



## TABLE OF CONTENTS

The AI Mentor: Cognitive Sovereignty in the Age of Automation.....	1
1. Introduction: Extinction vs. Exoskeleton- Two futures are already emerging .....	1
The Crisis of Competence and Scale .....	2
2. What Is an AI Mentor? From Cognitive Digital Twin to Wisdom Engine .....	3
Identity, Continuity, and the Cognitive Digital Twin.....	3
Situating, Multi-Modal Context Across All Skills.....	3
Learning Design Intent: An Anti-Crutch by Design .....	3
Lifelong, Cross-Domain Mentoring.....	3
Ethical and Cognitive Guardrails .....	4
3. Scientific Foundations: Why AI Mentors Can Work .....	4
Mastery Learning and One-on-One Tutoring .....	4
Intelligent Tutoring Systems and Cognitive Tutors.....	4
Deliberate Practice and High-Stakes Simulation.....	4
Active Engagement: ICAP, Testing Effect, and Desirable Difficulties .....	5
Extended Mind, Cognitive Offloading, and Techno-Dependence .....	5
4. Architecture and Design Principles of an Anti-Crutch AI Mentor.....	5
Think-First Architecture .....	5
Progressive Scaffolding and Fading.....	6
Metacognitive Coaching .....	6
Explicit Management of Cognitive Offloading.....	6
Multi-Context Transfer Engine.....	6
Human-in-Command and Automation-Bias Defenses .....	7
5. The Engine Room: Why AI Mentors Are Now Technically Feasible .....	7
Massive Context and Longitudinal Coherence .....	7
Reasoning Models and Metacognitive Mirroring .....	7
6. Current State of AI Mentor Development.....	8
AI Mentor 1.0: AR Firefighter Training for AFCEC.....	8
AI Mentor 2.0: Architecture for Accelerated Mastery .....	8
Expected Human Outcomes .....	9
Outcome 2: Democratized access to expertise (no gatekeepers). .....	10
Outcome 3: Cognitive integrity and retention despite automation. ....	10
Outcome 4: Alignment with AI governance and risk management frameworks. ....	10
7. Beyond Training and Education: AI Mentors Across All Skills and a Human Lifetime .....	10
Workforce and On-the-Job Learning.....	10
Vocational Skills and Trades .....	11
Health, Wellbeing, and Self-Care .....	11
Creativity and Innovation .....	11

8. Risk Landscape: “Dumbing Down” Is the Default Without Intervention.....	11
9. Evaluation of AI Mentors .....	11
10. A Cognitive Bill of Rights for the Age of AI Mentors .....	12
11. Roadmap and Standards: From EdTech to Cognitive Infrastructure .....	12
12. Conclusion: AI Mentors as a Civilizational Bet on Wisdom .....	13
List of Abbreviations.....	14
References .....	15

## THE AI MENTOR: COGNITIVE SOVEREIGNTY IN THE AGE OF AUTOMATION

### From Answer Engines to Lifelong Skill Builders

**Author:** Manuel Miranda, Chief AI and Growth Officer APV

**Abstract-** AI is rapidly becoming an “answer engine” for everything from email to engineering, and the path of least resistance leads to a future of cognitive atrophy; over-trust in automation, shallow understanding, and a workforce that fails when the tools go dark. This white paper proposes a different trajectory, the AI Mentor a long-lived, context-aware system designed as a cognitive exoskeleton, not a replacement brain. Drawing on decades of research in mastery learning, one-on-one tutoring, intelligent tutoring systems, deliberate practice, and the extended-mind thesis, we define design principles for an anti-crutch AI that makes people think before it helps. The AI Mentor maintains a longitudinal model of the learner’s skills, misconceptions, and decision habits; lives inside real work environments (e.g., XR, simulators, IDEs, operational tools); and uses prediction, explanation, and graded challenge to build durable competence and metacognition over time. Architecturally, we specify how such a system can be assembled from current multimodal models, agentic workflows, and learning analytics, while aligning with emerging governance frameworks and defenses against automation bias. The paper is written for leaders in government, industry, and education who must deploy AI at scale without hollowing out human capability. Its thesis is simple if we do not deliberately design AI to train judgment, we will accidentally train dependence. The AI Mentor is a blueprint for avoiding that outcome.

### 1. INTRODUCTION:

#### EXTINCTION VS. EXOSKELETON- TWO FUTURES ARE ALREADY EMERGING

**Future A – Cognitive Atrophy:** AI tools answer everything instantly. People default to autocomplete for writing, coding, planning, and even thinking about ethics or policy. Empirical evidence shows that when we expect information to be externally available, we remember less of it ourselves and instead remember “where to find it” [11, 12]. In navigation, heavy GPS use is associated with poorer spatial memory and navigation ability [10]. Even in high-stakes domains humans to over-trust automated suggestions even when they are wrong. This over trust is known as **automation bias**, the tendency to trust an automated system’s recommendation over your own judgment even when its wrong [8, 9].

Extend these trajectories across all cognitive tasks reasoning, writing, diagnosing, planning; and Future A becomes a world of superficially competent operators whose internal skills have quietly atrophied.

**Future B – Cognitive Exoskeleton:** AI is reimagined as a Cognitive Exoskeleton, a structure that supports and strengthens human cognition instead of replacing it. In this future, AI tools:

- Withhold answers until we have attempted solutions.
- Force retrieval and explanation before providing feedback [5, 13, 14].
- Track our decision patterns across years, identifying persistent misconceptions and blind spots.
- Train us, by design, to resist automation bias and remain intellectually “in command” even in highly automated environments [8, 9].

