

Scouting the battlefield: Will disputes over the interpretation of California's new law on AI data transparency render it meaningless?

By Jason Haas, Esq., Ervin Cohen & Jessup LLP

FEBRUARY 4, 2025

One of the most troubling features of generative AI models is their tendency to be a "black box," where governments and consumers lack any information about how the models work or were trained.

In September 2024, California's legislature took the first step to crack open that box. In its new law on "Artificial Intelligence Training Data Transparency," more commonly known as AB 2013, California imposes new disclosure requirements on developers.¹

While AB 2013 is relatively short, ambiguities in its text ensure significant litigation in coming years against developers of generative AI models and the businesses that use their models.

These rules take effect on January 1, 2026 and apply to all generative AI models made available to the California public after January 1, 2022. Unfortunately, the law is less than a model of clarity, leaving it uncertain whether the law will ever achieve its objective of improving consumer understanding of AI.

While AB 2013 is relatively short, ambiguities in its text ensure significant litigation in coming years against developers of generative AI ("GenAI") models and the businesses that use their models. Depending on the outcome of these conflicts, developers may ultimately be able to keep the black box closed and shield their AI models from significant scrutiny.

Yet it seems equally, if not more, likely that courts will resolve AB 2013's ambiguities in accord with the legislature's pro-transparency objectives, with significant negative implications for developers' competitive positions and their risks of future liability.

In the interim, developers and businesses using GenAI models must decide what disclosures, if any, to make as to their training data once the law goes into effect next year.

To provide some early reconnaissance, this article discusses the scope and purposes of AB 2013 and then identifies some of the key statutory terrain on which the upcoming battles over the meaning and application of AB 2013 will be fought.

A brief overview of AB 2013

On its face, AB 2013 presents a straight-forward list of information that developers of GenAI must disclose on their websites: a "high-level summary" of the datasets used to train their models, including twelve specific types of information.

Some of these twelve types, now codified at California Civil Code §3111(a), are the sources or owners of the datasets, a description of how the datasets "further the intended purpose" of the AI model or product, the general number and types of data in the datasets, and when the data was collected and first used to develop the model.

Whether the datasets contain protected intellectual property, consumer information or synthetic data, and whether the developer cleaned or processed the data (as well as the purpose for doing so) must also be disclosed.

The reach of AB 2013 is broad in some respects, but narrower in others. On one hand, it governs all AI systems and services that generate "derived synthetic content, such as text, images, video, and audio", regardless of the format of the output or the level of risk that the models' misuse might impose on others.²

Whether consumers are charged a fee for their use of the AI system or service does not matter.³ AB 2013 was, however, amended during the legislative process to exclude AI models using statistical, regression or other approaches. Only GenAI models are covered.

AB 2013 requires entities that "substantially modify" an AI system to disclose their datasets, so a company that fine tunes a GenAI model it obtained for use with its customers must also comply with AB 2013.

Yet AB 2013 does not apply to GenAI models developed only for use by corporate affiliates. AI systems and services used solely to improve computer security or the operation of aircraft are carved out

of the scope of AB 2013, as are ones intended for national security and defense purposes.

And while the text of AB 2013 does not limit the geographic location of developers, like all California laws, AB 2013 is subject to constitutional limits on California's ability to exercise personal jurisdiction over foreign AI developers.

The legislative purpose for AB 2013

The text of AB 2013 lacks an express statement of purpose, but a review of its legislative history reveals two primary objectives. Both aim to increase transparency into GenAI models.

First, AB 2013 seeks to improve the public's acceptance of AI, and its willingness and ability to use products and services incorporating GenAI. AB 2013's author observed that consumer confidence in AI lags well behind the rate of its adoption by business.

If greater transparency into training data can improve consumer understanding, the expectation is that this will increase consumer confidence in AI as well as improve consumers' ability to pick intelligently between models and to mitigate risks from using AI.⁴

Perhaps the most important issue to be resolved is the level of detail that a "developer" must provide under AB 2013 to constitute a "high-level summary of the datasets used" to train the AI system or service.

