

After the University

A proposal for a federated open learning commons in
an age of AI and saved hours

**The objective is not fewer teachers. It is to
expand learning across reach, language, practice
and human attention without reducing education
to scale alone.**

PUBLIC REVIEW DRAFT · CAMELOT

Author	Dennis Hedegreen
Series	A Saved Hours Doctrine subpaper
Version	v0.7 public review draft
Date	30 July 2026
Boundary	Public Camelot draft · no DOI · not final · open for criticism

PUBLIC REVIEW DRAFT. Published in the Hedegreen Research Camelot archive for criticism. This is not a final release or a peer-reviewed paper.

Abstract

Modern education commonly bundles access to knowledge, tutoring, mentorship, assessment, certification, physical facilities, professional regulation and social recognition into one institutional pathway. Networked publication and open-education programmes have expanded access to learning material, but the precedents reviewed for this paper still route recognition, credible assessment and sustained support through narrower institutional arrangements.

Generative AI does not abolish education's remaining human problems. Structured trials indicate that it can reduce time or improve short-run outcomes for some bounded educational interactions, while unguarded answer access can weaken later unaided performance. These interventions do not establish a generic AI tutor or a complete cost result.

This paper extends *The Saved Hours Doctrine* from the general distribution of productivity-created time to educational labour.[26] It asks what happens when AI reduces defined teaching tasks: whether the released hours are returned to mentorship, accessibility, curriculum maintenance, assessment and public learning capacity, or captured as staffing cuts, larger classes, denser workloads, lower budgets and platform rents. It proposes Saved Teaching Hours accounting inside a federated open learning commons. The proposal is an institutional design and pilot hypothesis, not evidence that AI tutoring has already produced educational abundance.

The proposed commons has four interoperable layers: a Knowledge Commons governed by teachers and domain experts; a Tutor Commons in which replaceable AI and non-AI systems operate over auditable material; a Competence Commons designed to carry evidence of learning between providers; and a Human Commons preserving mentorship, laboratories, community, safeguarding and high-stakes judgement. The design is not another online university. It asks whether learning infrastructure can become more forkable, multilingual, learner-controlled and teacher-governed while disclosing its full human, technical and physical cost.

The paper's *1,000x learning* phrase names a direction across reach, language, availability, practice, durability of teacher contributions and human attention where it matters. It does not mean one credential, course or test score multiplied by one thousand.

Keywords: open education; artificial intelligence; saved hours; teachers; competency assessment; open standards; open source; credentials; public infrastructure

The proposal in one sentence

Build teacher-governed open material; test replaceable tutors over it; let learners export evidence of demonstrated competence; and govern released teaching hours toward more learning rather than fewer teachers.

What “open” means in this proposal

The word *open* does not impose one licence or software model on every layer. It names four distinct commitments:

Layer	Required meaning of open
Knowledge	Educational material uses a named open licence whose permissions and restrictions are disclosed module by module. “Openly licensed” does not mean that all licences permit identical reuse.
Interoperability	Module, tutor, evidence and record contracts are documented as open standards that competing hosts can implement.
Reference infrastructure	The protocol's reference implementation is open-source software so its behaviour can be inspected, operated and forked.
Tutor implementations	Commercial, public-interest and local tutors may participate without being open source, but they must meet disclosure, privacy, export and switching requirements.

This terminology leaves room for commercial models while preventing a proprietary tutor from becoming the constitution of the curriculum.

Contents

1. The institution that is missing

2. Education is a bundle, not a single service

3. The internet opened knowledge but left recognition scarce

4. What AI changes — and what it does not

5. Where do the saved teaching hours go?

6. The four commons
7. The teacher after scarcity

8. The right to prove competence

9. Compute at cost; humans at real cost

10. Governance without a new gatekeeper

11. A one-course pilot

12. Positionality: why Hedegreen Research is writing this

13. Limitations and unresolved tests

14. Conclusion

1. The institution that is missing

Networked publication has made lectures, textbooks, primary sources, software repositories, research papers, simulations and learning communities reachable beyond the institutions that produced them. Access remains unequal across connectivity, language, disability, paywalls and local support, but the expansion documented by open-education programmes is substantial.

The Saved Hours Doctrine asked who receives the time released by productivity. *Post-Work Requires Power* asked what material capacity is required to inhabit that time. *After the University* asks what happens when the same accounting problem enters education. If AI reduces a defined part of teaching work, the released hours do not automatically become more learning. They can become staff cuts, larger classes, denser workloads, lower public budgets or platform rents. Returning those hours to learners, teachers and the public knowledge commons requires an institution.

The precedents reviewed for this paper do not provide one mechanism connecting open materials to personal explanation, durable teacher stewardship, hard independent assessment, portable evidence of competence and broad social recognition. They provide fragments. The missing institution proposed here is not simply a website with free courses. It is a commons in which teaching work can accumulate, learners can receive individual help, competence can be tested through independent bodies, and no single platform needs to own the result.

This distinction matters because access to information is not the same as an education. A learner may have the correct textbook and still need a different explanation, structured practice, feedback, motivation, a community, equipment, or someone capable of judging whether the apparent understanding is real. MIT OpenCourseWare, OSSU and OERu explicitly separate some of these functions: their boundaries should not be mistaken for evidence that the materials failed.

The proposal begins from a narrower observation. In the reviewed systems, learning materials became available beyond enrolment while assessment, credit and recognition remained with authorised institutions or partners. Structured AI systems now make some forms of explanation and practice easier to scale. The question is whether a new design can widen those interactions without weakening learning, labour, assessment or trust.

What “after” means

The title does not predict that campuses, universities or scholarship will disappear. “After” names a change in institutional rank. Universities may remain excellent places for laboratories, research communities, studios, debate and long-form human development. They should no longer be treated as the sole legitimate gateway to learning, assessment or intellectual contribution.

The proposal is therefore not an online imitation of four years on campus. It is a common layer beneath and beyond institutions. A university, library, trade union, community centre, employer, municipality or informal study group could all become nodes in the same learning commons. Existing institutions are invited to participate; they are not entitled to own the protocol.

Open publication widened access to knowledge. Recognition remained institutionally controlled. Structured AI makes some tutoring interactions more scalable. The commons is a proposal for testing the next step.

