Extractive AI

Mél Hogan & Théo Lepage-Richer

Although artificial intelligence (AI) models vary greatly in their concepts and applications, they all intensify extractivism and their environmental toll wherever harnessed or deployed. To be clear, AI's extractivism goes beyond data mining the internet and plundering the work of artists and writers: deep ties to mining and fossil fuel industries and to the capitalist system enable its extraction. Underlying these extractive extremes are long-term–and longtermist–speculations about human evolution towards transhumanism by means of science and technology. These speculations are largely enabled by expansive data storage and AI (the cloud) which itself requires massive amounts of energy and water to operate.

Headlines like "ChatGPT needs to 'drink' a water bottle's worth of fresh water for every 20 to 50 questions you ask, researchers say" (Business Insider, April 14, 2023) and "Warning AI industry could use as much energy as the Netherlands" (BBC, Oct 14, 2023) are examples of how AI, despite its close association with abstract machine learning techniques like reinforcement learning and neural networks, is ultimately a resource-heavy technology reliant on large-scale computing infrastructures. These headlines illustrate growing media coverage about the heavy energy and water demands made by artificial intelligence (AI), especially by newly popular generative models like ChatGPT. Stories on the matter have documented how the AI industry is greedily guzzling energy and resources, while also relaying the promise made by industry leaders that AI is, or will be, worth its environmental cost.

In recent public statements, tech CEOs have described generative AI (GenAI) as an ever-evolving technology geared towards re-engineering the economy for environmental and social good, with near-future applications including providing everyone with affordable, personalized healthcare and education and resolving complex issues ranging from cybersecurity to discrimination. As AI comes under scrutiny, it also becomes the object of greater concerted efforts by industry leaders to maintain the hype and momentum around that technology. However, despite these varying forms and applications, extractivism is still central to contemporary AI systems.

Training and using AI systems require large amounts of data and computing power, whose infrastructural requirements exceed the capacity of most individual companies. While the latter used to operate their own (called "enterprise") data centers to meet their computational needs, they now increasingly rely on external providers to access the remarkable processing power needed for AI applications. This service is offered thanks to hyperscale data centers, which are massive, specialized buildings leased by cloud computing companies and large digital firms like Amazon, Microsoft, and Google. These infrastructures are large and centralized, which allows major technology providers to meet the data processing, storage, and networking needs of their clients at a reduced price—albeit at a greater energy cost. These infrastructures currently use somewhere between 1-4% of global energy supplies, with cryptocurrencies and AI applications driving that number up year after year.

This ongoing trend signals a consolidation of critical computing infrastructures under a handful of both tech and investment giants who can afford

the high fixed costs of retrofitting, building, and maintaining hyperscale facilities. This centralization raises urgent questions about potential market monopolies, the control of, and access to, critical infrastructures, as well as the supply chain dependencies underpinning AI, which only become more urgent with time given the steady increase of global data and processing needs. Accounting for the environmental and social 'externalities' of these infrastructures is especially important as more and more companies outsource some of their core operations to the so-called 'cloud' in the name of convenience. But far from being a simple matter of efficiency, this shift from enterprise data centers towards hyperscale facilities increases the lobbying power and influence of the entities who own and operate these infrastructures. They now exercise enough pressure on government bodies and power companies to revive defunct coal plants and drive renewed investments in much debated power sources like nuclear fusion. This means that so-called natural resources and energy become increasingly managed, owned, and controlled by these expansive cloud companies.

Engineers and quantitative researchers such as Emma Strubell, Sasha Luccioni, Walid Saad, and Aimme van Wynsberghe have proposed several solutions to reduce the energy consumption associated with both training and employing AI systems. These include requiring companies to report on the carbon emissions of their AI models; use smaller datasets when possible; outsource their computational needs to data centres closer to where they are based; and train their models at times of day when energy demands are lower. By providing actionable solutions to reduce AI's energy needs, this type of optimization-based thinking promotes a vision of what has been called "frugal AI," an ideal "that promises the use of less data and less compute power while guaranteeing robustness within the intended field of use for a given "AI model." What these measures fail to consider are the wider social, political, and ecological stakes embedded in AI companies' ideologies. In other words, AI's extractivism is not an accounting problem. Focusing on AI's energy efficiency or emissions overlooks the many more subtle and pernicious ways AI disrupts the environments in which it operates–from where it is deployed to where its resources are extracted. This oversight can be partly attributed to the difficulty in quantifying other forms of environmental disruption, as well as their distant impact on communities and locations not typically considered in corporate and political decision-making circles.

