AH ANALECTA HERMENEUTICA
INTERNATIONAL INSTITUTE FOR HERMENEUTICS / INSTITUT INTERNATIONAL D'HERMÉNEUTIQUE

ISSN 1918-7351
Volume 15.1 (2023)

Artificial Intelligence: Thoughts from a Psychologist

Micheal J. Meitner
University of British Columbia, Canada

Abstract
In its current state, Artificial Intelligence (AI) is still very far from reaching the complexity of the human brain. Technological progress, however, might bring about AI as a general intelligence surpassing our own in important aspects. The concept of neurodiversity is introduced and it is suggested that the field of AI might benefit from the incorporating of this concept in the form of sets of neurodiverse AIs from which a diverse set of solutions can be generated. Issues of sustainability and equity are also discussed in light of rapid advances in the field.

Keywords: artificial intelligence, neurodiversity, AI risks, open access data, machine learning

Introduction

Artificial intelligence, or AI as it is commonly referred to, is a suite of technologies that are poised to change the world as we know it. The concept of AI has been with us throughout antiquity in the mythologies of the Greeks and in early conceptions of automata. Early work in cybernetics and eventually neural networks brought this concept out of the realm of fantasy and into the modern world. In 1952, Marvin Minsky and Dean Edmonds succeeded in creating the world's first functional neural network machine, the Stochastic Neural Analog Reinforcement Calculator or SNARC.[1] While this was no doubt an impressive achievement of the day, it did not really live up to the dreams of the ancients of a machine that would embody more human characteristics.

The holy grail of AI has always been to create a machine capable of generalized intelligence. In fact, the first know test of generalized intelligence in AI was posited by Alan Turing in his seminal paper "Computing Machinery and Intelligence" and to this day is known as the Turing Test.[2] The point of this test was to ascertain if a machine could fool a human into thinking it was conversing with an actual human being. When I was in graduate school, I had the pleasure to interact with ELIZA, a computer program written in 1966 by Joseph Weizenbaum to mimic the behavior of a Rogerian psychotherapist which some would argue was the first program to pass the Turing Test.[3] While this contention remains controversial to some I can personally attest to the convincing nature of the program. However, I can also attest to the fact that it was quite easy to trip up this software and therefore destroy the illusion. It could never imagine new or novel situations and often answered any question that required creativity with a question of its own. Some historians of the internet age might say that ELIZA represented the first "bot," a software program that imitates the behavior of a human, as in participating in chatroom or IRC discussions. As most of us know today, bots have become far more sophisticated and for many of us they seem quite human when we interact with them.

However, there remains a massive disconnect between imitating a human and creating an artificial human brain. The human brain contains approximately 86 billion neurons and each neuron has on average 7000 synaptic connections yielding nearly a quadrillion synapses.[4] In terms of simple computational power (measured in floating

---

[1] Marvin Minsky, "A neural-analogue calculator based upon a probability model of reinforcement," (Technical document, Harvard University Psychological Laboratories, Cambridge, Massachusetts January 8, 1952).
[2] Alan Turing, "Computing Machinery and Intelligence," *Mind* 59, no. 236 (1950): 433–60. http://www.jstor.org/stable/2251299.
[3] Joseph Weizenbaum, "ELIZA—a computer program for the study of natural language communication between man and machine," *Communications of the ACM* 9, no. 1 (1966): 36-45.
[4] David A Drachman, "Do we have brain to spare?," *Neurology* 64, no. 12 (2005): 2004-2005; Herculano-Houzel, Suzana. "The human brain in numbers: a linearly scaled-up primate brain," *Frontiers in human neuroscience* (2009): 31

point operations per second or FLOPS) the human brain is estimated to be capable of approximately 1 exaFLOP (10^{18}).[5] Modern technology still falls short of this degree of raw computational power. The world's fastest supercomputer, Fujitsu for Japan's RIKEN Center for Computational Science supercomputer, has currently achieved .422 exaFLOPS.[6] However, it should be noted that important architectural aspects of the human brain are even further from the realm of possibility currently. Simply having the ability to do the same number of calculations over time does not mean that the arrangement of those neuronal units is in anyway similar to that of a human brain. Even in the case of AI modeling of C. elegans, a common worm that has only has 302 neurons, researchers are still refining the architecture of that model based on new electron microscopy data.[7] Therefore, the goal of a generalized intelligence instantiated in a computer is likely very far in the future. One possible technological development that may change this calculus is quantum computing but this still has significant challenges to overcome to become relevant to this discussion. Computational power (quantum or not) will certainly close the gap but this belies the fact that human brain is not simply the sum of its abilities to do raw computations.

