ADAM S. SIEFF (CAL. BAR. # 302030)
DAVIS WRIGHT TREMAINE LLP
350 S. Grand Avenue, 27th Floor
Los Angeles, CA 90071
Tel.: (213) 633-8618
adamsieff@dwt.com

ERIN E. MURPHY (*pro hac vice* pending)
MATTHEW D. ROWEN (Cal. Bar #292292)
JAMES Y. XI (*pro hac vice* pending)
MITCHELL K. PALLAKI (*pro hac vice* pending)
ILAN J. POSNER (*pro hac vice* pending)
CLEMENT & MURPHY, PLLC
706 Duke Street
Alexandria, VA 22314
Tel.: (202) 742-8900
erin.murphy@clementmurphy.com

Attorneys for Plaintiff X.AI LLC

UNITED STATES DISTRICT COURT
FOR THE CENTRAL DISTRICT OF CALIFORNIA
WESTERN DIVISION

X.AI LLC *Plaintiff*, v. ROB BONTA, in his official capacity as Attorney General of the State of California, *Defendant*.	Case No.: 2:25-cv-12295-JGB-SSC *Honorable Jesus G. Bernal* **PLAINTIFF X.AI LLC'S NOTICE OF MOTION AND MOTION FOR PRELIMINARY INJUNCTION; MEMORANDUM OF POINTS AND AUTHORITIES** **ACTION SEEKING STATEWIDE RELIEF** Courtroom: 1 Date: February 23, 2026 Time: 9:00 a.m.

NOTICE OF MOTION AND MOTION FOR PRELIMINARY INJUNCTION

TO THE CLERK OF THE UNITED STATES DISTRICT COURT FOR THE CENTRAL DISTRICT OF CALIFORNIA, AND TO ALL PARTIES AND THEIR COUNSEL OF RECORD:

PLEASE TAKE NOTICE that, on Monday, February 23, 2026, at 9:00 a.m., or as soon thereafter as the matter may be heard by the Honorable Jesus G. Bernal in Courtroom 1 of the above-entitled Court, located at the George E. Brown, Jr. Federal Building and United States Courthouse, 3470 Twelfth Street Riverside, CA 92501, California, Plaintiff X.AI LLC ("xAI") will and hereby does move, this Court pursuant to Federal Rule of Civil Procedure 65(a), for an Order preliminarily enjoining Defendant Rob Bonta, in his official capacity as Attorney General of the State of California, and his employees, his agents, and successors in office from enforcing Assembly Bill 2013 ("A.B.2013" or the "Act"). xAI's motion is based on this Notice, the accompanying Memorandum of Points and Authorities, xAI's Complaint, documents incorporated by reference or subject to judicial notice, and any other materials presented at the time of the hearing.

STATEMENT OF RELIEF SOUGHT

xAI seeks a preliminary injunction pursuant to Federal Rule of Civil Procedure 65(a) enjoining Defendant Rob Bonta, in his official capacity as Attorney General of the State of California, and his employees, his agents, and successors in office from enforcing A.B.2013. As the accompanying memorandum explains, A.B.2013's provisions are unconstitutional on their face under the Due Process Clause of the Fourteenth Amendment to the U.S. Constitution and as applied to xAI under the First Amendment, the Takings Clause of the Fifth Amendment, and the

Due Process Clause of the Fourteenth Amendment to the U.S. Constitution.

L.R. 65-1 CERTIFICATE OF NOTICE TO DEFENDANT

Pursuant to Federal Rule of Civil Procedure 65(a)(1) and Local Rule 65-1, xAI, through undersigned counsel, certifies that xAI has notified Attorney General Bonta that it has filed this motion, which seeks to preliminarily enjoin the enforcement of A.B.2013.

Respectfully submitted,

s/Adam S. Sieff
ADAM S. SIEFF (CAL. BAR. # 302030)
DAVIS WRIGHT TREMAINE LLP
350 S. Grand Avenue, 27th Floor
Los Angeles, CA 90071
Tel.: (213) 633-8618
adamsieff@dwt.com

ERIN E. MURPHY (*pro hac vice* pending)
MATTHEW D. ROWEN (Cal. Bar #292292)
JAMES Y. XI (*pro hac vice* pending)
MITCHELL K. PALLAKI (*pro hac vice* pending)
ILAN J. POSNER (*pro hac vice* pending)
CLEMENT & MURPHY, PLLC
706 Duke Street
Alexandria, VA 22314
Tel.: (202) 742-8900
erin.murphy@clementmurphy.com
matthew.rowen@clementmurphy.com
james.xi@clementmurphy.com
mitchell.pallaki@clementmurphy.com
ilan.posner@clementmurphy.com

Attorneys for Plaintiff X.AI LLC

January 16, 2026

TABLE OF CONTENTS

TABLE OF AUTHORITIES iv

INTRODUCTION 1

BACKGROUND 3

A. Generative Artificial Intelligence 3

B. xAI Develops Its Own Generative AI Systems 6

C. Assembly Bill 2013 8

ARGUMENT 10

I. xAI Is Likely To Succeed On The Merits 10

A. A.B.2013 Violates the Takings Clause 10

1. xAI holds trade secrets in its dataset information 11

2. A.B.2013 effects *per se* takings of xAI's trade secrets 13

3. A.B.2013 effects regulatory takings of xAI's trade secrets 15

B. A.B.2013 Violates the First Amendment 18

C. A.B.2013 Is Unconstitutionally Vague 22

II. The Equitable Factors Overwhelmingly Support Preliminary Relief 24

CONCLUSION 25

L.R. 11-6.2 CERTIFICATE OF COMPLIANCE

PROOF OF SERVICE

TABLE OF AUTHORITIES

Cases

Abba Rubber Co. v. Seaquist,
286 Cal.Rptr. 518 (Cal. Ct. App. 1991) 12

Amgen v. Cal. Corr. Health Care Servs.,
260 Cal.Rptr.3d 873 (Cal. Ct. App. 2020) 11

Andersen v. Stability AI Ltd.,
2025 WL 1927796 (N.D. Cal. July 14, 2025) 16

Ariz. Dream Act Coal. v. Brewer,
757 F.3d 1053 (9th Cir. 2014) 25

Armstrong v. United States,
364 U.S. 40 (1960) 14

Baggett v. Bullitt,
377 U.S. 360 (1964) 23

Baird v. Bonta,
81 F.4th 1036 (9th Cir. 2023) 24, 25

Cedar Point Nursery v. Hassid,
594 U.S. 139 (2021) 10, 13, 14, 15

Cienega Gardens v. United States,
331 F.3d 1319 (Fed. Cir. 2003) 18

City of Austin v. Reagan Nat'l Advert.,
596 U.S. 61 (2022) 20

CTIA – The Wireless Ass'n v. City of Berkeley,
928 F.3d 832 (9th Cir. 2019) 24

E. Bay Sanctuary Covenant v. Biden,
993 F.3d 640 (9th Cir. 2021) 25

E. Enters. v. Apfel,
524 U.S. 498 (1998) 16

E.W. Bank v. Shanker,
2021 WL 3112452 (N.D. Cal. July 22, 2021) 12

Elrod v. Burns,
427 U.S. 347 (1976) 24

FCC v. Fox Television Stations,
567 U.S. 239 (2012) 22

Hartley Pen Co. v. U.S. Dist. Ct.,
287 F.2d 324 (9th Cir. 1961) 14

Hurley v. Irish-Am. Gay, Lesbian & Bisexual Grp. of Boston,
515 U.S. 557 (1995) 19

