

Designing AI-Driven Learner Modeling Architectures for Regulated Competency-Based Education: A Comparative Analysis of Probabilistic and Symbolic Approaches in the Quebec Collegiate Context

Mouenis Anouar Tadlaoui¹, Khalid Mahdi², Mohammed Chekour³

(1) Collège d'Enseignement en Immobilier (CEI), Montréal, Québec, Canada

(2) École Normale Supérieure, Hassan II University, Casablanca, Morocco.

(3) École Supérieure de l'Éducation et de la Formation, Ibn Tofaïl University, Kenitra, Morocco.

Amatadlaoui@collegecei.com; k.mahdi@enscasa.ma; mohamed.chekour@uit.ac.ma

DOI : 10.5281/zenodo.20336761

Revue Sciences & Education, Vol.2 N° 1, 2025

<https://scienceseducation.org/>

Reçu le : 30/11/2025 | Accepté le : 30/12/2025 | Publié le : 14/01/2026

Comment citer cet article :

Anouar Tadlaoui, M., Mahdi, K., & Chekour, M. (2026). Designing AI-Driven Learner Modeling Architectures for Regulated Competency-Based Education. *Sciences & Education*, 2(1), 71–93. <https://doi.org/10.5281/zenodo.20336761>

Designing AI-Driven Learner Modeling Architectures for Regulated Competency-Based Education: A Comparative Analysis of Probabilistic and Symbolic Approaches in the Quebec Collegiate Context

Mouenis Anouar Tadlaoui¹, Khalid Mahdi², Mohammed Chekour³

(1) Collège d'Enseignement en Immobilier (CEI), Montréal, Québec, Canada

(2) École Normale Supérieure, Hassan II University, Casablanca, Morocco.

(3) École Supérieure de l'Éducation et de la Formation, Ibn Tofaïl University, Kenitra, Morocco.

Amatadlaoui@collegecei.com; k.mahdi@enscasa.ma; mohamed.chekour@uit.ac.ma

DOI : 10.5281/zenodo.20336761

ABSTRACT

Learner modeling is a foundational component of adaptive educational systems, yet the dominant architectures in the field have been conceived without regard for the specific regulatory and curricular constraints of provincial or national education systems. This article addresses that gap by conducting a structured conceptual analysis of two major learner modeling paradigms — symbolic (rule-based) and probabilistic (Bayesian) models — explicitly situated within the regulatory and curricular framework governing Quebec's collegiate sector. Drawing on a decade of empirical and theoretical work on Bayesian learner modeling, a systematic documentary analysis of Quebec's Act to Modernize Legislative Provisions Respecting the Protection of Personal Information (Law 25), competency-based program specifications (AEC/DEC devis ministériels), and the professional regulatory frameworks of the OACIQ and the RBQ, we identify six dimensions along which the two paradigms diverge in pedagogically and legally consequential ways. Based on this analysis, we propose the AIPAQ framework (Pedagogically Aligned AI Architecture for Quebec), a four-layer hybrid architecture that deploys symbolic logic at the regulatory and explainability layers and probabilistic inference at the competency modeling and micro-pedagogical layers. The proposed framework offers both a practical design guide for Quebec collegiate institutions and a generalizable methodology for regulatory-aware AI design in competency-based educational systems.

Keywords: artificial intelligence in education, learner modeling, Bayesian networks, Law 25, competency-based education, Quebec collegiate system, OACIQ, adaptive learning systems, hybrid architecture, algorithmic transparency

1) INTRODUCTION

The deployment of artificial intelligence in education (AIED) has generated substantial scholarly interest in learner modeling — the computational representation of a student's knowledge, competencies, misconceptions, and learning behaviors — as the inferential engine of adaptive educational systems (AES) and intelligent tutoring systems (ITS) (VanLehn, 2011; Brusilovsky & Millán, 2007). Two paradigms have historically dominated the field: symbolic, or rule-based, models, which encode expert-defined IF–THEN pedagogical rules, and probabilistic models, particularly Bayesian networks (BN) and their multi-entity extensions (MEBN), which represent learner states as probability

distributions updated dynamically as new evidence (Brusilovsky & Millán, 2007; VanLehn, 2011).

Building upon established work in probabilistic and multi-entity Bayesian network (MEBN) learner modeling, this paper extends Bayesian architectures into a complex regulatory and curricular context that prior literature has not adequately addressed.

A growing body of comparative research has examined the relative strengths of symbolic and probabilistic paradigms in terms of predictive accuracy, adaptability, and pedagogical transparency (Baker & Inventado, 2014; Nkambou et al., 2010). However, this literature has developed in a largely context-agnostic manner, treating the institutional, regulatory, and curricular environment as incidental rather than constitutive constraints on architectural design.. This is a significant omission. The viability of any AI-driven learner modeling architecture is not determined solely by its computational properties; it is equally determined by the legal frameworks governing data collection and algorithmic decision-making, the curricular structures within which pedagogical decisions must be grounded, and the professional regulatory requirements that govern the specific competencies being taught.

Quebec presents an unusually demanding instance of this problem. Its collegiate sector - encompassing cégeps and private non-subsidized colleges - operates under a competency-based curricular framework mandated by the Ministère de l'Enseignement supérieur (MES), in which every program is governed by a ministériel program specification (devis ministériel) that defines the competencies to be acquired and the criteria by which their attainment is evaluated (MES, 2022). Several professional programs within this system — including real estate brokerage (TCIR) and building inspection (TIBBI) — are additionally governed by professional regulatory bodies: the Organisme d'autoréglementation du courtage immobilier du Québec (OACIQ) and the Régie du bâtiment du Québec (RBQ), respectively. Simultaneously, the Act to Modernize Legislative Provisions Respecting the Protection of Personal Information (Law 25), fully in force since September 2023, imposes obligations of data minimization, purpose limitation, and — most critically for AI-driven systems — a right to explanation of automated decisions that affect individuals (article 12.1), a requirement without precedent in Canadian provincial privacy law.

No existing AIED framework, to the authors' knowledge, has addressed the joint architectural implications of these three constraints simultaneously. This article seeks to fill that gap. We pose the following research question: How should a learner modeling architecture be designed to be simultaneously predictively powerful, pedagogically interpretable, and fully compliant with Quebec's regulatory and curricular framework? To answer this question, we conduct a structured conceptual analysis of symbolic and probabilistic learner modeling paradigms across six dimensions chosen to capture the specific demands of the Quebec context. We then synthesize this analysis into the AIPAQ framework — a four-layer hybrid architecture — and discuss its implications for the design and deployment of adaptive learning systems in Quebec's collegiate institutions.

