An Analysis of Machine Learning and Artificial Intelligence on Climate Change Through Capitalism

A Research Paper submitted to the Department of Engineering and Society

Presented to the Faculty of the School of Engineering and Applied Science

University of Virginia · Charlottesville, Virginia

In Partial Fulfillment of the Requirements for the Degree

Bachelor of Science, School of Engineering

Franceska Padilla Coo

Spring 2024

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

Pedro A. P. Francisco, Department of Engineering and Society

Introduction

Increasing carbon emissions and rising temperatures within the past century have elicited a response of concern from across the globe, and today, we face a global climate crisis. With everyone's habitat on the line, we hold global legislatures and our neighbors accountable to live consciously of our waste and production, establishing rules in order to preserve Earth and maintain its livability. However, the general population remains either unaware or unconcerned of how deeply corporations contribute to our planet's current crumbling conditions, and we very often overlook the contribution of engineers and scientists - especially those behind big data and machine learning. Although we have our basic calls for reusing, reducing, and recycling, education of environmentally conscious engineering is not part of many educational curricula. Engineers and scientists today, especially within the computing field, are geared towards rapid innovation and are very rarely called to incorporate sustainable approaches. Therefore, the lasting impact of their inventions and improvements tend to be the last thing on their mind. Rather, the prioritization is time and money. Oftentimes, engineers want minimum cost in production and maximum profit, and research and development centers - especially in corporate settings- encourage this. On the legal front, the information we have about corporate activity is complex and confusing. There is a lack of information and understanding, all the way up to our legislatures, both about how technology works and how to contain it to prevent further damage. It is one reason why we fail to either implement or generate actions within the legislature to restrict larger, corporate actions contributing to the climate crisis. Some of this ignorance is also intentional and protected. Ignorance allows the continued utilization of rapid data processing from big data and artificial intelligence and therefore faster profit.

I wish to address the environmental impact of machine learning and artificial intelligence (AI). Amidst the drastic and urgent calls for more environmentally friendly actions as climate change heightens, society has continued to overlook the strenuous demand of resources required to support machine learning and AI for the sake of their abilities and economic benefits. Across the globe, we see countries and leadership prioritize immediate economic profit, turning away from environmental and sustainability efforts in order to avoid even temporary monetary loss. While virtual from a consumer end, big data and machine learning are backed by hardware contributing to the global climate crisis and are major sources of energy consumption. We can see and measure economic gain, but we cannot see the implications of unsustainable machinery immediately. Computer technology and the methods in place to produce it are repeatedly accepted as a black box, making its contributions to the environmental crisis harder to address and allowing it to continue its environmental harm.

Background and Significance

The computing field has taken steps to remedy their contribution to climate change. In the 1990s, the Energy Star Programme demanded for more urgent environmental action. Green computing, which attempts to make computer production more cognizant of the environmental impact of computational technology production on the environment and therefore reduce the rapid usage of resources and provide a more sustainable approach to computer production, was a response to the Programme (Rai, et al., 2023). This study of green computing not only included more sustainable materials with which to build computers out of but also how to code and use computing materials efficiently and as necessary and how more sustainable computing approaches are financially beneficial to those involved. Despite the field of green computing emerging from the growing environmental crisis, it has not kept up with the pace of the rising temperatures. Because of it being a more recent field of study, it has yet to trickle down and completely pierce into the main computer science field as its own topic to be incorporated into educational practices, even higher education. Additionally, some scholars like Brevini and Lippert speculate that green computing and other attempts at reducing carbon emissions in technological production pose a threat to many producers as its initial implementation could potentially harm the immediate economy and lead to, though temporary, financial losses (Brevini, 2020; Lippert, 2016).

Even this one instance demonstrates a hesitance towards efforts making technology greener because it would require either a pause in production or a change in current systems, resulting in a financial loss for producers, though temporary. The attitude towards machine learning and artificial intelligence heavily mirror this hesitance to adopt sustainable methods at the risk of temporarily losing profits. The fascination surrounding machine learning and AI has skyrocketed their usage to levels that prevent the generation of effective plans of action to handle their environmental contribution. Because of the public's demand for machine learning and AI capabilities to maximize and grow the performance of the economy, society overlooks and diminishes the environmental impact of machine learning and AI.

