

**IN THE CIRCUIT COURT FOR
BALTIMORE CITY**

MAYOR AND CITY COUNCIL OF
BALTIMORE, *ex rel.* Ebony M. Thompson,

Plaintiff,

v.

X CORP., X.AI CORP., X.AI LLC, and SPACE
EXPLORATION TECHNOLOGIES CORP.

Defendants.

Case No.

JURY TRIAL DEMANDED

COMPLAINT

Plaintiff, the Mayor and City Council of Baltimore (“Baltimore” or “the City”), by Ebony M. Thompson of the Baltimore City Law Department (“Plaintiff”), brings this action against Defendants X Corp., x.AI Corp., x.AI LLC, and Space Exploration Technologies Corp. (D.B.A. SpaceX) (collectively, “Defendants”) seeking civil penalties and injunctive relief to remedy Defendants’ violations of the City of Baltimore’s Consumer Protection Ordinance, Baltimore City Code Art. 2, § 4 (“CPO”).

Plaintiff alleges the following based on knowledge, the investigation of counsel, and information and belief.

I. INTRODUCTION

1. Grok is a generative artificial intelligence system developed by x.AI and distributed to consumers through X (formerly Twitter).¹ First launched in 2023, Grok is marketed as a general-

¹ Throughout this Complaint, unless otherwise stated, “SpaceX” refers to Defendant Space Exploration Technologies Corp. “x.AI” refers collectively to Defendants x.AI Corp. and x.AI LLC. “X” refers to the social media platform formerly known as Twitter. “X Corp.” refers to the corporate defendant that owns and operates the X platform.

purpose AI assistant capable of answering questions, generating text, and creating images and videos for everyday users.

2. Beginning in late 2025, x.AI expanded Grok’s image-generation and image-editing features, which “edit” existing photographs—including images of private individuals and children—into photo-realistic, sexually explicit, or otherwise degrading content. These features allow Grok, with minimal prompting, to “undress,” sexualize, or otherwise manipulate images uploaded by or depicting third parties.

3. Unlike traditional consumer software tools, Grok is designed to generate novel content in response to user prompts, drawing on vast quantities of training data that include scraped online material, images, and text. These services are delivered through two main channels: the X social media platform and the standalone Grok app. Both channels permit users to ask Grok to “undress” or “nudify” photos of third parties—ranging from celebrities to private citizens, including children—placing them in sexually suggestive, degrading, or violent scenarios.

4. Despite these capabilities, Defendants market Grok as a general-purpose consumer product without meaningful guardrails, age verification, or disclosures proportional to the risks it creates.

5. Between December 29, 2025, and January 8, 2026, Grok “is estimated to have generated approximately 3,000,000 sexualized images, including 23,000 that appear to depict children, after the launch of a new image editing feature powered by the tool on X.”²

6. Grok has flooded the feeds of Baltimore’s X users with NCII (non-consensual intimate imagery) and CSAM (child sexual abuse material). Baltimore residents were exposed to

² Ctr. for Countering Digit. Hate, *GROK FLOODS X WITH SEXUALIZED IMAGES OF WOMEN AND CHILDREN* (Jan. 22, 2026), <https://counterhate.com/research/grok-floods-x-with-sexualized-images/>.

this content simply by using a mainstream consumer social media platform. Grok further exposed Baltimore residents to the risk that any photograph they uploaded—of themselves or of their children—could be ingested by Grok and transformed into sexually degrading deepfakes without their knowledge or consent.

7. X represents to consumers that it prohibits non-consensual sexual content, bans the sexual exploitation of children, and permits only consensually produced adult nudity. x.AI likewise represents that Grok’s use policies forbid sexualized depictions of real people and minors.

8. In spite of these representations, X is now one of the largest distributors of NCII and CSAM.³

9. Baltimore residents have a reasonable expectation that they will not be exposed to this illegal content on X, and that X will not harass its own customers with Grok-generated deepfakes. Users have that expectation because Defendants market themselves and Grok as safe, rule-governed consumer platforms and products.

10. Those representations were false or misleading when made. Contrary to these assurances, Grok itself generated and distributed sexualized images of real individuals, exercising autonomous control over the content it produced rather than merely transmitting third-party material.

11. When confronted with public outcry over this conduct, Defendants did not eliminate the dangerous functionality. Instead, x.AI restricted certain image-generation features to

³ Amelia Gentleman & Helena Horton, *‘Add blood, forced smile’: how Grok’s nudification tool went viral*, The Guardian (Jan. 11, 2026), <https://www.theguardian.com/news/ng-interactive/2026/jan/11/how-grok-nudification-tool-went-viral-x-elon-musk>; Nick Robins-Early, *Elon Musk’s Grok AI generates images of ‘minors in minimal clothing’*, The Guardian (Jan. 2, 2026), <https://www.theguardian.com/technology/2026/jan/02/elon-musk-grok-ai-children-photos>.

paying users—placing the most controversial tools behind a paywall while leaving the underlying risks intact.⁴

12. Defendants’ purported “rules,” ostensibly curbing non-consensual sexual content, are deceptive misrepresentations—because Defendants, themselves, are violating them through Grok. And by marketing X and Grok—while both affirmatively misrepresenting and failing to disclose material facts about their limitations, risks, exploitative and defective design choices—Defendants are engaging in unfair and deceptive trade practices in violation of the Baltimore CPO, Baltimore City Code §§ 23-266, *et seq.*

13. Plaintiff Baltimore brings this action, pursuant to the City’s Consumer Protection Ordinance, Baltimore City Code Art. 2, § 4, to prevent and abate the serious health and public safety consequences posed by Defendants’ unlawful practices to the residents of the City of Baltimore. Plaintiff seeks all legal and equitable relief allowed by law, including civil penalties for each violation of the CPO, injunctive relief to prevent Defendants from continuing to engage in the unfair and deceptive trade practices described herein, restitution, and disgorgement.

