Undergraduate Thesis Prospectus

Detecting Graph Invariant Sub-Features Using Graph Neural Networks

(Technical Report)

The Interplay of Social and Technological Factors in the Development of AI

(STS Research Project)

Peter Sailer

pms2jkc@virginia.edu

Fall 2023

University of Virginia

On my honor as a University student, I have neither given nor received unauthorized aid on this assignment as defined by the Honor Guidelines for Thesis-Related Assignments.

Advisor

Travis Elliot , Department of Engineering and Society

Introduction:

Machine learning (ML) and Artificial Intelligence (AI) have dramatically changed the world over the past few decades. New algorithms, hardware, available data, and expanding expertise continue to perpetually produce the rapid improvement of existing systems and promote new technical breakthroughs. The technical portion of this thesis will analyze new machine learning developments in bioinformatics, specifically speaking to how Graph Neural Networks (GNNs) can be used to learn graph invariant sub-features. Then this thesis will more broadly investigate the complex actor network that has underpinned the development of machine learning both historically and recently.

Technical Report:

The field of bioinformatics has been greatly expanded by the rise of sophisticated machine learning and AI algorithms. These algorithms contribute to new and more practical means of gathering biological information, particularly in the domains of genomics, the study of genomes, and proteomics, which attempts to understand the organizational principles within nucleic acid and protein sequences (Bioinformatics, 2022).

Graph Neural Networks are one of many techniques that have been put forth in an attempt to solve these intrinsically difficult and complex problems. GNNs stand at the forefront of modern machine learning, offering a specialized approach to analyzing data with complex relationships and dependencies in graph-structured data. Unlike traditional neural networks, GNNs operate by leveraging techniques that allow nodes to communicate and aggregate information from their neighbors. In addition, GNNs are designed to operate directly on preexisting graph data, which significantly reduces the necessary vectorization and preprocessing of data.

A recent paper published at the International Conference on Learning Representations (ICLR) proposed a new class of highly powerful or expressive GNNs called Graph Isomorphism Networks (GINs). Xu et al. demonstrated that GINs could better discriminate between a range of different kinds of graph data including, graph representations of chemical compounds and enzymes. In particular, GINs were able to outperform all of the state-of-the-art algorithms when classifying "chemical compounds screened for the ability to suppress or inhibit the growth of a panel of human tumor cell lines" from the NCI1 dataset (Xu et al., 2019). Further works on GNNs from ICLR 2023 have laid the groundwork for designing GNNs that can efficiently discriminate between graphs with different biconnective sub-trees or articulation points (Zhang et al., 2022).

In the technical report, we plan to investigate and engineer a new kind of GNN, specifically designed to identify graph invariant sub-features. At this stage, we are currently familiarizing ourselves with state-of-the-art techniques (Sato, 2020).

The Interplay of Actors within Factors in the Development of AI

Problem:

The field of AI has historically been fueled by overenthusiasm and hindered by false promises. Repeatedly, lofty and inaccurate claims have been made regarding the capabilities of learning algorithms leading to the failure of many programs and a loss of funding for the field as a whole. Today more than ever modern AI is touted as nearing superhuman capabilities and promises to revolutionize the global economy. Some experts in the field even shared concerns regarding near-future AI algorithms as posing an existential risk to humanity. As such, the purpose of this paper will be to explore the major agents acting within the complex network of AI development, how they have affected one another historically, and what their current coevolution looks like today.

STS Framework:

This paper will describe the history and current state of machine learning and AI development through the lens of Actor Network Theory (ANT). This framework, first proposed by Bruno Latour, provides a method of analyzing complex sociotechnical systems by prescribing agency to both human and non-human actants and then characterizing their roles within the networks that define their interactions. In sociology and the context of ANT, the interactions and alignment of various actants to support a common objective is referred to as the process of translation. ANT describes the process of translation through four stages, problematization, interessement, enrollment, and mobilization, which will later be discussed in the context of the current state of modern machine learning (Lolha, 2022). Using ANT as a framework will allow for the ability to examine the intricate contributions of several characteristically different but symmetrically important actants such as prominent scientists, leading academic institutions, advocacy groups, policies, algorithms, data, hardware, and the ideas surrounding the current and future state of AI, particularly in the case of Large Language Models (LLMs).