In general, I would say that AI, in its current form, is in no way like the human brain even though AI researchers use architecture developed from observations of neuroanatomy. Modern AI is mostly focused on "narrow, shallow or weak AI" tasks such as finding patterns in our purchases and suggesting new ones based on these patterns. Even those AI's considered "broad, deep or strong AI" do not really approach the complexity of the human brain. Deep AI consists of numerous neural networks often hierarchically arranged that allow for deeper levels of abstraction from the inputs in the model. In addition, deep AI techniques deal well with unstructured data and can analyze that data in an unsupervised fashion. These qualities have made deep learning techniques quite ubiquitous and they have been employed to tackle problems such as speech recognition and computer vision. Artificial general intelligence, on the other hand, will require substantial leaps in both hardware and software before this can be realized.

At this point I would like to compare and contrast the nature of artificial and human general intelligence as seen in table 1 below.

---

[5] A point should be made that direct comparison of the human brain's computational power and a computers is not technically possible to achieve. For more information see "Brain performance in FLOPS," aiimpacts.org, AI Impacts, January 13 2021, https://aiimpacts.org/brain-performance-in-flops/.

[6] Scott Fulton III, "Top500: Japan's Fugaku Still the World's Fastest Supercomputer," Data Center Knowledge, November 18 2020, January 26 2021, https://www.datacenterknowledge.com/supercomputers/top500-japan-s-fugaku-still-world-s-fastest-supercomputer.

[7] Steven J. Cook, Travis A. Jarrell, Christopher A. Brittin, Yi Wang, Adam E. Bloniarz, Maksim A. Yakovlev, Ken CQ Nguyen et al, "Whole-animal connectomes of both Caenorhabditis elegans sexes." *Nature* 571, no. 7763 (2019): 63-71.

Artificial	**Human**
Infinite sensors	Limited sensors (can be augmented)
Infinite dimensions	Dimensionally challenged
Infinite data storage	Limited
Technology bound	Organism bound
Fairly stable goals (can be made to evolve)	Changing and evolving goals
Ever increasing processing speed	Speed mostly fixed
Replication generally yields copies (unless a genetic algorithm is used)	Replication yields neurodiversity
Consciousness?	Multiple unconsciousness systems partially discovered by consciousness

Table 1: A comparison of artificial and human general intelligence

As is evident from table 1, artificial general intelligence holds much promise and will likely lead to the formation of an artificial superintelligence. Being able to surpass our limitations in data sensing, data storage (memory in humans) and in hyper-dimensional thinking at speed will allow AIs to make tractable those problems that have long eluded us. The goals of AI and the eventual architecture (and potential diversity of architecture) seem to be important turning points in our thinking about how we might make progress toward the creation of artificial general intelligence. Let us start by first turning our attention to diversity in AI.

Neurodiversity and AI

The concept of neurodiversity has been with us since 1998 and refers to the revelation that variation in the human brain is vast and while some variation may be detrimental other variations may represent significant strengths or improvements. In fact, I would go as far as saying that neurodiversity can in fact represent a competitive advantage. If true for humans, this surely would be true of diversity in AI as well. It has long been known that genetic algorithms (GAs) can be used to spawn novel architectures for neural networks that can be used to evaluate the degree of performance of its progeny on some fitness function or goal. This allows for competition between various forms of an AI algorithm and leads to better solutions to problems that the AI is tasked with.

This represents some degree of neurodiversity in AI already, albeit a weak form of it, as unsuccessful progeny are "killed off" and therefore diversity is not maintained. Ideally, neurodiverse AI systems would be persisted and alternate solutions could be investigated to allow for insights into divergent approaches that may help us to better define and build robust and resilient AIs in the future.

