The University of Texas at Arlington

Relative Generality — Clarifications on the Evolution of Artificial Intelligence

Jacob Valdez 8688

Professional Practices CSE 4314 - 002

Dr. Ron Cross

18 November 2021

Author Note:

This Assignment (assignment#5) is submitted towards and in support of the partial completion of the requirements for the Professional Practices Course. The reader is invited to immediately email the author after reading this paper with any comments on its thesis. Correspondence may be sent to jacob.valdez2@mavs.uta.edu.

Table of Contents

Relative Generality — Clarifications on the Evolution of Artificial Intelligence

Popular science and the broader general public often endorse a system II-type misconception: artificial intelligence (AI) is either 'narrow' to some particular problem domain or 'generalized' in its understanding and ability with a categorical divide between the two.[1] Such a binary and absolute measure of intelligence runs counter to the very foundation that modern AI including machine learning (ML) continually[2] evolves from. It perpetuates the mythologically-inspired picture of humans as being the most general of all intelligences[3]. Its type-II categorical mode of thought accentuates 'great leaps for AI' as if they were entirely responsible for its usefulness without recognizing thousands of other works leading up to the few which steal public recognition[4]. More fundamentally, just as Ptolemy's geocentric model of the universe stalled astronomical understanding[5], so too this anthropocentric perspective on intelligence generality blinds the general public from appreciating the present potential of humankind's greatest technological achievement of all time.

The issue here is an uneven common ground on the term "general". Machine learning theory and practice establishes various associated definitions: "Domain generalization" refers to the ability of an ML model that is pretrained in one domain to maintain its performance in an unseen one.[6] Examples include pretraining face recognition models on pre-COVID face datasets and measuring the model's performance on mask-wearing faces or pretraining large language models on next-word prediction and then assessing their performance on question answering.[7] "In-context learning"/generalization refers to the ability of large sequential processing models to make complex transient behavior changes at inference-time (i.e.: without "training"), and examples include gpt3's ability to adopt and continue the tone of an input prompt.[8] Then some ML researchers consider "shallow generalization" and "broad generalization" separately where in the former case, an image classifier might be expected to perform robustly when noise is added to the image or the object is rotated, recolored, or scaled, and the latter case resembles generalization at the human-level such as imagining a zebra with only having seen independent examples of stripes and horses.[9] Finally, artificial general intelligence (AGI) refers to the holy grail of AI/ML science research and engineering possessing "the extensive general intelligence possessed by humans within a computational system"[10]. That ambition "artificial general intelligence" draws ideas on generality from the psychological notion of "general intelligence" which was popularized over a century ago by statistical psychologist Charles Spearman in a work titled *"General Intelligence," Objectively Determined and Measured*[11]. Spearman noted a hierarchical correlation of individuals' performance across various psychometric tests within and between domains and therefore suggested the existence of an underlying g-factor to human subject intelligence which should predict their performance across all measured tasks[12].

This is where the common ground divides: although useful for comparing individuals within the psychometric task space, it is misleading to treat the g-factor as a universal measure of intelligence generality since it only assesses psychometric performance within the anthropocentric realm of tasks[13]. Humans do not possess "general intelligence"[9]. The human brain possess functional regions suited for acquiring a broad variety of adaptive and self-actualized skills such as navigation in both physical and cognitive space (involving the hippocampus), external speech and internal consciousness [1] (involving the temporal lobe), and real and planned action sequence execution (involving the executive mode network)[14], but it is fundamentally an organ of finite capacity, complexity, and capability and its resultant bounded rationality demands it to be specialized to sensing, processing information, and acting on the temporal scale relevant to humans[13]. The human brain, for instance, cannot generalize to simultaneously memorize 20 numbers flashed on a screen for 1 second as *maqueque* monkeys can[15] suggesting a trade-off between the sequential (human) and parallel (*maqueque*) processing cognitive architectures. Anatomical neuroscience has long noted that the brain utilizes separate dorsal and ventral pathways to carry distinct "what" and "where" information streams[14]. Reflecting on this contemporary neuroscientist Karl Friston describes the human brain as less adapted to engage environments where object and spatial -- or more generally -- class and relational -- information covary than in the human-scale rigid-body world[16]. It is also not equipped to adapt to regularly interact with information streams that are excessively disintegrated across time or space such as flashing lights[17], lottery machines[17], or internet surfing[18], and behavioral 'adaptations' in the latter two cases are recognized as harmful additions[17,18]. Much meaningful information is structured in a form outside the reach of the human modalities such as light at higher and lower than visible frequencies and micro- and telescopic distances[19]. Even after developing instruments that transform previously unknown signals into observable form, the human brain has difficulty perceiving the hidden structure and causal hierarchy of this information[19].

