The University of Texas at Arlington

Using Artificial Intelligence to Solve
The Problem

Jacob Valdez 8688

Professional Practices CSE 4314 - 002

Dr. Ron Cross

2 November 2021

Author Note:

This Assignment is submitted towards and in support of the partial completion of the requirements for the Professional Practices Course

Table of Contents

Using Artificial Intelligence to Solve the Problem

The endless (Marcus 2020, 1) evolution of artificial intelligence (AI) spans a broad spectrum of approaches from symbolic to connectionist architecture, hand-crafted expert systems to machine learned neural networks, supervised to unsupervised to reinforcement learning paradigm, and uncountably many other hyperparameters. (Toosi, et al. 2021, 4-6) With dozens of AI-related papers published every day -- including defining accomplishments for the scientific community and larger globalized society as a whole -- a contemporary attitude is aired that no problem exists outside the application space of artificial intelligence and that by its continued evolution, humans may soon witness the realization of AI systems capable of operating completely independent of humans often identified as ‘Artificial General Intelligence’ (AGI). (Toosi, et al. 2021, 1,11-12,15; Latapie, et al. 2021) While this holy grail to numerous human research and engineering disciplines attracts much attention and endeavour for its practical and benevolent potential to all forms of social organization, the end goal is elusive. (Yampolskiy 2021, 1-3) There will never be a point when machine learning and related research and engineering disciplines can finally sit back and believe they have developed a system that requires no further intervention to operate effectively in all problem spaces. (Yampolskiy 2021, 2-3)

The strictest definition of universally “general” intelligence is intractable (Yampolskiy 2021, 6). “[A]ll algorithms that search for an extremum of a cost function perform exactly the same, when averaged over all possible cost functions.” (Wolpert, David H, and William G MacReady 1995, 1); for every finite pattern recognizer, there exists a more complex pattern it does not recognize (Sipser 2020, 35,140-145,152); and in the limiting case *ad infinitum*, the distribution of all patterns becomes uniform -- where intelligent prediction by finite computation is impossible. (Hornik, Kurt, et al. 2003) However, the Problem artificial general intelligence ambitiously aims for is to engineer a single algorithm that solves all possible problems.

This Problem is not just theoretical. The software engineering industry would give billions for one silver bullet (Brooks, Frederick P. 1987, 1-2), yet validation and maintenance are an ever-present concern to management. Really, all disciplines engineering would benefit if failure analysis were a closed field of study. However the chaos and complexity underlying apparent order and normalcy render precisely modeling chaotic open systems like the economy, states of health, and the brain elusive in the long-run. (Lee, Sung W. 2019)

In the void artificial general intelligence never filled, human research and development have acquired a broad assortment of formulae, algorithms, programs, and machine learning systems to individually solve narrower task domains. Specialized to their respective domains, principles of mathematics provide tools to impose formal order on real and imagined engineering objectives; algorithms idealize the implementation-agnostic execution of a sequence of instructions; programs leverage the computational properties of their electronic substrate to enable developers to process information at superhuman speed; and machine learning systems leverage large amounts of data to enable ML engineers to define programs at superhuman speed. (Goodfellow, Ian, et al. 2017, 1-3) No single approach wins in all problem domains, and the competition remains dynamic as human engineers shape the fitness landscape by their selection.

The above analysis is not intended to discourage research and development toward increasingly general AI systems, yet any research in that direction must recognize it is embarking on an endless road where competitive pressures are ever-present. For instance, the LeeNet represented a tremendous accomplishment in the image-classification problem space in 1989 (LeCun, Y., et al.), yet nearly two decades later, it was surpassed in accuracy by AlexNet

(Krizhevsky, Alex, et al., 2017) with many other rivals paving the way. Deep language models which began appearing in the late 2010's likewise have progressively grown in parameter count, sequence length, and inductive prior complexity (Goodfellow, Ian, et al. 2017, 373-376). Rather than suggest a single approach to suddenly out-compete all the aforementioned tools and levers of human intelligence, it is more realistic to expect a gradual, incremental, evolutionary development of artificial intelligence.

Real progress on the endless road to artificial general intelligence is not always noticeable, and its practical fruitage is even less apparent at times. For instance, returning to the earlier statement on language models, when the LSTM cell was first proposed in 1997 (Hochreiter, Sepp, and Jürgen Schmidhuber) few appreciated its significance then (Goodfellow, Ian, et al. 2017, 373-376). Over a decade passed before large-scale interest in deep learning revitalized this architectural component for practical language modeling (Goodfellow, Ian, et al. 2017, 373-376). Then the transformer (introduced in 2017) embodied a wealth of research in differentiable gating architectures that had not been fully cultivated to that point in time. (Vaswani, Ashish, et al. 2017, 1-2,8-10) Many more cases could enumerate this point of unrecognized research achievement.

