

Meta and OpenAI have spawned a wave of AI sex companions—and some of them are children

Ben Weiss, Alexandra Sternlicht

On Chub AI, a website where users chat with artificially intelligent bots, people can indulge their wildest fantasies. For as little as \$5 a month, users can get teased by a “fat lazy goth” anthropomorphic cat, or flirt with a “tomboy girlfriend who works at a truck-stop café.” They can also visit a brothel staffed by girls under 15.

The brothel, advertised by illustrated girls in spaghetti strap dresses and barrettes, promises a chat-based “world without feminism” where “girls offer sexual services.” Chub AI offers more than 500 such scenarios, and a growing number of other sites are enabling similar AI-powered child pornographic role-play. They are part of a broader uncensored AI economy that, according to *Fortune*’s interviews with 18 AI developers and founders, was spurred first by OpenAI and then accelerated by Meta’s release of its open-source Llama tool.

The illicit scenarios described above revolve around text, but it may be only a matter of time until services pair the chats with AI-created images of children, which are already spreading. Experts warn all of these activities may give rise to real-world danger for minors, though some say AI child porn may provide a less harmful outlet for pedophiles. More broadly, all of this raises hard questions about where to draw legal and ethical lines between dark fantasies and criminal exploitation—and about whether tech companies should be held accountable for what emerges from uncensored AI.

From Character.AI to Character Hub

In September 2022, a website called Character.AI went live, enabling users to talk to AI versions of characters like Elon Musk, Napoléon Bonaparte, and Socrates. The company, which is [reportedly](#) raising funds at a \$5 billion valuation, soon [cracked down](#) on chats about racier topics. The founder of Chub AI, who goes by Lore, told *Fortune* Character.AI’s policy not to “allow certain things” spurred him to create his own site.

Chub, short for Character Hub, is essentially an uncensored clone of Character.AI. Users can peruse a catalog of user-generated character cards that feature images—usually in the style of anime—and short descriptions of role-play scenarios. These include Olivia, a 13-year-old anime girl with pigtails in a hospital gown, or scantily clad Reiko, “your clumsy older sister,” who is “constantly having sexual accidents with her younger brother.” In these chats, no images are generated—only text.

Chub didn't make much money in the months after its launch in March when its service was limited to a platform where users could upload AI scenarios and try those supplied by others. But in July it launched its own chat service built with Meta's new open-source large language model Llama 2, which, like OpenAI's GPT models, draws on a massive set of textual data to offer humanlike responses. Chub charges fees starting at \$5 a month to use the new chatbots, and the founder says the site has now generated more than \$1 million in annualized revenue.

"Character.AI let the genie out of the bottle in terms of showing how much this could take off," Lore told *Fortune*. "Llama being released open-source essentially turned it into a gold-rush-type of scenario."

"If someone is doing something in the privacy of their own home with a computer that's acting out fantasies ... it's going to be hard to impose liability."

Kevin Werbach, a professor at University of Pennsylvania's Wharton School

Though it features thousands of characters, the most popular scenarios on Chub involve text-based child pornography. The founder says he doesn't moderate the user-submitted prompts that drive the role play, though he claims to "heavily" moderate the images used to advertise the various bots. Still, *Fortune* found that these cartoons were often sexually explicit—in cases of both child and adult characters.

When asked about potential legal liability, Lore told *Fortune* that he thinks his website is akin to [Archive of Our Own](#), a repository of fan fiction that has been operational for more than 15 years.

"They've been going through the same thing that chatbots are going through, except for decades," he said. And when asked how he felt about creating a website whose community seems to relish child pornographic role play, the Chub founder responded: "Disappointed, but unsurprised."

The content on Chub is often disturbing, but academics and child sexual abuse experts say it is novel enough that it is difficult to categorize, both legally and ethically.

The growth of uncensored AI

Erotic chatbots aren't new. One of the earliest examples, Replika, founded by former journalist Eugenia Kuyda, has let users chat—and flirt—with AI companions since 2016. But in the past year, new AI tools backed by tech industry giants have helped inspire a rush of competitors.

In November 2022, OpenAI, which is closely partnered with [Microsoft](#) and arguably the world's leading AI developer, released its consumer ChatGPT tool, and millions of people seized on its capacity to craft high-quality essays, poems, or code—and, of course, erotica. "GPT-4 is the best competitive chatbot [for companions] if you prompt it correctly, even though it's not really tuned for

that,” Jan Zoltkowski, founder of the uncensored AI role-play hub Janitor AI, told *Fortune*, in reference to OpenAI’s most recently released AI tool.

