THE SIGNAL BRIEF / MEMORY & DECISION-MAKING / JULY 30, 2026

# Your Personal AI Needs a Change Log

## Memory can make advice fit your life. A review process helps it keep fitting when your life changes.

Don Long · Founder, SummitMark · Helping people and organizations become capable with AI

An illustrative still life: the record is visible, human-reviewed, and still only part of the decision.

> *The useful personal AI is not the one that remembers everything. It is the one that can show its work, notice a conflict, and ask what is current.*

Most AI advice fails in a quiet way. It sounds sensible for someone with a different schedule, a different risk tolerance, or a different history than yours.

That is the problem Don Long is trying to solve with Cora, a personal AI built around plain-text notes and instructions. The notes record what he is building, what is true now, what he will not trade away, decisions and their reasoning, paused projects, and corrections to the system's misunderstandings.

The idea is not to make a model magically wise. It is to give an existing model a record it can inspect, along with rules for how to use that record.

# Memory is only half the problem

A personal record can make a response more relevant. It can also make an old assumption more durable.

A boundary from six months ago may still be firm. Or a circumstance may have changed. A project may look active in a note even though it has been paused. A sentence written during a difficult week can be mistaken for a permanent preference. If the system quietly chooses which version to believe, personalization turns into a confident guess about who you are now.

Long's proposed answer is a deliberate review loop. The AI can suggest that a conversation revealed a new goal, correction, decision, or contradiction. The person reviews the exact wording and approves it before it enters the lasting record. A temporary feeling or AI interpretation does not become biography by accident.

That human approval is not clerical overhead. It is the control that keeps a useful record from becoming an unexamined authority.

# The record needs operating rules

Long's starter system separates current facts, goals, boundaries, decisions, and rules for the AI. It asks the assistant to distinguish what the person said from what it inferred, flag information that may be stale, and identify a missing fact before making a consequential recommendation.

Those are good standards for a decision partner. They are not guarantees. An instruction can ask an AI to retrieve relevant notes; it cannot prove that the system found every relevant note. A polished context receipt can still omit something. The person needs a way to inspect and correct the system's account.

Long's own example makes that limit visible. Cora compared several projects and assigned confidence-like scores. The scores gave him a quick view of where he saw stronger evidence, but he also had to correct the result. Those percentages are an author's description of one interaction, not calibrated measurements of project quality or proof that Cora makes better decisions.

The useful part is the inspection: which projects had artifacts behind them, and which mostly had activity? The number is a prompt for a closer look, not a verdict.

# A folder is not an automatic workflow

The files can stay ordinary Markdown on a computer, but a local folder does not mean a conversation is processed locally. When a cloud AI reads those files, their contents are sent to that provider for processing. Long explicitly recommends a dedicated folder and avoiding secrets, medical records, confidential work material, and other information a person would not knowingly share.

The connection method also matters. Product capabilities and plan availability change. As of September 13, 2026, Anthropic documents local desktop extensions for access to local files, available in Claude Desktop and Claude Code. Google's Gemini personalization features depend on eligibility and account type. OpenAI's current Memory FAQ describes a memory summary and response sources, while warning that those sources may not show every factor shaping a response. None of those product features, by themselves, guarantees that a hand-built personal record will be read completely or kept current.

So the durable part of Long's approach is not a specific app combination. It is the operating pattern: keep the record inspectable, tell the assistant how to reason with it, review proposed changes, and ask what could have changed since the last conversation.

# Try one reversible decision

Start with a question where a poor answer is inconvenient, not dangerous. Ask an AI for help choosing between two projects, planning a learning goal, or deciding what to prioritize next month.

Then ask it to show a short context receipt:

- Which notes did you use?
- Which outside sources did you check?
- What did you assume?
- What is unknown or possibly out of date?
- Which missing fact could most change the recommendation?

If the assistant names a material gap, answer that question before asking it to recommend. If it spots a conflict between an old boundary and a newer goal, decide which record is current. At the end, ask what it thinks should be updated, review the exact proposal, and approve only what you want saved.

This is not a test of whether an AI is wise. It is a test of whether you can see the context behind its advice, correct the record, and stay in control of what becomes permanent.

Cora is still an experiment. Long does not claim that a folder of notes reliably improves decisions, and a detailed process can still produce personalized mistakes. The value to test is narrower: does the system make its assumptions and missing context easier to catch before you act?

What would you want a personal AI to remember, and what should it ask you to confirm first? Reply with the line you would draw.

M.R.

Source note: This issue is based on Don Long's user-supplied article, “How to Build a Personal AI for Decisions That Actually Fit Your Life.” Cora's setup and project comparison are attributed to Long and were not independently reproduced. Current feature statements were checked against official documentation on September 13, 2026: Anthropic's guide to desktop and web connectors, Google's Gemini personalization guidance, and OpenAI's Memory FAQ. These product descriptions can change; they do not establish that any system retrieves a personal folder completely or improves decision quality.