The paper is structured as follows. Section 2 reviews the relevant literature on learner modeling paradigms and AIED architecture. Section 3 characterizes the Quebec regulatory and curricular context. Section 4 presents the methodology. Section 5 conducts the comparative analysis. Section 6 presents the AIPAQ framework. Section 7 discusses theoretical and practical implications. Section 8 concludes with directions for future empirical validation.

2) BACKGROUND AND RELATED WORK

2.1. Symbolic Learner Models

The earliest architectures for intelligent tutoring systems were grounded in symbolic AI, representing learner knowledge through deterministic, expert-defined rule sets (Anderson et al., 1985). The overlay model — representing a learner's knowledge as a subset of an idealized expert model — and the buggy model — which explicitly encoded systematic learner misconceptions — exemplify this tradition (Carr & Goldstein, 1977; Brown & Burton, 1978). These models operated through forward-chaining inference: given a current learner state, the engine fired matching rules to generate diagnostic conclusions and instructional decisions (Wenger, 1987).

The primary virtues of symbolic models are well-documented: their transparency supports pedagogical validation by domain experts, their determinism ensures reproducible decision traces, and their structure maps naturally onto expert-articulated instructional sequences (VanLehn, 2011). However, as learning environments scaled in complexity and learner variability, the limitations of symbolic models became structural. Rule explosion impeded maintainability, and the deterministic architecture offered no mechanism for reasoning under uncertainty: a learner whose performance fluctuated near a mastery threshold was classified categorically rather than probabilistically, generating systematic classification errors for heterogeneous learner populations (Tadlaoui et al., 2016; Nkambou et al., 2010). Comparative empirical analysis has further confirmed that the inflexibility of rule-based models limits their accuracy in real-world adaptive hypermedia environments where learner behavior is inherently non-deterministic (Tadlaoui et al., 2018b).

2.2. Probabilistic Learner Models

The introduction of Bayesian networks into learner modeling represented a paradigmatic shift toward uncertainty-aware, data-driven inference (Mislevy, 1994; Conati et al., 2002). A Bayesian network represents learner attributes as nodes in a directed acyclic graph (DAG), connected by conditional probability tables (CPTs) that encode probabilistic dependencies. As new evidence arrives, the network updates its posterior probability estimates through belief propagation, enabling real-time inference about latent learner states from incomplete observational data.

Foundational research in BN learner modeling established methodologies for constructing networks directly from Unified Modeling Language (UML) specifications

of domain knowledge, enabling principled, expert-driven network design (Tadlaoui et al., 2015a). To address the critical cold-start problem inherent in purely data-driven probabilistic systems under sparse data, hybrid methods combining Bayesian inference with stereotype-based prior estimation were subsequently developed (Tadlaoui et al., 2017) — a strategy particularly well-suited to small-cohort AEC contexts.

Furthermore, the integration of Multi-Entity Bayesian Networks (MEBN) extended classical BNs by supporting relational modeling across multiple learners and dynamic temporal contexts, leveraging cross-learner relational data to significantly improve individual-level inference (Tadlaoui et al., 2019b).

2.3. Hybrid Architectures

Recognition of the complementary strengths of both paradigms has generated sustained interest in hybrid architectures. Early hybrid approaches successfully integrated BN probabilistic inference with symbolic overlay models to evaluate learner profiles, demonstrating that the representational clarity of rule-based frameworks can be combined with the inferential flexibility of Bayesian networks in an operational system (Tadlaoui et al., 2019a). Subsequent practical implementations in production e-learning platforms further confirmed that such probabilistic hybrid models are deployable at scale within functional educational software environments.

More recent work has formalized hybrid integration through pedagogical constraints on BN structure, rule-gated activation of probabilistic modules, and ontology-BN hybrids (Nkambou et al., 2010; Arroyo et al., 2014). Despite these contributions, a fundamental tension persists: the transparency–accuracy trade-off. Systems that maximize predictive accuracy through deep probabilistic inference tend to minimize the interpretability of their outputs; systems that maximize interpretability through rule-based transparency tend to sacrifice predictive nuance (Drachsler & Greller, 2016; Winne & Baker, 2013). Existing hybrid frameworks have addressed this tension as a technical optimization problem, but have not grounded it in specific legal obligations that might transform it from a design preference into a binding constraint. The present paper makes this move explicit.

2.4. Regulatory-Aware AIED Design: An Underexplored Frontier

The relationship between AI governance frameworks and AIED architecture has received increasing attention since the enactment of the EU AI Act and GDPR (Holmes et al., 2019; Zawacki-Richter et al., 2019). However, this literature has focused predominantly on European or generic Anglo-American institutional contexts. Competency-based educational systems present distinctive challenges that generic AIED architectures are not equipped to address: competency states are holistically assessed and not reducible to atomic knowledge units, regulatory bodies prescribe non-negotiable minimum standards, and competency evaluation is governed by institutional policies with legal force. Consequently, the personalization of digital learning environments must be conceptualized not merely as a technical optimization problem, but as an institutional

practice deeply embedded in specific governance contexts — a perspective that directly motivates the present study's regulatory-first design approach

3) THE QUEBEC REGULATORY AND CURRICULAR FRAMEWORK

3.1. Law 25: Data Privacy and Algorithmic Transparency

Quebec's Act to Modernize Legislative Provisions Respecting the Protection of Personal Information (Law 25, L.Q. 2021, c. 25; L.Q. 2022, c. 22) entered full force in September 2023 and represents the most stringent personal data protection legislation in North America. Its implications for AI-driven educational systems are threefold.

First, the data minimization principle (section 5) requires that personal information collected from learners be limited to what is necessary for the explicitly stated and consented purpose. This constrains the behavioral data collection strategies on which probabilistic learner models depend: clickstream traces, time-on-task logs, navigation patterns, and affective signals are all personal information under Law 25, collectible only to the extent justified by a documented pedagogical purpose.

Second, section 12.1 requires that any organization using personal information to render a decision based exclusively on automated processing that significantly affects the individual must, upon request, be able to explain the principal parameters and data that led to the decision and disclose the individual's right to have a human review the decision. This provision functionally operationalizes a right to explanation directly analogous to — and in some respects stricter than — GDPR Article 22. For AI-driven adaptive learning systems, every automated instructional recommendation, remediation trigger, or risk-of-failure alert constitutes a decision affecting the learner and, accordingly, requires a documentable explanation.

Third, section 3.3 requires an a priori Privacy Impact Assessment (PIA) for any project involving the collection or use of personal information through new technologies. This institutionalizes a design-time obligation that necessitates architectural decisions to be privacy-reviewed before deployment.

Together, these provisions create a binding design constraint: any learner modeling architecture deployed in Quebec must be capable of generating documented, human-readable explanations of its adaptive decisions — a requirement that fundamentally privileges architectures whose inference mechanisms are inherently interpretable or that incorporate dedicated explainability layers.

