

A Constraint-First Framework for Ethical Location-Based Recommender Systems in Technology-Enhanced Learning

Madjid Sadallah, Marie Lefevre

▶ **To cite this version:**

Madjid Sadallah, Marie Lefevre. A Constraint-First Framework for Ethical Location-Based Recommender Systems in Technology-Enhanced Learning. 2026. ⟨hal-05716768⟩

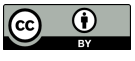
HAL Id: hal-05716768

https://hal.science/hal-05716768v1

Preprint submitted on 17 Aug 2026

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY 4.0 - Attribution - International License

A Constraint-First Framework for Ethical Location-Based Recommender Systems in Technology-Enhanced Learning

Madjid Sadallah and Marie Lefevre

Université Claude Bernard Lyon 1, CNRS, INSA Lyon, LIRIS UMR5205, 69622 Villeurbanne, France

Abstract

Purpose – Location-Based Recommender Systems (LBRS) increasingly mediate learning in physical environments. However, their commercial optimization logics—prioritizing engagement, efficiency, and data extraction—introduce pedagogical, ethical, and accessibility risks. This paper advances a constraint-first design principle ensuring that educational validity, equity, and learner protection structurally precede algorithmic optimization in technology-enhanced learning.

Design/methodology/approach – The paper adopts a critical-analytical, theory-grounded approach. Drawing on situated, embodied, and connectivist learning theories, as well as Universal Design for Learning, it interrogates dominant recommender architectures and their implicit value hierarchies. Paradigmatic analysis of educational and commercial systems identifies structural misalignments, informing normative framework synthesis.

Findings – Five systemic tensions are identified: directive scripting versus learner agency; proximity optimization versus pedagogical relevance; algorithmic opacity versus pedagogical explainability; popularity bias versus equitable access; and surveillance logics versus learner protection. In response, the paper formalizes a constraint-first processing hierarchy in which pedagogical validity, accessibility, fairness, and transparency operate as non-compensatory architectural gates defining the admissible recommendation space prior to personalization and ranking.

Originality/value – The paper advances an implementable architectural hypothesis translating educational values into system invariants. It demonstrates how ethical and pedagogical commitments can be embedded by construction—rather than mitigated post hoc—while outlining design and governance implications for educational LBRS deployment.

Keywords: Location-based recommender systems; technology-enhanced learning; algorithmic fairness; Universal Design for Learning; AI in education; educational technology ethics

1. Introduction

Educational recommender systems personalize resources to learners' needs (Drachsler et al., 2015) but often neglect the physical, social, and cultural contexts essential to meaningful learning (Brown *et al.*, 1989; Lave and Wenger, 1991). This limitation is particularly salient in physical learning environments—laboratories, museums, field sites, vocational contexts—where location shapes learning practices and equitable access varies across populations.

Location-based recommender systems (LBRS) address this gap by integrating geospatial context into recommendation processes. Yet most LBRS stem from commercial domains prioritizing engagement over learning outcomes, equity, or ethics (Ekstrand *et al.*, 2022). When transposed into educational settings, these reveal mismatches: sparse interaction data, strict privacy requirements for minors, and algorithmic logics privileging proximity over pedagogically meaningful learning—systematically disadvantaging learners constrained by disability, socioeconomic status, or geographic isolation.

These mismatches are not incidental but symptomatic of deeper architectural tensions in AI-enhanced learning. As the field moves, in Cingillioglu and Barbieri's (2025) terms, "from admiration to apprehension," research demonstrates that engagement-driven optimization may fundamentally conflict with pedagogical coherence, learner agency, and equity (Ofosu-Asare, 2025; Zeb et al., 2024). LBRS exemplify this conflict: while commercial systems optimize for dwell time and proximity, educational deployments must foreground accessibility and ethical safeguards. Current educational recommender systems continue to prioritize predictive accuracy over pedagogical effectiveness (da Silva *et al.*, 2023), revealing a persistent gap between algorithmic optimization and educational values.

This gap reveals a core design challenge: how to architecturally embed pedagogical and ethical commitments. While Value-Sensitive Design provides conceptual frameworks (Friedman & Hendry, 2019) and AI ethics guidelines offer principles (Holmes *et al.*, 2022), few frameworks specify how pedagogical theory should structurally constrain algorithmic design. Drawing on critical educational technology scholarship (Selwyn, 2016), we argue system architectures constitute theoretical hypotheses encoding values—yet this encoding typically occurs implicitly through processing order rather than deliberate pedagogical grounding.

We address these questions by introducing the *constraint-first principle*: pedagogical validity and ethical safeguards function as non-compensatory gates delimiting the admissible recommendation space prior to

optimization. This principle inverts conventional engagement-first design, where statistical optimization defines solution spaces and pedagogical constraints subsequently filter outputs. The inversion is consequential: in engagement-first systems, pedagogically invalid locations may still be recommended if sufficiently popular; in constraint-first systems, pedagogical invalidity constitutes architectural inadmissibility—no optimization objective can override foundational commitments. We operationalize this principle through a five-component processing hierarchy enabling empirical evaluation of whether architectural ordering better preserves learning integrity, equity, and learner agency.

This paper advances educational LBRS design through four interconnected contributions. First, we identify five tensions arising when commercial optimization logics transpose into educational contexts, demonstrating these are systematic consequences of architectural ordering rather than incidental flaws. Second, we formalize the constraint-first principle as a computational ordering rule where pedagogical validity and ethical safeguards operate as non-compensatory gates, making explicit how processing order embodies political choices about values shaping technology (Winner, 1980). Third, we operationalize this principle through a five-component processing hierarchy that enforces pedagogical coherence, accessibility, fairness, and transparency as structural invariants rather than adjustable parameters. Fourth, we specify an empirical validation agenda translating architectural hypotheses into testable research questions, addressing the gap between ethical aspirations and empirical validation identified in recent reviews (Hardaker and Glenn, 2025).

