

<https://internationalviewpoint.org/Trump-vs-Anthropic-Does-the-U-S-Want-Killer-Robots>



USA

# Trump vs. Anthropic: Does the U.S. Want Killer Robots?

- IV Online magazine - 2026 - IVP615 - April 2026 -



Publication date: Monday 27 April 2026

---

Copyright © International Viewpoint - online socialist magazine - All rights reserved

---

**“I fired them like dogs.” This is the formula that US president Donald Trump, with the elegance we know him for, used to sum up his battle with the American company Anthropic. He will have to wait a little longer before bragging: in a first order issued on 27 March 2026, the court suspended the blacklisting of the artificial intelligence (AI) giant, whose tools can therefore still be used by administrations, contrary to the president's wishes. But the tug-of-war is not over, and we must take the measure of what is at stake: the United States' use of AI in the service of a regime of terror.**

Founded by former employees of Open AI, the company that designed the famous chatbot ChatGPT, Anthropic is one of the main challengers in the field of large language models (LLMs). Its model, Claude, had reached 30 million users by mid-2025. In its communication, it emphasizes reliability and security and advances the concept of "constitutional AI", i.e. AI trained to act in accordance with founding texts such as the Universal Declaration of Human Rights. An approach that is intended to be wise and reasonable, but which did not prevent it from becoming, in November 2024, the official supplier of AI software for the US Department of Defense.

The Pentagon is using Claude, in particular, as part of a partnership that also involves the scandalous big data company Palantir, owned by far-right billionaire Peter Thiel. Palantir provides the tools to collect and process large amounts of data, and Anthropic makes it possible to use them to design action plans. In the context of the war in the Middle East, these tools have made it possible to automate the search for targets, which explains the exceptional pace of strikes. The *Wall Street Journal* also revealed that Claude had been used to plan the kidnapping of Venezuelan president Nicolás Maduro in January.

At the end of 2025, the Pentagon began negotiations to revise the terms of these contracts. At issue: restrictions on the frameworks of use, which Secretary of State Pete Hegseth wanted to sweep away in favour of the formula “any legitimate use”. Anthropic was then open to discussion, but set two red lines: mass surveillance of American citizens, and fully autonomous weapons. While claiming that these uses were prohibited by American law anyway (which is highly questionable), the Pentagon was offended, set an ultimatum, and then broke off collaboration with the company. On 4 March, Anthropic received a letter informing it of the punishment that Donald Trump had chosen for it: it would now be considered a “supply chain risk”, a status usually reserved for companies from enemy or unreliable countries, which prohibits any administration from using its services. And it was this decision that, three weeks later, was suspended by the courts – a welcome rescue for a company now in disgrace, designated as “woke radical left” by the US president.

## Settling of scores and petty cronyism

In the meantime, the competition has rubbed its hands. With calculated cynicism, OpenAI and Google published an *amicus curiae* at the beginning of the legal proceedings, an external opinion intended to enlighten the judge: in it, they defend their rival, and affirm the legitimacy of the concerns raised by Anthropic. At the same time, OpenAI was negotiating a contract to take over their place while it was still warm; and while it claims to have reaffirmed Anthropic's two red lines, and to have had them accepted by the US administration, the agreement actually reached seems rather flexible. In an internal message to his employees, Anthropic's CEO, Dario Amodei, accuses OpenAI's CEO, Sam Altman (who is also its former boss) of having engaged in a pure and simple staging to hide his opportunism. A little later, the sinister Elon Musk joined the game with his company XAI, which concluded another agreement with the Pentagon whose terms, this time, clearly no longer contain any restrictions.

Is Anthropic therefore a “woke radical left” company, committed to resisting Trumpian fascism? The chances are slim. In all the spaces where he is handed a microphone, Dario Amodei reaffirms his commitment to “national security”, and worries that the U.S. army will lose efficiency because of this affair (imagine that for several days, it could not bomb any school). His scruples about mass surveillance seem to concern only American citizens; and as far as fully autonomous weapons are concerned, his only argument is that AI is not yet reliable enough, clearly suggesting that it could become so tomorrow. In reality, setting limits on these hazardous uses is above all a way to protect himself from possible scandals, which could cause Anthropic's rating to plummet, in an AI market that is still very speculative.

From the viewpoint of view of the Trump administration, this operation is mainly about rewarding loyalties and punishing infidelities. We know the relationship, sometimes stormy but always very real, between Elon Musk and Donald Trump. What is less known is that Sam Altman, the CEO of OpenAI, was also among the Republican's donors; and Anthropic, for its part, made the mistake of supporting Kamala Harris. In the internal message already cited, Dario Amodei himself said that this was the main reason why the Pentagon was so closed in the negotiations – implying that he himself would have been ready to make a lot of concessions.

The subject is nevertheless serious, and merits being taken up by our social camp, without letting the Trumps and the Amodei define the terms of the debate. The rise of AI opens up new possibilities for mass surveillance, which the far right has already made a dogma. Palantir, Anthropic's former partner in collaborations with the Pentagon, has also distinguished itself in recent months for the help it has provided to ICE, the US immigration police: its software that has been used to identify and track migrants, with the consequences that we know. Amnesty International has also shown that Palantir's AIs have been used to identify leaders of the Palestine solidarity movement.

In the amicus curiae of OpenAI and Google, the two multinationals explain that their technologies have the potential to completely transform the type of surveillance that a state can put in place: “In 2018, there were about 70 million surveillance cameras in use in the United States, spread across airports, subway stations, parking lots, in front of stores, and on street corners. Each smartphone continuously transmits location data to carriers and dozens of apps. Credit and debit cards generate a time-stamped history of almost every business transaction made by Americans. [...] What doesn't yet exist is the AI layer that transforms this sprawling, fragmented data landscape into a unified, real-time surveillance apparatus.” How will social change activists cope?

