Agentic Learning Ecosystem for Open, Negotiated, and Accountable Education

Julio Cesar dos Reis
University of Campinas
Campinas, Brazil
dosreis@unicamp.br

ABSTRACT

Education is entering an agentic phase in which autonomous systems actively observe learning contexts, negotiate pathways, coordinate assessments, and broker credentials and learning resources. This article proposes the *Agentic Learning Ecosystem* (ALEc): an open, multi-agent educational ecosystem that orchestrates personalized learning trajectories while embedding economic, institutional, and societal constraints into auditable decision cycles. ALEc is positioned as a successor to platform-centric learning management systems, extending them with specialized learning agents, an economic layer for micro-credential negotiation and content licensing, and governance mechanisms supporting explainability, human oversight, and openness. An illustrative scenario highlights hybrid human–agent roles and trust infrastructure, while key challenges and research directions are identified across architecture, pedagogic, economics, governance, and human experience.

KEYWORDS

Agentic education, multi-agent orchestration, competency-based micro-credentials, semantic interoperability and governance

1 INTRODUCTION

The digital education stack has historically advanced by layering new capabilities onto stable institutional forms. Learning management systems stabilized content distribution and administrative workflows. Intelligent tutoring systems introduced adaptive dialogue and individualized feedback under well-bounded pedagogical scripts. Learning analytics connected interaction traces to dashboards and interventions.

A qualitative shift is underway as educational computer systems are starting to behave not only as tools, but as *agents* that plan, coordinate, negotiate, and act across multiple services and stakeholders. This transition motivates the concept of an *Agentic Learning Ecosystem* (ALEc), designed here as an ecosystem mediated by autonomous (LLM) agents that observe, interpret, negotiate, and orchestrate learning experiences in continuous cycles of action, evidence, reflection, and replanning.

An ALEc is not merely an "AI-powered LMS" and cannot be adequately captured by a chatbot interface placed above conventional courseware. The distinctive claim is architectural, pedagogical, and socio-economic: learning is treated as an ongoing, context-sensitive process in which the unit of organization is not "the course as a static product", but an evolving trajectory of competencies, evidence, and opportunities that must be negotiated among learners, educators, institutions, and labor markets.

Agents act on behalf of learners to propose learning pathways, acquire resources, coordinate assessments, and construct credential portfolios. Other agents represent content providers, universities, professional bodies, and employers, exposing machine-readable policies and negotiation. Learning becomes a form of multi-party coordination in which autonomy is distributed across specialized agents and continuously bounded by governance constraints such as transparency, auditability, fairness, and human oversight.

This agentic turn in education does not arise in isolation. It is entangled with the broader emergence of an agentic digital economy in which assistants and service agents reduce communication frictions and mediate transactions more fluidly than current platform-bound interfaces [8]. It is also shaped by an interoperability horizon in which agents communicate through open protocols and shared semantics, echoing the Semantic Web's foundational ambition to make the Web machine-interpretable [1, 6]. Recent work on open agent ecosystems argues that decentralized and semantically interoperable infrastructures can enable heterogeneous agents to coordinate across organizational boundaries, thereby reshaping markets, governance, and power relations. In education, this vision suggests a shift from siloed edtech platforms toward an open educational fabric in which learner agents negotiate across providers, credential authorities, and opportunity marketplaces.

This investigation suggests that the opportunity is therefore to treat the ALEc as both an educational infrastructure and a socio-technical institution. Architecturally, it requires agents equipped with semantic memory, tool-using and planning techniques, and coordination protocols that support negotiation and verifiable execution. Economically, it calls for fine-grained licensing and micro-credential mechanisms that can be composed into individualized portfolios supported by micro-transactions and conditional payments. From a governance standpoint, it demands human-in-the-loop safeguards, explainability, provenance, and measures that prevent new monopolies of discovery, identity, or orchestration. Socially, it raises questions about trust calibration, pedagogical agency, equity, and the changing role of educators in an environment where machines can propose and enact interventions at scale.