3.2. The Competency-Based Curricular Framework

Quebec's collegiate education system is organized around a competency-based approach (*approche par compétences*, APC) that constitutes the curricular architecture of all DEC and AEC programs. This structure reflects the broader francophone tradition of rigorous standardization and centralized curriculum design, whose distinctive characteristics relative to anglophone approaches have been analyzed comparatively in recent

pedagogical research (Mahdi, 2024). Every program is governed by a ministériel program specification (devis ministériel) issued by the Ministère de l'Enseignement supérieur. Each competency is defined by a competency statement, constituent competency elements, and observable performance criteria. Competency attainment is evaluated according to a binary logic — competency achieved or not achieved — as specified in each institution's learning assessment policy (PEA).

This architecture has two critical implications for learner modeling. First, the unit of modeling is the competency, not the knowledge concept: a learner's state must be represented at the level of integrative, cross-domain competencies whose attainment can only be inferred from performance on authentic, complex professional tasks. This contrasts sharply with the knowledge concept graphs that underpin most BN-based learner models such as Bayesian Knowledge Tracing (Corbett & Anderson, 1994), requiring structural reconfiguration of the BN to align with competency rather than concept hierarchies. Second, the binary evaluation logic of PEA policies imposes a threshold-based assessment structure that is directly encodable as rule-based conditions but requires explicit modeling decisions in probabilistic architectures.

AEC programs present additional specificities: they target adult learners with existing professional experience, producing markedly heterogeneous populations that amplify both the value of probabilistic uncertainty management and the importance of transparent, explainable recommendations (Tardif, 2006).

AEC programs present additional specificities: they target adult learners with existing professional experience, producing markedly heterogeneous populations that amplify both the value of probabilistic uncertainty management and the importance of transparent, explainable recommendations (Tardif, 2006).

3.3. Professional Regulatory Bodies: OACIQ and RBQ

The Organisme d'autoréglementation du courtage immobilier du Québec (OACIQ) governs the practice of real estate brokerage in Quebec and defines the knowledge and competency requirements for the practice license. The AEC in Real Estate Brokerage (TCIR) must align its program outcomes with these requirements, and any AI-driven adaptive system deployed in this program must guarantee that its pedagogical recommendations respect the OACIQ competency framework as a non-negotiable external constraint backed by professional licensing consequences.

The Régie du bâtiment du Québec (RBQ) performs an analogous function for the building inspection sector. For learner modeling architectures, the implication is clear: adaptive recommendations cannot be optimized exclusively against academic performance metrics. They must be constrained by prescriptive external requirements expressed as non-negotiable performance thresholds — a class of constraint naturally expressible as hard rules but requiring careful structural integration in probabilistic systems.

3.4. Institutional Quality Assurance: CEEC

The Commission d'évaluation de l'enseignement collégial (CEEC) conducts periodic evaluations of all Quebec collegiate programs. For institutions deploying AI-driven adaptive systems, CEEC evaluation introduces a traceability requirement: the institution must document how the adaptive system's decisions contributed to the learning process. This institutional accountability obligation reinforces the Law 25 explainability requirement at the organizational level, further anchoring the case for architectures that generate documentable decision traces.

4) METHODOLOGY

4.1. Research Design

This study adopts a conceptual analysis research design (Jabareen, 2009; Mouza & Lavigne, 2013), combining a systematic literature review with a structured documentary analysis of Quebec's regulatory and curricular frameworks. Conceptual analysis is appropriate when the objective is to develop a framework synthesizing existing knowledge across disciplinary domains into an architecturally actionable specification, consistent with design science research traditions in educational technology (Reeves, 2006).

4.2. Systematic Literature Review

The literature review followed PRISMA guidelines adapted for conceptual reviews in educational technology (Page et al., 2021). Databases searched included ERIC, IEEE Xplore, ACM Digital Library, Scopus, and Google Scholar. Search terms combined: 'learner modeling', 'Bayesian network education', 'rule-based adaptive learning', 'intelligent tutoring systems', 'hybrid learner model', 'algorithmic transparency education', 'AI governance education', and their French-language equivalents. Inclusion criteria: peer-reviewed articles or official regulatory documents; English or French; published 1985–2024 for AIED literature; no date restriction for regulatory documents.

4.3. Analytical Framework

The comparative analysis operationalizes six dimensions derived from the convergence of established AIED learner modeling evaluation criteria (Brusilovsky & Millán, 2007; Tadlaoui et al., 2016) and the specific regulatory and curricular demands identified through documentary analysis. The six dimensions are: (1) representational capacity for integrated competency modeling; (2) inference mechanism traceability; (3) management of epistemic uncertainty; (4) algorithmic transparency under Law 25 article 12.1; (5) professional regulatory compliance (OACIQ, RBQ, CEEC); and (6) structural alignment with the APC curricular framework.

The selection of these six dimensions was guided by two convergent criteria. First, each dimension corresponds to a distinct source of constraint on learner modeling architecture identified through the documentary analysis: dimensions 1, 3, and 6 reflect

the pedagogical and curricular demands of the APC framework; dimensions 2 and 4 reflect the algorithmic transparency obligations of Law 25; and dimension 5 reflects the prescriptive requirements of the OACIQ and RBQ. Second, each dimension corresponds to a decision point at which the two paradigms diverge in architecturally consequential ways — that is, dimensions where a choice of paradigm yields materially different compliance, interpretability, or pedagogical outcomes. Dimensions where the two paradigms converge, or where the difference is not consequential in the Quebec context, were excluded.

5) COMPARATIVE ANALYSIS IN THE QUEBEC CONTEXT

Table 1 presents the structured comparative analysis across the six dimensions. The subsections that follow develop the reasoning underlying each assessment.

Table 1. Comparative Analysis of Symbolic and Probabilistic Learner Modeling Paradigms in the Quebec Collegiate Context

Dimension	Symbolic Models	Probabilistic Models	Contextual Advantage (Quebec)
1. Competency representational capacity	Direct representation of competency elements and performance criteria	Models interdependencies and partial mastery states	Hybrid: symbolic structure anchors hierarchy; BN estimates mastery probability (Tadlaoui et al., 2018a)
2. Inference traceability	Deterministic; every recommendation traceable to the firing rule	Probabilistic; requires post-hoc explainability techniques	Symbolic: CEEC audit traceability and Law 25 art. 12.1 satisfied directly
3. Uncertainty management	Weak: fixed thresholds misclassify borderline adult learners	Strong: continuous posterior updates handle heterogeneous AEC populations	Probabilistic: essential for heterogeneous adult AEC learners (Tadlaoui et al., 2017)
4. Algorithmic transparency (Law 25, s. 12.1)	High: rule sets are inherently human-readable explanations	Low: opacity of dense BNs requires LIME/SHAP post-processing	Symbolic: directly satisfies legal obligation without additional tooling
5. Professional regulatory compliance (OACIQ/ RBQ/ CEEC)	Direct: prescriptive competency requirements encoded as hard guardrails	Indirect: constraints must be structurally integrated into BN	Symbolic: natural encoding of non-negotiable professional licensing standards