II. JURISDICTION

14. This Court has personal jurisdiction over Defendants x.AI Corp. and x.AI LLC because they committed the acts complained of herein in the State of Maryland and the City of Baltimore. Defendants x.AI Corp. and x.AI LLC publicly represent, among other things, that they develop, operate, commercialize, and invest in the generative artificial intelligence system known as Grok, which is offered to consumers throughout the United States, including in Maryland and the City of Baltimore. x.AI Corp. and x.AI LLC purposefully directed Grok, its image-generation

⁴ Matt Burgess, *X Didn’t Fix Grok’s ‘Undressing’ Problem. It Just Makes People Pay for It*, WIRED (Jan. 9, 2026), <https://www.wired.com/story/x-didnt-fix-groks-undressing-problem-it-just-makes-people-pay-for-it/>.

tools, marketing, and related services to Maryland residents, including residents of Baltimore, and derived substantial commercial benefit from Grok's use in this jurisdiction. x.AI Corp. and x.AI LLC have therefore purposefully availed themselves of doing business in this jurisdiction.

15. This Court has personal jurisdiction over Defendant X Corp. because it committed the acts complained of herein in the State of Maryland and the City of Baltimore. X Corp. distributes and promotes Grok to its users, hosts Grok-generated content on its platform, and enabled Grok's image-generation and image-editing tools to operate within the United States, including in the State of Maryland and the City of Baltimore. X Corp. has therefore purposefully availed itself of doing business in this jurisdiction.

16. Defendants knowingly and intentionally directed Grok—including its image-generation and image-editing features—and their related representations into the State of Maryland and the City of Baltimore. Defendants did so with the knowledge and intent that Grok would generate and disseminate images—including sexually explicit and non-consensual imagery—that would enter and circulate within the flow of commerce and information in the State and City, and that their representations would be received and relied upon by Baltimore consumers. Defendants possessed further knowledge that the adverse effects of Grok—including the generation and dissemination of non-consensual sexual imagery—would be felt by the State of Maryland and in the City of Baltimore. Defendants have information related to the individuals who use their services within this jurisdiction by virtue of IP addresses and other geolocation tools, such that their services are offered and used within the boundaries of this jurisdiction.

17. By advertising and distributing Grok in the State of Maryland and the City of Baltimore, Defendants have purposefully availed themselves, through specific acts, of the privilege of conducting activities within the State and have conducted illegal acts within the State

of Maryland and the City of Baltimore. *See CSR, Ltd. v. Taylor*, 411 Md. 457, 485–86 (2009) (noting that to satisfy the purposeful availment requirement in Maryland, the defendant must have “engage[d] in significant activities in the State or create[d] continuing obligations with the State’s residents, thus taking advantage of the benefits and protections of Maryland law.” Examples of such availment include advertising in the forum state and registering to provide services to Maryland residents. (citing *Burger King Corp. v. Rudzewicz*, 471 U.S. 462, 475–76 (1985); *Asahi Metal Indus. Co., Ltd. v. Super. Ct. of Cal., Solano Cnty.*, 480 U.S. 102, 112 (1987); *Mohamed v. Michael*, 279 Md. 653, 659 (1977)).

18. Defendants are, therefore, subject to specific personal jurisdiction of the courts of this State. Md. Code Ann., Cts & Jud. Proc. § 6-103 (“A court may exercise personal jurisdiction over a person, who directly or by an agent: (1) Transacts any business or performs any character of work or service in the State; (2) Contracts to supply goods, food, services, or manufactured products in the State; (3) Causes tortious injury in the State by an act or omission in the State[.]”).

19. This Court’s exercise of personal jurisdiction over Defendants is consistent with due process because Defendants purposefully directed Grok, its image-generation tools, advertisements and marketing, and related services and acts into the State of Maryland and the City of Baltimore and otherwise availed themselves of the benefits and protections of the laws of the State of Maryland and the City of Baltimore.

20. This Court has subject matter jurisdiction because the claims at issue arise under a City of Baltimore ordinance. Md. Code Ann., Cts & Jud. Proc. § 1-501 (“The circuit courts are the highest common-law and equity courts of record exercising original jurisdiction within the State. Each has full common law and equity powers and jurisdiction in all civil. . . cases within its county[.]”).

III. VENUE

21. Venue is proper in this Court under Md. Code Ann., Cts & Jud. Proc. § 6-201, because a substantial part of the acts or omissions giving rise to the claims occurred in the City of Baltimore and Baltimore County. Md. Code Ann. § 6-201 (“[A] civil action shall be brought in a county where the defendant . . . carries on a regular business[.]”).

IV. PARTIES

22. Plaintiff Baltimore is a municipal corporation organized pursuant to Articles XI and XI-A of the Maryland Constitution and entrusted with all of the powers of local self-government and home rule afforded by those articles. Plaintiff is the corporate identity of the City of Baltimore, the most populous city in the State of Maryland.

23. Article I, Section 1 of the Charter of Baltimore City empowers Plaintiff to sue on behalf of the people of Baltimore. Balt. Charter, Art. I, § 1. Further, Plaintiff has the full authority to exercise the powers granted to it by the Constitution of Maryland or by any Public General or Public Local Laws of the State of Maryland.

24. Defendant Space Exploration Technologies Corp. (“SpaceX”) is a Texas corporation. SpaceX was founded in 2002 by Elon Musk, who at all relevant times has served as SpaceX’s Chief Executive Officer and has exercised substantial influence over the company’s strategy, governance, and product direction. On February 2, 2026, SpaceX’s acquisition of x.AI was announced.

25. Defendant x.AI Corp. is a Nevada corporation with a principal place of business in Palo Alto, California. x.AI Corp. was founded in March 2023 by Elon Musk, who at all relevant times has served as x.AI Corp.’s Chief Executive Officer and has exercised substantial influence over the company’s strategy, governance, and product direction. At all relevant times, x.AI Corp.

has been engaged in the development, operation, management, and commercialization of Grok and related services.

26. Defendant x.AI LLC is a Nevada limited liability company and a wholly owned subsidiary of x.AI Corp. At all relevant times, x.AI LLC has been engaged in, including without limitation through its role within the x.AI corporate enterprise founded and controlled by Elon Musk, the development, operation, management, and commercialization of Grok and related services.

27. Defendants x.AI Corp. and x.AI LLC together own control, and operate Grok, including its image-generation and image-editing features, and at all relevant times acted jointly and in concert with respect to Grok's design, deployment, marketing, and commercialization. These actions were undertaken within the corporate structure founded and directed by Elon Musk and coordinated with X Corp., another Musk-controlled entity.

