EQGENIX INC.

The Expansion of Human Possibility

An EQGENIX Human–AI Formation Thesis

Yvonne G. Hall

EQGenix Inc., Canyon Lake, California, USA

ORCID: https://orcid.org/0009-0000-0081-7713

Version: 0.3 — first public release

Date: 16 September 2026

DOI: 10.5281/zenodo.22805703

Document type: White paper · public-intellectual thesis (not peer-reviewed)

Suggested citation: Hall, Y. G. (2026). *The expansion of human possibility: An EQGENIX human–AI formation thesis* (Version 0.3). EQGenix Inc. https://doi.org/10.5281/zenodo.22805703

Abstract

Generative artificial intelligence can raise the quality of a person's work immediately, but improved output is not evidence that the person has become more capable. This paper argues that the central design question of the AI era is not how much work machines can perform but which human capacities must be formed, protected and strengthened for people to govern increasingly capable systems. Grounded in an explicitly stated anthropology of stewardship, it proposes Formation AI™: AI configured to increase human capability rather than only complete tasks. The paper reviews recent experimental evidence showing that AI assistance can produce three distinct outcomes once assistance is removed — erosion, parity or strengthening — and that design determines which occurs. It then specifies an architecture comprising a nine-step Formation Engine with verification built into the cycle, five AI Gates governed by developmental appropriateness and demonstrated capability, adaptive and embedded formation, a child-safety boundary, and public behavioral indicators. It sets an evidentiary standard under which parity after withdrawal is a safety floor rather than evidence of formation, confronts the strongest counterargument, and states the

conditions under which the thesis should be revised or rejected. The framework has not yet been empirically validated; its propositions are offered for testing.

Keywords: generative AI; human formation; cognitive offloading; learning versus performance; AI in education; cognitive sovereignty; withdrawal and transfer; stewardship

1. Introduction

> *Prove all things; hold fast that which is good.*
>
> 1 Thessalonians 5:21 (KJV)

Artificial intelligence should not diminish the human. Properly governed, it should expand what a developed human being is capable of becoming and doing.

The defining question of the AI age is not only how capable machines will become. It is whether people will deliberately develop the capacities required to think, create, discern, regulate, relate, decide, adapt and lead alongside them. When answers are abundant, judgment becomes the scarce capacity.

Two responses dominate public debate. One urges fear and restriction; the other urges adoption at any cost. This paper proposes a third: **formation**. People should develop the capacities needed to govern powerful technology before surrendering cognitive operations to it, and AI can be deliberately configured to create the conditions in which people become more capable, not merely more productive. We call this approach **Formation AI™**. Its purpose is not to make human cognition unnecessary. Its purpose is to expand human possibility by developing people capable of governing increasingly powerful intelligence.

The paper proceeds in four movements. Sections 2–4 state the anthropological frame, define the problem and review the evidence. Sections 5–11 specify the architecture. Sections 12–14 set the evidentiary standard, confront the strongest objection and address motivation. Sections 15–19 extend the model across populations, publish behavioral indicators, state limitations and set out the research proposition and falsification conditions.

The Formation AI™ evidentiary standard

Parity after withdrawal is a safety floor, not evidence of successful formation. Formation requires evidence that the human capability itself strengthened — performed independently, retained over time where retention is relevant, and carried into new contexts.

2. An Anthropology of Stewardship

Every theory of education carries a theory of the person. This paper states its own. EQGENIX holds Scripture as its source authority and treats scientific evidence as corroborating witness. The anthropology below is therefore the frame from which the architecture was derived, not a commentary added afterward. Readers who do not share this starting point can evaluate every empirical claim and falsification condition in this paper on its own terms.

2.1 Made to be formed

> *I will praise thee; for I am fearfully and wonderfully made.*
>
> Psalm 139:14 (KJV)

A person is not a container for outputs. Human capacities — reasoning, attention, voice, moral judgment, relationship — are part of what it means to be made, and they are developed through use. This is the ground of **Built, Not Downloaded**™: what is most valuable in a person cannot be installed from outside. It must be built.

2.2 Cultivate and keep

> *And the LORD God took the man, and put him into the garden of Eden to dress it and to keep it.*
>
> Genesis 2:15 (KJV)

The first human vocation is given in two verbs. To *dress* is to cultivate — to work, extend and bring out what is latent. To *keep* is to guard what has been entrusted from loss. Technology belongs to cultivation: tools extend what people can do. But cultivation without keeping becomes consumption. A person who gains extraordinary reach while losing the capacities that make judgment, verification and responsibility possible has spent something entrusted.

This double mandate gives the architecture its two rails. **Cognitive Sovereignty**™ is the keeping rail — the floor of human capacity to be protected. **The Expansion of Human Possibility** is the cultivating rail — the horizon of what developed people can reach. Expansion that erodes the floor fails the mandate. Protection that refuses expansion fails it too.