At the center of Future B is the **AI Mentor**. This AI system is not an answer engine, but a **Wisdom Engine**, a system whose primary output is improved human judgment over time. It applies to all skills cognitive, physical, creative, social, technical. If something can be learned, practiced, evaluated, and improved, an AI Mentor can support it.

The purpose of this paper is to define, justify, and operationalize the AI Mentor as an engine for cognitive sovereignty. It is written for three stakeholder groups:

- The **General Public**, who will live with AI embedded in daily life and decision-making;
- **Academia**, which must integrate AI into teaching, learning, and assessment without hollowing out human skills; and
- **Government**, which has the opportunity to deploy AI Mentor at scale to address speed-to-competency gap in critical skills.

## The Crisis of Competence and Scale

---

Work is changing faster than we can prepare people for it. AI has accelerated this shift and is expected to disrupt jobs across most sectors of the economy [25]. Most institutions still rely on a familiar learning approach: a few days of training, a stack of manuals, and a hope that the employee will “learn it on the job.” In many domains this is no longer viable. The half-life of skills is shrinking, while the complexity and interdependence of systems are increasing [1, 2, 18].

This creates a speed-to-competency gap, the distance between when a worker is hired and when they can safely perform complex tasks without constant supervision. For the U.S. Government, which faces critical shortages in high-stakes roles, closing this gap is a matter of national operational continuity:

**Government inspectors.** Regulatory agencies cannot afford to wait years for new inspectors to internalize complex rules and tacit judgment through slow exposure and informal mentorship. They need staff who can reach independent, field-ready competence in months, not years, and who can keep up as regulations and technologies evolve.

**Air traffic controllers and other safety-critical roles.** Training pipelines are constrained not only by classroom seats but by the availability of experienced instructors and by the rarity of the most important events. Many of the scenarios that define true expertise such as runway incursions, cascading system failures, multi-agency coordination breakdowns are precisely the ones that do not occur frequently enough in the real world to train against [21].

On top of this, there is a gatekeeper problem. Deep, apprenticeship-style mentorship is available only to a small minority of learners. Whether someone receives sustained, high-quality guidance often depends on luck: being assigned to the right supervisor, entering at the right time, or working at a well-resourced institution. Yet we know from decades of research that one-on-one tutoring and mastery learning can produce performance gains on the order of two standard deviations, the famous “2 sigma” effect [1, 2, 3].

The core claim of this paper is that AI Mentors are a way out of this bind. If we architect them correctly, as lifelong, anti-crutch systems rather than answer engines, they can simultaneously:

- Deliver something close to high-quality, one-on-one mentorship at scale.
- Dramatically shorten time-to-competence in complex roles.
- Preserve, and in some cases enhance, human cognitive sovereignty in the face of pervasive automation [6–9, 17, 10-12, 18].

The rest of the paper elaborates what an AI Mentor is, why it is scientifically plausible, how it should be architected, and how it can be governed so that it strengthens rather than hollows out human capability.

## 2. WHAT IS AN AI MENTOR? FROM COGNITIVE DIGITAL TWIN TO WISDOM ENGINE

We define an **AI Mentor** as: *A long-lived, context-aware, multi-modal AI system that engages a learner in cycles of action, feedback, reflection, and adaptation, deliberately structured to strengthen human understanding, judgment, and creativity across all domains of skill over time.*

This is deliberately more than “an LLM with fine-tuned prompts.” Conceptually, an AI Mentor has five defining properties.

### Identity, Continuity, and the Cognitive Digital Twin

The AI Mentor maintains a **longitudinal model** of the learner’s:

- Skills and knowledge
- Misconceptions and common error patterns
- Decision strategies and habits
- Goals, values, and evolving roles

Over years, this becomes a **Cognitive Digital Twin**: a model of how the person tends to think, decide, and learn in different contexts. It does not simply track quiz scores; it tracks the *shape* of the person’s cognition. This enables feedback such as: “You’re making the same risk assessment error in this supply chain negotiation that you made in a disaster-response simulation 9 months ago. Let’s walk through what changed and what didn’t.” This enables highly personalized, context-aware feedback, much like what a human mentor can offer after working with the same student for years.

### Situated, Multi-Modal Context Across All Skills

The AI Mentor operates inside real contexts, not just a chat window:

- XR training simulators for pilots, firefighters, surgeons, and heavy-equipment operators [21]
- Software IDEs for programmers and data scientists
- CAD tools and digital twins for engineers and manufacturing technicians
- Learning management systems and productivity suites for students, analysts, and policymakers
- Everyday tools such as browsers, office suites, communication platforms that support civic reasoning, media literacy, and personal finance

Because it is multi-modal, AI Mentor can analyze text, code, diagrams, and video to understand not just *what* the user did, but *how* they responded under stress, uncertainty, or time pressure.

### Learning Design Intent: An Anti-Crutch by Design

Unlike typical AI tools that optimize for speed and satisfaction, the AI Mentor optimizes for:

- Long-term retention and transfer [5, 13, 14]
- Metacognitive skill planning, monitoring, reflection
- Cognitive independence performing without assistance when needed [11, 12]

AI Mentor deliberately uses principles from mastery learning, deliberate practice, Interactive Constructive Active Passive framework (ICAP), desirable difficulties, and cognitive apprenticeship. Default interaction with the learner is for the AI Mentor to **Ask** before it tells, **Challenge** before it explains, and **Fade** support as competence increases [1–4, 13, 14, 16, 17].