Iancu v. Brunetti,
588 U.S. 388 (2019) 20

Kashem v. Barr,
941 F.3d 358 (9th Cir. 2019) 23

Kelo v. City of New London,
545 U.S. 469 (2005) 18

Lingle v. Chevron U.S.A.,
544 U.S. 528 (2005) 13

McCullen v. Coakley,
573 U.S. 464 (2014) 21

Melendres v. Arpaio,
695 F.3d 990 (9th Cir. 2012) 24

Moody v. NetChoice,
603 U.S. 707 (2024) 20

Murr v. Wisconsin,
582 U.S. 383 (2017) 16

NAACP v. Button,
371 U.S. 415 (1963) 22, 23

Nat'l Inst. of Fam. & Life Advocs. v. Becerra,
585 U.S. 755 (2018) 19

O2 Micro Int'l v. Monolithic Power Sys.,
420 F.Supp.2d 1070 (N.D. Cal. 2006) 11

Penn Cent. Transp. Co. v. City of New York,
438 U.S. 104 (1978) 15

Pharm. Rsch. & Mfrs. of Am. v. Stolfi,
153 F.4th 795 (9th Cir. 2025) 16, 17

Pharm. Rsch. & Mfrs. of Am. v. Williams,
64 F.4th 932 (8th Cir. 2023) 15

Philip Morris v. Reilly,
312 F.3d 24 (1st Cir. 2002) 17, 18

Platt v. Moore,
15 F.4th 895 (9th Cir. 2021) 25

Quintara Biosciences v. Ruifeng Biztech,
149 F.4th 1081 (9th Cir. 2025) 13

Riley v. Nat'l Fed'n of the Blind of N.C.,
487 U.S. 781 (1988) 19

Ruckelshaus v. Monsanto Co.,
467 U.S. 986 (1984) 10, 14, 16, 17

Sheetz v. Cnty. of El Dorado,
601 U.S. 267 (2024) 10, 13

Sorrell v. IMS Health,
564 U.S. 552 (2011) 20

Tyler v. Hennepin Cnty.,
598 U.S. 631 (2023) 10, 15

United States v. Google,
2025 WL 2523010 (D.D.C. Sept. 2, 2025) 5

United States v. Google,
747 F.Supp.3d 1 (D.D.C. 2024) 3

United States v. Sec. Indus. Bank,
459 U.S. 70 (1982) 16

United States v. Williams,
553 U.S. 285 (2008) 22

W. Va. State Bd. of Educ. v. Barnette,
319 U.S. 624 (1943) 19

WeRide Corp. v. Kun Huang,
379 F.Supp.3d 834 (N.D. Cal. 2019) 11, 13, 24

Whyte v. Schlage Lock Co.,
125 Cal.Rptr.2d 277 (Cal. Ct. App. 2002) 12, 13

X Corp. v. Bonta,
116 F.4th 888 (9th Cir. 2024) 19, 21

Zepeda v. INS,
753 F.2d 719 (9th Cir. 1983) 25

Constitutional Provision

U.S. Const. amend. V 10

Statutes

15 U.S.C §9401(3) 3

18 U.S.C. §§1831, *et seq* 15

18 U.S.C. §1839(3) 11, 13

Cal. Civ. Code §§3426, *et seq.* 15

Cal. Civ. Code §3110(b) 9

Cal. Civ. Code §3110(f) 23

Cal. Civ. Code §3111 ..... 8, 9, 23

Cal. Civ. Code §3111(a) 9, 23

Cal. Civ. Code §3111(a)(1)-(12) 9

Cal. Civ. Code §3111(a)(1)-(4) 14

Cal. Civ. Code §3111(a)(2) 22

Cal. Civ. Code §3111(a)(5) 14, 22

Cal. Civ. Code §3111(a)(6) 14

Cal. Civ. Code §3111(a)(7)-(8) 14

Cal. Civ. Code §3111(a)(12) 14

Cal. Civ. Code §3111(b) 9, 17, 20

Cal. Civ. Code §3426.1(d) 11, 13

Other Authorities

Cal. S. Rules Comm., A.B.2013, *Senate Floor Analysis* (Aug. 20, 2024), https://perma.cc/6L5E-LXKC 8, 20

Order, *Tremblay v. OpenAI*, No. 3:23-cv-3223 (N.D. Cal. Sept. 24, 2024) 12

What Is Generative AI?, McKinsey (Apr. 2, 2024), https://perma.cc/GHW4-DZSE 3

xAI, *Grok 4 Fast Model Card* (Sept. 19, 2025), https://perma.cc/FEK8-CYQV 6, 21

xAI, *xAI Frontier Artificial Intelligence Framework* (Dec. 30, 2025), https://perma.cc/X45R-NM2N 9

INTRODUCTION

For as long as companies have been innovating, states have protected their trade secrets. The reason is simple: Companies are unlikely to invest the considerable time and money required to develop novel products if they cannot protect the methods, processes, and information that make their products innovative. Companies have relied on trade-secret protections to develop fertilizers, airplane parts, consumer beverage products, and everything in between. Trade-secret protections have been applied across industries, allowing innovation to thrive in all manner of sectors. The burgeoning artificial intelligence ("AI") industry is no different.

AI-focused companies like xAI develop generative AI models, which are software programs that simulate human intelligence to perform a wide range of tasks and that create original content based on user prompts. To train these models, companies utilize datasets—collections of information that derive from various sources. Companies input datasets into their AI models so that they can memorize and then extrapolate from the information to generate new content. By repeating the process, AI companies train their models to perform a variety of tasks—everything from answering trivia questions and summarizing documents to generating images and videos.

But those models are only as good as their training, which is only as good as the data the company utilizes. AI developers thus dedicate substantial resources to identifying high-quality training data, particularly from sources their competitors are not using. After all, if a developer's AI model receives unique training based on sources of data a rival lacks, it will be able to respond in unique ways, resulting in a competitive advantage. Unsurprisingly, then, businesses like xAI make

significant efforts to safeguard information about their datasets as trade secrets.