2 RELATED STUDIES

The ALEc concept designed in our proposal sits at the intersection of three trajectories: agentic AI in education, the evolution of language agents as cognitive architectures, and the semantic interoperability as a foundation for open multi-agent coordination.

Recent syntheses on agentic AI in education characterize a shift from AI that supports isolated tasks to AI systems that can autonomously pursue educational goals, coordinate sub-tasks, and adapt strategies over time [7]. Our study emphasizes both opportunities and risks: opportunities include personalized assistance, scalable tutoring, and new forms of reflective learning support; risks include misalignment with educational values, emergent bias, and the opacity of autonomous decision-making. Importantly for ALEc, agentic education is framed not only as a technological upgrade, but as a socio-technical transformation that requires participatory governance and alignment with pedagogical principles [7]. The educational AI lineage includes earlier intelligent tutoring systems that pioneered mixed-initiative dialogue and adaptive scaffolding. AutoTutor [5], for example, demonstrated how dialogue-based tutoring can engage learners through conversational strategies.

Applied studies in online learning and competency development increasingly incorporated agentic mechanisms into practical frameworks. For instance, an agentic AI-driven online learning model for digital industry entrepreneurs integrates work-integrated learning and coaching with adaptive AI components, reporting measurable gains in digital literacy and perceived entrepreneurial impact [11]. While such models may remain within a single institutional program, they reveal the demand-side logic of agentic education: learners benefit when learning activities are continuously aligned with evolving contexts, workplace constraints, and competency evidence, rather than being confined to fixed curricula.

A central technical enabler for ALEc is the emergence of language agents that integrate planning, tool use, memory, and multi-step reasoning. Cognitive architectures for language agents argue that "agentic" behavior requires more than a powerful model; it demands a structured loop that connects perception, memory, deliberation, action, and learning [9, 10]. This perspective is foundational for ALEc because education is intrinsically longitudinal: learner goals and constraints evolve, evidence accumulates over months or years, and interventions must be justified not only in the moment, but as part of an interpretable trajectory over time.

If ALEc is to move beyond a single-vendor platform, it must align with open, interoperable standards for agent communication. The idea of an open Web of interoperable agents extends the Web's decentralization principle to autonomous machine actors that communicate and coordinate through open standards This framing draws on the Semantic Web vision [1].

Recent discussions of semantic-driven agent communication emphasize that natural language alone is insufficient for robust interoperability and that agents need shared or translatable semantic structures, provenance, and mechanisms for reliable exchange [12] and negotiation [13]. Protocol initiatives for agent-to-agent interoperability similarly suggest that agent ecosystems depend on standardized methods for discovery, capability advertisement, and secure communication [3, 4].

3 AGENTIC LEARNING ECOSYSTEM

This section presents our proposal for an Agentic Learning Ecosystem (ALEc) as an open, multi-agent educational ecosystem. The core claim is that education can be re-institutionalized as an agent-mediated process in which learning trajectories, credential portfolios, and opportunity pathways are continuously negotiated and governed through auditable decision cycles.

3.1 Definition and Core Principles

In ALEc, autonomous agents act on behalf of learners, educators, institutions, and market stakeholders to orchestrate personalized learning trajectories, to validate competency evidence, and to negotiate credential and resource access under explicit governance constraints. The proposal is "agentic by design" in the sense that its primary objects are not "static course pages", but active loops: observation of learner context, hypothesis formation about pedagogical next steps, coordinated action across learning resources, evidence collection, reflective interpretation, and replanning over time. This loop is not a hidden optimization procedure but a transparent, contestable process in which human actors interrogate rationales and override decisions.

Our proposal implies that the education process becomes structurally similar to other agentic domains, in which autonomy is expressed through coordination and negotiation. In the same way that 'open agent ecosystems' envision assistants that negotiate travel bundles through service agents, the ALEc envisions learner assistants that negotiate learning bundles through institutional and content agents.