The primary purpose of this paper will be to identify the unignorable actants participating in the translation of new AI technologies and then, in parting with ANT, attempt to establish a minimal network that can maximally explain the dynamic evolution of the translation of new AI technologies.

Methods:

To understand the complex Machine Learning and Artificial Intelligence Actor Network (MLAIAN) underpinning the current state of computer science research, it will be informative to discuss the historical roles of different actants, with a particular emphasis on human and institutional contributions.

The mathematical origins of Machine Learning (ML) and Artificial Intelligence (AI) algorithms date back as far as the early 1800s with the discovery of linear regression by Adrien-Marie Legendre and later Johann Carl Friedrich Gauss. Both physicists uncovered systematic techniques for predicting orbital trajectories based solely on previous records (Schmidhuber, 2022). This characterizes the initial problematization of the MLAIAN, namely to find more efficient means of solving computationally or conceptually difficult problems. Later, the foundations of Neural Networks were laid in the 1940s when computer scientists began to explore the possibility of creating machines that could simulate human intelligence. The conceptual groundwork came originally in 1943 with McCulloch and Pitts' computational model called threshold logic (Beeman, 2001). This later inspired the experimental work of Frank Rosenblatt in 1957, who created the first neural network known as the perceptron, designed to simulate the thought processes of the human brain (Fabien, 2018). Rosenblatt, a student of psychology, described the device as "the first machine which is capable of having an original idea" (Lefkowitz, 2019). The success and intrigue of early machine learning models led to the interessment and enrollment of institutional actants. Excessive funding was provided from various departments and agencies within the national government such as the Defense Advanced Research Projects Agency (DARPA) and the CIA. In the same decade, the Georgetown experiment was able to make headlines in the field of machine translation by deciphering more than forty Russian sentences into English (CSE 490H History Exhibit, 2023). However, the sophistication of these algorithms was largely exaggerated with results that did not generalize to longer texts where context was needed to disambiguate otherwise equivalent translations (Hutchins, 1996; Garvin, 1967). This led to an investigation by the Automatic Language Processing Advisory Committee (ALPAC) and eventually resulted in the National Research Council (NRC) ending all support for the project (Hutchins, 2005). Likewise, in 1969 Marvin Minsky and Seymour Papert published "Perceptrons: An Introduction to Computational Geometry" in which they provably showed that single-layered perceptrons, such as the one Rosenblatt had created, could not perform at least one of the two fundamental logical operations severally upper bounding its theoretical capabilities (Swaine, 2023). This ushered in a period known as the AI winter where funding and interest in artificial intelligence became significantly reduced.

The initial rise and fall of machine learning as a field demonstrate a translational failure within the MLAIAN. The focal actant being the idea of the extreme usefulness of AI algorithms was not ably substantiated. Ultimately, the mobilization of actants such as Rosenblatt and the Georgetown experiment fell short due to the primitiveness of the theories and the need for more powerful hardware.

This lies in stark contrast, however, to the flourishing MLAIAN to come. There was a resurgence of machine learning in the 80s with a particular emphasis on analyzing arbitrary patterns in data. The pioneering of deep learning and back-propagation laid the groundwork for Convolution Neural Networks (CNN) and Long Short-Term Memory (LSTM) Recurrent Neural Networks, which revolutionized machine vision and speech recognition respectively (Schmidhuber, 2022). Stronger and more powerful graphics processors (GPUs) have been developed with specialized tensor cores to speed up the matrix calculus used in ML (Zhang, 2023). The creation of the Internet and large labeled datasets such as MNIST have enabled deep learning algorithms to become parameterized to a vastly superior extent. In the words of Goodfellow, “The learning algorithms reaching human performance on complex tasks today are nearly identical to the learning algorithms that struggled to solve toy problems in the 1980s…. The most important new development is that today we can provide these algorithms with the resources they need to succeed” (Goodfellow, 2019). In saying this, Goodfellow is suggesting the availability of data has made an extraordinary impact between the 80s and 2010s, overshadowing much of the previous algorithmic progress. Hence in addition, to algorithms and institutions, data, and hardware have become prominent actants aiding in the mobilization of MLAIAN.

