



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

Entropy-Band-Indexed Training Depth Profiles

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

Not all content should integrate into a model at the same depth. Reference material may warrant shallow integration that informs without shaping core parameters. Foundational knowledge may warrant deep integration that shapes the model's fundamental representations. Entropy-band-indexed depth profiles govern how deeply each class of content is permitted to influence model parameters.

What It Is

Training depth profiles define how deeply content from each entropy band and content class is permitted to integrate into model parameters. Depth is measured in terms of which model layers receive gradient updates from each content class. Shallow integration restricts updates to upper layers. Deep integration permits updates through to foundational layers.

The profiles are indexed by both entropy band (reflecting content complexity) and content class (reflecting content type and provenance), creating a two-dimensional governance surface for training depth.

Why It Matters

Uniform training depth means that ephemeral web content and peer-reviewed research both modify the model's deepest representations equally. This is inappropriate. Depth profiling ensures that content influence on model parameters is proportional to content quality, permanence, and governance status.

How It Works

Each admitted training example is classified by its entropy band and content type. The depth profile for that classification specifies which model layers are eligible to receive gradient updates from this example. Gradient routing directs updates only to the eligible layers, blocking updates to restricted layers.

The depth profiles are policy objects, subject to versioning, audit, and governed modification. Changes to depth profiles affect future training but do not retroactively alter already-applied updates.

What It Enables

Depth profiling enables fine-grained control over how training content shapes model behavior. Verified factual content can integrate deeply, forming reliable knowledge representations. Opinion content can integrate shallowly, informing surface-level responses without shaping core reasoning. Rights-restricted content can be limited to influence patterns without memorization. This graduated approach produces models with more appropriate knowledge hierarchies.

[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