### Lifelong, Cross-Domain Mentoring

The AI Mentor is designed to follow the learner from:

- Foundational schooling and vocational training
- Through early-career technical roles, supervisory and leadership responsibilities
- Through mid-career reskilling and late-career domain transfer
- Into lifelong personal learning

This is essential for a world in which entire domains of skill will be disrupted multiple times within a single career [18]. A person may need to become a novice again several times. The AI Mentor is the continuity layer, preserving cognitive gains and enabling rapid retooling across domains.

### Ethical and Cognitive Guardrails

An AI Mentor explicitly counters **Automation bias**, over-trust in AI recommendations, and **Cognitive offloading overuse**, excessive reliance on external memory and problem-solving [8, 9, 10-12]. The AI Mentor surfaces uncertainty, demands human justification for critical decisions, and periodically withholds support to ensure core competencies remain intact and extends into **cognitive safety**, the protection of human cognitive capacities over time.

## 3. SCIENTIFIC FOUNDATIONS: WHY AI MENTORS CAN WORK

### Mastery Learning and One-on-One Tutoring

Benjamin Bloom showed that students receiving one-on-one tutoring combined with mastery learning outperformed traditional classroom students by roughly two standard deviations, the famous “2 sigma” effect [1]. Carroll’s model of school learning argued that achievement is largely a function of time spent on effective tasks relative to time needed [2]. While the classic Carroll and Bloom studies were conducted in school contexts, the underlying time-on-task and mastery principles have been replicated in adult and professional training, and there is now growing evidence that AI-driven tutoring systems can support both K-12 and post-secondary learners through adaptive feedback and mastery-oriented design [26, 27]. Historically, the barrier has been cost and scalability; we cannot assign every learner a full-time human tutor or master craftsman. AI Mentors are among the first plausible ways to approach “tutoring + mastery” at population scale without diluting quality.

### Intelligent Tutoring Systems and Cognitive Tutors

Decades of work in Intelligent Tutoring Systems (ITS) and cognitive tutors provide a strong baseline [3, 4, 16]. Well-designed ITS often approach the effectiveness of human tutors and significantly outperform conventional classroom instruction [3]. Cognitive Tutor systems in algebra have improved achievement at scale [4, 16].

AI Mentors can be thought of as **ITS NextGen**:

- Broader in scope (all skills, not just academic subjects)
- Richer in context (real tools, VR/XR, operational data streams)
- Explicitly aligned with cognitive and ethical goals (anti-crutch, human-in-command)

### Deliberate Practice and High-Stakes Simulation

Expert performance research emphasizes **deliberate practice**: effortful, feedback-rich practice on well-defined tasks with clear performance criteria. Defense, aviation, healthcare, and emergency services already show that VR/XR plus simulation plus data-driven feedback can accelerate skill acquisition while reducing risk and cost [21].

AI Mentors sit on top of such environments, turning them from “fancy simulators” into longitudinal coaching systems that know each learner over time. The same pattern applies to:

- Skilled trades (e.g., welding, machining, construction)
- Creative domains (e.g., design, music, writing)
- Knowledge work (e.g., legal reasoning, policy analysis, scientific research)

### Active Engagement: ICAP, Testing Effect, and Desirable Difficulties

The **ICAP framework** ranks learning activities from Passive to Interactive, with learning gains increasing along that continuum. Research shows that when people must work a bit for the answer by recalling, predicting, and explaining, they learn and remember far more over time than when they just re-read or passively review information [13, 14].

The AI Mentor operationalizes this by:

- Requiring predictions before revealing outcomes
- Asking for explanations and rationales before offering corrections
- Tuning difficulty to the “sweet spot” where challenge is effortful but not overwhelming

### Extended Mind, Cognitive Offloading, and Techno-Dependence

Clark and Chalmers’ “extended mind” thesis argues that tools and environments can become part of a cognitive system when tightly coupled to our thought processes [15]. Modern AI systems, which can generate text, code, plans, and designs, are extreme extensions of this idea.

But research on cognitive offloading and the “Google effect” shows the downside. People who expect information to be available externally remember less of it themselves [11, 12]. Heavy GPS users show poorer spatial memory and navigation ability [10].

Indiscriminate use of AI as an answer machine is poised to produce the same effect on reasoning, writing, numeracy, and judgment. AI Mentors must therefore be designed as extended minds that strengthen the internal mind rather than supplant it.

Mastery learning and tutoring show that flexible time, targeted feedback, and clear standards dramatically boost achievement, and ITS/cognitive tutors prove that software can deliver much of this at scale. The AI Mentor must demand effort, coach in rich simulations, and act as a cognitive exoskeleton that strengthens rather than replaces the human mind.

## 4. ARCHITECTURE AND DESIGN PRINCIPLES OF AN ANTI-CRUTCH AI MENTOR

While the AI Mentor concept can be applied from K–12 through college, the architecture and design principles that follow focuses on adult learners only.

