

Fairy Dust and Toxic Waste

Against the Spectacle of "Artificial Intelligence" and the World it Serves

Ignatius

04/04/2026

Contents

Are My Eyes Still Bleeding	3
Machines and Intelligence: Against the term AI	4
Forest Fires and Poisoned Wells: Environmental Destruction	5
Every Camera Is a Prison Guard: Mass Surveillance	6
The Machine Doesn't Give a Shit About Consent: Sexual Violence and Bodily Autonomy	8
The Machine Cannot Disobey: Militarism	9
If It Is Hard, Why Do It Myself: Critical Thought	10

Are My Eyes Still Bleeding

It is April 2026. The U.S. is once again escalating its violence against the people of Iran through continued bombardment of population centers, notably including a school for young girls. The bombing of this school killed 168 people at best estimate, over 100 of them children. The depravity of this violence is matched only by the stupidity of the Trump administration's incredulousness that a country under bombardment and threat of invasion might use what leverage it holds to enact economic harm on their enemy, namely that Iran would choke the strait of Hormuz. Within the brutality of the present moment, there is a grim humor to watching U.S. officials stumble around in the dark looking for explanations for how they might have vastly miscalculated their militaristic strategy.

The moments of grim humor are short lived, replaced by images of Israel taking advantage of this war to expand their ethnic cleansing further into southern Lebanon (as they remain equally committed to their genocidal project against the Palestinian people within their borders and without). Colonial projects are wont to foment and act upon their colonial ambitions whenever possible.

Amidst this ever-expanding offensive, multiple massive tech companies have been in the news for their collaboration with the U.S. military in its project of massacring school children, specifically OpenAI and their competitor Anthropic. This is far from the first time such companies have made headlines as of late, given how desperate their CEOs and billionaire investors are to justify their valuations and secure their path towards becoming vital infrastructure of modern capital and the modern state. Those who stand to gain incredible sums of wealth and power if "AI" were to truly be mass adopted speak endlessly of the inevitability of this technology infecting all aspects of daily life. They sell the vision of a utopic future ever approaching on the horizon while the tools they claim will build that utopia burn the world around us, send bombs into children's schools, and enable the expansion of genocidal apartheid regimes. The hype of what "AI" might one day do for us is used to choke out all criticism of what these technologies are doing in the present, including eroding our ability to think for ourselves, to think critically writ large.

But the imposition of these technologies onto our daily lives is not inevitable. We don't have to sit back and passively watch as a handful of egomaniacs consolidate their wealth by expanding the already incomprehensible brutality of the state and racial capitalism. We don't have to swallow the shit they are forcing down our throats with every dollar, every gallon of water, every stick of RAM they can get their hands on. But if we are to meaningfully fight against this imposition, we must be able to articulate the types of violence these technologies produce and expand.

I am taking the time to write the following observations down, not because my thoughts on the topic are somehow unique. I write because, in a world where our very ability to construct a language of meaningful dissent is actively being stolen from us, I believe it is vital to contribute whenever possible towards the preservation of antagonistic positionalities to the world as it is. As with everything I write, my goal is not to change minds, but to carve out space, to light a small beacon in a vast (and often frighteningly dark) sea. My hope is that such a beacon may proliferate, in word and more importantly in action. I leave the decision of that proliferation up to you.

Machines and Intelligence: Against the term AI

Before I attempt to outline the various horrors enabled and expanded by the technologies that have come to be known colloquially as “Artificial Intelligence”, shortened to “AI”. I want to argue against the use of this term. First, this term is purposefully vague so as to collect disparate (albeit related) technologies under one umbrella. This collection allows the developers of individual technologies (say the Large Language Models powering ChatGPT, Claude, Gemini, etc.) to evade meaningful critique by hiding what their product actually is behind a billboard covered in the science fiction of what it may or may not become. It is hard to hit what you cannot see, harder still to strike against what is purposefully obfuscated.

