



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

Determined AI Orchestrates Compute, Not Learning Depth

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

Determined AI, now part of Hewlett Packard Enterprise, provides distributed training infrastructure that handles GPU cluster management, elastic resource allocation, fault-tolerant training, and adaptive hyperparameter search. The platform governs how compute resources serve the training process. But governing compute allocation and governing what the model learns at each layer are structurally different operations. The infrastructure ensures training runs efficiently. It does not ensure that learning occurs at the right depth, with the right provenance, under the right governance policies.

What Determined AI built

Determined AI's platform manages the operational complexity of distributed training. It handles GPU scheduling, elastic scaling when resources become available or are reclaimed, checkpoint management for fault tolerance, and distributed hyperparameter search across cluster resources. The platform

abstracts infrastructure complexity so researchers can focus on model development rather than cluster management.

The resource governance is sophisticated. The platform dynamically allocates GPUs, manages communication between distributed training processes, and recovers from hardware failures without losing training progress. Hyperparameter search runs are automatically scheduled and prioritized based on intermediate results. The infrastructure layer is well-governed. But the learning layer is not. The platform ensures the model trains on the allocated resources. What the model learns during that training is determined entirely by the training code, not by the infrastructure platform.

The gap between compute governance and learning governance

Compute governance determines which hardware runs which training job and how resources are shared. Learning governance determines which training examples influence which model layers and how that influence is tracked. The first is an infrastructure concern. The second is a model development concern. Current ML platforms govern the first and leave the second to the researcher's training code.

The gap becomes visible in distributed training. When training is distributed across multiple GPUs, gradient aggregation occurs across devices. The infrastructure ensures the aggregation is correct and efficient. But the governance question of whether certain gradient updates should be routed to specific layers based on their provenance or depth policy is not addressed. The infrastructure treats all gradients equally because its concern is communication efficiency, not learning governance.

Depth-selective training governance would operate within the distributed training pipeline. Gradient routing policies would be applied before or during the aggregation step. The infrastructure would not just communicate gradients efficiently but would route them according to governance policy. This requires the compute orchestration layer and the learning governance layer to be structurally integrated rather than independent.

What training governance enables for compute orchestration

With training governance integrated into distributed training infrastructure, Determined AI's compute orchestration gains learning-aware scheduling. Training jobs that require specific governance policies can be scheduled on resources configured to enforce those policies. Depth-selective gradient routing becomes a first-class infrastructure primitive rather than something implemented ad hoc in training code.

Fault tolerance gains governance semantics. When a training run recovers from a hardware failure, the governance layer verifies that the recovered state maintains provenance integrity. Checkpoints include not just model weights but governance state: which provenance chains are active, what depth policies are in effect, and what entropy profiles are expected at each layer.

The adaptive hyperparameter search becomes governance-aware. Instead of searching hyperparameter space based only on loss metrics, the search considers governance metrics: provenance coverage, depth profile conformance, and memorization detection. Training runs that achieve low loss through memorization are pruned in favor of runs that achieve governed learning.

The structural requirement

Determined AI solved distributed training compute orchestration. The structural gap is between governing compute resources and governing what models learn during training. Training governance provides depth-selective gradient routing integrated with distributed training infrastructure, governance-aware fault tolerance, and hyperparameter search that optimizes for governed learning rather than loss metrics alone.

[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