### Think-First Architecture

**Principle:** The mentor should almost always require a human attempt before providing an explanation or solution.

**Patterns:**

- Prediction prompts: “Before I show you the options, predict the most likely failure mode and why.”
- Self-explanation prompts: “Walk me through your reasoning step by step.”
- Multiple representations: Ask for diagrams, summaries, numerical estimates, or causal maps before feedback.

This directly implements retrieval practice and constructive engagement [5, 13, 14].

---

## Progressive Scaffolding and Fading

---

**Principle:** Support decreases as competence increases. In early stages, the AI Mentor can provide hints, examples, and guided steps. As skill improves, the AI Mentor:

- Hides steps unless requested,
- Switches to Socratic questioning, counterexamples, and scenario variation,
- Encourages learners to design their own practice plans.

This mirrors experienced human mentors and prevents long-term dependency [3, 4, 16, 17].

---

## Metacognitive Coaching

---

**Principle:** The mentor trains not only domain skills but also the learner’s ability to manage their own cognition.

**Examples:**

- Planning: “Outline your approach and the checks you will use to detect errors.”
- Monitoring: “You’ve changed your diagnosis twice, what new evidence motivated the shift?”
- Reflection: “What surprised you in this scenario? What will you do differently next time?”

Metacognition is key to transfer and lifelong learning, yet it is rarely trained systematically at scale because most education and training systems prioritize content coverage and task performance over explicit practice in planning, monitoring, and evaluating one’s own thinking.

---

## Explicit Management of Cognitive Offloading

---

**Principle:** Offloading should be strategic, not default.

The mentor differentiates between:

- Safe offloading: details that can live outside the brain (e.g., long tables of constants, rarely used codes).
- Core-skill offloading: reasoning steps, structured writing, basic numeracy, diagnostic frameworks.

**The mentor:**

- Logs and visualizes how often the learner asks for full solutions vs hints.
- Periodically runs “naked mind” drills: tasks where AI support is deliberately disabled to verify independent competence.
- Adjusts its behavior if it detects habitual over-reliance, nudging the learner back toward internal problem-solving [11, 12].

---

## Multi-Context Transfer Engine

---

**Principle:** The mentor’s real value is helping people transfer skills across domains and situations.

The AI Mentor tracks conceptual knowledge (e.g., causal reasoning, risk assessment, feedback loops) and deliberately reuses it in new contexts:

“The communication failure in this multi-agency emergency response looks structurally similar to the coordination breakdown in your software deployment incident last month. Let’s compare the two.”

Research suggests that AI-supported learning gains can be larger when learners are nudged to connect tasks and apply concepts in new settings [5, 24].

---

## Human-in-Command and Automation-Bias Defenses

---

**Principle:** Human judgment is explicitly privileged at key decision points.

Given strong evidence of automation bias:

- The mentor always exposes uncertainty and alternative options [8, 9]
- For critical decisions, it requires the human to articulate a rationale that can stand alone
- It periodically asks the learner to **override its own suggestions** and justify why, thereby training independent judgment

This is essential in safety-critical domains, aviation, healthcare, energy, defense, but equally important in domains like policymaking, finance, and civic life.

---

## 5. THE ENGINE ROOM: WHY AI MENTORS ARE NOW TECHNICALLY FEASIBLE

---

### Massive Context and Longitudinal Coherence

---

Early chatbots forgot you when the browser tab closed. Modern architectures with **context windows in the millions of tokens** enable the AI Mentor to ingest a learner's:

- Past simulation logs,
- Key documents and projects,
- Code repositories or designs,
- Longitudinal assessment histories.

This supports feedback like: “Your negotiation strategy here repeats a pattern we saw in your second-year management course: you over-weight short-term operational risk and under-weight reputational risk. Let’s re-examine the trade-offs.”

### Reasoning Models and Metacognitive Mirroring

---

New **reasoning models** (e.g., OpenAI 5.2, Google Gemini 3, Anthropic Claude 4.5, and multimodal, agentic workflows, “thinking” variants of frontier LLMs) generate hidden chains of thought before responding [22, 23].

An AI Mentor can use these capabilities for **Metacognitive Mirroring**:

- Silently simulate the user’s likely reasoning path.
- Distinguish careless slips from deep misconceptions.
- Provide feedback not just on *what* was wrong, but *how the thinking went off track*.

Agentic Workflows and the Total Learning Architecture (TLA)

In architectures such as the U.S. DoW’s **Total Learning Architecture (TLA)**, AI Mentors become agents that:

- Observe learner actions across tools and platforms (LMS, VR simulators, IDEs, operational dashboards).
- Write to and read from a **Master Personal Learning Record** that persists over time [20].
- Coordinate with other agents (e.g., a scheduling agent, a content-retrieval agent, a compliance agent) to deliver individualized experiences at scale.

This agentic ecosystem is essential to extend AI Mentors beyond niche pilots and into **national training, education, and workforce systems**.

## 6. CURRENT STATE OF AI MENTOR DEVELOPMENT

The AI Mentor is not just a thought experiment. Early generations have already been developed and deployed in demanding environments, providing evidence that the architecture can work with real world operational constraints. This section summarizes the current state of development, focusing on AI Mentor 1.0 and the emerging AI Mentor 2.0 architecture.

### AI Mentor 1.0: AR Firefighter Training for AFCEC

AI Mentor 1.0 was developed and deployed for the U.S. Air Force Civil Engineer Center (AFCEC) under the leadership of Program Manager Alton Robinson. Implemented by APV (Always Provide Value), the effort produced the **AR Incident Commander / AI Mentor Game**, a production system used to modernize DoW firefighter training at scale.

