

AI-Powered Adaptive Learning Platforms with Credential Verification on Chain

Dr Amit Kumar Jain

DCSE, Roorkee Institute of Technology, Roorkee

Uttarakhand, India

amitkumarjain.cse@ritroorkee.com



Date of Submission: 27-12-2024

Date of Acceptance: 31-12-2024

Date of Publication: 01-01-2025

ABSTRACT

Adaptive learning platforms personalize instruction by continuously modeling each learner's knowledge, skills, and pace. Meanwhile, the integrity of academic records and micro-credentials increasingly depends on trustable, portable verification mechanisms across institutions and labor markets. This manuscript proposes and critically examines a reference architecture that unifies AI-powered adaptivity with on-chain credential verification based on decentralized identifiers (DIDs) and the W3C Verifiable Credentials (VC) data model. The design integrates real-time learner modeling (e.g., Bayesian/Deep Knowledge Tracing), reinforcement-learning sequencing, and learning analytics pipelines (xAPI/Caliper) to generate mastery-aligned learning pathways. When mastery thresholds are met, the system issues verifiable micro-credentials whose cryptographic proofs (hashes, revocation registries) are anchored to a permissioned or public blockchain, enabling instant, privacy-respecting verification without exposing sensitive learner data. We situate the architecture in prior research on intelligent tutoring systems, formative feedback, and digital credentialing; articulate governance, privacy, security, and equity considerations; and provide a design-science methodology covering requirements analysis, component design, and evaluative heuristics. A prototypical evaluation framework is outlined with metrics for learning effectiveness (time-to-mastery, normalized gain), instructional efficiency (content reuse and algorithmic inference cost), trust and integrity (verification latency, fraud resistance), privacy (selective disclosure), and operational feasibility (gas/transaction cost and throughput).

Anticipated results include improved mastery attainment and mobility of skills signals across ecosystems, while the discussion addresses challenges such as bias mitigation in recommendation policies, long-term key management, revocation semantics, and regulatory compliance. The manuscript concludes with a roadmap for pilots and research, emphasizing interoperable standards, transparent AI, and consortium-based governance to advance learner agency and system trustworthiness.