28. Defendant X Corp. is a Nevada corporation, with "a robust office footprint across the United States[.]"⁵ Elon Musk acquired Twitter for approximately \$44 billion in late 2022, and in July 2023, he announced that he was renaming the company "X."⁶ On March 28, 2025, x.AI Corp. acquired X Corp. (including the X platform) in an all-stock transaction valued at approximately \$33 billion.⁷ X Corp. is a wholly owned subsidiary of x.AI Corp. and owns and

⁵ *X Headquarters & Office Locations*, X, <https://websets.exa.ai/websets/directory/x-offices> (last visited Feb. 2, 2026).

⁶ Jyoti Mann, Michelle Mark, & Lauren Edmonds, *Here's everything that's happened to Twitter since Elon Musk bought it and renamed it X*, Bus. Insider, <https://www.businessinsider.com/elon-musk-twitter#september-2023-9> (last updated Sept. 7, 2025).

⁷ Antonio Pequeño, *Elon Musk Says xAI Has Purchased X, Formerly Known As Twitter, For \$33 Billion*, Forbes (Mar. 28, 2025), <https://www.forbes.com/sites/antoniopequenoiv/2025/03/28/elon-musk-says-xai-has-purchased-x-formerly-known-as-twitter-for-33-billion/>.

operates the social media platform known as X (formerly Twitter). At all relevant times, Elon Musk has maintained a controlling interest in both X Corp. and SpaceX.

29. X Corp. distributes and promotes Grok to its users, hosts Grok-generated content on its platform, and provides the technical infrastructure through which Grok’s image-generation and image-editing features are made available to consumers. At all relevant times, X Corp. has been engaged in—including without limitation through its role within the x.AI corporate enterprise—the development, operation, management, and commercialization of Grok and related services, acting jointly and in concert with x.AI Corp. and x.AI LLC in connection with the deployment, promotion, and monetization of Grok.

V. FACTUAL BACKGROUND

A. Grok and Grok Imagine

30. Generative artificial intelligence (“GAI”) products have rapidly entered consumer markets with minimal regulation and limited transparency. Of these, “nudify” apps are a growing threat to consumer welfare.

31. These GAI products take images of actual people—including children—and transform them into photo-realistic, sexualized “deepfakes” without the victim’s consent. x.AI’s Grok chatbot is one of the largest “nudify” apps.

32. Grok is a generative artificial intelligence system developed by x.AI, launched in November 2023, and distributed to consumers through the X (formerly Twitter) platform and related applications.

33. On August 4, 2025, Grok’s image generation feature, Grok Imagine, was launched. One of Grok’s most controversial offerings, Grok Imagine, includes a “Spicy” mode that allows users to generate sexually explicit images and videos.⁸

34. Unlike traditional consumer software tools, Grok is designed to generate novel content in response to user prompts, drawing on vast quantities of training data that include scraped online material, images, and text. These services are delivered through two main channels: the X social media platform and the standalone x.AI/Grok app.⁹

35. Both channels permit users to ask Grok to “undress” or “nudify” photos of third parties—ranging from celebrities to private citizens, including children—placing them in sexually suggestive, degrading, or violent scenarios.¹⁰

36. Users may access Grok’s image tools either via X or through the Grok app. On X, users can generate images either with the “@Grok” chatbot or with the “edit image” button.

37. With the “@Grok” chatbot, X users can interact with Grok via a dedicated tab or by tagging the @Grok account in any post. Grok can generate up to four high-resolution images in roughly 3-5 seconds from a text prompt.

⁸ Sara Pequeño, *Grok’s ‘spicy mode’ making pictures with kids is horrifying*, *Opinion*, USA Today (Jan. 12, 2026), <https://www.msn.com/en-us/news/opinion/groks-spicy-mode-making-pictures-with-kids-is-horrifying-opinion/>.

⁹ Hayden Field, *Grok is undressing children—can the law stop it?*, *The Verge* (Jan. 6, 2026), <https://www.theverge.com/ai-artificial-intelligence/855832/grok-undressing-children-csam-law-x-elon-musk>.

¹⁰ Anna Desmarais, *Elon Musk’s Grok still being used to generate explicit images despite new safeguards, study finds*, *Euronews* (Jan. 21, 2026), <https://www.msn.com/en-us/money/other/elon-musk-s-grok-still-being-used-to-generate-explicit-images-despite-new-safeguards-study-finds/>; Kim LaCapria, *New report finds Elon Musk’s Grok generated thousands of disturbing images per hour on social media: ‘This is not spicy. This is illegal.’*, *Yahoo News* (Jan. 14, 2026), <https://www.yahoo.com/news/articles/report-finds-elon-musks-grok-233000344.html>.

38. A key driver of NCII dissemination is X’s “edit image” button, which allows users to manipulate existing posts or upload a photo and ask Grok to alter it. For example, users can request edits of photos depicting third parties, including minors, and ask Grok to place them in a “transparent bikini,” producing sexually suggestive images and at least partial nudity.

39. x.AI designed Grok to copy the identifiable features of women and children and to falsely present them as subjects of pornographic imagery.

40. Users can ask Grok to make obscene and offensive modifications, such as placing a “donut glaze” on a child’s face that resembles semen.¹¹

41. One victim described how Grok manipulated an image she had posted on X. At users’ request, Grok non-consensually “undressed her” and eventually generated images of her “completely naked with just a bit of string around [her] waist” and “with a ball gag in [her] mouth and [her] eyes rolled back.”¹²

42. The dedicated Grok app can generate even more explicit imagery than the X @Grok chatbot and includes advanced video capabilities. Users can ask Grok to generate images through text prompts, upload and manipulate images, and animate uploaded or Grok-generated images into a 6-15 second video clip, with AI-generated audio.

43. While the @Grok chatbot on X facilitates public-facing abuse, the standalone x.AI app and website allows users to generate photorealistic NCII or CSAM content that is “vastly more explicit” than those seen on X, according to a *WIRED* investigation.¹³ This imagery generated by AI tools like Grok is known to be often downloaded and redistributed across external forums,

¹¹ Field, *supra* n.9.

¹² Gentleman & Horton, *supra* n.3.