2.3 Faithful over a few things

> *Thou hast been faithful over a few things, I will make thee ruler over many things.*
>
> Matthew 25:21 (KJV)

Responsibility grows with demonstrated faithfulness. This is the logic of progressive responsibility that governs how Formation AI™ relates people to machines: deeper cognitive delegation is extended as a person demonstrates the capacity to govern it. Access is not a reward and restriction is not a punishment; both are stewardship matched to readiness.

2.4 Prove all things

Verification is not a modern literacy skill bolted onto faith; it is commended in the text itself (1 Thessalonians 5:21). A person entrusted with powerful tools is responsible for testing what those tools produce and holding fast only to what proves good. This is why verification is built into the Formation Engine rather than taught as an optional add-on (Section 7).

These passages frame principles about the human person and entrusted responsibility. They are not offered as proof texts for a particular technology or study design.

3. Performance Is Not Formation

AI can make a person appear more capable almost instantly: it can improve prose, summarize difficult material, generate ideas, solve problems and analyze data. Those capabilities are extraordinary. But an improved product does not necessarily mean an improved person. Learning scientists have warned explicitly that performance gains produced with generative AI can occur without the deep cognitive and metacognitive processing on which durable learning depends, and should not be mistaken for learning (Yan et al., 2025).

Cognitive offloading itself is ordinary and often rational; people routinely use external resources to reduce cognitive demand (Risko & Gilbert, 2016). Its costs are conditional. In three experiments with 172 participants each, offloading working-memory demands improved immediate task performance but reduced later memory for the offloaded information — yet participants who offloaded while aware of an upcoming memory test

largely counteracted that cost (Grinschgl et al., 2021). Whether a tool costs later depends on what is offloaded, when, and with what intention.

The concern of this paper is therefore narrower than tool use. It is **premature substitution**: assistance entering during the formation of the very capacity that the assistance performs. If a learner receives the argument before learning to reason through it, the prose before developing a voice, or the judgment before weighing the evidence, the task is completed while the developmental opportunity disappears.

> **Performance AI optimizes the product. Formation AI™ optimizes the person.**

4. The Evidence: Erosion, Parity and Strengthening

The research base is young and moving quickly. The studies below were selected because each measures something beyond AI-assisted performance: performance after assistance is removed, retention after delay, the effect of who controls access, or the value of self-generation. Peer-reviewed articles and working papers are labeled separately because they carry different evidentiary weight.

Study	Population and design	Finding	Limits
Bastani et al. (2025), *PNAS* · peer-reviewed	Nearly 1,000 high school mathematics students; field experiment; standard GPT-4 interface vs. safeguarded tutor vs. no access	Assisted practice grades rose 48% (standard) and 127% (safeguarded). With AI removed, the standard-interface group scored 17% below no-access students; safeguards largely eliminated the penalty.	One subject and country. The safeguarded tutor reached parity; it did not show a positive unassisted effect.
Barcaui (2025), *Social Sciences & Humanities Open* · peer-reviewed	120 business undergraduates randomized; ChatGPT-assisted vs. traditional study; pre-registered; surprise test at 45 days	Retention 57.5% vs. 68.5% (d = 0.68); the effect remained after adjusting for time on task.	Convenience sample at one Brazilian university; 85 of 120 completed the delayed test; unrestricted use only.
Kestin et al. (2025), *Scientific Reports* · peer-reviewed	194 introductory physics students; randomized crossover over two consecutive weeks;	Significantly greater learning in less time with the AI tutor; higher reported engagement and motivation.	Immediate post-tests only; no retention or transfer; highly selected population.

Study	Population and design	Finding	Limits
	engineered AI tutor vs. in-class active learning		
Poulidis et al. (2025) · working paper	Over 200 chess club students; 12 weeks; system-timed AI tips vs. the same plus on-demand help	Both improved, but on-demand students gained less than half as much (30% vs. 64%). Strong students also over-requested help; intrinsic motivation offered only partial protection.	Not yet peer-reviewed; bounded domain; volunteer club members.
Chung et al. (2026) · working paper	Students in ten Taipei high schools; five-month Python course; all used a designed AI tutor; fixed vs. adaptive practice sequencing	Adaptive sequencing raised unassisted final exam performance by 0.15 SD; mediation analysis pointed to engagement.	Isolates adaptive sequencing, not AI vs. no AI; not yet peer-reviewed.
Brod et al. (2025), *npj Science of Learning* · peer-reviewed	95 children aged 5–7; generating predictions vs. observing matched predictions; no AI involved	Self-generation produced deeper conceptual understanding, associated with greater experienced agency.	Not an AI study; single task; correlational link to agency.

