

Ambient Mentorship

AI that helps people become more capable while they work

John Koble | September 2026 | Version 1.0

The proposal

An AI assistant can help someone finish a task and use that experience to support their development. Ambient Mentorship makes that second purpose deliberate. It connects ordinary work conversations to learner-defined goals, remembers useful evidence of progress, and brings relevant learning back into later work.

The guiding ambition is simple: complete today's task while increasing tomorrow's capability. Across an ongoing relationship, explanations, practice and reflection should help the person become better able to reason, judge and act. The timing matters. A mentor must also recognize when the learner needs a direct answer and space to finish the work, a challenge learning research calls the assistance dilemma [1].

A persistent mentor also creates a problem of trust. The observations that make mentoring effective, such as honest questions, half-formed reasoning and misconceptions caught early, are the same observations that become dangerous once they can be used to judge someone. This is the tension between formative use of learner data, to teach, and summative use, to evaluate. Learning research recognizes the tension but has rarely treated it as an architectural problem. In a workplace it is sharper: if people suspect their uncertainty can reach their employer, they stop exposing it, and a mentor that cannot see uncertainty cannot teach.

This paper makes two contributions. The first is a five-component framework for a persistent mentoring relationship, built from established learning research. The second, and the central one, is a trust architecture that separates a private, continuously evolving developmental model, which is never attested to anyone, from stable, learner-authorized outcomes that can travel through existing credential mechanisms. The boundary is not a privacy feature added to the design; it is a condition for the mentoring to work. The paper describes intended behavior and a way to test it; its educational benefits remain to be demonstrated.




Figure 1. The five components form a continuing development loop governed by the learner's goals and boundaries.

Learning that remains connected to work

Structured training offers depth and sustained practice, but the opportunity to use it may arrive much later. Everyday work provides immediate context. Ambient Mentorship connects the two: a real problem reveals a learning need, focused study develops understanding, and subsequent work gives that understanding a place to take hold. The central design question is what the person becomes able to do over time, alongside what the assistant helps them produce today.

02 / AMBIENT MENTORSHIP

Five components of the relationship

The components have distinct responsibilities, but they operate together. The contract governs the relationship, the models guide teaching decisions, and experience provides evidence for revision.

Mentorship Contract

The learner defines what they want to become capable of doing and how they want to be helped. The contract captures development goals, preferred tone and teaching approaches, challenge level, communication preferences, humor and interaction boundaries. It also defines what may be remembered and shared. The learner can revise it as their responsibilities, confidence or available time change.

A practical contract might ask for short explanations during routine work, deeper questions when exploring a new concept, and direct assistance under deadline pressure. It should include an easy way to pause mentoring without losing access to help.

Learner Model

The model holds a provisional account of demonstrated strengths, unresolved misconceptions and areas that need more evidence. It records the context and amount of assistance behind an observation. A correct result produced entirely by the assistant must not be treated as proof that the learner understands it.

The learner should be able to inspect and challenge this account, a principle established in work on open learner models [2]. Teaching preferences are adjustable choices, not fixed labels about a person's mind. Inferences should carry uncertainty, be revised when contradicted, and lose weight when they become stale.

Curriculum Model

The curriculum is a living development pathway. It connects the learner's goals and current work to relevant concepts, prerequisites and practice. It can recommend courses, books, articles, labs, projects or approved enterprise knowledge. Recommendations should explain why a resource matters now and how it supports a stated goal.

Ambient Reinforcement

The mentor reinforces useful ideas inside ordinary conversations through concise explanations, analogies, questions and opportunities to apply earlier learning. A "Socratic Side Quest" is an optional, brief question that invites the learner to reason through a point. Its value depends on relevance and timing; it should not become a compulsory detour.

Reflection and Alignment

Periodic reflection compares evidence of growth with the contract. The learner and mentor review what has become easier, what still requires support, and whether the goals remain useful. That review updates both models and the mentoring approach. Progress should be explainable through concrete observations rather than an opaque score.

An example across several weeks

Consider an analyst building a regional sales dashboard. Their development goal is to understand data relationships well enough to maintain and extend the report. The following sequence illustrates the proposed behavior.