The educational case is arguably more complex because outcomes are both normative and instrumental: the value of a credential depends on social recognition, and the "best" learning pathway must be evaluated against educational values such as equity, learner autonomy, and meaningful pedagogical understanding. This is why the ALEc concept couples agentic orchestration with governance and explainability as defining requirements, rather than as optional.

The ALEc proposal is guided by four principles that shape its architecture and institutional role. **Learning trajectories** must be dynamic and longitudinal. In many educational settings, learners face interruptions, goal shifts, and changing constraints. Agentic orchestration should therefore treat a pathway as a living plan that can be renegotiated as evidence and context change.

Competencies become the basic unit of organization. Traditional course-centric models bundle multiple outcomes into single products, making it difficult to express granular achievement and to assemble flexible pathways. The ALEc treats competencies, evidence, and contextualized performance as the composable units from which credentials and opportunity matching (in the labor market) can be constructed. This reorientation is a prerequisite for micro-credential negotiation and for learning marketplaces that function at finer resolution than semesters or degrees.

Autonomy must be negotiated and hybrid in nature. Agentic systems should not replace educational judgment; they should redistribute cognitive labor while preserving meaningful human agency. In practice, learners and educators become co-negotiators of trajectories, with agents providing proposals, simulations, and evidence syntheses rather than unilateral decisions. The governance layer must therefore implement human-in-the-loop thresholds and explicit "escalation" policies, analogous to how agents in other domains must know when to act and when to defer [2].

Openness and semantic interoperability are architectural commitments, not integration features. If ALEc is to avoid becoming another siloed edtech platform, agents must communicate via shared semantics and open protocols to enable cross-provider negotiation [12]. This echoes the Semantic Web's vision of machine-readable meaning [1]. Within ALEc, semantic interoperability enables learner agents to interpret credential definitions and competency frameworks across institutions, making educational markets contestable by allowing new providers to join without centralized permission.

3.2 A Layered Architecture for ALEc

A publishable ALEc benefits from a clear architectural stack that separates concerns while supporting end-to-end accountability. The architecture must be compatible with open-agent ecosystem visions, while specializing in educational needs.

At the foundation, an *observation and context layer* collects signals about learning activity, engagement, performance, and constraints. Unlike traditional learning analytics, which primarily feed dashboards, this layer provides a continuous-state representation used by orchestration agents. It must be designed with privacy and consent as structural constraints, as educational traces are sensitive and the system's legitimacy depends on learner trust [7].

Above this sits a *semantic memory layer* that stores longitudinal representations of learner history, competency evidence, and contextual annotations. The role of semantics is not only to index content, but to make learner progress machine-interpretable across multiple agents and institutions. This motivates the use of ontological structures for competencies, evidence types, assessment conditions, and credential requirements [1, 6].

A third layer is *learning planning and orchestration*, implemented by learner-centered agents that propose actions, plan sequence learning objects, and trigger assessment opportunities. The system must support deliberation over goals and constraints, tool use for resource retrieval, and reflective updating based on evidence [10]. This layer must negotiate with external agents that represent content, credential, and opportunity learning services, and maintain a structured record of why particular actions were proposed.

A fourth layer consists of *communities of specialized agents*, distributing responsibilities across agents with explicit roles. A learner assistant agent focuses on self-regulation and goal refinement; an assessment agent manages evidence and validity; a curation agent evaluates resource fit; a pacing agent monitors cognitive load; a credential agent negotiates recognition; and an opportunity-matching agent aligns competency portfolios with projects, scholarships, internships, or jobs in the labor market.

At the top, a *governance and explainability layer* enforces constraints on agent actions, records decision rationales, and supports auditing. This layer operationalizes human-in-the-loop education by specifying when agent decisions require teacher or learner approval, how disagreements are handled, and how evidence is presented in interpretable forms [7].

4 ILLUSTRATIVE SCENARIO

To make the ALEc vision concrete, we illustrate its capabilities and implications by considering a single, comprehensive scenario that surfaces architectural requirements and socio-economic dynamics.

