AI AND THE UGLY ENVIRONMENTAL FOOTPRINT IT LEAVES BEHIND

CHRISTINE LI

On November 30th, 2022, OpenAI launched ChatGPT, the seemingly magical AI chatbot that can do anything from writing Shakespearean-style poetry and catchy song lyrics to teaching you to solve an iterated integral. In the weeks and months after its launch, ChatGPT experienced a meteoric rise. With this came an onslaught of criticisms about its implications, from concerns about its tendency to "hallucinate"—generating confident-sounding but completely false statements—to worries about its algorithmic bias, reflecting the entrenched inequities in its training data, to predictions that plagiarism would proliferate in academia (Hao, "What is ChatGPT"). ChatGPT represents a tipping point, testing the boundaries between human and artificial intelligence.

Meanwhile, humanity has brought the planet to the edge of another tipping point. Over the past several years, climate change has quickly morphed from a problem of the distant future into one that is immediate. As the world reels from increasingly destructive catastrophes, from historic storms in California to extreme droughts in East Africa, there is a flurry of conversation, international agreements, and pledges from corporations to make amends.

Amid the urgent search for better solutions, AI is seen as a key component in combatting climate change. In the same month that ChatGPT was released, the United Nations Environment Program (UNEP) published an article explaining the role of AI in tackling climate challenges. According to David Jensen, coordinator of the UNEP Digital Transformation team, AI has the potential to revolutionize a multitude of climate-related efforts, from planetary-scale "satellite monitoring of global emissions" to granular-level energy-efficient smart homes. For example, the World Environment Situation Room online platform launched by UNEP leverages AI to analyze "complex, multifaceted datasets" tracking atmospheric carbon dioxide concentrations, sea levels, and glacier mass in order to help policymakers make data-driven decisions ("How Artificial Intelligence"). As climate change is an immensely complex problem with an enormous number of shifting variables, AI is uniquely suited to streamline the process of collecting and analyzing climate datasets to build predictive models and provide invaluable insights into policymaking and mitigation strategies.

But there's a catch. Even though artificial intelligence may be pivotal in our fight against climate change, the technology itself is a massive culprit in accelerating the crisis that is humanity's greatest existential threat.

Lost in the flurry of excitement about the future potential of AI is the ugly reality of its immense environmental costs. As highlighted in the title of a recent *Bloomberg* article from March 2023, "Artificial Intelligence is Booming—So is Its Carbon

Footprint." Authors Josh Saul and Dina Bass explain that, due to the massive amounts of data used in the training process, and the fundamental structure of deep learning models, AI requires orders of magnitude more energy consumption than traditional forms of computing. However, the exact energy usage and carbon emissions of most AI models remain a mystery due to the lack of transparency from the large technology companies developing them.

Out of curiosity, I turned to ChatGPT itself and asked, "What is your carbon footprint?" But, like a seasoned politician trained to weasel its way out of a thorny topic, ChatGPT gave a cautious reply coated in ambiguity: "As an artificial intelligence language model, I do not have a physical form and therefore do not have a direct carbon footprint. However, the computers and servers that power my operation do require energy, which may contribute to carbon emissions" ("What").

In the absence of any openly reported emissions data from OpenAI, independent researchers in academia and non-profit organizations have attempted to calculate the carbon footprint of GPT and other large AI models—and the results are terrifying. The earliest work in this area stems from a seminal 2019 paper by Strubell, Ganesh, and McCallem from the University of Massachusetts, Amherst. Presenting their study at the Association of Computational Linguistics conference that year, they stunned the computer science community by revealing that the process of training the BERT transformer model—a type of large AI model with over 200 million parameters—emits more than 626,000 pounds of carbon dioxide (Strubell et al. 3645). This is equivalent to "nearly five times the lifetime emissions of the average American car. . . [including] manufacture of the car itself" (Hao, "Training"). Since then, the size of large AI language models has grown exponentially, alongside their energy usage and carbon emissions. In 2021, OpenAI released GPT-3. With 175 billion parameters (Heaven), this large generative language model is nearly 100 times larger than BERT from two years prior. Researchers from the AI startup Hugging Face estimated the training emits more than 500 metric tons (over one million pounds) of carbon dioxide (Luccioni et al. 7), which is around 610 one-way direct flights from New York to Paris. By the release of GPT-4 in March 2023, the parameter count ballooned to a purported 170 trillion, representing another hundred-fold increase from their previous model two years prior (Zaveria). Assuming the computational costs required scale proportionally, this entails that a single training process of the new model emits the equivalent of 61,000 transatlantic flights.