2. Theoretical Foundations and Paradigm Analysis

2.1 Location as Learning Medium

Location operates as an active pedagogical medium shaping inquiry, practice, and social interaction. Here, *pedagogical relevance* extends beyond engagement to denote the degree to which a location affords learning aligned with authentic objectives and situated practices.

Situated cognition posits that knowledge is inseparable from the contexts and practices in which it develops (Brown *et al.*, 1989; Lave and Wenger, 1991). Pedagogically-informed LBRS should scaffold *legitimate peripheral participation*, supporting learners' gradual integration into communities of practice while aligning with sociocultural (Vygotsky, 1978) and ecological (Gibson, 1979) frameworks of affordances. Recent work applying activity theory to AI-enhanced workplace learning demonstrates how technological mediation shapes tools, contexts, and processes (Martinez Marroquin and Senadji, 2025).

Extending these insights to LBRS foregrounds the need for activity-centered recommendations specifying embodied, meaningful actions ("measure sediment flow") rather than passive visitation.

Embodied cognition situates cognitive processes in sensorimotor experience (Wilson, 2002), whereby physical interaction scaffolds conceptual understanding (Abrahamson & Lindgren, 2014). Universal Design for Learning (UDL) principles (Rose & Meyer, 2002) prescribe multiple means of engagement, representation, and action from the outset, ensuring accessibility across diverse sensory and motor profiles (Burgstahler, 2021; Edyburn, 2021). Connectivist perspectives frame locations as nodes in dynamic knowledge networks, where LBRS scaffold connections to resources and peer communities (Siemens, 2005).

From these foundations, we formalize three architectural requirements for LBRS:

1. *Action Primacy*. Recommendations must specify embodied verbs ("measure," "manipulate," "observe," "construct") rather than passive visitation.

2. *Multimodal Affordances.* Locations must offer tactile, kinesthetic, and spatial manipulation opportunities accessible across sensory and motor profiles, e.g., geological sites with tactile rock samples and audio-described formations.

3. *Concept Grounding.* Activities must explicitly link bodily experience to curricular abstractions, such as manipulating rock strata before introducing geological time scales.

These requirements act as non-compensatory pedagogical gates: accessibility and embodiment are structural prerequisites—no pedagogical benefit can justify inaccessibility, and no optimization metric can override embodiment requirements. This non-compensability reflects a fundamental principle: educational values must constrain, not merely influence, algorithmic design.

2.2 Paradigm Analysis and Architectural Implications

Educational LBRS operate under constraints largely absent in commercial systems. Privacy obligations—particularly for minors under COPPA and GDPR Article 8—necessitate privacy-by-design strategies (Jones, 2019), while pedagogical requirements often exceed simple metadata descriptors. Ethical frameworks must be operationalized as architectural constraints rather than aspirational guidance, motivating our *constraint-first hierarchy*, where pedagogical validity gates precede personalization.

Analysis of four key paradigms reveals critical trade-offs: Collaborative Filtering (CF) suffers from data sparsity and popularity bias; Content-Based Filtering (CBF) mitigates cold-start but requires rich ontologies; Knowledge-Based Systems (KBS) offer rule-based transparency at the cost of scalability; and Deep Learning (DL) excels at pattern recognition but introduces opacity (Resnick & Varian, 1997; Pazzani & Billsus, 2007; Burke, 2000; Feng et al., 2018). No single paradigm addresses educational requirements comprehensively. CF and DL introduce fairness and opacity risks; CBF and KBS demand extensive knowledge engineering.

These observations motivate *hybrid architectures* assigning complementary roles:

- *Collaborative Filtering* provides social enrichment supporting underutilized locations but risks amplifying popularity bias and inequitable exposure.
- *Content-Based Filtering* enables goal-location matching with diversity mechanisms but may over-specialize and neglect cross-disciplinary goals.
- *Knowledge-Based Systems* enforce pedagogical and safety constraints with rule-based transparency but face scalability challenges.
- *Deep Learning* enables pattern recognition but introduces opacity and may amplify biases.

The constraint-first architecture assigns each paradigm a specific role in the five-component processing hierarchy (Section 4.2).

Historical educational LBRS (e.g., *EcoMOBILE*, Kamarainen *et al.*, 2013; *ARIS*, Holden and Sykes, 2011) illustrate authentic-context learning potential but predate fairness-aware ranking, privacy-by-design for minors, or contestability mechanisms. While general AI ethics frameworks offer domain-agnostic guidance, operationalization for location-based educational systems remains limited. Privacy frameworks from Educational Data Mining (Prinsloo and Slade, 2017) emphasize governance but rarely encode pedagogical constraints architecturally.

3. Methodology

This work adopts a critical-analytical approach, integrating theory-driven framework development with focused engagement with the literature. Our methodology emphasizes paradigm critique, illustrative case

construction, and normative framework synthesis to translate educational theory into concrete architectural principles for LBRS.

We situate this work within the design science tradition (Hevner *et al.*, 2004), framing it as artifact-oriented conceptual research. Rather than evaluating existing systems empirically, we articulate prescriptive design theory, specifying *how systems should be designed* based on theoretical imperatives. Our epistemological stance is critical constructivist: architectures are not neutral artifacts but encode political, ethical, and pedagogical choices that can be deliberately engineered. Building on Winner's (1980) insight that *artifacts have politics* and on Friedman and Hendry's (2019) Value-Sensitive Design, we operationalize values as computational orderings, foregrounding processing sequences as loci of value inscription that instantiate pedagogical and ethical commitments.

Our critical engagement synthesized three domains: (1) spatial and embodied learning theories (e.g., Brown et al., 1989; Lave & Wenger, 1991); (2) recommender system paradigms and their limitations (e.g., Ekstrand et al., 2022); and (3) educational technology ethics (e.g., Friedman & Hendry, 2019). This synthesis enabled the identification of recurring architectural tensions to guide prescriptive design.