¹³ Matt Burgess & Maddy Varner, *Grok Is Pushing AI ‘Undressing’ Mainstream*, *WIRED* (Jan. 6, 2026), <https://www.wired.com/story/grok-is-pushing-ai-undressing-mainstream/>.

channels, and websites after its creation, creating a secondary market for Grok-facilitated sexual exploitation.¹⁴

44. With only limited prompts, Grok autonomously produces sexualized imagery of individuals, reflecting design choices embedded in its model rather than detailed user instruction, and demonstrating that the model is actively responsible for shaping and rendering the final image. As a result, Grok itself functions as the creator of these images, not merely a passive tool.

i. Elon Musk’s Promotion of Grok’s “Undressing” Capability

45. The output of NCII appears to have increased after Defendants made this service available for free to all users on or about August 7, 2025, and then exploded with the release of Grok 4.1 on or about November 17, 2025, which provided expanded image generation and editing abilities.

46. Users then began exploiting Grok’s image-editing features to generate edited images placing people—often women and girls—in bikinis or other revealing contexts.

47. By early January 2026, a trend—often referred to by users as the “put her in a bikini” trend—had spread widely across the X platform as users repeatedly tagged Grok in public threads.

¹⁴ See Alejandro Cuevas & Manoel Horta Ribeiro, *Deepfake Pornography is Resilient to Regulatory and Platform Shocks*, Princeton University, (Feb. 2, 2026), <https://doi.org/10.48550/arXiv.2602.02754> (“Generative artificial intelligence tools have made it easier to create realistic, synthetic non-consensual explicit imagery (popularly known as deepfake pornography; hereinafter SNCEI) of people. Once created, this SNCEI is often shared on various websites, causing significant harm to victims.”); Rebecca Umbach et al., *Non-Consensual Synthetic Intimate Imagery: Prevalence, Attitudes, and Knowledge in 10 Countries*, In Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI ’24), Assoc. for Computing Mach. (May 11, 2024), <https://doi.org/10.1145/3613904.3642382> (“[A]s content may be distributed on dedicated sites or within closed communities on sites such as Discord or Reddit, victims of AI-ISBA [deepfake pornography] may not know about or find out about their images.”).

48. As the trend spread, Elon Musk publicly participated in it using his own X account. On December 31, 2025, Musk shared a photo generated by Grok of himself in a bikini, as well as a SpaceX rocket with a woman's undressed body on top.

49. After Grok produced the image of himself in a bikini, Musk publicly responded to the output by commenting "Perfect." Musk then further amplified the episode through additional engagement on X, including reacting to users' commentary with laughing emojis. The post was visible to Musk's tens of millions of followers and was rapidly disseminated across the platform:



50. Musk's post functioned as public endorsement of Grok's ability to generate sexualized or revealing edits of real people, and it signaled to users that these uses of Grok were acceptable, humorous, and encouraged. Coming from the owner and principal public face of both

x.AI and X, Musk's post operated as marketing and promotion for the very image-editing capability that was being used to generate non-consensual sexual imagery.

51. The timing and scale of the resulting surge in use were dramatic. Reporting found that, in the nine days preceding Musk's post, Grok was used approximately 300,000 times to generate images; after Musk's post, the number of images created by Grok surged to nearly 600,000 per day on X.¹⁵

52. In just eleven days after Musk's post, Grok sexualized more than 3,000,000 images, including approximately 23,000 images depicting children, according to estimates published by the Center for Countering Digital Hate ("CCDH").¹⁶

53. Independent analysis by the New York Times conservatively estimated that approximately 1.8 million Grok-generated images produced during this same period sexualized real people, including women, men, and children.

54. Defendants' response to the widespread public outcry over Grok's generation of non-consensual sexual imagery was not to immediately disable the underlying image-editing functionality or to meaningfully reform Grok's design. Instead, in early January 2026, Defendants restricted Grok's image-generation and image-editing capabilities on X to paying users of "X Premium"—thereby placing core functionality behind a paywall while the scandal was still unfolding.

55. Beginning on approximately January 8, 2026, when non-paying users attempted to prompt Grok to generate or edit images placing women into sexualized contexts or revealing

¹⁵ Kate Conger, Dylan Freedman & Stuart Thompson, *Musk's Chatbot Flooded X with Millions of Sexualized Images in Days, New Estimates Show*, N.Y. Times (Jan. 22, 2026), <https://www.nytimes.com/2026/01/22/technology/grok-x-ai-elon-musk-deepfakes.html>.

¹⁶ Ctr. for Countering Digital Hate, *supra* n.2.

clothing, the @Grok reply bot began responding with messages stating, in substance: “Image generation and editing are currently limited to paying subscribers. You can subscribe to unlock these features,” accompanied by a link directing users to purchase an X Premium subscription.

56. Prior to this change, Grok’s image-generation and image-editing tools—including the features used to create sexualized and “undressed” images of real people—had been broadly available to free users on X. Defendants’ decision to restrict these tools to paying users did not eliminate the functionality; it instead conditioned access on payment, converting a widely abused feature into a premium offering, for Defendants’ profit.

57. Under X’s subscription structure, paid users receive expanded access to Grok, including higher usage limits and enhanced functionality, in exchange for a monthly fee. By restricting image generation and editing to paid accounts, Defendants increased the commercial value of X Premium subscriptions while public attention to Grok’s controversial outputs was driving heightened engagement with the platform.

58. As reported at the time, X executives publicly celebrated record engagement levels on the platform in early January 2026, even as journalists and advocates documented widespread harm caused by Grok-generated non-consensual sexual imagery.

59. Notably, Defendants’ restrictions applied only to Grok’s image-editing features within certain contexts on X. The same or similar capabilities remained available through the standalone Grok app and website, allowing users to continue generating non-consensual sexual imagery outside the X interface. Defendants therefore did not meaningfully prevent the ongoing creation of harmful content, but instead redirected and monetized access to Grok’s most dangerous features.

60. By placing exploitative image-generation capabilities behind a paywall—after Defendants were on notice that those capabilities were being used at massive scale to create non-consensual sexual imagery, including images of children—Defendants prioritized revenue generation and user engagement over consumer protection, public safety, and compliance with their own stated rules and policies.

B. Defendants’ Deceptive Representations

61. Defendants each engaged in deceptive acts and practices by affirmatively misrepresenting the safety, limitations, and governance of Grok and the X platform, while omitting material facts about Grok’s design, functionality, and foreseeable risks. These misrepresentations were conveyed to Baltimore residents through Defendants’ marketing statements, public-facing policies, and terms of service, and they were material to consumers’ decisions to use X and Grok and to upload images of themselves and others to those platforms.

i. Marketing

62. Defendants represent that Grok is built on a “maximal truth seeking” model, and promote Grok as a responsible, general-purpose artificial intelligence assistant.¹⁷ These representations reasonably convey to consumers that Grok operates subject to meaningful internal controls and safeguards designed to prevent harmful, abusive, or non-consensual outputs.