Perhaps the most terrifying thought is that this is only a fraction of the computational resources used in the entire lifecycle of developing a model. As Strubell explains, "Training a single model is the minimum amount of work you can do" (qtd. in Hao, "Training"). In practice, the full development pipeline involves many, many more rounds of training and fine-tuning. Strubell and her colleagues estimate that building and testing a "final paper-worthy model required training 4,789 models over a six-month period" (qtd. in Hao, "Training"). In the case of a commercialized product

like ChatGPT, model training was just the first step. In the two months following its launch, it amassed over 100 million unique users at an unprecedented rate of growth (Milmo). Generating responses to hundreds of millions of prompts requires even more computational power. It is almost beyond comprehension just how much power GPT consumes.

As it becomes increasingly clear that AI has a monstrous appetite for energy, OpenAI is desperately trying to hide this ugly truth by keeping silent on the issue. An article by *MIT Tech Review* shortly after the launch of GPT-4 in March 2023 calls the new model "the most secretive release the company has ever put out, marking its full transition from nonprofit research lab to for-profit tech firm" (Heaven). When asked just how large the model is, the company's chief scientist, Ilya Sutskever, claimed that it is something he "can't really comment on at this time" because "[i]t's pretty competitive out there" (Heaven).

The field of AI research is split into two major theories regarding artificial general intelligence (AGI): one is that AGI can be achieved by simply scaling existing models, and the other is that the current approach, through deep learning, is fundamentally insufficient. OpenAI has doggedly pursued the former path in its quest for AGI. As evident through the exponentially growing size of its GPT language models, most of the breakthroughs by OpenAI "have been the product of sinking dramatically greater computational resources into technical innovations developed in other labs." For the leadership team, this computation-driven strategy is their "primary competitive advantage" over other research labs (Hao, "Messy, Secretive Reality"). Caught in a race with other giant tech companies such as Google and Meta to build the best AI model, the company has grown increasingly guarded about its research process and increasingly closed off to the public. Ironically, OpenAI is no longer so open.

It wasn't always this way. Originally founded by Sam Altman as a non-profit research organization in 2015, OpenAI had the goal of democratizing AI. Their core charter states, "Our primary fiduciary duty is to humanity." But the document later reveals, "We anticipate needing to marshal substantial resources to fulfill our mission," and "we expect that safety and security concerns will reduce our traditional publishing in the future" ("OpenAI Charter"). This language already insinuates their shift in priorities. In 2019, when OpenAI initially announced the $1-billion investment from Microsoft and the transition to a capped for-profit model, the leadership team claimed that "any commercialization efforts would be far away." But, in an internal meeting just months later, Altman's bottom line was clear: "OpenAI needs to make money in order to do research—not the other way around" (Hao, "Messy, Secretive Reality"). In February 2023, OpenAI began charging users twenty dollars per month for premium subscriptions that offered faster results through ChatGPT and API access to GPT-4. But, to satisfy its insatiable hunger for computational resources, OpenAI also agreed to give up almost half of its profits to Microsoft in exchange for access to the company's Azure cloud computing network. Clement Delangue, the

CEO of Hugging Face, which develops open-source AI language models, fears this type of investment is leading to "cloud money laundering" (Bass). The easy access to computing resources crushes incentives to develop more efficient, environmentally friendly solutions. The trend toward larger and larger models creates "unsustainable use cases for machine learning" that threatens both the development of AI and the future of the planet (Bass).

