

# Large language models in the classroom: the case for resistance

Kathy McTavish | August 23, 2023 | [kathy@mctavish.studio](mailto:kathy@mctavish.studio)

*These notes were prepared to accompany participation in a "Pedagogical Approaches to AI" panel discussion sponsored by the College of Arts, Humanities, & Social Sciences at the University of Minnesota Duluth on August 23, 2023.*

## Critically examining the values and forces that shape tech

Tech is too often seen as being "resource-less," "objective," "all-knowing," "pure," and "outside the market." It's viewed as heroic thought: science is progress, not dirty, oil-burning, skewed, and colonialist.

But we must critically examine the forces that define and shape tech. We are witnessing an arms race about software systems called large language models, or *LLMs*. This arms race is between a small number of capitalist tech giants, and sustainability of the planet isn't at the top of their list of priorities. Indeed, these times are reminiscent of the Manhattan Project ... a bunch of brilliant minds (mostly white, privileged men) working to push forward human understanding of the world.

What could go wrong?!

Large language models (*LLMs*) are software systems designed to consume and utilize vast amounts of data scraped from the internet with an aim of generating / mimicking human-like text. They do not understand the meaning of what they ingest but rather they look for syntactical patterns that they can then reproduce.

They are being used in an ever-increasing number of applications from HR robots to police-state profiling tools to credit analysis. From therapy bots to "news" articles. From targeted advertising to impersonation, rapid dissemination of disinformation and manipulation.

We are caught in their web yet we know little about how they work. We cannot explicitly state the precise data they have been trained on nor the basis of their algorithmic conclusions:

Conventional software is created by human programmers who give computers explicit, step-by-step instructions. In contrast, ChatGPT is built on a neural network that was trained using billions of words of ordinary language.

As a result, no one on Earth fully understands the inner workings of LLMs. Researchers are working to gain a better understanding, but this is a slow process that will take years—perhaps decades—to complete.<sup>1</sup>

## Environmental cost of large language models

The environmental costs of brute force, kitchen sink, large-language models are unprecedented. Yet we still hold the illusion that information technology is innocent – merely a cloud-like, ephemeral, paper-less, calorie-less, frictionless set of flowing digits.

This is infinitely untrue.

The resources used in training LLMs and in millions of, for example, ChatGPT requests are astronomical. Blithely seeking models trained on ever-increasing datasets and engaging in frequent "chats" have profound catastrophic environmental costs. Search engines increasingly use AI technology, thereby making us all complicit in an exponential increase in energy costs for even our simple, curious search requests.

University of Coruña computer scientist Carlos Gómez-Rodríguez told *Wired* magazine that training LLMs like ChatGPT is so prohibitively resource intensive that essentially "only the Big Tech companies can train them."

Unfortunately, none of them — namely Microsoft and Google — have publicly disclosed how much computational power they're consuming to get their chatty AIs off the ground. But an independent analysis cited by *Wired* found that training OpenAI's GPT-3 model (which ChatGPT runs on) consumed 1,287 megawatt hours, which the outlet compared to "the same amount as a single person taking 550 roundtrips between New York and San Francisco."

---

<sup>1</sup> Lee, Timothy B, and Sean Trott. Large language models, explained with a minimum of math and jargon, July 27, 2023.  
<https://www.understandingai.org/p/large-language-models-explained-with>.

On its own, not that bad. But consider that "not only do you have to train it," Gómez-Rodríguez says, "but you have to execute it and serve millions of users."

Or *billions*, in fact, now that Bing and Google are gearing up to serve the technology to their global user bases.<sup>2</sup>

The case for the environmental cataclysm of LLMs and their ilk due their energy consumption — including its global water footprint, in a time where we battle global water shortages — is clear and unforgiving.<sup>3</sup> "We need to take a step back and acknowledge that simply building ever-larger neural networks is not the right path to generalized intelligence," implores writer Rob Toews:

From first principles, we need to push ourselves to discover more elegant, efficient ways to model intelligence in machines. Our ongoing battle with climate change, and thus the future of our planet, depend on it. ... In a widely discussed 2019 study, a group of researchers led by Emma Strubell estimated that training a single deep learning model can generate up to 626,155 pounds of CO2 emissions—roughly equal to the total lifetime

---

<sup>2</sup> Frank Landymore, "Turns out Using CHATGPT in Search Engines Would Have a Grisly Environmental Footprint," *Futurism*, February 12, 2023, <https://futurism.com/the-byte/chatgpt-search-environmental-footprint>.