AI's water needs is one of such examples, despite having received a certain level of journalistic attention. Journalists have reported how training GPT-3, for example, required roughly 700,000 liters of clean freshwater to prevent the servers that were used from overheating, but very few have noted how greater that number would have been had GPT-3 been trained in warmer environments. Scaling AI and providing global access to its applications put remarkable pressure on the limited resources of several water-stressed regions, draining lakes and rivers while accelerating the displacement of vulnerable populations. The implications of AI's 'thirst' are far-reaching and include exacerbating droughts and desertification, disrupting ecosystems and fisheries, triggering conflict, and amplifying water inequalities by diverting water supplies towards technology hubs. As climate change intensifies water scarcity worldwide, unrestrained AI expansion threatens to siphon public waters towards private tech profits and further disrupt the environments and communities most affected by climate change.

Water, however, is only one of the many resources needed to build, operate, and maintain the infrastructures required to train and use AI systems. Minerals such as neodymium, cerium, and praseodymium, cobalt, nickel, and lithium possess material properties that are essential for meeting AI's computational needs, but their

extraction comes at a steep environmental—and human—price. Such minerals are often mined in conflict zones with lax labour regulations, using methods that ravage landscapes, contaminate groundwater, and destroy natural habitats. And once these minerals are outside the ground, their footprint gets hardly better. The plants in which these minerals are processed emit important amounts of carbon and toxins, which are ignored by most carbon emission reporting requirements. The same applies to silicon, which is one of the most energy-greedy industrial materials. Together, these minerals and materials highlight how AI as an industry rests upon highly destructive mining practices, whose effects span continents while extending beyond the reach of most regulatory frameworks.

What these more pernicious– albeit no less detrimental–environmental effects illustrate is that no amount of optimizing, measuring, or reporting can curtail the deeper ideological assumptions that underpin global investments in AI. It's important to note that our anti-extractive concerns around AI reflect those of various climate justice movements. What makes curtailing so-called advancements in AI so urgent is that AI is largely the marketing tactic of the billionaire and trillionaire class, which is infiltrating almost all systems and institutions. Once infiltrated, these institutions will become tethered to the technology and its ideological underpinnings – and at a steep price.

This is why the current emphasis on making AI more 'sustainable' misses this technology's more fundamental problem, namely, that it functions as a necrotic future- and agenda-setting technology, which justifies the mobilization of large amounts of resources for the economic gain of a few actors. AI's infrastructures, operations, and history are inexorably linked with specific ideological, financial, and political investments that have splintered AI into several instantiations – stock options, real estate deals, market monopolies, etc. – that are almost entirely in the control of the billionaire class. Through AI and other similar technologies, the billionaire class claims for itself the status of custodian and curator of key natural resources, which they frame as having no value and meaning when left in their original environment and pre-extracted state.

This is why, as Murray Bookchin argues, there is an urgency to address the relationship between environmental and social problems, especially when technology is presented as both the culprit for, and the solution to, these same problems. As he writes, "[u]nless we realize that the present market society, structured around the brutally competitive imperative of "grow or die," is a thoroughly impersonal, self-operating mechanism, we will falsely tend to blame technology as such [...] for environmental problems." Environmental issues stem fundamentally from social dynamics, wherein inequalities and market-driven policies shape a perspective where everything – including workers and resources – is seen primarily in terms of the surplus value that can be derived from them. In this context, technology serves primarily as another catalyst for channeling greater numbers of people and resources into an economic system it perpetuates, even though it didn't originate those principles. In his book Ways of Being, James Bridle adapts this argument to AI specifically by writing that, "[d]riven by the logic of contemporary capitalism and the energy requirements of computation itself, the deepest need of an AI in the present era is the fuel for its own expansion" (7). The environmental impact of AI can be significantly reduced not simply by optimizing existing processes or making incremental improvements, but by tackling the underlying social and economic factors that enabled the widespread mobilization of resources and talent for this technology in the first instance.

This is an important distinction given that AI in its current form is unthinkable outside of the larger industrial backdrop that supports it. AI first

emerged in the 1950s as a symbolically-inclined, heuristics-based academic discipline in response to the shortcomings of cybernetics; evolved in the 1970s into a field dedicated to knowledge representation with the rise of expert systems; and finally transformed in the 1990s into a highly operational approach to the extraction of patterns from large amounts of data. Given the sizable computational needs and data requirements of that latter iteration, AI's third form took shape at the intersection of both the academy and industry from the get-go, finding its validation in its capacity not so much to reproduce certain forms of human reasoning as to improve operationally relevant tasks and industrial processes. As Meredith Whittaker argues, AI's current meaning is indistinguishable from that technology's capacity to generate profits and concentrate wealth, at the expense of any meaningful relationship to human intelligence or any other form of cognition. With investors and tech CEOs holding the purse strings of the large amounts of data and computational power needed for AI, it is up to them to shape the development of that technology by granting access to the startups and labs aligned with the future they envision while withholding these key resources from those that are not.

