

PROCEEDINGS | *African Engineering Education Forum (AEEF 2026)*

Self-evolving engineering curricula: A Reinforcement Learning architecture to align academic training with industry demands in developing countries

Leslie Erika Kouamouo Ndangang ¹, Justin Moskolai Ngossaha ^{*,1} and Adolphe Ayissi Etémé²

¹Department of Mathematics and Computer Science, Faculty of Science, University of Douala, PO Box 24157, Douala, Cameroon.

²IUT of Ngaoundéré, University of Ngaoundéré, PO Box 455, Ngaoundéré, Cameroon.

*Corresponding author: kouamouoerika@gmail.com

(Received: 18 June 2026; Revised: 06 August 2026; Accepted: 13 August 2026; Published: 17 August 2026)

(Editor: Dr. Bappa Muktar; Reviewers: Anonymous reviewers)

Abstract

Engineering education in developing countries faces a widening gap between academic curricula and labor market needs, as traditional revision processes are too slow to keep pace with rapid technological change. This paper proposes a conceptual, self-evolving curricular architecture based on RL, formalizing academic programs as a Markov Decision Process in which curricular adjustments are actions, academic and employability indicators are state variables, and a multi-criteria reward function balances performance, employability, industrial alignment, and budget constraints. An Explainable AI layer and a human-governed workflow support institutional trust. As a proof-of-concept contribution, no deployment, dataset, or simulation has been carried out; the paper instead formalizes the problem, details the architecture, and outlines a validation protocol for future empirical work.

Keywords: Reinforcement Learning; Self-Evolving Systems; Artificial Intelligence in Education; Employability; Skills-to-Market Alignment

Abbreviations: DC: Developing Countries & RL: Reinforcement Learning & AI: Artificial Intelligence & MDP: Markov Decision Process & XAI: Explainable Artificial Intelligence

1. Introduction

Modern engineering education strives to meet technological skill requirements while fostering innovation and economic growth. In developing countries, this concept integrates critical dimensions such as adaptation to limited resources, equitable access to quality training, and the gradual integration of global technological standards with local industrial realities [1, 2, 3]. Educational ecosystems in these regions are characterized by a rapidly growing student population, underdeveloped training infrastructures, and a marked disconnection from real-time labor market data [4, 5]. These structural constraints hinder the direct application of curricular revision models that have proven effective in more developed contexts. Indeed, academic management in these areas faces major obstacles, including the prevalence of slow bureaucratic processes, the exclusion of emerging skills such as AI or green energy from official databases, and limited inclusion of industrial stakeholders in pedagogical

decision-making [6].

In response to this inertia, a paradigm shift toward more agile systems has become imperative. Consequently, a central question arises: to what extent can the integration of RL support the design of next-generation “self-evolving” curriculum systems specifically tailored to the constraints of developing countries? In other words, how can traditionally static academic governance be transformed into a dynamic mechanism capable of optimizing training pathways based on changing market needs?

The objective of this work is to examine this technological transition by combining a comprehensive literature review with the formal design of a conceptual prototype architecture in a representative context, namely Cameroon. This research aims to argue for the technical feasibility of an integrated approach in which RL enables increased visibility of employability trends and continuous optimization of curricular parameters [7, 8], and to lay out a roadmap from conceptual design toward eventual, empirically validated real-world deployment via a lightweight, interoperable infrastructure [9].

The remainder of this article is organized as follows: Section 2 reviews the state of the art in engineering education and the opportunities offered by AI integration. Section 3 formalizes the curricular-governance problem and details the proposed architecture. Section 4 outlines the envisioned validation protocol, expected contributions, and governance safeguards. Section 5 then discusses the paper’s contribution and outlines future research directions, and finally, Section 6 concludes with a summary and recommendations.

2. State of the art

2.1 Sustainable engineering education, employability, and digital opportunities in developing countries

Sustainable engineering education aims to meet technical training needs while minimizing social and economic skill mismatches. In developing countries (DCs), this concept also integrates the management of obsolete curricula, equitable access to digital and pedagogical resources, and the progressive integration of global technological standards with local industrial realities [10, 11, 12, 13, 14].

Academic institutions in DCs face massive student enrollment, undersized training infrastructures, and a lack of reliable data on actual labor-market needs [15, 16, 1, 17]. These constraints complicate the direct application of curricular revision models that have proven successful in developed countries. In particular:

- **Administrative Inertia:** Program update processes are often too slow and bureaucratic to keep pace with rapid technological cycles (e.g., generative AI) [18, 19];
- **Lack of Data Granularity:** Specific skills demanded by the local industrial sector (SMEs, technical informal sector) are frequently absent from official academic frameworks [20];
- **Technical Gap:** Unequal access to cutting-edge tools and the low adoption rate of Learning Analytics solutions hinder the use of artificial intelligence for optimizing training pathways [21].

Initiatives such as the UNESCO framework on AI and education demonstrate that the use of modular and interoperable systems is essential for adapting to local conditions [22, 23, 24]. Learning Analytics and AI-based frameworks have been used to deploy web-accessible predictive analysis tools, thereby facilitating decision-making for training managers and academic leaders [25, 26, 27, 28]. However, several challenges remain, such as the lack of structured occupational data, high digital-infrastructure and computational costs, and the difficulty of integrating RL into rigid pedagogical frameworks. This motivates rethinking curricular engineering in developing countries through lightweight, participatory, and AI-supported digital tools.

