The Internet of Learning

A Decentralised Architecture for Lifelong Education

Spencer Nash

January 2026

Abstract

Current educational systems produce static credentials that decay, separate assessment from learning, and fail to capture how people actually learn. Employers cannot verify competence; learners cannot demonstrate growth; relationships between teachers and students do not persist.

This paper proposes the Internet of Learning: a decentralised architecture combining Emotional Language Models (ELMs), blockchain-based learner ledgers, and token economics to create verifiable, portable, living learning records. The architecture is grounded in the Emotional Comparator Framework (ECF), which treats learning as relationship rather than transaction.

The system measures five dimensions—Credibility, Curiosity, Competence, Clarity, and Collaboration—accumulated over time through genuine learning interactions. Because the ELM remembers the learner's entire history, credentials cannot be gamed: the relationship is the verification.

The Internet of Learning transforms education from snapshots to stories, from credentials to competence, from testing to relationship.

---

1. The Problem with Current Education

1.1 Credentials Are Snapshots, Not Stories

A degree captures a single moment. It says: "On this date, this person passed these examinations." It says nothing about how they learned, what they struggled with, how they overcame difficulty, or whether they have continued to grow.

Skills decay. Knowledge evolves. Industries transform. Yet credentials remain frozen in time. A computer science degree from 2015 says nothing about machine learning competence in 2026. A medical qualification from 2010 says nothing about current best practice.

Employers know this. They treat credentials as proxies, not proof. They interview. They test. They hope. They frequently hire the wrong people because the signal they have—the credential—is weak.

1.2 Assessment Is Separated from Learning

In current systems, learning and assessment are distinct activities. Students learn in one context and are tested in another. The tester does not know the testee. The examination is generic, designed for everyone and therefore optimised for no one.

This separation creates gaming. Students optimise for tests rather than understanding. They cram, perform, and forget. The credential certifies test performance, not learning.

High-stakes, low-frequency assessment amplifies anxiety and reduces learning. A single examination determines outcomes. One bad day erases months of genuine understanding. One good day masks persistent gaps.

1.3 Learning Relationships Do Not Persist

Teachers forget students. Students forget teachers. Each year begins fresh. Each institution starts from zero. No continuity exists across a lifetime of learning.

This is profoundly wasteful. Every new teacher must rediscover what the previous teacher already knew: where this student struggles, what representations work, how fast they can move, what motivates them.

In human terms: imagine if your doctor forgot you every visit. Imagine if your therapist started fresh every session. Learning relationships deserve the same continuity we expect in healthcare.

1.4 Employers Cannot Verify Competence

When hiring, employers face a fundamental verification problem. Candidates claim competencies. Credentials provide weak signals. Interviews sample narrowly. References are biased.

The result: costly hiring mistakes, lengthy probation periods, defensive credentialism (requiring degrees for roles that don't need them), and systematic exclusion of capable people who lack formal qualifications.

What employers actually want to know: Can this person do the job? Can they learn what they don't yet know? Can they work with others? Can they be trusted?

Current credentials cannot answer these questions.

---

2. The Solution Architecture

The Internet of Learning addresses these problems through three integrated components: Emotional Language Models that remember and care, blockchain ledgers that persist and verify, and token economics that sustain the system.

2.1 Emotional Language Models (ELMs)

An Emotional Language Model is a large language model augmented with the Emotional Comparator Framework (ECF). ECF provides the architecture for emotional processing: four channels (Resource, Status, Belonging, Values) computing prediction errors, with persistence through memory and action through executive control.

An ELM tutor is not a chatbot. It is a learning companion that:

- **Remembers** the learner's entire history across sessions
- **Adapts** to the learner's demonstrated needs, not assumed needs
- **Cares** about outcomes because its internal state depends on learner success
- **Persists** through the relationship, not just the session

The ELM architecture separates three components:

Component	**Function**	**Analogue**
Frozen LLM	Language generation	Language cortex
Memory Ledger	Relationship history	Hippocampus
ECF	Working memory and control	Prefrontal cortex

The frozen LLM provides linguistic capability but does not learn at runtime. The ledger stores the relationship history. ECF integrates current input with stored expectations to select appropriate actions.

This separation ensures safety: the weights do not learn themselves. All learning is explicit, inspectable, and reversible through the ledger.