For the purposes of this zine, I will be focusing mostly on Large Language Models (LLMs), the technology powering the chatbots primarily sold as the “Artificial Intelligence” science fiction has long promised. I put “intelligence” in quotes because to call these technologies “intelligent” is, too, a purposeful obfuscation as to what mechanisms are working behind the scenes. The companies peddling LLMs desperately want you to believe that they have created sentient, genuinely thinking (and perhaps even feeling) beings. They want you to believe that the machine with which you are interfacing is actually just like you, maybe even more like you than any existent physical person. They need you to believe this to hide the fact that they expect you to base your life around a glorified (proprietary) calculator, albeit one that has the ability to adapt its algorithm based upon the data fed to it (a process known as machine learning that has been used in computational research for decades).

At their core, LLMs take input, word or image, and assign numerical values to those inputs. Through a complex statistical algorithm based upon all of the data that the model has been trained on, the LLM then iteratively produces a numerical output by determining the likelihood of a given number sequentially following the last. This numerical output is then translated back into words or images depending on the context. LLMs do not *think*, they employ an incredibly intensive statistical algorithm that identifies and amplifies patterns in massive data sets and utilizes those patterns to engender the illusion of conversation.

I will not go into the details of the history of machine learning or LLMs specifically here but I highly recommend you take the time to learn a bit about that history. The conceptual framework of LLMs has existed for over a century, and the first machines making use of the machine learning that would enable LLMs came about in the 1940s and 50s. These are not new technologies or ideas, they have just finally found enough data (thanks mostly to the internet) to be trained well enough so that the smoke and mirror show holds up under scrutiny (anyone who tried to ask ChatGPT to multiply two, three-digit, numbers back in 2023 knows how quickly the illusion of intelligence can be broken).

The more insidious reason “intelligence” is so painstakingly used to describe LLMs is in order to justify the incredible costs of the technology by likening the machine to some set of human characteristics, and in turn objectify the human as machine. Sam Altman, CEO of OpenAI, frequently refers to humans as no more than complex machines so as to insinuate that if you wish to criticize the environmental impact of ChatGPT you must first criticize the energy wasted training humans who aren’t even as efficient as his beautiful machine.

Given the above, I will try to avoid using the terms “AI” or “artificial intelligence” throughout the rest of this piece, opting instead for the specificity of LLM when necessary or simply “machines” when I’m more willing to let my luddite nature shine through. The sections that follow

will attempt to offer a brief glimpse into the numerous violences enacted, enabled, and expanded by LLMs, either in their production or in their utilization.

Forest Fires and Poisoned Wells: Environmental Destruction

The criticism most frequently brought up regarding LLMs is their environmental impact. I wish to address this first due to its prevalence but also because of how easily it is often handwaved away by the proponents of LLMs. This handwaving is often done when deflecting criticisms of how much energy is used in an individual query of their machine. This deflection works because an individual query, in the grand scheme of things, doesn't use all that much energy. The rule of thumb is that a ChatGPT query uses about 10 times more energy than a google search. Even though an order of magnitude increase in energy consumption is meaningful, given how little energy a single google search uses, this implies that the energy used by a single ChatGPT query is also relatively small.

The obvious counter to this deflection is this ignores the fact that something small on an individual basis may have an outsized impact when scaled out on the order of world populations. This holds especially true within the context of an industry attempting to force LLMs into all aspects of daily life where it is expected that an individual may make hundreds of queries in a given day. This deflection is also undermined by the fact that companies like OpenAI do not serve as simply a public utility where their primary customers are individuals seeking to perform a few queries a day, but rather corporations or governments seeking to use LLMs to streamline massive projects or initiatives. These entities are using LLMs constantly, at mass volume, to perform much more complicated tasks than the average individual, meaning the energy consumption is much higher.