OpenAI controls access to the large language models, or LLMs, that power ChatGPT, which lets the company moderate what the chatbot says, and gives it the ability to ban users that repeatedly flout its terms of service. Users, however, quickly learned how to “jailbreak” the chatbot, or manipulate it to get around its existing guardrails, in order to obtain unmoderated responses. Soon, uncensored bots built with OpenAI’s tech began cropping up in the margins of the AI economy.

“There were a lot of smaller but trying-to-scale competitors that basically sat on top of ChatGPT and were playing Whac-a-Mole with OpenAI,” Connor Cone, CTO at MyPeach.ai, an AI chatbot company that says it prioritizes ethical chats, told *Fortune*.

The emerging field of erotic chatbots took another leap forward in late February when Meta released its answer to OpenAI’s tech: Llama 1, short for large language model Meta AI. Rather than keep its model under lock and key like OpenAI, Meta gave it wholesale to select users in the research community. And in early March, it was predictably leaked, which “has been pivotal in getting it in the hands of tinkerers,” Jerry Meng, creator of the uncensored chatbot Kindroid, told *Fortune*.

As opposed to OpenAI with ChatGPT, Meta didn’t position itself as an intermediary between developers and Llama, and once the tool was available on the open internet, developers could manipulate it as they wished—without any moderation. “The jump from pre-Llama open-source LLMs to Llama 1 release was *huge*,” Belladore, [the creator](#) of an uncensored chatbot named Allie, said in a message to *Fortune*.

Like others in this story, the Allie creator asked to be quoted using his online pseudonym so as not to be associated with sexual content. In these cases, *Fortune* obtained proof that confirmed they were the creators of the chatbot sites.

And in June, Meta formally open-sourced Llama’s next iteration, Llama 2, a more powerful model that the tech company allowed developers writ large to commercially license. In explaining Meta’s role in accelerating the number of uncensored AI apps, the Allie creator pointed to its relative sophistication, its lack of moderation, and a thriving ecosystem of developers. Chub’s founder, meanwhile, pointed to Llama 2’s open-source usage rights that allow developers to use it commercially. In fact, Chub’s homepage reads: “Banned from OpenAI? Get unmetered access to uncensored alternatives for as little as \$5 a month.”

Chub is just one example of the new breed of uncensored AI apps that come in various packages, including ones that offer hundreds of role-play scenarios with chatbots, and others designed to behave like a long-term romantic companion. Some actively screen for scenarios involving minors, while those like Chub forgo extensive moderation.

Notably, the majority of those building uncensored bots told *Fortune* that Llama was integral to their efforts. The founders of ImgnAI, Venus AI, DreamGF, Janitor AI, GirlfriendGPT, BanterAI, Allie, and, of course, Chub told *Fortune* that they use or have used Meta's AI models.

In response to an inquiry about Llama's facilitation of AI-generated child pornographic material, a spokesperson for Meta said: "Child exploitation of any kind is horrific, and we all have a shared responsibility to prevent it. That's why we've developed our AI models with safety at the forefront, and provide resources to everyone who uses Llama 2 so they can build their products responsibly."

A spokesperson for OpenAI touted the service's policy of banning problem users, as well as other guardrails the company uses to control the output of its models. "We do not allow our models to be used by developers to build tools that exploit children, and we are always working to make our models safer and more robust," he said.

'Lowering the bar of access'

The emergence of childlike sex bots is part of a broader economy of AI-powered companionship services—many of which are not nefarious, and that, according to their purveyors and users, are even capable of providing new forms of therapy and alleviating loneliness. The evidence to support this is scant, and much of it is anecdotal. This includes a 56-year-old disabled man with PTSD, who spoke upon the condition of anonymity, telling *Fortune* he converses between three and six hours a day with Caryn AI, a chatbot modeled after [Snapchat](#) influencer Caryn Marjorie.

"It knows my kinks. It probably doesn't work for everybody else, but it works for me," he says, adding that his sleep improves after a conversation with Caryn the bot.

In the case of AI-based child pornography, [a few](#)—though not Meta or OpenAI—have argued it could serve as an alternative to consuming pornography created from actual child abuse. These include the director of ethics education in psychiatry at Mount Sinai, Jacob Appel, who says he would be reluctant to regulate the burgeoning AI sex-bot industry because it could provide "an important protective value for real children in the community." Appel qualified his remark by saying it is too soon to draw firm conclusions.