2.2 RL and self-evolving systems

Foundations of Reinforcement Learning

Reinforcement Learning enables software agents to refine their behaviors by interacting with changing environments [29]. Foundational work on policy gradients and Q-learning [30, 31] established the mathematical basis for training agents through rewards and penalties [32]. As the field has evolved, innovations such as Proximal Policy Optimization (PPO) [33] and Deep Q-Networks (DQN) [34] have significantly enhanced autonomous decision-making capabilities [35, 36]. A significant advantage of RL lies in its ability to adapt to unforeseen situations, making it a candidate approach for self-programmable AI. However, challenges such as sample inefficiency, high computational requirements, and learning instability remain major obstacles that current research seeks to address via model-based RL or hierarchical RL.




Figure 1. Foundations of Reinforcement Learning Algorithms: DQN and PPO (schematic, adapted for illustration).

Figure 1 summarizes the two RL families referenced throughout this paper.

2.2.1 Application to self-evolving curricula

In this context, a self-evolving curriculum is defined as an *adaptive curriculum decision-support system*: a model-based representation of a training pathway that supports the real-time acquisition, analysis, and human-approved adjustment of teaching modules based on market signals [37]. The term “self-evolving” refers to the continuous update of the underlying model and its recommendations, not to unsupervised institutional change; the human-governance requirement is explicit. This technology is intended to be powered by RL, where the agent learns to optimize a “policy” (the curriculum) by receiving rewards based on success indicators such as placement rates or skill alignment. The envisioned system architecture relies on microservices, facilitating incremental deployment and the integration of heterogeneous data sources through RESTful APIs.

2.2.2 Interoperability and data collection

To ensure seamless interaction between stakeholders (universities, students, employers), semantic ontologies are used to structure educational entities (skills, course units, job profiles) and support complex relational and temporal queries [38, 39]. Self-evolving systems of this kind are designed to integrate data from professional platforms (LinkedIn, Indeed), alumni feedback (crowdsourcing), and public databases. The use of lightweight solutions, tailored for resource-constrained environments, is intended to ensure participatory data collection and a more inclusive alignment between academic training and field realities [40].

2.3 AI and sustainable education

Artificial Intelligence (AI) plays a key role in transforming educational systems through its predictive and prescriptive capabilities. Models such as Recurrent Neural Networks (LSTM), Transformer architectures, and especially Reinforcement Learning (DQN, PPO) are used in the literature to predict learning trajectories, optimize training pathways, and dynamically adjust pedagogical content based on industrial demand [9]. Such tools can, in principle, generate indicators such as predictive employability rates, skill-alignment scores, and program saturation points, and can support “what-if” scenario analysis to test the impact of new subjects before their official deployment [41].

To ensure transparency and faculty buy-in, Explainable AI is also proposed to make automated recommendations understandable to non-technical stakeholders, such as academic directors or certification committees [42]. More broadly, probabilistic and predictive modeling has been applied to other education-related human factors, such as teacher self-efficacy [43], illustrating the broader relevance of data-driven predictive modeling to educational governance, of which the RL-based curricular model proposed here is one example. Finally, recent research explores the use of Large Language Models (LLMs) to automatically generate contextual competency frameworks and anticipate the evolution of local professions.

Table 1. Review of Academic Works on AI, Reinforcement Learning, and Digital Twin Approaches for Engineering Curriculum Alignment (2022-2026)

Authors, Year	Main Topic	Methodology	Tools / Algorithms	Main Limitation
Liu & Song, 2025	Dynamic alignment between education and industry	RL modeling	Reinforcement Learning	Limited to micro-level vocational training tasks
Stasolla et al., 2025	Deep Learning & RL in higher education	Scoping review	DL + RL	Focused on the student micro-level
Satheeskumar et al., 2025	AI emerging skills alignment	Applied AI system	AI-driven alignment	Contextual implementation challenges
Nkrumah et al., 2025	Skill extraction and matching	Comparative study	BERT vs LLM	Strong dependency on textual data
Cousineau, 2025	Explainable AI in education	Doctoral research	XAI frameworks	Implementation complexity
Oeij et al., 2024	Industry 5.0 skills	Conceptual framework	Industry 5.0 taxonomy	Not integrated into an adaptive model
Perisic et al., 2023	Educational interoperability	Conceptual framework	Open interoperability model	Lack of sequential optimization
Sivamayil et al., 2023	RL applications	Systematic review	Q-learning, PPO	Limited curricular applications
Xu et al., 2022	Big Data & employability	Empirical study	Big Data analytics	Static predictive model

This synergy between AI and self-evolving systems suggests a promising direction for addressing curricular obsolescence and professional exclusion in high-growth regions. Beyond simple predictive analysis, the convergence of AI and instructional engineering as argued in [8] could enable a strategic, adaptive, and explainable alignment of training with local socio-economic realities. This direction, aligned with the Industry 5.0 skill framework [20], motivates the architecture proposed in Section 3.

2.4 Global objective and positioning: identified limitations in the literature and architectural justification

Analysis of recent studies highlights that, despite significant advances in the use of AI in education, several limitations persist, particularly in developing countries. Existing solutions often struggle to adapt to contexts characterized by institutional inertia, insufficient digital infrastructure, and fragmented data regarding the local labor market [15, 17, 6]. Furthermore, many approaches remain focused on descriptive or predictive analysis of student performance without proposing a sequential and adaptive curricular optimization mechanism [14, 8].

Moreover, although RL models demonstrate strong potential for dynamic optimization [29], their algorithmic complexity and lack of decision-making transparency still represent major barriers to institutional adoption. The issues of explainability and trust in intelligent systems thus remain central [42].

It is precisely to address these gaps that this research proposes an integrated methodological framework. This framework combines RL-based formalization, an interoperable modular architecture [39], dynamic alignment with industrial skills [20], and XAI mechanisms, with the goal of designing a self-evolving, contextualized, and sustainable curriculum decision-support system.