2. Education is a bundle, not a single service

Debates about replacing or defending “the university” are often confused because the university is not one thing. It is a bundle of functions held together by law, history, finance and reputation. Different disciplines depend on different parts of that bundle.

Function	Why it matters	Can it be unbundled?
Knowledge distribution	Access to explanations, sources and examples.	Largely yes; open licensing and the web already demonstrate this.
Tutoring and feedback	Adapts explanation and practice to the learner.	Partly; AI can scale some functions, but quality and dependence remain open questions.
Mentorship and community	Builds judgement, belonging, motivation and professional identity.	Can be organised separately, but not reduced to content delivery.
Assessment	Tests whether competence has actually been demonstrated.	Partly; existing systems separate prior learning or examination from teaching, while authority and acceptance remain conditional.
Certification	Communicates a trusted result to others.	Technically portable; social trust and legal recognition remain harder.
Research	Creates new knowledge and trains researchers.	Requires expert communities, funding, methods and often physical infrastructure.
Physical practice	Laboratories, clinics, workshops, studios and field work.	Only partly; real supervision and equipment remain scarce.
Professional permission	Protects the public in regulated occupations.	No simple unbundling; law must distinguish learning from permission to practise.

The purpose of unbundling is not to declare each function cheap or interchangeable. The pilot would test whether a learner can combine open material, local mentoring, partner-based physical practice and an assessment body without purchasing every function from one provider. The functions can remain demanding while the proposed pathway becomes more open.

This also clarifies the boundary of the proposal. An open right to learn medicine is not an open right to practise medicine. A public curriculum in structural engineering does not remove the need for supervised experience or legal licensure. The commons can widen learning and evidence. It cannot responsibly abolish every safety boundary.

3. The internet opened knowledge but left recognition scarce

The open-education movement provides several policy and operating precedents for this proposal. UNESCO's 2019 Recommendation on Open Educational Resources treats openly licensed learning material as an international public-policy concern and names access, adaptation, capacity, quality and sustainability as public work.[1] It is a policy framework, not evidence that the four-commons architecture works.

Several projects show both open access and its institutional boundaries. MIT OpenCourseWare publishes material from thousands of courses without enrolment, start dates or fees and states that OCW use earns no MIT credit or certification.[2] The Open Source Society University organises free resources into undergraduate-style curricula but does not award a degree.[3] OpenStax operates a professionally maintained nonprofit textbook programme at large self-reported adoption scale. New and updated OpenStax editions use a [CC BY-NC-SA](#) licence, so “open” here must name the exact reuse boundary.[4]

OERu provides a lifecycle-bounded precedent for separating free self-directed study from optional paid assessment and credit under partner institutions' authority.[5] It has announced permanent closure on 31 July 2026. The network therefore demonstrates a real unbundling mechanism and a sustainability warning; it does not establish durable federation.

The separation of teaching from trusted assessment is older than the internet. The University of London began as an examining and degree-awarding body and, in 1865, made its examinations available remotely. Its official history describes this as a pathway to a degree without physical attendance at a university and records a global network of examination centres intended to preserve assessment integrity.[15] The precedent is imperfect and remained institutionally controlled, but it demonstrates that the teaching provider and the examining authority need not be identical.

That historical record shows that teaching location and examining authority can be separated. It does not show that independent assessment automatically creates trust or transfers across institutions. The proposed new element is to test whether open material and external assessment can be joined by structured, scalable interaction between them.

These examples show that the missing institution was never blocked only by file hosting. The hard problems were coordination, maintenance, feedback, assessment and trust. A repository can make knowledge available. It cannot by itself explain a difficult

concept differently to each learner, notice a recurring misconception, defend a curriculum against capture, or persuade an employer that an external candidate has genuinely demonstrated the relevant competence.

The lesson is not that open education failed. It is that the reviewed programmes retained assessment, credit or recognition authority outside the open-material layer. The proposed system must test whether open material can be joined to demanding evaluation and a sustainable human community without disguising those remaining authorities.

4. What AI changes — and what it does not

Generative AI can restate an explanation, translate it, ask questions, generate practice and compare an answer with a rubric. Capability is not the same as reliable availability, affordability or learning. These functions do not make a model a teacher in the human or professional sense; they make it a possible interface to teacher-governed knowledge.

Early evidence is promising but highly context-dependent. A 2025 randomised controlled trial in undergraduate physics reported substantially higher immediate learning gains and lower median time on task for students using a carefully designed AI tutor than for students in an in-class active-learning condition.[6] The study's system did not merely open a general chatbot: it used structured sequencing, pedagogical constraints and instructor-designed material. Its authors also reported limitations, including failures to scaffold complex multi-part problems reliably without additional design.

A World Bank randomised trial in Nigeria tested a six-week, teacher-led after-school programme using a GPT-4-powered assistant. The study reported a 0.31 standard-deviation improvement and described the learning gains as highly cost-effective in that specific setting.[7] Again, the intervention combined AI with teachers, training, curriculum and structured use. It does not establish that unsupervised access to a general model produces the same outcome.

The correct conclusion is therefore narrower than the marketing claim that every learner now has a perfect tutor. AI can lower the cost of certain interactions when the learning design is strong. It can also enable answer production without understanding, introduce errors, encourage dependency, and give confident form to weak material. UNESCO's guidance on generative AI in education accordingly emphasises human agency, inclusion, privacy, governance and the need for educational institutions to retain responsibility.[8]

AI should not be the constitution of the curriculum. It should be an inspectable and replaceable interface to a curriculum governed by people.

This is why the proposed pilot is model-neutral. It must test whether the same module can survive a switch between competing commercial models, public-interest or local systems and a non-AI mode without losing its learning record. Until that switch is demonstrated, model neutrality remains an architectural requirement rather than an established capability.

Augmentation is not outsourcing

The strongest current evidence does not support one undifferentiated category called “AI in education”. It points instead to a design distinction. Tutor CoPilot, a preregistered randomised trial involving 900 tutors and 1,800 K-12 students, gave human tutors real-time pedagogical suggestions. Students whose tutors had access to the system were four percentage points more likely to master topics; the reported benefit reached nine percentage points among students assigned to lower-rated tutors. The authors also reported that tutors used more guiding questions and were less likely to give away answers.[22] This is a working paper rather than a final verdict, but its architecture matters: the AI strengthened a human teaching relationship.