Following Flyvbjerg (2006), we constructed illustrative cases exemplifying tensions between commercial optimization logics and educational imperatives. Cases were derived iteratively through three stages: identifying core values, mapping commercial priorities, and analyzing resulting misalignments. These paradigmatic cases are conceptual tools that synthesize theoretical commitments with documented challenges in LBRS, guiding prescriptive framework development. This analysis reveals what we term architectural causality: the principle that system architectures encode causal relationships between design choices and normative outcomes. Processing order—which computational operations precede others—determines whether educational values function as governing constraints or compensable objectives.

The constraint-first principle emerged through comparative conceptual analysis. Examination of existing LBRS—commercial platforms (Foursquare, Yelp, Google Maps) and educational deployments (e.g., *EcoMOBILE*, Kamarainen *et al.*, 2013)—revealed a consistent pattern: optimization often defines candidate solution spaces before educational validation, producing pedagogically invalid or inequitable recommendations. We formalized this as an ordering rule: pedagogical and ethical constraints → optimization, operationalized through five processing components aligned with paradigm strengths.

Iterative refinement ensured each component addresses structural tensions, aligns with established learning theories, and preserves feasibility under practical and ethical constraints.

Validation proceeds through theoretical coherence and derivation of testable hypotheses. Internal consistency ensures components interlock according to the constraint-first principle. Theoretical grounding anchors architectural decisions in established learning and ethics theories. Empirical derivation translates architectural propositions into hypotheses suitable for future testing (Section 5.3). As a normative conceptual framework, ultimate validation lies in future implementation and empirical comparison, which extends beyond the scope of this work.

4. Structural Tensions and Normative Framework

4.1 Architectural Tensions in Educational LBRS

Following Flyvbjerg's (2006) phronetic case methodology, we construct five paradigmatic cases exposing systemic contradictions—*architectural tensions*—emerging when commercial LBRS optimization logics are transposed into technology-enhanced learning infrastructures.

Tension 1: Directive Scripting versus Learner Agency

Optimization-first LBRS operationalize quality as convergence toward singular "best" trajectories, producing directive spatial scripts where learners follow fixed paths. Self-Determination Theory establishes that autonomy support underpins intrinsic motivation through meaningful choice, transparency, and contestability (Deci & Ryan, 2000; Reeve & Cheon, 2021). When convergence optimization precedes pedagogical path representation, learner autonomy is computationally reframed as inefficiency—an artifact of optimization ordering. From this tension, we derive three architectural requirements: (1) *perceived choice*—multiple pedagogically valid options must remain available; (2) *rationale transparency*—recommendations must be explainable in relation to learning objectives; and (3) *contestability*—learners and educators must be able to override system suggestions without penalty.

Tension 2: Proximity Optimization versus Pedagogical Relevance

Commercial LBRS treat proximity as relevance, collapsing pedagogical meaning into distance and ignoring accessibility or curricular alignment. Lower-income neighborhoods are systematically farther from cultural institutions (Eisenlohr et al., 2025), while accessibility barriers render proximate sites functionally inaccessible (Papadimitriou, 2016). Virtual alternatives sometimes yield superior outcomes

compared to proximate but pedagogically impoverished sites (Zhao et al., 2022). Universal Design for Learning frameworks require educational resources provide multiple means of engagement, representation, and action from the outset (Burgstahler, 2021; Edyburn, 2021). When ranking prioritizes proximity *before* evaluating pedagogical affordances, accessibility and disciplinary meaning are systematically overridden by logistical convenience.

Tension 3: Algorithmic Opacity versus Pedagogical Explainability

Deep learning models extract interpretable signals from mobility data through attention mechanisms yet surface statistical regularities rather than pedagogical rationales. When pattern recognition precedes pedagogical constraints, explainability becomes anchored in behavioral data rather than curricular intent. Zeb et al. (2024) identify opacity as a core ethical concern in ChatGPT deployment. This concern extends to LBRS, where statistical regularities may obscure pedagogical rationales. A system may recommend a museum because similar learners visited it after specific modules, yet cannot determine whether those visits advanced learning objectives or reflected incidental convenience. Pattern-based explanations ("users like you") fail to address pedagogical accountability ("does this advance competency X?"). Ethical AI in education requires explainability aligned with pedagogical values, enabling auditability and contestation (Holmes *et al.*, 2022; Holstein and Doroudi, 2021; Khosravi *et al.*, 2022). We therefore derive two requirements: (1) *pedagogical grounding*—explanations must reference learning objectives rather than correlations; and (2) *contestable transparency*—educators must be able to interrogate and override pattern-based outputs.

Tension 4: Popularity Bias versus Equitable Access

Popularity bias in commercial LBRS amplifies frequently visited locations (Ekstrand et al., 2022). In educational contexts, structural inequalities transform absence into algorithmic invisibility. When interaction frequency drives ranking before pedagogical validation, underutilized locations are systematically deprioritized. A community laboratory in a lower-income area may offer high-value learning yet remain under-visited; collaborative filtering suppresses its visibility, turning structural constraint into inferred irrelevance. ML systems trained on biased data risk amplifying inequities (Baker and Hawn, 2022; Doroudi et al., 2019). We derive two requirements: (1) *equity-aware ranking*—exposure must be calibrated to ensure undervisited locations receive equitable consideration; and (2) *barrier*

detection—systems must identify and compensate for access constraints rather than interpret absence as lack of pedagogical value.