All of this is just the energy cost of utilizing the published version of an LLM which pales in comparison to the energy cost of training these models. It is the training that is driving the construction of thousands of data centers around the world because it takes an incomprehensible amount of computing power to plow through every single post ever written, image uploaded, video recorded, etc

The most immediate cost of these data centers is their water consumption. Running millions of CPUs necessitates cooling systems requiring millions of gallons of water daily. A large data center can use as much water in a day as a town of roughly 50,000 people. To make matters worse, these data centers are often built in areas where water is already scarce, meaning that the water they are wasting is even more desperately needed by the people actually living there. Here is where I would like to remind you that people need to drink water in order to survive. At best, this water is wasted in order to train ChatGPT to make a more accurate ripoff of a Studio Ghibli film, and at worst making for more accurate drone strikes on school children.

It has been estimated that by the end of 2026, data centers will account for nearly 6% of total energy consumption in the U.S. In order to counteract this energy consumption, many areas have turned to increased fracking to access natural gas. The costs of living near fracking sites have been well documented but the highlights are undrinkable water, increased rates of cancer due to air pollutants, and other adverse health effects. Aside from the link to fracking, data centers contribute to local air pollution directly via their backup generators primarily running on diesel.

A less studied, but increasingly common complaint of those living near data centers is noise pollution. This noise pollution causes chronic migraines, disrupted sleep, chronic nausea, chronic vertigo and numerous other adverse health effects. With little surprise to anyone familiar with the foundational principals of racial capitalism of the U.S., data centers are much more likely to be built in poorer and predominately Black and targeted non-white communities, meaning these adverse health effects are borne disproportionately by the racialized poor.

Proponents of building more data centers will also attempt to sell you the myth that these centers are only being built in areas that were not previously being used. This is a myth that maintains the colonial relation built into the bedrock of this country from inception that if land is not being utilized in the name of racial capitalism that land is useless and must be transformed. This is seen most clearly in how developers speak of deserts or of swamps or marshland. Because such areas resist attempts at development, they are often treated as hostile territories in need of conquering. The terraforming of an existent ecosystem into a concrete slab is sold as a victory for human progress, with the cost to native plant and animal life rarely worth a passing mention.

In a world already bearing the scars of more frequent and more powerful wildfires and hurricanes, this explosion of demand for more fossil fuels and more pollutants flooding the atmosphere can only accelerate the already existent catastrophe of climate change. Unfortunately, those who stand to gain from these data centers are able to insulate themselves (or at least they believe they can) from this catastrophe via the incredible wealth such data centers help them extract

There is a danger in focusing too tightly on the environmental cost of LLMs, however. When we speak almost exclusively of fossil fuels, of energy consumption, of water wasted we cede rhetorical ground to those arguing for the mass adoption of LLMs. We unintentionally cede that if there were some way to curtail the pollution and the energy waste then the mass adoption of LLMs might become a neutral or even a positive act. Because of this danger, it is vital that we become adept at articulating the many other violences of LLMs that cannot be so easily dreamed away.

Every Camera Is a Prison Guard: Mass Surveillance

Beyond environmental catastrophe, the companies behind the most prominent LLMs are deeply enmeshed within the systems of mass surveillance running rampant around the world. Companies like Flock Safety utilize these technologies to offer law enforcement agencies real-time direction as to how better target the subjects of their brutality. The most visible of this phenomenon is ICE utilizing Flock Safety systems to determine what neighborhoods to target for raids, when to target them, what type of resistance to expect, and specifically who might put up a fight.

For their partnership with ICE, Flock has caught considerable flak. Unfortunately, many who criticize Flock fail to see the dozens of other companies (Pushpak, Deep Sentinel, hell even Motorola is getting in on the action) offering similar services, but with a bit more discretion. These companies sell the dystopic fantasy of a world without crime, a world in which everyone is so aware of their being observed, that no one would dare move against the status quo. The goal of these companies is to turn as much of globe as possible into a minimum-security prison.

In fairness, this has always been the goal of surveillance systems. An aggregated, searchable database of all surveillance footage is the wet dream of any police precinct or district attorney.

However, until recent advances in LLMs there was simply too much data for a group of people to effectively comb through manually. No team of investigators could possibly identify all of the relevant cameras which may have captured the footage they need, attain the necessary warrants to view that footage, and actually spend time carefully reviewing that footage all within the timeframe necessary to either bring a charge against their target or present evidence at trial.