Until recently, California protected those trade secrets from disclosure. But Assembly Bill 2013 ("A.B.2013") threatens to eviscerate those protections. Though billed as a consumer-transparency statute, A.B.2013 is really a trade-secrets-destroying disclosure regime that hands competitors a roadmap for how companies like xAI develop and train their proprietary models. If enforced, A.B.2013 would compel xAI to publicly disclose critical details about its datasets that would reveal confidential information about how xAI develops, trains, and refines its unique generative AI models. That approach might be at least understandable—albeit still unconstitutional—if the information were valuable to consumers. But A.B.2013 does little to help consumers, who are far more interested in evaluating how a model performs than in obtaining technical details about the datasets used to train it. The only parties likely to benefit from A.B.2013 are competitors. Unlike consumers, xAI's rivals have the wherewithal and motivation to use this detailed information to replicate xAI's models or improve their own, thus robbing xAI of its competitive edge.

A.B.2013's disclosure regime is unconstitutional several times over. It violates the Takings Clause by forcing xAI to disclose its economically valuable and fiercely protected trade secrets without any promise of compensation—a *per se* taking and regulatory taking to boot. It violates the First Amendment by compelling xAI to disseminate specific information based on the generative AI models it develops, a content- and viewpoint-based requirement that triggers and fails strict scrutiny. But whatever scrutiny ultimately applies, A.B.2013's disclosure requirement fails because it is not at all tailored to further California's professed consumer-transparency goal. Finally, A.B.2013 is unconstitutionally

vague. It provides no guidance as to which AI systems or datasets it covers and how much detail developers must provide. Without such clarity, A.B.2013 invites arbitrary enforcement and forces developers to over-disclose to avoid violating the law. Due process demands more, especially when trade-secrets and speech rights are at stake.

In sum, A.B.2013 threatens to gut the AI industry, violating numerous constitutional provisions along the way. The Court should enjoin the enforcement of this deeply misguided and Constitution-flouting law.

BACKGROUND

A. Generative Artificial Intelligence.

"Artificial intelligence is the science and engineering of getting machines … to exhibit intelligent behavior." *United States v. Google*, 747 F.Supp.3d 1, 52 (D.D.C. 2024); 15 U.S.C §9401(3). Today, most developers focus on generative artificial intelligence, or gen AI for short. Unlike older AI systems, gen AI creates new content from user prompts. In the last few years, gen AI's capabilities and potential have progressed rapidly, potentially adding up to $4.4 trillion to the global economy annually. *What Is Generative AI?*, McKinsey (Apr. 2, 2024), https://perma.cc/GHW4-DZSE.

Datasets are essential to achieving content-generative capabilities. Gen AI developers use datasets to teach AI models to recognize patterns and make predictions. Stanley.Decl.¶7. Broadly defined, datasets are collections of stored information used to develop and train gen AI systems. Stanley.Decl.¶¶6-8. The information within a dataset can vary wildly, from a simple table of data points (e.g., spreadsheets containing addresses and emails) to a combination of text, images, and audio (e.g., notes, photos, and voicemails saved on a cellphone). Stanley.Decl.¶8.

Datasets can be gathered from a wide variety of sources. They can be acquired by purchasing pre-packaged datasets, as advertisers often do when they purchase customer information. They can be developed by using application programming interfaces, which allow different software applications to communicate with one another. And they can be compiled from a wide variety of bespoke sources, including by licensing data.

AI companies do not rely on all the same data; some acquire and utilize data that others do not. What differentiates one company's AI model from another depends in large part on the unique data it has acquired. Stanley.Decl.¶¶10-11. Take legal data. While many developers might use data from obvious sources like federal or state courts' websites, some may use less obvious sources, such as Justia or the Government Publishing Office. Those latter sources might have distinct data that provide unique training, giving an AI company that uses them a competitive advantage. Stanley.Decl.¶11. Companies may also acquire unique data from non-public sources. For example, a developer may assemble briefs addressing common legal questions to train its legal AI model to perform better than other models that lack such data. Thus, datasets—especially the specific sources and kinds of data they contain—are the linchpin in an AI model's development and success. Stanley.Decl.¶¶9, 24.

Before datasets can be used to develop AI models, they must be "cleaned" and converted into a format that a computer can understand. Stanley.Decl.¶12. That requires AI engineers to remove duplicate or incomplete entries (e.g., repeated phone numbers or those lacking area codes). The conversion, or "tokenization," process involves translating data into numerical representations, called tokens, that the model can process. Stanley.Decl.¶13. Tokens are essentially components of

larger datasets that represent words, characters, or phrases. Stanley.Decl.¶13. The tokenization process is critical in preparing data for further processing.

Creating a gen AI system proceeds in several stages. First, in pre-training, engineers develop a foundation model—a term used to describe gen AI models trained using a vast range of datasets and data formats. Stanley.Decl.¶14. These foundation models take data inputs, convert them to tokens, predict the most likely next token in a sequence, and then convert those predicted tokens back into language. *See United States v. Google*, 2025 WL 2523010, at *9-10 (D.D.C. Sept. 2, 2025). A model's ability to predict the next token depends on the quality, diversity, and quantity of the input data. *Id.* So to build effective foundation models, developers must input huge volumes of raw data, which is computationally intensive (requiring thousands of clustered graphics processing units), time-consuming (weeks of processing), and expensive (costing millions of dollars). Stanley.Decl.¶14.

With a foundation model in place, the AI system can be fine-tuned for specific applications. Developers use tools like reinforcement learning to transform a foundation model's unfiltered and potentially chaotic responses into more polished and reliable outputs. Stanley.Decl.¶15. That typically involves a smaller, yet higher-quality training dataset to target an AI system's specific capabilities. Stanley.Decl.¶16. To create such datasets, developers label data and curate targeted data, requiring more direct human involvement. Once the model is fine-tuned, it is regularly assessed, refined, and improved. That can occur through additional training focused on fine-tuning, or by analyzing its actual outputs to real-world user prompts.

B. xAI Develops Its Own Generative AI Systems.

In March and April 2023, xAI began developing its own AI models. Stanley.Decl.¶¶3-4. It invested substantial time and energy to acquire datasets that vary in source, data type, and format to develop and eventually train its AI models. *See* Stanley.Decl.¶¶4, 9, 27, 29.

As part of that process, xAI used the methodology outlined above. It first acquired and refined datasets to create a foundation AI model unique to xAI. Stanley.Decl.¶¶12-14. xAI engineers then used specially constructed datasets to tweak and refine the model. Stanley.Decl.¶¶14-16. They adjusted different aspects of the foundation model's code and tested the resulting variations using xAI's curated datasets to see which were most effective. Stanley.Decl.¶15.

In November 2023, after months of development, xAI executed a limited public release of its flagship AI model, Grok-1—with a full public release a few months later. Stanley.Decl.¶4. xAI followed up on that success by releasing Grok-2 (August 2024), Grok-3 (February 2025), and Grok-4 (July 2025). Stanley.Decl.¶4. While xAI is constantly developing new versions of its models, it also releases updates to existing models. Stanley.Decl.¶4. Despite starting more recently than many competitors, xAI's models have consistently topped performance benchmarks, demonstrating the success of its development process and the value of its datasets. Stanley.Decl.¶5.

xAI values consumer confidence. To facilitate that trust, xAI provides up-to-date information about how its models perform. xAI, *Grok 4 Fast Model Card* (Sept. 19, 2025), https://perma.cc/FEK8-CYQV. For example, xAI evaluates whether its models exhibit political bias and how they fare when asked to give the wrong answer to a question. *See id.* at 4. These tests reflect how well xAI's models

perform tasks consumers give them. Stanley.Decl.¶¶5, 26. Because the tests replicate how consumers interact with a model, they help consumers assess whether the AI model is worthwhile.