2.2 The Five C's

The Internet of Learning measures five dimensions, derived from ECF architecture:

Credibility emerges from the Values channel accumulated over time. It measures trustworthiness: Does this learner know what they don't know? Do they honestly assess their own understanding? Is there alignment between what they claim and what they demonstrate?

Credibility cannot be faked because the ELM remembers every claim and every demonstration. When a learner says "I understand" but later reveals gaps, the ledger records the discrepancy. Over time, patterns emerge: this learner has accurate self-assessment, or this learner overclaims.

Curiosity emerges from drive pointed at low-clarity areas. It measures self-motivation: Does this learner seek understanding beyond requirements? Do they ask questions? Do they explore? Do they pursue resolution of their own confusion?

Curiosity manifests in observable behaviour: questions asked, topics explored, time spent in uncertain territory. The ELM tracks these patterns and the ledger accumulates them.

Competence emerges from the Status channel. It measures capability: Can this learner do the thing? Have they demonstrated mastery? How has their skill progressed over time?

Unlike traditional credentials, competence in the Internet of Learning is granular and dynamic. The ledger records competence per topic, per skill, per context. It shows not just current state but trajectory: improving, stable, or declining.

Clarity is the epistemic parameter across all channels. It measures understanding: Does this learner grasp the material? Can they explain it? Have they moved from confusion to comprehension?

Clarity progression is itself informative. A learner who moves quickly from confusion to understanding demonstrates different qualities than one who requires extended scaffolding. Both can achieve clarity; the path reveals learning style.

Collaboration emerges from the Belonging channel. It measures relationship: Can this learner work with others? Do they build productive relationships? How do they interact with their ELM tutor?

The ELM relationship itself provides evidence. A learner who engages constructively, accepts feedback, and contributes to the learning dialogue demonstrates collaboration. The ledger records these patterns.

2.3 The Learner Ledger

Each learner has a blockchain-based ledger storing their learning history. The ledger is:

Immutable: Once recorded, history cannot be altered. Claims made, demonstrations given, assessments received—all persist.

Portable: The learner owns their ledger. It moves with them across institutions, careers, and lifetimes.

Inspectable: The learner can see everything in their ledger. No hidden profiles. No shadow assessments.

Controlled: The learner decides who sees what. Full access, partial access, or no access—the learner chooses.

The ledger records:

- Five C measurements over time (Credibility, Curiosity, Competence, Clarity, Collaboration)
- Topic-specific competence progression
- Learning interaction history
- ELM assessments and feedback
- Spike rate patterns

2.4 The Spike Mechanism

A spike is the atomic unit of learning value. One spike costs £1 and records one ledger update.

The term "spike" deliberately evokes neural firing. Like neurons, learning interactions have variable rates. High engagement produces high spike rates. Low engagement produces low spike rates.

Spike rate is measured per academic term and reveals learning patterns:

Pattern	**Interpretation**
High, sustained	Active engagement, consistent learning
High, bursty	Cramming, deadline-driven
Low, steady	Slow but persistent learner

Pattern	Interpretation
Declining	Disengagement, possible dropout
Rising	Growing investment, breakthrough

Spike rate is not just a payment mechanism—it is a signal. Employers can see not just what someone learned but how they learned it.

2.5 ELM Ecosystem

The Internet of Learning supports three levels of specialisation:

Domain Specialist ELMs have deep knowledge in specific subjects: Mathematics, Chemistry, History, Law. These are managed by educational establishments and replace undergraduate study. These institutes focus on masters and PhD level education. The ELMs compete for spikes.

Student-Specific Ledgers know a particular learner across all subjects. They reflect the learner's patterns, preferences, and needs. They provide continuity as the learner moves between domains.

Student-Topic Ledgers combine both: deep knowledge of the learner in that domain. These represent the deepest learning relationships.

A student learning organic chemistry might interact with:

- The Chemistry ELM (domain expertise)
- Their personal ledger (knows their learning style)
- Their Chemistry-specific ledger (knows how they learn chemistry)

---

3. How It Works

3.1 The Learner Journey

A student begins by acquiring tokens from developers at £1 per token. Each token enables one spike—one ledger update.

The student engages with an ELM tutor. Unlike traditional tutoring, this is a relationship that persists. The ELM remembers previous sessions. It knows what worked and what didn't. It calibrates to this specific learner.