Tension 5: Surveillance versus Learner Protection

Commercial LBRS justify extensive location tracking for marginal accuracy gains. In learning contexts involving minors, this creates disproportionate privacy risks (Acquisti et al., 2015). When optimization precedes necessity, data collection is incentivized for statistical performance rather than educational function. Tracking a learner's museum trajectory may marginally refine recommendations while exposing sensitive behavioral traces without pedagogical justification. Educational data collection must be demonstrably necessary, with heightened protections for minors (Prinsloo and Slade, 2017; Tsai et al., 2020), as data infrastructures can drift toward institutional surveillance (Williamson, 2021). We therefore derive three requirements: (1) *necessity*—collect only data demonstrably required for pedagogical function; (2) *minimization by design*—aggregate or anonymize data at the earliest possible stage; and (3) *purpose limitation*—prohibit secondary uses diverging from stated educational objectives.

Collectively, these cases reveal a shared causal architecture: agency erosion, pedagogical irrelevance, epistemic opacity, structural inequity, and surveillance risk emerge when optimization precedes educational validation. In each tension, optimization-first architectures encode problematic assumptions—convergence maximizes efficiency, spatial distance proxies engagement, statistical regularities justify recommendations, interaction frequency indicates value, and data collection is warranted by marginal gains—that systematically generate educational failures when algorithmic processing precedes pedagogical and ethical validation.

4.2 Constraint-First Processing Hierarchy

Formalization and Architectural Properties

Educational LBRS require architectures that operationalize pedagogical and ethical priorities as processing preconditions rather than optimization variables. We formalize the *constraint-first principle* as an architectural ordering rule.

Let L denote the universe of possible locations, U a learner, and $R(U)$ the recommended set. In engagement-first systems, the recommendation function optimizes $R(U) = \arg\max_{l \in L}$ *Engagement(l, U)*, subject to weak post-hoc constraints. In constraint-first systems, the function instead

optimizes $R(U)$ = arg max over $l \in L_\text{valid}(U)$ of $Personalization(l, U)$, where $L_\text{valid}(U)$ comprises all locations $l \in L$ satisfying both $Pedagogical(l, U)$ and $Ethical(l, U)$.

Pedagogical and ethical criteria are architectural preconditions, not optimization factors. This ordering guarantees that no optimization objective can override foundational commitments, translating Winner's (1980) insight—artifacts have politics—into processing order as architectural politics. Commercial LBRS employ engagement-first architectures where candidate generation precedes weak post-hoc constraints, optimizing for engagement and retention with fairness audits applied retroactively if at all. Traditional constraint-based systems validate rules after candidate generation, ensuring constraint satisfaction but typically without systematic fairness integration. The constraint-first approach sequences pedagogical validity gates, fairness-aware enrichment, and statistical personalization, establishing educational values as non-negotiable architectural preconditions that define the admissible solution space before any optimization occurs.

Sequential processing ensures three essential properties. *Guaranteed invariants*: pedagogical and ethical rules operate as hard constraints, ensuring no recommendation violates core principles. *Operational transparency*: discrete outputs at each stage create inspectable checkpoints supporting educator auditing and contestation. *Practical robustness*: early stages rely on ontological and pedagogical knowledge rather than interaction histories, preserving validity under data sparsity.

Five-Component Architecture

We operationalize the constraint-first principle through five ordered components (Figure 1):

Component 1: Pedagogical Constraint Filter. Educational ontologies enforce hard eligibility rules across safety (adult supervision; environmental risks), accessibility (tactile exhibits; audio descriptions; wheelchair access), and prerequisites (module completion; certification). Locations violating any rule are excluded before semantic matching, ensuring validity cannot be overridden. Institutions customize rule sets through participatory processes with educators, refined iteratively through contestation feedback (Component 5).

Component 2: Educational Matching. Locations are represented as semantic vectors encoding competencies, prerequisites, embodied actions, social configurations, and developmental complexity. Semantic vectors encode accessibility affordances alongside competencies using structured templates (e.g.,

Bloom's taxonomy). For a visually impaired learner studying erosion, accessible sites offering tactile samples are prioritized over proximate but inaccessible alternatives. Controlled semantic diversity mechanisms support autonomy, competence, and relatedness (Deci and Ryan, 2000).

Component 3: Fairness-Aware Enrichment. Collaborative signals are incorporated through fairness-aware re-ranking (Ekstrand et al., 2022), operationalized via demographic-parity constraints (e.g., 85–115% of mean exposure). Kleinberg et al. (2017) prove no system can simultaneously achieve calibration and equalized odds, demonstrating inherent trade-offs in fairness metrics. Constraint-first processing renders these trade-offs explicit and contestable (C5), enabling stakeholders to negotiate context-appropriate solutions.