Opposing the deployment of these tools is obviously the beginning of the answer. In the European Union, the AI Act adopted in March 2024 already prohibits states from using some of the most sordid forms of AI, such as real-time facial recognition, or software that claims to predict how likely an individual is to commit crimes – two safeguards that do not exist in the United States. But while these provisions are gains to be defended, their limits are obvious: there is actually no way to ensure that, in the secrecy of intelligence offices, these technologies are never really used. The AI Act also authorises the use of facial recognition for specific cases (search for missing children, human trafficking, terrorism), which means that the police have these tools – and can actually use them as they wish, as long as they don't do so in too obvious a way. The best way to avoid state repression boosted by facial recognition is still not to have cameras in the streets. And this is undoubtedly the reasoning that will now have to be applied to counter this “unified surveillance apparatus” that frightens even Google and Open AI: to fight step by step all the levers that states – or companies – have at their disposal to collect data, from restrictions on encrypted messaging to the systematization of card payments.

## **An international treaty against killer robots?**

As far as armaments are concerned, this is not a good time to welcome innovations in this sector with enthusiasm. At this stage, fully autonomous weapons, piloted by AI, still seem to be used only in very specific cases, generally

defensive and without human casualties (the interception of a missile, for example). However, the Anthropic affair shows that killer robots are no longer science fiction.

Characterizing the current situation is a difficult task, as autonomous weapons are difficult to characterize. Many weapons already deployed, in particular the drones massively used in the war in Ukraine, have a significant level of autonomy. Officially, the armies of the major powers, including those of the United States, all claim the doctrine of “man in the loop”. But this expression is subject to interpretation: what is the role of the human in question? To set a target? To validate the one proposed by the system? To monitor the system and regain control if it makes mistakes? To activate the system and let it engage in combat without a specific target? And of course, claims of this nature are not always verifiable in practice. In 2021, a UN report established that a Turkish drone in Libya had opened fire entirely autonomously.

These developments trigger concerns on several levels. The first level is mainly a fantasy: that of machines with a will of their own, beyond the control of their designers. Abundantly fed by science fiction, but also by ambiguous lexical choices (“killer robots”, or even the term “autonomous”), this apocalyptic figure makes it possible to deflect the debate and reassure populations at low cost – as the French Minister of the Armed Forces said, “Terminator will not parade on the Champs-Élysées”. No army, in reality, has an interest in developing a weapons system that sets its own objectives independently of the state's strategies and tactics: the autonomy in question always consists of following a program written in advance, with well-defined objectives, while introducing a certain degree of adaptability to changing conditions.

The most commonly accepted red line is the ability of a weapons system to choose a target itself – this is what Anthropic refuses to contribute to, with a technical argument: current AIs are not (yet) capable of making such a choice reliably. Clearly, the risk of killing civilians by mistake is too great. This is a second level of concern, largely legitimate; after all, it is for this reason that anti-personnel mines, which can be considered as a first form of autonomous weapon, were banned by an international treaty in 1997 (signed by 161 states, but not the United States, Russia or China). But arguments of this kind can also represent a trap, because they open the way to a deployment of these technologies once they have become efficient enough to make as many or fewer mistakes as human soldiers, which could well be possible tomorrow.

Our rejection of these systems must mobilize a third level: the automation of warfare, whether it is autonomous weapons or the applications of AI to intelligence, simply gives too much power to states. The shift, at the end of the twentieth century, from conscript armies to professional armies, was already a giant step towards the concentration of the power to kill: where the former, deeply linked to the population, were often the scene of protests and mutinies that were sometimes difficult to quell, the latter have become much more disciplined – and much more capable of committing atrocities without batting an eyelid. Far from the imaginary of the robot turning against its creator, military AIs are dangerous precisely because they are the ultimate disciplined soldier. Add to this the fact that the development of these technologies to their full potential, which requires gigantic resources, will probably only be accessible to a few great powers: we have a world where they will be able to take the decision to engage in totally asymmetrical conflicts, ravaging countries with very few human losses.

Humanity already has international treaties limiting the use of nuclear, chemical and bacteriological weapons. They are largely insufficient, and the horizon must remain that of the total dismantling of arsenals in these three areas; but they have the merit of existing, and it is reasonable to think that they have made it possible to avoid some disasters.

A few years ago, the UN began negotiations for a similar treaty on lethal autonomous weapons systems, following the positions taken by many countries (especially from the Global South), a broad coalition of NGOs, a large part of the AI research community, the UN Secretary-General, and even the Catholic Church. They came to nothing. And the list of countries that blocked the process will come as no surprise: mainly the United Kingdom, Australia, India,

## Trump vs. Anthropic: Does the U.S. Want Killer Robots?

---

the United States, Russia and Israel.

At a time when the imperialist powers are seeking to reassert their domination in blood and suffering, curbing the race for the most nightmarish weapons technologies is a political priority. And for this, it is better not to rely on private multinationals.

21 April 2026

Translated by **International Viewpoint** from [Gauche Anticapitaliste](#).

PS:

If you like this article or have found it useful, please consider donating towards the work of International Viewpoint. Simply follow this link: [Donate](#) then enter an amount of your choice. One-off donations are very welcome. But regular donations by standing order are also vital to our continuing functioning. See the last paragraph of [this article](#) for our bank account details and take out a standing order. Thanks.