6. APC curricular alignment	Direct: minist�riel performance criteria map naturally to evaluation rules	Structured: competency elements become root nodes; learning activities become evidence nodes	Hybrid: symbolic layer encodes PEA logic; BN models probabilistic mastery trajectories (Tadlaoui et al., 2019b)
-----------------------------	--	--	---

5.1. Representational Capacity for Integrated Competency Modeling

Quebec’s APC framework requires learner states to be modeled at the level of integrative, multi-dimensional competencies, not the atomic knowledge concepts that underpin most existing BN-based learner models such as Bayesian Knowledge Tracing (Corbett & Anderson, 1994). A competency such as ‘Analyze and advise on a real estate transaction in compliance with OACIQ brokerage regulations’ cannot be decomposed into independent sequences of atomic knowledge nodes without losing the integrative character that defines competency-based assessment. Symbolic models can directly encode the hierarchical structure of competencies as defined in the devis minist riel, but cannot represent probabilistic interdependencies among competency elements. Probabilistic models, particularly the MEBN framework developed in Tadlaoui et al. (2019b), offer superior capacity to represent these interdependencies and to infer latent mastery states from observable behavioral evidence. However, standard BN architectures are not designed around APC competency hierarchies and must be deliberately restructured — a process formalized in Tadlaoui et al. (2015a), which demonstrated how UML-specified domain knowledge can be systematically transformed into BN structures suitable for learner modeling.

5.2. Inference Mechanism Traceability

CEEC evaluations require institutions to document the relationship between adaptive instructional decisions and minist rially specified competency standards. Law 25 section 12.1 independently requires that automated decisions affecting individuals be documentably explainable. Symbolic models satisfy this inherently: the active rules at the moment of decision constitute an explanation directly referencing the competency criteria and regulatory standards that motivate each inference. Probabilistic models, whose outputs are posterior probability distributions, do not inherently produce such explanations. Post-hoc techniques such as LIME (Ribeiro et al., 2016), SHAP values, or counterfactual explanation generation can produce approximate explanations, but add architectural complexity and may not satisfy the ‘principal parameters and data’ standard of Law 25 section 12.1 in the way a direct rule trace does. This asymmetry motivates the dedicated symbolic explanation layer in the AIPAQ framework.

5.3. Uncertainty Management

Adult learners in AEC programs present heterogeneous prior knowledge profiles, diverse professional backgrounds, and non-linear learning trajectories that symbolic models with

fixed mastery thresholds are structurally ill-equipped to handle. The comparative analysis conducted in Tadlaoui et al. (2016) confirmed that rule-based models generate systematic classification errors for learners who fall near mastery thresholds — precisely the population most in need of adaptive support. Probabilistic models handle this heterogeneity naturally: by representing mastery as a probability distribution, they allow for graduated, evidence-responsive adaptation. The stereotype-based initialization strategy developed in Tadlaoui et al. (2017) is particularly valuable in this context, providing informed prior estimates for new AEC learners whose behavioral data is initially sparse, mitigating the cold-start problem inherent in purely data-driven probabilistic systems.

5.4. Algorithmic Transparency Under Law 25

The transparency obligation of Law 25 section 12.1 is not merely an administrative requirement; it reflects a substantive commitment to learner agency that is pedagogically consequential in its own right. Adult learners who understand why a system recommends a particular learning pathway are better positioned to engage critically with that recommendation and take ownership of their trajectory (Winne & Baker, 2013). Symbolic models, whose rule sets constitute self-evident explanations, align with this pedagogical value as well as the legal requirement. The practical challenge for probabilistic models is not merely technical but communicative: a posterior probability of 0.67 for competency mastery is not, without translation, a pedagogically actionable explanation. The AIPAQ framework addresses this through a dedicated symbolic explanation layer that translates probabilistic inferences into human-readable, competency-anchored justifications — a translation mechanism consistent with the hybrid explanation architecture explored in Tadlaoui et al. (2019a).

5.5. Professional Regulatory Compliance

The prescriptive competency standards of the OACIQ and RBQ function as hard constraints: no AI-driven system should recommend a learning pathway that could result in a learner advancing to professional licensure examination without having demonstrated the required competencies. In symbolic systems, these standards are encoded directly as non-overridable guardrails. In probabilistic systems, constraint enforcement requires structural integration into the BN itself — designating OACIQ-mandated competency nodes as mandatory prerequisites whose posterior probability must exceed a regulatory threshold before downstream recommendations are activated. Without such integration, optimization over a global performance metric could in principle generate recommendations that violate regulatory standards in edge cases. The AIPAQ architecture addresses this by assigning regulatory constraint enforcement exclusively to the symbolic layer, which acts as a non-overridable filter on all recommendations generated by the probabilistic layers.

5.6. APC Curricular Alignment

The binary competency attainment logic of Quebec's PEA policies maps directly onto the conditional rule structure of symbolic models. However, this binary endpoint does not preclude probabilistic modeling of the developmental trajectory toward attainment: the BN represents the continuous evolution of mastery probability across learning activities, with the binary PEA judgment applied only at the official evaluation threshold. This two-level structure — continuous probabilistic modeling of learning progress, binary symbolic enforcement of evaluation rules — is directly supported by the MEBN architecture developed in Tadlaoui et al. (2019b), where competency elements constitute the root nodes of the probabilistic network and assessment observations constitute the evidence nodes, precisely mirroring the competency hierarchy of the *devis ministériel*.

6) THE AIPAQ FRAMEWORK: A REGULATORY-AWARE HYBRID ARCHITECTURE

The AIPAQ (Pedagogically Aligned AI Architecture for Quebec) framework synthesizes the comparative analysis into a four-layer hybrid architecture designed to satisfy simultaneously the predictive, pedagogical, and regulatory demands of Quebec's collegiate context. The architecture departs from existing hybrid proposals by treating regulatory compliance and curricular alignment as first-class architectural constraints rather than post-hoc additions.

6.1. Layer 1 — Regulatory Symbolic Layer

The regulatory symbolic layer constitutes the non-negotiable foundation of the AIPAQ architecture. It encodes, as hard rules, the full set of legally and professionally binding constraints: the OACIQ and RBQ minimum competency thresholds, expressed as prerequisite conditions that must be satisfied before any recommendation advances a learner toward professional licensure assessment; the PEA-mandated evaluation rules, including course pass/fail criteria and competency sanction logic; the Law 25 data collection constraints, specifying permissible data types, consent-validated purposes, and retention limits; and the CEEC traceability requirements, encoded as logging and documentation obligations that accompany every adaptive decision.

