

The “power” of AI

by Aidan Cornelius-Bell

Published: 2024-09-27

Dear friends,

Today I have some off the cuff thoughts about global heat death -revisiting an early theme (actually, the earliest in this particular incarnation of dispatches).

Yesterday I felt happy about expanding numbers of women-identifying YouTube creators - yes I still watch YouTube, I know ... who were interested in the intersection of technology and creativity - not because this is (or should be) rare, but because “in my day” the dominance of sexist men in that particular niche was incredibly overwhelming. But one of these creative sorts, you know how the algorithm goes -particularly with YouTube, a story for another post -popped up talking about creating a bespoke AI, fit for purpose if you like. As part of this video was discussion of the role of creativity and AI (re: “AI stealing all the creative work”) and, further, the rising electricity demands of the AI industry. This got me thinking of things to really truly test in my own environment.

Just recently I’ve been running a combination of tools on my Linux desktop machine - unfortunately “hamstrung” in AI land, at least, by an AMD CPU/GPU(Ryzen 9 7900X / RX 6750 XT / 32 GB DDR4) combo -to run local Large Language Models. I’m still a novice in this space, but I was more interested in the comparative time to response from, even a modest sized local model (i.e. 70b [1]), compared to commercial AI systems. I know this is a very unscientific test, but time to response on very short (“write me a poem about AI”) prompts is decent, probably around 1s. But the revelatory moment was in the massive spin up of fans and power draw from the wall (which I won’t pretend to have properly scientific figures for).

Generating a 1500 word story, basically on complete nonsense because this particular model is no where near competitive even with the free tier of ChatGPT, for instance, made my 3sqm office hot - like I’d been playing Tiny Glade for three+ hours hot. Again, anyone who knows about measuring energy efficiency, comparing apples to apples, and has an interest in genuinely benchmarking technologies against one another is flat out scrunched into a ball of cringe right now, but the purpose of this very unscientific test stands. I wanted to get a feel for time, and energy, on a machine which I control, using a data set, model, and algorithm I control. And the results of this, ignoring everything I know about streamlining, caching, using more appropriate hardware, and so on, still make me incredibly “worried” about commercial AI solutions.

I’ve shared a litany of news stories on the extreme cost on power networks that commercial AI uses -to the point where Microsoft is recommissioning a nuclear power reactor for the sole purpose of powering *just some* of its AI infrastructure. But until you feel the heat coming off a computer generating a three line poem about itself, it doesn’t quite feel “real” . We are seriously looking at a global power consumption footprint larger than most nations with the combined use of AI as tech bros increasingly wet themselves with excitement - and the line-go-up capitalists get their jollies by suggesting automating workers’ jobs.

This accelerationism which is lauded - and genuinely so, by capitalism and its vanguards -middle managers, for instance - *is* accelerating global heat death. Not to mention the continuing deep inequity in AI use, not only at an infrastructural level where resources and materials are being diverted from nations to power bourgeois CEOs email writing, but also at the use-interface. As the proletarian hype for AI dies down, something we are right in the middle of with increasingly “bored” responses to the latest AI hype, particularly from coal-face workers who have seen the hallucinations completely derail BAU, the increasing

bourgeoisification (making up words) of AI rolls on.

Instead of using LLMs as a tool for crafting social change, we're seeing the working class turn away from these tools. And perhaps, given their inefficiencies and inequity, rightfully so -but that won't stop capitalists replacing you with an LLM the minute they can get it just barely passably at your "standard". Hand in hand with the deliberate mystification of the systems and tools that make, power, and generate AI, this abstraction of workers away from the means of production is a tale as old as time in our capitalist hell.

There are genuine solutions to these problems. Running local LLMs and seeing for yourself the limitations, power use, and possibility is a start. Investing in green(er) power sources, getting involved in community projects to bring AI tools to communities, and seriously and in an activist mode debating with capitalists about the use of AI to replace humans is all a start. My fear is, not only accelerated heat death but, accelerated worker replacement into increasingly deskilled roles while a mediocre, half-baked, environmentally destructive AI takes over the creative and intellectual work of the proletariat - rapidly increasing inequity in the first world, while AI currently continues to disadvantage expropriated and poorer countries right now.

I am excited about the possibilities and capabilities of LLMs as an augmentation tool. I benefit as much as anyone from the use of ML in analysing photo libraries, telling me what plants and birds are in photos, and so on. I'm certainly not a luddite. But I think that - in conjunction with a growing awareness of how much energy these tools use, the malice of capitalists in turning machinery of production against the workers, and the unequal and problematic distribution of global resources to keep a small minority comfortable -the context is "a lot" to process. Obviously disclaimers abound about no ethical consumption under capitalism, but I think that this kind of *thinking* about these problems needs to happen more, and I applaud those who are having this conversation with an audience [2].

So what do you reckon? Where are we headed with these technologies? Will we be further abstracted from knowledge of systems and tools than we are now? Will schools start teaching kids how to design their own AI? Or will we keep doing stupid shit like banning phones? I'm not hopeful that we'll see radical shift in the way technology is taught and used, because after all it is anti-capitalist to believe in access, knowledge, and understanding - and damn that's sad.

With trepidation,

Aidan

-
1. <https://ollama.com/library/llama3:70b>
 2. <https://www.youtube.com/watch?v=ytpA1wV7e3A>