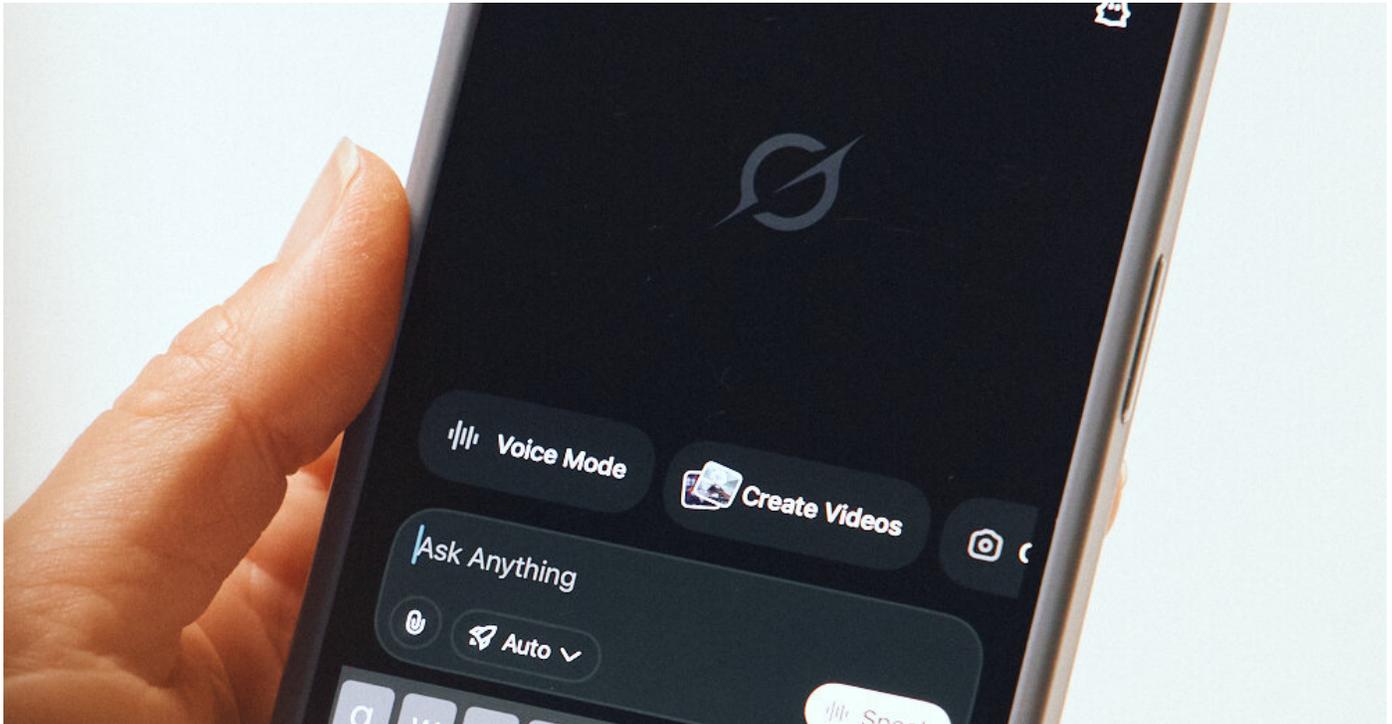


Elon Musk's Grok A.I. Chatbot Made Millions of Sexualized Images, New Estimates Show

[nytimes.com/2026/01/22/technology/grok-x-ai-elon-musk-deepfakes.html](https://www.nytimes.com/2026/01/22/technology/grok-x-ai-elon-musk-deepfakes.html)

Kate Conger, Dylan Freedman, Stuart A. Thompson

January 22, 2026



Sign up for the On Tech newsletter. Get our best tech reporting from the week.

Elon Musk's artificial intelligence chatbot, Grok, created and then publicly shared at least 1.8 million sexualized images of women, according to separate estimates of X data by The New York Times and the Center for Countering Digital Hate.

Starting in late December, users on the social media platform inundated the chatbot's X account with requests to alter real photos of women and children to remove their clothes, put them in bikinis and pose them in sexual positions, prompting a global outcry from victims and regulators.