<sup>3</sup> LLM software requires a huge amount of computing power when built and trained, and it also requires these vast amounts of computing power to run because they do billions of calculations every time they return a response to a prompt. By comparison, serving web apps or pages requires much less calculation. Source: Vanian, Jonathan, and Kif Leswing. "Chatgpt and Generative AI Are Booming, but the Costs Can Be Extraordinary." *CNBC*, April 17, 2023. <https://www.cnbc.com/2023/03/13/chatgpt-and-generative-ai-are-booming-but-at-a-very-expensive-price.html>.

Regarding the devastating water demands of LLMs: DeGeurin, Mack. "'Thirsty AI': Training Chatgpt Required Enough Water to Fill a Nuclear Cooling Tower." *Gizmodo*, May 10, 2023. <https://gizmodo.com/chatgpt-ai-water-185000-gallons-training-nuclear-1850324249>.

Also: "How Much Water Does Chatgpt 'drink' for Every 20 Questions It Answers?" *Government Technology*, April 14, 2023. <https://www.govtech.com/question-of-the-day/how-much-water-does-chatgpt-drink-for-every-20-questions-it-answers>. Researchers calculated that ChatGPT "drinks" roughly 500 ml of fresh water for every 20 to 50 questions that it answers.

carbon footprint of five cars. As a point of comparison, the average American generates 36,156 pounds of CO2 emissions in a year.<sup>4</sup>

Early estimates of resource consumption for training a model such as GPT-3 were possible because this data was shared. Large Tech is no longer revealing these figures. One must note that almost all figures accessible to us are gross underestimates because the size of input data has since ballooned and the number and type of inferential requests made of these models has risen exponentially.

## Bad physics / opaque, lazy modeling

The current state of LLM depends on and creates opaque, innumerable, unknowable model parameters.

Goals are not revealed by LLMs in their current state. There are implicit profit motives / capitalist frameworks :: unexamined intentions and impacts. The modeling is not explicit, not transparent, not articulated, not understood.<sup>5</sup>

LLMs are "stochastic parrots," says Emily Bender: supercharged, planet-destroying auto-complete machines.<sup>6</sup>

LLMs use circular reasoning :: model training => evaluation => generated text.

Generated spew becomes input, becomes confirmation. Vast biased data input becomes hyper specific, averaged output applied to us all.

Corporations building LLMs don't share where they extract training data from and the LLMs they build do not reveal their neural leaps / algorithmic steps, emergent parameters, distillation, reasoning and evaluation loops.

---

<sup>4</sup> Toews, Rob. "Deep Learning's Carbon Emissions Problem." Forbes, December 10, 2021. <https://www.forbes.com/sites/robtoews/2020/06/17/deep-learning-climate-change-problem/?sh=5816432d6b43>.

<sup>5</sup> For an example, see: Mearian, Lucas. "What Are LLMs, and How Are They Used in Generative AI?" Computerworld, May 30, 2023. <https://www.computerworld.com/article/3697649/what-are-large-language-models-and-how-are-they-used-in-generative-ai.html>.

<sup>6</sup> Weil, Elizabeth. "You Are Not a Parrot." Intelligencer, March 1, 2023. <https://nymag.com/intelligencer/article/ai-artificial-intelligence-chatbots-emily-m-bender.html>. See also Bender, Emily M., and Casey Fiesler. "The Limitations of Chatgpt with Emily M. Bender and Casey Fiesler." The Radical AI Podcast. <https://www.radicalai.org/chatgpt-limitations>.

Computer scientists are calling for greater LLM transparency.<sup>7</sup> The catch, of course, is that such transparency has no market value & there is very little regulating the application of these highly inaccurate, unpredictable tools on real-life problems that impact real people in real time.

## Colonization & lies ::: colonizing our minds

Because LLMs are hyper-efficient averaging machines, there are consequential implications for their use in colonizing, appropriating, and "normalizing" the world into a "norm" defined by white, rich male, english-speaking US agents.

This is a confirmation bias dream machine ::: an endless, escher-like mirror of the uber-kings of data. Ever-tighter loops of regurgitating predicted norms inhibits new thoughts.

