



Scaling Educational AI at the Edge: MLOps Strategies for Low-Code Platforms in Resource-Constrained Environments

By Lalitha Potharalanka

Jawaharlal Nehru Technological University

Conf42.com MLOps 2025

Agenda

Today we'll explore MLOps strategies specifically designed for educational environments with resource constraints, focusing on low-code platforms that democratise AI deployment.

O1

The Educational MLOps Landscape

Unique challenges faced by schools and current technology gaps

O2

Low-Code Platforms as Enablers

Empowering non-technical stakeholders in AI deployment

O3

MLOps Architectures for Education

Balancing deployment speed with governance requirements

O4

Implementation Strategies

From containerisation to federated learning in practice

O5

Key Takeaways & Future Directions

Actionable strategies to implement in your educational setting

The Educational MLOps Challenge

Implementing AI systems in educational institutions presents distinct challenges:

- **Infrastructure Disparity** – From well-connected urban environments to bandwidth-limited rural settings, infrastructure varies widely.
- **Technical Expertise Gaps** – Most educators lack the specialized DevOps or ML engineering backgrounds required for AI deployment.
- **Strict Governance & Privacy** – Handling student data necessitates heightened privacy measures and adherence to stringent regulations.
- **Resource Constraints** – Limited budgets often restrict investment in dedicated ML infrastructure.
- **Heterogeneous Device Ecosystems** – Diverse hardware configurations exist across school districts, complicating standardized deployments.

These unique factors demand a specialized approach to MLOps that differs significantly from typical enterprise implementations.



Data Distribution Challenges in Educational Settings

Seasonal Behavioral Patterns

Student interactions with educational platforms follow academic calendars, leading to cyclical data patterns. Without proper management, these patterns can trigger false model drift alerts.

Age-Based Distribution Shifts

As students progress through different year groups, their interaction patterns naturally evolve. This necessitates specialized monitoring to distinguish between natural data shifts and genuine model drift.

Limited Data Volume

Smaller rural schools generate significantly less data than their urban counterparts. This poses a challenge for traditional drift detection mechanisms that rely on large, statistically significant datasets.

Effective monitoring systems for educational ML models must account for these unique distribution patterns. This ensures genuine performance issues are detected while avoiding unnecessary retraining cycles.

Low-Code Platforms: Democratising Educational MLOps

Low-code platforms are transforming educational MLOps by lowering barriers and enabling wider AI adoption within schools:

- Empowering educators to participate directly in model deployment without extensive coding expertise.
- Offering intuitive visual interfaces for configuring inference parameters and monitoring thresholds.
- Automating complex MLOps workflows through user-friendly interfaces.
- Integrating governance guardrails directly into the platform architecture.
- Significantly reducing the technical barrier to entry for AI implementation.
- Fostering seamless collaboration between technical and pedagogical teams.

This democratisation is crucial for scaling AI across educational settings, particularly where dedicated ML engineers are scarce or unavailable.

Privacy-Preserving MLOps for Educational Data



Data Minimisation Pipelines

Automated workflows that preprocess and reduce data, ensuring only essential information is used for model training.



Differential Privacy Implementation

Adding calibrated noise to training data to prevent individual student identification while preserving aggregate insights for model development.



Federated Learning Architectures

Training models locally on school devices, sharing only model updates instead of raw student data across the district network.



On-Device Inference Prioritisation

Implementing edge deployment strategies that prioritize keeping sensitive predictions and student data local to school infrastructure.

These privacy-preserving approaches are not optional in education; they are fundamental requirements that must be integrated into MLOps workflows from the outset.

Automated Bias Detection in Educational AI

Ensuring fairness and equity in educational AI systems is paramount. This requires specialized bias detection pipelines that address unique concerns, focusing on:

- **Multi-dimensional fairness metrics:** Monitoring model performance across crucial dimensions such as socioeconomic status, linguistic background, and accessibility needs.
- **Age-appropriate evaluation frameworks:** Developing and applying bias assessment methods that account for different developmental stages of learners.
- **Cultural context awareness:** Implementing region-specific bias detection mechanisms that consider local educational norms and cultural nuances.
- **CI/CD integration:** Embedding automated testing gates within continuous integration/continuous deployment pipelines to proactively prevent the deployment of biased models.

Low-code platforms are crucial for seamlessly integrating these essential checks directly into AI deployment workflows, thereby making robust bias detection accessible even to non-technical stakeholders.

MLOps Architecture for Educational Environments

An adaptable MLOps architecture is crucial for AI in diverse educational settings. It balances central oversight with local autonomy, accounting for varying connectivity, expertise, and strict data privacy requirements, moving beyond traditional cloud-centric models to a more distributed approach.



Central Training Infrastructure

District-level model development, integrating privacy-preserving techniques and automated governance, ensures secure data aggregation, robust model versioning, and compliance.