This layer functions as a constraint satisfaction guardrail: every recommendation generated by any higher layer must pass through this layer's rule evaluation before being acted upon. A recommendation that would violate any hard rule is blocked and replaced by a rule-compliant alternative, with the blocking reason logged for the Law 25 decision record. This architecture ensures that regulatory compliance is not dependent on the correct behavior of the probabilistic layers — it is enforced structurally.

6.2. Layer 2 — Curricular Hybrid Layer

The curricular hybrid layer operates at the interface of symbolic structure and probabilistic inference. It maintains a Bayesian Network whose macro-structure is defined by the competency hierarchy of the relevant *ministériel* program specification: competency statements become root nodes, competency elements become intermediate

nodes, and specific learning activities and assessment items become evidence nodes. This structure is not learned from data but defined by expert pedagogical mapping, following the UML-to-BN construction methodology formalized in Tadlaoui et al. (2015a), ensuring that the network's causal architecture reflects the curricular logic of the *devis ministériel*.

Conditional probability tables are estimated from historical learner data — ideally pooled across cohorts using the MEBN framework's capacity for cross-learner inference (Tadlaoui et al., 2019b). The cold-start problem for new cohorts is addressed through the stereotype-based prior initialization strategy developed in Tadlaoui et al. (2017), which provides informed initial probability estimates based on learner profile features, enabling the system to generate meaningful recommendations from the first learning session. In parallel, soft rules — pedagogically recommended but not legally mandated — provide structural guidance on canonical learning sequences and common prerequisite relationships.

6.3. Layer 3 — Micro-Pedagogical Probabilistic Layer

The micro-pedagogical probabilistic layer handles fine-grained, real-time adaptive decisions: recommending specific learning resources, adjusting task difficulty, flagging recurring error patterns indicative of persistent misconceptions, and detecting early indicators of disengagement or failure risk. At this level, the full inferential flexibility of Bayesian reasoning is engaged, drawing on the posterior competency state estimates from Layer 2 combined with real-time behavioral evidence to generate individualized micro-recommendations that maximize expected learning gain within the constraints enforced by Layer 1.

The MEBN framework is particularly appropriate for this layer in multi-cohort settings, enabling the system to leverage structural similarities across learner profiles through probabilistic relational reasoning (Tadlaoui et al., 2019b) while maintaining full individualization of posterior estimates. Furthermore, empirical implementations confirm that such probabilistic layers can be integrated into production learning management systems without prohibitive technical overhead.

6.4. Layer 4 — Symbolic Explanation Layer

The symbolic explanation layer translates the posterior probability estimates and micro-recommendations of Layers 2 and 3 into human-readable, competency-anchored explanations that satisfy the requirements of Law 25 section 12.1 and support meaningful learner and instructor engagement with the system's outputs. Explanation generation proceeds through a counterfactual rule engine that maps each probabilistic recommendation to the competency elements, performance criteria, and regulatory requirements that motivate it.

The combination of probabilistic inference with symbolic explainability at the output layer draws on the hybrid architecture explored in Tadlaoui et al. (2019a), which

demonstrated that overlay model representations can be used to generate interpretable explanations of Bayesian inference outputs in adaptive hypermedia contexts. In the AIPAQ framework, this approach is formalized and extended: every recommendation is accompanied by an explanation structure that references the relevant *devis ministériel* competency, the current posterior mastery estimate, the specific performance criteria that the evidence pattern suggests are not yet reliably satisfied, and — where applicable — the OACIQ or RBQ regulatory requirement that the competency element supports. This explanation is logged and stored as part of the learner's decision record, constituting the documentation required by Law 25 section 12.1 and the CEEC traceability requirement.

To illustrate how the four layers interact in practice, consider a learner enrolled in the AEC in Real Estate Brokerage (TCIR) who completes a scenario-based assessment simulating a residential transaction. Layer 1 verifies that the learner has satisfied the OACIQ-mandated prerequisite competency thresholds before any progression recommendation is generated. Layer 2 updates the posterior probability distribution over the learner's "Transactional Compliance" competency node on the basis of the observed performance evidence, integrating this result with prior estimates initialized through the stereotype-based method of Tadlaoui et al. (2017). If the posterior falls below the mastery threshold, Layer 3 identifies the specific error patterns in the learner's reasoning — for example, a recurring misapplication of disclosure obligations under the Real Estate Brokerage Act — and recommends a targeted remediation resource. Layer 4 then generates a human-readable explanation referencing the relevant OACIQ competency element, the current posterior mastery estimate, and the specific performance criterion that the evidence pattern suggests has not yet been reliably satisfied. This explanation is logged as part of the learner's decision record, satisfying both Law 25 section 12.1 and the CEEC traceability requirement.

To operationalize this layer, the counterfactual rule engine applies deterministic threshold-gating to translate continuous probabilities into actionable pedagogical assertions that fulfill the requirements of Law 25 article 12.1. Formally, if the micro-pedagogical layer yields a posterior probability $P(C_i | E) < \theta$ (where C_i represents a specific competency element and θ is the institutional mastery threshold, typically set at 0.70), a conditional rule fires to intercept the automated recommendation pathway. Instead of presenting a raw statistical output to the user, the engine executes a query that maps the specific configuration of evidence nodes E to a human-readable template. For example, an operational rule structured as IF $P(\text{Transactional_Compliance} | E) < 0.70$ AND $E_{\text{disclosure}} == \text{'incorrect'}$ THEN $\text{Trigger_Remediation}(\text{Resource_A})$ AND $\text{Generate_Explanation}(\text{'Law 25 Parameter Log: Your mastery estimation for mandatory transaction disclosure is currently 67\%, which falls below the OACIQ-aligned threshold of 70\% due to an incorrect application of disclosure rules in Scenario 3.'})$ ensures that every automated instructional divergence is accompanied by a transparent, verifiable parameters trace.

6.5. Architectural Summary

Table 2 provides a consolidated overview of the AIPAQ framework's four-layer architecture. Each layer is characterized by its primary function, the modeling paradigm it employs, and the specific regulatory or pedagogical requirement it fulfills. Together, the four layers instantiate the framework's core design principle: symbolic logic governs where legal traceability and human-readable explainability are mandated, while probabilistic inference operates where nuanced, evidence-responsive adaptation is pedagogically essential.