Machine learning and AI have enabled society to accomplish many things. The ability to predict outcomes has drastically improved, and informed decision-making recently often stems from these two rising fields. Machine learning and AI base themselves on mass amounts of data, which is incredibly powerful in the metaphorical sense for enabling society with new capabilities, and in the literal sense, in terms of resource demands. Like other computer technologies, production contributes a significant amount to carbon emissions. However, machine learning and AI eat up even more energy resources due to the databases that store the data used for teaching algorithms and due to the power required to run such high-energy actions. “Data mining and computation evaluations of persons and corporations have far-reaching environmental costs” (Brevini, 2020, 1). The production of machine learning and artificial intelligence mirrors that of computer production in that it generates a significant amount of carbon emissions, but because of how resource-demanding machine learning and artificial intelligence are in terms of power and physical storage, machine learning and AI contribute significantly more to climate change and carbon emission production. Even more, data collection increases exponentially, and the growth requires the construction of physical storage and more physical computer production. The contributions of machine learning and AI are so great that they eclipse their negative environmental effects, whether intentionally or not.

The importance placed on machine learning and AI means that society will probably forgo the health of the environment for the sake of AI’s powers and also because caring about the environmental effects would slow down or minimize the rapid profit growth they have enabled. Society wants to maximize their benefit rather than scrap the technology. This is reflected in Velkova’s research about the accessibility of machine learning and AI with a paywall (Velkova, 2016). Society’s high regard of AI relies on the easy access to that type of machinery and function, where easy access especially refers to pricing. Society holds free and accessible AI at a higher position than AI locked behind a price. This indicates two things: that society tends to want the most value from their investments and that economy/financials, artificial intelligence, and then the environment, best reflect society’s current order of priorities. Notions like these impact our future as the climate crisis worsens, leaving society to “again find themselves at an important crossroad, where decisions taken today and behavioral patterns taking hold now may shape developments and environmental impacts of the technology over the next decades" (Konig, et al., 2022, 1).

Methodology

In order to analyze society's treatment of machine learning and AI, I use the black box theory. The black box theory covers the issues of ignorance about machine learning and AI's environmental impact (Brevini, 2020). From the engineering perspective, they lack an understanding of the environmental impact of the technology currently in use and of future innovations. If we remove this black box by incorporating sustainable practices and calls to action throughout education, we might witness a change in attitude from computer engineers and scientists. However, with this black box hindering the details of computers and their impact on the environment, we cannot expect engineers and scientists to understand their impact and know how to improve their methods to be less environmentally harmful.

The black box theory also explains corporate actions. As mentioned before, corporations can take advantage of the black box of machine learning and AI. Through it, they can avoid accountability and redirect attention to profits. They also are not obligated to demonstrate their environmental impact, though public pressure highly encourages it. Through this way, a general lack of knowledge from either the corporate world prevents people from understanding what is at stake, what can be done, and who to hold accountable.

Based on initial research on the topic, the value of AI depends on how much the public thinks it will help them and how useful the results of machine learning and AI can be to them. Evidence will be gathered using literature, reports, and media. Literature, reports, and media can better capture the different feelings across the globe over surveys and interviews which can only really be beneficial locally. Literature, reports, and media also provide information on the economy, AI, and the environment, which may be too much to ask of one person. The societal perspective is preferred in this context as the environment cannot change with just one person's actions, but also value is derived from a collective audience. The information gathered on AI and its value in the context of the economy and of its impact on the environment will be analyzed using case studies, sustainability, and consequences. The case studies provide context on values. Sustainability corresponds with the environmental aspect, especially as this topic pervades over some time, and consequences refer back to utilitarian ethics. Consequences will allow an analysis of what society might find valuable.