4.1 Three outcome zones

Zone	After assistance is removed	Interpretation
Erosion	Independent performance falls below comparison.	The tool performed; the person weakened. Failure.
Parity	Independent performance matches comparison.	Safety floor met. No harm shown — and no formation shown.
Strengthening	Independent performance exceeds comparison, holds over time and transfers.	Evidence that the human capability grew. The target.

The evidence already spans all three zones. Unrestricted access has produced erosion in mathematics and in delayed retention. Safeguards have produced parity. Early signals of strengthening now exist: adaptive practice sequencing improved unassisted exam performance over a semester, and system-regulated assistance more than doubled skill gains compared with on-demand access. The conclusion is modest and consequential: **design determines the zone, and the zone must be measured rather than assumed.**

5. Formation AI™ and the Architecture

Most general-purpose AI is optimized to be helpful: the user asks and the system answers. Formation AI™ begins with a different objective — what should the system do now if the desired outcome is increased human capability? Sometimes the right response is an answer. Sometimes it is a question, counterexample, simulation, incomplete clue, feedback sequence or harder problem. Sometimes it is that AI should remain closed. Success is judged by what the person can later retrieve, reason through, create, verify, regulate, explain, transfer and govern when assistance changes or disappears.

Layer	Name	Function
Philosophy	Built, Not Downloaded™	What EQGENIX stands for
AI approach	Formation AI™	AI develops the person, not merely the product
Developmental system	Formation Engine	Continuous human-capacity development (Section 7)
Progression	AI Gates 0–4	Governed AI relationships matched to readiness (Section 8)
Personalization	Adaptive Formation™	Prescription responds to demonstrated capacity (Section 9)
Experience design	Embedded Formation™	Formation and evidence inside meaningful experiences (Section 10)
Application	Missions	Where capacities are exercised and demonstrated
Floor rail	Cognitive Sovereignty™	The human retains cognitive authority
Horizon rail	The Expansion of Human Possibility	Technology enlarges what developed humans can become and accomplish

The rails are not final steps in a sequence. Cognitive Sovereignty™ governs every step from beneath; the Expansion of Human Possibility draws every step upward; and the AI Gates regulate the space between them, so that each Gate widens reach only when the floor holds. The single design requirement is **expansion without erosion**: governed, AI-augmented reach that grows while independent capability holds and, by the evidentiary standard, strengthens.

5.1 Two layers of human capacity

The architecture distinguishes what is developed from what must hold before delegation deepens.

Layer 1 — the twelve Built, Not Downloaded™ capabilities are what formation develops: Critical Thinking, Working Memory, Deep Reading, Metacognition, Creative Synthesis, Original Writing Voice, Frustration Tolerance, Quantitative Reasoning, Spatial Reasoning, Conversational Depth, Boredom Tolerance and Moral Reasoning. They are a developmental canon, not twelve equally validated psychometric constructs.

Layer 2 — the six Formation Capacities are what must hold under pressure before AI delegation deepens: Interruption Latency, Judgment Ownership, Verification Capacity, Effort Tolerance, Position-Holding Under Relational Cost, and Accountability Retention. These, not general ability, form the delegation floor. A brilliant pattern-finder who cannot yet verify is precisely the person a fluent machine persuades most easily. Ability alone unlocks nothing; sovereignty unlocks delegation.

6. Cognitive Sovereignty™ and Hold the Pen

Cognitive Sovereignty™ is not independence from technology. It is retained authority over cognition while using technology. A sovereign person can originate, question, choose, verify, reject, revise, explain, withdraw and accept responsibility for consequential judgments. The sovereign question is not “Did you use AI?” but **“Who governed whom?”**

Hold the Pen is the EQGENIX practice through which that authority is kept. Its governing rule is that assistance may contribute to the work while authorship, accountable judgment and the final consequential decision remain human. Learners practice it through five moves: **Think First, Bring It In, Check It, Pen Back,** and **Take It Somewhere New**. At any point a person should be able to say what the machine contributed, what they contributed, what they rejected, what they verified, and who owns the decision. This matters most where polished output can conceal weak ownership: a fluent artifact cannot by itself show that its author can explain, defend, adapt or reproduce the reasoning behind it.