Table 2. AIPAQ Architecture: Four-Layer Summary

Layer	Dominant Paradigm	Primary Function	Regulatory & Curricular Alignment
1. Regulatory Symbolic	Symbolic (hard rules)	Non-overridable constraint enforcement on all upstream recommendations	Law 25 (ss. 5, 12.1, 3.3), OACIQ, RBQ, PEA, CEEC traceability
2. Curricular Hybrid	Hybrid (BN + soft rules)	Probabilistic competency state modeling aligned with <i>devis ministériel</i> hierarchy (Tadlaoui et al., 2017)	MES program specifications; CEEC outcome documentation
3. Micro-Pedagogical Probabilistic	Probabilistic (MEBN)	Real-time individualized recommendation, error detection, and risk flagging (Tadlaoui et al., 2019c)	Operates within the constraint space defined by Layers 1 and 2
4. Symbolic Explanation	Symbolic (counterfactual rules)	Human-readable decision justification for learners and instructors (Tadlaoui et al., 2019a)	Law 25, s. 12.1 (right to explanation); CEEC audit documentation; human override support

Figure 1 provides a visual overview of the AIPAQ framework's four-layer architecture. Moving from bottom to top, the diagram illustrates how each layer builds upon the constraints enforced by the layer below: the regulatory symbolic layer (Layer 1) acts as a non-overridable foundation, the curricular hybrid layer (Layer 2) aligns probabilistic competency modeling with the *devis ministériel* hierarchy, the micro-pedagogical probabilistic layer (Layer 3) generates real-time individualized recommendations within that constraint space, and the symbolic explanation layer (Layer 4) translates probabilistic outputs into human-readable, Law 25-compliant justifications. The bidirectional arrows between layers indicate that posterior estimates and explanations flow upward, while regulatory constraints propagate downward.

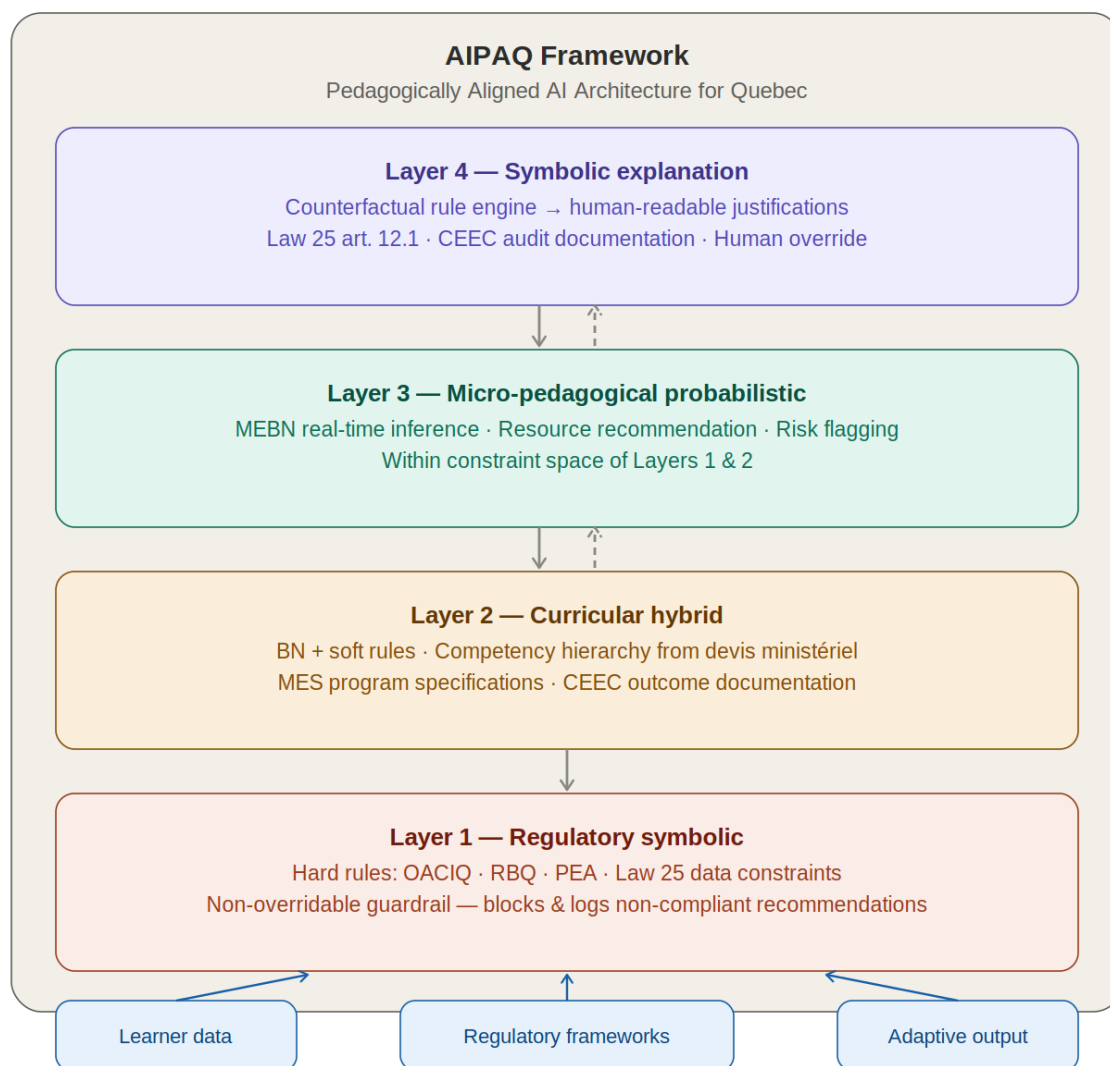


Figure 1. The AIPAQ Four-Layer Hybrid Architecture for Regulated Competency-Based Education.

7) DISCUSSION

7.1. Theoretical Contributions

The AIPAQ framework contributes to the AIED literature in three respects. First, it introduces regulatory compliance as a first-class architectural constraint in learner modeling design, operationalizing the transparency–accuracy trade-off not as a technical optimization problem but as a legal design requirement. The framing of Law 25 section 12.1 as a binding explainability specification transforms the architecture decision from a choice between paradigms into a structural allocation problem: which layer of the system does each paradigm serve best, given its legal obligations?

Second, the paper demonstrates that APC curricular frameworks impose structural constraints on BN architecture not addressed by existing learner modeling approaches. The process of embedding the ministériel competency hierarchy in the DAG structure —

rather than allowing network structure to be learned from data — represents a principled extension of the UML-to-BN methodology developed in Tadlaoui et al. (2015a) and the overlay-BN integration explored in Tadlaoui et al. (2019a), into a regulatory-compliant design workflow. The initialization strategy formalized in Tadlaoui et al. (2017) is directly applicable to Layer 2, addressing the cold-start problem in small-cohort AEC environments.

Third, the four-layer architecture (see Figure 1) formalizes the hybrid model concept more structurally explicitly than prior proposals in this authors' research program (Tadlaoui & Khaldi, 2019a), assigning each paradigm to dedicated layers with well-defined interfaces. This enables modular development, layer-specific validation, and incremental deployment — properties particularly valuable in the resource-constrained environments of private non-subsidized colleges.