Second, AB 2013's disclosures are intended to increase visibility into potential problems with GenAI training data as well as expose developers' use of others' intellectual property and private information. The Legislature found there is "very little transparency" about training data and that the lack "hamstrings efforts to address and adequately identify many of the issues being raised by AI's rapid development."⁵

AB 2013 is supposed to help "identify and mitigate biases, address[] hallucinations and other problematic outputs, and shine[] the light on various other issues, such as privacy and copyright concerns."⁶ One issue of particular concern is the use of synthetic data, data generated by AI models rather than by humans. When synthetic data is used to train models, it can sometimes lead to poor performance, biased models and even "model collapse."⁷

AB 2013 can be enforced through private action

The California Chamber of Commerce sought to have AB 2013 expressly exclude a private right of action, but that effort was unsuccessful.⁸

Indeed, the legislative record expressly anticipates that enforcement will occur through private action, likely under California's unfair

competition law, California Business & Professions Code §17200 *et seq.*⁹ This means courts will have to address several major ambiguities in the text of the statute. The resolution of those ambiguities will determine the ultimate impact of California's new data transparency law.

The content and format of required disclosures

Perhaps the most important issue to be resolved is the level of detail that a "developer" must provide under AB 2013 to constitute a "high-level summary of the datasets used" to train the AI system or service.

Notably, this "high-level summary" standard was not part of the bill from its beginning. AB 2013 was amended to add it in place of a more burdensome "description of *each* dataset" standard. The Legislature thus recognized that a developer's unique selection of data sets might itself be proprietary and acknowledged that overly detailed disclosures may allow competitors to infer the precise data used in a model.¹⁰

Yet a single collective summary of all the datasets, which some developers may claim is all that is required by a "high-level summary of the datasets" standard, is arguably contrary to the purpose and text of AB 2013.

For example, can the first of the twelve mandated elements of AB 2013 — that "[t]he sources or owners of the datasets" be disclosed, California Civil Code §3111(a)(1) — be satisfied by a simple list of one or more sources, without any indication as to which entity owned or supplied each dataset?

If so, then very little information might have to be disclosed to comply with the statute, particularly if a developer obtained multiple datasets through a data broker. Knowing that data was obtained from a broker, without any understanding of *what* data was supplied, would not further the objectives of AB 2013. It would neither improve consumer confidence in AI nor allow consumers to intelligently assess whether to trust any resulting model.

Similarly, for California Civil Code §3111(a)(5), which addresses intellectual property, could a developer just disclose that one or more of the datasets used to train its model contained protected intellectual property, without disclosing which dataset contained that information or even what form of intellectual property was used?

Such a skeletal disclosure would not aid the government's regulation of AI nor allow private parties to assess whether any of their intellectual property may have been infringed.

To further cloak their models from scrutiny, developers might couch required disclosures in highly technical language. Such language might prevent any outsiders, except, perhaps, other AI experts, from even understanding what information is being disclosed.

Finally, despite AB 2013 listing twelve specific categories of required information, those categories may not cover the full range of necessary disclosures. AB 2013 states that the "high-level" summary of the datasets must include, but is expressly not limited to, the twelve categories. This language opens the door for plaintiffs

to argue that additional information must be disclosed and that a developer violated the statute by omitting it.

These are just a few examples, but a multitude of court battles over the content of the disclosures required by AB 2013 can be expected before the new law's full contours will become apparent.

The datasets to be disclosed

Perhaps inadvertently, an inconsistency exists in AB 2013 as to which datasets must be included in the "high-level summary." In California Civil Code §3111, the statute requires information to be disclosed regarding the "data used by the developer to *train*" the GenAI system or service, a term which the Legislature defined to include "testing, validating, or fine tuning" data.

A multitude of court battles over the content of the disclosures required by AB 2013 can be expected before the new law's full contours will become apparent.

But subpart (a) of §3111 requires the "high-level summary" to include datasets "used in the *development*" of the GenAI system, arguably imposing a broader scope of disclosure. With a model that has had multiple iterations, or has been trained on synthetic data, there could be multiple datasets that were used along the way in its "development," that were not used to train, test or validate the final deployed model.

This inconsistency will create fertile ground for new claims to be raised, particularly during litigation once a plaintiff has obtained some initial discovery.

