Generated by author's AI

Human Model — Research Status

24 October 2026

Current status: Secondary-data research closed. Next phase requires real participants.

Research journey

Phase	What we did	What we learned
1. Thesis & literature	Explored personal models, memory, preference learning and personalization research.	Significant prior art exists. A persistent Human Model is a hypothesis, not an assumed solution.
2. Benchmark landscape	Audited existing personalization benchmarks and identified missing comparisons.	No verified unified benchmark established incremental real-user value against strong retrieval and clarification baselines.
3. Synthetic experiments	Built a controlled PAHF experiment covering 900 synthetic decisions and seven architectures.	Retrieval + targeted clarification reached **79.4% accuracy**, versus 32.6% for context alone. The tested persistent note memory underperformed retrieval.
4. Identity-residual experiments	Used Amazon data to investigate whether the correct person's history carries unique predictive information.	Donor matching and similarity measurement proved insufficiently reliable. Experiments stopped before outcome evaluation. Hypothesis remains untested.
5. Dataset exploration	Audited 49 datasets/instruments, including Amazon, Google Local, Ciao, Douban and KuaiRand.	Public behavioral datasets could answer narrower predictive questions but lacked key evidence about actual user benefit and interaction burden.
6. Real-behavior replication	Attempted to compare retrieval and compact persistent state using KuaiRand.	First execution failed a temporal protocol requirement; prospective replication failed its data-coverage gate. No valid confirmatory comparison emerged.

What we established

- Relevant retrieval matters substantially more than merely possessing historical information in our synthetic test.
- A richer persistent person representation has not demonstrated incremental value beyond strong retrieval and selective clarification.
- Existing datasets did not provide a sufficiently reliable instrument to answer our remaining real-world questions.
- We have not disproved the Human Model thesis, but neither have we validated the necessity of a compact persistent person model.

Why we now need actual participants

The question has shifted from *"Can AI use information about a person?"* to *"Does maintaining a richer understanding of that person actually improve their experience and decisions?"*

Historical datasets cannot adequately measure whether people endorse an intervention, how much effort clarification requires, whether memories become inappropriate, or how preferences change through real interactions.

Our next research phase needs consented participants making real decisions, with outcomes and friction measured prospectively.

Next experiment

Compare two systems using equivalent prior information:

A — Strong baseline	Current context + relevant episodic retrieval + selective clarification.
B — Human Model challenger	The same system plus editable, compact persistent person-specific state.

Measure real user-endorsed outcomes, effort reduction, correction burden, and performance on future decisions.

Bottom line

We have completed the controlled synthetic and available secondary-data phases. The next meaningful scientific advance requires a purpose-built human study—not another round of dataset searching.

Research summary derived from the Human Model canonical research state and the closed PAHF, identity-residual, dataset landscape and KuaiRand experiment records.