Due to the critical role datasets play in the development process, details about those datasets are highly valuable—and kept secret. Barrett.Decl.¶¶15-16. So long as xAI uses datasets that its competitors do not, its models will be better trained. But if xAI revealed details about those datasets, competitors would immediately move to acquire the sources they are missing to ensure their models are equally effective. Stanley.Decl.¶¶22-23; Barrett.Decl.¶16. The same is true of the amount of data xAI uses; the breadth of information xAI uses to train an AI model affects its ability to perform various tasks. Stanley.Decl.¶24. So too for xAI's method for cleaning and refining its datasets, as it highlights what data xAI believes best improves its model's capabilities. Stanley.Decl.¶25. There is thus significant value in xAI keeping details about its datasets secret; disclosing that information would undercut any competitive edge xAI holds in the AI development processes.

Given the economic importance of this confidential information, xAI takes many measures to prevent disclosure and maintain secrecy. All employees sign confidentiality provisions to work for xAI and develop its models. Barrett.Decl.¶4. These provisions, along with xAI's robust confidentiality policy, emphasize that all parts of the development process are non-public and proprietary, that information must be used solely for developing xAI's models, and that nothing should be publicly disclosed. Barrett.Decl.¶¶4-5. xAI also securely stores its datasets in locations known only to those that need access for approved purposes and that can be accessed only by those with the appropriate level of access. Barrett.Decl.¶¶7, 10. xAI further protects its datasets with security alerts, by encouraging employees

to report unauthorized access, and through periodic tests and exercises that simulate security breaches. Barrett.Decl.¶¶6, 11, 14. xAI has introduced role-based access requirements, which ensure that an employee's access is limited to the datasets that employee actually needs. Barrett.Decl.¶8. And xAI has implemented time-limited access controls, so that access does not extend beyond the needs of particular projects. Barrett.Decl.¶9. In short, xAI takes safeguarding information about its datasets extremely seriously.

C. Assembly Bill 2013.

In September 2024, the Governor of California signed A.B.2013 into law. The bill, entitled "Artificial Intelligence Training Data Transparency," imposes substantial information-disclosure requirements on gen AI developers. It took effect on January 1, 2026. Cal. Civ. Code §3111.

Though A.B.2013 lacks a statement of purpose, the legislative history describes it as "providing transparency to consumers of AI systems … by providing important documentation about the data used to train" such systems that will purportedly "help[] identify and mitigate biases." *See* Cal. S. Rules Comm., A.B.2013*, Senate Floor Analysis* 3 (Aug. 20, 2024), https://perma.cc/6L5E-LXKC. However, A.B.2013 does nothing to accomplish that goal. It does not require AI companies to disclose the kinds of information consumers find useful—e.g., how well an AI model actually performs real-world tasks. Instead, A.B.2013 requires companies to disclose proprietary dataset information that is useful only to parties with the requisite technical expertise—*viz.*, competitors. Stanley.Decl.¶¶20-25; Barrett.Decl.¶¶15-16.

Specifically, A.B.2013 requires any "developer" of a "generative artificial intelligence system or service" made "publicly available to Californians for use"

after January 1, 2022, to "post" on its "website documentation regarding the data used … to train" the gen AI system. Cal. Civ. Code §3111. Those disclosures must include, but are not limited to, a "high-level summary of the datasets used in the development" of the gen AI model. *Id.* §3111(a). A.B.2013 also lists 12 categories of dataset-related information that developers must disclose, including their sources; a description of how the datasets further the model's intended purpose; their size; whether the datasets include intellectual property or personal or aggregate consumer information; whether there was any cleaning, processing, or other modification of the datasets; and more. *Id.* §3111(a)(1)-(12).

California appears to recognize that its disclosure requirements demand the disclosure of valuable, confidential information because A.B.2013 contains three exceptions: AI models (1) used solely to help ensure security and integrity; (2) used solely for the operation of aircraft; or (3) developed for federal national security purposes are exempt from its onerous requirements. *Id.* §3111(b).

Despite imposing burdensome disclosure obligations, A.B.2013 leaves key terms undefined. It never defines what constitutes a "dataset" or how "high-level" "summar[ies]" must be—even though it defines "developer" extremely broadly to cover anyone who "designs, codes, produces, or substantially modifies an artificial intelligence system or service." *Id.* §3110(b). Given this uncertainty about A.B.2013's obligations, xAI would need to locate, collect, summarize, and disclose extensive proprietary information about datasets it has used to train and develop all of its AI models released since 2022 to comply.[1] Stanley.Decl.¶¶27-28.

[1] As a precaution, xAI provided a high-level, limited disclosure that does not reveal its trade secrets. xAI, *xAI Frontier Artificial Intelligence Framework* (Dec. 30, 2025), https://perma.cc/X45R-NM2N. But due to A.B.2013's apparent broad

ARGUMENT

To obtain a preliminary injunction, xAI must show: (1) a substantial likelihood of success on the merits; (2) a substantial threat of irreparable injury; (3) the threatened injury outweighs the threatened harm to the state; and (4) injunctive relief is in the public interest. *See Apartment Ass'n of L.A. Cnty. v. City of Los Angeles*, 10 F.4th 905, 912 (9th Cir. 2021). xAI satisfies all these factors.

I. xAI Is Likely To Succeed On The Merits.

A. A.B.2013 Violates the Takings Clause.

The Takings Clause provides that "private property [shall not] be taken for public use, without just compensation." U.S. Const. amend. V. "By requiring the government to pay for what it takes, the Takings Clause saves individual property owners from bearing 'public burdens which, in all fairness and justice, should be borne by the public as a whole.'" *Sheetz v. Cnty. of El Dorado*, 601 U.S. 267, 273-74 (2024). As the Supreme Court has underscored, these protections safeguard intangible property rights, like trade secrets, as well. *See Ruckelshaus v. Monsanto Co.*, 467 U.S. 986, 1003-04 (1984).

Government action that triggers the Takings Clause is divided into "*per se* takings" and "regulatory takings." *See Cedar Point Nursery v. Hassid*, 594 U.S. 139, 146-49 (2021). By requiring xAI to disclose information protected as trade secrets under the federal Defend Trade Secrets Act ("DTSA") and California's Uniform Trade Secrets Act ("CUTSA"), A.B.2013 accomplishes both. *See Tyler v. Hennepin Cnty.*, 598 U.S. 631, 638 (2023) (property protected by the Takings

---

sweep, there is a material risk that California will find xAI's disclosure insufficient. If California were permitted to force xAI to disclose further details and reveal its trade secrets, it would violate xAI's constitutional rights as outlined below.