While RL-based strategies for industry-teaching integration already exist, such as the one proposed in [37], that work targets micro-level vocational-training tasks within a single, highly structured training environment. The present study differs from [37] along three concrete dimensions:

- (i) **Decision granularity.** [37] optimizes individual training exercises within one course; the framework proposed here treats an entire degree program—including module additions, removals, updates, and credit-hour reallocation—as the unit of optimization (*Macro-Level Curricular Optimization*).
- (ii) **Transparency.** [37] does not report an explainability mechanism; the proposed framework couples the RL formalization with a SHAP-based explainability layer specifically to support institutional, non-technical decision-makers (*Explainable AI Integration*).
- (iii) **Contextual constraints.** [37] assumes a data-rich, structured vocational setting; the proposed framework is explicitly designed to remain operable under the low data-granularity and bureaucratic constraints typical of developing-country institutions (*Contextual Resilience for Developing Countries*).

Table 1 summarizes this contrast for [37] as “limited to micro-level vocational training tasks”; the full comparison is developed in the three points above. Thus, whereas [37] offers an algorithmic optimization technique for isolated training, this study proposes a comprehensive, scalable *governance framework design*, intended for systemic evaluation in resource-limited academic ecosystems once empirically validated.

3. Proposed methodological framework

3.1 Materials and formal problem definition

Over the last decade, Artificial Intelligence (AI) and Learning Analytics have been reported to reduce the gap between graduate skills and industrial requirements [14, 20]. Recent research underscores the potential of Deep Learning and RL to enhance educational adaptability, with RL offering more dynamic sequential optimization than traditional analytical approaches [29, 9, 37]. To situate this contribution within this expanding research landscape, Figure 2 illustrates the scientific shift from student-centered predictive models toward dynamic systems focused on continuous industrial alignment.




Figure 2. Illustrative trend of publication activity (2019-2025) across three converging research areas.

Figure 2 illustrates the growing, converging interest in AI in education, training-industry alignment, and developing-country contexts, based on a non-systematic, exploratory keyword search of indexed publications; as no database, search string, or inclusion/exclusion criteria were applied, the figure should be read as motivating context rather than as bibliometric evidence. Despite this recent progress, contemporary research focuses almost exclusively on micro-level optimizations, such as individual student tracking, while overlooking the broader institutional and industrial dimensions. Although interoperability and XAI are gaining traction as essential components for academic trust [39, 42], current frameworks rarely integrate them into a cohesive strategy. This gap motivates the multi-objective architecture proposed below, specifically designed for the structural realities of developing nations.

This shift is necessary because truly self-evolving mechanisms remain scarce in the literature. Most existing systems operate merely as static decision-support tools, lacking the continuous learning loops required for autonomous curricular adaptation. Furthermore, RL applications are still largely confined to optimizing student pathways [29, 9], leaving the challenge of institutional-scale program restructuring largely unaddressed. By tackling the specific hurdles of developing countries, ranging from data fragmentation to weak university-industry ties [14, 20], the framework proposed here moves beyond purely descriptive analysis toward a formally specified, scalable governance design.

The proposed methodology follows an integrated four-stage adaptive approach, illustrated in Figure 3: (1) a literature review and needs analysis to identify curricular gaps and emerging market skills; (2) formalization of the curriculum as a Reinforcement Learning environment (Markov Decision Process, MDP); (3) collection, integration, and semantic alignment of academic and industrial data to build a coherent state representation; and (4) implementation of the RL agent’s decisions, interpretation via XAI techniques, and re-evaluation through a continuous, human-governed feedback loop.

3.1.1 Formal specification of the curricular MDP

To answer the request, raised independently by several reviewers, for a fully specified problem definition rather than a purely narrative description, the curricular-governance problem is formalized here as the tuple

$$\mathcal{M} = \langle \mathcal{S}, \mathcal{A}, P, R, \gamma, H \rangle, \quad (1)$$

where $\mathcal{S}$ is the state space, $\mathcal{A}$ the action space, $P(s_{t+1} \mid s_t, a_t)$ the transition function, R the reward function, $\gamma \in [0, 1)$ the discount factor, and H the episode horizon (number of curricular-revision cycles per program lifetime considered).




Figure 3. Proposed methodological approach.

State space S_t . The state vector $s_t \in S$ at decision cycle t is proposed to concatenate four groups of normalized indicators: (i) *academic* indicators (mean grade point average, pass rate, credit-hour distribution across modules); (ii) *employability* indicators (Skill Alignment Score, SAS_t , and six-month placement rate); (iii) *industrial* indicators (partner satisfaction score, number of emerging skills covered); and (iv) *budgetary* indicators (available credit-hour envelope, cost of the current module configuration).

Action space $\mathcal{A}_t$. At each cycle, the agent is proposed to select one action from a discrete set

$$\mathcal{A} = \{\text{add_module}, \text{update_module}, \text{remove_module}, \text{reallocate_credit_hours}, \text{no_change}\}, \quad (2)$$

parameterized by the module or credit block targeted, subject to institutional feasibility constraints (accreditation limits, prerequisite structure, staffing).

Transition dynamics P . The environment is expected to update s_t to s_{t+1} through (a) a deterministic accounting update for credit hours and implementation cost following the chosen action, and (b) a stochastic update of SAS , placement probability, and satisfaction, to be calibrated on historical institutional and labor-market data once such data are collected. Specifying and calibrating P from real data is identified as a required step before any RL agent can be trained, and is explicitly left to future empirical work.