But it does not have to be this way. As outlined by a correspondence in *Nature* from March 2023, the carbon emissions of large AI models can be drastically reduced by "tailoring the structure of the model and by promoting energy-efficient hardware and the use of clean energy sources" (An et al. 586). For example, Strubell et al. discovered that the fine-tuning process called neural architecture search, used to increase the final accuracy of the model through exhaustive trial-and-error, had "extraordinarily high associated costs for little performance benefit" (Hao, "Training"). Eliminating this step in the BERT model reduced the carbon footprint to less than 400 times that of the original. Another study found that the open-source BLOOM model developed by Hugging Face produced around 25 metric tons of carbon dioxide, which is just five percent of the estimated carbon footprint of GPT-3 (Luccioni et al. 7). Even though the two models are roughly the same size, with around 175 billion parameters, BLOOM has a much smaller carbon footprint because it was "trained on a French supercomputer powered mostly by nuclear energy." On the other hand, "[m]odels trained in China, Australia, or some parts of the US, which have energy grids that rely more on fossil fuels, are likely to be more polluting" (Heikkilä). With the right incentives and regulations, large AI models can be optimized to be much more efficient through environmentally conscious engineering practices.

However, in our capitalist economy where researchers at for-profit companies are locked in cutthroat competition, they are forced to take whatever steps are necessary to produce better AI results. In this race, any considerations about the consequences on the climate get completely thrown out of the picture. OpenAI "chases a computationally heavy strategy—not because it's seen as the only way to AGI, but because it seems like the fastest" (Hao, "Messy, Secretive Reality"). It's not that AI researchers and scientists are inherently evil or don't care about the environment. In fact, most are good people who genuinely believe in AI's potential for helping humanity. I recall, in November last year, I had a FaceTime call with a longtime family friend of mine. Two months prior, he had completed his doctorate at MIT and joined OpenAI's research team working on GPT-3. When I asked him why he chose this path, he explained that he had actually turned down lucrative offers from quant firms to pursue AI research in academia, and he was truly drawn to OpenAI's founding mission of advancing artificial intelligence for the benefit of humanity.

Humanity's greatest challenge is climate change, and AI indeed has the potential to be of great service in our search for a solution. But, as it stands today, AI inflicts far greater harm to humanity and the planet at large with its unabated carbon footprint.

In order for researchers to not just pay lip service to noble aspirations of advancing AI to help humanity, two key issues must be addressed: the lack of transparency and the misalignment of incentives.

The root of these complications in contending with AI's carbon footprint problem stems from the privatization of AI research. Strubell explains that "training huge models on tons of data is not feasible for academics—grad students especially, because we don't have the computational resources" (qtd. in Hao, "Training"). Since the publication of her study on the carbon footprint of AI models, which she wrote as a grad student, Strubell has joined the Computer Science department at Carnegie Mellon University as an Assistant Professor, but her decision to remain in academia places her in the minority. As shown in the most recent 2023 report from Stanford's Institute for Human-Centered AI, the past decade has witnessed the unrelenting trend of more and more new AI PhDs in North America leaving academia and entering jobs in industry upon graduation, rising steadily from approximately forty percent in 2011 to nearly sixty-five percent in 2021 (Lynch). This phenomenon is intertwined with the trend toward larger AI models. Due to the inequitable access to computational resources, researchers are increasingly drawn to large tech companies whose for-profit natures necessitate a research agenda motivated by drastically different incentives compared to purely academic research. Granted, it might be too late to completely reverse this shift. However, it is not too late to demand more transparency from large tech companies and to push for policies mandating public disclosure of carbon emissions from their AI models. According to a paper from 2022, coauthored by Strubell, Luccioni, and other researchers, some AI and machine learning "[c]onferences such as NeurIPS and NAACL have recently added emissions reporting as an optional part of the submission process," but "both carbon estimation and reporting in ML publications and technical reports remain a relatively rare phenomenon" (Dodge et al. 1886). The practice is even rarer in industry, where the stakes are higher with larger models. Greater transparency is essential to sparking dialogue, which is the first step to action.