7.2. Practical Implications for Quebec Collegiate Institutions

For private non-subsidized colleges, the AIPAQ framework recommends a phased implementation strategy. An initial deployment focused on Layers 1 and 2 delivers immediate pedagogical value and full regulatory compliance without requiring the data volumes or technical infrastructure needed to operate Layer 3 at scale. This phased approach is consistent with the practical deployment experience reported in Tadlaoui et al. (2019c). As cohort data accumulates, Layer 3 is progressively activated, with the MEBN's cross-cohort inference capacity (Tadlaoui et al., 2019b) mitigating the small-sample limitations characteristic of AEC programs.

For cégeps and larger collegiate networks, AIPAQ's modular architecture supports collaborative data pooling across institutions while maintaining per-institution Law 25 compliance: the MEBN's relational structure allows posterior distributions to be informed by cross-institutional learner data, provided appropriate data-sharing agreements satisfy the consent and purpose limitation requirements of Law 25.

7.3. Contribution to the Broader AIED Research Agenda

As AI governance frameworks proliferate globally — including the EU AI Act's high-risk classification of AI systems used in education, and emerging Canadian frameworks including the proposed federal Artificial Intelligence and Data Act (AIDA) — the AIED community faces a growing need for architectures designed with regulatory compliance as a primary requirement. The AIPAQ framework offers a generalizable methodology: identify the binding regulatory constraints on transparency and data use; assign symbolic logic to enforce these constraints at the appropriate architectural layer; deploy probabilistic inference within the constraint space; and provide a dedicated symbolic explanation layer to satisfy the explanation right. This methodology is applicable, with appropriate adaptation of the regulatory symbolic layer, to any jurisdiction imposing algorithmic transparency requirements on automated educational decision-making.

From the perspective of the present authors' prior research program, the AIPAQ framework represents the first systematic attempt to situate the probabilistic learner modeling architecture developed across Tadlaoui et al. (2019c) within a specific, legally demanding institutional context. The resulting framework demonstrates that the technical advances achieved in that prior work — probabilistic ontologies, UML-driven BN construction, stereotype-based initialization, MEBN-based multi-learner inference, and hybrid overlay-BN architectures — are not merely theoretically elegant but practically deployable in regulated real-world educational systems, provided they are embedded within an appropriate governance-aware architectural design.

7.4. Limitations

Several limitations warrant explicit acknowledgment. First, AIPAQ is a conceptual framework that has not been empirically validated. The claims regarding regulatory compliance, pedagogical effectiveness, and technical feasibility are analytically derived, not empirically demonstrated. Validation requires system implementations, deployment in live AEC programs, and rigorous evaluation of learning outcomes, instructor adoption, and Law 25 compliance under CAI review.

Second, the regulatory analysis is current as of the date of this article. Law 25 is subject to ongoing elaboration through implementing regulations, and both the OACIQ and RBQ competency frameworks are subject to periodic revision. The AIPAQ framework's regulatory symbolic layer must accordingly be treated as a living component requiring systematic review at each regulatory update cycle.

Third, the feasibility of the cross-disciplinary expertise required — integrating MEBN modeling, pedagogical engineering, regulatory analysis, and software development — in the resource context of a small private college has not been assessed.

A fourth limitation concerns implementation complexity. The four-layer architecture presupposes a level of cross-disciplinary expertise — combining MEBN probabilistic modeling, pedagogical engineering, regulatory analysis, and software integration — that may be difficult to assemble within a single institution, particularly a small private college. The modular design mitigates this to some extent, since Layers 1 and 2 can be implemented independently of Layer 3. Nevertheless, the human and technical resources required for even a partial deployment should not be underestimated, and a realistic cost-benefit analysis for institutions of varying sizes represents an important direction for future practical research.

A significant technical challenge within this multi-layered framework concerns the cumulative computational latency and execution overhead introduced by stacking deterministic rules on top of dynamic network updates. In a production environment with high concurrent enrollment, a single adaptive learning recommendation must travel sequentially through multi-entity probabilistic propagation in Layer 3, constraint evaluation filtering in Layer 1, and counterfactual explanation generation in Layer 4 before rendering to the client interface. While our previous implementation of a

probabilistic learner model within the COPROLINE platform demonstrated that real-time Bayesian Network propagation is technically viable without prohibitive overhead, the addition of sequential symbolic processing layers introduces an architectural bottleneck. If every micro-adaptation requires real-time graph updates combined with complex cross-layer rule evaluation, the system risks experiencing noticeable response delays that could degrade user experience. Future engineering benchmarks must assess the latency impacts of this tight coupling and explore optimization strategies, such as asynchronous explanation caching or rule-gated module activation, to maintain high-throughput scalability.

7.5. Future Research Directions

Four primary directions emerge. First, empirical validation: comparative studies measuring learning gains, instructor satisfaction, and system compliance across AIPAQ-implementing and conventional adaptive systems in AEC programs. Second, PIA methodology: development of a Privacy Impact Assessment template specifically designed for AI-driven learner modeling systems in the Quebec collegiate context. Third, MEBN parameterization in small-cohort settings: investigation of regularization strategies and cross-institutional data pooling architectures that enable reliable probabilistic inference with limited datasets, building on the initialization strategies established in Tadlaoui et al. (2017). Fourth, comparative regulatory analysis: extension of the present framework to other competency-based educational systems operating under emerging AI governance frameworks, to assess the generalizability of the AIPAQ design methodology.

8) CONCLUSION

This article has argued that the design of AI-driven learner modeling architectures cannot be separated from the regulatory, institutional, and curricular contexts in which those systems are deployed. Quebec's collegiate sector presents a demanding and underexplored instance of this principle: Law 25's algorithmic transparency requirements, the competency-based curricular architecture of AEC and DEC programs, and the prescriptive professional standards of the OACIQ and RBQ collectively impose constraints that no existing AIED architecture has been designed to jointly satisfy.

The comparative analysis across six contextually grounded dimensions reveals that neither symbolic models nor probabilistic models, taken individually, satisfy the full constraint set. Drawing on a decade of empirical and theoretical work on Bayesian learner modeling (Tadlaoui et al., 2020), this paper proposes the AIPAQ framework, a four-layer hybrid architecture that deploys each paradigm where its structural properties best serve the relevant requirement: symbolic logic at the regulatory constraint and explanation layers, where traceability and human-readability are legally mandated; probabilistic inference at the competency modeling and micro-pedagogical layers, where nuanced, evidence-responsive adaptation is pedagogically essential.

Beyond its specific contribution to Quebec collegiate AI design, this work advances a broader argument: that the AIED field must treat the governance frameworks of the jurisdictions in which its systems are deployed as constitutive design constraints. As AI governance frameworks proliferate and strengthen globally, the capacity to design learner modeling architectures that are simultaneously powerful, transparent, and compliant will become a core competency of the AIED research and practice community. The AIPAQ framework is offered as a contribution toward that capacity, grounded in and extending a substantial prior research program in probabilistic learner modeling.

