Can machine learning accelerate sustainability research?

Perspective Piece

Written by Lucy Kim on May 2, 2024

Abstract

As of 2024, machine learning and artificial intelligence is accelerating progress in many fields, including science. In sustainability research (i.e. research with the goals of meeting the needs of the future as well as the current generations), where the bottleneck is quite often in acquiring and analyzing sufficient amounts of data, machine learning may be used to push the current boundaries of research. This paper highlights transfer learning, a machine learning method developed in 1995 to help machines learn from available, reliable data to perform new tasks on unfamiliar data sets. I explore the ways transfer learning can be applied in sustainability research by looking at two prime examples of agricultural remote-sensing research and alternative energy research in large language models. The paper concludes with commentary on the 2023 petition to freeze AI research, all to answer the question of how machine learning can interact with research in climate change, ecology, natural resources and the most basic needs of humankind.

Introduction

a. The Climate Crisis and the Scientific Solution

The planet is warming significantly, and a majority of the temperature increase is caused by human action. These actions include fuel consumption, transportation, agriculture, and the existence and construction of buildings. Moreover, despite global efforts to cut emissions, humans continue to emit *increasing* rates of greenhouse gasses which cause global warming. The IPCC's annual results on climate change in 2023, noted that warming has already caused extreme heat and weather events, loss of biodiversity, food insecurity, diseases, and displacement. At the time of the report's publication in 2020, humans are estimated to have already caused a 1.07 degree increase from pre-industrial levels (pre 1850). At this rate, warming is projected to exceed 1.5°C by the end of this century (2100), making "harder to limit warming below 2°C" (the Paris Agreement tipping point). Thus, it is critical to change, on behalf of our current and future generations' wellbeing, now.

As a species, humans have two options moving forward: continue (scientific, economic, and other) development or regress to a simpler mode of living. I outlined these options because I do not think it is a quite obvious decision as most academics make it out to be. The sources of emissions outlined in the previous paragraph are, at the same time as they are potential downfalls of humanity, milestones of the industrial revolution, scientific progress, and societal development. In fact, the IPCC results paper clearly states that the highly developed countries cause the most warming, and that the disparity exists even at the household level, where the top 10% of highest-per-capita emissions are responsible for 34-45% of total household emissions [12]. Prosperity, economic growth, acceleration, and development, should be examined with suspicion in their alignment to climate change.

Sustainable development, therefore, sounds like an oxymoron, as some critics have and continue to argue: as development is linked with more consumption [13]. Referring to the 1987 Brundtland Commission's definition, sustainable development is "development that meets the needs of the present without compromising the ability of future generations to meet their own needs". Thus, the term envisions a mode of development that does not deplete natural resources and cause environmental damage. The role of academia in this vision is to push forward with sustainability research: an umbrella term that includes the interdisciplinary research areas of ecology, alternative energy, earth sciences, policy, systems engineering, and policy.

Sustainability research assumes that the same mindsets and skill sets that brought humanity to this point of warming will help reverse the effects. This paper does not go further in depth to challenge this assumption, but instead lays out how a recent technological discovery might be a promising key turning point that is uniquely capable of bringing about change. This technology has been a multiplier of progress and discovery across all industries and is critical to align with sustainability goals.

b. Machine Learning

Research by learning from data has been achieved using statistical models since the 1800s. However, with the improvements of GPU's (i.e. computing power) and development of backpropagation algorithms in the 1980s, the world of statistical modeling became changed forever. Machines can now detect patterns in massive amounts of data and predict trends in unseen data, sometimes in ways that are literally incomprehensible to humans. This is the field of machine learning and artificial intelligence (AI), accessible now to the public in the forms of AI chatbots, ChatGPT, and computer generated products like art, music, and text.

Machine learning methods have already been applied effectively in energy consumption and grid distribution optimization, solar and wind energy prediction, and climate modeling. However, most AI models require large amounts of data to train and establish their settings (i.e. parameters), which is a problem in the space of sustainability, where data collection is inherently slow, difficult, or expensive. Thus, this paper will dive into specific mechanisms that would help adjust machine learning to sustainability research, namely: transfer learning for general machine learning applications, and few-shot prompting for Large Language Models (LLMs). This section is written to provide a brief introduction to the science behind the two, before delving into their applications in sustainability in the case-study section.

i. Transfer Learning

As a thought experiment, imagine you are learning the guitar for the first time. In this situation, you have no prior knowledge of guitar but you have the choice to import into your brain the knowledge of Yo-Yo Ma, the acclaimed cellist. Would you take that choice? Hopefully so! Although Yo-Yo Ma may not have been pre-trained to play rock-and-roll, he *has* been pre-trained to understand pitch, rhythm, sight-reading, finger-coordination, among other skills that are related to the target task. Using his knowledge base would be much more efficient than starting from scratch, especially if you only have access to a few classes with a teacher.