Regional Deployment Hubs

These distribution nodes stage models and manage deployments across school clusters, optimizing delivery for local network conditions, caching for quicker access, and simplifying updates.



School-Level Inference

Edge computing infrastructure enables offline capabilities and local data processing, with AI models executing directly on school devices to provide real-time feedback and prioritize data privacy.

This distributed architecture balances centralized governance with local adaptation, ensuring greater scalability, enhanced data security, and reliable, equitable AI services. By segmenting responsibilities, it facilitates efficient resource allocation and maintenance, supporting diverse educational needs while safeguarding student data.

Containerised Model Serving for Offline Capability



Offline capabilities are essential in educational environments where connectivity can be unreliable. Key implementation strategies include:

- Lightweight container images optimised for educational hardware
- Progressive model loading that prioritises core functionality
- Intelligent synchronisation protocols for intermittent connectivity
- Local data caching with privacy-preserving encryption
- Graceful degradation pathways when full model functionality is not available

These approaches ensure AI educational tools remain operational even when internet access is limited or unavailable.

Automated Retraining Pipelines for Educational Models

Scheduled Data Collection

Privacy-compliant aggregation of model performance data from school infrastructure.

Graduated Deployment

Phased rollout, starting with pilot schools before district-wide implementation.



Distribution Analysis

Education-specific drift detection, accounting for academic calendar effects.

Contextual Retraining

Selective model updates targeting only affected components or student segments.

Multi-dimensional Validation

Comprehensive testing across various educational contexts and student demographics.


This cyclical process ensures models remain accurate while respecting the unique operational constraints of educational environments.

Federated Learning in Practice: District-Wide Collaboration

Federated learning enables powerful cross-school collaboration while maintaining data privacy:

- **Local model training** on school-specific devices using only local student data
- **Encrypted parameter sharing** between schools via district coordination servers
- **Differential privacy integration** to prevent reverse engineering of student information
- **Adaptive aggregation strategies** that account for varying school sizes and demographics
- **Implementation via low-code interfaces** that abstract away federated learning complexity

This approach has enabled district-wide model improvements without centralizing sensitive student data, addressing both privacy regulations and bandwidth constraints.



Monitoring Strategies for Heterogeneous Device Ecosystems

Resource-Aware Metrics Collection

Adaptive telemetry systems that adjust data collection based on device capabilities, ensuring comprehensive metrics on high-end hardware and lightweight monitoring on resource-constrained devices.

Device-Specific Performance Baselines

Tailored performance baselines that define appropriate thresholds for each hardware profile, eliminating false alerts and ensuring accurate evaluations across diverse school equipment.

Centralized Monitoring Dashboards

Intuitive low-code dashboards that transform complex monitoring data into educator-friendly visualizations, providing actionable insights without requiring technical expertise for interpretation.

Automated Intervention Workflows

Proactive intervention workflows that empower non-technical staff to address detected issues, enabling immediate actions like model rollbacks or dynamic parameter adjustments.

Implementation Roadmap for Educational Institutions

1 Phase 1: Foundation

Lay the groundwork by establishing low-code MLOps platform infrastructure and robust privacy frameworks, coupled with comprehensive training for technical and educational stakeholders.

2 Phase 2: Pilot Deployment

Pilot containerized model serving in diverse schools to assess adaptability across varying resource profiles. Concurrently, establish foundational monitoring baselines and refine governance workflows.

3 Phase 3: Scaled Rollout

Execute a district-wide deployment, integrating localized adaptations. Integrate federated learning to foster collaborative model improvement across the district.

4 Phase 4: Optimization

Continuously refine automated retraining pipelines based on real-world performance data. Develop bespoke monitoring solutions tailored to education-specific metrics and insights.

This phased approach strategically balances immediate value realization with sustainable, long-term operationalization.

Key Results from Educational MLOps Implementations

- **Offline Reliability**

AI functionality maintained during connectivity disruptions

- **Resource Utilisation**

Computing resource reduction via edge-optimized deployment

- **Non-Technical Engagement**

Increased educator participation in AI governance through low-code interfaces

Implementation Impact:

- Successful deployment across districts serving millions of students with diverse infrastructure
- Maintained model performance across bandwidth-constrained environments
- Enabled educator-driven model refinement without requiring technical intervention
- Achieved compliance with educational data privacy regulations across jurisdictions
- Reduced technical support requirements through self-healing deployment architectures

Key Takeaways



Low-Code Empowers Educators

Intuitive low-code platforms remove technical barriers, enabling broader AI adoption and ensuring robust governance within educational settings.



Privacy-First by Design

Educational MLOps demands privacy-preserving architectures, integrating solutions like federated learning and differential privacy directly into deployment workflows.



Offline Functionality is Key

Containerized serving with local inference ensures AI tools perform reliably, even in bandwidth-constrained school environments without consistent internet access.

Thank You !