Consider a learner, Ana, who is a mid-career professional transitioning into data-intensive roles. She has limited weekly time, variable work demands, and a specific goal: to qualify for a project-based internship that requires demonstrable competencies in data preparation, model evaluation, and responsible AI practices. Ana enters the ALEc via a learner assistant agent that first elicits high-level intent and converts it into a structured representation of goals, constraints, and preferences. Rather than fixing a curriculum, the agent treats Ana's learning pathway as a negotiable plan whose shape depends on evidence and on external opportunity requirements.

The learner assistant agent queries an open ecosystem of educational service agents. A university credential agent advertises a micro-credential in "Applied Machine Learning Foundations" with evidence requirements expressed in machine-readable form. A professional body agent advertises a competency badge aligned with industry practice. Several content provider agents advertise granular learning modules, including interactive notebooks, short assessments, and project templates. An employer opportunity agent advertises the internship's competency profile and acceptable credential combinations. In an open agent ecosystem, these interactions are enabled by shared semantics and interoperability protocols, reflecting the broader vision of open, semantically interoperable agent networks.

Negotiation begins immediately. Ana's learner assistant agent proposes a pathway that combines two micro-credentials and a portfolio project. It suggests substituting one assessment format for a workplace artifact, given Ana's prior experience in analytics. The university credential agent responds with constraints: it can accept prior learning evidence if accompanied by a proctored evaluation or a structured portfolio review. The learner assistant agent counterproposes a portfolio review and a short oral defense conducted by an assessment agent under institutional supervision. The negotiation is not purely technical; it is about legitimacy. The institution must trust that evidence is valid, while Ana must trust that requirements are fair and achievable.

At this stage, the teacher's role becomes visible. Ana is assigned a mentor teacher who reviews the proposed plan through a governance interface that displays the negotiation trace: what evidence was proposed, what constraints were invoked, and what trade-offs were made. The teacher does not micromanage every step, but validates the pedagogical coherence of the plan and checks whether it aligns with institutional values, such as avoiding excessive workload and ensuring meaningful assessment. This illustrates the hybrid human–agent model: agents propose and coordinate, whereas human educators remain responsible for educational legitimacy and ethical oversight [7].

As Ana progresses, a learner progress agent aggregates evidence from multiple sources: quiz results, code submissions, reflection notes, and time-on-task signals. When the agent detects declining engagement, it does not simply push reminders. It interprets the pattern as a possible overload and proposes a pacing adjustment. It also suggests a change in learning strategy, shifting from video explanations to hands-on notebook exercises. Because these are consequential interventions, the governance layer prompts Ana for consent and provides an evidence-based explanation. The system's legitimacy depends on Ana experiencing the agent as a collaborator rather than a controller.

Meanwhile, micro-transactions occur in the background. Ana's agent licenses an advanced dataset for a week because her portfolio project requires realistic data. It purchases a small amount of expert feedback from an educator agent for reviewing her model evaluation report. These are conditional, fine-grained transactions.

As the internship deadline approaches, Ana's learner agent negotiates with the employer opportunity agent regarding acceptable evidence. The employer agent specifies that the portfolio must demonstrate not only accuracy metrics but also error analysis and responsible AI considerations. The learner assistant agent then coordinates a targeted assessment sequence with the credential agent and proposes a final capstone review.

5 CHALLENGES AND DIRECTIONS

The ALEc vision is ambitious and would be irresponsible without a clear articulation of the interdisciplinary challenges it entails. This section highlights research directions across technical architecture, pedagogic aspects, economics, governance, and social experience.

Technical Challenges. A first challenge is reliable orchestration under uncertainty. Education is noisy: learners may disengage for reasons not visible in logs, performance measures may be context-dependent, and goals may be underspecified. Agent cognitive architectures emphasize the need for structured loops and memory to sustain coherent behavior [10]. However, ALEc demands additional guarantees because errors can harm learners. Research is needed on uncertainty-aware orchestration that can recognize when the system's confidence is insufficient.