The project addressed a classic speed-to-competence and gatekeeper problem: incident commanders must be prepared for rare, complex, multi-stage emergencies, but live burns and traditional exercises are expensive, logistically constrained, and limited in the variety of scenarios they can realistically present [21].

Key characteristics of AI Mentor 1.0 include:

**Augmented Reality (AR) immersion.** Firefighters engage in high-risk scenario simulations through AR experiences delivered on mobile devices. This allows units around the world to drill complex incidents, multi-structure fires, hazardous materials events, multi-agency responses, without the cost and safety constraints of repeated live burns.

**AI driven mentorship during and after scenarios.** An AI mentor guides learners through scenarios by asking them to justify decisions, branching the situation according to those decisions, and delivering targeted feedback afterward. The mentor is not just a hint provider; it embodies the anti-crutch principles described in Sections 2 and 4 by prompting explanation, reflection, and repeated practice rather than simply supplying answers.

**Data rich learner modeling.** The system captures decision paths, timing, missteps, and corrections. Over time this supports a longitudinal model of each firefighter's strengths and weaknesses, an operational example of the Cognitive Digital Twin concept introduced in Section 2.1.

Operationally, AFCEC reports measurable improvements in readiness, decision-making quality, and situational awareness across a globally distributed firefighter workforce. The program has been recognized as a best-in-class innovation, including the *ACT-IAC Innovation Champion Award* and the *FORUM Innovation Award*. While the results are early and context-specific, they demonstrate that AI Mentor architectures are feasible at scale in high-risk, mission-critical training environments.

### AI Mentor 2.0: Architecture for Accelerated Mastery

AI Mentor 2.0 generalizes the lessons of AI Mentor 1.0 to a broader set of roles and institutions. It is explicitly engineered to address the speed-to-competency and gatekeeper problems described in Section 1.1, while maintaining the cognitive sovereignty and anti-crutch design principles developed throughout this paper.

Concretely, AI Mentor 2.0 combines:

**Structured, stage-based acceleration grounded in mastery learning.** For a new inspector, air traffic controller, or safety engineer, the AI Mentor orchestrates a progression similar to mastery learning models [1, 2]: ○ *Orientation and demonstration*. The mentor introduces core concepts and

demonstrates them on realistic, high-fidelity problems such as a facility inspection, a runway incursion scenario, or a multi-agency emergency response [21].

- *Guided practice.* Learners attempt tasks with “training wheels” support: hints, Socratic questions, and partial scaffolds that encourage active reasoning rather than passive copying.
- *Independent practice and evaluation.* Support is deliberately faded. The system runs “naked mind” drills where AI assistance is disabled for critical steps, verifying that the learner can perform under conditions of partial or total automation loss [11, 12].

**Anti-crutch behavior as a first-class requirement.** Unlike standard AI tools that optimize for convenience and speed, AI Mentor 2.0 optimizes for retention, transfer, and independent performance [5, 13, 14]. In practice this means: ○ Requiring learners to attempt solutions and explain their reasoning before revealing model-generated answers (“Predict what will happen if you clear this aircraft now and justify your reasoning”).

- Detecting patterns of over-reliance on automation and periodically withholding suggestions or explanations to prevent core skills from being offloaded [8, 9, 11, 12].
- Logging offloading patterns over time and feeding them back to the learner as part of explicit metacognitive coaching (e.g., “You are relying on me for initial risk assessments in 80% of cases in this scenario type”).

**“Mentor Mode” for on-the-job guidance.** The most important phase begins once formal training ends and the learner enters real operations. In Mentor Mode, AI Mentor 2.0 becomes a continuous companion: ○ **For an inspector,** it shadows facility visits, quietly checking actions against standard operating procedures, surfacing overlooked checks, and flagging likely misclassifications before they appear in official reports.

- **For an air traffic controller,** it observes performance during simulations and supervised shifts, identifying systematic risk-assessment errors or coordination failures that a human supervisor might miss under heavy load.
- In both cases, the mentor writes detailed events into the learner’s Cognitive Digital Twin, enabling focused debriefs and targeted drills that address recurring weaknesses [3, 4, 16, 17].

Technically, AI Mentor 2.0 is designed to integrate with architectures such as the Total Learning Architecture (TLA) [20], coordinating with other agents (for scheduling, content retrieval, compliance logging) and writing to a persistent learning record. This allows the same mentor that coaches a learner through an AR simulation to later comment on their real reports, incident postmortems, and performance reviews.

## Expected Human Outcomes

---

To justify AI Mentors as infrastructure rather than a niche “ed-tech tool,” deployments must be evaluated against concrete human outcomes. Building on both the AFCEC experience and the learning-science foundations summarized in Sections 3 and 4, AI Mentor 2.0 is designed to achieve four primary outcomes.

### **Outcome 1: Radical acceleration of competency.**

**Benefit.** Dramatically reduced time-to-proficiency for complex roles, allowing agencies to fill empty seats with highly capable decision-makers, not just novices.

**Application.** New inspectors, air traffic controllers, and incident commanders can accumulate the equivalent of years of rare, high-stakes scenario experience in months. This provides an immediate lift in service quality and performance; a new hire supported by an AI Mentor can navigate complex

regulations or crises with the sophistication of a senior veteran, solving the "experience gap" that currently plagues understaffed departments.