The official abstract of a 2026 CEPR discussion paper reports a contrasting pattern in a 30-month panel of 26,811 Chinese secondary students. In its difference-in-differences design, AI adoption coincided with higher homework scores, shorter completion time and lower later closed-book and high-stakes examination scores. Reported losses concentrated among users whose unusually short completion time and high homework scores were classified as behaviour consistent with outsourcing.[23] That classification is a proxy, not direct observation. The full methods and robustness material were not accessible during this review, so the paper remains abstract-level warning evidence rather than a verified causal result.

The pilot should therefore test effort-preserving constraints: ask before answering, expose reasoning selectively, retrieve from governed material, request learner explanations, include unaided transfer tasks and preserve points for human intervention. These controls are hypotheses drawn from contrasting intervention designs, not proven safeguards. A system that improves assignments while weakening unaided capability should fail the pilot's quality floor even if its productivity metrics improve.

5. Where do the saved teaching hours go?

The language of educational AI is full of promises to save teachers time. Time saved is often treated as if it disappears. It does not. A reduction in paid human effort always creates a transfer. The released time may become smaller workloads, more students, fewer employees, lower public expenditure, higher profit, better support, additional feedback, more research or a shorter working week.

Hedegreen Research's earlier Saved Hours work argues that acknowledgement is not a distribution mechanism. When the same output requires fewer paid human hours, an institution should make the destination of the released time and value visible.[9] Applied to education, the principle is simple: AI-supported efficiency must not be evaluated only by whether the institution spent fewer hours. It must be evaluated by what the learner and teacher received in return.

A Saved Teaching Hours statement

Every pilot course should publish a Saved Teaching Hours statement answering at least these questions:

1. What defined human work was performed before the intervention?
2. What human work was reduced, and how was the reduction estimated?
3. What changed in learning volume, quality, accessibility and service level?
4. How much human time was reinvested in mentorship, difficult feedback, support or curriculum maintenance?
5. Who received the remaining financial value?
6. What happened to teacher employment, workload, autonomy and learner contact?
7. What new supervision, verification and correction work did the system create?
8. Through which channel was released time or value captured, and by whom?

The statement requires a defined baseline service, a quality floor and comparable old and new human hours. As a hypothetical example, if automated routine feedback removes five teacher hours but new supervision, correction or additional students consume six, the institution has not demonstrated a returned hour. It has changed the work or the target.

Hypothetical Saved Teaching Hours ledger

The following row shows the accounting form, not an expected pilot result. It holds the service scope at 40 learners, one 12-week course run and the same assignment set. Its quality floor requires delayed unaided performance and the feedback service level to be no worse than the baseline.

Ledger field	Illustrative entry
Baseline human time	120 hours under the previous delivery model.
AI-assisted human time before reinvestment	92 hours, including 12 new hours for tutor supervision, correction and coordination.
Gross released time	28 hours: 120 minus 92.
Reinvested human time	16 hours: 8 in mentorship, 4 in accessibility support and 4 in paid curriculum maintenance.
Net released value	12 hours at a disclosed illustrative loaded rate of EUR 50 per hour: EUR 600.
Destination	EUR 300 to the shared maintainer account and EUR 300 returned through a lower learner fee.
Scope and output check	Still 40 learners, the same assignments and no failure of the pre-registered quality floor.

The numbers are deliberately hypothetical. A real ledger must publish the measurement window, role definitions, loaded-cost rule and uncertainty; it cannot count work shifted to learners, unpaid maintainers or hidden correction as a saving.


TALIS 2024 provides a lower-secondary task baseline across teaching, preparation, marking, administration and other work, with substantial cross-system variation.[10] It is not a university workload model or an AI-savings coefficient. An AI system that removes a measured task may create value; whether it improves education depends on the quality floor, new human work and allocation rule.

Automate repetition. Socialise the released time.

The paper's most important political claim is therefore not technological. Saved teaching hours should have a presumptive destination in the learning commons: more learners reached, smaller mentorship groups, paid maintenance of open material, translation, accessibility, independent assessment, public compute and time for teachers to do the work that brought them into teaching.

6. The four commons

The proposal is a federation of four interoperable commons. They may be hosted by different organisations and implemented with different technologies. The pilot must test whether the stack can operate without one actor owning every layer.



6.1 Knowledge Commons

The proposed Knowledge Commons is a network of openly licensed, versioned repositories maintained by teachers, researchers and practitioners. A candidate module would contain learning objectives, concept maps, sources, explanations, examples, exercises, assessment standards, known misconceptions, unresolved disputes, accessibility versions and a public change history.

The design is not a single canonical textbook. Legitimate schools of thought could coexist. A controversial history curriculum could preserve majority and minority interpretations, while a scientific module could mark claims as established, contested or under revision. Forking is intended to preserve continuity when maintainers stop or governance is captured, but provenance survives only if lineage is technically and institutionally enforced.

6.2 Tutor Commons

The proposed Tutor Commons consists of interfaces intended to be interchangeable over the Knowledge Commons. A participating tutor would disclose its model and version, the curriculum version, sources retrieved, the distinction between sourced material and generated explanation, its data policy and its measured cost.

Retrieval over a maintained knowledge layer should precede domain-specific model training in the initial design because it may make correction faster and sources more inspectable. The pilot must test whether the same module remains usable through two model implementations and a static low-compute mode; the paper has not demonstrated that switch.

6.3 Competence Commons

The proposed Competence Commons would define what must be demonstrated and record how it was demonstrated. A candidate competence record should name the standard and version, evidence, assessment methods, conditions, assessors, identity checks, limitations, appeals and any expiry or recertification requirement.

Technical standards already exist for portable digital credentials. W3C's Verifiable Credentials Data Model specifies a three-party ecosystem of issuer, holder and verifier, while 1EdTech's Open Badges packages metadata about achievements and competencies.[11][12] These standards can carry records. They do not establish that the assessment was good. The social and professional trust must come from transparent standards, assessors, audit history and evidence.