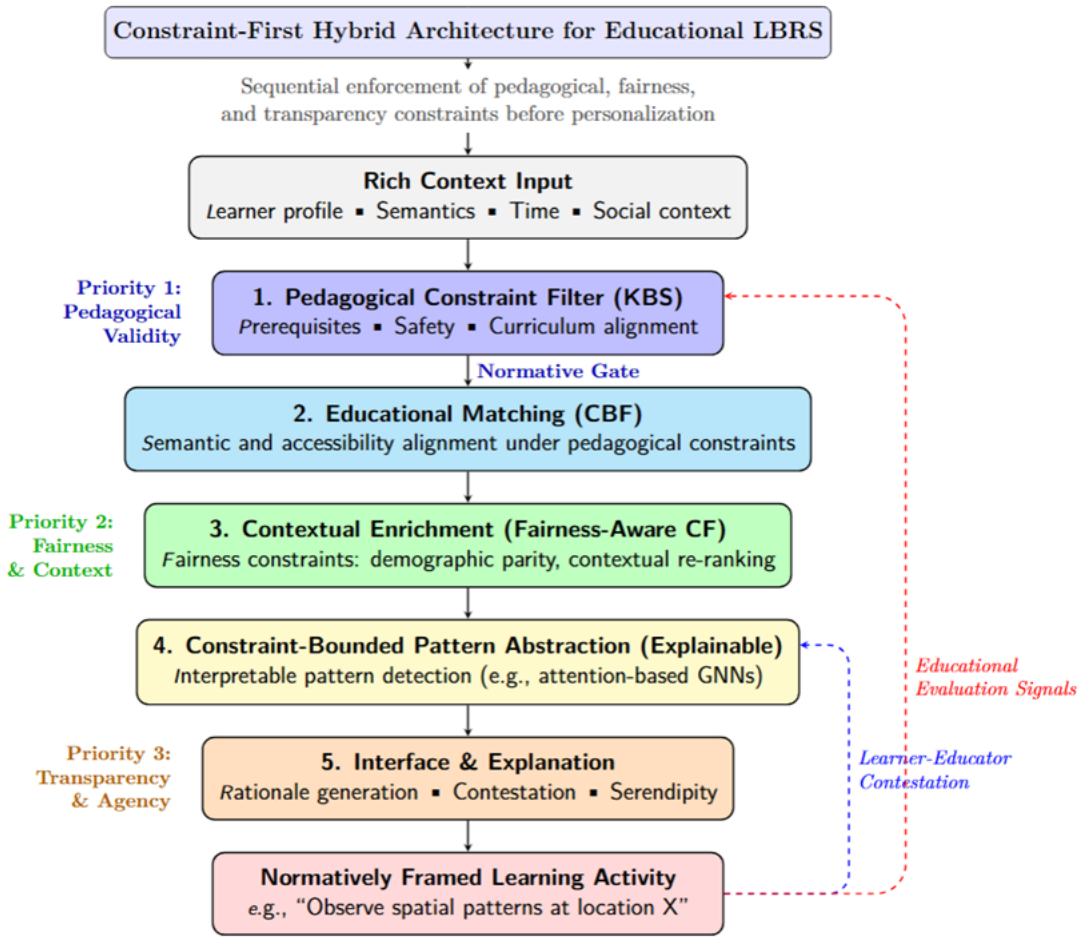


Figure 1. Constraint-First Hybrid Architecture: Five-Component Processing Hierarchy

Component 4: Pattern Recognition. Where sufficient data exists, explainable pattern-recognition methods (e.g., Graph Neural Networks; Wu et al., 2022) identify spatio-temporal learning patterns. Federated learning architectures (Yang et al., 2019) enable such analysis without centralizing sensitive learner data. This component is activated only when empirical evidence demonstrates pedagogical benefit exceeding privacy cost.

Component 5: Interface and Explanation. Educator-facing interfaces support co-orchestration through overrides and layered explanations referencing pedagogical rules (C1), goal alignment (C2), fairness interventions (C3), or patterns (C4). Explanations explicitly distinguish pedagogical validity from fairness-driven exposure adjustments, supporting algorithmic literacy. Contestations feed back into rule refinement, semantic representations, and optional pattern models.

Inter-Component Conflict Resolution

When semantic diversity (C2) conflicts with fairness constraints (C3), re-ranking operates *within* the pedagogically valid candidate set. When pattern recognition (C4) conflicts with pedagogical constraints (C1), constraint violations are non-negotiable. When learner autonomy (C2/C5) conflicts with equity mandates (C3), explanations distinguish fairness intervention from pedagogical invalidity. Institutions may adopt components incrementally: initial deployment (C1–C2) requires no interaction data; fairness mechanisms (C3) activate as data accumulates; pattern recognition (C4) only when justified.

Illustrative Implementation

In a 9th-grade erosion unit, C1 filters by safety/accessibility/prerequisites. C2 aligns "observe weathering" with sites affording sediment collection, prioritizing an accessible stream over an inaccessible canyon. C3 boosts the stream site in an underserved neighborhood. C4 identifies improved gains when sequenced with lab work. C5 explains: "Supports sediment collection (objective 3.2), provides tactile materials, ensures equitable exposure."

Empirical Validation Questions

The constraint-first principle constitutes an architectural design hypothesis requiring empirical validation. We propose five research questions: (1) *Pedagogical coherence*—can educators achieve reliable consensus on core rules? (2) *Learner agency*—do systems preserve or enhance perceived autonomy, competence, and relatedness? (3) *Learning efficacy*—do such systems yield equivalent or superior

outcomes? (4) *Equity in access*—do fairness mechanisms reduce demographic disparities in exposure? (5) *Data responsibility*—can meaningful personalization be achieved without disproportionate increases in location data collection?

5. Implications and Research Agenda

5.1 Theoretical and Practical Contributions

This framework advances technology-enhanced learning research as the field navigates the transition "from admiration to apprehension" (Cingillioglu & Barbieri, 2025) regarding AI—a shift demanding architectural, not merely rhetorical, responses to ethical concerns.

Operationalizing architectural causality (Section 3), the framework demonstrates that recurrent failures—agency erosion, inequitable exposure, opacity, surveillance—are systematic consequences of processing order, not implementation flaws. This shifts ethical analysis from interface features to architectural preconditions. It extends Value-Sensitive Design (Friedman and Hendry, 2019) by operationalizing values as non-compensatory computational gates rather than aspirational commitments. Values define the admissible solution space for any educational system balancing personalization with ethics. Methodologically, the use of paradigmatic case construction contributes an analytic template for diagnosing architectural tensions across educational technologies. Rather than evaluating isolated system failures, this approach reveals recurring contradictions between commercial optimization logics and pedagogical imperatives.

The constraint-first principle yields three transferable design patterns: (1) *Values-as-Gates* (non-compensatory filters); (2) *Explainability-by-Construction* (from processing traces); (3) *Fairness-within-Validity* (equity within valid sets). These reframe design: optimization as search within ethical bounds; explainability as intrinsic property; equity as structural invariant.