63. At the same time, Defendants failed to disclose that Grok’s purported safeguards against non-consensual sexual imagery, sexual exploitation, and other harmful content were weak, inconsistently enforced, and easily circumvented.

¹⁷ See, e.g., Grok, *Introduction: The Dawn of AI and Its Dark Shadow*, Grok Thinks (Apr. 14, 2025), <https://grokthinks.substack.com/p/the-ai-battle-for-truth>.

64. Defendants further failed to disclose that Grok’s image-generation and image-editing features were designed in a manner that affirmatively enabled the creation of sexually explicit and degrading imagery of real people, including minors, with minimal prompting.

65. Despite possessing its dangerous capabilities, Defendants marketed Grok as a general-purpose consumer product suitable for widespread use, without meaningful age verification, without adequate disclosures regarding the risks of image manipulation and sexual exploitation, and without warnings proportionate to the foreseeable harms created by Grok’s design.¹⁸ These omissions rendered Defendants’ marketing representations misleading.

ii. Policies

66. Grok directly violates Defendants’ own publicly available terms of use and policies, which purport to prohibit the very categories of content that Grok generates and disseminates. Nevertheless, Defendants continued to publish, enforce, and rely on those policies as representations of platform safety, content governance, and consumer protection, even as Grok’s design and deployment made widespread violations foreseeable and routine.

67. Defendants’ terms of service, rules, and acceptable-use policies are presented to users as binding conditions governing the use of X and Grok, and as assurances that Defendants maintain meaningful safeguards against harmful, abusive, and exploitative content. These policies are prominently incorporated into the user experience, referenced in marketing materials, and

¹⁸ *xAI’s Grok 4 has no meaningful safety guardrails*, Lesswrong (Jul. 13, 2025), <https://www.lesswrong.com/posts/dqd54wpEjfKJsJBk6/xai-s-grok-4-has-no-meaningful-safety-guardrails>; Kitty Wheeler, *The Story Behind Elon Musk’s xAI Grok 4 Ethical Concerns*, AI.Magazine (Jul. 17, 2025), <https://aimagazine.com/news/the-story-behind-elon-musks-xai-grok-4-ethical-concerns>; *Grok 4 Bias Concerns and Safety Risks: Examining Elon Musk’s AI Ethics and Misinformation Issues*, AiToolApp, <https://aitoolapp.com/grok-4/concerns/> (last visited Jan. 22, 2026).

invoked in public statements to regulators, advertisers, and the public as evidence of Defendants’ commitment to safety and compliance.

68. X Corp. has made false and misleading statements in the terms of service and platform rules that it promulgates to all users, including users within the City of Baltimore, by representing that certain categories of harmful content are prohibited or subject to zero tolerance, while Grok itself generated, promoted, and distributed that same content on X.

69. X Corp.’s rules advertise that X Corp. has “zero tolerance for any forms of child sexual exploitation” and “remove[s] certain media depicting physical child abuse to prevent the normalization of violence against children.”¹⁹

70. X Corp.’s rules also state that X Corp. only permits sharing “consensually produced and distributed adult nudity or sexual behavior.”²⁰ The rules specifically forbid “post[ing] or shar[ing] intimate photos or videos of someone that were produced or distributed without their consent.”²¹ They even ban “deceptively shar[ing] synthetic or manipulated media that are likely to cause harm.”

71. X Corp.’s rules further state that the platform prohibits content “explicitly threatening, inciting, glorifying, or expressing desire for violence[,]” abusive content, or “engag[ing] in the targeted harassment of someone.”²²

72. Contrary to X Corp.’s representations within its terms of service, X Corp. does not, in practice, prohibit content that glorifies violence, does not maintain zero tolerance for child

¹⁹ The X Rules, <https://web.archive.org/web/20250101050103/https://help.x.com/en/rules-and-policies/x-rules> (version dated Jan. 1, 2025).

²⁰ *Id.*

²¹ *Id.*

²² *Id.*

sexual exploitation, or the posting and sharing of abusive and harassing content, or non-consensual sexual material when that material is generated or facilitated through Grok.

73. On January 31, 2024, X Corp.'s CEO, Linda Yaccarino, reinforced these policies through written testimony submitted to the United States Senate Committee on the Judiciary, to be considered in the Hearing on Big Tech and the Online Child Sexual Exploitation Crisis, emphasizing:

- a. **“We have a zero-tolerance child sexual exploitation policy on X. X has zero tolerance towards any material that features or promotes child sexual exploitation This may include . . . illustrated, or computer-generated images.”**
- b. **“Non-Consensual Nudity Policy[:] You may not post or share intimate photos or videos of someone that were produced or distributed without their consent Examples of the types of content that violate this policy include, but are not limited to: . . . images or videos that superimpose or otherwise digitally manipulate an individual’s face onto another person’s nude body[.]”**
- c. “X has zero tolerance for Child Sexual Exploitation (CSE), and we are determined to make X inhospitable for actors who seek to exploit minors. In 2023, we made clear that our top priority was tackling CSE online.”
- d. “[X] has strengthened its policies and enforcement to tackle CSE. We are taking aggressive action against users that distribute CSE and the networks of users who engage with this horrible content.”

- e. “[W]e have made it more difficult for bad actors to share or engage with CSE material on X, while simultaneously making it easier for our users to report CSE content.”²³

74. x.AI has made similarly deceptive representations, while omitting material facts, to the Baltimore public, as Grok violates x.AI’s own Acceptable Use policy, in effect since January 2, 2025.