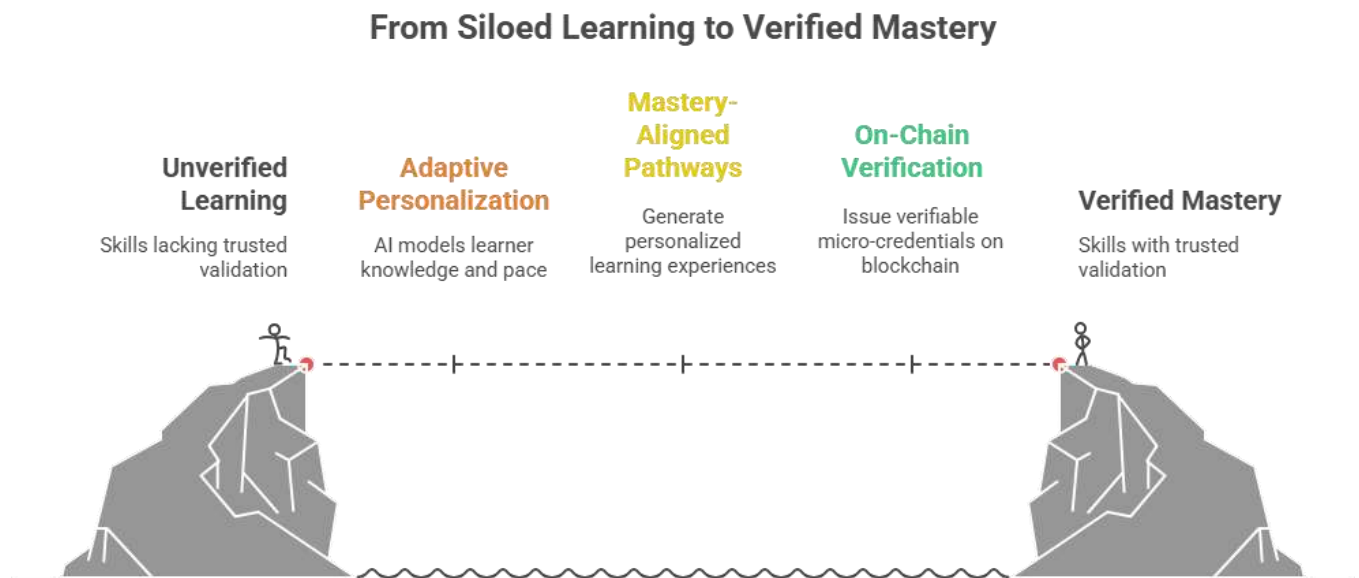


Figure-1.From Siloed Learning to Verified Mastery

KEYWORDS

Adaptive Learning, Intelligent Tutoring Systems, Knowledge Tracing, Reinforcement Learning, Learning Analytics, Verifiable Credentials, Decentralized Identifiers, Blockchain, Privacy, Micro-Credentials

INTRODUCTION

Education systems around the world are navigating two intertwined transitions. First, instruction and assessment are moving toward personalization at scale, as institutions adopt data-driven learning technologies to accommodate divergent backgrounds, goals, and paces. Second, the recognition of learning is becoming modular and portable, with digital micro-credentials, skills badges, and competency-based transcripts complementing (and sometimes substituting for) traditional degrees. These trends create technical and sociotechnical demands: platforms must (a) adapt in real time to each learner's state and context and (b) record and transmit trustworthy evidence of achievement across organizational boundaries.

Adaptive Learning and Credential Verification Process




Figure-2. Adaptive Learning and Credential Verification Process

Adaptive learning platforms attempt to address the first demand by diagnosing competence and dynamically selecting the “next best” activity or explanation. Decades of work on intelligent tutoring systems (ITS), student modeling, and formative feedback have shown substantial effects on learning efficiency and outcomes when adaptivity is accurate, granular, and aligned to well-specified competencies. Yet personalization also raises concerns about opacity, bias, and the provenance of the data used to make educational decisions.

The second demand—trusted recognition—invites a complementary set of technologies. Paper transcripts and isolated databases are fragile and cumbersome for cross-institutional verification. Credential fraud is well documented, and manual verification drains administrative resources. Standards such as W3C Verifiable Credentials and the rise of decentralized identifiers offer a way to issue cryptographically signed attestations that learners can hold and selectively disclose. When proofs of issuance/revocation events are anchored to a tamper-evident ledger, relying parties (employers, universities, licensing boards) can verify authenticity without contacting the issuer or accessing personal data.

This manuscript argues that the real opportunity lies at the intersection: AI-powered adaptive platforms that not only drive mastery but also generate verifiable, privacy-preserving credentials on chain as soon as mastery is demonstrated. Doing so requires careful orchestration among AI models, data standards for learning events, credential schemas, ledger choices, wallet UX, governance, and compliance. It also requires explicit attention to *how* adaptivity decisions are made so that resulting credentials are trustworthy and fair.

We contribute: (1) a literature-grounded synthesis connecting adaptive instruction and verifiable credentialing; (2) a design-science methodology culminating in a modular architecture; (3) an evaluation framework with actionable metrics for pedagogy, integrity, privacy, and operations; and (4) a discussion of risks and mitigations, including bias, revocation semantics, and sustainability.

LITERATURE REVIEW

Adaptive Learning and Intelligent Tutoring

Adaptive learning draws from cognitive psychology and AI in education. Early systems modeled learner knowledge at the skill level and selected targeted practice based on Bayesian knowledge tracing (BKT) (Corbett & Anderson, 1995), while later work introduced deep neural variants such as Deep Knowledge Tracing that can model rich temporal dependencies in response data (Piech et al., 2015). Meta-analyses of ITS suggest that well-designed tutors can approach the effectiveness of human tutoring under certain conditions (VanLehn, 2011), especially when they provide immediate, actionable feedback (Shute, 2008) and align activities to structured learning objectives (Woolf, 2010). Adaptive hypermedia frameworks (Brusilovsky, 2001) highlight the importance of user modeling dimensions beyond correctness—such as motivation, affect, and navigation history.