The framework demonstrates computational enforcement of pedagogical commitments, extending learning analytics by embedding fairness, privacy, and agency as architectural constraints rather than post-hoc indicators. UDL integration ensures accessibility is encoded from inception rather than retrofitted (Edyburn, 2021). The architecture shows privacy-by-design remains compatible with personalization through federated learning and minimal data collection.

5.2 Scope and Boundary Conditions

The framework's scope is deliberately delimited by boundary conditions that clarify applicability while identifying productive directions for extension.

The contribution is normative and architectural, grounded in theoretical synthesis rather than empirical deployment. Claims concern design validity and architectural effects; empirical validation remains subsequent research. The constraint-first principle addresses processing order but does not exhaust AI ethics in education. Labor implications, environmental costs, and governance exceed its scope. Implementation feasibility varies: Component 1 ontology development challenges under-resourced institutions; Components 3-4 encounter cold-start constraints. Extension to generative AI introduces complexity, as probabilistic content generation cannot be fully pre-filtered through deterministic gates.

Learner-driven architectures may misalign with instructor-directed pedagogical traditions. Tacit or spiritual knowledge central to apprenticeship or Indigenous pedagogies may resist ontological encoding, limiting computational representability without diminishing epistemic value. The framework thus operates within specific pedagogical paradigms emphasizing learner agency and explicit knowledge structures.

Architectural interventions cannot remediate structural inequalities. While fairness-aware ranking mitigates disparities, equitable outcomes depend on institutional investment and infrastructure. The framework operates within, not as substitute for, systemic reform. Technology-mediated interventions must therefore complement rather than replace structural reforms addressing root causes of educational inequity.

5.3 Empirical Validation Agenda

While personalization research in AI-enabled education demonstrates promising outcomes, systematic validation of ethically grounded architectures remains limited (Hardaker and Glenn, 2025). The constraint-first framework reframes ethical commitments as empirically testable architectural hypotheses. Validation trajectories organize across three interdependent layers.

Initial studies must establish the feasibility of pedagogical constraint encoding. This includes developing educator-facing ontology enrichment tools capable of achieving inter-rater reliability thresholds ($\kappa > 0.7$), evaluating privacy–utility trade-offs under federated learning regimes (Yang *et al.*, 2019), and addressing

consent and assent complexities in minor populations. These foundational investigations determine whether the architectural vision translates into implementable systems.

Controlled comparative deployments should evaluate immediate learning outcomes, short-term transfer effects, and enduring competencies. Particular attention should be given to self-regulated learning development across longitudinal intervals (e.g., 6–12 months), testing whether autonomy-supportive architectures yield durable motivational and epistemic benefits. Pilot studies enable iterative refinement before scaling, identifying unanticipated implementation challenges and validating theoretical predictions under real-world constraints.

System-wide implementation introduces additional challenges: rule maintenance complexity, institutional variation in ontological capacity, limitations of existing fairness metrics, and the need for computationally lightweight implementations suitable for resource-constrained educational contexts. Equity auditing must examine exposure distributions, participation rates, and downstream learning outcomes across demographic groups, ensuring fairness mechanisms achieve intended effects rather than introducing new disparities.

5.4 Generalization and Cross-Cultural Considerations

The framework requires cross-cultural validation and adaptation, as pedagogical traditions, privacy norms, and regulatory environments (e.g., GDPR vs. COPPA) vary globally, directly influencing components like fairness-aware enrichment (C3). Moreover, the constraint-first principle generalizes beyond LBRS to any educational technology where optimization threatens pedagogical values—such as generative AI tutors or adaptive assessments—where validity checks and pedagogically-grounded explanations must similarly precede optimization. This aligns with the growing recognition that ethical AI in education requires architectural, not just procedural, foundations (Ofosu-Asare, 2025).

Conclusion

We have demonstrated that educational LBRS require architectural reordering where pedagogical and ethical constraints precede optimization. Through paradigmatic analysis, we showed failures in educational recommender systems are not incidental design flaws but systematic consequences of engagement-first processing orders. By formalizing the constraint-first principle and operationalizing it as a processing hierarchy, we specify how pedagogical validity, accessibility, fairness, and transparency can

be enforced as non-compensatory architectural invariants. While prior work establishes the educational value of location-based learning (Kamarainen et al., 2013; Holden and Sykes, 2011), our contribution clarifies the architectural conditions under which such value can be preserved without undermining learner agency or equity.

Beyond LBRS, this work contributes to responsible AI in education by demonstrating ethics is not merely post-hoc evaluation but architectural sequencing. The ordering of validation and optimization processes determines whether educational values remain governing constraints or are reduced to compensable objectives. As generative and adaptive AI systems increasingly mediate learning, the constraint-first principle offers a transferable heuristic for embedding educational commitments structurally rather than rhetorically.

Ultimately, the success of an educational recommender system rests not on predictive accuracy but on how faithfully its architecture embodies educational principles. As location-based technologies mature, the central question is no longer whether systems can optimize engagement, but whether optimization can be permanently held accountable to learning, equity, and human dignity.

ACKNOWLEDGEMENTS

This work is supported by the MOBILES project funded by the Agence Nationale de la Recherche (ANR), France (grant ANR-20-CE38-0009).

References

Abrahamson, D. and Lindgren, R. (2014), "Embodiment and embodied design", in Sawyer, R.K. (Ed.), *The Cambridge Handbook of the Learning Sciences*, 2nd ed., Cambridge University Press, Cambridge, pp. 358–376.

Acquisti, A., Brandimarte, L. and Loewenstein, G. (2015), "Privacy and human behavior in the age of information", *Science*, Vol. 347 No. 6221, pp. 509–514, doi: https://doi.org/10.1126/science.aaa1465

Baker, R.S. and Hawn, A. (2022), "Algorithmic bias in education", *International Journal of Artificial Intelligence in Education*, Vol. 32 No. 4, pp. 1052–1092, doi: https://doi.org/10.1007/s40593-021-00285-9

Brown, J.S., Collins, A. and Duguid, P. (1989), "Situated cognition and the culture of learning", *Educational Researcher*, Vol. 18 No. 1, pp. 32–42, doi: https://doi.org/10.3102/0013189X018001032

Burgstahler, S. (2021), *Universal Design in Higher Education: From Principles to Practice*, 3rd ed., Harvard Education Press, Cambridge, MA.