Transfer learning takes this principle to machine learning: given a scarcity of training data, take a model originally trained on some dataset for some task and apply it on *your* target dataset for *your* target task. In more technical words, the pre-training of the model will differ from its test environment by: source, task, or distribution. The field adopts the language of statistics, to define domain and task space, and divides and classfies transfer learning methods depending on whether it is the marginal, conditional, or label probabilities that differ between the source and target.

The motivations for transfer came about in a workshop in 1995 (NIPS-95), specifically, the need for machines to retain and reuse learned information. Transferability was developed in parallel under many other names and tangential ideas, including multi-task learning. Hallmark survey papers, for reference, are Pan and Yang in 2010 and a followup in 2015 by Weiss et al, which summarized the 700 papers published since the first, signifying the importance and rapid growth of the field [4,7]. The 2010 paper highlighted research on what to transfer and how to transfer, and applications to image classification, sentiment classification, wifi localization, cross-language text labeling. The 2015 paper identified areas of interest in the popularity of one-stage over two-stage solutions (i.e. doing domain adaptation and final classifier learning at the same time) as well as emerging interests in: heterogeneous transfer learning, negative transfer, and reinforcement learning (unlabeled source and target data).

ii. Large Language Models and Few-shot Prompting

Large language models (LLMs) are large. They are a type of artificial neural network trained on millions of text files to set billions of variable parameters. Simply, they work by representing words as groups of numbers, and finding patterns in the training data number sequences to help them predict the next sequence of numbers and thus words. Thus, machines gain a human-like understanding of grammar and syntax. A common application of LLMs is text generation in commercially available products: OpenAI's ChatGPT, Google's Gemini, and Meta's LLaMa.

However, LLMs can be fine-tuned and trained for a specific task for users with specific goals. Here, the model can also face the problem of few training data. A common method used to address this lack, separate from but related to transfer learning, is few-shot prompting. A user can prompt a pre-trained LLM like ChatGPT4 to do a task, providing a few (less than 12) training examples. For example, a student is assigned the task of summarizing news articles in the style required by a picky college professor. Although her GPT of choice gives a very basic summary at first, she discovers that she can feed a couple successful student-examples (full article along with exemplary student summaries provided by the professor) to her GPT, and it will spit out the results she needs. The GPT will adjust its parameters using supervised techniques like gradient descent to minimize the difference between its first-intuition output and the student examples.

Because of their efficiency and accuracy, both transfer learning and few-shot prompting shows promise in helping solve some current gaps in sustainability research. A few applications will be highlighted next.

Case Study Discussion

Last month, I conducted stake-holder interviews with experts in machine learning with the hopes of better understanding its role in sustainability. In one of these interviews, I talked to a data scientist at Climax Foods, a spinoff from UC Berkeley that used machine learning models to invent new plant-based cheeses. In his words, the field of food science did not yet seem "ripe" for machine learning, due to the large bottle-neck in data collection. The wet lab's creation and testing of cheeses could not reach the throughput and size needed in the dry lab for ML modeling. The startup closed earlier this year. This discussion shaped my search for areas of application that were ripe for machine learning and transfer learning methods: areas that could provide massive amounts of relevant training data even if there were scarcities in test data. I landed on precision agriculture, drone instruction, and nuclear fusion.

a. Precision Agriculture

When agriculture and drained peatlands are responsible for 10% of U.S. greenhouse gass (GHG) emissions and 5% of global GHG emissions respectively, responsible land use and farming practices will be key to the sustainability discussion, and monitoring the use is thus a required middle-step. Whereas collecting physical measurements of soil quality, crop coverage, and water levels is expensive, looking at satellite imagery of vegetation is cheap and scalable. One popular application of machine learning is to analyze images from remote sensors (i.e. long distance, no physical contact devices such as drones and satellites). A 2024 study used machine learning to predict vegetation coverage based on levels of groundwater and other environmental factors, combining 10 datasets for a total of 18 variables, some of which are based on satellite imagery and optimized sampling. The magnitude of data needed to feed this model is indicative of the bottleneck in using machine learning in sustainability: the "impracticality and time-consuming nature of field observations" [1].