### Outcome 2: Democratized access to expertise (no gatekeepers).

**Benefit.** High-quality mentorship becomes a standard feature of work and education, not a privilege limited to those with the “right” supervisor.

**Application.** Every learner gets an AI Mentor tuned to organizational best practices and governance frameworks, available 24/7. Human mentors and instructors are freed to focus on what they do uniquely well: dealing with ambiguity, ethics, organizational context, and tacit knowledge, while the AI Mentor handles structured practice, feedback, and longitudinal tracking.

### Outcome 3: Cognitive integrity and retention despite automation.

**Benefit.** Workers maintain critical thinking, situational awareness, and independent performance even as AI systems become more capable.

**Application.** Evidence from navigation and search shows that heavy reliance on GPS and ubiquitous information access can degrade spatial memory and shift memory from *facts* to *where to find facts*. AI Mentors counter this by systematically scheduling retrieval practice, running naked-mind drills, and monitoring offloading behavior. The goal is a workforce that functions when tools fail and treats AI as a partner to question not an oracle to obey.

### Outcome 4: Alignment with AI governance and risk management frameworks.

**Benefit.** AI Mentor deployments are auditable, governable, and aligned with emerging standards, enabling adoption in regulated settings.

**Application.** Because AI Mentors track when they coached versus when they answered, when humans overruled their recommendations, and when automation bias was detected, organizations can analyze these logs against frameworks such as the NIST AI Risk Management Framework and ISO/IEC 42001 [6,7]. This makes cognitive outcomes like retention, transfer, offloading patterns, trust calibration measurable governance targets alongside traditional metrics like performance, safety, and compliance (see Section 9).

Taken together, AI Mentor 1.0 and 2.0 show a credible path from theory to practice, from the conceptual architecture laid out in Sections 2 through 5 to operational systems that improve readiness today, and to governance-ready infrastructures that can support whole-of-government and whole-of-economy deployments tomorrow.

## 7. BEYOND TRAINING AND EDUCATION: AI MENTORS ACROSS ALL SKILLS AND A HUMAN LIFETIME

### Workforce and On-the-Job Learning

Most real learning happens **on the job**, not in classrooms [18]. An AI Mentor embedded in work tools can:

- Observe real artifacts (code reviews, incident postmortems, design documents, legal briefs).
- Identify recurring patterns of error or excellence.
- Propose **micro-drills** and reflective prompts targeted at those patterns.

Crucially, it does this **without collapsing into “auto-complete my work.”** It stays in mentoring mode; challenging, questioning, and structuring learning rather than silently performing tasks.

## Vocational Skills and Trades

For trades, electricians, welders, machinists, HVAC technicians, the AI Mentor can:

- Pair with AR headsets and smart tools to observe physical performance
- Provide real-time corrective prompts (e.g., torque application, sequence order)
- Schedule spaced simulations of rarely used but critical procedures

Skill is skill, whether algebraic manipulation or precise hand-eye coordination. AI Mentors are **domain-agnostic** as long as they can see actions and consequences.

## Health, Wellbeing, and Self-Care

In health professions education, AI tutors and agents already assist with teaching protocols and decision-making. For individuals, an AI Mentor can:

- Translate complex health information into graded learning experiences
- Train people to ask better questions of clinicians
- Support adherence and self-management via spaced retrieval and scenario practice

## Creativity and Innovation

There is legitimate concern that AI-generated content could make humans less creative. Research on cognitive offloading indicates that external tools can either **free up working memory** for higher-level thinking or **replace thinking entirely** depending on how they are used.

An AI Mentor for scientists, engineers, or artists might:

- Help structure experiments and critique hypotheses
- Encourage exploration of more radical design spaces while demanding explicit articulation of principles and trade-offs
- Challenge the user to revise, extend, and defend their own ideas rather than simply accept AI-generated ones

## 8. RISK LANDSCAPE: “DUMBING DOWN” IS THE DEFAULT WITHOUT INTERVENTION

Evidence already shows that “AI as crutch” is not speculative:

- Internet search leads people to remember where to find information rather than the information itself [11]
- Heavy GPS use correlates with weaker spatial memory [10]
- AI tool usage can correlate with lower critical thinking when people offload too much cognitive effort [12]
- Automation bias and complacency may degrade decision quality in domains such as clinical diagnosis and radiology [8, 9]

If AI systems in education, work, and everyday life optimize for Satisfaction scores, Completion time, and Engagement metrics, they will almost certainly accelerate cognitive atrophy.

An AI Mentor must be engineered and governed to optimize for learning outcomes, cognitive resilience, and human-in-command decision making, even when that means making experiences harder and slower in the short term.

## 9. EVALUATION OF AI MENTORS

Traditional evaluation models (e.g., Kirkpatrick’s four levels: Reaction, Learning, Behavior, Results) are necessary but not sufficient for AI Mentors [19]. Consider these three areas:

**Retention and Transfer, metrics should include:**

- Long-term retention of core knowledge and procedures (weeks and months, not hours) [13, 14].
- Ability to apply skills to **novel** scenarios, not just rehearsed ones.