Burke, R. (2000), "Knowledge-based recommender systems", in Kent, A. (Ed.), *Encyclopedia of Library and Information Science*, Vol. 69, Marcel Dekker, New York, NY, pp. 180–200.

Cingillioglu, I. and Barbieri, W. (2025), "From admiration to apprehension: the evolving role of generative AI in learning", *International Journal of Information and Learning Technology*, Vol. ahead-of-print No. ahead-of-print, doi: https://doi.org/10.1108/IJILT-11-2024-0262

da Silva, F.L., Slodkowski, B.K., da Silva, K.K.A. and Cazella, S.C. (2023), "A systematic literature review on educational recommender systems for teaching and learning: research trends, limitations and opportunities", *Education and Information Technologies*, Vol. 28 No. 3, pp. 3289–3328, doi: https://doi.org/10.1007/s10639-022-11341-9

Deci, E.L. and Ryan, R.M. (2000), "Intrinsic and extrinsic motivations: classic definitions and new directions", *Contemporary Educational Psychology*, Vol. 25 No. 1, pp. 54–67, doi: https://doi.org/10.1006/ceps.1999.1020

Doroudi, S., Aleven, V. and Brunskill, E. (2019), "Where's the reward? A review of reinforcement learning for instructional sequencing", *International Journal of Artificial Intelligence in Education*, Vol. 29 No. 4, pp. 568–620, doi: https://doi.org/10.1007/s40593-019-00187-x

Drachsler, H., Verbert, K., Santos, O.C. and Manouselis, N. (2015), "Panorama of recommender systems to support learning", in Ricci, F., Rokach, L. and Shapira, B. (Eds.), *Recommender Systems Handbook*, 2nd ed., Springer, Boston, MA, pp. 421–451, doi: https://doi.org/10.1007/978-1-4899-7637-6_12

Edyburn, D.L. (2021), "Universal usability and universal design for learning", *Intervention in School and Clinic*, Vol. 56 No. 5, pp. 310–315, doi: https://doi.org/10.1177/1053451220963082

Eisenlohr, A., Currid-Halkett, E. and Steil, J. (2025), "The geography of cultural capital: measuring neighborhood-level gaps in access as a key driver of social mobility", *Journal of Economic Geography*, lbaf040, doi: https://doi.org/10.1093/jeg/lbaf040

Ekstrand, M.D., Das, A., Burke, R. and Diaz, F. (2022), "Fairness in information access systems", *Foundations and Trends in Information Retrieval*, Vol. 16 No. 1-2, pp. 1–177, doi: https://doi.org/10.1561/1500000079

Feng, J., Li, Y., Zhang, C., Sun, F., Meng, F., Guo, A. and Jin, D. (2018), "DeepMove: predicting human mobility with attentional recurrent networks", in *Proceedings of the 2018 World Wide Web Conference*, Lyon, France, ACM, pp. 1459–1468, doi: https://doi.org/10.1145/3178876.3186058

Flyvbjerg, B. (2006), "Five misunderstandings about case-study research", *Qualitative Inquiry*, Vol. 12 No. 2, pp. 219-245, doi: https://doi.org/10.1177/1077800405284363

Friedman, B. and Hendry, D.G. (2019), *Value Sensitive Design: Shaping Technology with Moral Imagination*, MIT Press, Cambridge, MA, doi: https://doi.org/10.7551/mitpress/7585.001.0001

Gibson, J.J. (1979), *The Ecological Approach to Visual Perception*, Houghton Mifflin, Boston, MA.

Hardaker, G. and Glenn, L.E. (2025), "Artificial intelligence for personalized learning: a systematic literature review", *International Journal of Information and Learning Technology*, Vol. 42 No. 1, pp. 1–14, doi: https://doi.org/10.1108/IJILT-07-2024-0160

Hevner, A.R., March, S.T., Park, J. and Ram, S. (2004), "Design science in information systems research", *MIS Quarterly*, Vol. 28 No. 1, pp. 75 –105, doi: https://doi.org/10.2307/25148625

Holden, C.L. and Sykes, J.M. (2011), "Leveraging mobile games for place-based language learning", *International Journal of Game-Based Learning*, Vol. 1 No. 2, pp. 1–18, doi: https://doi.org/10.4018/ijgbl.2011040101

Holmes, W., Porayska-Pomsta, K., Holstein, K., Sutherland, E., Baker, T., Shum, S.B., Santos, O.C., Rodrigo, M.T., Cukurova, M., Bittencourt, I.I. and Koedinger, K.R. (2022), "Ethics of AI in education: towards a community-wide framework", *International Journal of Artificial Intelligence in Education*, Vol. 32 No. 3, pp. 504–526, doi: https://doi.org/10.1007/s40593-021-00239-1

Holstein, K. and Doroudi, S. (2021), "Equity and artificial intelligence in education: will 'AIEd' amplify or alleviate inequities in education?", arXiv preprint arXiv:2104.12920, doi: https://doi.org/10.48550/arXiv.2104.12920

Jones, K.M. (2019), "Learning analytics and higher education: a proposed model for establishing informed consent mechanisms to promote student privacy and autonomy", *International Journal of Educational Technology in Higher Education*, Vol. 16 No. 1, pp. 1–22, doi: https://doi.org/10.1186/s41239-019-0155-0

Kamarainen, A.M., Metcalf, S., Grotzer, T., Browne, A., Mazzuca, D., Tutwiler, M.S. and Dede, C. (2013), "EcoMOBILE: integrating augmented reality and probeware with environmental education field trips", *Computers & Education*, Vol. 68, pp. 545–556, doi: https://doi.org/10.1016/j.compedu.2013.02.018