A second challenge is semantic interoperability at scale. Even if the ALEc begins within a single institution, the long-term value of agentic negotiation depends on cross-institutional interaction. This requires shared or aligned ontologies for competencies, evidence, credential conditions, and opportunity requirements. Classic ontology principles remain relevant [6], but ALEc must address the modern reality that meaning is negotiated in language and that LLM-based agents may need to translate between vocabularies dynamically [12]. A promising direction is hybrid semantics: ontologies provide stable anchors for machine-verifiable constraints, while LLMs provide flexible mapping and explanation layers.

A third challenge is negotiation protocols and the enforcement of commitments. Agentic evaluation through negotiation tasks highlights that the ability to negotiate is a core measure of agency [2]. In education, negotiation involves commitments with legal and institutional implications. Micro-credentials, assessment outcomes, and paid feedback services create obligations that must be enforceable and interpretable. Research on real-time negotiation agents suggests that protocols must balance dialogue flexibility with structured contract terms [13]. ALEc must, therefore, develop educational contract primitives that encode assessment validity, identity verification, evidence provenance, and appeal rights.

A fourth challenge concerns interoperability protocols for agent-to-agent communication and discovery. Emerging initiatives aim to standardize how agents advertise capabilities and interact across ecosystems [4]. ALEc would benefit from aligning with such protocols to avoid reinventing closed integration layers. Yet education adds unique requirements: child protection, consent management, enforcement of institutional policies, and long-term data stewardship.

Economic Challenges. The economic layer of ALEc is not optional if the proposal aims to enable granular learning components and micro-credentials. However, agent-mediated markets create new incentive problems by introducing intermediaries for discovery, reputation, and compliance. In education, discovery intermediaries could become gatekeepers, determining which learning resources and credentials are visible to learner agents, thereby recreating platform dominance in a new form. A key research direction is to design decentralized discovery and reputation systems that remain contestable and resistant to capture, consistent with open web-of-agents governance concerns.

Governance Challenges. They are central because ALEc delegates educational decisions to autonomous systems. One challenge is ensuring that explainability is operational rather than rhetorical. Explanations must be tied to evidence traces and policy constraints, not merely post hoc narratives. The ALEc governance must therefore produce auditable decision knowledge graphs that allow stakeholders to easily inspect how data, rules, and agent reasoning produced outcomes. This is consistent with broader calls in agentic education for participatory governance and ethical safeguards [7].

Human and Social Challenges. The ALEc vision changes what it means to be a learner and teacher. A central challenge is trust calibration. Learners must trust agents enough to delegate, but not so much that they become passive. Teachers must trust the system's evidence traces while maintaining professional judgment. In ALEc, controllability is achieved not by limiting agent capabilities, but by negotiating autonomy and providing contestable decision processes. A second challenge is the risk of cognitive outsourcing. If agents handle planning, resource selection, and reflective synthesis, learners may lose opportunities to develop metacognitive skills. ALEc must therefore design for *productive agency*: agents should scaffold learners' capacity to set goals, evaluate evidence, and reflect on strategies.

6 CONCLUSION

This article proposed the vision of an Agentic Learning Ecosystem as an open, negotiated, and accountable educational frame. We argued that the future of digital education lies in LLM-based multi-agent infrastructures that continuously orchestrate learner trajectories through cycles of observation, semantic memory, action, evidence, reflection, and replanning. By aligning with open agent ecosystem principles, integrating micro-credential negotiation and fine-grained educational economics, and elevating governance to first-class architectural concerns, the ALEc paradigm positions trust, interoperability, and human oversight as foundational. Realizing this vision requires interdisciplinary advances spanning technically agentic cognitive architectures, agentic semantic interoperability, agent negotiation protocols, and open standards, alongside institutional innovation in governance and pedagogical roles.

