

JANUARY 16, 2026

# California's AB 2013 Takes Effect: Navigating AI Training Data Transparency and Trade Secret Risk

---

*Bethany P. Withers, Martha Schreck*

California's AB 2013, also known as the Generative Artificial Intelligence: Training Data Transparency Act (TDTA), took effect on January 1, 2026. In our June 2025 alert, "California's AB 2013: Generative AI Developers Must Show Their Data," we discussed the statute's core requirements and the challenges it presents for developers of generative artificial intelligence (AI) systems.

The TDTA requires AI developers to publicly post a "high-level summary" of the datasets used to train generative AI systems or services made available to the public since January 2022. The statute enumerates 12 categories of information that must be disclosed:

1. The sources or owners of the datasets.
2. A description of how the datasets further the intended purpose of the [AI] system or service.
3. The number of data points included in the datasets, which may be in general ranges, and with estimated figures for dynamic datasets.
4. A description of the types of data points within the datasets (including the types of labels used or general characteristics).
5. Whether the datasets include any data protected by copyright, trademark, or patent, or whether the datasets are entirely in the public domain.
6. Whether the datasets were purchased or licensed by the developer.
7. Whether the datasets include personal information, as defined in California Civil Code Section 1798.140(v).
8. Whether the datasets include aggregate consumer information, as defined in California Civil Code Section 1798.140(b).
9. Whether there was any cleaning, processing, or other modification to the datasets by the developer, including the intended purpose of those efforts in relation to the [AI] system or service.
0. The time period during which the data in the datasets were collected, including a notice if the data collection is ongoing.
1. The dates the datasets were first used during the development of the [AI] system or service.
2. Whether the generative [AI] system or service used or continuously uses synthetic data generation in its

## Transparency Obligations Versus Trade Secret Protection

One of the key challenges posed by the TDTA is how AI developers can provide disclosures sufficient to satisfy the statute without compromising valuable trade secrets. For many companies, the selection, composition, and use of training datasets, and the methods by which that data is cleaned, processed, and modified for training purposes, constitute confidential information that is the “secret sauce” or, more appropriately, the “secret recipe” of their proprietary AI models. AI developers invest significant resources and experimentation in selecting and using datasets in their development processes to seek competitive advantages over other AI models.

The statute does not define how much detail is required to meet the “high-level summary” standard, nor has there been official guidance that distinguishes compliant disclosure from the revelation of trade secrets. Industry groups and developers have raised concerns that requiring overly granular public disclosures could be anticompetitive, in that they would enable competitors to use such information to reverse engineer training strategies or replicate model development approaches using similar datasets.

## Scope of the TDTA and Affected Companies

The TDTA applies broadly to any company that designs, codes, produces, or creates a new version, release, or update that materially changes functionality or performance, including through retraining or fine-tuning, of a generative AI system or service for use by members of the public. Only a subset of exemptions to the statute are available (i.e., AI systems or services for which the “sole purpose is to help ensure security and integrity” or “the operation of aircraft in the national airspace” and systems or services “developed for national security, military or defense purposes made available only to federal” entities). As a result, the statute is seemingly not limited to developers of large “foundation” models such as OpenAI, Anthropic, Google, and Meta. Companies that leverage generative AI through internally developed systems or by materially modifying third-party models should carefully assess whether their activities arguably fall within the statute’s scope.

## Initial Compliance Signals From OpenAI and Anthropic

OpenAI and Anthropic have each posted new training data transparency materials in response to the TDTA going into effect, offering early guideposts for compliance. It appears that many AI developers have been waiting for such guideposts from the foundation developers, and they will likely base their approaches on these initial efforts.

OpenAI’s disclosure, “Training Data Summary Pursuant to California Civil Code Section 3111,” provides a short summary that touches on each of the 12 categories of information required to be disclosed under the statute. Anthropic’s “Training Data Documentation Pursuant to California Civil Code Section 3111 (AB 2013)” adopts a more structured, enumerated format, providing additional contextual explanations for each required item.

However, neither company identifies specific datasets used to train their models. Instead, each disclosure remains at a high level, disclosing only generalized categories of training data — including publicly available

information, nonpublic data obtained from third-party partners, data from users (subject to opt-out mechanisms) or human evaluators, and synthetic data. These disclosures appear designed to satisfy the statute's requirements (applying the statute's "high-level summary" requirement to each item of information) while avoiding any revelation of dataset-level details that could undermine trade secret protections.

With respect to intellectual property (IP) rights, OpenAI disclosed only that its training datasets include "data that may be protected by copyright," as well as "data in the public domain." Anthropic similarly stated that its training datasets include materials with varying IP statuses, including publicly available content that may be subject to third-party IP rights and some content in the public domain. Anthropic's disclosure provides information regarding its data acquisition methods for publicly available data, stating that it obtains publicly available internet content through a general-purpose web crawler, but OpenAI did not specify how it collects publicly available data. The generality of these disclosures reflects the difficulty in verifying the licensing status of third-party or open-source data, which is often obtained through publicly available sources but may nonetheless be subject to copyright or other IP rights. Anthropic further noted that its approach with respect to public content is consistent with what it characterizes as "standard industry practice" in training large language models.

Both Open AI and Anthropic disclosed that their training datasets may contain personal information and aggregate consumer information, as those terms are defined in California Civil Code Section 1798.140. OpenAI stated that it "take[s] steps to reduce the amount of" such information included in its training datasets. Anthropic provided additional context, noting that personal information is "incidentally" present in internet-sourced training data and used solely to enable models to learn and respond to language and that aggregate consumer information "is not used to identify or target individual consumers." Anthropic further disclosed that, when appropriate, it employs "tools and processes, including privacy-preserving analysis tools, to obfuscate sensitive data" and "post-training techniques" to "minimize the amount of personal information included in model outputs." Anthropic's more detailed disclosure, beyond the fact it uses such data, may be intended to mitigate potential privacy- and consumer protection-related concerns that could trigger broader scrutiny of its training data practices.

How other major AI developers will approach compliance (and how California intends to enforce the TDTA) remains to be seen. Notably, xAI has directly challenged the statute as unconstitutional in a recently filed lawsuit against the California attorney general. Among other claims, xAI alleges that the law compels disclosure of trade secrets in violation of the Fifth Amendment's Takings Clause. For more on that lawsuit, please read our recent client alert "xAI Challenges California's Training Data Transparency Act" (January 15, 2026).

## Looking Ahead

On the one hand, California (and other states) is likely to continue pushing for transparency about training datasets to identify and mitigate biases and cull AI developers' efforts to mine for data in ways that may violate privacy or intellectual property rights (e.g., copyright). On the other hand, AI developers appear unwilling to compromise their proprietary and competitively sensitive information, as indicated by OpenAI's and Anthropic's generalized disclosures and the lack of published documentation responsive to the TDTA by other AI developers to date. Absent official statutory guidance or federal regulation in the AI space, other AI

developers may look to early disclosures — or take a “wait and see” approach with respect to litigation developments — in establishing their own compliance approach to this law and other state (and international) AI transparency laws.

*This informational piece, which may be considered advertising under the ethical rules of certain jurisdictions, is provided on the understanding that it does not constitute the rendering of legal advice or other professional advice by Goodwin or its lawyers. Prior results do not guarantee similar outcomes.*

## CONTACTS

### **Bethany P. Withers**

Partner

Technology

[bwithers@goodwinlaw.com](mailto:bwithers@goodwinlaw.com)

### **Martha Schreck**

Associate

[mschreck@goodwinlaw.com](mailto:mschreck@goodwinlaw.com)