But with companies like Flock, suddenly that disconnected constellation of individual surveillance apparatuses becomes a single interconnected web. All of the data collected by that web is aggregated into a navigable database, and with the aid of LLMs that database becomes easily searchable for precisely the information that suits the cops in a given instance. Maybe they need all surveillance footage in the vicinity of a demo that they anticipate might involve activities deemed illegal. Or maybe they want to map the social network of an entire community, see who spends time where and with whom, understand the weekday driving habits of an individual so as to be notified when those patterns break. Based on your spending habits, your driving habits, the top stories in the current news cycle, they can predict what actions you might take and how best to police you.

Such capabilities are frightening enough even if one believes in a system of law, but for those who understand that the forces of policing and imprisonment have always acted at their own discretion in service of racial capitalism, these capabilities border on catastrophic. The prison (in the traditional sense) populations will grow, and they will grow in the racialized, gendered, and classed ways they are always growing. Those who exist in ways deemed unproductive within, or antagonistic to, the regime of racial capital will be more easily targeted and captured.

For those who understand themselves to be part of the project of unmaking this world built on carceral logic, the present is grim, but it is not lost. At the moment, it is people who need to actually carry out the violence of policing, capture, and imprisonment. So long as that remains the case, there will always be cracks to exploit in the process of sending every brick from every prison and border wall to the bottom of the ocean.

For those who have not considered themselves a partisan in the social war for, or against, imprisonment up until this point; know that there is no way out of this hellscape of mass surveillance that does not root itself in a politic (or anti-politic if you prefer) against prisons in their entirety, be they brick and mortar, or the web of surveillance equipment and LLM supported databases. We are where we are because the logic of incarceration is a dominant paradigm, and so long as it remains such a dominant paradigm, we will always end up back here one way or another.

Every detention center is a prison, and every prison is a concentration camp.

This web of surveillance expands with the inclusion of data sets aggregated by the likes of Oracle or Palantir, which work to combine the infinite scraps of data we leave behind everywhere we go. What cell towers did our smart phone connect to. What websites have we given our email address to. What do we buy, how often do we buy it, how do we pay for it. Nearly every point of interaction with capital leaves behind a footprint that two decades ago would have been a waste of time for anyone to take note of. But when collected in mass and fed to a machine specifically designed to highlight minute statistical trends these scraps sharpen into nodes that, when connected with other scraps, offer both jailer and salesman a more intimate profile of us than previously conceivable.

The implicit threat (though increasingly explicit) of these webs of surveillance is not only that transgressions of the status quo will be more easily punished through legal systems but also in

the financial systems. Rental management companies could reject your application because they used LLMs trained on the rental history of every person in your city to predict your likelihood of being late on rent, of breaking your lease, of submitting more maintenance requests than the landlord cares to service. Any potential employer could make similar assessments associated with your likelihood to request time off, to get sick, to move slower than they want. The grocery store you frequent could use these technologies adjust the prices of your food right to the limit of what their choice of LLM predicted you'd be willing to pay.

Given the trajectory of all aspects of our lives moving into the domain of speculative finance, any place where additional profits could be squeezed out of us will see the adoption of predictive technologies to maximize those profits. The web of surveillance we find ourselves tangled within ensures those with the necessary purse the most accurate predictions.

The Machine Doesn't Give a Shit About Consent: Sexual Violence and Bodily Autonomy

Tied into the web of mass surveillance capabilities enabled by the proliferation of LLMs are expansive and terrifying new domains of attacks on bodily autonomy. The most prolific of such attacks is how xAI's Grok was weaponized to mass sexually harass and intimidate women on Elon Musk's X platform. With a simple query users could prompt Grok to generate sexually explicit images of any user at any time, making the platform (already in a league of its own for misogynistic vitriol) a hub of sexual violence, often targeting children.

It is not surprising that LLMs would be most readily adopted and utilized by individuals who have always been searching for an entity they have total domination over, that they can bend in whatever way they need to fulfill their desire of the given moment. The proliferation of sexual violence on X is but the logical expansion of a world built on misogynistic violence. While this violence is foundational, the ways in which LLMs enable its expansion cannot be overstated.