Literature Review

Brevini covers the topic of the lack of knowledge behind machine learning and AI in terms of functionality but also its environmental consequences (Brevini, 2020). In her paper "Black boxes, not green: Mythologizing artificial intelligence and omitting the environment", Brevini explains how the use of the Black Box theory, which is often used in terms of electronic technology, applies to not only the functionality of machine learning and artificial intelligence but also extends to their environmental impact, which is what makes it dangerous within the current economic climate (Brevini, 2020). For Brevini (2020), there is a lack of information surrounding how machine learning and AI work, what kind of environmental waste and consequences arise from their usage, and also how corporations and businesses using this machine learning and AI fail to report on how they plan to be held accountable for their usage. "Unfortunately, the carbon footprint of AI-powered algorithms is not only largely absent from public discourses on AI developments, but often it is neglected in the academy," Brevini says

(Brevini, 2020, 3). In her article, Brevini (2020) mentions that even academia lacks coverage of green computing and talks about environmental impact surrounding development and usage of machine learning.

Velkova also touches on the same sustained ignorance as Brevini (2020) but does so on a corporate and legislative level in order to demonstrate how corporations and legislations can manipulate the public's view on machine learning and AI (Velkova, 2016). Velkova (2016) focuses on case studies of efforts to decentralize data centers and their waste production in order to imbed the importance of data generation in households by creating urban areas dependent on data centers and data furnaces to power their local electricity and heating. Unlike Brevini (2020), Velkova (2016) does not talk much about how academia plays an important role in teaching data engineers and the tech industry about proper environmental efforts. Instead, Velkova (2016) pins the blame on data engineers, corporations, and legislatures as powerful figures who manipulate the perception of waste heat, "the major agent of disruption from within the data center and [threatens the existence of data centers'] power through the threat of network instability, transience, and decay, into something valuable in order to profit from it (Velkova, 2016, 5). Although not mentioning Brevini's (2020) perspective, Velkova's (2016) argument complements the later half of a timeline in terms of people involved with the process of data production and management. The lack of coverage for negative consequences of machine learning finally starts affecting the surrounding area, and in order to maintain this technological advancement and continue moving forward instead of sparking criticism and an overhaul of a relatively new implementation, those in power, the data engineers, corporations, and legislators, must restructure the narrative. They failed to account for the consequences and instead pivoted.

According to Velkova, they “[redefine] waste into a desirable commodity” in order to “extend the sphere of capitalist production and the digital economy” (Velkova, 2016, 2).

Lippert meanwhile touches on this power play and its effects (Lippert, 2016). In “Failing the market, failing deliberative democracy: How Scaling up corporate carbon reporting proliferates information asymmetries”, Lippert (2016) explains how corporations and their usage of data has become so interlaced and deeply embedded in society. Like Velkova, Lippert (2016) seconds this idea that this corporate use of data and their lack of transparency about their environmental footprint has ended up hurting society. Lippert (2016) combines both Brevini’s (2020) and Velkova’s (2016) arguments as he talks about a lack of transparency, the lack of knowledge from the public about corporations and their carbon footprint usage. Lippert, however, also included a specific breakdown and analysis of data from a company he refers to as GFQ (Lippert, 2016, 3). While Velkova (2016) talked generally about procedures and attempts at mitigating waste heat, she never fully analyzed a specific company and their energy production. This is where Lippert’s paper stands out. Lippert (2016) analyzes GFQ’s carbon footprint report and also compares it with their environmental policies and plans to hold them accountable, another type of case analysis that Velkova (2016) and Brevini (2020) fail to undertake. However, the other two articles were very theory heavy and expounded on topics, which Lippert’s (2016) article cannot afford to do as it focuses on unhashing a specific case. Lippert (2016) also does not focus on machine learning and AI based corporations. Still, Lippert’s (2016) article proves important in showing an actual case in which a company lacks transparency and twists the narrative in order to avoid accountability for their carbon footprint.

Together, Brevini (2020), Velkova (2016), and Lippert (2016) all cover how those with influence can hide information surrounding machine learning and AI in order to keep harnessing its power. All three authors talk about how those with influence can twist the narrative in their favor and avoid environmental accountability. The three authors are essential in shaping my argument as Brevini (2020) explains how the hiding of information has allowed the emergence of machine learning and AI's continued contribution to the climate crisis, Velkova (2016) demonstrates how public opinion can be manipulated to believe that machine learning and AI's contribution to the climate crisis is minimal and being addressed, and Lippert (2016) demonstrates the lack of corporate responsibility and how it hurts both consumers and citizens.