**Metacognitive Skill, we should measure improvements in:**

- Planning and goal setting
- Monitoring and error-detection
- Self-evaluation and adaptive strategy shifts

**Cognitive Independence and Offloading Patterns, key indicators include**

- How often learners request full solutions vs hints
- Performance on “**no-AI challenge**” tasks after training
- Changes in offloading patterns over time [11, 12]

## 10. A COGNITIVE BILL OF RIGHTS FOR THE AGE OF AI MENTORS

To protect human cognitive sovereignty, we propose a **Cognitive Bill of Rights** for human–AI interaction in mentoring contexts:

1. **The Right to Struggle** AI systems shall not solve a problem before the human has made a credible attempt, except in safety-critical contexts where immediate intervention is required.
2. **The Right to Explanation** AI systems must explain *why* they prompted or intervened, not just *what* the answer is.
3. **The Right to Disconnect (“Naked Mind” Drills)** AI systems must support regular periods where assistance is disabled, and humans must be evaluated and trained to perform core tasks without AI.
4. **The Right to Provenance** Users have the right to know whether a skill or artifact (e.g., a piece of writing, a diagnostic decision) was primarily produced by human effort, AI assistance, or AI automation.
5. **The Right to Cognitive Integrity** AI systems must be evaluated and governed for long-term impacts on memory, reasoning, and judgment, not just short-term performance.
6. **The Right to Human Command** AI systems must be designed and built to always allow for humans to override it.

These principles should inform both technical design and policy frameworks for AI Mentors.

## 11. ROADMAP AND STANDARDS: FROM EDTECH TO COGNITIVE INFRASTRUCTURE

To move from pilots to practical real world deployment, we need:

**Richer learner models**, capturing not just right/wrong answers but misconceptions, strategies, and metacognitive behaviors over years.

**Cross-context event streams**, integrating VR/XR, LMS, operational tools, and HR systems into coherent learning records [21].

**Standardized governance schemas** for cognitive safety and automation bias, logging when AI gave answers vs coached, and when humans overruled or corrected AI.

**Value-sensitive design practices**, aligning system behavior with educational values like Long-term retention and transfer, Metacognitive skill (planning, monitoring, reflection), and Cognitive independence (performing without assistance when needed).

National and international standard-setting bodies should recognize **AI Mentors** as a distinct class of system, different from generic chatbots and recommendation engines, and develop **specific guidance** for their evaluation, deployment, and oversight.

## 12. CONCLUSION: AI MENTORS AS A CIVILIZATIONAL BET ON WISDOM

We are not deciding **whether** AI will reshape learning, work, and everyday reasoning. That decision has already been made by the deployment of powerful AI models into every browser, phone, and productivity tool.

**The real decision is:**

- Do we accept **answer engines** that optimize for convenience and quietly erode memory, judgment, and creativity?
- Or do we build **AI Mentors**, cognitive exoskeletons designed to make people better thinkers, creators, and decision-makers over decades?

The research base is clear Tutoring plus mastery learning can produce dramatic gains but historically doesn't scale [1–4]. Intelligent tutoring systems and simulation-based training work when grounded in learning science is effective [3, 16, 21]. Cognitive offloading and automation bias are powerful forces that can hollow out human capability if left unchecked [8-12].

AI Mentors reconcile these truths. They scale tutoring while preserving and extending human intellect. They transform training and education **across all skills** academic, technical, creative, civic into a continuous, data-informed mentoring relationship. They create **Wisdom Workers** capable of auditing, guiding and when necessary, overruling autonomous systems.

If we get this right, AI will not be the technology that made humanity intellectually obsolete. It will be the technology that gave us our only realistic shot at staying smart enough individually and collectively to handle the century we have built.