The downstream application of AB 2013

AB 2013 defines a "developer" to include corporations, partnerships and other persons that "substantially modify" an AI system or service. California Civil Code §3110(b). Yet a "substantial modification" sweeps broadly to include any "new version, new release, or other update ... that materially changes [a model's] functionality or performance, including the results of retraining or fine tuning."¹¹

As such, any business that obtains a GenAI model from a tech company and then alters or refines the model could itself become subject to disclosure under AB 2013. For this reason, the additional requirement for "developer" status — that an AI system or service be "for use by members of the public,"¹² — can be expected to be heavily litigated.

To be sure, a chatbot or GenAI application that consumers use to generate text or video would clearly fall within the scope of this requirement. Conversely, a business that obtains a GenAI system only for internal purposes, such as finance or human resources, would likely not count as a developer.¹³

The ambiguity rests with GenAI models that a business might use internally but whose outputs or products are shared with consumers or that affect a business's treatment of its customers. If, for example, a business uses a GenAI system to generate forms or other documents it sends to consumers, would this render the business subject to AB 2013?

Or if an insurance company uses a GenAI system to develop a succinct summary of a lengthy claim file, which the company then uses in deciding whether to grant a consumer's claim, is the model itself "used by" the consumers for purposes of the statute?

Should the application of the statute depend on whether consumers have any direct access to the AI system, such as filling out an online claim form that is used by the model? These fact-intensive inquiries will be for the courts to decide.

The future evolution of AB 2013

As this discussion illustrates, there are substantial opportunities for developers to narrow the effective scope of AB 2013 or perhaps even water down its requirements to a point where no useful information is disclosed. Developers' claims that training data is proprietary may also sway some courts.

But the larger court battles will likely focus on the proper interpretation of AB 2013's text. There, the legislative record's clear pro-transparency language should tilt most courts towards mandating increased disclosures.

Developers' best legal arguments to counter this dynamic will be that the statute's language is clear, making the legislative record of AB 2013 irrelevant. "If the [statutory] language is clear and unambiguous there is no need for construction, nor is it necessary to resort to indicia of the intent of the Legislature. ..."¹⁴

For example, §3111(a)(7) requires the disclosure of "[w]hether the datasets include personal information" (as defined elsewhere in the code). A plain statement that the datasets used for a GenAI model contain personal information under the relevant definition might thus satisfy the letter of this provision, even though it would not reveal what kind of information was at issue, from where the information came, or even which data set contained that information.

These technical arguments may (or may not) convince many courts, but the battles over these provisions will likely determine the extent to which AB 2013 ultimately leads to greater disclosures about the training data for GenAI models used in California. Until then, the black box will remain closed.

Notes:

¹ 2024 Cal. Legis. Serv. Ch. 817 (A.B. 2013).

² See California Civil Code §§ 3110(c), 3111.

³ *Id.*

⁴ CA B. An., A.B. 2013 Assem., 4/30/2024 at 9.

⁵ CA B. An., A.B. 2013 Assem., 6/25/2024 at 1.

⁶ 2023 CA A.B. 2013 (NS) 8/20/2024 at 3.

⁷ CA B. An., A.B. 2013 Assem., 4/30/2024 at 8-9. A model can collapse when an AI model is repeatedly trained on synthetic data, as biases in the earlier data get amplified in successive iterations, eventually rendering the model useless.

⁸ Cf. CA B. An., A.B. 2013 Assem., 4/30/2024 at 10-11.

⁹ See *id.*

¹⁰ See CA B. An., A.B. 2013 Assem., 4/30/2024 at 12.

¹¹ California Civil Code §3110(d).

¹² California Civil Code §3110(b).

¹³ The business would, however, itself be a “member of the public,” requiring the developer from whom the model was obtained to comply with AB 2013, unless that developer was an affiliate of the purchaser.

¹⁴ *Lungren v. Deukmejian* (1988) 45 Cal.3d 727.

About the author



Jason Haas is a partner in the litigation department at **Ervin Cohen & Jessup LLP** in Beverly Hills, California. He has spent more than 20 years litigating a wide variety of complex business disputes for his clients in federal and state courts and in private arbitration. Haas focuses his practice on domestic and international litigation at the intersection of law, business and technology. He can be reached at Jhaas@ecjlaw.com.

This article was first published on Westlaw Today on February 4, 2025.