The immediate problem

A dashboard total looks wrong. The analyst asks for help, and the assistant helps trace the problem to the way records are related. It explains the specific issue and supports a correction. Because the analyst is preparing for a meeting, it keeps the explanation short and offers to revisit the underlying concept later.

The learner model records a tentative learning opportunity around relationships and aggregation. It does not conclude that the analyst lacks the skill based on one question. The curriculum identifies a relevant concept and an appropriate practice exercise.

A useful opportunity to learn

During a later conversation, the analyst asks why another table behaves differently. With permission to explore, the mentor asks them to predict what will happen when one record matches several others. The analyst's reasoning gives the mentor better evidence of what is understood and where the explanation needs to change.

When the analyst has time for focused study, the mentor recommends a short lesson and a small exercise using unfamiliar sample data. The selection follows the actual gap. It avoids sending the learner through a broad introductory course merely because that course contains the relevant topic.

Application and feedback

The analyst later adds a new data source. Before offering a solution, the mentor invites them to identify the intended relationship and explain how they would check the totals. If they need help, it supplies a hint and records that support as part of the evidence.

The system looks for transfer: can the analyst recognize the same underlying issue in a different setting? A successful explanation and a correctly designed check provide stronger evidence than accepting a suggested fix. If the concept remains unclear, the mentor adjusts the pathway without turning a mistaken attempt into a permanent judgment.

Reflection and an optional outcome

At a review, the analyst sees a concise account of progress: which decisions they can now explain, what they handled independently, and what still needs practice. They can correct the record or change the goal. If they want to share an achievement, a separate assessment process can produce a narrowly stated outcome with its evidence and conditions.

The employer might receive an authorized statement that the analyst demonstrated a particular data-modeling capability. The private questions, tentative misconceptions and conversational history that supported that growth remain within the developmental relationship.

How to recognize genuine growth

A useful deliverable proves that work was completed. It does not, by itself, establish how much the person learned. Ambient Mentorship therefore needs to distinguish the capability of the person from the combined performance of the person and the assistant.

Assessment can draw on ordinary work, but it must be deliberate about attribution. Each observation should identify the task, relevant goal, assistance provided and reason for the inference. The model should retain “not enough evidence” as a legitimate result.

Evidence worth looking for

Evidence	What the learner demonstrates
Explanation	Explains why an approach works and where it may fail, using their own reasoning.
Transfer	Applies a concept to an unfamiliar problem with different surface details.
Judgment	Recognizes an error, questions a plausible answer or chooses between competing approaches.
Retention	Uses the concept again after time has passed, without reconstructing it entirely through assistance.
Independence	Needs less support on comparable tasks, with task difficulty and available tools recorded.

Continuous observation with focused checks

Conversational evidence can suggest progress, but important claims deserve a focused demonstration. An occasional explanation, new scenario or short exercise can test an inference. The learner should know when they are practicing and when they are being assessed, especially if an outcome could be shared.

Fluent language, technical vocabulary and confident answers are insufficient on their own. Nor should a request for help automatically count against the learner. Choosing when to seek assistance can itself reflect sound judgment.

Testing the proposal

A pilot should compare ordinary AI assistance with Ambient Mentorship in one defined skill area. Use a baseline task, a later unfamiliar task and a delayed check, with comparable tools and a consistent rubric. Track completion quality and time alongside retained understanding, transfer, assistance required and interruption burden. These measures would test whether the mentoring improves learning without making useful assistance harder to obtain.

The trust architecture

A mentor needs room to observe uncertainty, and a learner who suspects their uncertainty could reach an employer will stop revealing it. The developmental model is private, continuously evolving and non-attestable: it informs teaching and is never issued as a claim about the learner. Anything shared externally is a separate, stable outcome the learner reviews and authorizes.



Figure 2. The disclosure boundary. A shared outcome does not carry the developmental record behind it.

Meaningful learner control

“The learner owns the learner model” is a design commitment that must be implemented through specific controls. The learner should be able to inspect observations, correct mistaken inferences, set retention preferences, remove information and understand who can access it. Any limits to these controls must be visible before the relationship begins.

Employer access should not be the default destination for the private model. Participation and disclosure choices must be meaningful, and learning goals should remain distinguishable from organizational reporting requirements. Emotional observations and inferred personality characteristics should not become routine inputs merely because they can be inferred.