As transfer learning methods are refined, they may be applied to fill in the gaps. A Stanford study in 2016 learned the socio economic/infrastructure indicators of poverty regions in Africa by using two transfer learning steps: one transferring a popular object classification model to a regression model predicting light intensity from daytime satellite images, and one transferring this light intensity prediction module to the final poverty prediction module [8]. Coupling transfer learning with remote sensing would allow for studies to overcome the slowness of field-observations and data collection.

b. Drones and World Sensing

Thus, remote sensors like drones may play an increasing role in climate sciences: not only can you get a better picture of land use, but drones can also help monitor trends in the weather, oceans, and arctics. In general, drones can help with sensing the multiple, chaotic patterns of the world. And interestingly, drones themselves are a place where machine learning practices of few-shot learning have a promising application.

Researchers at Yale are developing a drone system that can take natural language instructions: that is human speech. The drones use an image processor to create a text representation of their surroundings and combine the text with their human operators spoken instructions to prompt an LLM models for code that would control the drone's next steps. This system, called TypeFly, found that a critical part of the accuracy of the code was few-shot prompting, without which the accuracy would drop by more than 50% [10].

c. Nuclear Fusion

Dozens of years of papers, databases, and experiment logs are much easier for machines to sift through and process than human researchers. Researchers from Princeton, MIT, and Carnegie Mellon are using LLMs to help make real-time parameter changes to devices and reactors, doing their decision-making based off of the LLMs understanding of all previous runs of the same device or related uses of other devices [2]. Achieving nuclear fusion would generate energy without outputting greenhouse gasses, solving one of the largest sustainability problem areas. Such a discovery would easily offset the energy consumption needed by computation and machine learning models, which is a major issue both for the field of machine learning and for sustainability.related to the final topic of today's paper.

Conclusion

This paper would not be complete without including the context of the AI's own alignment with sustainability, and the ongoing tug-of-war between accelerated development of AI capabilities and the push for more AI regulation against potential risks.

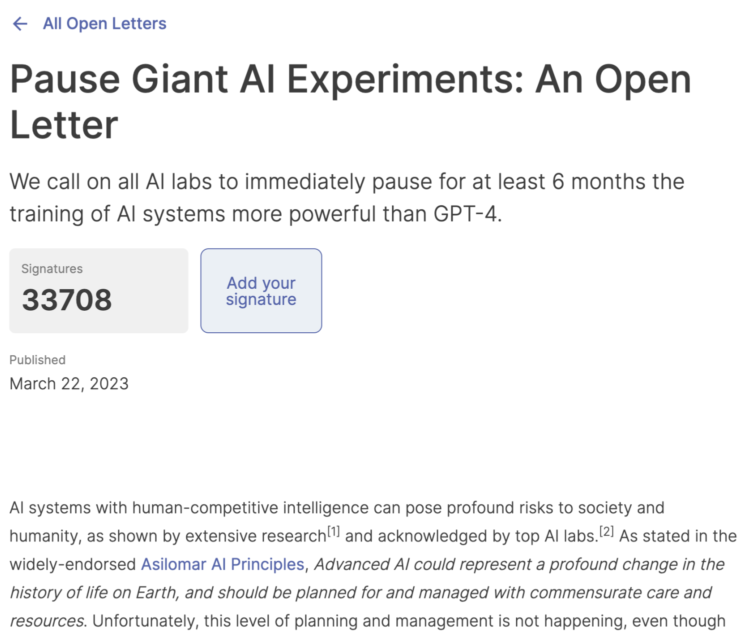

← All Open Letters

Pause Giant AI Experiments: An Open Letter

We call on all AI labs to immediately pause for at least 6 months the training of AI systems more powerful than GPT-4.