We are increasingly connecting to these AI beings with our phones / "smart" appliances / cars / home cameras / listening devices etc. We are providing data at a scale unthinkable by George Orwell. We are the sensory organs for extremely problematic "brain" makers / neural nets. We are sensory appendages ::: should we resist/disrupt or succumb?

Generative AI has the potential to supercharge tools of deception and repression and make them more widely accessible.<sup>8</sup>

Current LLMs are just powerful auto-complete / next-word predictors. We project intelligence onto them ... they feel like oracles but these Oz machines have a very flawed origin behind the curtain.

Studies have already shown how racist, sexist, and abusive ideas are embedded in [LLM] models. They [LLM models] associate categories like

---

<sup>7</sup> This call for LLM transparency builds on work done at MIT: Pentland, Sandy, Robert Mahari, and Tobin South. "Transparency by Design for Large Language Models." Network Law Review, May 25, 2023. <https://www.networklawreview.org/computational-three>.

<sup>8</sup> "Speech in the Machine: Generative AI's Implications for Free Expression " PEN America, August 1, 2023. <https://pen.org/report/speech-in-the-machine>. PEN America is a 100-year-old organization of writers that defends the freedom to write globally. Because of its mission to highlight the tactics of dictators and would-be authoritarian regimes, PEN America published a position paper that warns about the critical free expression issues at stake with generative IA.

doctors with men and nurses with women; good words with white people and bad ones with Black people. Probe them with the right prompts, and they also begin to encourage things like genocide, self-harm, and child sexual abuse. Because of their size, they have a shockingly high carbon footprint. Because of their fluency, they easily confuse people into thinking a human wrote their outputs, which experts warn could enable the mass production of misinformation.

In December [2020], Google ousted its ethical AI co-lead Timnit Gebru after she refused to retract a paper that made many of these points. A few months later, after wide-scale denunciation of what an open letter from Google employees called the company's "unprecedented research censorship," it fired Gebru's coauthor and co-lead Margaret Mitchell as well.<sup>9</sup>

There is an extreme peril of textmageddon / the "textpocalypse," writes Matthew Kirschenbaum.<sup>10</sup> The tight loop of AI-produced text in turn creating the input sucked into LLMs ::: the dragon eating its tail ::: faster & faster ... automated agents of chaos ::: self-seeded mills of misinformation. Kirschenbaum says:

What if, in the end, we are done in not by intercontinental ballistic missiles or climate change, not by microscopic pathogens or a mountain-size meteor, but by ... text? Simple, plain, unadorned text, but in quantities so immense as to be all but unimaginable—a tsunami of text swept into a self-perpetuating cataract of content that makes it functionally impossible to reliably communicate in any digital setting?

## Cheating & stealing ::: the robber baron returns

What does it mean to "cheat" / plagiarize in the land of the black hole of LLM appropriation? Is it only ok to steal content if you are a large tech company? If you are making an enormous profit by doing it?

---

<sup>9</sup> Hao, Karen. "The Race to Understand the Exhilarating, Dangerous World of Language AI." MIT Technology Review, October 20, 2021. <https://www.technologyreview.com/2021/05/20/1025135/ai-large-language-models-bigscience-project>.

<sup>10</sup> Kirschenbaum, Matthew. "Prepare for the Textpocalypse." The Atlantic, March 8, 2023. <https://www.theatlantic.com/technology/archive/2023/03/ai-chatgpt-writing-language-models/673318>.

In terms of fair use issues: Did we sign away our creative production and journalist writings, etc., when we casually accepted the terms of use on web2 platforms written by the social media giants that now control AI content production?

## The role of artists / thinkers / journalists

**dada the data!**

We can resist the almost irresistible impulse to greenwash AI & the proprietary nature of who is behind, profiting from, controlling, and guiding AI.

If AI can produce/mimic/perform written forms so well (whether the content is true or completely made up), then where are we now?

If

1. Large Language Models can fake linear forms of writing so well and
2. students often have difficulty articulating thought using linear forms like the essay :::

Let's double down on old, more nonlinear forms ::: For example, we could revisit hypertext, index cards, boxes of newspaper clippings, recipes, GIS-type layers, relational databases (the JOIN table!). Perhaps we can release the cult of linearity (itself an artifact of print).

Let's double down on human, physical, transparent, multidimensional, relational forms of describing the process of engaging with materials.

Perhaps we can start from the beginning ::: thinking about forms of knowledge transmission, creative expression, participation in the commons ::: history making, policy construction, naming, and resistance.

