



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

Training Governance for Medical AI

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

Medical AI models are trained on clinical data that carries regulatory requirements, patient privacy constraints, and varying levels of clinical evidence quality. Current training pipelines treat all training data uniformly, learning from randomized controlled trials and case reports with the same depth. Training governance provides depth-selective gradient routing that governs what the model learns, at what depth, and with what provenance, enabling medical AI training that is auditable, evidence-weighted, and compliant with clinical data governance requirements.

The ungoverned training problem in medical AI

A medical AI model trained on a corpus of clinical literature absorbs everything uniformly. A finding from a large-scale randomized controlled trial and a finding from a single case report receive training signal proportional to their representation in the corpus, not proportional to their clinical evidence grade.

The model cannot distinguish between well-established clinical knowledge and preliminary findings because the training process does not encode this distinction.

For regulatory approval, this is problematic. A regulator reviewing a medical AI system needs to understand what the model learned, from which sources, and with what confidence. Current training pipelines can show the training data but cannot show the relationship between specific training examples and specific model behaviors. The training process is a statistical aggregate that resists decomposition into auditable components.

Why data curation is not training governance

Medical AI teams curate training data carefully: excluding low-quality sources, balancing representation across conditions, and removing personally identifiable information. Curation governs what enters the training pipeline. It does not govern how the training pipeline processes what enters it. A curated dataset still trains uniformly. The model learns from every included example with equal gradient depth.

The gap between data curation and training governance is the gap between selecting ingredients and governing the cooking process. Both matter. Current medical AI governance focuses almost exclusively on ingredient selection.

How training governance addresses medical AI

Training governance inserts depth-selective gradient routing into the training loop. Each training example carries metadata: evidence grade, source provenance, patient population, and clinical context. The gradient routing mechanism controls how deeply the example influences model parameters based on this metadata.

Findings from large-scale RCTs route gradients to deeper model layers, establishing foundational clinical knowledge. Findings from observational studies route to intermediate layers with moderated gradient magnitude. Case reports and preliminary findings route to surface layers with minimal gradient depth, informing pattern recognition without establishing deep clinical commitments.

Entropy-based training profiles detect when the model is memorizing specific clinical examples rather than learning generalizable patterns. Memorization detection prevents the model from encoding patient-specific details that should remain private, addressing a privacy risk that data de-identification alone does not fully eliminate.

Provenance tracing connects model behaviors to specific training examples. When the model produces a clinical recommendation, the provenance trace identifies which training examples most influenced that recommendation, the evidence grade of those examples, and the gradient depth at which they were learned. This trace provides the audit documentation that regulatory review requires.

What implementation looks like

A medical AI development team deploying training governance annotates training data with evidence grade, source provenance, and clinical context metadata. The training pipeline routes gradients based on this metadata, producing a model whose clinical knowledge is depth-stratified by evidence quality.

For FDA premarket review, training governance provides the documentation that demonstrates the model's clinical knowledge is grounded in high-quality evidence. The provenance trace shows regulators exactly how the model learned its clinical behaviors.

For ongoing post-market monitoring, training governance enables targeted model updates. When new clinical evidence emerges, the update can be routed to the appropriate depth, integrating new findings without destabilizing established clinical knowledge at deeper layers.

[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
- 72 28 14 36 01



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