6.4 Human Commons

The proposed Human Commons names people and places that the design does not treat as replaceable by a language model: mentors, teachers, examiners, laboratories, studios, workshops, clinics, field sites, local learning circles, safeguarding services and professional communities. Exact requirements must remain discipline- and jurisdiction-specific.

The design goal is not to minimise this layer. It is to stop wasting it on scalable repetition. Human time should be used where human judgement, relationship, responsibility and physical presence change the outcome.

The Human Commons is larger than campus

The social answer is not to simulate a university campus inside a platform. Campus is one valuable human arrangement among many. The Human Commons can include public libraries, community workshops, learning houses, laboratories, unions, museums, studios, schools after hours, local associations and existing universities.

UNESCO's learning-cities framework describes lifelong learning as a city-wide activity spanning education, families, communities, workplaces and modern learning technologies.[16] Its work on Community Learning Centres similarly treats inclusive, flexible and locally owned spaces as infrastructure for people who are outside formal education or underserved by it.[17] These precedents demonstrate alternatives to a residential campus; they do not show that the proposed federation prevents isolation.

A federated learning commons should test whether global knowledge and evidence standards can support dense local human environments rather than displace them. Local nodes would carry friendship, argument, tools, care, culture and physical practice.

7. The teacher after scarcity

The common fear is that AI will make teachers economically redundant. That outcome is neither technically predetermined nor controlled by technology alone. Labour institutions, budgets, policy and governance determine how any released work is allocated.

Teacher workload includes teaching, preparation, marking, administration, collaboration and support.[10] Some explanations and materials are repeated or remain institutionally bounded, but repeated work is not automatically waste. An open commons would test whether selected contributions can persist, be cited, translated, corrected and adapted without hiding the maintenance required.

The pilot should test and fund teacher roles including:

- a maintainer of a public body of knowledge;
- a designer of explanations and authentic problems;
- a researcher of recurring misconceptions;
- a mentor for difficult decisions and stalled learners;
- an examiner capable of testing whether understanding is real;
- an accessibility and localisation steward;
- a moderator of disciplinary disagreement;
- an auditor of the models operating over the curriculum;
- a builder of communities in which learning has human meaning.

These roles are a design hypothesis, not a prediction of the profession. The pilot must ask teachers whether durable public contributions improve autonomy and educational value, and it must measure the new supervision, verification and correction work created by AI.

AI should not automate teachers out of education. It should automate the scarcity around teaching.

The *1,000x* direction defined at the beginning of this paper is a manifesto target, not a forecast or a composite score. Its dimensions must be reported separately so that an increase in reach cannot conceal weaker retention, autonomy or human contact.

A vector, not a slogan

The phrase should never be converted into one vanity number. Learning has no single denominator. A pilot should publish a vector of expansion and quality measures so that gains in one dimension cannot hide losses in another.

Dimension	Proposed measure	Failure signal
Reach	Distinct learners served, including previously excluded groups.	Higher enrolment with falling persistence or support.
Availability	Hours and response latency across time zones.	Always-on access that substitutes for every human contact.
Language	Learners completing in their strongest language; human-reviewed localisation.	Machine translation creates lower-quality parallel tracks.
Accessibility	Use and outcomes across declared access needs.	Accessibility is counted as a feature but not tested with users.
Practice	Meaningful attempts, retrieval, transfer and delayed retention.	More generated answers and fewer learner-produced explanations.
Teacher leverage	Learners reached per teacher-maintainer hour, alongside workload and autonomy.	The same teacher is simply assigned more cases.
Durability	Reuse, correction, translation and forks of teacher contributions.	Material accumulates but becomes stale or ungoverned.
Recognition	Assessments accepted for interviews, progression or credit.	Learners complete but the record has no social value.

The objective is expansion without amputation. The proposed dashboard pairs volume with competence, access with quality, automation with human contact, and cost with the distribution of saved hours. That pairing cannot prevent gaming by itself; it creates visible tests against a thousand-fold increase in low-trust outputs being misreported as learning.

8. The right to prove competence

This is proposed as a normative access principle, not a presently enforceable legal right.

This paper treats recognition as a central gate, but the reviewed evidence does not establish that it is always the strongest one. Degrees, professional licences and institutional credits perform different signalling and permission functions. Existing prior-learning and external-examination systems show that teaching and assessment can sometimes be separated; they do not show that a new record will be trusted.

The European Union has long recognised the principle that learning acquired outside formal education can be identified, documented, assessed and certified. Its 2012 recommendation explicitly included learning through open educational resources and described those four phases as separable elements of validation.[13] This is not a complete alternative university, but it is an important institutional precedent.

The proposed right is therefore specific: a person should be able to seek a transparent and appropriately supervised assessment regardless of where learning occurred, subject to jurisdictional, safety and professional-authority boundaries.

Assessment must become harder, not softer

An open pathway cannot survive if it becomes a credential mill. Low-stakes learning can use self-tests, automated checks and peer review. Medium-stakes assessment can use project review, timed work and oral defence. High-stakes competence may require controlled examination, live practical performance, laboratory sign-off, clinical supervision or independent panels.

AI may help organise evidence, generate test variants and check rubric consistency. It should not be the sole authority for high-stakes credentials. Oral defence and transfer tasks are especially important in an AI-rich environment because they test whether the learner can reason beyond the produced artefact.

The resulting record should be more informative than a decorative badge. It should state exactly what was demonstrated, under which conditions, against which standard, with what evidence and limitations. Current standards make export and verification technically plausible; the pilot must test whether another provider accepts the record without requiring surrender of the underlying work.

The anti-tier principle

The most dangerous failure would be a two-tier system in which wealthy learners keep campuses, networks and prestigious degrees while everyone else receives an AI tutor and an alternative badge that employers quietly discount. The commons must not become discount education for people who cannot buy institutional membership.

The proposed answer is same standard, richer evidence. Where two pathways claim the same competence, the pilot should test them against the same public standard. Cross-pathway access to assessment remains an institutional hypothesis, not a right this paper can already confer. The record should reveal practical artefacts, oral defence, conditions and assessor judgement rather than asking the public to trust a new logo.

