



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

## Training Examples as Proposed Semantic Mutations

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

In governed training, every training example is not simply fed to the optimizer. It is treated as a proposed semantic mutation to the model's parameters, subject to the same admissibility evaluation, policy compliance, and lineage recording that governs all mutations in the architecture. This transforms training from an ungoverned optimization process into a governed parameter evolution.

---

### What It Is

Training examples as proposed semantic mutations means that each data point presented during training is evaluated by the governance framework before its gradients are applied. The evaluation considers the data point's provenance, its content classification, its entropy characteristics, and its compliance with the training policy. Only data points that pass the admissibility evaluation contribute to model parameter updates.

## Why It Matters

Ungoverned training treats all data points equally. Poisoned data, biased samples, and rights-violating content all update model parameters without distinction. By treating each training example as a proposed mutation, the architecture applies the same governance rigor to training that it applies to all other semantic operations.

## How It Works

Before gradient computation, each training example is evaluated by the admissibility gate. The gate checks provenance (is the data point from an authorized source?), content classification (does the data point comply with training policy for its content type?), and entropy characteristics (does the data point fall within acceptable entropy bounds?). Only admitted data points proceed to gradient computation.

Rejected data points are recorded in the training lineage with rejection reasons, creating an auditable record of what the model was not trained on and why.

## What It Enables

Treating training examples as governed mutations enables training that is auditable, rights-compliant, and resistant to data poisoning. Every parameter change in the model can be traced back to the specific training examples that caused it, through which governance policies they were admitted, and with what provenance. This is the foundation for accountable machine learning.

[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