



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

Curriculum-Integrated Depth Scheduling

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

Training depth does not remain constant throughout the training process. Curriculum-integrated depth scheduling coordinates training depth profiles with the curriculum engine's progression stages. Early training may permit deep integration of foundational content while restricting advanced content to shallow layers. As the model progresses, depth profiles evolve to match the curriculum's progression.

What It Is

Curriculum-integrated depth scheduling links the training depth profiles to curriculum progression stages. At each stage, the depth profiles are configured to match the training objectives: early stages focus deep integration on foundational content, middle stages gradually deepen integration of more complex content, and advanced stages permit deep integration of specialized content.

Why It Matters

Training a model on all content at maximum depth from the beginning can produce unstable learning dynamics. Advanced concepts integrated before foundational knowledge is established may not integrate properly. Curriculum-depth scheduling ensures that the model builds knowledge in an appropriate order, with each layer of knowledge properly grounded before the next is introduced.

How It Works

The curriculum engine defines training stages with associated depth profile configurations. At each stage transition, the depth profiles are updated to reflect the new stage's requirements. The transition is governed: the model must demonstrate sufficient competence at the current stage before depth profiles are relaxed for the next stage.

Competence evaluation uses the same evidence-based gating mechanisms applied to skill gating in the agent architecture.

What It Enables

Curriculum-depth scheduling enables training processes that build knowledge systematically. Models trained under curriculum-depth governance develop more robust foundational representations and more appropriate knowledge hierarchies. The resulting models generalize better because their knowledge is built on properly established foundations rather than accumulated haphazardly.

[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