superminds-are-redefining-the-future-of-work/

Manyika, J., & Sneader, K. (2018, June). Ai, Automation, and the Future of Work: Ten Things to Solve for (tech4good). Retrieved July 9, 2018, from https://www.mckinsey.com/featured-insights/future-of-organizations-and-work/ai-automation-and-the-future-of-work-ten-things-to-solve-for

McKinsey & Company. (2017, July). Ask the AI Experts: What Advice Would You Give to Executives About AI? Retrieved July 15, 2017, from http://www.mckinsey.com/business-functions/mckinsey-analytics/our-insights/ask-the-ai-experts-what-advice-would-you-give-to-executives-about-ai

Metz, C. (2017, November 5). Building A.I. That Can Build A.I. *The New York Times*. Retrieved from https://www.nytimes.com/2017/11/05/technology/machine-learning-artificial-intelligence-ai.html

Muehlhauser, L., & Bostrom, N. (2014). Why We Need Friendly AI. *Think*, *13*(36), 41–47. https://doi.org/10.1017/S1477175613000316

Mukherjee, S. (2017, March 13). The Future of Humans? One Forecaster Calls for Obsolescence. *The New York Times*. Retrieved from https://www.nytimes.com/2017/03/13/books/review/yuval-noah-harari-homo-deus.html

Ogilvy, J. (2018, January 3). The Promise and the Threat of AI. Retrieved January 4, 2018, from https://worldview.stratfor.com/article/promise-and-threat-ai

O'Neil, C. (2017, July 16). How Can We Stop Algorithms Telling Lies? *The Observer*. Retrieved from http://www.theguardian.com/technology/2017/jul/16/how-can-we-stop-algorithms-telling-lies

Piekniewski, F. (2018, May 29). AI Winter Is Well on Its Way. Retrieved June 2, 2018, from https://blog.piekniewski.info/2018/05/28/ai-winter-is-well-on-its-way/

Ransbotham, S. (2017, June 20). Romantic and Rational Approaches to Artificial Intelligence. Retrieved July 6, 2017, from http://sloanreview.mit.edu/article/romantic-and-rational-approaches-to-artificial-intelligence/

Rosenberg, S. (2017, November 1). Why AI Is Still Waiting For Its Ethics Transplant. Retrieved January 23, 2018, from https://www.wired.com/story/why-ai-is-still-waiting-for-its-ethics-transplant/

Rotman, D. (n.d.). The Relentless Pace of Automation. *MIT Technology Review*, *120*(2), 92–95.

Simonite, T. (2017, July 21). IBM's Watson Won Jeopardy, but Can It Win the New AI Biz? Retrieved July 22, 2017, from https://www.wired.com/story/ibm-watson-won-jeopardy-but-is-it-smart-enough-to-spin-big-blues-ai-into-green/

Snow, J. (2018, January 5). DeepMind's Cofounder Thinks AI Should Get Ethical in 2018. Retrieved January 6, 2018, from https://www.technologyreview.com/the-download/609895/deepminds-cofounder-thinks-ai-should-get-ethical-in-2018/

Stiglitz, J. (2018, September 8). Joseph Stiglitz on Artificial Intelligence: "We're Going Towards a More Divided Society." Retrieved from https://www.theguardian.com/technology/2018/sep/08/joseph-stiglitz-on-artificial-intelligence-were-going-towards-a-more-divided-society

Tegmark, M. (2017, August 29). Friendly AI: Aligning Goals. Retrieved January 23, 2018, from https://futureoflife.org/2017/08/29/friendly-ai-aligning-goals/

The Algorithm Kingdom. (2017, July 15). Retrieved July 25, 2017, from https://www.economist.com/news/business/21725018-its-deep-pool-data-may-let-it-lead-artificial-intelligence-china-may-match-or-beat-america

The cost of innovation has risen, and productivity has suffered. (2017, September 30). Retrieved October 2, 2017, from https://www.economist.com/news/finance-and-economics/21729746-exploitation-currently-available-knowledge-far-complete-cost

The Download from MIT Rech Review. (2018, May 4). Nigerian princes are still raking in the dough.