ACKNOWLEDGMENTS

We thank the National Council for Scientific and Technological Development (CNPq), Brazil (grants #301337/2025-0, #400456/2025-7). This research was funded by the São Paulo Research Foundation (FAPESP) (grant #2025/10403-2).[1]

[1]The opinions expressed in this work do not necessarily reflect those of the funding agencies.

REFERENCES

[1] Tim Berners-Lee, James Hendler, and Ora Lassila. 2001. The Semantic Web. *Scientific American* 284, 5 (2001), 34–43.

[2] Thomas R. Davidson, Vladimir Veselovsky, Martin Josifoski, Maxime Peyrard, Antoine Bosselut, Michal Kosinski, and Robert West. 2024. Evaluating Language Model Agency through Negotiations. *arXiv preprint* (2024). arXiv:2401.04536 [cs.CL]

[3] Abul Ehtesham, Aditi Singh, Gaurav Kumar Gupta, and Saket Kumar. 2025. A Survey of Agent Interoperability Protocols: Model Context Protocol (MCP), Agent Communication Protocol (ACP), Agent-to-Agent Protocol (A2A), and Agent Network Protocol (ANP). *arXiv preprint arXiv:2505.02279* (2025). https://arxiv.org/abs/2505.02279

[4] Google. 2025. Agent2Agent Protocol: A New Era of Agent Interoperability. Online. Developer blog post.

[5] Arthur C. Graesser, Paul Chipman, Brian C. Haynes, and Andrew Olney. 2005. AutoTutor: An Intelligent Tutoring System with Mixed-Initiative Dialogue. *IEEE Transactions on Education* 48, 4 (2005), 612–618.

[6] Thomas R. Gruber. 1995. Toward Principles for the Design of Ontologies Used for Knowledge Sharing. *International Journal of Human-Computer Studies* 43, 5–6 (1995), 907–928.

[7] Georgios Kostopoulos, Vasileios Gkamas, Maria Rigou, and Sotiris Kotsiantis. 2025. Agentic AI in Education: State of the Art and Future Directions. *IEEE Access* 13 (2025). https://doi.org/10.1109/ACCESS.2025.3620473

[8] Daniel M. Rothschild, Markus Möbius, Jake M. Hofman, Emma Dillon, Daniel G. Goldstein, Nicole Immorlica, Sonya Jaffe, Brendan Lucier, Aleksandrs Slivkins, and Meir Vogel. 2025. The Agentic Economy. *arXiv preprint* (2025). arXiv:2505.15799 [cs.CY]

[9] Eryck Silva, Frances A Santos, Pedro Thompson, and Julio C dos Reis. 2025. LLM-Powered Conversational Multi-Agent Cognitive System for Collaborative Task Solving. In *Workshop-School on Agents, Environments, and Applications (WESAAC)*. SBC, 59–70.

[10] Theodore R. Sumers, Shunyu Yao, Karthik Narasimhan, and Thomas L. Griffiths. 2024. Cognitive Architectures for Language Agents. *arXiv preprint* (2024).

[11] P. Tantrarungroj and T. Thangkabutra. 2025. A Responsive and Agentic Artificial Intelligence-driven Online Learning Framework for Personalized Digital Competency Development among Digital Industry Entrepreneurs. *Journal of Research in Innovative Teaching & Learning* (2025). https://doi.org/10.1108/JRIT-08-2025-0236

[12] K. Yu, M. Sun, G. Wu, et al. 2025. Semantic-driven AI Agent Communications: Challenges and Solutions. *arXiv preprint* (2025). arXiv:2510.00381 [cs.CL]

[13] Matej Zust, Marko Grobelnik, Andrej Sittar, et al. 2025. Towards AI-powered Real-time Negotiation Agent. In *Proceedings of the 1st Workshop on Semantic Generative Agents on the Web (SemGenAge) at ESWC 2025 (CEUR Workshop Proceedings, Vol. 3977)*. Paper #5 in the SemGenAge 2025 proceedings (CEUR-WS).