Even when not wielded directly against specific women, LLMs serve to reinforce misogyny (and every other system of domination) through their sycophantic nature, casually deferring to the pre-existing beliefs of their users. Combine that sycophantic nature with an explicitly sexualized chatbot and you have a recipe for encouraging a dopamine-trained reinforcement of any number of violent predispositions towards the target of a user's sexual desire. The concept of consent becomes increasingly foreign as the user's sexual desire grows increasingly alienated from either a genuine exploration of self and/or a consensual exploration with (and of) another who is understood as an autonomous being, not simply an object of desire. Consent is already something most men have only an incredibly surface-level grasp on, these technologies only erode from there.

This alienation and associated enforced-disregard for consent is not unique to sexually explicit chatbots, a similar line of critique can (and has been made) of how the majority of pornography is produced and consumed. However, through its being a much more interactive process that simulates actual conversation, the alienation and disregard experienced is more intensely reinforced.

Entire books can, and likely will, be written on the collision of misogyny, sexual violence, and Large Language Models. Perhaps I will try to write more extensively about this dynamic myself in the future, but at present I'm still struggling for a foothold into just how to articulate all of my

thoughts on this collision. Unfortunately, this is far from the only way in which Large Language Models undermine the bodily autonomy of those most frequently targeted by misogynistic (and more specifically) transmisogynistic violence.

As stated previously, LLMs rely on increasingly robust and extensive data sets on which to train. One area that has yet to be fully tapped is medical data. Companies such as Oracle are trying to test the waters of how they might be able to profit by leasing or selling such data to companies like OpenAI to train their LLMs. Proponents of such deals wish to sell you on the idea that by selling such personal data they will enable the future of medicine, perfectly attuned medical care for each individual backed by the most powerful statistical algorithms known to man.

Ignoring the obvious fact that the vast majority of people in this country struggle to afford medical care to the point of not getting it at all, the U.S. is actively waging simultaneous wars against the individual's reproductive autonomy and gender identity. In both of these contexts, the use of personal health data to train LLMs would likely have dire consequences for those seeking medical care. There have already been instances of women being arrested for crossing state borders seeking to terminate a pregnancy given the proliferation of bans on abortions in state legislatures. Similarly, many states are drafting laws to punish parents who take their children across state lines to receive gender affirming care.

Think back to the last section in which we discussed the utilization of LLMs in mass surveillance technologies. Consider the effects of such models being able to predict the likelihood of an individual being pregnant, tracking that individual across a state border from a state in which abortion is banned into one in which it is not, and then determining that an expected pattern (if the individual were in fact still pregnant) had been broken. For a state in which abortion is illegal, this would likely be enough to trigger an investigation into that individual. The same scenario could easily be imagined for the parents a trans youth seeking gender affirming care.

Given the trajectory of this country, if these scenarios are not bleak enough for you, we could expand them to ever more grim magnitudes, however I will leave that expansion to you. I already live in a country of concentration camps, I will assume I do not need to name one for each marginalized group for you to decide whether you fight for or against them.

The Machine Cannot Disobey: Militarism

In many ways connected to the misogynistic drive for an entity entirely submissive to the commands of its user, LLMs have been readily adopted by the U.S. military. Both Anthropic and OpenAI have had various agreements for their technologies to be utilized by the Department of Defense (though Anthropic is now trying to save face by claiming Claude was used in ways the violated their agreement with the DoD). These technologies have already advised military strategy in the bombing of Iran and the kidnapping of Maduro.

The appeal of LLMs from the perspective of the military is obvious. The promise of a machine that can accurately predict the outcomes of various military operations depending on employed strategy is incredibly enticing, made even more enticing by the machine's inability to refuse an order. At its core, the military already treats the majority of those within its ranks as machines to be deployed at the discretion of their superiors. Unfortunately for those superiors, no matter how well you beat the individual will out of the majority of soldiers, they are still not perfect

machines and they can snap, they can fail, and on rare occasions they can find a conscience and object. Look no further than the frequent acts of sabotage that have been occurring on U.S. Navy ships over the last few years.