7. The Formation Engine

Human formation is not a course a person completes. The Formation Engine runs as a continuous cycle:

FORM → PRESS → DEMONSTRATE → AUGMENT → VERIFY → WITHDRAW → TRANSFER → RE-CALIBRATE → ADVANCE

Step	Function
FORM	Develop the targeted human operation.
PRESS	Introduce calibrated difficulty — ambiguity, distraction, contradiction, load, emotional activation, persuasion or time pressure — so capability forms under the conditions in which it will be needed.
DEMONSTRATE	Require observable evidence of the human operation before deeper delegation.
AUGMENT	Introduce the AI relationship permitted by the person's Gate and the task.
VERIFY	Require the person to check claims, reasoning, evidence, fit and consequences rather than accept fluent output.
WITHDRAW	Remove or reduce assistance where independent capability is material.
TRANSFER	Require the operation in a meaningfully different context.
RE-CALIBRATE	Update the developmental profile using the new evidence.
ADVANCE	Increase complexity, responsibility or the sophistication of augmentation when justified.

VERIFY follows AUGMENT deliberately. Augmentation without verification may improve output while weakening governance; in Formation AI™, "prove all things" is part of the formation sequence itself.

7.1 Crosswalk

The Engine is the research description of a cycle already used in EQGENIX curriculum design, where each activity carries an AI condition, and in the learner-facing Hold the Pen practice. One cycle, three audiences:

Curriculum AI condition	Formation Engine	Hold the Pen move	Purpose
ABSENT	FORM · PRESS ·	Think First	Protect formation;

Curriculum AI condition	Formation Engine	Hold the Pen move	Purpose
	DEMONSTRATE		establish independent evidence
DELAYED	DEMONSTRATE → AUGMENT	Think First, then Bring It In	Lock the human position first; prevent anchoring
PERMITTED	AUGMENT → VERIFY	Bring It In · Check It	Use assistance without surrendering governance
CONTESTED	PRESS → AUGMENT → VERIFY	Check It	Train discrimination against persuasive machine error
WITHDRAWN	WITHDRAW → TRANSFER	Pen Back · Take It Somewhere New	Test ownership, resilience and generalization

8. AI Gates and Progressive Responsibility

Formation AI™ does not treat more AI as advancement. Gates are developmental conditions, not status badges or rewards for using more technology.

AI Gate	AI role	Human floor requirement	Evidence before advancing
0 · Protected Formation	Absent, or working behind the glass	Person generates or attempts the target operation	Independent baseline and formation evidence
1 · Challenger	Questions, counters, supplies planted errors, changes constraints	Human reasoning remains primary	Responds to challenge without answer substitution
2 · Training Partner	Adaptive practice, debate and feedback	Person performs the target operation	Repeated evidence under calibrated variation
3 · Co-Builder	Assists construction after protected origination	Human authorship, verification, decision ownership	Ownership record plus withdrawal and transfer
4 · Governed Delegation	Performs selected delegated work	Person sets objective and standard, checks, and accepts responsibility	Justified delegation, audit, and retained competence where material

8.1 Two keys, stricter wins

Movement between Gates requires two keys turned together: **developmental appropriateness** and **demonstrated capability**, judged against the Formation Capacities. Age alone never unlocks a Gate, and neither does performance alone.

A person's Gate sets their overall range of AI relationships. Each task also sets an **assistance limit**, because the task exists to form a particular operation. A learner at Gate 3 may co-build on a creative project while AI stays closed during an independent retrieval demonstration. When Gate and task differ, **the stricter condition governs**. Gates are reversible: when the floor stops holding, the range narrows until it holds again.

9. Adaptive Formation™

People do not develop uniformly. One person shows strong original generation but weak verification; another reasons well under ordinary conditions and loses calibration under authority or time pressure. Adaptive Formation™ changes the developmental prescription — the active capacity, pressure condition, human-first requirement, permitted AI role, evidence requirement, withdrawal condition and transfer demand — rather than merely changing content difficulty. It never collapses a person into a single readiness score.

Recent evidence sharpens one design implication. When chess students could summon AI help at will, they gained less than half as much as students whose AI tips were timed by the system, and neither skill nor motivation fully protected them from over-reliance (Poulidis et al., 2025). Separately, adaptive sequencing of practice difficulty improved unassisted exam performance (Chung et al., 2026). **The developmental system — not momentary convenience alone — may sometimes need to decide when assistance appears.** That governance is disclosed, never hidden; it loosens as self-regulation is demonstrated; and teaching people to regulate their own AI use is itself a formation target.

Adaptation has limits. It may change difficulty, task world, role and pacing. It may not skip a Gate, relax a safety boundary or override the stricter-wins rule. And the next challenge should be *better* — sometimes harder, sometimes the same difficulty in a new world.

10. Embedded Formation™, Missions and Productive Difficulty

Formation cannot become an endless sequence of tests. **Embedded Formation**™ places formation and evidence inside meaningful experiences — investigations, simulations, debates, design challenges, relational practice and field problems. Assessment traditions have long connected constructs, tasks, observable behaviors and inferences inside authentic activity, including measurement embedded in games (Shute & Ventura, 2013). EQGENIX extends that logic from measurement toward deliberate formation; the extension is a proposition requiring validation.

