Regulated Industry Model Governance With Provenance

by Nick Clark | Published March 27, 2026 | PDF

Financial services, healthcare, and pharmaceutical companies deploying AI models face regulatory requirements that current training processes cannot satisfy. Regulators require demonstrable knowledge of what data trained the model, how training was validated, and that the training process met regulatory standards. Training governance with structural provenance tracing provides this: a verifiable chain from training data through gradient updates to model parameters, enabling model certification backed by auditable training lineage.

The regulatory gap in model training

The Federal Reserve's SR 11-7 guidance on model risk management requires that financial institutions understand and document the data, assumptions, and limitations of every model they use. The EU AI Act requires documentation of training data, training methodology, and validation procedures for

high-risk AI systems. FDA guidance for AI in medical devices requires traceability from training data to model behavior.

Current training practices produce models that cannot satisfy these requirements. A foundation model trained on internet-scale data cannot document every training example. A fine-tuned model inherits the provenance gap of its base model. The model exists, and it performs, but the chain from training data to model behavior is opaque.

Why model cards and data sheets are insufficient

Model cards describe a model's intended use, performance characteristics, and known limitations. Datasheets for datasets describe training data composition and collection methodology. Both are documentation artifacts, not structural properties of the model. A model card can be inaccurate. A datasheet can be incomplete. Neither provides verifiable provenance that a regulator can independently audit.

The documentation describes the model. It does not structurally connect the model to its training lineage. A regulator reviewing a model card must trust that the card accurately describes the model. There is no mechanism for independent verification.

How training governance addresses this

Training governance maintains structural provenance through the entire training process. Each training example's contribution is tracked: which data entered the training loop, what governance profile it carried, which layers received gradient updates, and what the impact on model behavior was.

Entropy-depth profiles provide a structural characterization of how training data influenced the model at each layer. A regulator can examine the entropy-depth profile to understand whether the model's behavior in a specific domain is driven by deep integration of appropriate training data or by shallow pattern matching on potentially inappropriate data.

The governed training loop ensures that every training step complies with governance constraints. Data that should not be used for training in a specific domain is excluded structurally, not by filtering. Fine-tuning provenance tracks exactly what data was used for each fine-tuning run, enabling the institution to demonstrate that fine-tuning did not introduce data that violates regulatory requirements.

Knowledge retention monitoring detects when training causes the model to lose capability in previously validated domains, a common failure mode in fine-tuning. The governance layer can prevent training steps that would degrade validated capabilities, maintaining the model's certified performance characteristics.

What implementation looks like

A regulated institution deploying training governance wraps its training pipeline with a provenance layer that records every training step, every data contribution, and every governance constraint evaluation. The resulting model carries a verifiable training lineage that regulators can audit independently.

For pharmaceutical companies training models for drug discovery, training governance provides the data lineage documentation that FDA submission requires. Every training example's source, rights status, and integration depth is recorded and verifiable.

For financial institutions, training governance satisfies SR 11-7 model risk management requirements by providing the training documentation, validation evidence, and ongoing monitoring infrastructure that regulators expect, backed by structural provenance rather than self-reported documentation.

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

Last updated: 2026-03-03

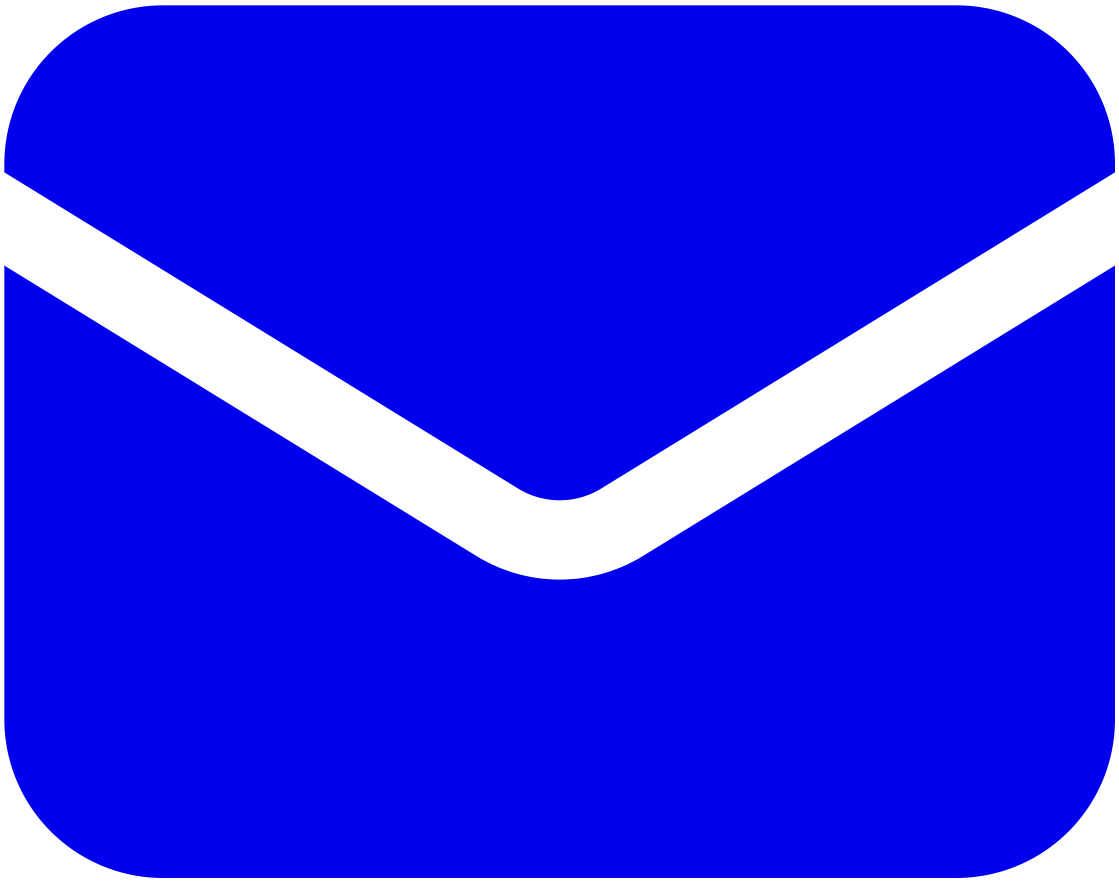

-
- nick@qu3ry.net
- 72 28 14 36 01

Invented by Nick Clark | Founding Investors: Devin Wilkie