Discussion and Results

My thesis aims to suggest that the lack of common knowledge and academic coverage of how machine learning and artificial intelligence (AI) are executed, especially at a corporate level, allows much of the tech industry to twist public opinion and either minimize the importance of sustainable and eco-friendly computing or trick the public into thinking that efforts towards reducing the carbon footprint and destructive environmental effects are being taken. There are several components to this argument that fit together in order for it to take shape. First, machine learning and AI eat up at many energy resources and produce heat and a carbon footprint. Second, this carbon footprint is not well known or acknowledged academically or in many public forums because of a lack of knowledge and coverage in academia and a lack of transparency from the tech industry. Third, human opinions about machine learning and AI are incredibly positive. Fourth, this incredibly positive view of AI is due to the powerful systems and infrastructures that most financially benefit from machine learning and AI usage swaying public opinion to continue the cycle of profit.

Firstly on the topic of machine learning and artificial intelligence eating up at resources, there are several case studies demonstrating this. In 2019, a study at the University of Massachusetts Amherst reported that a single big language model's carbon footprint amounts to the same amount of carbon footprint as 125 round-trip flights from New York to Beijing (Hao, 2019). Other figures include how a data centers' energy usage averages 200TWh, which is more than Iran's energy consumption (Brevini, 2020). Velkova (2016) also mentions how waste heat, heat generated from data centers, is significant enough to power cities. This demonstrates that machine learning and artificial intelligence specifically contribute a significant amount to carbon emissions.

Though these numbers demonstrate a significant contribution of carbon emissions from machine learning and artificial intelligence, there is also the indication of a lack of institutional coverage and knowledge surrounding the topic. Though the climate crisis and machine learning have been around for a long period of time, around 1990s at the earliest, the most recent movements for environmental accountability began in the recent five years of around 2019. Green computing emerged around the 1990s, but the incorporation of these ideas and even the very teaching of green practices in computer science practice is rare (Rawat et al., 2023). Like Brevini (2020) mentioned, the discussion of machine learning and AI's environmental impact is rarely touched on in academia. However, Brevini (2020) also mentions a lack of public discourse surrounding the topic of machine learning and AI, which probably also ties into the lack of coverage in schools and vice versa. The black box theory is incredibly important in explaining this phenomenon. How machine learning and AI works is already relatively misunderstood or very simplified. This extends to its environmental implications.

The lack of knowledge and transparency around machine learning and AI's environmental impact is a necessary layer in the positive public perception of them. With their production, pricing, use of resources, and environmental impact to be hidden, machine learning and AI can continue their progress unhindered. For the most part, much of the public views machine learning and AI as incredibly useful.

The positive reputation of AI, however, is also tied very closely to the black box surrounding technology. Velkova (2016) states, "Simply put, those who have the power to define what is valuable in a society … are those who produce and maintain difference." With big tech, corporations, and legislatures leading economies and the world, they can greatly influence the narrative of positive machine learning and AI. Brevini (2020) mentions how corporations release statements and goals to meet environmental goals and needs and align with reputable human rights organizations like the United Nations in their goals. However, using GFQ, Lippert (2016) demonstrates how, like Brevini (2020) mentions, these institutions are not truly transparent of their carbon impact. "Environmental accounting provides decision-makers with data. However, whether that data actually represents what they claim to represent is obscured. Only because data is provided, we cannot assume that the data approximates well what they supposedly represent" (Lippert, 2016, 6). This means that even if companies do show a portion of their carbon emission data, they can still manipulate the narrative and therefore avoid accountability. For example, Amazon only recently started being transparent about their carbon footprint. Amazon is a double edged sword in terms of carbon emissions as their main business is to sell, contributing to consumption, which goes against efforts for climate change that aim at reducing consumption. Additionally, Amazon has their technology sector Amazon Web Services (AWS), which is known for having high compute power for machine learning. From 2010 to 2021, Amazon refused to report to Carbon Disclosure Project's carbon questionnaire, one instance of how Amazon hid their carbon emissions despite being one of the biggest producers and corporations in America (West, 2022). While Amazon has improved since then, becoming one of the leading corporations with the Climate Pledge, we also witness a lack of complete transparency with their AWS carbon emissions. While Amazon's 2022 carbon report talks about AWS carbon reduction efforts, they lack the detailed charts and numbers that their manufacturing front includes (Amazon, 2022). With major powers being able to implement this strategy of cherry-picking what they want, they can give people the illusion of progress. This illusion of progress is incredibly dangerous as the body to hold these powers accountable are deceived and cannot check that environmental concerns are addressed. With such a time sensitive issue like climate change, this deception proves incredibly detrimental.