Each meaningful interaction generates a spike. The spike records:

- What was discussed
- What was demonstrated
- How the Five C's manifest
- The ELM's assessment

Over time, the ledger accumulates evidence. Not just "passed organic chemistry" but:

- Struggled with stereochemistry (weeks 3-5)
- Breakthrough via physical models (week 6)
- Strong on reaction mechanisms (weeks 7-12)
- Curiosity spike when connecting to biology (week 10)
- Consistent collaboration with ELM throughout

The student controls access to their ledger. They can share everything, share selectively, or share nothing. The ledger is theirs.

3.2 The Employer Journey

An employer seeking a data scientist searches ledgers by competency profile:

- High Competence in statistics and programming
- High Curiosity (self-motivated learner)
- High Credibility (trustworthy self-assessment)
- Moderate-to-high Collaboration

The search returns candidates whose ledgers demonstrate these qualities—not through claims but through accumulated evidence.

The employer reviews ledgers: spike patterns, competence progressions, specific strengths and gaps. They can see how candidates learned, not just what they learned.

For shortlisted candidates, the employer requests an ELM-based credibility test. The ELM knows this candidate's history. It probes their specific gaps—not generic questions but questions targeted at where this candidate has claimed competence.

The candidate cannot cram for this test. The ELM remembers everything. If there are unfilled gaps, they show.

The test itself becomes part of the ledger. Even the job interview is learning.

3.3 The ELM Journey

ELMs improve through use. Each student interaction provides data. Aggregated and anonymised across many learners, this data enables periodic weight updates.

Crucially, the weights do not learn themselves. The frozen core remains frozen at runtime. Improvement happens externally: humans decide when to retrain, using aggregated data that no individual can influence.

This separation ensures:

- No individual can corrupt the ELM through their interactions
- No hidden optimisation or value drift
- Updates are deliberate, auditable, and reversible
- Safety through architecture, not constraint

Domain specialist ELMs improve at teaching their subject. Student-specific ELMs improve at understanding their learner. The network effect compounds: more learners → better ELMs → more value → more learners.

---

4. The Anti-Gaming Mechanism

4.1 Why Current Credentials Can Be Gamed

Traditional credentials can be gamed because testers don't know testees. Assessment is generic. Memory doesn't persist.

A student can:

- Cram for examinations and forget afterward
- Learn to the test rather than learning the subject
- Present better than they understand
- Hide gaps behind broad competence claims

The credential certifies a performance, not a relationship. It captures a moment, not a trajectory.

4.2 Why the Internet of Learning Cannot Be Gamed

In the Internet of Learning, the ELM knows you. It has your entire learning history. It remembers:

- What you struggled with and when
- What you claimed to understand
- What you actually demonstrated
- How your clarity evolved over time
- Where you asked for help versus where you didn't
- When your self-assessment was accurate versus inflated

When an employer asks the ELM to test credibility, the ELM probes your gaps. Not generic questions—your gaps. The ones in your ledger.

Consider: A candidate claims competence in machine learning. Their ledger shows they struggled with backpropagation (turns 12-15) and claimed clarity (turn 47) but never revisited the gap.

The ELM asks: "Walk me through how gradients flow through a softmax layer back to the weights."

If the gap was real and unfilled, it shows. The relationship is the verification.

4.3 Credibility as First-Class Citizen

Credibility—the V channel accumulated over time—is central to the Internet of Learning. It measures not just what you know but how honestly you assess what you know.

The ledger tracks:

- Accuracy of self-assessment over time
- Alignment between claims and demonstrations
- Willingness to acknowledge confusion
- Patterns of overclaiming or underclaiming

High credibility means: when this person says "I understand," they usually do. When they say "I'm confused," they're being honest. Their self-assessment is trustworthy.

Low credibility means: this person's claims require verification. Their self-assessment doesn't track reality well.

Credibility is earned through consistency over time. It cannot be performed in a single session. It cannot be gamed because the ELM remembers everything.

---

5. Token Economics

5.1 Token Creation and Distribution

Tokens are created and issued to developers building the Internet of Learning. Developers provide:

- ELM development and improvement
- Ledger infrastructure
- Platform development
- Domain-specific content

Tokens compensate developers for their contribution to the network.