Recognition should be earned empirically. Pilot credentials must be tested with independent assessors and employers, and their predictive value should be published. If the record does not create enough trust to secure an interview, enter a further course or perform the stated work, the assessment mechanism has failed regardless of its technical elegance.

A trust ladder

Not every learning record should pretend to carry the same authority. The Competence Commons should expose a trust ladder in which the assessment conditions become progressively stronger and more expensive.

Level	Evidence	What it can honestly claim
0 — Learning trace	Self-checks, tutor interactions and practice history.	Evidence of engagement only; not a credential.
1 — Verified tasks	Automated tests, signed artefacts and reproducible outputs.	A defined result was produced under disclosed technical conditions.
2 — Human-reviewed portfolio	Independent review, rubric, feedback and revisions.	A reviewer judged the work against a public standard.
3 — Challenge assessment	Identity-verified practical work, oral defence, transfer tasks and controlled conditions.	The learner demonstrated the competence without relying on the submitted artefact alone.
4 — Supervised practice	Laboratory, clinical, workshop or field sign-off by an authorised body.	Competence under real or regulated practice conditions; legal permission remains jurisdiction-specific.

The ladder is designed to stop a technical record from being presented as social trust. Whether learners, assessors and employers understand and respect its levels must be tested.

The European micro-credential debate is a useful warning. A 2025 Joint Research Centre report based on 15 exploratory interviews, including only one learner, identifies unresolved problems involving long-term finance, quality assurance, formal recognition, marketisation and qualification frameworks.[24] It does not establish effectiveness or general recognition. The proposed commons should publish fewer, harder standards and reveal exactly where each record sits on the trust ladder.

9. Compute at cost; humans at real cost

The hypothesis is that some scalable tutor interactions may approach their measured technical delivery cost. No pilot has established that result. The educational system as a whole cannot operate at compute cost: curriculum maintenance, support, assessment, mentorship, safeguarding, software, administration, infrastructure and governance require labour and capital, while physical disciplines also require buildings and equipment.[27][28]

The economic proposal is therefore to separate cost layers rather than pretend they vanish.

- Scalable machine costs: inference, storage, bandwidth, routine practice and basic transformation.
- Scarce human costs: maintenance, mentorship, oral defence, appeals, specialist support and safeguarding.
- Scarce physical costs: laboratories, clinics, studios, workshops, secure assessment sites and field training.
- Institutional and lifecycle costs: software, administration, governance, compliance, capital, depreciation and service continuity.

The proposal would allow public funding or subsidy for material and bounded tutor interactions, while pricing scarce human and physical services transparently. Whether an assessment fee can cover real assessment without recreating an opaque bundle requires pilot accounting.

The commons must not depend permanently on unpaid teacher idealism. The pilot should compare paid maintenance with public funding, institutional contributions, philanthropy, membership, employer support and a possible Saved Hours contribution. Existing precedents show partial funding patterns, not that this mix is sufficient.

The accounting must also expose hidden subsidies. A platform that advertises almost-free tutoring while relying on volunteer curriculum labour, public research, underpaid moderation or proprietary learner data may have moved costs rather than removed them.

A separate financing hypothesis: Commons Contribution Covenant

The covenant is one optional financing hypothesis, not a constitutional requirement of the four-commons architecture. The architecture and one-course pilot must remain testable if this funding mechanism fails.

The hypothesis is a pilot covenant attached to trusted participation in one course, not a proposed global levy and not a promise that philanthropy will appear. The knowledge would remain accessible under its licence. During the pilot, operators seeking an

interoperability listing, assessment-registry status or pilot trust mark would accept disclosed contribution duties, subject to legal and competition review. Procurement eligibility is outside the first covenant test.

A pilot contribution account would record money, audited paid maintainer labour and public compute separately; it would not convert them into one score. Learner volume, commercial revenue and documented teaching hours saved may be observed as possible future variables, but v0.1 would not use them to calculate a mandatory rate. Valuation, gaming and audit cost are themselves pilot questions.

Existing open projects show several partial funding patterns. OpenStax explicitly credits philanthropic partners for financing free educational resources.[18] Wikimedia combines volunteer contribution with a professionally staffed foundation and reports that most foundation revenue comes from donors spread across more than 200 countries.[19] Apache demonstrates durable, vendor-neutral and community-governed infrastructure, but its reliance on active volunteers also warns against treating unpaid contribution as a complete labour model.[20] None is a ready-made solution; together they show that open access and organised financing are compatible.

The covenant must not be smuggled into a licence in a way that destroys openness or creates a single permission authority. A non-compliant operator may still use genuinely open material under its licence. It should not automatically receive the commons' trust marks, governance privileges, procurement status or assessment recognition. This separates the right to learn and reuse from the right to present oneself as a responsible steward.

Contribution in three currencies

A first covenant does not need a global tax rate. It needs a one-course contribution account that tests three categories: money for common infrastructure; paid maintainer time under public governance; and public compute for tutors, translation, accessibility or assessment infrastructure. Each category remains visible in its native unit. The pilot must determine whether the account is understandable and auditable, and contributions must never purchase curriculum outcomes.

Participation status	Access	Reciprocal obligation
Open user	May read, learn, reuse and fork according to the open licence.	No fee or permission. Attribution and licence compliance remain required.
Steward operator	May use the commons trust mark and participate in designated governance.	Publishes costs, model dependencies and contributions in money, paid labour or compute.
Registered assessor	May issue records listed in the assessment registry.	Accepts audit, conflict disclosure, appeals and contribution to shared assessment infrastructure.

This structure preserves unconditional access while making trusted stewardship conditional on reciprocity. The first pilot should test the accounting burden as carefully as the revenue: a covenant that costs more to audit than it returns will become theatre.

10. Governance without a new gatekeeper

The proposal will fail if it reproduces the university's gatekeeping under a different logo. It will also fail if openness is treated as the absence of standards. The governing problem is to combine public entry with hard accountability.

Federation and forkability

The commons should be a protocol and federation, not a single platform. Different organisations can host curricula, tutors, assessment bodies or physical facilities. A small interoperability core should define module identity, version, provenance, licensing, competencies and assessment records. It should not dictate one pedagogy or one culture.