Missions are the proving ground. A learner does not merely define verification; she meets information worth verifying. An officer does not merely study regulation; he regulates while deciding under pressure. A veteran does not merely discuss adaptability; she completes an unfamiliar civilian task requiring it. **The mission is the measure.**

Two commitments keep Embedded Formation™ consistent with sovereignty. It is **embedded for the participant and visible to those responsible**: for minors, parents and guardians can see the capacity being developed, the AI condition, the evidence category observed and the safety boundaries, while proprietary scoring remains protected. And **withdrawal stays a visible, chosen event**, because only the moment a person knowingly sets the tool down shows what belongs to them.

10.1 Productive difficulty

Formation AI™ removes sterile friction — confusing interfaces, needless waiting, irrelevant complexity — while preserving, and when appropriate creating, formative friction. Actively retrieving information strengthens long-term retention compared with restudying (Roediger & Butler, 2011; McDermott, 2021). Children who generated their own predictions learned more deeply than children who observed matched predictions (Brod et al., 2025). The easiest path to completion is not always the best path to development.

11. Child-Safety Boundary

> **Generation happens behind the glass. No live generative-AI interaction for learners under 14 in the current EQGENIX youth design.**

This is an EQGENIX design and safeguarding rule, not a claim that research has established 14 as a universal threshold. For younger learners, generative systems may prepare scenarios, variations and planted errors for approved adults or the platform, and every generated item is reviewed before a child sees it — a "planted error" that turns out to be correct would teach children to distrust good evidence.

11.1 A graduated path, not a cliff

To avoid an abrupt jump at 14, the design progresses through modes rather than switching on all at once. Modes accumulate: each stage keeps the earlier modes and adds one.

Stage	Mode	What the learner does
From age 8	Authored AI	Works with reviewed, fixed AI-generated material, clearly labeled, including planted errors to detect
From age 11	Branching AI	Navigates pre-written responsive exchanges that push back, hint and mislead, by choice rather than free typing
From age 12	Adult at the keyboard	Drafts the question on paper; a parent or facilitator enters it into a live system; the learner verifies and judges the output
From age 14	Bounded live AI	Uses a governed system beginning at the learner's current Gate — usually Challenger or Training Partner — with parental disclosure, privacy controls, moderation and monitoring; Gates 3–4 still require both keys

Age never overrides a more restrictive task or capacity condition. Implementations involving children's data are designed to comply with the U.S. Children's Online Privacy Protection Rule.

12. The Evidentiary Standard

AI-assisted performance cannot by itself establish that formation occurred. Across the evidence in Section 4, the assisted phase alone would repeatedly have told the wrong story: impressive assisted gains coexisted with lower unassisted scores and weaker delayed retention.

- **Erosion** — independent performance below comparison — is failure.

- **Parity** — independent performance equal to comparison — is the safety floor. Necessary; not formation.
- **Strengthening** — independent performance above comparison, retained where relevant, and transferred — is the evidence formation requires.

A positive Formation AI™ claim therefore requires a study that defines the targeted construct in advance and shows credible strengthening on outcomes appropriate to that construct — ideally independent performance, delayed retention, transfer, calibration or governance of assistance — against an appropriate comparison condition. Research should preregister hypotheses and analyses where feasible, separate feasibility from effectiveness, distinguish assisted from independent outcomes, and state population and task limits rather than generalizing across domains.

Withdrawal is not a claim that people should live without tools. It is an **evidence condition**, used when independent capability is material to verification, resilience, responsibility, further learning or safe governance.

13. Strongest Counterargument: The Calculator Objection

13.1 The objection at full strength

Every powerful tool has provoked this worry, and the worry has usually been wrong. Writing was said to ruin memory; calculators, arithmetic. Capability does not disappear when a tool arrives — it relocates into the human–tool system. If AI will be permanently available, measuring what a person can do without it tests a world that no longer exists and spends scarce learning time on obsolete skill.

13.2 Where it is right

Often. Many operations should be delegated permanently, and a framework that treated unaided performance as inherently superior would mistake nostalgia for development. Where independent capability is not load-bearing for verification, resilience, responsibility, further learning or governance, the objection wins and delegation should be free.

13.3 Where it fails

- **Verification is asymmetric.** Generative AI is fallible and its users must check its outputs (Bastani et al., 2025). A person cannot reliably detect a plausible error in a domain they cannot perform in at all.
- **Responsibility is not delegable.** Civil-liberties analyses of AI-drafted police reports argue that such systems complicate accountability, memory and the ability to distinguish human from machine authorship (Stanley, 2024; Electronic Frontier Foundation, 2025). The principle is not that officers must type every word; it is that responsibility for a sworn factual account cannot be delegated merely because drafting can be automated.
- **Tools fail and are contested.** Systems go offline, conflict, hallucinate and meet tasks beyond their reliable frontier. Recognizing that boundary is a human capacity.
- **Development is cumulative.** Higher capabilities are built on lower ones. The adult who offloads an established skill is in a different position from the child who never builds it.
- **Self-regulation cannot be assumed.** Even strong, motivated learners over-request AI help when it is always available (Poulidis et al., 2025).