It is difficult to find examples that are humanly-relevant, but the light of observation and experiment continually identify instances of the above where human intelligence does not generalize to every task, skill set, or problem domain, and the retrospective analysis should invite one to consider whether any objective measure of intelligence exists outside of the context of a given task and metered amount of experience. The No-free-lunch theorems for search and optimization imply that no universal objective exists for finite-horizon analysis[20]. Godel's incompleteness theorems establish that no formal proof system can both be complete and closed under itself[21]. Turing's Halting Problem gives one such example where even after an infinite number of computation steps, the solution to a problem remains unknown[22]. Though not identical to intelligence, complex tasks such as general search, optimization, proof, and decidability are minimally a baseline to 'General Intelligence', and yet they are impossible! If bounded human mathematical rationality can establish those impossibility proofs, we should likewise begin to entertain the corollary of a less categorical notion of whether AI is 'narrow' or 'general'.

There is no finish line to open-ended discovery, and the endless evolution of artificial intelligence from symbolic to connectionist approaches emphasizes this point[23]. Discussing generalization in the context of various human and artificial intelligences begs abandoning the anthropocentric conception of Intelligence as a magical, uniquely human, or even absolutely measured quality. This contemporary perspective optimally guides human research, development, business activity, and common understanding. Only once humans precisely and unambiguously formulate their objectives for the nebulous ambition of "artificial general intelligence" can they expect to reach this lofty goal, yet the reality for common person and ML professional alike is incremental growth in understanding and achievement. Incremental thought begets incremental development. Rather than waiting for a 'critical mass' of ML engineering to suddenly explode into some mythologically-inspired Artificial 'General' Intelligence, ML engineers actively contribute to its evolution in complexity, autonomy, and generality recognizing that no artificial intelligence is objectively 'better' than another outside the local context of task and experience -- which when measured relative to human objectives delivers feedback for endless improvement.

Appendix A: References

1. Emmert-Streib, Frank, et al. “A Clarification of Misconceptions, Myths and Desired Status of Artificial Intelligence.” ArXiv.org, Cornell University, 3 Aug. 2020, https://arxiv.org/abs/2008.05607.
2. Marcus, Gary. “The Next Decade in AI: Four Steps Towards Robust Artificial Intelligence.” ArXiv.org, Cornell University, 3 Aug. 2020, arxiv.org/pdf/2002.06177.pdf.
3. Latapie, Hugo, et al. “A Metamodel and Framework for Artificial General Intelligence from Theory to Practice.” ArXiv.org, Cornell University, 3 Aug. 2020, 11 Feb. 2021, https://arxiv.org/abs/2102.06112.
4. Toosi, Amirhosein, et al. “A Brief History of AI: How to Prevent Another Winter (a Critical Review).” ArXiv.org, Cornell University, 8 Sept. 2021, https://arxiv.org/abs/2109.01517.
5. Crrlstaff. “Early Astronomers: From the Babylonians to Galileo.” Central Rappahannock Regional Library, Central Rappahannock Regional Library, 19 Aug. 2021, https://www.librarypoint.org/blogs/post/early-astronomers/.
6. Ghifary, Muhammad, et al. “Scatter Component Analysis: A Unified Framework for Domain Adaptation and Domain Generalization.” ArXiv.org, Cornell University, 26 July 2016, https://arxiv.org/abs/1510.04373.
7. Vaswani, Ashish, et al. “Attention Is All You Need.” ArXiv.org, Cornell University, 6 Dec. 2017, https://arxiv.org/abs/1706.03762.
8. Brown, Tom B., et al. “Language Models Are Few-Shot Learners.” ArXiv.org, Cornell University, 22 July 2020, https://arxiv.org/abs/2005.14165.
9. Yampolskiy, Roman V. “Human ≠ AGI.” ArXiv.org, Cornell University, 11 July 2020, https://arxiv.org/abs/2007.07710.