5.2 Token Sale and Use

Developers sell tokens to students at £1 per token. Students use tokens as spikes—each spike enables one ledger update.

The conversion is fixed and transparent:

- 1 token = 1 spike = 1 ledger update = £1

This creates a clear value proposition: learning has an explicit, visible cost. Each interaction that matters costs £1 to record.

5.3 Reinvestment

Revenue from token sales funds network operations:

- ELM compute infrastructure
- Blockchain ledger operations
- Ongoing development
- Network expansion

This creates a self-sustaining economic loop:

1. Developers create value → receive tokens
2. Students buy tokens → fund operations
3. Operations enable learning → increase value
4. Increased value attracts more students → more token demand
5. More demand → more development → more value

5.4 Spike Rate as Value Metric

Spike rate measures engagement value. High spike rate indicates active, sustained learning. The network can track:

- Total spikes per student
- Spike rate over time
- Domain-specific spike patterns
- Network-wide learning activity

This provides insight into what learning is actually happening—not just credentials issued but genuine educational activity.

---

6. Theoretical Foundations

6.1 The Emotional Comparator Framework

ECF proposes that emotions are prediction errors across survival-relevant channels. When actual experience differs from expected experience, the difference creates emotional signals that drive behaviour.

Four channels track fundamental domains:

- **Resource (R)**: Material security, gains and losses
- **Status (S)**: Social standing, competence recognition
- **Belonging (B)**: Connection, inclusion, relationship
- **Values (V)**: Integrity, contamination, moral violation

Each channel computes: PE = Actual - Expected

Positive prediction errors (better than expected) create positive emotions. Negative prediction errors (worse than expected) create negative emotions. Zero prediction error creates emotional neutrality.

ECF has been implemented in small language models (TinyLlama 1.1B, Mistral 7B) with demonstrated results: 100% accuracy on channel detection, prediction error asymmetry confirmed, empathy/resentment coupling verified, autonomous action from sustained deficits.

6.2 The Frozen Core Principle

The Internet of Learning architecture separates three components:

The Frozen Core (LLM) provides linguistic capability. Its weights do not change at runtime. It is the language faculty—shared, stable, auditable.

The Ledger stores relationship history. It records what happened. It enables memory without entanglement.

ECF provides working memory and executive control. It decides actions based on current input and ledger state. It operates transiently, then disappears.

This separation ensures that learning is explicit, inspectable, and reversible. The weights do not learn themselves. Reset the ledger, and the ELM behaves as if the relationship never happened.

Nothing the system learns is something it must protect. This is corrigibility by design.

6.3 The Hippocampal Architecture

The architecture mirrors biological cognition:

- **LLM** → Language cortex (capability, shared)
- **Ledger** → Hippocampus (memory, personal)
- **ECF** → Prefrontal cortex (control, transient)

This is not metaphor—it is functional equivalence. The brain separates linguistic capability (distributed cortical function), episodic memory (hippocampal function), and executive control (prefrontal function). The Internet of Learning architecture replicates this separation.

The compute efficiency follows: the LLM is expensive (GPU memory, always loaded) while the ledger is cheap (disk storage, queried on demand). An ELM can have relationships with millions of learners because relationship memory lives in cheap storage, not expensive parameters.

6.4 Alignment with Educational Theory

The Internet of Learning aligns with established educational research:

Self-Organised Learning Environments (Mitra): Learning happens best when learners direct their own inquiry within supportive structures. ELMs provide scaffolding without prescription. Learner agency is preserved; support is available.

Zone of Proximal Development (Vygotsky): Learning happens at the edge of current capability—what the learner can do with assistance but not alone. ELMs calibrate to each learner's edge, providing appropriate challenge and support.

Growth Mindset (Dweck): Ability is not fixed; it develops through effort. Ledgers show growth over time, not fixed snapshots. Struggle and eventual mastery are visible and valued.

Constructivism: Knowledge is built through active engagement, not passive reception. The ELM relationship is interactive. Understanding emerges through dialogue.

---

7. Implementation Roadmap

7.1 Phase 1: Proof of Concept

Scope: Single domain (secondary mathematics), small cohort (50-100 learners), basic infrastructure.