Forking is proposed as a constitutional safeguard. It can preserve an exit option when maintainers become inactive, governance is captured, localisation is needed or scholarly disagreement exists. It also creates fragmentation risk. The pilot must demonstrate lineage preservation and record continuity across a real fork or host switch.

A provisional constitutional model

Federation is not an answer unless authority is named. A first pilot should therefore publish a small constitution rather than promise perfect decentralisation. The following bodies are proposed as testable starting points, not permanent offices.

Body	Limited authority	Constraint
Protocol Council	Maintains the smallest interoperability core and release process.	Cannot define disciplinary truth or own curricula.
Domain Assemblies	Maintain named subject modules, sources and competency standards.	Decisions, conflicts and minority reports remain public and forkable.
Teacher-Learner Chamber	Reviews workload, accessibility, data rights and Saved Hours distribution.	No curriculum veto solely on popularity; must state evidence or rights basis.
Assessment Registry	Lists assessors and methods that meet published audit conditions.	Cannot monopolise assessment; competing qualified bodies may coexist.
Independent Ombud	Handles procedural complaints, exclusions and conflicts of interest.	Publishes rulings and cannot silently rewrite standards.

The Apache Software Foundation describes open communications, consensus-oriented decisions, delegated oversight, vendor neutrality and a priority of community over code. [20] This is a relevant but self-described organisational precedent. Its volunteer structure cannot simply be copied into education, and it does not establish that the proposed bodies will be legitimate or independent.

The constitutional core should contain its own expiration rule. Before the pilot, it must define who participates in ratification, how representation and quorum work, how officers are removed, and what happens to assets and learner records if reference

governance lapses. Founders may write version 0.1; they should not inherit permanent sovereignty.

Capture risks

- A commercial model provider controls access or silently changes tutor behaviour.
- A government turns curriculum governance into censorship.
- Prestigious institutions dominate standards through reputation rather than evidence.
- English-speaking and wealthy communities define the global default.
- A small maintainer class becomes an unelected oligarchy.
- Open contribution becomes permanent unpaid labour.
- Portable records become lifelong surveillance files.
- Forkability produces incompatible standards and confusion.

No design eliminates these risks. The paper therefore proposes public decision logs, declared conflicts, time-limited powers, appeals, minority reports, data minimisation, selective disclosure, paid maintenance, model switching and the right to fork. Creative Commons has similarly argued that the governance systems that enabled open sharing have not kept pace with AI and that stronger mechanisms for attribution, agency, transparency and public infrastructure are required.[21]

A public digital-good baseline

The Digital Public Goods Alliance offers a useful baseline for the technical layer. Its standard requires, among other things, open licensing, clear ownership, platform independence, documentation, privacy compliance, open standards and safety protections. It also states an important boundary: recognition of the core solution does not evaluate every local implementation.[25]

The learning commons should adopt the same distinction. A protocol, module or tutor may meet the public-good baseline while a particular school, assessor or deployment still performs badly. Trust marks must therefore be scoped. “Compliant module” cannot mean “effective teaching”, and “registered assessor” cannot mean that every judgement is correct.

- Open licence and a documented right to fork.
- Clear ownership and contributor provenance.
- Platform and model independence.
- Public technical and governance documentation.
- Data minimisation and lawful privacy practice.

- Open interoperability standards.
- Safety, harassment and inappropriate-content procedures.
- A review process with visible issues and revision history.

Metrics also require restraint. A system can count completion, response time and test scores while missing exhaustion, curiosity, confidence, belonging or the ability to recognise when a model is wrong. As Hedegreen Research has argued elsewhere, a metric can reveal and also amputate.[14] The pilot must measure what it can without treating the measured layer as the whole learner.

11. A one-course pilot

The proposal should not begin by declaring a global university. It should begin with one course whose claims can be tested.

Provisional pilot: Open Software Systems 1

Open Software Systems 1 is the provisional first case because its materials can be open, practical work can be stored in repositories, some tests can be automated and understanding can be challenged through oral defence. This selection remains conditional on willing maintainers, independent assessors, learner safeguards and a feasibility review.

The pilot would include:

1. a public, versioned curriculum maintained by a small teacher-practitioner group;
2. at least two interchangeable AI tutor implementations and a static non-AI mode;
3. source-bound explanations and visible correction mechanisms;
4. practical repository tasks with automated tests;
5. scheduled human mentorship checkpoints;
6. a final oral defence and independent assessor;
7. a portable competence record linking to reproducible artefacts;
8. a full cost report and Saved Teaching Hours statement;
9. an explicit red-team and stop conditions;
10. learner consent, data-minimisation, deletion and selective-disclosure tests;
11. assessor calibration, disagreement and inter-rater reliability rules.

What success would mean

Success would not mean that the cohort completed a chatbot course. Before enrolment, the pilot must pre-register thresholds and measurement times for transfer, oral defence, delayed retention, model-error detection, record export and human contact. Teacher outcomes must include workload, autonomy and all new supervision or correction work rather than only a report of higher-value activity.

What failure would teach

The pilot protocol must distinguish immediate stop conditions from revisable failures. Repeated core misteaching, unsafe data handling, invalid assessment, hidden labour or uncorrectable governance failures should stop the affected component. Dependence, unstable cost, weak support or failed switching may trigger redesign only if the pre-registered quality floor remains intact. A failed pilot would still be useful if its evidence remains public.

A minimum viable federation

The first pilot must contain more than one actor in each place where monopoly is being challenged. Otherwise it will only demonstrate that Hedegreen Research can operate a small online course.

Phase	Minimum mechanism	Public artefact
0 — Constitution	Three teacher-practitioner maintainers, learner representation, conflict rules and an expiry date for founder control.	Curriculum constitution and decision log.
1 — Knowledge	One versioned module with sources, misconceptions, exercises, accessibility variants and an open licence.	Repository v0.1 and source ledger.
2 — Tutors	At least two model implementations plus a static low-compute route.	Tutor contracts, prompts, costs and error reports.
3 — Human node	A local or synchronous learning group with named mentors and safeguarding boundaries.	Mentorship protocol and contact-hours ledger.
4 — Assessment	Two independent assessor panels using the same public standard and trust-ladder level.	Rubrics, calibration results, appeals and anonymised records.
5 — Audit	Pre-registered outcomes, Saved Teaching Hours accounting, red-team and publication of null or negative results.	Pilot report, cost statement and revision decision.