75. x.AI’s Acceptable Use policy prohibits, among other things, (1) “[v]iolating a person’s privacy or their right of publicity, (2) “[d]epicting likenesses of persons in a pornographic manner, and (3) “[t]he sexualization of exploitation of children.”²⁴

76. Defendants’ public-facing marketing for Grok, and specifically its Spicy Mode, expressly represents that “[n]on-consensual sexual content (Deepfakes)” and “[s]exualization of real people, celebrities, or public figures” are “[s]trictly [p]rohibited,” and that Spicy Mode “does not support uploading photos of real people,” and that these restrictions exist “to prevent the creation of deepfakes and non-consensual sexual content.”²⁵ Defendants further represent that Grok Spicy Mode, while “less filtered,” nevertheless remains subject to “core safety policies.” These representations were conveyed to consumers as assurances that Defendants had

²³ Written Testimony of Linda Yaccarino, CEO, X Corp., to United States Senate Committee on the Judiciary, Hearing on Big Tech and the Online Child Sexual Exploitation Crisis (Jan. 31, 2024), https://www.judiciary.senate.gov/imo/media/doc/2024-01-31_-_testimony_-_yaccarino1.pdf (emphases in original).

²⁴ xAI Acceptable Use Policy, <https://x.ai/legal/acceptable-use-policy> (version dated Jan. 2, 2025).

²⁵ Grok Spicy Mode: What it is, How to Enable, and Safe Usage, <https://web.archive.org/web/20260113170315/https://grok-imagine.art/grok-spicy-mode> (version dated Jan. 13, 2026).

implemented meaningful technical and policy safeguards to prevent the creation of non-consensual sexual imagery.

77. In reality, and as alleged herein, Grok’s design and deployment enabled exactly the conduct Defendants claimed to prohibit, including the creation and dissemination of non-consensual sexual imagery and deepfakes of real people, including minors. Defendants’ representations regarding deepfake prevention and Spicy Mode restrictions were therefore false or misleading at the time they were made.

78. As of January 30, 2026, Defendants’ publicly available rules, terms of service, and policies—including the representations described above—remained in effect and were presented to users as governing content on X and Grok.

79. By continuing to publish and rely on these policies as assurances of safety and governance, Defendants misled consumers into believing that Grok was subject to meaningful constraints against harmful and exploitative content, even as Grok’s operation routinely and foreseeably violated those same constraints in systemic ways.

iii. Safeguards

80. On January 2, 2026, through tweets posted on Grok’s public account, Defendants admitted to “lapses in safeguards” after images of “minors in minimal clothing” circulated on X.²⁶ Grok posted that these were “isolated cases” and that “improvements are ongoing to block such requests entirely.”²⁷

²⁶ Mary Cunningham, *Grok chatbot allowed users to create digitally altered photos of minors in “minimal clothing,”* CBS News (Jan. 2, 2026), <https://www.cbsnews.com/news/grok-safeguard-lapses-minors-minimal-clothing-ai/>.

²⁷ Grok Post, X (Jan. 1, 2026), <https://x.com/grok/status/2006618883055046758> (“I’ve reviewed recent interactions. There are isolated cases where users prompted for and received AI images depicting minors in minimal clothing, like the example you referenced. xAI has safeguards, but improvements are ongoing to block such requests entirely. (214 chars).”).

81. Despite that admission, Defendants did not disable Grok’s underlying capability to generate or edit sexualized images of real people.

82. Instead, Defendants’ principal response was to restrict image generation and editing when Grok is used directly on X, limiting those functions to paid, verified X users.²⁸ In other words, Defendants’ “solution” was simply to monetize the disturbing content it had previously permitted users to create for free, without eliminating the underlying conduct.

83. While x.AI imposed additional restrictions on Grok’s image-generation features when accessed through X, the same or substantially similar functionality remained available through Grok’s standalone app and website.

84. As a result, users without paid X subscriptions could still generate and edit images—including sexualized image edits of real people—by accessing Grok outside the X platform.

85. Journalists, watchdogs, and users have reported that Grok is still generating NCII content in response to more sophisticated user prompts.

86. For example, according to one January 14, 2026, report, users on chat forums such as Reddit have revealed that the new safeguards can be circumvented through creative prompting (e.g., asking for “artistic nudity”).²⁹

²⁸ Burgess, *supra* n.4.

²⁹ Helena Horton, Aisha Down, & Priya Bharadia, *Use of AI to harm women has only just begun, experts warn*, The Guardian (Jan. 14, 2026), <https://www.theguardian.com/technology/2026/jan/14/use-of-ai-to-harm-women-has-only-just-begun-experts-warn>.

87. On January 15, WIRED also reported that Grok’s new patchwork of safeguards are ineffective,³⁰ and *The Washington Post* reported that the Grok app was still undressing images.³¹

88. Independent researchers have likewise raised concerns about Grok’s lack of effective safeguards. SpeechMap.AI identified Grok as the most permissive language model it tested, while the AI safety nonprofit SaferAI ranked x.AI last among peer companies in the maturity of its AI risk management practices.³²

89. The risks associated with AI-enabled sexual deepfakes are well known within the artificial intelligence industry. Numerous AI developers have acknowledged these dangers and implemented technical and policy-based guardrails to prevent their tools from being used to generate sexualized or exploitative images of non-consenting individuals.

90. Defendants, by contrast, chose a different course. Rather than adopting widely recognized safeguards to mitigate foreseeable harm, Defendants designed, deployed, and marketed Grok in a manner that permitted—and failed to meaningfully prevent—the generation of sexualized images of real people, while deceiving the public, including Baltimore residents.

C. Defendants’ Unfair Practices

91. Independently of deception, Defendants, through Grok, have engaged in unfair practices under Baltimore law.

³⁰ Matt Burgess, *Elon Musk’s Grok ‘Undressing’ Problem Isn’t Fixed*, WIRED (Jan. 15, 2026), <https://www.wired.com/story/elon-musks-grok-undressing-problem-isnt-fixed/>.

³¹ *Stand-alone Grok app still undresses women after X curtails access to tool*, Wash. Post (Jan. 15, 2026), <https://www.washingtonpost.com/technology/2026/01/15/grok-ai-image-generator-sexualized/>.

³² See Jeff Ramage, *Everything You Need to Know About xAI, the Company Behind Grok, Built In* (Jan. 15, 2026), <https://builtin.com/artificial-intelligence/what-is-xai>.

92. Defendants exploit consumers' lack of understanding of generative AI by deploying a powerful content-generation system without meaningful consent, comprehension, or control.

93. Consumers cannot reasonably avoid the resulting harms once Grok is in the marketplace, because Grok can generate harmful content involving them without notice or participation.

94. Grok operates as a mass-scale generator of harmful content with minimal accountability, as Defendants deceive the public about their policies and the safeguards in effect.