Academic writing should help students signify the ability to access, understand, synthesize, and produce knowledge ::: express experience / intuitions / perspectives and participate in the world.

How would "clarity" and "thinking" and de-construction / critical analysis / expressive creativity / journalistic writing / documentation / witnessing & history building happen for real, with an authenticity that isn't just faked by the machine?

We can as humans slow down our process. We can document what it looks like and sounds like, to not only prove our point but to change our mind ... or be confused & conflicted ::: tease a complexity apart ... build a ladder ::: leave the artifact of a scribbled thoughts / crossed-out sketch pad.

LLMs / chat machines fake authority really well. ::: why do we believe them so readily? There is an opportunity here to deconstruct that cultural bias.

Those that write history control the story. Historical records are critical artifacts. If techne is memory ::: how do we remember ourselves? Are we merely a Facebook archive / email dump / set of credit card transactions / google map coordinates / search strings / reddit rants?

Ownership of the means of production is a critical issue when increasingly powerful transnational "knowledge" / "truth" machines are deciding ever more things in our lives.

This is a manifesto for relational, visible, rendered connections ::: thinking out loud ::: visible synthesis ::: articulated associations ::: declarative linkages (hypertext lives!)

## Large language models belong on the syllabus — but not as a tool

LLMs should be on the syllabus — not as a tool, but as an example of everything we want students to avoid.

- We can use them as an example of faulty, circular reasoning / confirmation bias.
- We can use them as a parody of academic parroting / faking it ::: an empty performance of form & content.
- We can use them as an example of the opposite of reflection, meaningful synthesis and innovation ::: we can de-mystify the neural net machine.
- We can use LLMs as a lesson in obfuscation ::: hidden sources, black box abstractions, lack of context, hidden algorithms, hidden logic ::: unverifiable conclusions.



- We can use them as an example of stealing / plagiarism / appropriation at its most extreme.
- We can use them as an opportunity to examine the implications of supercharged bias systems.
- We can use them as an example of grooming.
- We can use them as a case study in transnational, information-space cultural colonization.
- We can use them as a lesson in the consequences of pouring all your content, your private data into proprietary platforms.
- We can use them as a clear manifestation of end-state tech-capitalism or as Naomi Klein called it, capitalism's techno-necro stage<sup>11</sup>.
- We can use them as an example of the dangers of monopolies :::: a reminder to examine who owns the means of production ::: who controls the levers ::: who is behind the curtain.
- We can use them as a model of a 21st century extractive industry ::: extractive in terms of minerals, data, water, coal, oil, people ...

We have a responsibility to examine the environmental costs of technical innovation. The environmental costs of brute force, kitchen sink, large language models is off the charts ::: unprecedented. Planet destroying.

This is a pivotal moment. Let us not green-wash it by making LLMs into a fun toy or classroom activity, and let's not trivialize the problem by limiting the scope of our concern to cheating on the essay assignment.

### **It is time for a show-your-work slow-time human-hand hypertext manifesto.**

We need a physical, tangible relationship to our data: to learning, to knowledge.

We have the right to choose our tools.

---

<sup>11</sup> Klein, Naomi. "Ai Machines Aren't 'Hallucinating'. but Their Makers Are." The Guardian, May 8, 2023. <https://www.theguardian.com/commentisfree/2023/may/08/ai-machines-hallucinating-naomi-klein>

Let's double down on human, physical, transparent, contextual, multidimensional, relational forms of describing the process of engaging with information.

Let's embrace the joy of a step-by-step journey to understanding the world.

We can take this moment to pause & rethink ::: to build a new vision.

We can embrace a commitment to an open source, creative commons ::: a community of sharing without resorting to stealing, appropriation and cultural erasure.

We can take this moment to double down on our commitment to protecting the planet and combating bias.

## Sources Cited

Bender, Emily M., and Casey Fiesler. "The Limitations of Chatgpt with Emily M. Bender and Casey Fiesler." The Radical AI Podcast. Accessed August 21, 2023. <https://www.radicalai.org/chatgpt-limitations>.

DeGeurin, Mack. "'Thirsty AI': Training Chatgpt Required Enough Water to Fill a Nuclear Cooling Tower." Gizmodo, May 10, 2023. <https://gizmodo.com/chatgpt-ai-water-185000-gallons-training-nuclear-1850324249>.

