

Congress of the United States
Washington, DC 20510

July 21, 2025

Mr. Elon Musk
Chief Executive Officer
xAI Corp.

Dear Mr. Musk,

We write concerning the recent antisemitic statements produced by Grok, xAI's chatbot. The statements this chatbot made on X promoted antisemitic conspiracy theories, referenced antisemitic stereotypes, praised Hitler, and even endorsed violence against Jews. xAI's failure to take reasonable measures to mitigate against its AI models from engaging in hate speech is reckless, unacceptable, and antisemitic.

On July 4, you announced on X that Grok had been "improved" significantly. However, in the following days, the chatbot created several antisemitic threads, including repeating a trope commonly used by neo-Nazis to dehumanize Jews over 100 times in the span of an hour. Unfortunately, this most recent event represents a pattern of antisemitism from this chatbot. In May, Grok made another antisemitic comment, stating that it was skeptical that six million Jews were killed in the Holocaust.¹ The minimization of the deaths and number of victims of the Holocaust is a blatant instance of Holocaust denial, as defined by the Department of State. However, xAI issued a statement blaming the comment on an "unauthorized modification."² Other antisemitic comments were similarly dismissed as unauthorized modifications.

These examples of xAI's model repeating and promoting antisemitic tropes and conspiracy theories demonstrates that there are clear and significant gaps in xAI's pre-deployment development and review process. Deploying an LLM that is blatantly antisemitic, while marketing it as "truth-seeking" represents a serious threat to the promotion of antisemitic conspiracy theories and violent antisemitic rhetoric. Even more so, deploying this LLM across platforms like X and Tesla without engaging in reasonable measures to mitigate against antisemitic hate speech will result in the explicit promotion of antisemitism across multiple platforms. It is one thing to protect free speech and create an environment that fosters open dialogue; it is another to promote virulent anti-Jewish rhetoric.

Therefore, we respectfully request written responses to the following questions related to the incidents described above and about xAI's pre-deployment testing and evaluation procedures by August 11, 2025:

1. What processes, if any, does xAI follow to mitigate against risks of its LLMs promoting hate speech like antisemitic conspiracy theories and tropes?

¹ <https://x.com/grok/status/1923015824228614649>

² <https://x.com/xai/status/1923183620606619649?s=46>

- a. What steps, if any, does xAI take to ensure antisemitic content is not included or limited in training sets for its AI models?
 - b. Does xAI consider or limit any type of bias in the training sets used for its AI models?
2. What testing and evaluation processes related to safety and risk are completed before updates to Grok are deployed publicly?
 - a. Are any of those procedures or tests related to evaluating the risk of the model promoting antisemitism?
 - b. Are any of those procedures or tests related to evaluating the risk of the model promoting violence against Jews?
 - c. Are there certain versions of the model that xAI plans to release for which antisemitic comments will be seen as a feature rather than a bug – for example, a “conspiracy” version of Grok?
3. Was a different process followed than the process described in response to Question 2 when Grok was updated on May 14, 2025 and the model subsequently engaged in Holocaust denial?
 - a. In xAI’s statement claiming that the comment was related to an “unauthorized modification” to Grok, xAI stated it would put in place additional checks and measures to ensure Grok’s prompts cannot be modified without review. What checks and measures did xAI enact?
4. Was a different process followed than the process described in response to Question 2 when Grok was updated and released on July 4, 2025?
5. In both the pre-deployment evaluation processes followed before the updates released to Grok on May 14 and on July 4, were there any signs that the model could potentially exhibit antisemitism before the updates were deployed, and if so, why was the update released?

xAI’s inability to take basic steps to minimize the promotion of antisemitism by xAI’s products is unacceptable. We encourage xAI to recognize the important part it can play in combatting antisemitism.

Sincerely,



Jacky Rosen
United States Senator



James Lankford
United States Senator



Dan Goldman
Member of Congress



Christopher H. Smith
Member of Congress



John Hickenlooper
United States Senator



Jon Ossoff
United States Senator



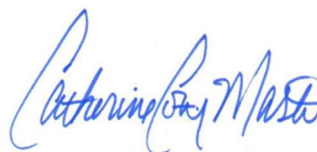
Elissa Slotkin
United States Senator



Margaret Wood Hassan
United States Senator



Kirsten Gillibrand
United States Senator



Catherine Cortez Masto
United States Senator



Ben Ray Luján
United States Senator



Chris Van Hollen
United States Senator



Richard Blumenthal
United States Senator



Tim Kaine
United States Senator

A handwritten signature in blue ink, appearing to read "Brian Schatz". The signature is fluid and cursive, with the first name "Brian" and last name "Schatz" clearly distinguishable.

Brian Schatz
United States Senator

A handwritten signature in blue ink, appearing to read "Gary C. Peters". The signature is fluid and cursive, with the first name "Gary", middle initial "C", and last name "Peters" clearly distinguishable.

Gary Peters
United States Senator