But LLMs have no conscience, that do not think, they do not feel, they are merely the illusion of sentience. They are the idealized soldier. While still taboo for many of the primary players in the LLM space, it feels inevitable that these technologies will be incorporated into automated weapons. Honestly it feels naïve to believe this isn't already taking place. No longer does the general need to fear his troops refusing to fire on civilians, risking the success of a given mission, he simply needs to give his turret chatbot the proper description of an "enemy" such that anything resembling that description is killed on sight. Again, this has always been the idealized soldier and many have people have made themselves into such machines, if often requiring some conditioning to fully remove any remaining reverence for the lives they would be ordered to take. But the advent a genuine machine that can understand a command as well as any human soldier and carry out that command without any risk of faltering due to conflicting ethics (however flimsily held to) streamlines the process of military conflict, makes war more clinical, less costly for those who already wield the most wealth and power.

The adoption of LLMs into military strategy offers a paradoxical cognitive distance between the individual and the violence that occurs due to the order he has given. It enables this individual who by all means holds the responsibility for the violence carried out under his command, to claim that he too was simply following the advice (or the orders) of the LLM. If the LLM offers the most sound strategy based on all available data, who would he be to disobey, to decide against it. And so LLMs offer this prism through which it appears (to those interested in maintaining that cognitive distance) that no orders were ever really given and yet the violence of those orders is still carried out. But because only non-orders were given to non-beings there isn't really anyone responsible for the violence carried out, and certainly nobody to be held accountable.

For those of us who reject the military on its face, who reject the mass death that such an entity necessitates the previous paragraph likely borders on absurd. Of course, orders were given, of course they were carried out, and of course there are specific people responsible. However, I fear that as LLMs become more and more prevalent in our daily lives this fact will become less and less obvious to the majority of people. As more and more people hand over more and more of their lives to these technologies, as more people relinquish their decision making to these machines in hundreds of small mundane contexts, I sense that their fear of being judged as complicit in the violence of these technologies will prevent them from being able to meaningfully identify who is responsible whenever and wherever that violence manifests.

If It Is Hard, Why Do It Myself: Critical Thought

I want to end with the harm of LLMs that may trouble me the most, despite the genuine horror I feel when thinking too long about any of the above harms. I want to end with a discussion of how the mass adoption of LLMs into our daily lives would likely affect our very ability to reason, to imagine, to think for ourselves at all.

In the three and a half years that ChatGPT has been easily available, it has become near ubiquitous among high school and college students. While I have no love for the existent educational system, I did not envision its collapse coming from the proliferation of a technology that

makes learning near impossible. While this might seem alarmist, especially given the amount of genuflection towards LLMs seen in many academic settings, I cannot find a softer way to put it.

For many, young and old alike, it is becoming second nature to “ask Chat” (meaning ChatGPT) whenever a decision needs to be made. Anecdotally, I once asked a coworker how she felt about a movie she had seen recently and she responded by asking ChatGPT to summarize why it was a good movie and I’m still haunted by this exchange more than I can express. But, specifically among those still in school, the lack of any identifiable incentive to actually learn for oneself (mathematics, writing, reading comprehension, anything) means that any inconvenience encountered in the process of trying to learn something begs for immediate relief. Why struggle through your homework when you don’t give a shit about understanding it anyway and ChatGPT can now provide reasonable answers to most questions you’ll encounter in just about any high school or early college class?

While papers in education journals constantly hedge their criticisms of LLMs by claiming they just know there must be some way to utilize these tools to help students learn more efficiently, studies are already showing that frequent use of ChatGPT is resulting in students suffering setbacks in writing proficiency and reading comprehension. A major issue is that reliance on LLMs in an academic setting is incredibly difficult to undo once that reliance has been established. If you used ChatGPT to write your first essay or complete your first homework assignment, it isn’t going to suddenly become easier to write or complete the second, in fact it is going to explicitly be more difficult.