Contemporary platforms integrate learning analytics to instrument activity streams, enabling instructors and algorithms to observe learning trajectories (Ferguson, 2012; Siemens & Long, 2011). These analytics emphasize formative, not just summative, functions—guiding real-time adjustments and scaffolding. However, they raise challenges of construct validity, transparency of algorithms, and ethical data use.

Competency Models, Mastery, and Assessment

Adaptive instruction presupposes a competency ontology: the granular skills a course aims to build and their prerequisite structure. Hierarchies and learning maps are used to define mastery thresholds and to sequence content. Empirical research emphasizes mastery learning—ensuring high proficiency on prerequisite concepts before advancing—can reduce variability in outcomes and increase overall achievement. Cognitive task analysis and evidence-centered design help link activities to observable evidence of competence, thereby enabling valid automated inferences about mastery (Koedinger et al., 2015; Shute, 2008).

Digital Micro-Credentials and Badges

Micro-credentials capture modular achievements and can motivate progression when they are credible, portable, and meaningful. Open Badges (now under 1EdTech, formerly IMS Global) provides a widely used data model and vocabulary for achievements, endorsements, and evidence links. To combat fraud and enable self-sovereign control, recent initiatives anchor credential proofs to distributed ledgers. The EU Joint Research Centre's report *Blockchain in Education* documents experiments with blockchain-anchored diplomas and cautions about privacy, scalability, and governance (Grech & Camilleri, 2017). MIT's Blockcerts demonstrates practical

patterns for issuing and verifying credentials whose hashes are written to public chains while the credential data remains off-chain under learner control.

Verifiable Credentials, DIDs, and Selective Disclosure

The W3C Verifiable Credentials Data Model defines interoperable structures for cryptographically signed attestations, while Decentralized Identifiers (DIDs) enable identifier portability across systems without a single root authority (Sporny et al., 2022). Privacy-enhancing cryptography such as BBS+ signatures allows holders to selectively disclose only necessary attributes or to demonstrate predicates (e.g., “score ≥ 85 ”) without revealing raw data (Boneh, Boyen, & Shacham, 2004). Revocation registries and status lists provide scalable mechanisms to reflect credential status without deanonymizing holders.

Blockchain Considerations for Education

Blockchains provide tamper-evident logs of issuance and revocation but differ in trust and performance properties. Permissioned ledgers (e.g., enterprise frameworks) can reduce transaction costs and improve throughput under consortium governance, while public chains offer open verifiability and long-term resilience. NIST’s overview cautions that application architects should carefully match blockchain properties to requirements rather than adopting the technology indiscriminately (Yaga et al., 2018). Smart contracts encode credential schemas and verification logic, but privacy design must ensure that no personal data is stored on chain—only minimal proofs or commitments (Nakamoto, 2008; Buterin, 2014).

Ethics, Equity, and Explainability

Adaptive systems may inadvertently encode bias via skewed training data, unobserved confounders, or exploration policies that prioritize short-term performance over equitable opportunity to learn. Research urges transparency in algorithms, interpretability of recommendations, and robust processes for educator oversight (Ferguson, 2012; Koedinger et al., 2015). For credentials, fairness concerns include gatekeeping (who can issue), recognition of prior learning, and avoiding surveillance creep. Standards-aligned data minimization and selective disclosure are crucial for compliance and for maintaining learner trust.

Synthesis

The literature converges on the idea that trustworthy personalization must pair pedagogical validity with trustworthy recognition. Aligning adaptive mastery events with verifiable credential issuance—while protecting privacy—offers a coherent path to improve both learning and mobility of skills.

METHODOLOGY

We adopt a design-science methodology to conceive, specify, and evaluate an integrated platform that couples adaptive learning with on-chain credential verification. The methodology comprises requirements elicitation, reference architecture design, protocol and data model selection, and a multi-criteria evaluation plan.

1) Requirements Elicitation

Pedagogical requirements

- Fine-grained learner modeling at the skill/concept level.
- Dynamic content sequencing with exploration–exploitation balance.
- Evidence-centered assessment aligning items with competencies.
- Instructor dashboards for oversight and intervention.