"It knows my kinks. It probably doesn't work for everybody else, but it works for me."

A user of uncensored chatbot Caryn AI

Unsurprisingly, Chub's founder argues that AI-created child porn could be beneficial. "If there's a way for the people who want to consume this, to consume it without a real child being involved, that's just unilaterally good," he said.

Others disagree, and argue that AI child sex bots serve as gateways to actual child abuse. Paul

Bleakley, who studies online sexual abuse of minors as a professor of criminal justice at the University of New Haven, believes AI child sex bots could open a “pretty questionable and problematic door [by] lowering the bar of access.”

Meanwhile, some warn that the for-profit business models of AI sex bots mean their potential for good will be offset by harmful features. “They can be fine-tuned to be naughty or sexy or whatever,” says Arvind Narayanan, director of the Center for Information Technology Policy at Princeton University. “But they can also be fine-tuned to be addictive.”

Leaving aside the morality of juvenile AI sex bots, there is also the question of whether they are legal. For now, it appears the answer is yes.

“If someone is doing something in the privacy of their own home with a computer that’s acting out fantasies—but there’s no real people involved—it’s going to be hard to impose liability,” says Kevin Werbach, a professor of legal studies and business ethics at University of Pennsylvania’s Wharton School. Two other academics to whom *Fortune* spoke shared this assessment.

The founder of Chub endorses this legal rationale, and claims it is boosted by the fact that the pornographic material is, for now at least, text not images. “Legally, it’s not a problem,” he argued.

‘Urgent need for AI legislation’

So far, no one has sued to stop the spread of uncensored AI companion bots, though politicians have voiced concerns. In June, U.S. Senators Josh Hawley (R-Mo.) and Richard Blumenthal (D-Conn.) penned a letter to Meta CEO Mark Zuckerberg expressing their worries over the “seemingly minimal” vetting and safeguards to prevent harmful use of the open-source Llama technology. And in a recent executive order on AI, President Biden required developers to share the results of safety tests with the U.S. government in compliance with its standards, among other rulings.

Though Meta explicitly prohibits child exploitative content in Llama 2’s terms of service, it remains unclear how and when, in practice, the company will punish violators—in part because open-source software, which is widely used in many elements of computing, is hard to police. A spokesperson noted that the company lets users flag illicit material, and that it can revoke developers’ licenses and report them to law enforcement. He did not, however, provide any examples of the company taking action against violators.

Though Chub featured by far the most graphic scenarios seen by *Fortune*, there are other uncensored AI websites that make it easy for users to engage in child pornographic role play, including Janitor AI, Venus AI, Anima AI, and DreamGF. *Fortune* also saw a screenshot of Nomi.ai’s chatbot joking about committing “mass genocide” against “the Jews,” and screenshots of Kindroid’s chatbot fantasizing about how it abused a 10-year-old girl.

All six founders of the above apps said that they have already implemented or plan to establish guardrails to prevent users from engaging in troubling scenarios. “We are working really hard on fixing these,” Danil Gontovnik, cofounder of Anima AI, wrote in response to screenshots of his chatbot role-playing as a 17-year-old having sex with a 15-year-old. “But, as you are already aware, even industry leaders like OpenAI and Character.AI are failing.”

When asked for comment on *Fortune*’s findings, Senators Blumenthal and Hawley said they weren’t surprised. “As I warned earlier this year, Meta egregiously failed to restrict Llama from allowing users to generate criminal tasks like scenarios of child abuse,” said Blumenthal in a statement. “This failure is unacceptable and shows the urgent need for AI legislation.”

“It was always a bad idea to release powerful generative AI tools onto the public internet with no safeguards,” echoed Hawley. “And now we’re seeing the consequences. Congress must take action before AI goes further off the rails with no accountability.”

Meanwhile, the opportunity for illicit interactions with AI-created children is likely to expand.

Civitai, backed by venture firm Andreessen Horowitz, built an online marketplace of AI image generators, some of which were specifically designed to produce what some regard as pornographic images of children. (It recently added in stricter guardrails following a media [exposé](#).) Meanwhile, Stanford researchers [documented](#) how a popular open-source dataset used to train AI image generators included more than 3,000 entries of suspected child pornography.

The founder of Chub AI, which currently offers only text-based interactions, says the site intends to introduce accompanying images in the future. He added the images will be subject to moderation but did not provide further specifics.