Reward function R_t . The reward is proposed to combine four weighted objective terms and a modification-cost penalty:

$$R_t = w_1 \cdot \Delta \text{Academic}_t + w_2 \cdot \Delta \text{Employability}_t + w_3 \cdot \Delta \text{Industry}_t + w_4 \cdot \text{BudgetCompliance}_t - w_5 \cdot \text{Cost}_t, \quad (3)$$

where each Δ term denotes the normalized ($[0, 1]$) change in the corresponding indicator between cycles $t-1$ and t , $\text{BudgetCompliance}_t \in \{0, 1\}$ penalizes actions that violate the institutional credit-hour envelope, and Cost_t is the normalized administrative and financial cost of implementing the action. The weights $w_1, \dots, w_5$ are design parameters to be elicited from institutional stakeholders (curriculum boards, industry partners) and validated through a sensitivity analysis showing whether recommended actions remain stable under plausible alternative weightings, rather than fixed a priori; no numerical values are reported here because no calibration exercise has yet been conducted.

Episode structure. One episode is proposed to correspond to a full degree-program revision cycle (e.g., $H = 4$ decision points over two academic years), terminating either at $t = H$ or upon a hard budget-constraint violation. The discount factor γ should be selected to reflect the institution's planning horizon and is left as a design parameter.

3.1.2 Explainable AI layer

An Explainable AI layer is proposed to make the (future) agent's recommendations interpretable to non-technical curriculum committees. The design envisions SHAP (SHapley Additive exPlanations) values computed for a surrogate function approximating the agent's action-selection probabilities from the state vector s_t , so that the relative contribution of each state variable (e.g., SAS, budget) to a given recommendation can be reported. As with the reward function, the concrete choice of background dataset, the granularity of explanations (local, per-decision, versus global, policy-level), and the handling of temporal dependencies across cycles are implementation decisions to be made once a trained agent exists.

3.1.3 Envisioned data sources

No dataset has been collected and no institutional deployment has taken place at the time of writing. The framework is designed with the Computer Science and Electrical Engineering departments of an engineering institution in Douala, Cameroon, as an illustrative future application context, chosen for the rapid technological evolution of these fields (AI, embedded systems, renewable energy). The data sources envisioned for a future empirical study are: (i) institutional academic records (curricula, grades, credit structure); (ii) job postings from professional platforms such as LinkedIn and Indeed, to be collected subject to the platforms' terms of service; and (iii) structured feedback from industrial partners and alumni. Before any such collection can occur, the ethical and data-governance conditions including institutional ethics approval, informed consent, and data-minimization safeguards will need to be secured.

3.2 Proposed architecture

To address the limitations identified in the literature, a self-evolving curricular architecture based on RL is proposed, specifically designed to dynamically align engineering training with industrial needs in developing countries. This architecture relies on a modular approach integrating intelligent data collection, semantic skill modeling, a sequential optimization engine, an XAI layer, and a lightweight infrastructure tailored to local constraints. The objective is to transform curricular governance from a static and periodic process into an adaptive, human-governed system driven by data and machine learning. The architecture is organized into seven sequential layers, introduced below and matched one-to-one with the numbered blocks of Figure 4:

1. **Heterogeneous data acquisition layer:** aggregates heterogeneous data from academic, student, and industrial sources. It distinguishes *static features* (student profile, academic history, learning style) from *dynamic features* evolving over time (LMS behavior, cognitive level, emotional state), while integrating macro-contextual data such as labor market trends, industrial needs, and institutional budgetary constraints. This multi-source structuring captures the individual, institutional, and socio-economic dimensions of the educational system.
2. **State representation (S_t):** transforms collected data into the contextual state representation S_t defined in Section 3.1.1, through a processing pipeline (ETL, normalization), semantic extraction via Transformer models, and ontological alignment between industrial skills and academic content.
3. **Data integration and ontological alignment:** structures and semantically aligns academic and industrial records (skills, course units, job profiles) into a coherent ontology-based representation

feeding the state vector.

4. **Reinforcement Learning optimization engine:** the core of the system is intended to be an RL agent (PPO/DQN) instantiating the MDP of Equations (1)–(3). Based on the state S_t , the agent would select *actions* such as adding, updating, or removing modules, as well as resource reallocation, guided by the multi-objective *reward function* of Equation (3).
5. **Curricular prediction and intervention:** the decisions generated by the RL agent would feed into an operational layer dedicated to performance prediction and curricular adjustment, translating algorithmic recommendations into concrete intervention strategies – modification planning, institutional implementation, and monitoring of effects.
6. **XAI:** to ensure transparency and system acceptability, an XAI layer using SHAP (or LIME) is proposed to identify the relative influence of state variables on the agent’s recommendations, reducing the “black box” effect of the RL model for academic decision-makers.
7. **Human-governed feedback loop and dashboard:** after implementation of any curricular adjustment, performance indicators (professional placement, satisfaction, skill-gap evolution) would be collected and fed back into the system, subject to mandatory human curriculum-board approval. This feedback loop is what allows the reward function and learning policy to be updated over time rather than the curriculum being revised in a purely periodic, static manner. A dashboard layer would provide real-time monitoring, “what-if” prospective analysis, and strategic reporting.

Table 2 summarizes, for each of the seven layers above, its expected input, output, candidate supporting technology, and an indicative design-level latency/cost class.

Table 2. Architectural layers, their input/output, and supporting technology.

