

# CONTINUITY WITH *consent.*

*You decide.*

Designing AI memory people can permit, inspect  
and correct.

---

Steph Campagnoli · Product direction, interaction design, requirements  
StephOS · January–October 2026

# AI KEPT SAYING IT REMEMBERED.

Every new chat started from zero. And when an assistant did "remember," it sometimes invented details, claimed work was saved when it wasn't, or replaced my original words with its own summary. For long-running work, that's worse than forgetting: you can no longer tell what to trust.

FORGETTING IS AN  
INCONVENIENCE.  
FALSE MEMORY IS A  
PROBLEM OF *Trust*.

## SO I SET THREE FOUNDING REQUIREMENTS

### ORIGINAL WORDS ARE EVIDENCE

Summaries and decisions link back to what was actually said.

### APPROVAL IS EXPLICIT


Nothing becomes approved on its own.  
Discussion is not permission.

### DONE MEANS I CAN OPEN IT

A receipt or a claim is not proof. The saved record has to open and check out.

# I DESIGN THE BEHAVIOR. AI WRITES THE CODE.

Sole designer and tester, January–October 2026. I set the goals, requirements and interaction direction, compared models, and decided what met acceptance. AI wrote the code and assisted research and documentation.



# ONE HOME FOR MEMORY *decisions.*

## SEPARATE PERMISSIONS


Searching on this Mac and sending to a cloud model are different grants. Allowing one never implies the other.

## DISTINCT CHOICES, NOT ONE SWITCH

Turn memory off, keep a chat out, or skip just the next message.

## REACHABLE, NOT REPEATED

Controls live in one place instead of under every message.



Interface capture · synthetic records · scripted replies

# CONVERSATION FIRST. EVIDENCE ON REQUEST.


At rest, one quiet line: "Memory · 1 source." Open it and you see the exact excerpt, who said it and when, and how the assistant used it.

## THE INTERFACE STATES ITS OWN LIMIT


"StephOS checked only that each label refers to an excerpt that was supplied; it did not verify the judgement."

HONESTY IS A DESIGN SURFACE, NOT A FOOTNOTE.

## AT REST



## ON REQUEST



# RETRIEVAL ONLY WITH *permission.*

Every answer that draws on memory passes the same path. Each step is a boundary the person can rely on.



TRADE-OFF

Automatic recall removes repeated manual choices but adds consent, relevance and index upkeep. I chose to carry that complexity in the system, not hand it to the person.

# RIGHT ANSWERS. UNINVITED MEMORY.

18/18

facts recalled correctly

2 of 4

unrelated questions still drew on  
memory

MEMORY YOU DIDN'T  
NEED IS STILL  
DISCLOSURE.

RELEVANCE IS A *consent*  
PROBLEM.

## WHAT I CHANGED

The assistant now reads the recent conversation to resolve references, topic changes and corrections, searches with a specific subject, and asks when something is unclear. Sources reused from earlier answers are rechecked and capped. Shipped October 19; now being judged in everyday use.

24 synthetic cases across two versions. Not a claim of 100% accuracy.

# NOTHING SHIPS ON ITS OWN WORD.

My rule: no release installs without a passing independent review. The conversation release needed one fix before it passed. The next feature, durable progress notes, went through three reviews and five fixes before it was confirmed and installed.

PROGRESS NOTES · CONFIRMED, INSTALLED OCT 19

A DEFECT WORTH LEARNING FROM

Withdrawing a later correction could make an older value look current again.

NEW RULE

Undo never silently restores. Bringing a value back takes an explicit new correction.

REVIEW 1

Three defects found · all fixed

REVIEW 2

One subtler defect · fixed

REVIEW 3

One more edge case · fixed

CONFIRMATION

Passed · installed Oct 19

# WHAT'S PROVEN. WHAT ISN'T YET.

## PROVEN IN MY OWN USE

- Permitted automatic recall across chats
- Memory controls in one place
- Original words preserved and linked to answers
- Every release independently reviewed before install

## NOT YET

- No external users yet
- Relevance is improved, not solved
- Imported chat history is out of scope
- The conversation release is still being judged in use

## WHAT I'D TEST WITH REAL PEOPLE

I haven't run this yet. The study I'd run: a few people resuming real work. Can they tell where an answer came from, choose what memory may use, and recover from an irrelevant recall? Success defined before the first session.

# AI THAT *people* RELY ON.

## TURN AI FAILURES INTO REQUIREMENTS

Find the problem people work around and specify what the system must do instead.

## DESIGN CONSENT AND EVIDENCE IN

Make permission, sources and limits part of the interface, not the fine print.

## SHIP THROUGH REVIEW

Independent checks before release, and plain words about what isn't proven.

DISCUSS A ROLE OR COLLABORATION

hello@stephos.studio · stephos.studio · linkedin.com/in/steph-campagnoli

*Steph Campagnoli*