Depth-Selective Gradient Routing for Governed Training

by Nick Clark | Published March 27, 2026 | PDF

Depth-selective gradient routing is the mechanism that enforces training depth profiles. Rather than allowing gradients from each training example to flow through all model layers, the routing system directs gradients only to layers authorized by the content's depth profile. Unauthorized layers receive zero gradient from that example, structurally preventing deeper integration than the governance policy permits.

What It Is

Gradient routing intercepts the standard backpropagation process and applies a layer mask derived from the training example's depth profile. Layers included in the mask receive gradients normally. Layers excluded from the mask receive zero gradient for that example. The routing operates at the individual training example level, not the batch level.

Why It Matters

Without gradient routing, depth profiles would be advisory: the system could specify desired depth but would have no mechanism to enforce it. Gradient routing provides the structural enforcement mechanism that makes depth profiling a guarantee rather than a guideline.

How It Works

During backpropagation for each training example, the routing system consults the depth profile associated with the example's content classification. It generates a binary layer mask that permits gradient flow through authorized layers and blocks it through unauthorized ones. The mask is applied before gradient accumulation, ensuring that the unauthorized layers' parameters are not affected.

The routing system records which layers received updates from which examples in the training lineage, creating a complete provenance record of parameter-level influence.

What It Enables

Gradient routing enables training governance that operates at the parameter level rather than the data level. It is not sufficient to filter training data; the system must also control how admitted data influences the model. Gradient routing provides this control with mathematical precision, ensuring that governance intent is structurally enforced during the optimization process itself.

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

Last updated: 2026-03-03

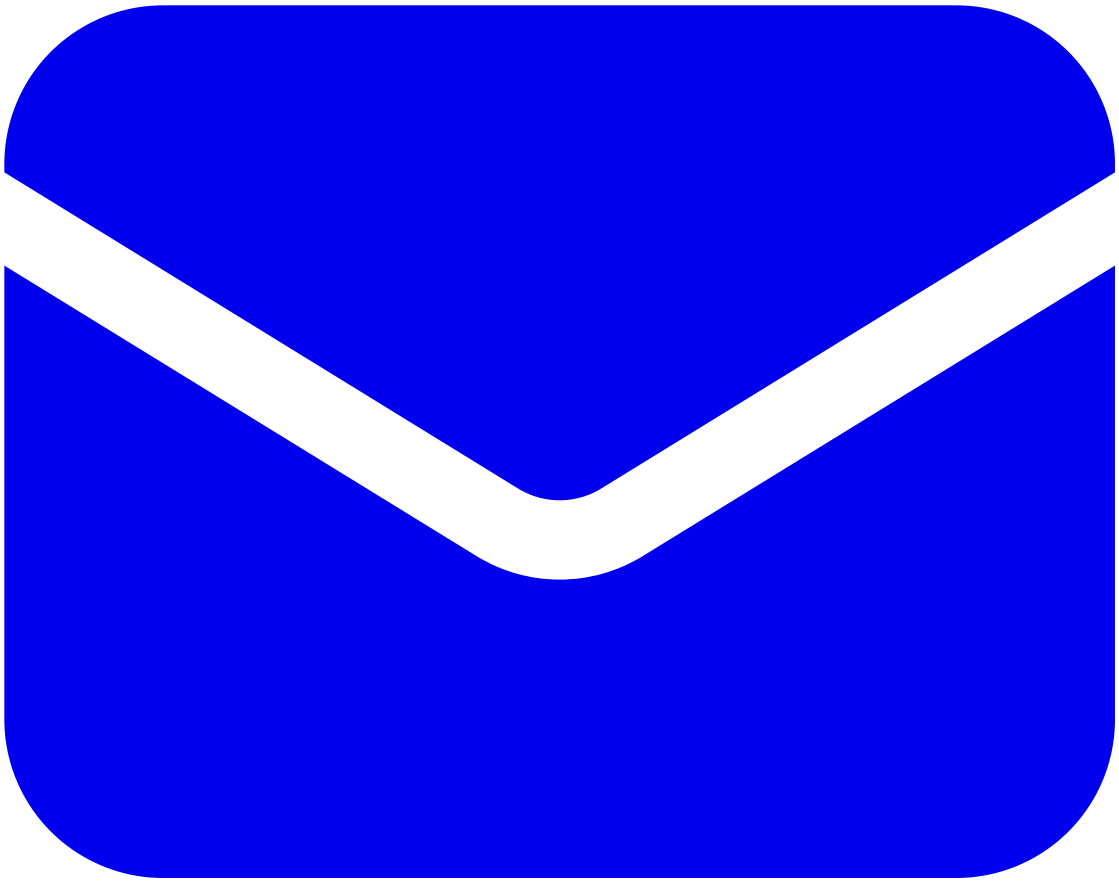

-
- nick@qu3ry.net
- 72 28 14 36 01

Invented by Nick Clark | Founding Investors: Devin Wilkie