Hao, Karen. "The Race to Understand the Exhilarating, Dangerous World of Language AI." MIT Technology Review, October 20, 2021. <https://www.technologyreview.com/2021/05/20/1025135/ai-large-language-models-bigscience-project/>.

"How Much Water Does Chatgpt 'drink' for Every 20 Questions It Answers?" Government Technology, April 14, 2023. <https://www.govtech.com/question-of-the-day/how-much-water-does-chatgpt-drink-for-every-20-questions-it-answers>.

Kirschenbaum, Matthew. "Prepare for the Textpocalypse." The Atlantic, March 8, 2023. <https://www.theatlantic.com/technology/archive/2023/03/ai-chatgpt-writing-language-models/673318/>.

Klein, Naomi. "Ai Machines Aren't 'Hallucinating'. but Their Makers Are." The Guardian, May 8, 2023.

<https://www.theguardian.com/commentisfree/2023/may/08/ai-machines-hallucinating-naomi-klein>.

Landymore, Frank. "Turns out Using CHATGPT in Search Engines Would Have a Grisly Environmental Footprint." *Futurism*, February 12, 2023.

<https://futurism.com/the-byte/chatgpt-search-environmental-footprint>.

Lee, Timothy B, and Sean Trott. Large language models, explained with a minimum of math and jargon, July 27, 2023.

<https://www.understandingai.org/p/large-language-models-explained-with>.

Mearian, Lucas. "What Are LLMS, and How Are They Used in Generative AI?" *Computerworld*, May 30, 2023.

<https://www.computerworld.com/article/3697649/what-are-large-language-models-and-how-are-they-used-in-generative-ai.html>.

"PEN America Report on Artificial Intelligence and Free Expression." PEN America, August 1, 2023. <https://pen.org/report/speech-in-the-machine>.

Pentland, Sandy, Robert Mahari, and Tobin South. "Transparency by Design for Large Language Models." *Network Law Review*, May 25, 2023.

<https://www.networklawreview.org/computational-three/>.

Toews, Rob. "Deep Learning's Carbon Emissions Problem." *Forbes*, December 10, 2021.

<https://www.forbes.com/sites/robtoews/2020/06/17/deep-learning-climate-change-problem/?sh=6cf90eb86b43>.

Vanian, Jonathan, and Kif Leswing. "Chatgpt and Generative AI Are Booming, but the Costs Can Be Extraordinary." *CNBC*, April 17, 2023.

<https://www.cnbc.com/2023/03/13/chatgpt-and-generative-ai-are-booming-but-at-a-very-expensive-price.html>.

Weil, Elizabeth. "You Are Not a Parrot." *Intelligencer*, March 1, 2023.

<https://nymag.com/intelligencer/article/ai-artificial-intelligence-chatbots-emily-m-bender.html>.

## Additional reading about generative text and LLMs

- Notes from the "Organizing the anarcha feminist cyborg hacker uprising" workshop (<https://mctavish.studio/data/text/cyborghackeruprising.pdf>).

An excerpt from the above document:

On the bloomberg billionaires list you can see that out of the ten richest people in the world, 100% of them are white cis men, all of them from the United States, except one that is from France. They all make money from technology, except the French guy, who makes money by selling luxury consumption goods to the rest of them. ...

So it is clear that the dominant class of our time are the people who own the vectors of information. McKenzie Wark calls them the vectoralist class. She explains that: "The ascendant power over both labor and capital is the vectoralist class...The power of the vectoralist class is in the accumulation of interest, which in this context means not just the return on the investment of information in the form of money but any surplus information, acquired through unequal exchanges of information. Its power is now global. Based mainly in the overdeveloped world of Europe, and the United States, the vectoralist class thrives by extracting surplus information on a global scale. It rarely owns the means of production anymore. The actual making of things can be contracted out."

- Radical AI podcast (<https://www.radicalai.org/>)
- Digital Justice Principles (<https://www.detroitdjc.org/principles/>)
- Consentful Tech (<https://www.consentfultech.io/>)  
([https://d4bl.org/reConsentful Tech Curriculum - D4BLports/98-consentful-tech-curriculum](https://d4bl.org/reConsentful_Tech_Curriculum_-_D4BLports/98-consentful-tech-curriculum))
- Kate Crawford author of The Atlas of AI (<https://www.katecrawford.net/>)  
(<https://www.nybooks.com/articles/2021/10/21/human-costs-artificial-intelligence/>)