Signatures
33708

Add your signature

Published
March 22, 2023

AI systems with human-competitive intelligence can pose profound risks to society and humanity, as shown by extensive research[1] and acknowledged by top AI labs.[2] As stated in the widely-endorsed Asilomar AI Principles, *Advanced AI could represent a profound change in the history of life on Earth, and should be planned for and managed with commensurate care and resources*. Unfortunately, this level of planning and management is not happening, even though

Fig 1. Open Letter formed by The Future of Life Institute to pause AI research

In March of 2023, a group of high-profile academics, tech industry leaders, and AI experts signed a petition called "Pause Giant AI Experiments," which called for all AI labs to pause work for at least six months and collaborate with policy-makers during this pause to write regulation for the safe development of AI against potential risk scenarios [5]. Signatories included Professor Max Tegmark (MIT), Elon Musk, and Yoshua Bengio (2018 Turing Award recipient). Although, the petition has not yet been successful in pausing OpenAI and other AI labs, it brought to the attention of the public what's called the AI alignment, which petition's the corresponding policy paper defines as: the "development of technical mechanisms for ensuring AI systems learn and perform in accordance with intended expectations, intentions, and values"

[6]. Simply, the alignment problem puts a name to the concern that AI development may not be aligned with improving the human condition.

In this vein, I asked the question: does the act of developing AI itself pose a risk to future generations? Assuming that there is non-zero risk of AI models prioritizing any goal over human life (for example, through recursive self-improvement methods derailing models from their original values), there is a valid security concern with the ideas outlined in this very paper. Giving LLM-controlled robots all knowledge of the natural world from AI's perfected world-modeling algorithms would make them extremely agile in the natural world [9].

To conclude, sustainability is a promising area of application of machine learning that could vastly improve the human condition. However, many experts seem to agree that although we should be excited about continuing to make discoveries with AI tools, we should encourage the development of AI safety research and consequences.

Bibliography

1. Bai, Taiya, et al. "Controls of groundwater-dependent vegetation coverage in the Yellow River basin, China: Insights from interpretable machine learning." *Journal of Hydrology*, vol. 631, Mar. 2024, p. 130747, https://doi.org/10.1016/j.jhydrol.2024.130747.
2. Garnier, Darren, and Viraj Mehta. "Towards LLMs as Operational Copilots for Fusion Reactors." *NeurIPS2023 Workshop*, 2023.
3. Leike, Jan, et al. "Our Approach to Alignment Research." *OpenAI*, 24 Aug. 2022, openai.com/blog/our-approach-to-alignment-research.
4. Pan, Sinno Jialin, and Qiang Yang. "A survey on Transfer Learning." *IEEE Transactions on Knowledge and Data Engineering*, vol. 22, no. 10, 16 Oct. 2009, pp. 1345–1359, https://doi.org/10.1109/tkde.2009.191.
5. "Pause Giant AI Experiments: An Open Letter." *Future of Life Institute*, 21 Feb. 2024, futureoflife.org/open-letter/pause-giant-ai-experiments/.
6. *Policymaking in the Pause*, Future of Life, 12 Apr. 2023, futureoflife.org/wp-content/uploads/2023/04/FLI_Policymaking_In_The_Pause.pdf.
7. Weiss, Karl, et al. "A survey of Transfer Learning." *Journal of Big Data*, vol. 3, no. 1, 28 May 2016, https://doi.org/10.1186/s40537-016-0043-6.
8. Xie, Michael, et al. "Transfer learning from deep features for Remote Sensing and poverty mapping." *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 30, no. 1, 5 Mar. 2016, https://doi.org/10.1609/aaai.v30i1.9906.
9. Yudkowsky, Eliezer. "The Only Way to Deal with the Threat from Ai? Shut It Down." *Time*, Time Magazine, 29 Mar. 2023, time.com/6266923/ai-eliezer-yudkowsky-open-letter-not-enough/.
10. Chen, Guojun, et al. "TypeFly: Flying Drones with Large Language Model." *arXiv*, 8 Dec. 2023, https://doi.org/arXiv:2312.14950.
11. "Climate Change 2023." *AR6 Synthesis Report*, Intergovernmental Panel on Climate Change, 19 Mar. 2023, www.ipcc.ch/report/ar6/syr/.
12. "The Sustainable Development Goals Report." *United Nations*, United Nations, 10 July 2023, unstats.un.org/sdgs/report/2023/.
13. https://www.tandfonline.com/doi/full/10.1080/13504509.2016.1235624