The story over the last few years has been slightly different, however, with the development of the modern attention-based transformer stealing much of the spotlight (Vaswani, 2017). The transformer is a machine-learning architecture that became a popular choice for natural language processing due to its ability to utilize parallelization on large bodies of sequential data. However, as with many machine learning architectures, the success is poorly understood. Google researcher Noam Shazeer stated, “We offer no explanation as to why these architectures seem to work; we attribute their success, as all else, to divine benevolence” (Shazeer, 2020). Yet, despite what little is known, these black box algorithms form the backbone of many modern generative AIs such as OpenAI’s ChatGPT, the fastest-growing application in history with over 100 million users only two months after launch (Hu, 2023). However, the success of technologies such as these has come with much controversy. Berkeley researcher Dan Hendriks in a paper studying the existential risks of AI noted that alpha fold, another generative algorithm used to analyze protein folding, has discovered new chemical weapons (Hendrycks et al., 2022). Action has even been taken by Open AI’s CEO Sam Altman to halt the development of AI, who practically begged Congress to regulate the industry stating on behalf of the company, “if this technology goes wrong, it can go quite wrong and we want to be vocal about that” (Kang, 2023). Clearly, there is a push to have a greater level of oversight for the development of modern AI algorithms but it remains an open question how this will come to fruition and what effect it will have.

Within the fervor of the last year, the transience of an old idea becomes more prominently mentioned. The unreasonable effectiveness of transformers as LLMs has caused many to question the nature of their intelligence. Deep learning experts such as Turing medalist Yoshua Bengio have cautioned that current AI systems might be on the cusp of becoming conscious stating: “Our analysis suggests that no current AI systems are conscious, but also suggests that there are no obvious technical barriers to building AI systems which satisfy these indicators” and further that, “conscious AI systems could realistically be built in the near term” (Butlin et al., 2023). Such claims are bound to evoke larger ethical conversations. In more recent news, OpenAI chief scientist Ilya Suskever announced the development of a new “Q*” algorithm which allegedly has problem-solving capabilities superior to humans, however, the company has decided that the details will remain secret for the time being (Radauskas, 2023). Stories like these bring to light serious ethical and safety concerns regarding the development of generative AI and investigating them will provide greater insight into the rapidly developing network of actants within the MLAIAN.

Conclusion:

While AI remains in a state of rapid development, the need to better understand how the underlying black box algorithms operate remains clear. As a whole, the field needs to place greater emphasis on studying the theories on which modern architectures are built as opposed to blindly navigating the frontier as has been done in recent years. Studying the problem-solving capabilities of GNNs serves only as one small part of a large issue, but will inevitably shed light on how to analyze these kinds of hard problems within machine learning.

References:

Beeman, D. (2001, October 30). *McCullogh-Pitts and Perceptron Models*. Www.cs.cmu.edu. https://www.cs.cmu.edu/afs/club/user/cmccabe/ecee.colorado.edu/~ecen4831/lectures/NNet2.html#:~:text=McCullogh%2DPitts%20Model&text=Each%20neuron%20has%20a%20fixed

Butlin, P., Long, R., Elmoznino, E., Bengio, Y., Birch, J., Deane, G., Fleming, S., Xu, C., Kanai, J., Klein, C., Lindsay, G., Michel, M., Mudrik, L., Peters, M., Schwitzgebel, E., Simon, J., & Vanrullen, R. (2023). *Consciousness in Artificial Intelligence: Insights from the Science of Consciousness*. https://arxiv.org/pdf/2308.08708.pdf

