



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

## Weights & Biases Tracks Experiments, Not Learning Governance

by [Nick Clark](#) | Published March 28, 2026 | [PDF](#)

Weights & Biases provides experiment tracking, model versioning, dataset management, and hyperparameter optimization for machine learning teams. The platform records metrics, gradients, model checkpoints, and system performance throughout training runs. The observation is comprehensive. But observing training and governing training are structurally different operations. W&B records what happened during learning. It does not control what the model learns at what depth, which examples influence which representations, or whether the resulting knowledge is governed by policy. The gap is between tracking and governance.

---

**What Weights & Biases built**

W&B's experiment tracking captures loss curves, gradient distributions, activation patterns, model weights, hyperparameters, and custom metrics across training runs. The platform enables comparison across experiments, identification of which hyperparameter configurations produced the best results, and reproduction of successful training runs. The model registry tracks model versions with lineage from training data through to deployed artifacts.

The tracking is detailed and useful for understanding training dynamics after the fact. A researcher can examine why one training run outperformed another by comparing loss trajectories, gradient statistics, and hyperparameter choices. But the tracking is observational. It records the training process. It does not intervene in the training process to govern how learning occurs.

## The gap between experiment tracking and training governance

Experiment tracking answers the question: what happened during training? Training governance answers a different question: what should happen during training? Tracking is passive observation. Governance is active control. A training run tracked by W&B may show that gradients in certain layers became pathologically small or that memorization occurred in later epochs. The tracking records these events. Training governance prevents them.

W&B's hyperparameter sweeps represent a step toward governance: they systematically explore hyperparameter space to find configurations that produce better outcomes. But the governance operates at the hyperparameter level, not the learning dynamics level. Adjusting learning rate or batch size changes the global training dynamics. Depth-selective gradient routing controls which layers absorb influence from which training examples. The granularity of control is fundamentally different.

The provenance gap is also significant. W&B tracks which dataset version and which hyperparameters produced a model. It does not track which specific training examples influenced which specific learned representations. Provenance-traceable training governance maintains this finer-grained connection, enabling accountability at the representation level rather than the experiment level.

## What training governance enables for experiment tracking

With depth-selective training governance, W&B's tracking infrastructure gains governance semantics. Instead of passively recording gradient distributions, the governance layer actively routes gradients based on policy. The tracking records both the governance decisions and their effects, creating an auditable record of not just what happened but what was intended and whether the intention was achieved.

Entropy-based depth profiles provide a governance metric that W&B can track: the complexity state at each layer throughout training. When entropy at a layer deviates from the governance profile, the system intervenes rather than merely recording the deviation. The tracking becomes a governance dashboard rather than a historical log.

Provenance-traceable training extends W&B's experiment lineage to the representation level. Instead of tracking that this model was trained on this dataset with these hyperparameters, the system tracks that this specific learned behavior was influenced by these specific training examples through these specific gradient paths. The experiment tracking evolves from run-level provenance to representation-level provenance.

## The structural requirement

Weights & Biases solved experiment tracking and model versioning for machine learning teams. The structural gap is between observing training dynamics and governing them. Training governance provides depth-selective gradient routing that actively controls learning, entropy-based profiles that govern layer-level dynamics, and representation-level provenance that extends experiment tracking from run lineage to learned-behavior accountability.

[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