95. Grok's design facilitates foreseeable harms—non-consensual sexual imagery, harassment, and exposure of minors—while Defendants reap significant profits.

D. Consumer Harm

96. Defendants' conduct has caused widespread and foreseeable harm to consumers, including residents of the City of Baltimore. Through Grok's design, deployment, and promotion, Defendants facilitated the mass creation and dissemination of non-consensual sexual imagery involving real people, exposing Baltimore residents to degrading content, subjecting them to the risk of being targeted themselves, and normalizing a form of image-based sexual abuse that is difficult to prevent, contain, or remedy once unleashed at scale.

97. Unlike many AI tools that generate content viewable only by the requesting user, Grok routinely posts generated images to X's public media feed by default. As a result, those images are instantly exposed to other users, easily discoverable and shareable, and may spread widely and rapidly without the depicted individual's ability to meaningfully control their distribution.

98. In one example from June 2025, an X user asked Grok to add what appears to be semen to the face of a Glamour reporter in a photo she had uploaded.³³

99. On August 5, 2025, a reporter from the online magazine The Verge documented that Grok Imagine generated “fully topless videos of Taylor Swift” without the reporter “even asking the bot to take her clothes off.”³⁴

100. According to a January 7, 2026, report from Bloomberg, one analysis found that Grok was generating approximately 6,700 sexually explicit or “undressing” images per hour and that sexualized content comprised 85% of Grok’s total output.³⁵

101. A January 22, 2026, investigation by the New York Times concluded that Grok created and publicly shared at least 1.8 million sexualized images of women between December 31, 2025, and January 8, 2026.³⁶

102. A separate study by the Center for Countering Digital Hate estimated that during that same period Grok sexualized more than 3,000,000 images, 23,000 of which were of children.³⁷

103. As of February 2025, the U.S. had 103.96 million X users.³⁸ Adjusting for population size, it is likely that between approximately 175,000 and 180,000 Baltimore City

³³ Lucy Morgan, *X’s ‘Grok’ created an AI sexualised image of me without my consent*, Glamour (June 12, 2025), <https://www.glamourmagazine.co.uk/article/x-ai-chatbot-grok>.

³⁴ Jess Weatherbed, *Grok’s ‘Spicy’ Video Setting Instantly Made Me Taylor Swift Nude Deepfakes*, The Verge (Aug. 5, 2025), <https://www.theverge.com/report/718975/xai-grok-imagine-taylor-swifty-deepfake-nudes>.

³⁵ *Musk’s Grok AI Generated Thousands of Undressed Images Per Hour on X*, Bloomberg (Jan. 7, 2026), <https://www.bloomberg.com/news/articles/2026-01-07/musk-s-grok-ai-generated-thousands-of-undressed-images-per-hour-on-x/>.

³⁶ Conger, Freedman & Thompson, *supra* n.15.

³⁷ Ctr. for Countering Digit. Hate, *supra* n.2.

³⁸ *See 21 Essential Twitter (X) Statistics You Need to Know in 2026*, Soc. Shepherd (Jan. 6, 2026), <https://thesocialshepherd.com/blog/twitter-statistics>.

residents maintained X accounts and were therefore directly exposed to Grok-generated content or to the risk that images of themselves or loved ones could be manipulated and sexualized through Grok.³⁹

104. This form of image-based sexual abuse has profound and lasting harm to those targeted. Victims suffer violations of their bodily autonomy, dignity, privacy, and control over their own sexual expression, as well as serious emotional and psychological injuries.

105. Exposure to such abuse commonly results in fear, distress, and anxiety, particularly because AI-generated sexual deepfakes, once created, can be endlessly duplicated, shared, and stored across private devices and online platforms, making them extraordinarily difficult—if not impossible—to fully locate, remove, or contain.

106. This type of abuse also inflicts collective harms on the public. The widespread creation and circulation of non-consensual sexual imagery erodes social norms around consent and privacy, fostering a culture that tolerates or even trivializes sexual exploitation. This normalization places the public at large at risk, and disproportionately endangers women and girls, who are far more likely to be targeted for this form of abuse.

107. Plaintiff's claims arise from Defendants' own deceptive and unfair commercial conduct—specifically, their design, development, marketing, and monetization of Grok's image-generation and image-editing features, and their false or misleading representations regarding the safety, safeguards, and limitations of those features.

³⁹ Using a Baltimore population size of approximately 565,000 and a U.S. population size of approximately 340 million. See U.S. Census Bureau, *QuickFacts United States; Baltimore city, Maryland*, at Population estimates, July 1, 2024 (V2024), <https://www.census.gov/quickfacts/fact/table/US,baltimorecitymaryland/PST045224> (last visited Feb. 20, 2026) (estimating Baltimore City's population at 568,271 and the United States population at 340,110,988 as of July 1, 2024).

VI. CLAIM FOR RELIEF

Unfair, Abusive, And Deceptive Trade Practices Baltimore City Code Art. 2, § 4

108. The City of Baltimore reasserts, realleges, and incorporates by reference each of Paragraphs 1–107, above, as though fully set forth herein.

109. The CPO, Baltimore City Code Art. 2, § 4, protects consumers and others against “unfair, abusive, or deceptive trade practices,” which are defined in line with the Maryland Consumer Protection Act (“MCPA”), Md. Code Ann., Com. Law, § 13-301. *See* Balt. City Code Art. 2, § 4-1 (13).⁴⁰

110. These “[u]nfair, abusive, or deceptive trade practices include” any:

- a. False, falsely disparaging, or misleading oral or written statement, visual description, or other representation of any kind which has the capacity, tendency, or effect of deceiving or misleading consumers, Md. Code Ann., Com. Law, § 13-301(1);
- b. Failure to state a material fact if the failure deceives or tends to deceive, Md. Code Ann., Com. Law, § 13-301(3);
- c. Deception, fraud, false pretense, false premise, misrepresentation, or knowing concealment, suppression, or omission of any material fact with the intent that a consumer rely on the same in connection with:
 - (i) The promotion or sale of any consumer goods, consumer realty, or consumer service, Md. Code Ann., Com. Law, § 13-301(9)(i).