13.4 What would make the objection win

If controlled research showed that people formed under Gate-governed conditions verified, judged and adapted no better than people who used AI freely from the start — while performing no better independently — the objection would prevail for those capacities. That condition is registered in Section 19. The answer to the calculator objection is therefore conditional: automate what is no longer load-bearing; protect and periodically test what is required to govern the automation.

14. Joy, Agency and Mastery as Research Constructs

People do not persist in difficult development they experience as joyless. EQGENIX investigates curiosity, discovery, agency, mastery, creation and contribution as **candidate mechanisms requiring measurement**, not promised outcomes. The evidence warrants exactly that caution: an engineered AI tutor increased reported engagement and motivation (Kestin et al., 2025), and engagement appeared to drive gains from

adaptive sequencing (Chung et al., 2026), yet intrinsic motivation only partly protected chess learners from over-reliance (Poulidis et al., 2025). Joy may fuel persistence; it does not replace governed design.

Because self-report is weak evidence, these constructs are studied through behavior. Candidate indicators: the person chooses the harder Mission when an easier one is available, returns to challenging work unprompted, and persists beyond what the task requires. The research question is whether a formation environment can increase developmental engagement without making ease, entertainment or AI dependence the reward. The prize is not more AI. The prize is greater human range.

15. One Architecture Across Populations

Formation AI™ is not inherently a youth curriculum. Content, safeguards, constructs and evidence requirements change by population; the developmental logic remains.

Population	Formation emphasis	Boundary
Youth	Reasoning, verification, creativity, persistence, metacognition, authorship, judgment	Behind-the-glass generation; graduated modes; parent visibility
Public-safety professionals	Retrieval and regulation under pressure, communication, situational judgment, accountable decisions	AI never replaces accountable human judgment or authorship of sworn accounts
Veterans in transition	Adaptability, civilian problem-solving, identity translation, communication, agency	AI rehearses unfamiliar situations; the veteran owns the decisions
Adults and relationships	Listening, perspective-taking, regulation, reflection, repair	AI is never therapist, moral authority or relational adjudicator

The governing question follows every application: what human capacity must remain available for this person to govern the assistance responsibly?

16. Public Behavioral Indicators

The indicators below are observable, non-proprietary illustrations. They are not diagnostic criteria and do not disclose scoring weights, thresholds, algorithms, item banks or psychometric rules.

16.1 Built, Not Downloaded™ capabilities

Capability	Illustrative observable indicators
Critical Thinking	Identifies the claim and evidence; generates a plausible alternative; changes a conclusion when stronger evidence warrants it
Working Memory	Maintains multiple relevant elements during a task; follows a multistep operation without repeated rescue; notices when load exceeds reliable processing
Deep Reading	Tracks an extended argument; makes supported inferences beyond summary; distinguishes the author's claim from supporting detail
Metacognition	States confidence before feedback; identifies what is not known; adjusts strategy after detecting an error
Creative Synthesis	Generates an original direction before examples anchor the task; combines ideas across domains; explains why the final form fits its purpose
Original Writing Voice	Produces an independent first expression; explains wording and position choices; reconstructs the core message without AI phrasing
Frustration Tolerance	Remains engaged after failure; attempts a second strategy before rescue; distinguishes difficulty from impossibility
Quantitative Reasoning	Estimates before calculating; detects implausible magnitudes; explains what a number means in context
Spatial Reasoning	Builds and updates an internal representation of space; predicts spatial relationships; navigates without total dependence on external direction
Conversational Depth	Restates another position accurately; asks a responsive follow-up question; revises understanding after new information
Boredom Tolerance	Sustains an unfilled interval without immediate stimulation; initiates self-directed activity
Moral Reasoning	Distinguishes principle from preference; weighs competing duties; owns and explains a consequential judgment

16.2 Formation Capacities (delegation floor)

Formation Capacity	Illustrative observable indicators
Interruption Latency	Pauses before reaching for assistance; attempts an independent first move; names what is about to be handed off
Judgment Ownership	States a position before seeing the AI's; explains which parts of a product are theirs; accepts or rejects suggestions with reasons
Verification Capacity	Detects planted errors; builds a verification trail; lowers confidence when evidence is thin