Unparalleled investments in creating new learning algorithms, building AI-first hardware, and developing new applications for neural networks and large languages models have seen the field transform from one spearheaded by research institutes and universities—albeit private, well-funded ones—to one led by corporations with the necessary funds to build, maintain, and operate the necessary infrastructures to train AI applications. As summed up by Whittaker, current "advances" in AI are "not due to fundamental scientific breakthroughs in AI techniques" but are instead "the product of significantly concentrated data and compute resources that reside in the hands of a few large tech corporations." This is why AI's current meaning is primarily defined by actors such as Microsoft (OpenAI), Meta, Google, Amazon (Anthropic), Palantir, IBM, and Tesla, which collect, broker, and store the necessary data to train GenAI systems and have the power to establish what applications these systems should be used for.

This is partly why both policymakers and the public at large should be skeptical of tech CEOs' appeals to dubious ideas like "artificial general intelligence" (AGI) and AI's existential threat. Such ideas, which are infused with a peculiar combination of quasi-religious, eschatological aspirations and racially loaded fantasies about transhumanism, do little but distract governments, news outlets, and the general public from AI's current environmental impact. Focusing on AI's long-term, and highly hypothetical, risks is an effective way to channel policymaker's attention away from the other industries (fossil fuels, mining, real estate, energy, military, etc.) AI thus avoids policies and regulations that might affect the (AI) industry's bottom line. In that context, it is thus especially important that scholars and policymakers attend to the many ways AI exacerbates issues–ranging from human rights violations to climate change–instead of falling for the type of long-term speculations favoured by tech CEOs. Any potential socioeconomic gain GenAI might offer should be considered in relation to its deep entanglement with resource extraction and non-renewable energy production. If AI can be said to be automating anything at this point, it might very well be the climate crisis itself, as Brian Merchant argues, or a certain form of coded gaze, to quote Joy Buolamwini, which gives primacy to a settler-colonial, capitalist view of the future over all else.

Conceiving of AI as an agenda-setting industry as opposed to a discrete technology is essential to develop policy frameworks that are truly adapted to AI's social, environmental, and economic impact. Corporate appeals to making AI more environmentally friendly and mitigating its impact on seafloors, lands, water, soils, and the environment at large overlooks how investments

in that industry systematically translates into investments into extractive industries like mining, oil, gas, and even the military.

Companies known for specific applications like online search (Google), e-commerce (Amazon), social media (Meta), and so on are in fact direct partners of, if not investors in, fossil fuel, mining, and even defense companies, which are essential to power AI's essential infrastructures. While absent from tech companies' public-facing communications, this proximity is in fact one of the stated incentives of power companies' renewed investments in oil and gas extraction. The arrival of Google Cloud and other cloud services like those offered by Microsoft, Amazon, Oracle in Saudi Arabia, for instance, is the product of major partnerships with oil companies like Aramco, which provide the necessary energy to power cloud computing infrastructures in some of the earth's hottest and more water-deprive regions. Or, as a similar case in Kazakhstan shows, documented as a first-hand account in the Logic Magazine essay "Oil is the New Data," merging Big Oil and Big Cloud imposes a panopticon onto the workers and field. The author proposes that workers rights is the main site of friction and resistance: "**Resource extraction is an ancient imperial practice. As tech workers join the movement for climate justice, we must also find ways to undo the legacies of Big Oil's imperialism, and bring countries like Kazakhstan fairly and safely into a carbon-free future.**"

Such partnerships between cloud and oil companies are only one of the many conditions that allow AI to capture environments everywhere as privileged sites for both data and resource extraction. Despite the great conceptual distinctions that differentiate AI's various models and the many applications they power, there remains a common denominator in AI's many forms—namely, that they all intensity extractivism wherever they are deployed and amplify the environmental consequences of that economic system. Far from metaphorical, AI's extractive qualities point to more than the pilfering of user data and copyrighted works online—which is without a doubt yet another pressing issue in its own rights—and in fact extend to that technology's deep entanglement with the extractive industries and the broader economic system that support and validate them.