The pilot should be time-bounded and modest enough to fail. The current 20–50 learner and 12–16 week ranges are planning hypotheses, not evidence-based sample or duration decisions. They must be justified against learning objectives, recruitment, assessor capacity, statistical use and cost before registration.

The federation is real only if a learner can change tutor, a maintainer can be replaced, an assessor can disagree, and the work remains usable when one host disappears. Interoperability must be demonstrated as behaviour, not asserted in a diagram.

12. Positionality: why Hedegreen Research is writing this

This proposal is written from outside the traditional university system. Hedegreen Research emerged through open sources, AI-assisted learning, public work logs, tools, articles and auditable artefacts rather than through institutional appointment.

That history is context, not proof. HR does not demonstrate that universities are unnecessary, and its work should not be accepted because the founder lacks a degree. The work should be judged through sources, methods, corrections, tools and whether others can reproduce or defeat the claims.

The narrower point is that institutional membership and intellectual contribution are not identical. Prior-learning, credit-transfer and external-examination mechanisms already assess some learning acquired elsewhere, but they remain jurisdictional and provider-dependent. This paper asks whether that fragmented possibility can become a more portable and transparent pathway without overstating its trust.

Hedegreen Research was built because the open knowledge institution its founder needed did not exist.

HR may prepare a protocol and candidate reference implementation. Any pilot or publication remains a separate decision. HR must not own the resulting commons; licences, succession rules and governance controls should keep materials and standards usable, inspectable and forkable without Hedegreen Research.

Build before permission

This paper is not primarily a request for existing institutions to convene another discussion about the future. Low-stakes components can be tested before universal agreement: prepare one module, run more than one tutor over it, convene a local learning group, disclose costs and invite hostile review. Assessment, learner data, safeguarding, public procurement and professional recognition remain subject to their actual legal and institutional gates.

Policy matters when the pilot reaches recognition, public finance, professional regulation or procurement. It should not be used as a reason to postpone the first working mechanism. The old institutions are welcome as contributors and hosts. They are not required as founders and should not receive ownership by default.

What this revision accepts — and rejects

- Accepts: governance, financing and assessment trust are mechanisms that must be specified and tested, not postponed behind the word “commons”.

- Accepts: without an anti-tier principle, open AI education could become a lower-status track for people excluded from elite institutions.
- Rejects: the Human Commons must reproduce four residential years on a university campus. Campus is one possible node, not the universal template.
- Rejects: the proposal's success depends first on winning permission from policymakers. It depends first on teachers, learners and builders producing a credible pilot.

13. Limitations and unresolved tests

This paper is an institutional proposal, not a completed systematic review, a costed implementation or a pilot result. Its reviewed precedents separate individual functions; none validates the complete four-commons architecture. One recent CEPR source remains abstract-only because its full methods and robustness material were not accessible.

The proposal does not yet establish:

- that a teacher-governed curriculum remains current and reliable at meaningful scale;
- that learners can switch tutors, hosts or assessors without losing records;
- that independent assessors agree sufficiently or that employers trust the resulting evidence;
- that money, paid labour and public compute can be valued comparably without excessive audit cost or gaming;
- that Saved Teaching Hours can be measured without surveillance, hidden new work, output redefinition or target capture;
- that the proposed governance bodies recruit legitimate teacher and learner representation or resist institutional, commercial and state capture;
- that privacy, deletion, selective disclosure, accessibility and safeguarding work across learner populations and jurisdictions;
- that low-stakes experimentation can advance into regulated education or professional permission.

The proposed *1,000x* vector is a manifesto direction, not a forecast. Open standards do not create social recognition, open licences do not fund maintenance, and AI availability does not establish learning. Hedegreen Research's position outside a university motivates the question but does not validate the answer.

14. Conclusion

The case for testing a federated open learning commons does not depend on believing that AI is wise, neutral or sufficient. Bounded trials show that some structured explanation and practice can become more scalable. That changes the available institutional choices; it does not make every remaining educational scarcity artificial.

Universities bundle knowledge, tutoring, assessment, facilities and recognition through law, finance, history and reputation. The reviewed precedents show that some functions can be separated, while authority, safety, trust and physical practice remain conditional. The paper proposes testing how far that separation can responsibly go.

The alternative proposed here is not a future without teachers or universities. It is a pilot in which teachers would govern the knowledge layer, model switching would be demonstrated, learners would export evidence, assessors would apply public standards, and Saved Teaching Hours would disclose whether released time returned to learning. None of those mechanisms should be written as accomplished before the pilot.

The objective is not fewer teachers. It is to expand learning across reach, language, practice and human attention without reducing education to scale alone.

Open-education and open-source communities provide partial precedents for shared infrastructure, professional maintenance, distributed contribution and forkability. Structured AI trials make a bounded personal-interaction layer plausible. The next step is not to declare a replacement institution, but to build a mechanism small enough to fail publicly.

Automate repetition. Socialise the released time. Leave teaching to those who came to teach.

References and source boundary

This is a public review draft, not a completed systematic review. References below support specific factual boundaries and precedents. They do not establish the full architecture or its expected effects. Source adjudication was updated on 29 July 2026.