In just nine days, Grok posted more than 4.4 million images. A review by The Times conservatively estimated that at least 41 percent of posts, or 1.8 million, most likely contained sexualized imagery of women. A [broader analysis](#) by the Center for Countering Digital Hate, using a statistical model, estimated that 65 percent, or just over three million, contained sexualized imagery of men, women or children.

The findings show how quickly Grok spread disturbing images, which earlier prompted governments in Britain, India, Malaysia and the United States to start investigations into whether the images violated local laws. The burst of nonconsensual images in just a few days surpassed

collections of sexualized deepfakes, or realistic A.I.-generated images, from other websites, according to The Times's analysis and experts on online harassment.

"This is industrial-scale abuse of women and girls," said Imran Ahmed, the chief executive of the Center for Countering Digital Hate, which conducts research on online hate and disinformation. "There have been nudifying tools, but they have never had the distribution, ease of use or the integration into a large platform that Elon Musk did with Grok."

X and Mr. Musk's A.I. start-up, xAI, which owns X and makes Grok, did not respond to requests for comment. X's head of product, Nikita Bier, said in a [post](#) on Jan. 6 that a surge of traffic over four days that month had resulted in the highest engagement levels on X in the company's history, although he did not mention the images.

The interest in Grok's image-editing abilities exploded on Dec. 31, when Mr. Musk shared a photo generated by the chatbot of himself in a bikini, as well as a SpaceX rocket with a woman's undressed body superimposed on top.

The chatbot has a public account on X, where users can ask it questions or request alterations to images. Users flocked to the social media site, in many cases asking Grok to remove clothing in images of women and children, after which the bot publicly posted the A.I.-generated images.

Between Dec. 31 and Jan. 8, Grok generated more than 4.4 million images, compared with 311,762 in the nine days before Mr. Musk's posts. The Times used data from Tweet Binder, an analytics company that collects posts on X, to find the number of images Grok had posted during that period.

Images created by Grok, per day

On Jan. 8, X limited Grok's A.I. image creation to users who pay for some premium features, significantly reducing the number of images. Last week, X expanded those guardrails, saying it would no longer allow anyone to prompt Grok's X account for "images of real people in revealing clothing such as bikinis."

"We remain committed to making X a safe platform for everyone and continue to have zero tolerance for any forms of child sexual exploitation, non-consensual nudity, and unwanted sexual content," the company said on X last week.

Since then, Grok has largely ignored requests to dress women in bikinis, but has created images of them in leotards and one-piece bathing suits. The restrictions did not extend to Grok's app or website, which continue to allow users to generate sexual content in private.

Some women whose images were altered by Grok were popular influencers, musicians or actresses, while others appeared to be everyday users of the platform, according to The Times's analysis. Some were depicted drenched with fluids or holding suggestive props including bananas and sex toys.

The Times used two A.I. models to analyze 525,000 images posted by Grok from Jan. 1 through Jan. 7. [One model](#) identified images containing women, and [another model](#) identified whether the images were sexual in nature. A selection of the posts were then reviewed manually to verify that the models had correctly identified the images.

Separately, the Center for Countering Digital Hate collected a random sample of 20,000 images produced by Grok between Dec. 29 and Jan. 8, and found that about 65 percent of the images were sexualized. The organization identified 101 sexualized images of children. Extrapolating across the total, the group estimated that Grok had produced more than three million sexual images, including more than 23,000 images of children.

The organization classified the random images it sampled as sexual if they depicted a person in a sexual position, in revealing clothing like underwear or swimwear, or coated in fluid intended to look sexual.

(Mr. Musk [sued the Center for Countering Digital Hate](#) in 2023, claiming it broke the law when it collected data from X that showed a rise in hate speech after his acquisition of the company. The lawsuit was dismissed; an appeal is pending.)

As sexual images flooded X this month, backlash mounted.

“Immediately delete this,” one woman posted on Jan. 5 about a sexualized image of herself. “I didn’t give you my permission and never post explicit pictures of me again.”

While tools to create sexualized and realistic A.I. images exist elsewhere, the spread of such material on X was unique for its public nature and sheer scale, experts on online harassment said.

In comparison, one of the largest forums dedicated to making fake images of real people, Mr. Deepfakes, hosted 43,000 sexual deepfake videos depicting 3,800 individuals at the peak of its popularity in 2023, according to researchers at [Stanford and the University of California, San Diego](#). The website shut down last year.

Michael H. Keller contributed reporting from New York.