Will we understand why driverless cars do what they do? (2016, July 9). Retrieved July 10, 2016, from http://marginalrevolution.com/marginalrevolution/2016/07/will-we-ever-understand-why-driverless-cars-do-what-they-do.html

Wilson, H. J., Daugherty, P. R., & Morini-Bianzion, N. (2017, Summer). The Jobs That Artificial Intelligence Will Create. Retrieved July 6, 2017, from http://sloanreview.mit.edu/article/will-ai-create-as-many-jobs-as-it-eliminates/

Wladawsky-Berger. (2018, June 11). The Impact of AI on R&D and Innovation. Retrieved June 11, 2018, from http://blog.irvingwb.com/blog/2018/06/the-impact-of-artificial-intelligence-on-rd-and-innovation.html

Wladawsky-Berger, irving. (2017, June 26). Serious AI Challenges that Require Our Attention. Retrieved July 31, 2017, from http://blog.irvingwb.com/blog/2017/06/how-can-we-ensure-that-our-increasingly-powerful-ai-systems-benefit-humanity.html

Wladawsky-Berger, I. (2018, July 23). What Machine Learning Can and Cannot Do. Retrieved July 23, 2018, from http://blog.irvingwb.com/blog/2018/07/what-can-machine-learning-do.html

Wong, J. I., & Wong, J. I. (2016, March 25). Google's Ai Won the Game Go by Defying Millennia of Basic Human Instinct. Retrieved September 17, 2017, from https://qz.com/639952/googles-ai-won-the-game-go-by-defying-millennia-of-basic-human-instinct/

References

2017 Social Progress Index. (2017). Retrieved January 1, 2018, from https://www.socialprogressindex.com/

Arthur, W. B. (2011). *The Nature of Technology: What It Is and How It Evolves* (1. Free Press pbk. ed). New York: Free Press.

Columbus, L. (2017, October 16). 80% Of Enterprises Are Investing In AI Today. Retrieved October 14, 2018, from https://www.forbes.com/sites/louiscolumbus/2017/10/16/80-of-enterprises-are-investing-in-ai-today/

Data Needs Cleaning Before Machine Learning Can Find Meaning. (2018, July 4). Retrieved July 6, 2018, from /blog/data-needs-cleaning-machine-learning-can-find-meaning

Drogan, J. (2009). *Data, Information, and Knowledge - Relevance and Understanding*. SUNY Maritime College. Retrieved from http://jmsdrgn.squarespace.com/storage/Data%20Information%20and%20Knowledge%20-%20Relevance%20and%20Understanding.pdf

Gall, J. (1975). *Systemantics: How Systems Work and Especially How They Fail*. San Antonio, TX: Quadrangle.

Go and Mathematics. (2017). In *Wikipedia*. Retrieved from https://en.wikipedia.org/w/index.php?title=Go_and_mathematics&oldid=781220113

Hansen, B., Tikka, K., & Docherty, G. (2018, April). *Autonomous Shipping and the Regulations Needed To Make It Happen*. Webinar presented at the IHS Markit Fairplay.

Knight, W. (2017, April 11). The Dark Secret at the Heart of AI. Retrieved April 1, 2018, from https://www.technologyreview.com/s/604087/the-dark-secret-at-the-heart-of-ai/

Machiavelli, N. (1995). *The Prince (1513)*. (D. Wootton, Trans.). Indianapolis/Cambridge: Hackett Publishing Company, Inc.

Milanovic, B. (2016, July 1). The Greatest Reshuffle of Individual Incomes Since the Industrial Revolution. Retrieved October 14, 2018, from https://voxeu.org/article/greatest-reshuffle-individual-incomes-industrial-revolution

Pate, S. K., & Sweo, R. (2016). The Social Progress Index in International Business Site Selection: Three Case Studies. *Journal of International Education and Leadership*, *6*(2).

Shoham, Y., Perrault, R., Brynjolfsson, E., Clark, J., & LeGassick, C. (2017). *Artificial Intelligence Index 2017 Annual Report* (p. 101). Stanford University. Retrieved from https://aiindex.org/2017-report.pdf

Wong, J. I., & Wong, J. I. (2016, March 25). Google's Ai Won the Game Go by Defying Millennia of Basic Human Instinct. Retrieved September 17, 2017, from https://qz.com/639952/googles-ai-won-the-game-go-by-defying-millennia-of-basic-human-instinct/