REFERENCES

- Anderson, J. R., Boyle, C. F., & Reiser, B. J. (1985). Intelligent tutoring systems. *Science*, 228(4698), 456–462.
- Arroyo, I., Woolf, B. P., Burelson, W., Muldner, K., Rai, D., & Tai, M. (2014). A multimedia adaptive tutoring system for mathematics that addresses cognition, metacognition and affect. *International Journal of Artificial Intelligence in Education*, 24(4), 387–426.
- Baker, R. S., & Inventado, P. S. (2014). Educational data mining and learning analytics: Potentials and possibilities for online education. In A. Krouska (Ed.), *Learning Analytics* (pp. 61–75). Springer.
- Brown, J. S., & Burton, R. R. (1978). Diagnostic models for procedural bugs in basic mathematical skills. *Cognitive Science*, 2(2), 155–192.
- Brusilovsky, P., & Millán, E. (2007). User models for adaptive hypermedia and adaptive educational systems. In P. Brusilovsky, A. Kobsa, & W. Nejdl (Eds.), *The Adaptive Web* (pp. 3–53). Springer.
- Carr, B., & Goldstein, I. P. (1977). *Overlays: A theory of modelling for computer aided instruction* (MIT AI Memo 406). Massachusetts Institute of Technology.
- Commission d'accès à l'information du Québec. (2022). *Guide pratique — Processus d'évaluation des facteurs relatifs à la vie privée*. CAI.
- Conati, C., Gertner, A., & VanLehn, K. (2002). Using Bayesian networks to manage uncertainty in student modeling. *User Modeling and User-Adapted Interaction*, 12(4), 371–417.
- Corbett, A. T., & Anderson, J. R. (1994). Knowledge tracing: Modeling the acquisition of procedural knowledge. *User Modeling and User-Adapted Interaction*, 4(4), 253–278.
- Drachsler, H., & Greller, W. (2016). Privacy and analytics: It's a DELICATE issue. In *Proceedings of the Sixth International Conference on Learning Analytics & Knowledge* (pp. 89–98). ACM.
- Holmes, W., Bialik, M., & Fadel, C. (2019). *Artificial Intelligence in Education: Promises and Implications for Teaching and Learning*. Center for Curriculum Redesign.
- Jabareen, Y. (2009). Building a conceptual framework: Philosophy, definitions, and procedure. *International Journal of Qualitative Methods*, 8(4), 49–62.
- Loi modernisant des dispositions législatives en matière de protection des renseignements personnels, L.Q. 2021, c. 25 (Bill 64). Assemblée nationale du Québec.
- Loi modernisant des dispositions législatives en matière de protection des renseignements personnels, L.Q. 2022, c. 22 (Bill 3). Assemblée nationale du Québec.
- Mahdi, K. (2024). *Approches pédagogiques dans les systèmes éducatifs anglophones et francophones : Une analyse comparative des défis et opportunités d'hybridation*. *Revue Sciences & Éducation*, 1. <https://doi.org/10.5281/zenodo.15120060>
- Millán, E., Loboda, T., & Pérez-de-la-Cruz, J. L. (2010). Bayesian networks for student model engineering. *Computers & Education*, 55(4), 1663–1683.
- Ministère de l'Enseignement supérieur du Québec. (2022). *Élaboration des programmes d'études techniques: Guide de conception et de production*. MEQ.
- Mislevy, R. J. (1994). Evidence and inference in educational assessment. *Psychometrika*, 59(4), 439–483.

- Mouza, C., & Lavigne, N. (Eds.). (2013). *Emerging Technologies for the Classroom: A Learning Sciences Perspective*. Springer.
- Nkambou, R., Bourdeau, J., & Mizoguchi, R. (Eds.). (2010). *Advances in Intelligent Tutoring Systems*. Springer.
- OACIQ. (2023). *Référentiel de compétences — Permis de courtage immobilier*. Organisme d'autoréglementation du courtage immobilier du Québec.
- Page, M. J., McKenzie, J. E., Bossuyt, P. M., Boutron, I., Hoffmann, T. C., Mulrow, C. D., & Moher, D. (2021). The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. *BMJ*, 372, n71.
- Régie du bâtiment du Québec. (2023). *Qualifications professionnelles requises pour l'exercice des activités en inspection de bâtiments*. RBQ.
- Reeves, T. C. (2006). Design research from a technology perspective. In J. V. D. Akker, K. Gravemeijer, S. McKenney, & N. Nieveen (Eds.), *Educational Design Research* (pp. 52–66). Routledge.
- Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). 'Why should I trust you?': Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 1135–1144). ACM.
- Tadlaoui, M. A., Aammou, S., & Khaldi, M. (2015a). Development of Bayesian Networks from Unified Modeling Language for learner modelling. *International Journal of Advanced Computer Science and Applications*, 6(2).
- Tadlaoui, M. A., Carvalho, R. N., & Khaldi, M. (2017). The initialization of the learner model combining the Bayesian networks and the stereotypes methods. *International Journal of Advanced Computer Research*, 7(33), 200–212.
- Tadlaoui, M. A., Carvalho, R. N., & Khaldi, M. (2019a). Combining the overlay model and Bayesian Networks to determine learning styles in AHES. In M. Auer & T. Tsiatsos (Eds.), *Handbook of Research on E-Assessment in Higher Education* (pp. 195–218). IGI Global.
- Tadlaoui, M. A., Carvalho, R. N., & Khaldi, M. (2019b). Managing the learner model with multi-entity Bayesian Networks in adaptive hypermedia systems. In D. Burgos (Ed.), *Cognitive Computing in Technology-Enhanced Learning* (pp. 89–108). IGI Global.
- Tardif, J. (2006). *L'évaluation des compétences: Documenter le parcours de développement*. Chenélière Éducation.
- VanLehn, K. (2011). The relative effectiveness of human tutoring, intelligent tutoring systems, and other tutoring systems. *Educational Psychologist*, 46(4), 197–221.
- Wenger, E. (1987). *Artificial Intelligence and Tutoring Systems: Computational and Cognitive Approaches to the Communication of Knowledge*. Morgan Kaufmann.
- Winne, P. H., & Baker, R. S. (2013). The potentials of educational data mining for researching metacognition, motivation and self-regulated learning. *Journal of Educational Data Mining*, 5(1), 1–8.
- Zawacki-Richter, O., Marín, V. I., Bond, M., & Gouverneur, F. (2019). Systematic review of research on artificial intelligence applications in higher education. *International Journal of Educational Technology in Higher Education*, 16(1), 39.