Ultimately the discussion of the concept of neurodiversity in the context of AI causes us to question our ideas about goals. Goals in AI must be made explicit in some way and often represent the most challenging aspect of creating a functional AI. For many AIs there are more than one goal that the algorithm is trying to maximize or balance amongst. However, all of these goals have a context and perspective. From a user's perspective, a common goal might be increasing the relevance of information retrieved based on a query. From the company's perspective a similar goal might be user engagement. These differences in defining goals can have significant effects on the outputs of an AI. In fact, they define them. Variability in goal definition over time allows a model to adapt to changing system conditions.

As referenced in table 1, human goals seem to be ever changing and evolving as our understanding of the world progresses. This is especially true in the case of "wicked" problems. Wicked problems are those that defy simple solutions and are often comprised of multiple interacting systems. They are wicked because they are typically poorly understood, include contradictory information and are highly variable over time. Wicked problems do not have an optimal solution, rather they have temporary or partial solutions that are likely themselves to change over time. The changing nature of wicked problems and the large uncertainties in their predictions mean we have to take an adaptive approach to the problem. Like wicked problems, adaptive problems are where the problem definition is mostly unknown. Adaptive problems often require the locus of control for solving the problems to be decentralized. Stakeholders become the focus rather than disciplinary experts and as we well know stakeholders often have a variety of perspectives on a problem. This is the type of diversity needed if we hope to be able to conceptualize the system properly. From that one might argue that this means that multiple AIs might be needed to focus on various specificities of a problem in order for a larger definition of the problem to occur.

AI in the Environmental Sciences

Many of the issues of our day are in fact adaptive problems, such as most of our environmental current problems. The World Economic Forum report titled "Harnessing Artificial Intelligence for the Earth" states that there are 6 priority action areas for addressing environmental issues: 1) climate change, 2) biodiversity and conservation, 3) healthy oceans, 4) water security, 5) clean air and 6) weather and

disaster resilience.[8] Each of these areas has a series of sub areas that AI could be applied to in order to create a more sustainable future. In the case of climate change they refer to: clean power, smart transport options, sustainable production and consumption, sustainable land-use, smart cities and homes. AI can be applied to all of these areas and in certain cases have the potential to transform these sectors. Consider a modern energy grid that can use AI to adapt to changing supply and demand, incorporate traditional power sources with clean energy source and to make distributed energy possible at scale. This would seriously improve our ability to meet our climate change targets. As well, significant improvements in transportation, agriculture, and water management systems can also be realized by the application of AI technologies.

AI has already been applied to many environmental problems. Monitoring endangered species,[9] tracking diseases,[10] crop optimization,[11] smart buildings and associated IoT to increase efficiency,[12] predicting storms,[13] and managing traffic[14] are but a few of the many applications of AI in the environmental domain. In all of these cases, AI offers us a method to deal with the massive degrees of complexity that represent these wicked environmental problems. This is made possible by the vast quantities of data that we are currently collecting to support decision making in these areas.

The world of "big data" has arrived and no technology is better poised to make use of this plethora of data than AI. In fact, without computer aided decision making, I would venture to guess that we would not be able to effectively navigate, understand or even utilize the amount of data that is currently available. AI, however, has a special relationship with big data and becomes better when provided with increasing data volumes. AI is especially good at detecting anomalies in massive data sets, determining the probabilities of future outcomes and it can recognize patterns that human cannot.

---

[8] Celine Herweijer, Benjamin Combes, Pia Ramchandani, Jasnam Sidhu, "Harnessing Artificial Intelligence for the Earth," www3.weforum.org, World Economic Forum, January 2018, January 17 2021, www3.weforum.org/docs/Harnessing_Artificial_Intelligence_for_the_Earth_report_2018.pdf

[9] Antoine M. Dujon, and Gail Schofield, "Importance of machine learning for enhancing ecological studies using information-rich imagery," *Endangered Species Research* 39 (2019): 91-104.

[10] Zoie SY Wong, Jiaqi Zhou, and Qingpeng Zhang, "Artificial intelligence for infectious disease big data analytics," *Infection, disease & health* 24, no. 1 (2019): 44-48.