A separate path for shared outcomes

A shared outcome should identify the capability demonstrated, assessment conditions, date and accountable issuer. The learner reviews the proposed claim and authorizes its disclosure. A translation layer selects only the necessary outcome information; it must not export the private model or repackage speculative judgments as credentials.

The W3C Verifiable Credentials Data Model provides a standard way to express issuer claims and exchange them among issuers, holders and verifiers [3]. Verifying a credential’s integrity does not establish that an assessment was valid or that the recipient should trust its issuer.

The boundary must hold in practice

Access controls, separate storage and disclosure records should enforce the boundary, but they are not sufficient: the private model can leak through inference even when no raw record leaves it. A summary or small-group report can reveal individuals without including any conversations. Absence is informative: when peers share an outcome and one learner does not, the missing credential says something. The timing of disclosures can reveal how long a capability took to develop, and even pausing mentorship can be read as a signal. Mitigations include minimum group sizes for aggregate reporting, outcomes that state the capability without the path taken to reach it, disclosure windows that separate when a capability is shared from when it was demonstrated, and organizational rules against drawing conclusions from non-disclosure.

A practical path to implementation

Begin with a bounded mentoring relationship

An initial implementation should focus on one skill domain, a small set of learning goals and an explicit contract. It needs a learner-visible record of evidence, a resource map, controls for intervention timing and a reflection process. A readable contract can be paired with machine-readable settings so that permissions and teaching preferences govern behavior consistently.

Learning decisions should be traceable. The learner should be able to ask why an explanation appeared, why a resource was recommended or why the model changed. A correction must influence subsequent behavior, and pausing mentorship must stop optional teaching interventions.

Connect work with structured learning

A first implementation could coordinate resources the learner already has access to: courses, books, labs and approved enterprise knowledge. The mentor does not need to own the content. Its value lies in deciding what to teach and when, using whatever resources fit the learner's goals and current work.

The resource pathway can also include generated explanations, practice scenarios and exercises where their quality can be checked. More extensive personalized learning experiences would require reliable content validation and assessment design. The five responsibilities can remain recognizable as the implementation changes; the technical architecture will still need to evolve.

Build on established learning ideas

Ambient Mentorship draws on established work: intelligent tutoring systems and their evidence base [4], the assistance dilemma of when to give help and when to hold it back [1], learner models that people can inspect and challenge [2], and cognitive apprenticeship, in which skills are learned inside authentic work [5]. The components are not new in isolation. What this paper adds is their arrangement inside a persistent AI relationship governed by a learner-authored contract, and a trust architecture designed for the workplace setting where that relationship would operate.

The trust architecture is the central contribution. Developmental observations support growth and are never attested, while any external claim of achievement follows an explicit assessment and disclosure process. Self-sovereign identity, verifiable credentials and learning and employment records address the portability of outcomes [3]; none of them protects the private developmental state that produced those outcomes. That is the gap the boundary is designed to close.

The intended result

The measure of success is whether the learner becomes more capable over time and can see credible evidence of that growth. The relationship should support curiosity without penalizing uncertainty, offer challenge without obstructing work, and preserve the learner's authority over their developmental record. If these conditions hold, everyday assistance can become a sustained opportunity to learn.

References

- [1] Koedinger, K. R., and Aleven, V. (2007). Exploring the assistance dilemma in experiments with Cognitive Tutors. *Educational Psychology Review*, 19(3), 239–264.
- [2] Bull, S., and Kay, J. (2007). Student models that invite the learner in: The SMILI open learner modelling framework. *International Journal of Artificial Intelligence in Education*, 17(2), 89–120.
- [3] World Wide Web Consortium. (2025). Verifiable Credentials Data Model v2.0. W3C Recommendation, 15 May 2025.
<https://www.w3.org/TR/vc-data-model-2.0/>
- [4] VanLehn, K. (2011). The relative effectiveness of human tutoring, intelligent tutoring systems, and other tutoring systems. *Educational Psychologist*, 46(4), 197–221.
- [5] Collins, A., Brown, J. S., and 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.

© 2026 John Koble. All rights reserved. Ambient Mentorship, Version 1.0, September 2026. JohnKoble.com