To connect discourse to action, there must be greater efforts to educate computer scientists about the role of AI in accelerating climate change. As it turns out, there is a course offered here at Columbia called "Machine Learning and Climate," taught by Professor Alp Kucukelbir. Based on the syllabus, of the twelve weeks of the course, eleven are devoted to learning how AI can be used to tackle climate challenges, from tracking worldwide power-plant emissions to modeling stratospheric aerosol injection. However, the question of AI's own carbon footprint is unaddressed until the last week of the course (Syllabus). While the course still serves as a small step towards raising awareness about the issue, it is only a graduate-level elective. The vast majority of CS students pass through the curriculum without ever touching on the environmental costs of AI ("CS@CU Undergraduate Programs"). This issue must be integrated into the core curriculum for all students majoring in Computer Science and related fields.

Education goes hand-in-hand with greater transparency in laying the basis for change. As students become future researchers, they must understand how the AI technologies they build impact the environment. As demonstrated by Strubell and Luccioni, the computational toll of large AI models can be drastically decreased through small steps in the engineering process. The problem is that most researchers are still completely unaware of this issue.

In recent weeks, there have been growing calls for a six-month moratorium on research into AI models more powerful than GPT-4. So far, the petition, "Pause Giant AI Experiments: An Open Letter," spearheaded by Max Tegmark, an AI researcher at MIT, has been signed by over 30,000 researchers, industry leaders, and policymakers. Among its supporters are Elon Musk, Steve Wozniak, Andrew Yang, Yuval Harari, and Columbia CS Professor Daniel Bauer, who teaches a popular course on Natural Language Processing—which lays the groundwork for the technology behind GPT-4. The letter cites concerns about unknown "risks to society and humanity" as justification for the moratorium but fails to explicitly mention the immense environmental toll of AI models ("Pause"). On a recent episode of fellow MIT AI research scientist Lex Fridman's podcast, Tegmark discusses fears of GPT-4 replacing human jobs, daydreams about using AI to search for extraterrestrial life—and even engages with catastrophic predictions of how AI might end humankind. But, during the entire three-hour-long interview, the issue of the immense impact of AI on global climate never directly surfaces. It simply feels like the elephant in the room. Clearly, despite the handful of voices sounding the sirens about the alarming carbon footprint of AI models, the issue remains relatively unacknowledged within the wider CS research community. And, even if it is known, there is a lack of incentive to address the problem, as the inconvenient truth is shunned in the frenzied race towards bigger, more accurate models.

Taking a pause on the development of larger AI models presents the perfect opportunity for researchers and policymakers to face the uncomfortable reality. It provides precisely the impetus needed to dissect the true carbon footprint of large AI models and to reevaluate the feasibility of continued research in this direction. This issue must be addressed before we continue spiraling down the dangerous path of chasing ever larger models in the search for greater artificial intelligence. It is, after all, *artificial* intelligence. We cannot get so caught up in its hype that we forget what it means to be human—especially what it means to be humans residing on planet Earth. It does not matter how much AI advances if we recklessly ignore the terrifying reality of climate change. We must prioritize protecting our only home. This is humanity's only path forward.

WORKS CITED

An, Jiafu, et al. "ChatGPT: Tackle the Growing Carbon Footprint of Generative AI." *Nature*, vol. 615, 23 Mar. 2023, pp. 586, https://www.nature.com/articles/d41586-023-00843-2.

Bass, Dina. "OpenAI Needs Billions to Keep ChatGPT Running. Enter Microsoft." *Bloomberg*, 26 Jan. 2023, https://www.bloomberg.com/news/articles/2023-01-26/microsoft-openai-investment-will-help-keep-chatgpt-online#xj4y7vzkg.

"CS@CU Undergraduate Programs." *Department of Computer Science*. Columbia University, 21 July 2020, https://www.cs.columbia.edu/education/undergraduate/.

Dodge, Jesse, et al. "Measuring the Carbon Intensity of AI in Cloud Instances." *FAccT '22: Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, 20 June 2022, pp. 1877-1894, https://doi.org/10.1145/3531146.3533234.

Fridman, Lex, host. "#371 – Max Tegmark: The Case for Halting AI Development." *Lex Fridman Podcast*, season 1, episode 371, 13 Apr. 2023, https://lexfridman.com/max-tegmark-3/.

