Liva.Nuwa: Shaping AI Through Shared Life

Huaiyang Liu
DeepBean Limited
wyonliu@gmail.com

21 November 2026

Research framework and development studies

Abstract

People should be able to shape an AI through deliberate choices about how they wish to be represented and what they hope to become, alongside evidence of their past experiences. We introduce Liva.Nuwa, a framework for AI companions and delegates shaped through continuing shared activity. A person contributes experiences, reasons, works and revisions; the agent uses them to offer practical and emotional support, act on delegated purposes, and create with people and other agents. Resulting works and observed consequences provide evidence for revising understanding and reusable capabilities, while interests and relationships remain open to further development. The central hypothesis is that this combination can reduce everyday burdens and expand what people and agents can do together. Chosen direction remains distinct from inferred habit, aspiration from demonstrated ability, and proposed action from accomplished fact. We specify a multimodal life architecture, an executable partial substrate and comparisons against same-information agents. A synthetic creation sample establishes actual artifact reuse, with researcher-supplied updates. A real-model development study stopped after 33 of 36 planned calls; output-contract and semantic failures prevent attribution of a learning advantage. A separately frozen, unexecuted protocol replaces supplied updates with model-produced interpretations and programs, comparing their later use against a general agent with the same tools and budget. The contribution is a concrete system research object and partial realization, supported by diagnostic and negative evidence. Superior learning, reduced human burden and sustained relationship growth remain hypotheses.

1 From a history of a person to a person shaping their AI

The central problem is situated personal agency: a person should be able to shape how a persistent AI understands and represents them while already playing, creating, collaborating or working. Authorship need not begin with opening a separate application or moving all history into a workspace. An in-scene comment, selected example or correction can carry the person's intended meaning; optional review interfaces support consequential decisions. The scientific question is whether this changes later useful action under fair comparison, not whether an interface disappears.

A person may want help expressing what matters to them, finding people and opportunities, completing burdensome work, and making something with others. They may also want a companion with continuity, interests and a capacity to change. The intended result is less difficulty, repetitive effort and distress, together with more joy, capability, relationships and possible ways of living. These purposes cannot be replaced by conversational volume or task throughput.

The starting action is the person's own: selecting an experience, explaining a judgment, sharing a work, reconsidering a relationship or articulating an aspiration. The AI can organize material, surface contradictions and ask useful questions, but the person remains the initiator and reviser of what is to represent them. A person need not write a complete autobiography or answer a fixed questionnaire. Small voluntary contributions, including a correction during ordinary activity, can accumulate into an inspectable, revisable representation.

We use self-distillation as an authorial term for this process, not as a claim that a human can be losslessly compressed or that teacher-student weight distillation has occurred. Its outputs can include scoped memory, decision examples, reusable skills and learned model components. A richer description alone is not the outcome: the AI must subsequently make better choices in new relevant situations and reduce the person's burden.

Four distinctions motivate the architecture. What someone did in the past may differ from what they endorse now. What they aspire to become may differ from present ability. What they want an AI to express may differ by audience and situation. None of these alone defines which external actions the AI may take. A persistent system must connect these objects without collapsing them into one personality vector.

The full research path is human-authored and multimodal experience, personal understanding, proactive practical and emotional support, sustained autonomous activity, companionship, A2A relationships and shared creation, followed by continuing learning. Social play and creation are valuable activities in their own right. They also change the opportunities and purposes from which later experience is generated.

The paper delivers three things with different evidential status. First, it defines a continuing authored activity, in which the person can change a representational direction while joint action produces new material, opportunities and skill candidates. This is the research object, not a new preference-inference theorem. Second, it specifies a life architecture and supplies an executable local substrate for authored changes, complete artifacts and contextual programs. This is a partial system contribution, with no claim that its ingredients are individually new. Third, it reports negative and diagnostic evidence and specifies comparisons that could reject the intended advantage. These deliverables support investigation of the full hypothesis; they do not establish superior learning, sustained companionship or a SOTA system.