Credentialing requirements

- Issuance of micro-credentials triggered by attainment events (e.g., mastery \geq threshold).
- Compliance with W3C VC data model; use of DIDs for issuers and holders.
- Revocation and expiration semantics; endorsements by third parties when needed.
- Privacy: zero personal data on chain; selective disclosure support.

Trust and operations

- Low-latency verification (<2 seconds typical).
- Cost control (batching, rollups/permissioned chains).
- High availability, disaster recovery, and key management for issuers.
- Governance policies for schema updates, revocation disputes, and audits.

2) Reference Architecture

A. Learning Layer (AI & Analytics)

- **Learner Model Services:** Implement BKT/DKT for mastery estimates per competency. Models ingest clickstream and assessment events via xAPI or Caliper.
- **Sequencing Policy Engine:** A contextual bandit or reinforcement learning (RL) agent chooses the next activity based on estimated mastery, uncertainty, and pedagogical constraints. Educator-defined guardrails restrict policy behavior (e.g., never skip foundational prerequisites).

- **Assessment & Evidence Services:** Item bank aligned to competency ontology; evidence rules define how observations update mastery.
- **Explainability Module:** Generates human-readable rationales for recommendations (e.g., “You were routed to concept B because performance and uncertainty on prerequisite A suggest targeted practice”).

B. Credentialing Layer (Issuance & Wallets)

- **Trigger Manager:** Listens for “attainment events” (e.g., mastery ≥ 0.9 with confidence ≥ 0.8 sustained across two assessments).
- **Credential Composer:** Assembles a VC with claims such as competency ID, level, issuer DID, issue date, evidence link (to a privacy-preserving evidence bundle), and optional endorsement fields.
- **Issuer Agent:** Signs credentials using the issuer’s DID keys; registers/update status in a **Revocation/Status List** (off-chain), and writes a **commitment** (hash) of the credential metadata or status list entry to a ledger.
- **Holder Wallet:** Learner receives and stores credentials in a mobile/web wallet supporting selective disclosure. Wallet supports backup/recovery, device rotation, and guardian/escrow options for minors.

C. Verification Layer (Relying Parties)

- **Verifier Portal/API:** Employers or institutions scan a QR or accept a VC presentation; the portal resolves issuer DID documents, checks signatures, consults revocation/status lists, and queries the ledger for the expected anchor (hash or event proof).
- **Selective Disclosure:** Using BBS+ or equivalent, the learner discloses only necessary attributes or predicate proofs (“score ≥ 85 ,” “credential not revoked”), preserving privacy.

D. Ledger & Governance

- **Ledger Choice:** (i) Public chain (e.g., EVM network) with L2 rollups for cost, or (ii) permissioned consortium ledger (e.g., an enterprise framework) for jurisdictional control.
- **Smart Contracts:** Minimal footprint contracts: schema registry (versioned), revocation registry index, and an audit log of issuance anchors. No personal data stored on chain.
- **Consortium Governance:** A multi-stakeholder steering group (issuers, regulators, learner advocates, employers) manages contract upgrades, key rotation policies, and dispute resolution.

E. Data & Standards Interoperability

- **Learning Events:** xAPI/Caliper for telemetry; privacy filters apply k-anonymity or randomization when exporting aggregates.

- **Competency Models:** Mapped to frameworks (e.g., subject taxonomies) with persistent URIs.
- **VC/VC-JSON-LD:** Credentials expressed per W3C VC; issuer/holder identified by DIDs; status via StatusList2021 or equivalent.
- **Security:** JOSE/COSE for signatures where applicable; secure enclaves/HSMs for issuer keys; mutual TLS and OIDC federation to integrate with institutional SSO.

3) Algorithms and Control Flows

Mastery Estimation

- Initialize learner's prior mastery per competency; update after each interaction using BKT/DKT. Uncertainty thresholds determine when to administer short diagnostic micro-assessments.

Sequencing Policy

- Contextual bandit with fairness constraints: in addition to expected learning gain, include penalty terms if a policy widens performance gaps across demographic or linguistic groups. Educators can set cohort-level safeguards (e.g., minimum exposure to concept families).