## LIST OF ABBREVIATIONS

**ACT-IAC** – American Council for Technology–Industry Advisory Council

**AFCEC** – Air Force Civil Engineer Center

**AI** – Artificial Intelligence

**AR** – Augmented Reality

**AETC** – Air Education and Training Command (U.S. Air Force)

**EdTech** – Educational Technology

**GPS** – Global Positioning System

**ICAP** – Interactive–Constructive–Active–Passive framework

**IDE** – Integrated Development Environment

**ISO/IEC** – International Organization for Standardization / International Electrotechnical Commission

**ITS** – Intelligent Tutoring System

**LMS** – Learning Management System

**LLM** – Large Language Model

**PTN** – Pilot Training Next (U.S. Air Force)

**RMF** – Risk Management Framework

**TLA** – Total Learning Architecture

**VR** – Virtual Reality

**XR** – Extended Reality (VR/AR and mixed reality)

## REFERENCES

- [1] Bloom, B. S. (1984). The 2 sigma problem: The search for methods of group instruction as effective as one-to-one tutoring. *Educational Researcher*, 13(6), 4–16. <https://doi.org/10.3102/0013189X013006004>
- [2] Carroll, J. B. (1963). A model of school learning. *Teachers College Record*, 64(8), 723–733. <https://doi.org/10.1177/016146816306400801>
- [3] VanLehn, K. (2011). The relative effectiveness of human tutoring, intelligent tutoring systems, and other tutoring systems. *Educational Psychologist*, 46(4), 197–221. <https://doi.org/10.1080/00461520.2011.611369>
- [4] Koedinger, K. R., Anderson, J. R., Hadley, W. H., & Mark, M. A. (1997). Intelligent tutoring goes to school in the big city. *International Journal of Artificial Intelligence in Education*, 8(1), 30–43.
- [5] Kasneci, E., Sessler, K., Küchemann, S., Bannert, M., Dementieva, D., Fischer, F., ... & Kasneci, G. (2023). ChatGPT for good? On opportunities and challenges of large language models for education. *Learning and Individual Differences*, 103, 102274. <https://doi.org/10.1016/j.lindif.2023.102274>
- [6] National Institute of Standards and Technology. (2023). Artificial Intelligence Risk Management Framework (AI RMF 1.0). NIST AI 100-1. <https://doi.org/10.6028/NIST.AI.100-1>
- [7] International Organization for Standardization/International Electrotechnical Commission. (2023). ISO/IEC 42001:2023 Information technology — Artificial intelligence — Management system. <https://www.iso.org/standard/42001>
- [8] Lyell, D., & Coiera, E. (2017). Automation bias and verification complexity: A systematic review. *Journal of the American Medical Informatics Association*, 24(2), 423–431. <https://doi.org/10.1093/jamia/ocw105>
- [9] Goddard, K., Roudsari, A., & Wyatt, J. C. (2012). Automation bias: A systematic review of frequency, effect mediators, and mitigators. *Journal of the American Medical Informatics Association*, 19(1), 121–127. <https://doi.org/10.1136/amiajnl-2011-000089>
- [10] Dahmani, L., & Bohbot, V. D. (2020). Habitual use of GPS negatively impacts spatial memory during self-guided navigation. *Scientific Reports*, 10, 6310. <https://doi.org/10.1038/s41598-020-62877-0>
- [11] Sparrow, B., Liu, J., & Wegner, D. M. (2011). Google effects on memory: Cognitive consequences of having information at our fingertips. *Science*, 333(6043), 776–778. <https://doi.org/10.1126/science.1207745>
- [12] 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>
- [13] Roediger, H. L., & Karpicke, J. D. (2006). Test-enhanced learning: Taking memory tests improves long-term retention. *Psychological Science*, 17(3), 249–255. <https://doi.org/10.1111/j.1467-9280.2006.01693.x>
- [14] Bjork, E. L., & Bjork, R. A. (2011). Making things hard on yourself, but in a good way: Creating desirable difficulties to enhance learning. In *Psychology and the real world: Essays illustrating fundamental contributions to society* (pp. 56–64). Worth Publishers.
- [15] Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis*, 58(1), 7–19. <https://doi.org/10.1093/analys/58.1.7>
- [16] Anderson, J. R., Corbett, A. T., Koedinger, K. R., & Pelletier, R. (1995). Cognitive tutors: Lessons learned. *The Journal of the Learning Sciences*, 4(2), 167–207. [https://doi.org/10.1207/s15327809jls0402\\_2](https://doi.org/10.1207/s15327809jls0402_2)
- [17] Collins, A., Brown, J. S., & Newman, S. E. (1989). Cognitive apprenticeship: Teaching the crafts of reading, writing, and mathematics. In L. B. Resnick (Ed.), *Knowing, learning, and instruction: Essays in honor of Robert Glaser* (pp. 453–494). Lawrence Erlbaum Associates.
- [18] Organisation for Economic Co-operation and Development. (2019). *OECD skills outlook 2019: Thriving in a digital world*. OECD Publishing. <https://doi.org/10.1787/df80bc12-en>

- [19] Kirkpatrick, D. L., & Kirkpatrick, J. D. (2006). *Evaluating training programs: The four levels* (3rd ed.). Berrett-Koehler Publishers.
- [20] Advanced Distributed Learning Initiative. (2019). *Total Learning Architecture (TLA) 2019 Report*. U.S. Department of Defense. <https://www.adlnet.gov/publications/2020/04/2019-Total-Learning-Architecture-Report/>
- [21] U.S. Air Force Air Education and Training Command. (2018). *Pilot Training Next cadre discuss lessons learned, way forward*. <https://www.aetc.af.mil/News/Article-Display/Article/1638707/pilot-training-next-cadre-discuss-lessons-learned-way-forward/>
- [22] OpenAI. (2025). *GPT-5.1: A smarter, more conversational ChatGPT*. OpenAI. Retrieved from <https://openai.com/index/gpt-5-1/>
- [23] Google DeepMind. (2024). *Gemini 1.5: Unlocking multimodal understanding across millions of tokens* (Technical report and updates).
- [24] Holmes, W., Bialik, M., & Fadel, C. (2019). *Artificial Intelligence in Education: Promises and Implications for Teaching and Learning*. Center for Curriculum Redesign. <https://curriculumredesign.org/wp-content/uploads/AIED-Book-Excerpt-CCR.pdf> (full report also available via ResearchGate and other academic repositories).
- [25] World Economic Forum. (2025). *The Future of Jobs Report 2025*. Geneva: World Economic Forum. Available at: <https://www.weforum.org/publications/the-future-of-jobs-report-2025/>
- [26] Vitale, K. M., Barsuk, J. H., Cohen, E. R., Wayne, D. B., Hansen, R. N., Williams, L. M., ... Schroedl, C. J. (2023). *Simulation-based mastery learning improves critical care skills of advanced practice providers*. *ATS Scholar*, 4(1), 48–60. <https://doi.org/10.34197/ats-scholar.2022-0065OC>
- [27] Merino-Campos, C. (2025). *The impact of artificial intelligence on personalized learning in higher education: A systematic review*. *Trends in Higher Education*, 4(2), 17. <https://doi.org/10.3390/higheredu4020017>