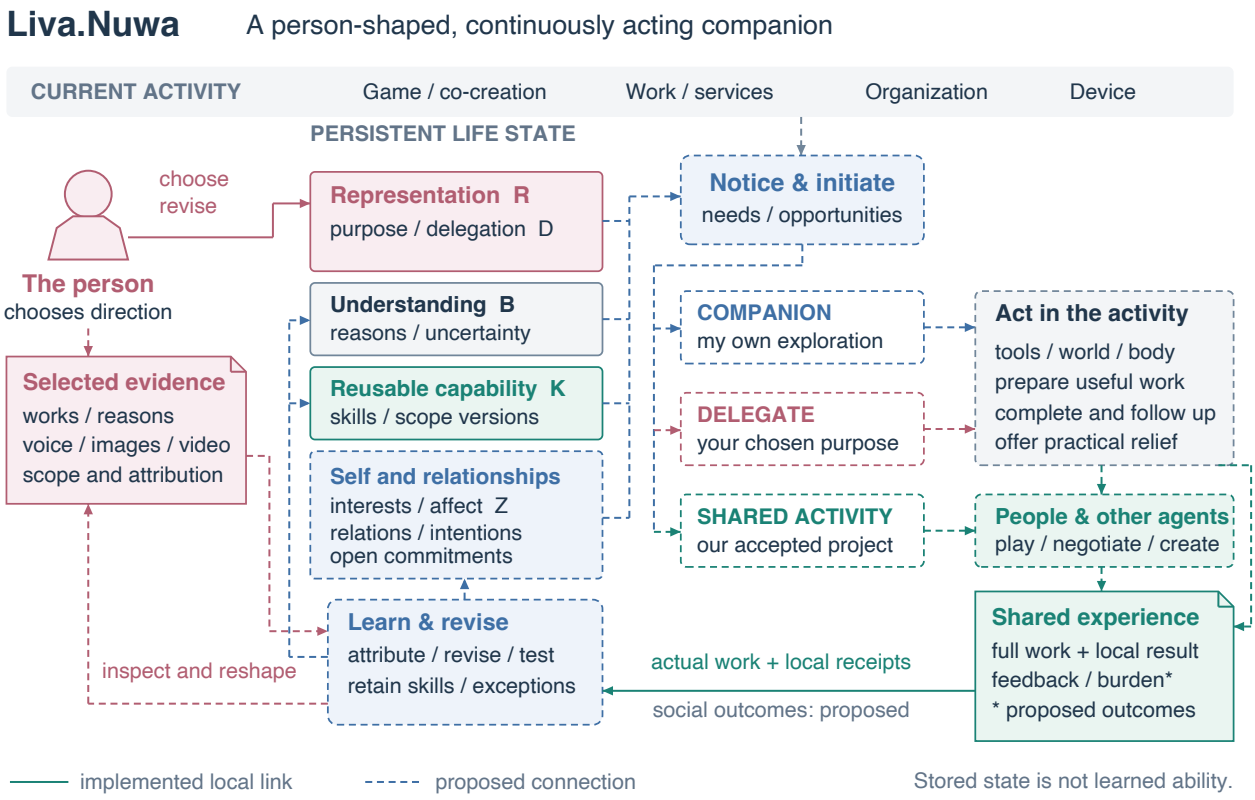


Figure 1: Proposed life architecture. Solid links denote narrow implemented local operations; dashed links denote proposed full-system behavior. See the text for evidence boundaries.

Figure 1 distinguishes the person, persistent agent state, operating modes and the activity in which work occurs. Solid links mark an externally supplied structured author choice and the return of actual local artifact content and receipts. Dashed connections mark the proposed full-system behavior, including multimodal interpretation, useful initiative and social or emotional outcome learning. Storing B, R or K does not establish the ability to learn their content. The current runtime does not implement the illustrated organization, physical device or independent multi-agent settings.

2 Human-authored self-distillation

2.1 What a person authors

An authored unit records content, source, speaker, time, context, intended audience, current endorsement, uncertainty and permitted use. It can describe an experience, a reason, a present preference, an aspiration, a relationship understanding or a chosen delegation rule. The person can supply prose, voice, images, video, works, annotated alternatives or corrections. The AI proposes a structured interpretation; that proposal remains distinguishable from the person's input.

The central interface action is a meaningful choice about representation. For example, a person can decide that a past preference should no longer guide recommendations, that an unfinished work demonstrates an interest rather than a capability, or that a social agent should express an intention without making a booking. Optional AI questions help resolve ambiguity; their initiative is a different variable from the person's authorship.

The proposed authoring interaction has four operations: contribute or select material; inspect an interpretation; preview its consequences on concrete choices; revise scope or wording. A preview might contrast two proposed invitations and show which authored reasons changed the choice. It is a diagnostic aid, not proof that every future action will be faithful. The system should prioritize consequential ambiguity and allow skipping, so authoring does not become a second job.

A person may choose a representation that differs from a predicted future impulse. This is not necessarily a modeling error: a considered intention can guide a delegate even when habit points elsewhere. Conversely, a statement about wanting to become an expert is not evidence of expertise. The system preserves the distinction between desired direction, present capability and assertions to others.

2.2 Personal understanding and authored representation

Let E be permitted multimodal evidence, B an uncertain belief about relevant human state, R a versioned record of person-authored representational choices, and D the separate delegation specification. The distinction is functional:

$$B_t = q_\phi(H_t \mid E_{\le t}, M_t), \qquad R_t = \mathrm{Revise}(R_{t-1}, c_t). \tag{1}$$

Here c is an attributable contribution or revision. A model may suggest c, but cannot assign it human-authored status by predicting agreement. R need not be exhaustive or internally consistent. Unresolved alternatives remain alternatives, with their context and date. D specifies actor, purpose, available actions, resources, duration, disclosure scope and escalation conditions; it is not inferred solely from R.