[11] Tanha Talaviya, Dhara Shah, Nivedita Patel, Hiteshri Yagnik, and Manan Shah, "Implementation of artificial intelligence in agriculture for optimisation of irrigation and application of pesticides and herbicides," *Artificial Intelligence in Agriculture* 4 (2020): 58-73.

[12] Rav Panchalingam, and Ka C. Chan, "A state-of-the-art review on artificial intelligence for Smart Buildings," *Intelligent Buildings International* 13, no. 4 (2021): 203-226.

[13] Amy McGovern, Kimberly L. Elmore, David John Gagne, Sue Ellen Haupt, Christopher D. Karstens, Ryan Lagerquist, Travis Smith, and John K. Williams, "Using artificial intelligence to improve real-time decision-making for high-impact weather," *Bulletin of the American Meteorological Society* 98, no. 10 (2017): 2073-2090.

[14] Rusul Abduljabbar, Hussein Dia, Sohani Liyanage, and Saeed Asadi Bagloee, "Applications of artificial intelligence in transport: An overview," *Sustainability* 11, no. 1 (2019): 189.

AI and Risk

The same World Economic Forum report that was mentioned above also identifies 6 areas of risk for AI. They are: performance, security, control, economic, social, and ethical. Performance risks refer to problems in deciphering the "black box" inner workings of an AI. Because we have little insight into what an AI is actually doing we have difficulties in knowing if its performance is accurate or even desirable. Issues of model fit are also complicated by this. If an AI is inferring future trends based on historical records then we need to wonder if those records contain enough information to support such prediction. If we don't know what an AI is doing internally then this problem is certain exacerbated.

Security risks, mentioned in this report, are also of concern. They reference "hackers" and the problems of bad actors manipulating algorithms to take control of them. This brings to light a more serious concern of who has control over these algorithms. Most AIs are in the hands of governments or large private sector companies. Neither of these has a great track record of acting for the social good. Private companies have a fiduciary duty to act in the best interests of stakeholders and while they may make efforts to address social issues this will never be their primary concern. However, one could argue that a government's main interest is the public good but as we all know this can be perverted in service of other goals that do not in fact create nor maintain social good. Even if these actors had social good in mind, how is it defined? Would those actions taken by these actors result in increased social good? This is an open question and certainly needs more thought and discussion to determine how to fully define this risk.

Control risks are some of the most blown out of proportion but are also some of the most worrying. This is where common narratives of post apocalyptic worlds governed by intelligent machines that have decided that humans represent a threat come in. However, this does not really represent a credible threat because you would need an AI capable of general intelligence and we have already determined that the likelihood that this will materialize in my lifetime is remote at best. What is of more concern are AIs that have direct control of various systems that might make decisions that lead to unintended consequences. One example of this is the flash crash of the US stock market in 2010 which was likely caused by interaction of multiple AI bots all speed trading at the same time.[15]

Economic risks are also potentially significant for AI as it moves forward. Companies that do not have access to AI or the associated data to drive them run the risk of being out competed. This in turn creates the risk that the business landscape will continue to shrink, creating increased inequity of wealth distribution and

---

[15] Tom Lauricella, Kara Scannell, and Jenny Strasburg, "How a Trading Algorithm Went Awry," The Wall Street Journal (New York, NY), October 2, 2010.
https://www.wsj.com/articles/SB10001424052748704029304575526390131916792

consolidating power with a few multinational companies. This may lead to a circumstance where a few companies begin to exert more power over global progression.

Social risks of AI are often defined as adaptation to increased automation pressures created by increased use of AI. Job loss and increased unemployment are real possibilities in a world where AI takes over much of the work of running the systems that we rely on. Additionally, AI algorithms can potentially be biased against certain factions of society, underpinning historic social inequities. New inequities can also be created by AI as it fundamentally changes the sector with in which it is being applied. Take as an example the transportation sector where autonomous vehicles are poised to massively disrupt people's lives who rely on this sector for employment.

The last risk that this report discusses is ethical risks. What choices will an AI make? Will they be beneficial choices? What about fairness and human rights? Privacy concerns are also discussed here. While all of these risks are important and represent an excellent attempt to get us all thinking about how AI will shape our future I believe that there is a significant omission in the risks associated with continued development and application of AI technology.