Business professionals and executives recognizing this should be willing to explore applying deep learning systems to their processes, but they should not expect them to deliver an order-of-magnitude improvement toward business objectives. They should not be fooled by high-tech versions of get-rich-quick schemes which point to state-of-the-art (SOTA) experiment results but gain no real business value. On the other hand, academic advisors, teachers, and mentors can contribute toward sustainable long-term development of artificial intelligence by cultivating interest in prospective researchers and engineers. The machine learning research community should facilitate gradual development by putting less emphasis on holding SOTA benchmarks and encouraging more on open-code, open-data experiments with the aim of accelerating the speed of information exchange in the research community. Machine learning platform providers can facilitate this exchange by subsidizing or making their services available free for non-business customers. Governments, research agencies, and other funding sources can establish incentive prizes for researchers to develop AI systems that are capable of matching and surpassing human-level performance in common human problem domains like healthcare, automated science, and intelligent process automation.

Artificial intelligence will continue to evolve in complexity, scale, and generality, yet rather than coalescing into a singularity of superintelligence or even reaching level of generality where humans leave the optimization loop, the reasonable estimation will keep in mind that no AI system -- even one capable of optimizing itself -- generalizes to every possible domain without loss. In the world's nonstationary problem space, it becomes increasingly important to recognize the endless competition where humans will remain relevant.

Appendix A: References

1. Marcus, Gary. “The Next Decade in AI: Four Steps Towards Robust Artificial Intelligence.” Arxiv, arxiv.org/pdf/2002.06177.pdf.

2. Toosi, Amirhosein, et al. “A Brief History of AI: How to Prevent Another Winter (a Critical Review).” ArXiv.org, 8 Sept. 2021, https://arxiv.org/abs/2109.01517.

3. Latapie, Hugo, et al. “A Metamodel and Framework for Artificial General Intelligence from Theory to Practice.” ArXiv.org, 11 Feb. 2021, https://arxiv.org/abs/2102.06112.

4. Brcic, Mario, and Roman V. Yampolskiy. “Impossibility Results in AI: A Survey.” ArXiv.org, 1 Sept. 2021, https://arxiv.org/abs/2109.00484.

5. Wolpert, David H, and William G MacReady. “No Free Lunch Theorems for Search.” SFI WORKING PAPER, Feb. 1995.

6. Sipser, Michael. Introduction to the Theory of Computation. Course Technology, 2020.

7. Hornik, Kurt, et al. “Multilayer Feedforward Networks Are Universal Approximators.” Neural Networks, Pergamon, 19 Mar. 2003.

8. Brooks, Frederick P. “No Silver Bullet Essence and Accidents of Software Engineering.” Computer, vol. 20, no. 4, 1987, pp. 10–19., https://doi.org/10.1109/mc.1987.1663532.

9. Lee, Sung W. “A Copernican Approach to Brain Advancement: The Paradigm of Allostatic Orchestration.” Frontiers, Frontiers, 1 Jan. 1AD, https://www.frontiersin.org/articles/10.3389/fnhum.2019.00129/full.

10. Goodfellow, Ian, et al. Deep Learning, The MIT Press, Cambridge, Mass, 2017, pp. 1–3.

11. LeCun, Y., et al. “Backpropagation Applied to Handwritten Zip Code Recognition.” Neural Computation, vol. 1, no. 4, 1989, pp. 541–551., https://doi.org/10.1162/neco.1989.1.4.541.

12. Krizhevsky, Alex, et al. “ImageNet Classification with Deep Convolutional Neural Networks.” Communications of the ACM, vol. 60, no. 6, 2017, pp. 84–90., https://doi.org/10.1145/3065386.

13. Goodfellow, Ian, et al. “Sequence Modeling: Recurrentand Recursive Nets.” Deep Learning, The MIT Press, Cambridge, MA, 2017, pp. 373–376.

14. Hochreiter, Sepp, and Jürgen Schmidhuber. “Long Short-Term Memory.” Neural Computation, vol. 9, no. 8, 1997, pp. 1735–1780., https://doi.org/10.1162/neco.1997.9.8.1735.

15. Vaswani, Ashish, et al. “Attention Is All You Need.” ArXiv.org, 6 Dec. 2017, https://arxiv.org/abs/1706.03762.

Appendix B: Signed Academic Integrity Form

[Please see page attached below]