Clause is drawn from "existing rules or understandings" of property, including state law). Because California provides no mechanism to compensate xAI for that lost property, it cannot enforce A.B.2013 against xAI consistent with the Takings Clause.

1. xAI holds trade secrets in its dataset information.

xAI undeniably holds trade secrets in the sources of its datasets, their sizes, and the methods used to clean and refine them. Under California law, a party has a trade-secrets property right in information (1) that is valuable because it is unknown and (2) that the owner has attempted to keep secret. *Amgen v. Cal. Corr. Health Care Servs.*, 260 Cal.Rptr.3d 873, 886 (Cal. Ct. App. 2020); Cal. Civ. Code §3426.1(d). The same is true under the federal Defend Trade Secrets Act. *See* 18 U.S.C. §1839(3). xAI's dataset information qualifies as trade secrets under both.

First, xAI's dataset information is economically valuable precisely because it is unknown. xAI's sources of data reveal the information it uses to train its models, which in turn sheds light on how useful that information is to the models' performance. *See supra* pp.4, 7. xAI's decisions about how to craft the best mixture of data from public and non-public sources, or how best to balance different combinations of text, images, and audio data, would be extremely valuable to competitors seeking to replicate xAI's top-notch models and improve their own training. Stanley.Decl.¶¶24-25; *see O2 Micro Int'l v. Monolithic Power Sys.*, 420 F.Supp.2d 1070, 1089 (N.D. Cal. 2006) ("Combinations of public information from a variety of different sources when combined in a novel way can be a trade secret."); *WeRide Corp. v. Kun Huang*, 379 F.Supp.3d 834, 847 (N.D. Cal. 2019).

The same is true of the amount of data xAI uses, as the differences between the datasets each company uses is what gives xAI a competitive edge. Stanley.Decl.¶¶11, 23-24. Competitors that know the size of xAI's datasets can

evaluate what information xAI has, what they lack, and how they can fill any gaps in their own training. *See Abba Rubber Co. v. Seaquist*, 286 Cal.Rptr. 518, 527 (Cal. Ct. App. 1991) (information that tells competitors "fact[s] which they previously did not know" are valuable trade secrets). Likewise, if competitors had access to xAI's processes for cleaning, modifying, and refining that data, that information would lose its primary value—to prevent rivals from using that insight to pinpoint deficiencies in their training regimen and replicate xAI's successes. *E.W. Bank v. Shanker*, 2021 WL 3112452, at *9 (N.D. Cal. July 22, 2021) (parties can hold trade-secrets property in "roadmaps related to confidential technology" that enable competitors to recreate its products).

xAI is not alone in understanding the trade-secrets rights inherent in such information. Courts around the country have recognized as much by entering protective orders safeguarding AI-training data against disclosure. *See, e.g.*, Order at 2, *Tremblay v. OpenAI*, No. 3:23-cv-3223 (N.D. Cal. Sept. 24, 2024), Dkt.182. As those courts have recognized, dataset information has "independent economic value"—thereby satisfying the first prongs of the CUTSA and DTSA—because it would be "valuable to a competitor" if disclosed. *Whyte v. Schlage Lock Co.*, 125 Cal.Rptr.2d 277, 287 (Cal. Ct. App. 2002).

Second, xAI has undertaken extensive efforts to keep this information secret. xAI employees must sign confidentiality provisions reinforcing that all parts of the development process are non-public and not subject to disclosure. *See id.* at 286-87. xAI likewise utilizes various methods to physically and digitally safeguard the information, including securely storing it, providing role-based, time-limited access only on a need-to-know basis, implementing a security-alert system, running security tests and exercises, and encouraging employees to report unauthorized

access. Barrett.Decl.¶¶4-14. These security efforts demonstrate that xAI takes protecting its trade secrets extremely seriously and seeks to ensure that xAI alone has access to this highly valuable information that is critical to its competitive edge in innovating AI models. *See Whyte*, 125 Cal.Rptr.2d at 286-87 (reasonable efforts to maintain secrecy include "advising employees of the existence of a trade secret, limiting access … on 'need to know basis,' and controlling plant access"); *WeRide*, 379 F.Supp.3d at 847 (similar).

xAI thus plainly holds trade secrets in the dataset information under both federal and California law.

2. A.B.2013 effects *per se* takings of xAI's trade secrets.

Because xAI holds trade-secrets property rights in its dataset information, the Takings Clause protects such information against *per se* takings. A *per se* taking occurs whenever government effects "a direct government appropriation or physical invasion of private property," *Lingle v. Chevron U.S.A.*, 544 U.S. 528, 537 (2005), including via regulations that appropriate "a fundamental element of the [owner's] property right," *Cedar Point*, 594 U.S. at 149-50. One such fundamental element is "the owner's right to exclude others from [its property]." *Sheetz*, 601 U.S. at 274 (2024). Laws that appropriate the right to exclude—"one of the most treasured rights of property ownership"—thus work *per se* takings. *Cedar Point*, 594 U.S. at 149, 155.

A.B.2013's disclosure obligations, which require companies like xAI to reveal their trade secrets, do just that. The key feature of the trade-secret property right is *secrecy*. Cal. Civ. Code §3426.1(d); 18 U.S.C. §1839(3); *accord Quintara Biosciences v. Ruifeng Biztech*, 149 F.4th 1081, 1085 (9th Cir. 2025) ("By definition, trade secrets derive their value from nondisclosure."). It therefore

follows that "[t]he property in a trade secret is the power to make use of it to the exclusion of the world," *Hartley Pen Co. v. U.S. Dist. Ct.*, 287 F.2d 324, 328 (9th Cir. 1961), making the "right to exclude" others from accessing that information the "*sine qua non*" of that property interest, *Cedar Point*, 594 U.S. at 150. "With respect to a trade secret, the right to exclude others is central to the very definition of the property interest." *Ruckelshaus*, 467 U.S. at 1011.

By appropriating xAI's dataset-related trade secrets, A.B.2013's disclosure obligations effect *per se* takings. They appear to compel xAI to reveal the sources of its datasets, their sizes, the type of data xAI uses, and how the datasets further the intended purpose of xAI's models—all information that would otherwise be protected as confidential trade secrets. Cal. Civ. Code §3111(a)(1)-(4), (6). Moreover, A.B.2013 requires xAI to disclose whether its datasets contain intellectual property or personal or aggregate consumer information, and whether xAI uses cleaning or modification of its datasets. *Id.* §3111(a)(5), (7)-(8), (12). To the extent these vague provisions demand anything more than a "yes" or "no" answer, they too would appropriate xAI's trade secrets, as such granular details have all the hallmarks of protected trade secrets. *See supra* Section I.A.