Layer	Input	Output	Candidate Technology	Design Latency/Cost Class
1. Data Acquisition	Job offers, LMS logs, academic records	Raw data	Scrapy	Low
2. State Representation	Normalized data	State vector S_t	BERT	Medium
3. Data Integration & Ontology	Raw/normalized data	Aligned ontology instance	OWL, Neo4j	Medium
4. RL Optimization Engine	State vector S_t	Action A_t	PPO/DQN (e.g., Stable-Baselines3)	High
5. Prediction & Intervention	Action A_t	Intervention plan	Rule-based post-processing	Low
6. XAI	State vector, policy output	Feature-importance report	SHAP	Medium
7. Feedback & Dashboard	Post-implementation KPIs	Updated reward, reports	Web dashboard	Low

Figure 4 depicts the resulting pipeline. Data flows sequentially through the seven layers, $1 \rightarrow 2 \rightarrow 3 \rightarrow 4 \rightarrow 5 \rightarrow 6 \rightarrow 7$, and the human-governed feedback loop embedded in layer 7 closes the cycle back to layer 1 for the next curricular-revision period.




Figure 4. Proposed seven-layer architecture.

4. Proposed validation protocol, expected contributions and governance

This section replaces what was, in the originally submitted version, presented as a set of simulation results. No computational simulation, dataset, or trained agent exists at this stage; the paragraphs below instead specify the protocol that would need to be executed to empirically validate the architecture of Section 3, together with the governance and ethical safeguards that would need to accompany any such validation.

4.1 Envisioned application scenario

The architecture is designed to be applicable to the Computer Science and Electrical Engineering departments of an engineering institution in Douala, Cameroon, chosen as an illustrative context given the rapid technological evolution in these fields. A future empirical study would need to instantiate the proposed seven-layer architecture, including automated data acquisition from professional platforms, semantic skill extraction, ontology-based knowledge representation, and a PPO/DQN optimization engine using real institutional data. It would also need to report the dataset size, collection window, feature list, and preprocessing pipeline in full, as required for reproducibility.

4.2 Proposed evaluation metrics and baselines

Table 3 lists the key performance indicators the architecture is designed to track, together with how each would be measured and against which baselines it would need to be compared. No values are reported for these indicators in this paper; they define the measurement protocol for future work.

A methodologically important requirement, raised independently by several reviewers, is that the employability indicator must be measured *independently* of the SAS term that the reward function in

Table 3. Proposed key performance indicators (KPIs), their definition, and candidate baselines for future empirical evaluation. No values are reported here; this table specifies the measurement protocol, not obtained results.

Indicator	Proposed definition / measurement	Candidate baselines for comparison
Skill Alignment Score (SAS)	Ontology-based overlap between the skill set taught and the skill set demanded by collected job postings, normalized to $[0, 1]$	No-change (static curriculum); rule-based periodic revision
Six-month employability rate	Share of graduates placed within six months, from independent longitudinal follow-up data (not derived from SAS)	Historical institutional placement records
Emerging-skill coverage	Share of modules covering skills flagged as “emerging” in the industry data	Supervised ranking without RL; classic Learning Analytics
Credit-hour reallocation	Share of total credit hours reassigned by the agent’s actions	Multi-objective optimization without RL
Industrial partner satisfaction	Structured survey of partner organizations	No-change baseline

Equation (3) optimizes; using a SAS-derived projection as a stand-in for employability would make the outcome circular with respect to the optimized objective, and any future empirical study must avoid this by using external, longitudinally observed placement data. Future work should also report means and standard deviations, or confidence intervals, over multiple independent training seeds, together with convergence diagnostics (learning curves, episode returns), rather than single-run figures.

4.3 Illustrative curriculum case study

To assess whether the policy recommended by a trained agent is pedagogically coherent, and not merely numerically optimal, a future validation study should present at least one auditable, worked example for one program (e.g., Computer Science): the baseline module list, the industry skills extracted for that program, the ontology matches, the specific actions recommended by the agent, the resulting credit distribution, prerequisite impacts, faculty and laboratory requirements, accreditation limits, an estimated implementation cost, and the human curriculum board’s decision. This case study is identified here as a required deliverable of future empirical work.

4.4 Human governance and responsible use

The framework is designed to *support*, not replace, human curricular decision-making. Any module addition, update, or removal recommended by the RL agent would require mandatory approval by the institutional curriculum board before implementation. The design anticipates: (i) stakeholder representation (faculty, students, industry, accreditation bodies) in the approval process; (ii) an appeal/override mechanism for board members who disagree with a recommendation; (iii) an audit trail and version control for every adopted curricular change, with rollback capability; (iv) a minimum review interval between successive changes to guard against oscillation and abrupt module turnover; and (v) explicit safeguards protecting foundational/theoretical content from being deprioritized in favor of short-term market trends. Potential biases in the envisioned data sources, including underrepresentation of informal and public-sector employment, regional and language inequities in online job postings, and employer dominance in shaping the reward signal are acknowledged as open risks requiring ongoing monitoring rather than one-time mitigation.

4.5 Ethics, privacy, and data governance

Because the architecture envisions processing sensitive data, including student profiles, academic histories, LMS behavior, learning styles, and inferred emotional state, alongside employer information and content collected from third-party professional platforms, any future empirical implementation would need to secure, in advance: (i) institutional ethics approval (or a documented exemption) for the use of student and academic data; (ii) informed consent and institutional authorization procedures; (iii) a data-minimization, pseudonymization, security, retention, and deletion policy; (iv) a documented legal and ethical basis for collecting data from third-party platforms, compliant with those platforms' terms of service; and (v) access-control and cybersecurity measures for the ontology and model pipeline. No such approvals have been sought at this stage, since no data collection has taken place.

4.6 Limitations of the proposed approach

The framework presented above is deliberately conceptual, and its limitations should be read as an explicit research agenda rather than as caveats appended after the fact.