Additional Risks

I would add two additional risks to this list; access to both data and the knowledge needed to make sense of it. Let us first tackle data access. AI does not represent a valuable technology without the data that drives it and data is not generally freely available. Of course there are open data sets but the vast majority of meaningful data being generated today is in the hands of private corporations or governments. In 2020, every minute of every day we collectively generate 500 hours YouTube video, WhatsApp users share 41,666,667 messages, Facebook users upload 147,000 photos, Instagram users post 347,22 stories, and TikTok is installed 2,704 times.[16] The amounts of data being generate currently is staggering and for the most part we create this data. It is estimated that in 2020 each person on earth generates 1.7 MB of data per second.[17] Because access to this proprietary data is in the hands of the few, and by all accounts, the powerful, we run the risk of increasing inequity in society. Not just in terms of wealth, which is certainly an issue worth discussing, but also in terms of access to the information being derived by various AIs. How are common people supposed to keep up when knowledge about our behaviour, actions, purchases, interests, beliefs and values are being used to manipulate us? To control our purchases, our information feeds, our attention, our very lives. Something must be done to level the playing field.

[16] Domo, "Data never sleeps 8.0," Domo.com, Domo Inc., January 25, 2020, https://www.domo.com/learn/data-never-sleeps-8

[17] Domo "Data never sleeps 6.0," Domo.com, Domo Inc., 2018 https://www.domo.com/learn/infographic/data-never-sleeps-6

At a bare minimum we should have access to information about specifically how this data is being used to influence us.

Leveling the playing field however is not an easy task. There are many issues that need to be dealt with before we can hope to begin to bring us closer to balance. One of the first is the fact that this data is often privately owned. By agreeing to the licensing agreements (that honestly we don't really have much of a choice about), we have given up our rights to this data (as per the individual agreements). Additionally, we might also be concerned about privacy. No one really wants their neighbor to have access to their search history. This later problem however, is a far more tractable problem. Data can always be anonymized and abstracted to hide individuals within the masses as is commonly done with census data. The real sticky wicket is the ownership issue. Companies will not give up this data without a fight. This data represents real value to these companies and access to these data sets is often sold to third party companies for a variety of reasons. If this data were freely accessible to all it would significantly alter the business model for many companies that specialize in this area. If this is not remedied however, we can expect the knowledge divide in society to grow and eventually this may in turn weaken the functioning of civil society in the future.

One possible solution to this is to consider making companies that supply services that are critical to civil discourse, public utilities and regulate them as such. This would ensure fair and equal access to these platforms that give citizens voice. No one can tell you that you can't have a phone and as well no one censors what you say when you are participating in a phone call. Why should digital communication services be any different? One argument would be that today's digital communication platforms are in the public sphere rather than a private communication between individuals but this simply changes the scope of the communication and who can see it. Currently our approach to this is one of censorship. We disallow those things that we find offensive and label it hate speech, striking any record of it from our collective discourse. To some this is seen as necessary to ensure a peaceful and equitable society, to others this is seen as top down control by those in power to limit personal freedoms. The real question about censorship is not whether we should to it but who is doing it? Who gets to decide what appropriate speech is? If you are in charge of this then I would imagine that you would be quite happy with the rules but others might think of you as intolerant. In my opinion speech should be protected unless it directly leads to action that is prohibited: violence, harassment, etc., otherwise you have to decide what speech is acceptable and what is not and as history has taught us, this is a slippery slope. Once the precedent is set then even if the previous government enacted censorship laws that we consider ethically correct, the next party in power could use this same power to rewrite the laws in their favor and impose restrictions on speech that may not be as ethically centered. Take for instance the case of the National Union of Students who in 1973 got racist speech banned at universities in England. This ban was supported at the time by an organization of Zionist students. For a while this

seemed like a win but a few years later a different group of students was in power at the National Union of Students and they decided ban a Zionist speaker from speaking on campus because they now considered Zionism a form of racism. As you might have imagined, the group of Jewish students likely did not see how this might be turned against them as the leadership of that organization changed over time.[18]

If these companies are considered public utilities, we could also mandate that the proceeds of all analysis (knowledge) of our collective data should be freely available to everyone. This could come in the form of information dissemination and outreach on the part of the companies involved or it could mandate free and equal access to this data for the purposes of analysis. Both approaches have their strengths and weakness but it seems to me that allowing companies to be the sole arbiter of what gets published is a bad idea. If we pursued the later idea then we would need to find ways to make these vast amounts of data available in real time. Additionally, there are numerous barriers to fair and equitable access to this data even if provided freely. Access to sufficient computing hardware and software is required for anyone to begin the process of data analysis of these massive data sets. This is certainly not equitably distributed either. As well the knowledge required to not only conduct such an analysis but to comprehend it as well.