That appropriation is a quintessential *per se* taking. *See Armstrong v. United States*, 364 U.S. 40, 48-49 (1960) (government action eviscerating mechanic's liens effected a *per se* taking because, "[b]efore the liens were destroyed, the lienholders admittedly had compensable property," and "afterwards, they had none"). If xAI must comply with A.B.2013, its trade secrets will be "disclosed to others," and xAI will "ha[ve] lost [its] property interest" in those trade secrets entirely. *Ruckelshaus*, 467 U.S. at 1011. Accordingly, California can impose A.B.2013's disclosure requirements only if it provides just compensation. Yet A.B.2013 does not

contemplate any means to do so. California thus cannot enforce A.B.2013 against xAI consistent with the Takings Clause and should be enjoined from doing so. *See Pharm. Rsch. & Mfrs. of Am. v. Williams*, 64 F.4th 932, 942-46 (8th Cir. 2023).

3. A.B.2013 effects regulatory takings of xAI's trade secrets.

Because A.B.2013 effectuates a *per se* taking, there is no need to consider the law under the factors set forth in *Penn Central Transportation Co. v. City of New York*, 438 U.S. 104 (1978). *See Cedar Point*, 594 U.S. at 149 (*Penn Central*'s balancing test has "no place" when the government "appropriates for the enjoyment of third parties the owners' right to exclude"). That said, each *Penn Central* factor confirms that A.B.2013 accomplishes a classic taking.[2]

Take, for example, the question of whether A.B.2013 interferes with investment-backed expectations. It obviously does. xAI began investing in and developing its trade secrets in early 2023 (well before A.B.2013 was introduced in January 2024), with the expectation that those trade secrets would be protected under both California and federal law. *See* Cal. Civ. Code §§3426, *et seq.*; 18 U.S.C. §§1831, *et seq*. xAI had no reason to expect that it would lose the value of its dataset trade secrets months after releasing its first AI model. After all, the government cannot "manipulate[]" or "extinguish a property interest that it recognizes everywhere else to avoid paying just compensation when it is the one doing the taking." *Tyler*, 598 U.S. at 645. Yet by bucking both state and federal law and compelling the disclosure of xAI's valuable trade secrets, A.B.2013 does precisely that. Worse, because it applies retroactively to January 2022, A.B.2013 compounds that disregard for settled expectations by targeting investments in trade

[2] xAI reserves the right to ask the Supreme Court to overrule *Penn Central*.

secrets at a time when xAI had no notice that the highly confidential information it was developing could be subject to sweeping disclosure obligations. *See E. Enters. v. Apfel*, 524 U.S. 498, 533-34 (1998) (plurality op.) ("Our Constitution expresses concern with retroactive laws through several of its provisions, including the ... Takings Clause[]."); *accord United States v. Sec. Indus. Bank*, 459 U.S. 70, 76-78 (1982). xAI could not reasonably have expected California to destroy this longstanding property interest by legislative fiat. *See Murr v. Wisconsin*, 582 U.S. 383, 396 (2017) (states lack "unfettered authority to 'shape and define property rights and reasonable investment-backed expectations,' leaving []owners without recourse against unreasonable regulations").

A.B.2013 would also substantially interfere with the economic value of xAI's datasets. The quality of xAI's datasets, their sources, and size all contribute to ensuring xAI's models are competitive. The economic value of this information "lies in the competitive advantage over others that [xAI] enjoys by virtue of its exclusive access to the data, and disclosure or use by others of the data would destroy that competitive edge." *Ruckelshaus*, 467 U.S. at 1012; *Pharm. Rsch. & Mfrs. of Am. v. Stolfi*, 153 F.4th 795, 839 (9th Cir. 2025) ("[D]isclosure of that data entirely extinguish[es] the value of the trade secret[s]."). Even if the data could still be used for training purposes, forcing xAI to disclose such information would allow rival developers to replicate xAI's models, eliminating the value that information derives from being kept secret. *Andersen v. Stability AI Ltd.*, 2025 WL 1927796, *2 (N.D. Cal. July 14, 2025) (forcing AI developer to disclose training data "raises serious competitive concerns").[3] A.B.2013's compelled disclosure requirements

[3] California clearly recognizes this information as valuable. It specifically

thus would have a severe economic impact on xAI. Stanley.Decl.¶¶20-25; *Philip Morris v. Reilly*, 312 F.3d 24, 41 (1st Cir. 2002) (en banc) ("potentially tremendous" economic impact from compelled disclosure of information that "make[s] it much easier to reverse engineer" trade secrets weighed in favor of finding a regulatory taking). The nascent AI industry is rife for substantial growth, so *any* competitive advantage is key to xAI's long-term success. But A.B.2013's disclosures would eviscerate any such advantage and destroy the economic value of xAI's trade secrets.

Third, the character of the government action here supports finding that A.B.2013 effects a taking. xAI has a property right in its dataset information, but that "property right is extinguished" once it is forced to comply with A.B.2013. *Ruckelshaus*, 467 U.S. at 1002. A.B.2013's categorical disclosure obligation—without *any* mechanism for xAI to designate information as confidential—confirms that it effects a regulatory taking. Unlike more narrowly focused disclosure statutes requiring case-specific determinations about the public need for disclosure, A.B.2013 broadly compels disclosure of xAI's trade secrets regardless of whether that information benefits the public. *Cf. Stolfi*, 153 F.4th at 839-40. In fact, it is not clear what consumers are supposed to do with the information companies must disclose under A.B.2013. A.B.2013 does not compel developers to alert the public to risks or dangers associated with using AI products. It just requires them to disclose dataset information. Without the requisite technical expertise, consumers are in no position to assess whether particular datasets are sufficiently comprehensive or whether one AI model is better than another based solely on the

---

exempts models used for security, aircraft operations, or national security from the disclosure obligations. Cal. Civ. Code §3111(b).

sources or sizes of datasets used for training. Stanley.Decl.¶26; *cf. Reilly*, 312 F.3d at 44 ("[F]or a state to be able to completely destroy valuable trade secrets, it should be required to show more than a *possible* beneficial effect."). Consumers' best metric—and the only one most care about—is the end product.

The principal beneficiaries of these disclosure requirements are instead competitors, who will use the information to bolster their own products. Because A.B.2013 essentially forces xAI to transfer its trade secrets to its private-company competitors without providing any discernable benefit to consumers, it plainly lacks a valid "public purpose." *See Kelo v. City of New London*, 545 U.S. 469, 478 (2005). Without a valid "public purpose," A.B.2013 cannot even begin to satisfy the threshold requirement that the taking be for a "public use." *See id.* at 477-78. But that aside, that A.B.2013 principally benefits xAI's competitors confirms that the character of government action weighs in favor of finding a taking. xAI would bear the entire burden of that boon to competitors through the loss of its trade-secrets property rights. *Cienega Gardens v. United States,* 331 F.3d 1319, 1338-39 (Fed. Cir. 2003) (finding a taking because the public program involved "the kind of expense-shifting to a few persons that amounts to a taking").

All three *Penn Central* factors therefore confirm that A.B.2013's mandated disclosure of xAI's trade secrets, at a bare minimum, accomplishes an uncompensated regulatory taking.