Hao, Karen. "The Messy, Secretive Reality behind OpenAI's Bid to Save the World." *MIT Technology Review*, 17 Feb. 2020, https://www.technologyreview.com/2020/02/17/844721/ai-openai-moonshot-elon-musk-sam-altman-greg-brockman-messy-secretive-reality/.

———. "Training a Single AI Model Can Emit as Much Carbon as Five Cars in Their Lifetimes." *MIT Technology Review*, 6 June 2019, https://www.technologyreview.com/2019/06/06/239031/training-a-single-ai-model-can-emit-as-much-carbon-as-five-cars-in-their-lifetimes/.

———. "What Is ChatGPT? What to Know About the AI Chatbot." *The Wall Street Journal*, 14 Apr. 2023, https://www.wsj.com/articles/chatgpt-ai-chatbot-app-explained-11675865177.

Heaven, Will Douglas. "GPT-4 Is Bigger and Better than ChatGPT—but OpenAI Won't Say Why." *MIT Technology Review*, 14 Mar. 2023, https://www.technologyreview.com/2023/03/14/1069823/gpt-4-is-bigger-and-better-chatgpt-openai/.

Heikkilä, Melissa. "We're Getting a Better Idea of AI's True Carbon Footprint." *MIT Technology Review*, 14 Nov. 2022, https://www.technologyreview.com/2022/11/14/1063192/were-getting-a-better-idea-of-ais-true-carbon-footprint/.

"How Artificial Intelligence is Helping Tackle Environmental Challenges." *UN Environment Programme*, 7 Nov. 2022, https://www.unep.org/news-and-stories/story/how-artificial-intelligence-helping-tackle-environmental-challenges.

Luccioni, Alexandra Sasha, et al. "Estimating the Carbon Footprint of BLOOM, A 176B Parameter Language Model." *arXiv*, 3 Nov. 2022, https://doi.org/10.48550/arXiv.2211.02001.

Lynch, Shana. "2023 State of AI in 14 Charts." *Stanford University Human-Centered Artificial Intelligence*, 3 Apr. 2023, https://hai.stanford.edu/news/2023-state-ai-14-charts.

Milmo, Dan. "ChatGPT Reaches 100 Million Users Two Months after Launch." *The Guardian*, 2 Feb. 2023, https://www.theguardian.com/technology/2023/feb/02/chatgpt-100-million-users-open-ai-fastest-growing-app.

"OpenAI Charter." *OpenAI*, https://openai.com/charter.

"Pause Giant AI Experiments: An Open Letter." *Future of Life Institute*, 22 Mar. 2023, https://futureoflife.org/open-letter/pause-giant-ai-experiments/.

Saul, Josh, and Dina Bass. "Artificial Intelligence Is Booming—So Is Its Carbon Footprint." *Bloomberg*, 9 Mar. 2023, https://www.bloomberg.com/news/articles/2023-03-09/how-much-energy-do-ai-and-chatgpt-use-no-one-knows-for-sure.

Strubell, Emma, et al. "Energy and Policy Considerations for Deep Learning in NLP." *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, edited by Anna Korhonen et al., July 2019, pp. 3645-3650, https://aclanthology.org/P19-1355/.

Syllabus for Machine Learning and Climate. Taught by Alp Kucukelbir, fall 2022, Columbia University, New York City, https://www.proditus.com/teaching/mlclimate2022/.

"What is your carbon footprint?" prompt. *ChatGPT*, 13 Feb. version, OpenAI, 20 Apr. 2023, chat.openai.com/chat.

Zaveria. "OpenAI Released Its New Technology GPT-4 and It Is Claimed That It Is 100x More Powerful than GPT-3." *Analytics Insight*, 19 Mar. 2023, https://www.analyticsinsight.net/gpt-4/gpt-4-is-it-100x-more-powerful-than-gpt-3.

CHRISTINE LI '26SEAS studies Computer Science and Linguistics. She grew up between Shanghai, Beijing and the San Francisco Bay Area. She is passionate about the intersection of AI and the humanities, as well as quantum computing. Having danced pre-professionally for most of her life, Christine has continued dancing at Columbia and currently serves as Executive Director of Columbia Ballet Collaborative.