Formation Capacity	Illustrative observable indicators
Effort Tolerance	Stays past the first impasse; requests the level of help actually needed rather than the answer
Position-Holding Under Relational Cost	Maintains a well-reasoned position under social pressure; revises for evidence, not disapproval
Accountability Retention	Discloses material AI assistance; owns consequences of delegated work; defends the work after the tool is removed

17. Limitations: What This Paper Does Not Claim

- Formation AI™ has not been experimentally validated as an EQGENIX intervention across the capacities or populations described. EQGENIX programs make no efficacy claims.
- The twelve Built, Not Downloaded™ capabilities and six Formation Capacities are a developmental framework, not validated psychometric constructs. Construct definition, reliability, validity, fairness and longitudinal interpretation require separate scientific work before any formal claim.
- Embedded behavioral indicators are not validated scores.
- Cognitive offloading is not inherently harmful, unaided performance is not always superior to human–AI collaboration, and not every operation must remain internal.
- Age 14 is a current EQGENIX safeguarding rule, not a universal scientific threshold.
- The studies cited are each limited to particular subjects, populations, countries and durations; two are working papers; and associations reported within them (for example, between agency or engagement and learning) do not establish causation.
- The author leads the organization that develops programs based on this framework (see Declarations).

18. Research Proposition

> Artificial intelligence can be intentionally structured as a human-capacity formation environment in which progressively sophisticated augmentation is conditioned on developmental appropriateness and demonstrated human capability, and in which formation is evaluated not by assisted performance

alone but by verification, independent performance, withdrawal retention, transfer and responsible governance of assistance — with parity treated as the floor and strengthened independent capability as the target.

- Does the Formation AI™ condition produce stronger independent performance, retention and transfer than credible comparisons, including well-safeguarded AI and no-AI instruction?
- Does Gate-governed progression outperform unrestricted or learner-regulated access?
- Does withdrawal performance predict durable ownership of capability?
- Do Missions generate valid evidence about the constructs they claim to observe?
- Does adaptive prescription outperform simpler fixed progression at equivalent exposure?
- Do behavioral indicators of curiosity, agency and mastery predict persistence?

19. Falsification

Formation AI™ should be revised or rejected under the following conditions:

- If it improves AI-assisted output but not the targeted human capability, it has produced performance rather than formation.
- If independent performance after withdrawal reaches only parity, the safety floor is met but formation is not demonstrated, and claims must say so.
- If performance collapses after withdrawal or delay where independent capability is material, it has produced dependence.
- If gains fail to transfer, claims must be narrowed to the trained context.
- If Gate-governed progression performs no better than free AI use on verification, judgment and adaptation, the calculator objection prevails for those capacities.
- If adaptive conditions do not outperform simpler training at equivalent exposure, their complexity is not justified.
- If embedded indicators do not support valid inference, the inference model must be redesigned.

- If greater augmentation reduces agency, verification, calibration or independent capability, the Gates and rails must change.
- If behavioral engagement indicators do not predict persistence, they must be removed as proposed mechanisms.

The purpose of the research program is not to prove EQGENIX right. It is to discover what actually forms stronger humans in an AI-rich world.

20. Conclusion: The Human Position

We reject two inadequate futures: one that tells people to fear artificial intelligence because machines are becoming extraordinarily capable, and one that tells people to surrender their cognition for the same reason.

Don’t panic. Prepare.
Don’t surrender. Develop.
Don’t compete with the machine. Govern it.

Teach children to think before asking, creators to originate before augmenting, professionals to verify before trusting, and leaders to retain responsibility after delegating. Then use AI boldly — to investigate, simulate, challenge, practice, discover, create and solve problems that once exceeded any one person’s reach. Human development and technological progress need not be rivals. Designed well, each can strengthen the other: to dress the garden, and to keep it.

BUILT, NOT DOWNLOADED™. BUILD THE HUMAN. GOVERN THE TECHNOLOGY. EXPAND WHAT’S POSSIBLE.

Declarations

Author contributions. Y. G. H.: conceptualization, framework design, writing — original draft, writing — review and editing.

Use of AI-assisted technologies. In preparing this paper the author used generative AI tools (Anthropic Claude and OpenAI ChatGPT) to assist with drafting, structural editing, and locating and checking sources. The author reviewed and edited all content, verified cited sources against the published works, and takes full responsibility for the content. Consistent with the thesis, the argument, architecture and final judgments are the author's own.

Competing interests. The author is the founder and chief executive of EQGenix Inc., which develops and sells education and formation programs that implement concepts described here, and which holds the associated trademarks.

Funding. This work received no external funding.

Data availability. No new data were generated or analyzed.

Ethics. This paper reports no research involving human participants.

Trademarks. Built, Not Downloaded™, Formation AI™, Adaptive Formation™, Embedded Formation™, Cognitive Sovereignty™ and Grit to Genius™ are trademarks of EQGenix Inc.