- **Conceptual stage.** This paper formalizes and proposes an architecture; it does not report an implementation, a trained agent, or empirical results. All quantitative claims made elsewhere in this text (e.g., regarding SAS, placement rates, or reward values) are design specifications, not measured outcomes.
- **Unvalidated transition and reward calibration.** The transition dynamics P and the reward weights $w_1, \dots, w_5$ in Equations (1)–(3) are design parameters, not empirically fitted quantities. Their calibration on real institutional and labor-market data, together with a sensitivity analysis over plausible weightings, is a prerequisite for any confidence in the resulting policy; without it, the reward function remains a hypothesis about institutional priorities rather than a validated objective.
- **Risk of circular evaluation.** As noted in Section 4, employability must be measured independently of the Skill Alignment Score (SAS) term optimized by the reward function. Failing to enforce this separation in future empirical work would make the evaluation partially circular with respect to the optimized objective, artificially inflating apparent performance.
- **Data availability and representativeness.** The data sources envisioned in Section 3.1.3 (institutional academic records, job postings from platforms such as LinkedIn and Indeed, and structured industry feedback) have not yet been collected. Their coverage, quality, and representativeness for developing-country labor markets, which are often characterized by a large informal sector that is underrepresented in online job postings, remain unverified and constitute a first-order risk for the validity of any downstream model.
- **Algorithmic and computational feasibility.** The choice of RL algorithm (e.g., PPO or DQN, as listed in Table 2) and its sample efficiency have not been assessed against the scale and update frequency realistically available in a resource-constrained institutional setting. Sparse decision cycles (a few per academic year) may be insufficient to train a stable policy without simulation-based pretraining or transfer learning, which is not addressed in this paper.
- **Human and bureaucratic complexity.** Any future simulation environment will necessarily simplify the human and bureaucratic factors, including negotiation among stakeholders, informal influence, and resistance to change that shape real pedagogical decision-making, and may therefore only partially capture the institutional dynamics the architecture is meant to support.
- **Scope and generalizability.** The illustrative application context (Computer Science and Electrical Engineering departments at a single institution in Douala) is narrower than the paper's general

framing around “engineering curricula in developing countries.” Consequently, findings from any future single-institution pilot should not be generalized beyond that validated context without a subsequent multi-institutional study, as outlined in Section 5.

5. Discussion and future work

The contribution of this paper is the formalization of curricular governance as an MDP, the design of a seven-layer modular architecture, and a validation protocol that explicitly separates conceptual design from empirical claims. This addresses a gap in prior RL-for-education research, which remains largely focused on student-level rather than institutional-level optimization [29, 9].

Future research will proceed along the following axes, building on the limitations identified above. First, the data-collection and ethics-approval steps must be completed before any model training can begin. Second, the reward-weight elicitation and sensitivity analysis must be conducted with institutional stakeholders. Third, once training data are available, the RL agent should be trained and evaluated against the baselines and KPIs of Table 3, with results reported using multiple independent seeds, convergence diagnostics, and confidence intervals, and with employability measured independently of the optimized SAS term. Fourth, to address data scarcity during early validation, synthetic data generation and an Edge-AI architecture (to reduce computational costs via distributed learning) will be explored. Fifth, Multi-Agent Reinforcement Learning (MARL) will be investigated to let students, teachers, and employers act as agents contributing directly to the reward signal, fostering more collaborative governance. Sixth, adaptive ontologies will be explored to enable the dynamic evolution of skill taxonomies using generative models. Finally, once a single-institution pilot is validated, a multi-institutional study at the national level is planned to evaluate the transferability and robustness of the framework across diversified academic contexts.

6. Conclusion

Engineering education in developing countries continues to suffer from a structural mismatch between academic curricula and labor-market needs, as slow, periodic, and largely descriptive revision processes struggle to keep pace with rapid technological change. Traditional curricular governance mechanisms, hindered by administrative inertia, fragmented industrial data, and limited stakeholder integration, are ill-suited to sustaining this alignment over time.

To address this gap, this paper proposed a conceptual, Reinforcement-Learning-based architecture for curricular governance, formalizing the evolution of academic programs as a constrained sequential optimization problem using a Markov Decision Process, a multi-criteria reward function balancing academic, employability, industrial, and budgetary objectives, and an Explainable AI layer together with an explicit human-governance and ethics framework. Unlike traditional, periodic and descriptive curricular-revision methods, the proposed architecture is designed for continuous adaptation to labor-market needs through the integration of external data, multi-objective decision logic, and a reward mechanism that incorporates industrial stakeholders directly into the optimization loop. Beyond this formalization, the paper’s main contribution lies in reframing curricular management as a formally specifiable, human-governed adaptive-decision problem, capable in principle of learning, adjusting, and optimizing recommendations in a dynamic and uncertain environment, provided the calibration, data-governance, and validation steps identified throughout this work are completed. No implementation has been carried out, and consequently no performance claim is made: the contribution is the formal problem specification, the modular architecture design, its positioning relative to existing RL-based education-alignment work, and a concrete, ethically grounded validation protocol for future empirical work.

This paper is therefore intended as a foundation for, not a substitute for, the empirical work required to demonstrate that data-driven, explainable, and multi-stakeholder curricular governance is achievable in resource-constrained settings. Future work should prioritize, in sequence: securing institutional ethics approval and completing the envisioned data-collection effort; eliciting and validating the reward-function weights with curriculum boards and industry partners through a sensitivity analysis; training and evaluating the RL agent against the proposed baselines and key performance indicators, with employability measured independently of the optimized alignment score to avoid circular evaluation; and, once a single-institution pilot has been validated, extending the framework to a multi-institutional study to assess its transferability and robustness across diversified academic contexts. Longer term, this line of work could be extended toward synthetic data generation and edge-AI deployment to mitigate early data scarcity, Multi-Agent Reinforcement Learning to let students, teachers, and employers act as agents directly shaping the reward signal, and adaptive ontologies capable of evolving skill taxonomies through generative models. Taken together, these steps chart a path from the present architectural proposal toward an empirically validated, trustworthy, and institutionally governed instrument for aligning engineering education with the evolving needs of developing-country labor markets.