This brings me to my second point regarding risks; access to education and the knowledge that it brings to the individual is critical for individuals to have sufficient skill and training to approach this analysis with rigor and accuracy. To some degree we are far closer to this goal than we are to the goal of equal access to the data itself. Online education has exploded over the years and many topics such as computer programming skills are currently freely available to those that have the inclination to pursue them. This does not mean that they have access to the best and brightest minds on the subject but they do at least have enough access to learn most of what would be required of an AI researcher today. This would allow many more minds to be focused on common problems that we face today but also to potentially uncover new and previously unknown ideas at a far greater rate. It is in this exploratory space that I see this type of citizen science as being most directly applicable. With more minds come more perspectives, potentially allowing us to see a greater degree of the underlying "Truth" of the world. This is certainly in line with the ideas presented earlier related to neurodiversity.

A criticism of this approach might be that there is little control over the preparedness of individuals that seeks to undertake this type of work. However as one can plainly see this has always been the case. Even today not all researchers are considered equal. Some have tremendous knowledge and insights into the complexities of this undertaking and it is highly likely that contributions by these

[18] Ira Glasser, "How Freedom of Speech Protects You from Rulers like Trump," www.thedailybeast.com, The Daily Beast, October 4, 2020, January 26, 2021, https://www.thedailybeast.com/aclu-hero-ira-glasser-on-how-freedom-of-speech-protects-you-from-rulers-like-trump

individuals would be of more import. The solution to this is as it always was. Peer reviewed publication practices can go a long way to maintaining a high standard when it comes to the quality of our collective scientific efforts. However, we also have to aware of the fact that the academic-industrial complex does not have exclusive license to seek the truth. Many minds of great importance do not get the chance in life to contribute to their full potential. Creating a strong program of citizen science, free and open data sources, and access to the knowledge required to pursue such endeavors is paramount for our society to move toward a collective vision of a future where discourse is alive and well, we share that which has the potential to collectively move us forward, and allow all voices to participate in the creation of said future. It is my great hope that we can find new ways to set our collective table in such a way that all leave nourished in mind, body and spirit.

Conclusion

AI is quickly redefining our world and if we continue along our current tack we will likely exacerbate social inequalities and eventually make a less stable world for our children. This is the challenge for the science of AI. Can it mature quickly enough to provide us with insights and abilities that may help us to create a more sustainable and equitable future for all? Climate change and associated global risks are the challenge of our time and human nature is likely the root cause of this dilemma. AI offers us the potential to turn the light of science on our interior nature and the ramifications that this has for our collective future and our future actions within it. Make no mistake, trying to understand our collective behaviour, is the most wicked problem of all. Made even more so because we are the both the cause and solution to the problem. We are on the dance floor of our own perceptions, emotions and thoughts and we need to get on the balcony to be able to see the patterns that are emerging. AI offers us this vantage point. Granted we have a long way to go to improve the science of AI to allow us this potential but it exists none the less. I would hope that we could find a way to act in the collective good. To create a digital world where the rights to participate in society are inviolate, where access to data critical to said discourse is guaranteed, and diversity of perspective is the only requirement for entry.

Social networking companies must begin to think about the world they are allowing us to flow into. We are all on a river of time and the topography that underlies that river is the very nature of our digital (and physical, etc.) world that we have created to date. But just as topography yields to the bulldozer, our digital landscape is ours to remake. Let us start a discourse on this topic. Let all the world's peoples participate. We now have the tools to make this possible. It is an amazing world, but it is also one that must continue to improve if we are to hope to engineer ourselves out of the current environmental trajectory.