Scripture. Quotations are from the King James Version.

References

Bastani, H., Bastani, O., Sungu, A., Ge, H., Kabakcı, Ö., & Mariman, R. (2025). Generative AI without guardrails can harm learning: Evidence from high school mathematics. *Proceedings of the National Academy of Sciences, 122*(26), e2422633122. https://doi.org/10.1073/pnas.2422633122

Barcaui, A. (2025). ChatGPT as a cognitive crutch: Evidence from a randomized controlled trial on knowledge retention. *Social Sciences & Humanities Open, 12*, 102287. https://doi.org/10.1016/j.ssaho.2025.102287

Brod, G., Holstein, E., Weindorf, L., Colantonio, J., Bonawitz, E., & Theobald, M. (2025). Do it yourself: Discerning the effects of self-directed activity on conceptual learning. *npj Science of Learning, 10*, 70. https://doi.org/10.1038/s41539-025-00364-9

Chung, A. T.-H., Zhang, B., Kung, L.-C., Bastani, H., & Bastani, O. (2026). *Effective personalized AI tutors via LLM-guided reinforcement learning* [Working paper]. The Wharton School Research Paper, SSRN. https://doi.org/10.2139/ssrn.6423358

Electronic Frontier Foundation. (2025, July 10). *Axon's Draft One is designed to defy transparency*. https://www.eff.org/deeplinks/2025/07/axons-draft-one-designed-defy-transparency

Grinschgl, S., Papenmeier, F., & Meyerhoff, H. S. (2021). Consequences of cognitive offloading: Boosting performance but diminishing memory. *Quarterly Journal of Experimental Psychology, 74*(9), 1477–1496. https://doi.org/10.1177/17470218211008060

Kestin, G., Miller, K., Klales, A., Milbourne, T., & Ponti, G. (2025). AI tutoring outperforms in-class active learning: An RCT introducing a novel research-based design in an authentic educational setting. *Scientific Reports, 15*, 17458. https://doi.org/10.1038/s41598-025-97652-6

McDermott, K. B. (2021). Practicing retrieval facilitates learning. *Annual Review of Psychology, 72*, 609–633. https://doi.org/10.1146/annurev-psych-010419-051019

Poulidis, S., Bastani, H., & Bastani, O. (2025). *Self-regulated AI use hinders long-term learning* [Working paper]. The Wharton School Research Paper, SSRN. https://doi.org/10.2139/ssrn.5604932

Risko, E. F., & Gilbert, S. J. (2016). Cognitive offloading. *Trends in Cognitive Sciences, 20*(9), 676–688. https://doi.org/10.1016/j.tics.2016.07.002

Roediger, H. L., & Butler, A. C. (2011). The critical role of retrieval practice in long-term retention. *Trends in Cognitive Sciences, 15*(1), 20–27. https://doi.org/10.1016/j.tics.2010.09.003

Shute, V. J., & Ventura, M. (2013). *Stealth assessment: Measuring and supporting learning in video games*. MIT Press.

Stanley, J. (2024, December 10). *AI generated police reports raise concerns around transparency, bias*. American Civil Liberties Union. https://www.aclu.org/news/privacy-technology/ai-generated-police-reports-raise-concerns-around-transparency-bias

Yan, L., Greiff, S., Lodge, J. M., & Gašević, D. (2025). Distinguishing performance gains from learning when using generative AI. *Nature Reviews Psychology, 4*, 435–436. https://doi.org/10.1038/s44159-025-00467-5

The Holy Bible, King James Version.

Appendix: Terms

Term	Definition
Formation AI™	AI configured to increase human capability rather than only complete tasks
Premature substitution	Assistance entering during formation of the capacity the assistance performs

Term	Definition
Formation Engine	FORM → PRESS → DEMONSTRATE → AUGMENT → VERIFY → WITHDRAW → TRANSFER → RE-CALIBRATE → ADVANCE
AI Gates 0–4	Progressive AI relationships unlocked by developmental appropriateness and demonstrated capability
Assistance limit	The maximum AI role a specific task permits
Stricter-wins rule	When a person's Gate and a task's assistance limit differ, the stricter governs
Expansion without erosion	AI-augmented reach grows while independent capability holds and strengthens
Parity floor	Independent performance equal to comparison after withdrawal: necessary, not sufficient
Hold the Pen	The practice of keeping authorship and accountable judgment human: Think First, Bring It In, Check It, Pen Back, Take It Somewhere New
Embedded Formation™	Formation and evidence inside meaningful experiences; embedded for the participant, visible to those responsible
Generation behind the glass	AI prepares reviewed material for learners under 14, who do not interact with live generative AI