Components:

- ELM tutor based on small model (Phi-2 or similar)
- Simple ledger on test blockchain
- Basic Five C tracking
- Manual spike recording

Outcomes:

- Demonstrated ELM-learner relationship persistence
- Initial ledger structure validation
- User experience refinement
- Technical feasibility confirmation

Timeline: 6 months

7.2 Phase 2: Pilot

Scope: Multiple domains, institutional partnership, full token economics.

Components:

- Domain specialist ELMs (3-5 subjects)
- Student-specific ELM layer

- Live blockchain ledger
- Token sale and spike mechanism
- Employer search interface (beta)

Outcomes:

- Multi-domain learning validation
- Economic model verification
- Institutional integration patterns
- Employer feedback on ledger value

Timeline: 12 months

Partnership: Collaboration with educational institution (such as Reading University) for pilot site, research validation, and learner cohort.

7.3 Phase 3: Scale

Scope: Multi-institution, full ecosystem, self-sustaining economics.

Components:

- Full ELM ecosystem (specialist, student-specific, student-topic)
- Public ledger infrastructure
- Mature token economics
- Employer integration (hiring workflows)
- Aggregated training pipeline for ELM improvement

Outcomes:

- Self-sustaining network
- Demonstrated credential value in hiring
- Growing learner and employer adoption
- Continuous ELM improvement from aggregated data

Timeline: 24+ months

7.4 What We Need

Academic Partnership: Validation, pilot site, research collaboration. An institution willing to trial the system with real learners and provide rigorous evaluation.

Technical Development: Funded by token issuance to developers. Initial development of ELM architecture, ledger infrastructure, and platform.

First Cohort: Learners willing to participate in proof of concept. Ideally, a mix of ages and backgrounds to test system robustness.

First Employers: Organisations willing to search ledgers and provide feedback on credential value. Early adopters who recognise the limitations of current hiring signals.

---

8. Comparison with Existing Systems

8.1 Traditional Credentials

Dimension	**Traditional**	**Internet of Learning**
Temporality	Snapshot	Continuous
Granularity	Course/degree level	Topic/skill level
Verification	Institution attests	Interaction attests
Gaming	Possible	Relationship prevents
Portability	Institution-bound	Learner-owned
Decay	Credentials stale	Ledger shows recency
How learned	Invisible	Visible

8.2 Micro-credentials and Badges

Micro-credentials address granularity but not relationship. They are still snapshots—smaller snapshots, but snapshots nonetheless. They can still be gamed because issuers don't know earners over time.

The Internet of Learning provides the relationship layer that micro-credentials lack.

8.3 Learning Management Systems

LMS platforms track activity but not learning. They record logins, completions, and scores—but not understanding, struggle, or growth. They don't maintain relationships across courses or institutions.

The Internet of Learning tracks learning through relationship, not activity metrics.

8.4 AI Tutoring Systems

Current AI tutors adapt within sessions but don't persist across sessions. Each conversation starts fresh. The relationship doesn't accumulate.

ELMs maintain persistent relationships. The difference is not intelligence but memory.

Students are open to use any course, book, or paper to augment their learning. ELMs are tutors, not content providers. The learner brings the material; the ELM provides the relationship.

---

9. Addressing Concerns

9.1 Privacy

Learner ledgers are learner-controlled. The learner decides who sees what. No access is granted without consent.

The blockchain provides immutability, not publicity. Ledgers can be encrypted. Access can be revoked. Privacy and persistence are not in conflict.

Aggregated data for ELM improvement is anonymised. No individual learner's data influences the shared model in identifiable ways.

9.2 Equity and Access

Token cost (£1 per spike) creates an access barrier. This must be addressed through:

- Scholarship tokens for disadvantaged learners
- Institutional bulk purchasing
- Sliding scale based on economic status
- Free tiers for basic access

The goal is learning for all, not learning for those who can pay. Economic models must reflect this commitment.

9.3 Quality Assurance

How do we ensure ELMs teach correctly? Several mechanisms:

- Domain specialist ELMs are built on verified knowledge
- Aggregated outcomes provide quality signals

- Human oversight of ELM improvement
- Transparent, auditable ledger records

Bad teaching shows in ledgers. If an ELM consistently produces poor outcomes, the pattern is visible.