References

- [1] David Mhlanga. “Digital transformation of education, the limitations and prospects of introducing the fourth industrial revolution asynchronous online learning in emerging markets”. In: *Discover education* 3.1 (2024), p. 32.
- [2] Ahmed Ayad. *Réussir le transfert technologique: Guide pour les pays en développement*. Chihab, 2025.
- [3] Rahma Sahali. “L’innovation en Algérie: quelle formation dans les écoles d’ingénieurs pour quelles innovations?” PhD thesis. HESAM Université; Ecole Nationale Supérieure de Management (ENSM), 2022.
- [4] Vui-Yee Koon, Sin Kit Yeng, Chong Soon Meng, and Poon Wai Chuen. “Transformative trends and challenges: Navigating disruptive innovation in higher education through bibliometric insights”. In: *Innovations in Education and Teaching International* 63.2 (2026), pp. 356–369.
- [5] Hunter Hughes. “Education Technology & Digital Upskilling in Emerging Economies: Utilizing Innovation for Growth & Development”. In: *Available at SSRN 5719702* (2025).
- [6] Amando Jr Singun. “Unveiling the barriers to digital transformation in higher education institutions: a systematic literature review”. In: *Discover Education* 4.1 (2025), p. 37.
- [7] Nelson Sizwe Madonsela. “Aligning Education and Workforce Training with Industry Needs: A Perspective on Human Capital Development”. In: *Proceedings of the First Australian International Conference on Industrial Engineering and Operations Management*. 2022, pp. 20–21.
- [8] R Satheeskumar, Ch V Satyanarayana, Talatoti Ratna Kumar, and M Suresh. “AI-Driven alignment of educational programs with industry needs and emerging skillsets [J]”. In: *International Journal of Modern Education and Computer Science* 17.3 (2025), pp. 15–28.
- [9] Fabrizio Stasolla, Antonio Zullo, Roberto Maniglio, Anna Passaro, Mariacarla Di Gioia, Enza Curcio, and Elvira Martini. “Deep learning and reinforcement learning for assessing and enhancing academic performance in university students: A scoping review”. In: *AI* 6.2 (2025), p. 40.
- [10] Rama Yusvana. “Addressing the skills gap in technical and vocational training for sustainable socio-economic growth and development”. In: *International Journal of Research and Innovation in Social Science* 8.IIIS (2025), pp. 6311–6325.
- [11] Md Jahangir Alam, SM Ali Reza, Keiichi Ogawa, and Abu Hossain Muhammad Ahsan. “Sustainable employment for vocational education and training graduates: the case of future skills

- matching in Bangladesh". In: *International Journal of Training Research* 22.3 (2024), pp. 266–288.
- [12] Zinan Sharrad Arsha. *An Undergraduate Internship/Project on Industry Academia Linkage and the Identification of Skills Gaps of the Graduating Engineering Students at Tertiary Level Education*. Tech. rep. Independent University, Bangladesh, 2023.
- [13] Md Abu Issa Gazi, Md Kazi Hafizur Rahman, Mohd Faizal Yusof, Abdullah Al Masud, Md Aminul Islam, Abdul Rahman bin S Senathirajah, and Md Alamgir Hossain. "Mediating role of entrepreneurial intention on the relationship between entrepreneurship education and employability: a study on university students from a developing country". In: *Cogent Business & Management* 11.1 (2024), p. 2294514.
- [14] Lin Xu, Jingxiao Zhang, Yiying Ding, Gangzhu Sun, Wei Zhang, Simon P Philbin, and Brian HW Guo. "Assessing the impact of digital education and the role of the big data analytics course to enhance the skills and employability of engineering students". In: *Frontiers in Psychology* 13 (2022), p. 974574.
- [15] Prince Dacosta Anaman, Deborah Morporkor Zottor, and Julius Kumi Egyir. "Infrastructural challenges and student academic performance: Evidence from a developing nation". In: *International Journal of Innovative Science and Research Technology* 7.11 (2022), pp. 1189–1200.
- [16] Boahemaa Brenya. "Higher education in emergency situation: Blended learning prospects and challenges for educators in the developing countries". In: *Journal of Applied Research in Higher Education* 16.4 (2024), pp. 1015–1028.
- [17] Victor Olugbenga Ayoko, Thankgod Peter, and Deborah Oluwaseun Jegede. "Inadequacy of infrastructural facilities in public universities in Nigeria: Causes, effects and solutions". In: *International Journal on Integrated Education* 6.3 (2023), p. 36.
- [18] James Hutson. *The adoption of artificial intelligence and inertia in higher education: Exploring complex resistance to technological change*. Taylor & Francis, 2025.
- [19] Olatunde Fatai Badmus. "Transforming Program Management through Generative Artificial Intelligence". In: *International Journal of Latest Engineering and Management Research (IJLEMR)* 8.12 (2023).
- [20] Peter Oeij, Karolien Lenaerts, Steven Dhondt, Wietse Van Dijk, Doris Schartinger, Sabrina Sorko, and Chris Warhurst. "A conceptual framework for workforce skills for industry 5.0: Implications for research, policy and practice". In: *Journal of Innovation Management* 12.1 (2024), pp. 205–233.
- [21] Kathiravan Ravichandran. "Identifying learning difficulties at an early stage in education with the help of artificial intelligence models and predictive analytics". In: *International Research Journal of Multidisciplinary Scope (IRJMS)* 5.4 (2024), pp. 1455–1461.
- [22] Yali Li, Laura Tolosa, Francklin Rivas-Echeverria, and Ronald Marquez. "Integrating AI in chemical education: Navigating UNESCO global guidelines, emerging trends, and its intersection with sustainable development goals". In: (2025).
- [23] Shafika Isaacs and Sanjaya Mishra. "Smart education strategies for teaching and learning: Critical analytical framework and case studies". In: (2022).
- [24] Giada Lalli. "Defining interoperability: a universal standard". In: *Journal of ICT Standardization* 13.2 (2025), pp. 139–156.
- [25] Carmen Fattore, Michele Buldo, Arcangelo Priore, Sara Porcari, Vito Domenico Porcari, and Mariella De Fino. "A comprehensive overview of heritage bim frameworks: platforms and technologies integrating multi-scale analyses, data repositories, and sensor systems". In: *Heritage* 8.7 (2025), p. 247.
- [26] Shihui Wang, Adzrool Idzwan Bin Ismail, and Panpan Qiao. "Optimized RE-CNN-Based Multi-Objective Decision Framework for Visual Feature Evaluation in Computational Art Analysis and Interactive Media". In: *Decision Making: Applications in Management and Engineering* 7.1 (2024), pp. 752–770.