B. A.B.2013 Violates the First Amendment.

A.B.2013 likewise violates the First Amendment by impermissibly compelling xAI's speech on the basis of both content and viewpoint.

As the Supreme Court has long recognized, the First Amendment's guarantee of free speech "necessarily compris[es] the decision of both what to say and what

not to say." *Riley v. Nat'l Fed'n of the Blind of N.C.*, 487 U.S. 781, 796-97 (1988). That includes not just statements bearing particular messages, *see W. Va. State Bd. of Educ. v. Barnette*, 319 U.S. 624, 642 (1943), but also "statements of *fact* the speaker would rather avoid," *Hurley v. Irish-Am. Gay, Lesbian & Bisexual Grp. of Boston*, 515 U.S. 557, 573 (1995) (emphasis added). Laws compelling speech are thus generally treated no differently from laws restricting speech, *see Nat'l Inst. of Fam. & Life Advocs. v. Becerra* (*NIFLA*), 585 U.S. 755, 766-67 (2018), even when the government does not compel speakers to express a particular message, *Riley*, 487 U.S. at 795.

A.B.2013 triggers strict scrutiny because it is both content- and viewpoint-based. A.B.2013 is content-based because it compels xAI to disclose specific *content* about its AI models. *X Corp. v. Bonta*, 116 F.4th 888, 900 (9th Cir. 2024); *NIFLA*, 585 U.S. at 766. In *NIFLA*, for example, the Supreme Court held that a law compelling crisis pregnancy centers to disseminate a government-drafted message that California provides free or low-cost abortion services triggered strict scrutiny because it compelled them to speak "a particular message" that "alte[red] the content of [their] speech." 585 U.S. at 766. Similarly, in *Riley*, the Supreme Court applied strict scrutiny to a law requiring professional fundraisers to disclose the percentage of their charitable contributions "actually turned over to charity." 487 U.S. at 795. As the Court explained, "[m]andating speech that a speaker would not otherwise make necessarily alters the content of the speech." *Id.* And in *X Corp.*, the Ninth Circuit held that a law requiring social-media companies to submit reports about their terms of service and content-moderation policies was a content-based regulation of speech subject to strict scrutiny. 116 F.4th at 903. Like those content-based disclosures, A.B.2013 forces xAI to speak a particular message it does not

want to convey to the public.

A.B.2013 compounds its First Amendment problems by discriminating based on viewpoint. *See Iancu v. Brunetti*, 588 U.S. 388, 394 (2019). A.B.2013 exempts developers of AI models with certain favored "purpose[s]" related to network security, aircraft operations, and national security from its onerous requirements. Cal. Civ. Code §3111(b). Those "purpose"-based distinctions are proxies for viewpoint discrimination. *See City of Austin v. Reagan Nat'l Advert.*, 596 U.S. 61, 74 (2022). The First Amendment does not permit California to compel private speech based on its perception that certain ideas (i.e., training data for security, aviation, or military AI models) are important enough to be kept secret, and that other, less-favored ideas (e.g., training data for creative-writing AI models) are not. A.B.2013's asserted purpose—compelling AI developers' speech to "identify and mitigate biases"—confirms its discriminatory design. Cal. S. Rules Comm., *supra*, at 3. It forces developers to publicly disclose their data sources in an attempt to identify what California deems to be "data riddled with implicit and explicit biases." *Id*. at 7; *see Moody v. NetChoice*, 603 U.S. 707, 741 (2024) ("[A] State may not interfere with private actors' speech to advance its own vision of ideological balance."). Given datasets are the knowledge base that AI models use to reach conclusions, A.B.2013's compelled disclosure of those sources indirectly attempts to influence the viewpoints espoused by xAI's models (i.e., their outputs) by targeting the data that goes into them. "Given the legislature's expressed statement of purpose, it is apparent that [A.B.2013] imposes burdens that are based on the content of speech and that are aimed at a particular viewpoint," triggering strict scrutiny several times over. *Sorrell v. IMS Health*, 564 U.S. 552, 565 (2011).

But whatever form of heightened scrutiny applies, A.B.2013 cannot survive.

Even under intermediate scrutiny, California must show that A.B.2013 is "narrowly tailored to serve a significant governmental interest," *McCullen v. Coakley*, 573 U.S. 464, 486 (2014), and A.B.2013 is plainly not narrowly tailored to advance any legitimate interest California could assert. Its disclosure obligations "are more extensive than necessary" to help "consumers make informed decisions" about AI models. *X Corp.*, 116 F.4th at 903. After all, it is far from clear how the trade secrets A.B.2013 would force xAI to disclose are of any value to consumers at all. *See supra* pp.8, 17-18. Reports from developers or outside certifiers who test AI models' effectiveness are far more useful to consumers than raw data. xAI releases "Model Cards" for its AI models regarding its models' *outputs* for precisely that reason. *See Model Card*, *supra*. By contrast, A.B.2013 is directed at the model's *inputs*—i.e., the data used during development. An exhaustive list of the raw datasets used to develop and refine xAI's models does not give the average consumer any meaningful information about those models; it simply enables competitors to replicate a valuable model's success.

In short, A.B.2013 imposes onerous disclosure requirements while providing no meaningful benefit to consumers—the opposite of narrow tailoring. *X Corp.*, 116 F.4th at 903. That A.B.2013 compels disclosures regarding all AI models released since 2022, even those no longer regularly used by consumers, confirms the disconnect between the law's obligations and the legislature's consumer-transparency goal. Because A.B.2013 imposes onerous disclosure obligations that do nothing to advance its professed consumer-protection interests, it cannot survive any level of heightened scrutiny. A.B.2013 accordingly violates xAI's First Amendment rights.

C. A.B.2013 Is Unconstitutionally Vague.

A.B.2013 is also unconstitutionally vague both facially and as applied to xAI.

"A fundamental principle in our legal system is that laws which regulate persons or entities must give fair notice of conduct that is forbidden or required." *FCC v. Fox Television Stations*, 567 U.S. 239, 253 (2012). A law is unconstitutionally vague if it "fails to provide a person of ordinary intelligence fair notice of what is prohibited, or is so standardless that it authorizes or encourages seriously discriminatory enforcement." *United States v. Williams*, 553 U.S. 285, 304 (2008). "When speech is involved, rigorous adherence to those requirements is necessary to ensure that ambiguity does not chill protected speech." *Fox*, 567 U.S. at 253-54. Such laws must speak "with narrow specificity." *NAACP v. Button*, 371 U.S. 415, 433 (1963).