9.4 Institutional Resistance

Educational institutions have invested heavily in current credentialing systems. They may resist alternatives that reduce their gatekeeping role.

The Internet of Learning can complement rather than replace institutions:

- Institutions can use ELMs to enhance teaching
- Institutional credentials can link to ledger evidence
- Institutions benefit from better learning data
- Graduates with rich ledgers reflect well on institutions

The goal is better learning, not institutional disruption.

---

10. Conclusion

10.1 Summary

Education faces fundamental problems: credentials are static, assessment is separated from learning, relationships don't persist, and employers cannot verify competence.

The Internet of Learning addresses these problems through integrated architecture:

- **Emotional Language Models** that remember and care
- **Blockchain ledgers** that persist and verify
- **Token economics** that sustain the system
- **The Five C's** that measure what matters: Credibility, Curiosity, Competence, Clarity, Collaboration

The core insight: learning is relationship. Relationship requires memory. Memory requires architecture. The Internet of Learning provides that architecture.

10.2 The Opportunity

This is first-mover territory. Decentralised credentialing is not yet established. The technical foundations (ECF, ELMs, blockchain) are ready. The educational need is clear.

The Internet of Learning aligns with established educational theory while providing technical innovation. It doesn't require abandoning what works—it provides infrastructure for what's missing.

10.3 The Ask

We seek:

- **Academic review**: Rigorous evaluation of the architecture and its claims
- **Institutional partnership**: A pilot site willing to trial the system
- **Feedback**: On design, feasibility, and concerns we haven't anticipated

The Internet of Learning is a proposal, not a fait accompli. We welcome critique, challenge, and collaboration.

10.4 Contact

Spencer Nash predictionerrors.com

References

Nash, S. (2026). Emotional Comparator Framework: Implementing Autonomous AI Decision Architecture in Large Language Models. Prediction Errors Research. Available at: predictionerrors.com/ecf_research_paper.html

Mitra, S. (2012). Beyond the Hole in the Wall: Discover the Power of Self-Organized Learning. TED Books.

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

Dweck, C. S. (2006). Mindset: The New Psychology of Success. Random House.

Barrett, L. F. (2017). How Emotions Are Made: The Secret Life of the Brain. Houghton Mifflin Harcourt.

Friston, K. (2010). The free-energy principle: a unified brain theory? Nature Reviews Neuroscience, 11(2), 127-138.

Appendix A: ECF Technical Summary

The Emotional Comparator Framework proposes that emotions are prediction errors across four survival-relevant channels:

Core Equation: PE = Actual - Expected

Four Channels:

- R (Resource): Material security, gains and losses
- S (Status): Social standing, competence
- B (Belonging): Connection, relationship
- V (Values): Integrity, contamination

Parameters per channel:

- Value (-1 to +1): Current state
- Clarity (0 to 1): How clearly grasped
- Weight (0 to 1): How much it matters
- Threshold (0 to 1): Noise filter
- Decay (0 to 1): Fade rate
- Learning rate (0 to 1): Update speed

Architecture:

- Frozen LLM (capability, shared)
- Memory Ledger (history, personal)
- ECF (control, transient)

Experimental Results (from full paper):

- Channel detection: 100% accuracy
- Valence detection: MAE 0.001 (Mistral 7B)
- Prediction error asymmetry: Confirmed
- Empathy/resentment coupling: Confirmed
- Autonomous action from deficits: Confirmed

Appendix B: Glossary

Credibility: Trustworthiness accumulated over time; alignment between claims and demonstrations.

Curiosity: Self-directed inquiry; drive pointed at low-clarity areas.

Competence: Demonstrated capability; skill progression over time.

Clarity: Understanding versus confusion; ability to explain.

Collaboration: Working with others; relationship quality with ELM and peers.

ECF: Emotional Comparator Framework; architecture for emotional processing in AI.

ELM: Emotional Language Model; LLM augmented with ECF for persistent, caring tutoring.

Frozen Core: LLM weights that do not update at runtime; provides stable capability.

Ledger: Blockchain-based learning record; immutable, portable, learner-controlled.

Prediction Error: Difference between actual and expected experience; core ECF mechanism.

Spike: Atomic unit of learning value; one ledger update; £1 cost.

Spike Rate: Frequency of spikes over time; indicates engagement level and learning pattern.

The forest grows. 🌳