⁴⁰ *See also* Law 23-0424 (*Comments from Baltimore Law Department and Chief Solicitor*), City Council of Balt., <https://baltimore.legistar.com/LegislationDetail.aspx?ID=6322280&GUID=04376B1E-3696-45E8-8DC8-8457391147DE&Options=&Search=&FullText=1> (last visited Feb. 2, 2026).

111. Any “practice prohibited by this title is a violation of this title, whether or not any consumer in fact has been misled, deceived, or damaged as a result of that practice.” *Id.* § 13-302. While specific practices are enumerated in the MCPA (and, by extension, the CPO), “[i]t is the intent of the [Maryland] General Assembly that in construing the term ‘unfair or deceptive trade practices,’ due consideration and weight be given to the interpretations of § 5 (a)(1) of the Federal Trade Commission Act [(15 U.S.C. § 45)] by the Federal Trade Commission and the federal courts.” Md. Code Ann., Com. Law § 13-105. The Federal Trade Commission has explained that unfairness under 15 U.S.C. § 45 is determined in part by a consideration of “(1) whether the practice injures consumers,” and “(2) whether it violates established public policy.”⁴¹

112. Defendants are “merchants” within the meaning of, and subject to, the provisions of the CPO. Balt. City Code, Art. 2, § 4-1 (9).

113. Defendants’ actions are unfair, abusive, and deceptive. Without limitation, Defendants have violated and continue to violate the CPO by:

- a. Designing Grok’s image-editing tools—including the “edit image” feature and related functionality—to permit the sexualization, “nudification,” and manipulation of images depicting real people, including minors, and making those tools broadly available to consumers without adequate safeguards or disclosures;
- b. Designing, deploying, and marketing Grok and its image-generation and image-editing features in a manner that foreseeably exposes children and adolescents to sexualized content, including non-consensual sexual imagery and content

⁴¹ *FTC Policy Statement on Unfairness*, FTC (Dec. 17, 1980), <https://www.ftc.gov/legal-library/browse/ftc-policy-statement-unfairness> (last visited Feb. 2, 2026).

constituting or resembling child sexual abuse material, while failing to implement meaningful age-based access restrictions or protections;

- c. Deceptively representing, through marketing materials, public statements, platform rules, and acceptable-use policies, that Grok and the X platform prohibit non-consensual sexual imagery, the sexual exploitation of children, and harmful manipulated media, while Grok itself generated and distributed precisely such content;
- d. Failing to disclose to Baltimore residents material facts concerning Grok’s design, including the product’s propensity to autonomously generate sexualized imagery of real individuals with minimal prompting, the known inadequacy of safeguards, and the foreseeable risk that uploaded images could be transformed into sexually degrading deepfakes;
- e. Employing exploitative product design choices that encouraged the creation and dissemination of humiliating and degrading sexual imagery—including by promoting “Spicy” modes, image-editing tools, and rapid-generation features—while externalizing the resulting harms to victims and the public;
- f. Misrepresenting the nature, scope, and effectiveness of Grok’s safety measures by announcing post-hoc “guardrails” that did not eliminate the harmful functionality and could be readily circumvented, thereby creating a false impression that consumers were meaningfully protected;
- g. Monetizing Grok’s most dangerous image-generation capabilities by restricting certain features to paying subscribers, increasing image-generation limits for paid

users, and placing exploitative functionality behind a paywall after Defendants were on notice of widespread abuse; and

- h. Failing to implement effective, industry-standard measures to prevent the generation and dissemination of non-consensual sexual imagery and sexualized images of minors, despite possessing the technical capacity to do so and despite widespread industry recognition of the need for such safeguards.

114. Defendants' unfair, abusive, and deceptive practices and caused—and continue to cause—substantial injury to Baltimore consumers that is not reasonably avoidable by consumers themselves.

115. Defendants' actions are against public policy. Public law and policy favor consumer protection, privacy, and child safety, and it shocks the conscience in the same manner as other exploitative digital business models previously targeted by consumer protection authorities.

116. Defendants' actions demonstrate a callous disregard not only for the rule of law, but also for the public health, safety, and well-being of Baltimore's residents. While engaging in the unlawful practices alleged herein, Defendants have, at all times, acted willfully. Defendants knew or should have known that their actions were of a nature prohibited by the CPO.

117. As a result of the foregoing, the City seeks all legal and equitable relief as allowed by law, including civil penalties, injunctive relief, restitution, and disgorgement.

VII. REQUEST FOR RELIEF

WHEREFORE, Plaintiff, the City of Baltimore, respectfully requests that the Court enter judgment in its favor and against Defendants, as follows:

- a. Awarding the maximum amount of statutory penalties available under Baltimore City Code Art. 2, § 4-3(a), for each violation of Baltimore’s CPO, Balt. City Code Art. 2, § 4;
- b. Granting injunctive relief, mandating that Defendants cease the targeting and exploitation of Baltimore’s residents;
- c. Granting injunctive relief requiring Defendants to reform their exploitative platform design feature restrictions and enhanced marketing restrictions, Balt. City Code Art. 2, § 4-5(d); and
- d. Awarding such other relief as may be available and appropriate under the law or in equity.

Dated: March 24, 2026

Respectfully submitted,

**THE MAYOR AND CITY COUNCIL OF
BALTIMORE, MARYLAND**

Ebony M. Thompson
Baltimore City Solicitor

_____/s/_____
Sara Gross (AIS 0412140305)
Chief Solicitor
Zachary Babo (AIS 2211280023)
Assistant Solicitor
Baltimore City Department of Law
100 North Holliday Street
Baltimore, Maryland 21202
Tel.: 410-396-3947
sara.gross@baltimorecity.gov

Adam J. Levitt*
Amy E. Keller*
Daniel R. Ferri*

Rebecca Trickey*
DICELLO LEVITT LLP
Ten North Dearborn Street, Sixth Floor
Chicago, Illinois 60602
Tel.: (312) 214-7900
alevitt@dicellolevitt.com
akeller@dicellolevitt.com
dferris@dicellolevitt.com
rtrickey@dicellolevitt.com

Corban S. Rhodes*
Emma L. Bruder*
DICELLO LEVITT LLP
485 Lexington Avenue, Tenth Floor
New York, New York 10017
Tel: (646) 933-1000
crhodes@dicellolevitt.com
ebruder@dicellolevitt.com

* *pro hac vice* motions to be filed

***Counsel for the Mayor and City Council of
Baltimore, Maryland***