- [1] UNESCO. Recommendation on Open Educational Resources (OER), adopted 25 November 2019. <https://www.unesco.org/en/open-educational-resources>
- [2] MIT OpenCourseWare. Get Started / About: open access, no enrolment, and no MIT credit or certificates for OCW use. <https://ocw.mit.edu/pages/get-started/>
- [3] Open Source Society University. Computer Science FAQ: OSSU organises undergraduate-style learning resources but does not offer a degree. <https://cs.ossu.dev/FAQ.html>
- [4] OpenStax. About, licensing and organisational records: nonprofit, professionally maintained textbooks and current CC BY-NC-SA licensing for new and updated editions; adoption figures remain operator-reported. <https://openstax.org/about/>
- [5] OERu. How It Works, learner records and closure notice: free self-directed learning with optional fee-for-service assessment under partner authority; permanent closure announced for 31 July 2026. <https://oeru.org/>
- [6] Kestin, G. et al. AI tutoring outperforms in-class active learning: an RCT introducing a novel research-based design in an authentic educational setting. Scientific Reports 15 (2025). <https://doi.org/10.1038/s41598-025-97652-6>
- [7] De Simone, M. E., Tiberti, F., Barron Rodriguez, M. R., Manolio, F. A., Mosuro, W., and Dikoru, E. J. From Chalkboards to Chatbots: Evaluating the Impact of Generative AI on Learning Outcomes in Nigeria. World Bank Policy Research Working Paper 11125 (2025). <https://doi.org/10.1596/1813-9450-11125>
- [8] UNESCO. Guidance for generative AI in education and research (2023). <https://www.unesco.org/en/articles/guidance-generative-ai-education-and-research>
- [9] Hedegreen, D. The Missing Mechanism. Hedegreen Research (2026). <https://hedegreenresearch.com/articles/the-missing-mechanism/index.html>
- [10] OECD. The demands of teaching: Results from TALIS 2024 (published 2025). https://www.oecd.org/en/publications/results-from-talis-2024_90df6235-en/full-report/the-demands-of-teaching_0e941e2f.html
- [11] W3C. Verifiable Credentials Data Model v2.0, W3C Recommendation (15 May 2025). <https://www.w3.org/TR/vc-data-model-2.0/>
- [12] 1EdTech Consortium. Open Badges standard. <https://www.1edtech.org/standards/open-badges>
- [13] Council of the European Union. Council Recommendation of 20 December 2012 on the validation of non-formal and informal learning (2012/C 398/01). https://www.cedefop.europa.eu/files/Council_Recommendation_on_the_validation_20_December_2012.pdf
- [14] Hedegreen, D. Do Not Worship the Measurable. Hedegreen Research (2026). <https://hedegreenresearch.com/articles/do-not-worship-the-measurable/index.html>
- [15] University of London. History of Distance Education: remote examinations, examining-body history and assessment-centre network. <https://www.london.ac.uk/centre-online-distance-education/about/history-distance-education>
- [16] UNESCO Institute for Lifelong Learning. UNESCO Global Network of Learning Cities: lifelong learning across education, families, communities, workplaces and technology. <https://www.uil.unesco.org/en/learning-cities>
- [17] UNESCO Institute for Lifelong Learning. Community learning: Community Learning Centres as inclusive, flexible and locally owned lifelong-learning spaces. <https://www.uil.unesco.org/en/community-learning>
- [18] OpenStax. Philanthropic Support To Provide Free Textbooks: supporter and partnership model for open educational resources. <https://openstax.org/supporters>
- [19] Wikimedia Foundation. 2023-2024 Annual Report: volunteer contribution, professional infrastructure and donor-supported financing. <https://wikimediafoundation.org/annualreports/2023-2024-annual-report/>

- [20] Apache Software Foundation. Briefing: The Apache Way: public communication, consensus, responsible oversight, vendor neutrality and community over code. <https://www.apache.org/theapacheway/>
- [21] Creative Commons. Strengthening the Commons for the AI Era: governance, attribution, agency and reciprocity challenges for shared knowledge used by AI. <https://creativecommons.org/2026/05/13/from-signals-to-infrastructure-strengthening-the-commons-for-the-ai-era/>
- [22] Wang, R. E., Ribeiro, A. T., Robinson, C. D., Loeb, S., and Demszky, D. Tutor CoPilot: A Human-AI Approach for Scaling Real-Time Expertise. Preregistered working paper / arXiv:2410.03017 (2024, revised 2025). <https://arxiv.org/abs/2410.03017>
- [23] Strömberg, D., Lei, V., and Wu, Y. The Generative AI Learning Penalty: Evidence from Chinese Secondary Education. CEPR Discussion Paper 21577 (2026). Abstract-level use only; full methods and robustness were not accessible during review. <https://cepr.org/publications/dp21577>
- [24] European Commission, Joint Research Centre. *Micro-credentials in higher education: Are they boosting digital and green skills for the twin transition?* JRC142086 (2025). Exploratory qualitative study with 15 interviews and one learner participant. <https://publications.jrc.ec.europa.eu/repository/handle/JRC142086>
- [25] Digital Public Goods Alliance. Digital Public Goods Standard: baseline indicators for open software, data, AI systems and content collections. <https://www.digitalpublicgoods.net/standard>
- [26] Hedegreen, D. *The Saved Hours Doctrine: Productivity, AI and the Question of Who Captures Time*. Hedegreen Research / Zenodo, version 0.94 (2026). <https://doi.org/10.5281/zenodo.20400880>
- [27] OECD. *Education at a Glance 2025 — Key system-level indicators of education finance*. https://www.oecd.org/en/publications/education-at-a-glance-2025_1c0d9c79-en/full-report/key-system-level-indicators-of-education-finance_178e2c40.html
- [28] The Open University. *Financial Statements for the year ended 31 July 2025*. https://university.open.ac.uk/foi/main/sites/www.open.ac.uk.foi.main/files/Financial_Report_2025_online.pdf

Questions that must be answered before v1.0

- Which existing open-education and assessment systems most closely approximate the four-commons architecture?
- What is the strongest evidence that structured AI tutoring improves durable learning, not only immediate task performance?
- What evidence shows harm from unstructured or dependent AI use?
- How can Saved Teaching Hours be measured without making teacher work more surveilled?
- Which one-course pilot has willing maintainers and independent assessors?
- What legal form can issue an open competence record without implying regulated professional permission?
- How should learner data be minimised while still preserving auditable evidence?
- What funding formula pays maintainers before volunteer burnout becomes the hidden business model?
- How can forks remain interoperable without one central curriculum authority?
- What would make an employer trust the record enough to interview or hire a learner?
- Which tutor behaviours preserve learner effort, and which become disguised answer outsourcing?
- Can two independent assessor panels reach acceptable agreement on the same public standard?
- What is the smallest contribution covenant that funds maintenance without creating costly bureaucracy?
- Which dimensions of the 1,000x vector can improve without reducing retention, autonomy or human contact?