10. Yamakawa, Hiroshi. “The Whole Brain Architecture Approach: Accelerating the Development of Artificial General Intelligence by Referring to the Brain.” ArXiv.org, Cornell University, 6 Mar. 2021, https://arxiv.org/abs/2103.06123.

11. Spearman, C. “‘General Intelligence," Objectively Determined and Measured.” The American Journal of Psychology, vol. 15, no. 2, 1904, p. 268., https://doi.org/10.2307/1412107.

12. Panizzon, Matthew S et al. “Genetic and Environmental Influences of General Cognitive Ability: Is g a valid latent construct?.” Intelligence vol. 43 (2014): 65-76. doi:10.1016/j.intell.2014.01.008

13. Chollet, François. “On the Measure of Intelligence.” ArXiv.org, Cornell University, 25 Nov. 2019, https://arxiv.org/abs/1911.01547.

14. Buzsáki, György. Rhythms of the Brain. Oxford University Press, 2006.

15. Templer, Victoria L., et al. “Co-Operation of Long-Term and Working Memory Representations in Simultaneous Chaining by Rhesus Monkeys (Macaca Mulatta).” Quarterly Journal of Experimental Psychology, vol. 72, no. 9, Sept. 2019, pp. 2208–2224, doi:10.1177/1747021819838432.

16. Friston, Karl. “Free Energy Principle — Karl Friston.” 16 June 2017, https://www.youtube.com/watch?v=NIu_dJGyIQI. Accessed 18 Nov. 2021.

17. Francisco, Alexandra de, et al. “Effect of Illumination Level [18F]FDG-Pet Brain Uptake in Free Moving Mice.” PLOS ONE, Public Library of Science, 13 May 2021, https://journals.plos.org/plosone/article/comments?id=10.1371%2Fjournal.pone.0251454.

18. Darnai, Gergely et al. “Internet addiction and functional brain networks: task-related fMRI study.” Scientific reports vol. 9,1 15777. 31 Oct. 2019, doi:10.1038/s41598-019-52296-1

19. Eagleman, David. Incognito: The Secret Lives of the Brain. Canongate, 2011.

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

21. Gödel Kurt, and B. Meltzer. On Formally Undecidable Propositions of Principia Mathematica and Related Systems. Dover Publications Inc, 1992.

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

23. Goodfellow, Ian, et al. Deep Learning, The MIT Press, Cambridge, Mass, 2017, pp. 16–18.

Appendix B: Endnotes:

[1] I hesitate using the term “consciousness” to describe our phenomenal experience. Mythologically and philosophically derived ideas like the above have long plagued neuroscience and bring us nowhere closer to real-world answers or solutions. *Buszaki*[14] asserts that the conversations people entertain today about “consciousness”, “attention”, “perception”, and related terms are nearly the same as they were in the days of Aristotle. We cannot expect to make progress towards understanding or imitating the brain while embracing the deceptive bias carried with this archaic terminology. I have not observed anything magical about the brain -- besides how hard it is for people to look beyond their own superstitions about it to real observations. Once contemporary neuroscience establishes a precise definition of what is actually meant by “consciousness”, such vague terms like it can be discarded and actually impede scientific progress.

Appendix C: Signed Academic Integrity Form

[Please see the attached page below.]