



[Home](#) [Licensing](#) [Patents](#) [Articles](#)

## Training Governance for Legal AI

by [Nick Clark](#) | Published March 27, 2026 | [PDF](#)

Legal AI models trained on case law corpora treat every judicial opinion as equal training signal. A Supreme Court majority opinion and a trial court dicta contribute to the model's legal knowledge without distinction. Training governance provides depth-selective gradient routing that encodes the legal authority hierarchy into the training process itself, ensuring the model learns binding precedent more deeply than persuasive authority, current law more deeply than overruled decisions, and holdings more deeply than dicta.

---

### The authority-flat training problem

Legal knowledge is inherently hierarchical. Supreme Court holdings bind all lower courts. Circuit court holdings bind district courts within the circuit. State supreme court holdings bind state lower courts. Persuasive authority from other jurisdictions carries weight but not binding force. Overruled

decisions are historically informative but legally invalid.

A model trained uniformly on all case law does not encode these distinctions. It learns patterns from all sources equally, and when asked to apply law to facts, it may cite overruled authority, weight persuasive authority as heavily as binding precedent, or treat dicta as holding. These are not hallucinations. They are structural consequences of authority-flat training.

## Why post-training citation checking does not fix authority confusion

Legal AI platforms add citation verification after generation: checking whether cited cases exist, whether they have been overruled, and whether the holding matches the proposition for which they are cited. This catches surface errors but does not address the underlying problem. The model's internal representation of legal knowledge does not encode authority hierarchies. Citation checking corrects outputs without correcting the knowledge structure that produces them.

A model that has learned overruled reasoning at the same depth as current law will consistently generate arguments that echo overruled reasoning, even when citation checking prevents it from citing the overruled case by name. The reasoning pattern persists because the training depth was equal.

## How training governance addresses legal AI

Training governance routes gradients based on legal authority metadata. Supreme Court holdings route to deep layers with full gradient magnitude. Circuit court holdings route to intermediate layers. Persuasive authority routes to surface layers with reduced gradient depth. Overruled decisions route with negative or near-zero gradients that prevent deep encoding while preserving historical awareness.

Jurisdictional scoping enables training specialized models. A model for New York practice routes New York Court of Appeals holdings to deep layers and treats other state supreme court holdings as persuasive authority at surface layers. The same training corpus produces different knowledge structures for different jurisdictional deployments.

Temporal governance ensures that the model's legal knowledge reflects current law. When a decision is overruled, the training governance mechanism can selectively reduce the influence of the overruled reasoning without retraining from scratch. Knowledge retention mechanisms preserve the doctrinal context while reducing the authority weight of superseded holdings.

Provenance tracing connects the model's legal reasoning to specific training sources. When the model generates a legal analysis, the provenance trace identifies which cases most influenced the reasoning, their authority level, and their current status. This trace enables attorneys to verify that the model's reasoning is grounded in valid, binding authority.

## What implementation looks like

A legal AI company deploying training governance annotates its case law corpus with authority metadata: court level, jurisdictional scope, current status, and holding versus dicta classification. The training pipeline routes gradients based on this metadata, producing models with authority-aware knowledge structures.

For law firm deployments, training governance enables jurisdiction-specific model variants that are deeply knowledgeable in the firm's practice jurisdictions while maintaining surface awareness of persuasive authority from other jurisdictions.

For legal education platforms, training governance produces models that can explain why certain authorities carry more weight than others, because the authority hierarchy is encoded in the model's knowledge structure rather than applied as a post-hoc filter.

[Training Governance All 21 steps →](#)

Govern what the model learns, at what depth, with what provenance.

Primary Technical Disclosure

◦ [Depth-Selective Training Governance for Machine Learning Systems](#)

Secondary Technical

◦ [Training Examples as Proposed Semantic Mutations](#) ◦ [Entropy-Band-Indexed Training Depth Profiles](#) ◦ [Depth-Selective Gradient Routing for Governed Training](#) ◦ [Training-Level Memorization Detection](#) ◦ [Differential Privacy Through Depth-Selective Routing](#) ◦ [Governed Fine-Tuning With Verifiable Provenance](#) ◦ [The Training Loop as a Governed Execution Environment](#) ◦ [Policy-Governed Knowledge Retention and Suppression](#) ◦ [Provenance-Traceable Training Dynamics](#) ◦ [Curriculum-Integrated Depth Scheduling](#) ◦ [Affect-Modulated Training Depth](#) ◦ [Training-Inference Governance Integration](#) ◦ [Training Governance for Human-Relatable Agents](#)

Applications (General)

◦ [Rights-Compliant Model Training Through Depth-Selective Routing](#) ◦ [Regulated Industry Model Governance With Provenance](#) ◦ [Training Governance for Medical AI](#) ◦ [Training Governance for Legal AI](#) ◦ [Training Governance for Financial Model Training](#) ◦ [Training Governance for Defense AI](#) ◦ [Training Governance for Educational AI Models](#) ◦ [Training Governance for Creative AI](#)

Applications (Specific)

◦ [OpenAI's Training Pipeline Has No Depth-Selective Governance](#) ◦ [Constitutional AI Training Lacks Depth-Selective Control](#) ◦ [Stable Diffusion's Training Has No Provenance Layer](#) ◦ [Midjourney Trains Aesthetics Without Governed Depth](#) ◦ [Scale AI Labels Data Without Governing What Models Learn](#) ◦ [Labelbox Manages Annotation Workflows, Not Learning Dynamics](#) ◦ [Snorkel AI Programs Labels but Does Not Govern Gradient Depth](#) ◦ [Weights & Biases Tracks Experiments, Not Learning Governance](#) ◦ [Determined AI Orchestrates Compute, Not Learning Depth](#) ◦ [MosaicML Optimizes Training Efficiency, Not Learning Governance](#)

[Training Governance overview →](#)

AQ

deterministic

autonomy

Legal

Subject to one or more pending U.S. and international patent applications, see [Patents](#) for the current list and status. No license, express or implied, is granted. Any use requires a separate written agreement—see [Licensing](#). Patent applications referenced on this site are pending. Claim scope, if any, is subject to examination and may issue in altered form or not at all. See [Legal](#) for terms and conditions.

Adaptive Query™ is a trademark of Nicholas Clark. U.S. federal registration is pending. federal registration. AQ™, AQ Inside™, Adaptive Index™, Adaptive Network™, Semantic Agent™, @AQ™, AQID™, and Adaptive Coin™ are used as trademarks in connection with the Adaptive Query platform and brand. Other names may be trademarks of their respective owners.

Platform operated by Adaptive Query LLC, which provides patent and trademark licensing services. Copyright © 2025-2026 Nicholas Clark. All rights reserved.

Last updated: 2026-03-03



- [Inventive Steps](#)
- [Licensing](#)
- [Patents](#)
- [Articles](#)
- [Legal](#)
- [Opportunities](#)
- [Sitemap](#)



- 
- [nick@qu3ry.net](mailto:nick@qu3ry.net)
- 72 28 14 36 01



[Invented by Nick Clark](#) | Founding Investors: Devin Wilkie