Conclusion

I hope to show the connection between the capitalist systems in place and the impending environmental emergency and their specific relationship with the machine learning and artificial intelligence industry. In order for machine learning and artificial intelligence to progress at the pace it has been progressing the past thirty years, it will require a lot of resources. It will require natural resources, energy, manpower, and finances. It is an incredibly profitable field, but with public outrage, its momentum could stall. Public outrage for this industry can come in many ways, but one way we see taking shape is its lack of environmentally friendly practices. The capitalist system in place with engineers, corporations, and legislatures profits from machine learning and AI, and to maintain their power, they must protect the perception of machine learning and AI. We must view the many moving parts of the machine learning and AI industry as interconnected because we need to address this as soon as possible to make a difference, and the environmental impact of machine learning and AI has far too much to catch up with to continue delaying the lack of awareness or urgency on this issue.

Reference List

Amazon. (2022). Amazon 2022 sustainability report. Retrieved from https://sustainability.aboutamazon.com/2022-report.

Awasthi, K., & Sameer Awasthi. (2023). Green Computing: A Sustainable and Eco-friendly Approach for Conservation of Energy (A Contribution to Save Environment). Sustainable Computing: Transforming Industry 4.0 to Society 5.0. 319-333.

Brevini, B. (2020). Black boxes, not green: Mythologizing artificial intelligence and omitting the environment. Big Data & Society, 7(2). https://doi.org/10.1177/2053951720935141.

Gough, C., OAPEN, O'Reilly Online Learning: Academic/Public Library Edition, & DOAB: Directory of Open Access Books (2015). Energy Efficient Servers: Blueprints for Data Center Optimization. New York, NY: ApressOpen.

Hao, Karen. MIT Technology Review. (2019, June 6). Training a single AI model can emit as much carbon as five cars in their lifetimes. Retrieved from https://www.technologyreview.com/2019/06/06/239031/training-a-single-ai-model-can-emit-as-much-carbon-as-five-cars-in-their-lifetimes/.

König, P. D., Wurster, S., & Siewert, M. B. (2022). Consumers are willing to pay a price for explainable, but not for green AI. Evidence from a choice-based conjoint analysis. Big Data & Society, 9(1). https://doi.org/10.1177/20539517211069632.

Lippert, I. (2016). Failing the market, failing deliberative democracy: How scaling up corporate carbon reporting proliferates information asymmetries. Big Data & Society, 3(2). https://doi.org/10.1177/2053951716673390.

Rawat, Y., & Priti Rai. (2023). Trends and Challenges in Green Computing. Sustainable Digital Technologies: Trends, Impacts, and Assessments (1st ed.). CRC Press. https://doi-org.proxy1.library.virginia.edu/10.1201/9781003348313.

Sharma, R., Śarmā, D., Bhatt, D., & Pham, B. T. (Eds.) (2022). Big Data Analysis for Green Computing: Concepts and Applications. Boca Raton, FL: CRC Press.

Strubell, E., Ganesh, A., & McCallum, A. (2020). Energy and Policy Considerations for Modern Deep Learning Research. Proceedings of the AAAI Conference on Artificial Intelligence, 34(09), 13693-13696. https://doi.org/10.1609/aaai.v34i09.7123.

Takiguchi, H., & Morita, K. (2009). Sustainability of silicon feedstock for a low-carbon society. Sustainability Science, 4(1), 117–131. https://doi.org/10.1007/s11625-009-0069-1.

Velkova, J. (2016). Data that warms: Waste heat, infrastructural convergence and the computation traffic commodity. Big Data & Society, 3(2). https://doi.org/10.1177/2053951716684144.

West, E. (2022). “Decoding Amazon’s Climate Pledge: Public Relations and the Platformization of Governance.” In Amazon: At the Intersection of Culture and Capital, edited by Paul Smith, Alexander Monea, & Maillim Santiago, pp.269–286. Rowman & Littlefield.