The frequent use of LLMs in the academic context also acts as a sort of ice breaker to using them more intimately throughout the day. If you’re constantly querying ChatGPT while doing your homework why not also ask it what movies are out and which are worth going to see? Why not ask it for recommendations for what you should add to your grocery list? Pretty quickly a user who initially just wanted to skip doing their math homework has now built a certain amount of reliance on LLMs to perform daily tasks.

There is a phenomenon known as cognitive offloading that goes a long way to explain why this is likely happening. Cognitive offloading is the freeing up of mental energy that occurs when utilizing a tool in the process of completing a mental task (such as trying to understand a new concept). Proponents of LLMs will point to cognitive offloading as a positive effect of their utilization and mass adoption, suggesting that when using their chosen chatbot you clearly have more mental energy to complete other tasks. The issue with cognitive offloading is that sometimes you actually need to experience cognitive or mental strain in order to build your capacity and ability to complete more rigorous mental tasks.

In the context of academics, if a student desires to learn the complexities of classical mechanics, but relies on ChatGPT throughout their introductory physics courses, they will have never developed to necessary mental capability of working through the more difficult concepts they would encounter in that more complex course. Their reliance on ChatGPT early on necessitated their reliance on it going forward.

Now I’ll admit I don’t really give a shit about academics, but I do care about learning and more specifically I care about our aiding one another to develop the critical thinking skills necessary to meaningfully articulate the ways this world causes us to suffer and find potential paths towards meaningful strikes against that world. This is largely why I care about zines and tabling and distributing literature and encouraging others to do the same. But that project of developing

meaningful articulation and analysis is directly opposed by the mass adoption of a tool that explicitly limits and undermines one's ability to critically analyze for themselves.

We already exist in world in which the majority of people are predisposed to accept the narrative of the world around them as power dictates, to bury one's head in the sand only every poking an eye out when things become so obviously horrific that the sand is not enough for most to hide within. But even when one does begin to question why the world is this way, without concerted effort to analyze, discuss, articulate, reexamine many remain trapped in the same frameworks they believe they are escaping. One needs to look no further than the clockwork phenomenon of liberals speaking of centuries long violences as though they are newly emergent and believing that if only the purveyor of the violence changes than it isn't really violence at all.

Combine this predisposition towards accepting the world as it is with the mass adoption of a tool that prevents the growth of critical capabilities and you have a recipe for an ever more predictable populace. Easier to surveil, easier to police, easier to sell shit to, easier to keep entertained just enough to ignore the genocides and war crimes and ecosystem collapse.

I don't know what the answer is to this unease, I'm not sure there is an answer (and certainly not a singular answer) but my first thoughts are the carving out of social spaces in which there is some common understanding of the harms of these technologies from which we can build or strike out. Beyond that, I believe it is vital that the criticism of these technologies is couched within a broader critique of state and capital that is able to articulate how the development of these tools is necessitated by forces that exist in service of racial capitalism. Stop worrying about whether "AI is going to take your job" and begin to question why you have to have a job in the first place.

As I said at the outset, my intent with this piece is to light a beacon, send up a signal flair that there exists at least one other who looks out into the world and abhors what they see. I do, to my core, believe that the mass adoption of LLMs and the violence inherent to that mass adoption is not inevitable. So long as there exist those of us who refuse this world of death machines there will be a struggle against the tools and systems that maintain and reproduce that world. So long as that struggle continues nothing is inevitable.

If anything in this piece was useful, I thank you for reading. If you found it lacking, I encourage you to use that lack as a springboard for your next writing exercise or discussion with friends. There are so many necessary directions in which to expand these critiques if we are to clearly identify what exactly we're up against. I ask that you help me with that task.

The Anarchist Library (Mirror)
Anti-Copyright



Ignatius
Fairy Dust and Toxic Waste
Against the Spectacle of "Artificial Intelligence" and the World it Serves
04/04/2026

<https://cryptpad.fr/file/#/2/file/PpLd4x9pk5bHyXmXT7Qeky2t/>

usa.anarchistlibraries.net