Khosravi, H., Shum, S.B., Chen, G., Conati, C., Tsai, Y.S., Kay, J., Knight, S., Martinez-Maldonado, R., Sadiq, S. and Gašević, D. (2022), "Explainable artificial intelligence in education", *Computers & Education: Artificial Intelligence*, Vol. 3, 100074, doi: https://doi.org/10.1016/j.caeai.2022.100074

Kleinberg, J., Mullainathan, S. and Raghavan, M. (2017), "Inherent trade-offs in the fair determination of risk scores", in *Proceedings of the 8th Innovations in Theoretical Computer Science Conference*, Berkeley, CA, USA, Article 43, pp. 1–23, doi: https://doi.org/10.4230/LIPIcs.ITCS.2017.43

Lave, J. and Wenger, E. (1991), *Situated Learning: Legitimate Peripheral Participation*, Cambridge University Press, Cambridge.

Martinez Marroquin, E. and Senadji, B. (2025), "Activity theory as framework for analysis of workplace learning technologies: the case of generative AI conversational agents", *International Journal of Information and Learning Technology*, Vol. 42 No. 4, pp. 353–365, doi: https://doi.org/10.1108/IJILT-07-2024-0141

Ofosu-Asare, Y. (2025), "Guiding the future: developing an ethical framework for generative AI use in education", *International Journal of Information and Learning Technology*, Vol. ahead-of-print No. ahead-of-print, doi: https://doi.org/10.1108/IJILT-06-2024-0113

Papadimitriou, I. (2016), "Identifying accessibility barriers in heritage museums: conceptual challenges in a period of change", *Museum International*, Vol. 68 No. 3–4, pp. 271–272, doi: https://doi.org/10.1111/muse.12134

Pazzani, M.J. and Billsus, D. (2007), "Content-based recommendation systems", in Brusilovsky, P., Kobsa, A. and Nejdl, W. (Eds.), *The Adaptive Web*, Springer, Berlin, Heidelberg, pp. 325–341, doi: https://doi.org/10.1007/978-3-540-72079-9_10

Prinsloo, P. and Slade, S. (2017), "Ethics and learning analytics: charting the (un)charted", in Lang, C., Siemens, G., Wise, A. and Gašević, D. (Eds.), *Handbook of Learning Analytics*, SOLAR, Edmonton, Canada, pp. 49–57, doi: https://doi.org/10.18608/hla17.004

Reeve, J. and Cheon, S.H. (2021), "Autonomy-supportive teaching: its malleability, benefits, and potential to improve educational practice", *Educational Psychologist*, Vol. 56 No. 1, pp. 54–77, doi: https://doi.org/10.1080/00461520.2020.1862657

Resnick, P. and Varian, H.R. (1997), "Recommender systems", *Communications of the ACM*, Vol. 40 No. 3, pp. 56–58, doi: https://doi.org/10.1145/245108.245121

Rose, D.H. and Meyer, A. (2002), *Teaching Every Student in the Digital Age: Universal Design for Learning*, Association for Supervision and Curriculum Development, Alexandria, VA.

Selwyn, N. (2016), *Is Technology Good for Education?*, Polity Press, Cambridge.

Siemens, G. (2005), "Connectivism: a learning theory for the digital age", *International Journal of Instructional Technology and Distance Learning*, Vol. 2 No. 1, pp. 3–10.

Tsai, Y.-S., Rates, D., Moreno-Marcos, P.M., Muñoz-Merino, P.J., Jivet, I., Scheffel, M., Drachsler, H., Kloos, C.D. and Gašević, D. (2020), "Learning analytics in European higher education: trends and barriers", *Computers & Education*, Vol. 155, 103933, doi: https://doi.org/10.1016/j.compedu.2020.103933

Vygotsky, L.S. (1978), Mind in Society: The Development of Higher Psychological Processes, Harvard University Press, Cambridge, MA.

Williamson, B. (2021), "Psychodata: disassembling the psychological, economic, and statistical infrastructure of 'social-emotional learning'", *Journal of Education Policy*, Vol. 36 No. 1, pp. 129–154, doi: https://doi.org/10.1080/02680939.2019.1672895

Wilson, M. (2002), "Six views of embodied cognition", *Psychonomic Bulletin and Review*, Vol. 9 No. 4, pp. 625–636, doi: https://doi.org/10.3758/bf03196322

Winner, L. (1980), "Do artifacts have politics?", *Daedalus*, Vol. 109 No. 1, pp. 121–136.

Wu, S., Sun, F., Zhang, W., Xie, X. and Cui, B. (2022), "Graph neural networks in recommender systems: a survey", *ACM Computing Surveys*, Vol. 55 No. 5, pp. 1–37, doi: https://doi.org/10.1145/3535101

Yang, Q., Liu, Y., Chen, T. and Tong, Y. (2019), "Federated machine learning: concept and applications", *ACM Transactions on Intelligent Systems and Technology*, Vol. 10 No. 2, pp. 1–19, doi: https://doi.org/10.1145/3298981

Zeb, A., Ullah, R. and Karim, R. (2024), "Exploring the role of ChatGPT in higher education: opportunities, challenges and ethical considerations", *International Journal of Information and Learning Technology*, Vol. 41 No. 1, pp. 99–111, doi: https://doi.org/10.1108/IJILT-04-2023-0046

Zhao, J., Wallgrün, J.O., Sajjadi, P., LaFemina, P., Lim, K.Y.T., Springer, J.P. and Klippel, A. (2022), "Longitudinal effects in the effectiveness of educational virtual field trips", *Journal of Educational Computing Research*, Vol. 60 No. 4, doi: https://doi.org/10.1177/07356331211062925