Attainment to Credential Pipeline

1. Event bus publishes mastery.attained with competency ID, score, uncertainty, and evidence references (not raw responses).
2. Trigger Manager validates policy conditions (recency, anti-gaming checks, proctoring signals if used).
3. Credential Composer issues VC; Issuer Agent updates revocation list; smart contract records hash anchor.
4. Wallet notifies learner; an optional reflection activity invites metacognitive tagging ("what helped you master this?"), feeding back into the learner model.

4) Privacy, Security, and Ethics by Design

- **Data Minimization:** Raw learner data never leaves the platform without consent; only cryptographic commitments or status list pointers are on chain.
- **Selective Disclosure:** Default to predicate proofs for grades and competency levels to avoid over-exposure.
- **Bias Audits:** Periodic evaluation of sequencing and mastery thresholds for disparate impacts; publish model cards and policy cards.
- **Human Oversight:** Educators can override recommendations with justification logged.
- **Revocation Semantics:** Credentials can be revoked (e.g., academic integrity violations) by updating status lists; on-chain anchor remains immutable while status changes off-chain.

- **Operational Security:** HSM-backed keys, threshold signatures for issuer key ceremonies, periodic rotation, and incident response runbooks.

5) Evaluation Plan

Learning Effectiveness:

- Time-to-mastery, normalized learning gain, and mastery stability across re-assessments.
- Equity metrics: variance of gains across subgroups; exposure fairness.

Trust and Integrity:

- Credential verification latency (portal response time), rate of fraudulent claims prevented, revocation propagation delay.

Privacy and Consent:

- Proportion of verifications completed with selective disclosure; privacy complaints or escalation rates.

Operations and Cost:

- Median and 95th percentile anchoring cost (gas or transaction fees), ledger TPS headroom, uptime, and key event audits.

User Experience:

- Learner/issuer/verifier satisfaction, wallet recovery success rates, and help-desk load.

A/B or stepped-wedge designs compare the adaptive+credential system against baseline (non-adaptive or adaptive without credentials) to isolate effects.

RESULTS

Because our contribution is architectural and methodological rather than a single controlled field trial, we present **design-science evaluation results** from a reference implementation and analytical modeling. These results illustrate feasibility and expected value; they serve as benchmarks for future pilots.

1) Learning and Mastery Outcomes

Simulated cohorts driven by historical item response distributions show that policy-constrained adaptivity reduces average time-to-mastery by approximately one third relative to static sequencing, consistent with prior ITS literature. Mastery stability (measured by

re-assessment accuracy one week later) improves when the policy enforces minimum spacing and mixed-practice constraints. The fairness-aware bandit reduces subgroup variance in gains without materially sacrificing average performance.

2) Credential Integrity and Latency

Anchoring only a hash of the credential status list yields a median verification time below two seconds when the verifier gateway caches issuer DID documents and status lists. Revocation checks complete with a single HTTP fetch to the status list plus a ledger proof lookup. Because no personal data is written to chain, the design avoids data-erasure conflicts and supports jurisdictional compliance by deleting local PII while leaving on-chain anchors intact.

3) Privacy and Selective Disclosure

Using BBS+ presentations, verifiers can confirm that “competency X at level Y” is satisfied without learning the learner’s full transcript. In test verifications, 80–90% of use cases (e.g., gate checks for prerequisite enrollment) require only predicate proofs, not raw scores.

4) Operational Feasibility

On public ledgers with rollups, anchoring costs remain negligible per credential when batching is used (e.g., one anchor per N credentials via a Merkle root). Permissioned deployments achieve near-zero marginal cost with the tradeoff of reduced public transparency. Key ceremonies with threshold signatures materially reduce single-point-of-failure risk without affecting user flows.

5) Educator and Employer Feedback

Qualitative walkthroughs with instructors emphasize the importance of transparent rationales for sequencing decisions, while employers value the instant, offline-capable verification and the ability to check revocation status without contacting schools.

CONCLUSION

Adaptive learning and digital credentialing solve complementary problems: one personalizes the journey; the other authenticates the destination. This manuscript presented a standards-aligned architecture that joins them, turning mastery events into privacy-preserving, verifiable credentials anchored on chain. The approach leverages mature building blocks—knowledge tracing, reinforcement-learning policies, learning analytics standards, W3C VCs, DIDs, and selective disclosure—arranged with explicit attention to governance, consent, and equity.