A.B.2013 fails to do so. It does not define the key terms "dataset" or "data point." A.B.2013 never explains whether they refer to each individual dataset developers might retrieve from a broad source (e.g., websites like Creative Commons) and each byte of information there, or only broad categories, like publicly available information on the Internet writ large. Plus, A.B.2013 provides no guidance on *how much* information must be disclosed to constitute a "high-level summary." Take subsection (a)(2), which requires a description of how the datasets further a model's "intended purpose." Cal. Civ. Code §3111(a)(2). A.B.2013 does not detail whether xAI must disclose its internal strategies regarding how it values individual datasets, or if it can simply state that such information improves AI models' effectiveness. Subsection (a)(5) likewise never states whether developers can provide a "yes" or "no" answer to whether its datasets include intellectual property. *Id*. §3111(a)(5). In a field as technically complicated and intricate as gen

AI, there is simply no way for xAI to know whether its "high-level summar[ies]" must be 100 words or 100 pages. A.B.2013 plainly lacks "narrow specificity." *Button*, 371 U.S. at 433.

Even if xAI could deduce what disclosures suffice under A.B.2013's enumerated categories, it would presumably need to provide more information to meet A.B.2013's "high-level summary" mandate. *See* Cal. Civ. Code §3111(a) (summary must "includ[e], but [is] not limited to," the enumerated list). That open-ended mandate "invite[s] arbitrary enforcement." *Kashem v. Barr*, 941 F.3d 358, 364 (9th Cir. 2019). Moreover, A.B.2013 does not even clearly identify which datasets it covers. Its operative provision requires developers to disclose information about datasets used to train gen AI models, *see* Cal. Civ. Code §3111—i.e., "testing, validating, or fine tuning" the model, *id.* §3110(f). But A.B.2013 also refers to "datasets used in the development" of such models, *id.* §3111(a), even though datasets used to *train* an AI model are far fewer than those used to *develop* that AI model. The boundaries of A.B.2013's disclosure requirements are not "clearly marked," as A.B.2013 leaves developers guessing whether they must provide information only about training datasets or the broader universe of datasets they might have sourced. *Baggett v. Bullitt*, 377 U.S. 360, 372 (1964). Against that uncertainty and threat of arbitrary enforcement, developers may feel compelled to disclose more information than A.B.2013 requires or eliminate certain datasets they fear that California may deem too "bias[ed]," thereby exacerbating A.B.2013's infringement on free speech—and xAI's trade secrets, to boot.

Simply put, A.B.2013 is plagued with vagueness problems at every turn. Due process demands more, especially when xAI's First Amendment and trade-secret rights are at stake. A.B.2013 is thus unconstitutionally vague.

II. The Equitable Factors Overwhelmingly Support Preliminary Relief.

The remaining preliminary injunction factors favor maintaining the status quo while this case is litigated to judgment. "It is well established that the deprivation of constitutional rights 'unquestionably constitutes irreparable injury.'" *Melendres v. Arpaio*, 695 F.3d 990, 1002 (9th Cir. 2012) (quoting *Elrod v. Burns*, 427 U.S. 347, 373 (1976)). Because xAI has shown that A.B.2013 likely violates the Constitution several times over, it will plainly "suffer[] irreparable harm" if California enforces A.B.2013 against xAI for failing to provide sufficiently detailed disclosures and deprives it of those constitutional rights—"no matter how brief the violation." *Baird v. Bonta*, 81 F.4th 1036, 1040 (9th Cir. 2023). That is especially true given that A.B.2013's enforcement in that manner would violate xAI's First Amendment rights. *See supra* pp.18-21; *Elrod*, 427 U.S. at 373.

That "[i]rreparable harm is relatively easy to establish in a First Amendment case," *CTIA – The Wireless Ass'n v. City of Berkeley*, 928 F.3d 832, 851 (9th Cir. 2019), underscores why the irreparable injury xAI faces is particularly acute. A.B.2013 compels xAI's speech regarding its own trade secrets. Those disclosures will not only destroy xAI's property interest in such information but also enable its competitors to gain an unfair advantage in this burgeoning, trillion-dollar industry—clear irreparable harm. *WeRide*, 379 F.Supp.3d at 853-54 ("It is well established that the loss of market position and the disclosure of trade secrets can constitute irreparable harm."). And if A.B.2013's obligations are as sweeping as they appear to be, xAI would need to devote substantial time and resources to locating, collecting, cataloguing, and summarizing the information related to its AI datasets—costs that only increase the older the model. *See* Stanley.Decl.¶¶27-29. Such substantial compliance costs, which xAI has no hope of recovering from the

state, *see Platt v. Moore*, 15 F.4th 895, 910 (9th Cir. 2021), further show that xAI will suffer substantial irreparable harm if A.B.2013 is not enjoined. *See E. Bay Sanctuary Covenant v. Biden*, 993 F.3d 640, 677 (9th Cir. 2021).

Because xAI has "establish[ed] a likelihood that [A.B.2013] violates the U.S. Constitution," it has "also established that both the public interest and the balance of the equities favor a preliminary injunction." *Ariz. Dream Act Coal. v. Brewer*, 757 F.3d 1053, 1069 (9th Cir. 2014). California "cannot reasonably assert that it is harmed in any legally cognizable sense by being enjoined from constitutional violations," as the public has no interest in allowing the state to enforce an unconstitutional law—especially one that violates three separate constitutional provisions. *Baird*, 81 F.4th at 1042 (quoting *Zepeda v. INS*, 753 F.2d 719, 727 (9th Cir. 1983)). Indeed, even apart from the law's constitutional infirmities, California has no discernible interest in forcing xAI to disclose its trade secrets, while failing to help consumers. *See supra* pp.17-18, 21. In short, the remaining preliminary-injunction factors overwhelmingly favor preliminarily enjoining A.B.2013.

CONCLUSION

The Court should grant xAI's motion.

Respectfully submitted,

s/Adam S. Sieff

ADAM S. SIEFF (CAL. BAR. # 302030)
DAVIS WRIGHT TREMAINE LLP
350 S. Grand Avenue, 27th Floor
Los Angeles, CA 90071
Tel.: (213) 633-8618
adamsieff@dwt.com

ERIN E. MURPHY (*pro hac vice* pending)
MATTHEW D. ROWEN (Cal. Bar #292292)
JAMES Y. XI (*pro hac vice* pending)
MITCHELL K. PALLAKI (*pro hac vice* pending)
ILAN J. POSNER (*pro hac vice* pending)
CLEMENT & MURPHY, PLLC
706 Duke Street
Alexandria, VA 22314
Tel.: (202) 742-8900
erin.murphy@clementmurphy.com
matthew.rowen@clementmurphy.com
james.xi@clementmurphy.com
mitchell.pallaki@clementmurphy.com
ilan.posner@clementmurphy.com

Attorneys for Plaintiff X.AI LLC

January 16, 2026

L.R. 11-6.2 CERTIFICATE OF COMPLIANCE

The undersigned, counsel of record for Plaintiff X.AI LLC, certifies that this brief contains 25 pages and 6,986 words, which complies with the page limit set by this Court's Standing Order and the word limit of L.R. 11-6.1.

Dated: January 16, 2026

s/Adam S. Sieff
ADAM S. SIEFF
DAVIS WRIGHT TREMAINE LLP
Counsel for Plaintiff X.AI LLC