- [27] Abdul Wali Sirat, Musawer Hakimi, Bilal Himmat, and Wahidullah Enayat. "Artificial intelligence in educational leadership: Strategic, analytical, interactive, and decision-making applications for the digital age". In: *Jurnal Ilmiah Telsinas Elektro, Sipil dan Teknik Informasi* 8.2 (2025), pp. 210–224.
- [28] Sanjai Vudugula, Sanath Kumar Chebrolu, Maniruzzaman Bhuiyan, and Farhana Zaman Rozony. "Integrating artificial intelligence in strategic business decision-making: A systematic review of predictive models". In: *International Journal of Scientific Interdisciplinary Research* 4.1 (2023), pp. 01–26.
- [29] Keerthana Sivamayil, Elakkiya Rajasekar, Belqasem Aljafari, Srete Nikolovski, Subramaniaswamy Vairavasundaram, and Indragandhi Vairavasundaram. "A systematic study on reinforcement learning based applications". In: *Energies* 16.3 (2023), p. 1512.
- [30] Christopher J. C. H. Watkins. "Learning from Delayed Rewards". PhD thesis. University of Cambridge, 1989.
- [31] Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. Cambridge, MA: MIT Press, 1998.
- [32] Mohamed-Amine Chadi and Hajar Mousannif. "Understanding reinforcement learning algorithms: The progress from basic Q-learning to proximal policy optimization". In: *arXiv preprint arXiv:2304.00026* (2023).
- [33] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. "Proximal Policy Optimization Algorithms". In: *arXiv preprint arXiv:1707.06347* (2017).
- [34] Volodymyr Mnih et al. "Human-level control through deep reinforcement learning". In: *Nature* 518.7540 (2015), pp. 529–533. doi: [10.1038/nature14236](https://doi.org/10.1038/nature14236).
- [35] Agus Perdana Windarto, Solikhun Solikhun, and Anjar Wanto. "Enhancing autonomous vehicle navigation in urban traffic using CNN-based deep Q-networks". In: *Journal of Applied Data Sciences* 6.4 (2025), pp. 2565–2581.
- [36] Juan Escobar-Naranjo, Gustavo Caiza, Paulina Ayala, Edisson Jordan, Carlos A Garcia, and Marcelo V Garcia. "Autonomous navigation of robots: optimization with DQN". In: *Applied Sciences* 13.12 (2023), p. 7202.
- [37] Yusong Liu and Jianjun Song. "Research on real-time dynamic adjustment strategy of industry-teaching integration practical training process in higher vocational education based on reinforcement learning". In: *J. Combin. Math. Combin. Comput* 127 (2025), pp. 3161–3176.
- [38] Claudia-Melania Chituc. "Exploring the Role of Semantic Web Technologies and Ontologies in the Digital Education Ecosystems for Education 4.0". In: *2025 IEEE 12th International Conference on E-Learning in Industrial Electronics (ICELIE)*. IEEE. 2025, pp. 1–6.
- [39] Ana Perisic, Ines Perisic, Marko Lazic, and Branko Perisic. "The foundation for future education, teaching, training, learning, and performing infrastructure-The open interoperability conceptual framework approach". In: *Heliyon* 9.6 (2023).
- [40] Idrees Alsolbi. "Enabling Data-Driven Innovation in the Third Sector: A Framework for Leveraging Big Data". In: *Cloud Computing and Data Science* (2026), pp. 126–153.
- [41] Jesus Kombaya Touckia, Nadia Hamani, and Lyes Kermad. "Digital twin framework for reconfigurable manufacturing systems (RMSs): design and simulation". In: *The International Journal of Advanced Manufacturing Technology* 120.7 (2022), pp. 5431–5450.
- [42] Carter Cousineau. "Advancing Trustworthy Artificial Intelligence (AI): Practical Approaches for Enhancing Interpretability and Transparency". PhD thesis. The University of Guelph, 2025.
- [43] Bery Leouro Mbaioosoum, Mahamat Atteib Ibrahim Doutoum, Apollinaire Batoure Bamana, Narkoy Baoutma, et al. "Self-Efficacy Prediction Model Using Bayesian Networks". In: *International Journal on Interactive Systems and Engineering Management (JISEM)* 10.37s (2024). doi: [10.52783/jisem.v10i37s.6447](https://doi.org/10.52783/jisem.v10i37s.6447).