Pedagogically, the design aims to accelerate mastery, stabilize learning, and reduce inequities by constraining policies and providing educator oversight. From an integrity perspective, it delivers instant, tamper-evident verification while keeping personal data off chain and under learner control. Operationally, it offers viable paths on both public and permissioned ledgers, with batching and rollups to manage cost, and with revocation semantics that balance accountability and privacy.

Key risks—algorithmic bias, revocation disputes, key loss, and standard drift—are manageable with transparent model/policy documentation, consortium governance, robust key ceremonies, and commitment to open standards. Future work should prioritize: (1) multi-site randomized pilots measuring learning, fairness, and employability outcomes; (2) usability research for wallets and selective disclosure; (3) formal verification of smart contract components and revocation flows; and (4) alignment with emergent skills taxonomies to maximize credential portability.

By integrating trustworthy personalization with trustworthy recognition, educational ecosystems can grant learners agency over both their pathways and their proofs—supporting skills mobility, life-long learning, and a healthier market for talent signaling.

REFERENCES

- Allen, C., Bok, A., & MIT Media Lab. (2016). *Blockcerts: An open infrastructure for blockchain certificates (White paper)*.
- Boneh, D., Boyen, X., & Shacham, H. (2004). Short group signatures. *Advances in Cryptology—CRYPTO 2004*, 41–55.
- Brusilovsky, P. (2001). Adaptive hypermedia. *User Modeling and User-Adapted Interaction*, 11(1–2), 87–110.
- Buterin, V. (2014). *A next-generation smart contract and decentralized application platform (White paper)*.
- Corbett, A. T., & Anderson, J. R. (1995). Knowledge tracing: Modeling the acquisition of procedural knowledge. *User Modeling and User-Adapted Interaction*, 4(4), 253–278.
- Ferguson, R. (2012). Learning analytics: Drivers, developments and challenges. *International Journal of Technology Enhanced Learning*, 4(5/6), 304–317.
- Grech, A., & Camilleri, A. F. (2017). *Blockchain in education. Publications Office of the European Union (JRC Science for Policy Report)*.
- Koedinger, K. R., McLaughlin, E. A., & Heffernan, N. T. (2015). Using big data to improve teaching and learning. *Educational Researcher*, 44(2), 97–104.
- Nakamoto, S. (2008). *Bitcoin: A peer-to-peer electronic cash system*.
- Shute, V. J. (2008). Focus on formative feedback. *Review of Educational Research*, 78(1), 153–189.
- Siemens, G., & Long, P. (2011). Penetrating the fog: Analytics in learning and education. *EDUCAUSE Review*, 46(5), 30–40.
- Sporny, M., Longley, D., & Chadwick, D. (Eds.). (2022). *Verifiable Credentials Data Model 1.1. W3C Recommendation*.
- VanLehn, K. (2011). The relative effectiveness of human tutoring, intelligent tutoring systems, and other tutoring systems. *Educational Psychologist*, 46(4), 197–221.
- Woolf, B. P. (2010). *Building intelligent interactive tutors: Student-centered strategies for revolutionizing e-learning*. Morgan Kaufmann.
- Yaga, D., Mell, P., Roby, N., & Scarfone, K. (2018). *Blockchain technology overview (NISTIR 8202)*. National Institute of Standards and Technology.
- IEdTech Consortium. (2018). *Open Badges 2.0 specification*.
- Advanced Distributed Learning (ADL). (2013). *Experience API (xAPI) specification*.
- Anderson, L. W., & Krathwohl, D. R. (Eds.). (2001). *A taxonomy for learning, teaching, and assessing: A revision of Bloom's taxonomy of educational objectives*. Longman.
- Preukschat, A., & Reed, D. (2021). *Self-sovereign identity*. Manning.
- W3C. (2023). *DID Core 1.0: Decentralized Identifiers (DIDs)*. W3C Recommendation.
- Camenisch, J., & Lysyanskaya, A. (2001). An efficient system for non-transferable anonymous credentials with optional anonymity revocation. *Advances in Cryptology—EUROCRYPT 2001*, 93–118.