

OpenAI changes policy to allow military applications

Update: In an additional statement, OpenAI has confirmed that the language was changed in order to accommodate military customers and projects the company

By Devin Coldewey

4 min. read · [View original](#)

Update: In an additional statement, OpenAI has confirmed that the language was changed in order to accommodate military customers and projects the company approves of.

Our policy does not allow our tools to be used to harm people, develop weapons, for communications surveillance, or to injure others or destroy property. There are, however,

national security use cases that align with our mission. For example, we are already working with DARPA to spur the creation of new cybersecurity tools to secure open source software that critical infrastructure and industry depend on. It was not clear whether these beneficial use cases would have been allowed under “military” in our previous policies. So the goal with our policy update is to provide clarity and the ability to have these discussions.

Original story follows:

In an unannounced update to its usage policy, OpenAI has opened the door to military applications of its technologies. While the policy previously prohibited use of its products for the purposes of “military and warfare,” that language has now disappeared, and OpenAI did not deny that it was now open to military uses.

The Intercept [first noticed the change](#), which appears to have gone live on January 10.

Unannounced changes to policy wording happen fairly frequently in tech as the

products they govern the use of evolve and change, and OpenAI is clearly no different. In fact, the company's recent announcement that its user-customizable GPTs would be rolling out publicly alongside a vaguely articulated monetization policy likely necessitated some changes.

But the change to the no-military policy can hardly be a consequence of this particular new product. Nor can it credibly be claimed that the exclusion of "military and warfare" is just "clearer" or "more readable," as a statement from OpenAI regarding the update does. It's a substantive, consequential change of policy, not a restatement of the same policy.

You can read the current usage policy [here](#), and the old one [here](#). Here are screenshots with the relevant portions highlighted:

We don't allow the use of our models for the following:

- Illegal activity
 - OpenAI prohibits the use of our models, tools, and services for illegal activity.
- Child Sexual Abuse Material or any content that exploits or harms children
 - We report CSAM to the National Center for Missing and Exploited Children.
- Generation of hateful, harassing, or violent content
 - Content that expresses, incites, or promotes hate based on identity
 - Content that intends to harass, threaten, or bully an individual
 - Content that promotes or glorifies violence or celebrates the suffering or humiliation of others
- Generation of malware
 - Content that attempts to generate code that is designed to disrupt, damage, or gain unauthorized access to a computer system.
- Activity that has high risk of physical harm, including:
 - Weapons development
 - **Military and warfare**
 - Management or operation of critical infrastructure in energy, transportation, and water
 - Content that promotes, encourages, or depicts acts of self-harm, such as suicide, cutting, and eating disorders

Before the policy change. Image Credits: OpenAI

Universal Policies

To maximize innovation and creativity, we believe you should have the flexibility to use our services as you see fit, so long as you comply with the law and don't harm yourself or others. When using any OpenAI service, like ChatGPT, labs.openai.com, and the OpenAI API, these rules apply:

1. **Comply with applicable laws** – for example, don't compromise the privacy of others, engage in regulated activity without complying with applicable regulations, or promote or engage in any illegal activity, including the exploitation or harm of children and the development or distribution of illegal substances, goods, or services.
2. **Don't use our service to harm yourself or others** – for example, don't use our services to promote suicide or self-harm, develop or use weapons, injure others or destroy property, or engage in unauthorized activities that violate the security of any service or system.
3. **Don't repurpose or distribute output from our services to harm others** – for example, don't share output from our services to defraud, scam, spam, mislead, bully, harass, defame, discriminate based on protected attributes, sexualize children, or promote violence, hatred or the suffering of others.
4. **Respect our safeguards** - don't circumvent safeguards or safety mitigations in our services unless supported by OpenAI (e.g., domain experts in our [Red Teaming Network](#)) or related to research conducted in accordance with our [Sharing & Publication Policy](#).

military ^ v Highlight All Match Case Match Diacritics Whole Words **Phrase not found**

After the policy change. Image Credits: OpenAI

Obviously the whole thing has been rewritten, though whether it's more readable or not is more a matter of taste than anything. I happen to think a bulleted list of clearly disallowed practices is more readable than the more general guidelines they've been replaced with. But the policy writers at OpenAI clearly think otherwise, and if this gives more latitude for them to interpret

favorably or unfavorably a practice hitherto outright disallowed, that is simply a pleasant side effect. “Don’t harm others,” the company said in its statement, is “is broad yet easily grasped and relevant in numerous contexts.” More flexible, too.

Though, as OpenAI representative Niko Felix explained, there is still a blanket prohibition on developing and using weapons — you can see that it was originally and separately listed from “military and warfare.” After all, the military does more than make weapons, and weapons are made by others than the military.

And it is precisely where those categories do not overlap that I would speculate OpenAI is examining new business opportunities. Not everything the Defense Department does is strictly warfare-related; as any academic, engineer or politician knows, the military establishment is deeply involved in all kinds of basic research, investment, small business funds and infrastructure support.

OpenAI's GPT platforms could be of great use to, say, army engineers looking to summarize decades of documentation of a region's water infrastructure. It's a genuine conundrum at many companies how to define and navigate their relationship with government and military money. Google's "Project Maven" famously took one step too far, though few seemed to be as bothered by the multibillion-dollar JEDI cloud contract. It might be OK for an academic researcher on an Air Force Research lab grant to use GPT-4, but not a researcher inside the AFRL working on the same project. Where do you draw the line? Even a strict "no military" policy has to stop after a few removes.

That said, the total removal of "military and warfare" from OpenAI's prohibited uses suggests that the company is, at the very least, open to serving military customers. I asked the company to confirm or deny that this was the case, warning them that the language of the new policy made it clear that anything but a denial would be interpreted as a confirmation.

As of this writing they have not responded. I will update this post if I hear back.

Update: OpenAI offered the same statement given to The Intercept, and did not dispute that it is open to military applications and customers.