CSE 490H History Exhibit. (n.d.). Courses.cs.washington.edu. Retrieved December 16, 2023, from https://courses.cs.washington.edu/courses/cse490h1/19wi/exhibit/nlp.html

Fabien, M. (2018, November 20). *Maël Fabien*. Maël. https://maelfabien.github.io/deeplearning/Perceptron/

Garvin, P. (1967). *THE GEORGETOWN-IBM EXPERIMENT OF 1954: AN EVALUATION IN RETROSPECT**. https://aclanthology.org/www.mt-archive.info/Garvin-1967.pdf

Hendrycks, D., & Mazeika, M. (2022). *X-Risk Analysis for AI Research*. https://arxiv.org/pdf/2206.05862.pdf

Hu, K. (2023, February 2). ChatGPT Sets Record for Fastest-Growing User Base. *Reuters*. https://www.reuters.com/technology/chatgpt-sets-record-fastest-growing-user-base-analyst-note-2023-02-01/

Hutchins, J. (1996). From the Archives... ALPAC: the (in)famous report. *MT News International*, *14*, 9–12. https://aclanthology.org/www.mt-archive.info/90/MTNI-1996-Hutchins.pdf

Hutchins, J. (2005). *The history of machine translation in a nutshell*. https://aclanthology.org/www.mt-archive.info/10/Hutchins-2014.pdf

Kang, C. (2023, May 16). OpenAI's Sam Altman Urges A.I. Regulation in Senate Hearing. *The New York Times*. https://www.nytimes.com/2023/05/16/technology/openai-altman-artificial-intelligence-regulation.html

Lefkowitz, M. (2019, September 25). *Professor's perceptron paved the way for AI – 60 years too soon*. Cornell Chronicle. https://news.cornell.edu/stories/2019/09/professors-perceptron-paved-way-ai-60-years-too-soon

Lolha, E. (2022, May 30). *A Brief Explanation About Actor-Network Theory*. Medium. https://aprilliale.medium.com/a-brief-explanation-about-actor-network-theory-d8e9aad41a18

National Human Genome Research Institute. (2022, September 6). *Bioinformatics*. Genome.gov. https://www.genome.gov/genetics-glossary/Bioinformatics

Radauska, G. (2023, November 28). *OpenAI's Q* mystery: was coup sparked by major and concerning technological milestone?* Cybernews. https://cybernews.com/editorial/openais-artificial-intelligence-sam-altman-qstar-singularity/

Sato, R. (2020, October 16). *A Survey on The Expressive Power of Graph Neural Networks*. ArXiv.org. https://doi.org/10.48550/arXiv.2003.04078

Schmidhuber, J., Blog, A., & Ch, J. (2022). *Annotated History of Modern AI and Deep Learning*. https://arxiv.org/ftp/arxiv/papers/2212/2212.11279.pdf

Shazeer, N. (2020). *GLU Variants Improve Transformer*. Google. https://arxiv.org/pdf/2002.05202.pdf

Swaine, M. (2023, March 31). *Perceptron and the AI Winter*. Medium. https://medium.com/@michaelswaine/perceptron-and-the-ai-winter-c465d47da85

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A., Kaiser, Ł., & Polosukhin, I. (2017). *Attention Is All You Need*. https://arxiv.org/pdf/1706.03762.pdf

Xu, K., Hu, W., Leskovec, J., & Jegelka, S. (2019). How Powerful are Graph Neural Networks? *ArXiv:1810.00826 [Cs, Stat]*. https://arxiv.org/abs/1810.00826

Zhang, A., Lipton, Z., Li, M., & Smola, A. (2023). *13.4. Hardware — Dive into Deep Learning 1.0.3 documentation*. D2l.ai. https://d2l.ai/chapter_computational-performance/hardware.html

Zhang, B., Luo, S., Wang, L., & He, D. (2022, September 29). *Rethinking the Expressive Power of GNNs via Graph Biconnectivity*. Openreview.net. https://openreview.net/forum?id=r9hNv76KoT3