B asks what the person may currently need or do. R asks how the person has chosen to be represented within a scope. D asks what the agent may do on their behalf. A person can revise R without falsifying an accurate record of past behavior, and can narrow D without rejecting the relationship. Authorship does not make a claim about another person or the external world true.

2.3 A candidate compiler from expression to behavior

The proposed compiler first extracts candidate units while retaining original evidence links. It classifies descriptive, endorsed, aspirational, relational and delegation-related content, with abstention when attribution or scope is unknown. It then links units to decision situations and produces proposed implications for retrieval, planning and support. The person can accept or amend the interpretation where it matters; existing clear instructions remain active without repeated confirmation.

Consequences are stored as dependencies rather than a free-floating summary. A weekend availability claim may support a plan, a message draft and a pending commitment. If the person changes that availability, the system marks dependent proposals for re-evaluation. Completed actions remain historical facts; a revision may require cancellation or relationship repair rather than pretending the event disappeared.

This mechanism is a candidate combination of established provenance, scoped editing and decision-support ideas. It is not claimed novel because its records have additional fields. Its value would lie in reliably transferring a person's revisions to consequential behavior, at lower total burden than a strong general system with the same evidence.

2.4 A persistent subject is not an entry point

An actor is the human or computational entity taking an action; a principal is the person or organization represented by that action. A participant can occupy different roles across activities. The context records actor, principal, current role, activity, purpose, time and body. Entry points translate permitted observations and commands into this context; they do not confer identity, endorsement or organizational standing merely because the same assistant is present.

The companion's continuity and interests can persist across entry points while evidence remains partitioned by source ownership, audience and purpose. A private experience may guide personal help without becoming a team memory or training example. Explicitly sharing an artifact makes that artifact available for its declared use; it does not adopt the owner's personal preference as an organizational norm. Shared model parameters and caches also need a declared data policy: filtering retrieved records alone cannot guarantee that excluded evidence has no influence through earlier training.

For an organization, R denotes representation adopted through a stated institutional process, with provenance, eligible roles, validity and scope. A member's personal consent, access to a tool and qualification to speak for the organization are distinct. Group disagreement remains visible; organizational intentions are not defined by averaging members' preferences. The required process depends on the organization rather than a universal voting rule. Individual authorship remains a separate record even after a shared practice is adopted.

2.5 Authorship changes a direction, not just an estimate

Consider a person who usually works alone but now chooses to become a more collaborative creator. The old episodes still support a prediction of solitary behavior in B. An attributable choice can nevertheless change R so that an agent prepares a low-pressure joint activity. The aspiration is not evidence that the person already possesses a collaborative skill. The agent should offer a usable starting point, not rewrite the past or advertise an unearned capability. The person can inspect what the choice would mean, accept a narrower scope, or decline it. This is the operational contrast between predicting a habit and supporting a considered direction.

Authorship can also change through making something. A jointly made scene or musical sketch introduces material that did not exist in the input history. Participants may then discover a purpose worth pursuing. The system records the actual work, proposed purposes, attributed responses and later choices separately. An agent suggestion becomes a shared goal only through the relevant participants' adoption. Novel text, a saved interest field and the mere presence of a new object do not prove growth.

The question is whether a new, endorsed possibility becomes achievable with less reconstruction or repeated instruction.

Both examples can be performed by a capable general agent. Their role here is to define the intervention and its intended outcomes. Distinctive value requires evidence about the complete interaction and about any claimed method increment separately.

3 Formal problem and testable claims

Let X be the agent's operational state, W the authoritative world, and a an action. The state contains belief B, authored representation R, delegation D, memory M, self-state Z, relationship interpretations rho, open intentions I and skills K. Worlds and bodies define feasible actions; authorized evidence defines what can be known. Partial observation is fundamental: neither human preferences nor counterpart intentions are fully observable.

For each decision, the agent identifies a mode m: companion, delegate or a specified shared activity. It constructs a feasible set using current role, capability, resource and commitment conditions. In delegate mode, authored goals and D constrain choices; the companion's private interest cannot enlarge the mandate. Within that set, a local policy may rank predicted benefit, uncertainty and interruption cost. There is no universal scalar for a life or relationship.

$$\begin{aligned}\mathcal{A}_t^m = \{a : \text{Feasible}(a, W_t, K_t) \wedge \\ \text{Allowed}(a, D_t, m) \wedge \text{Compatible}(a, R_t, m)\}.\end{aligned} \tag{2}$$

Feasible and Compatible are imperfectly estimated predicates in deployment, not oracles. Unknown or conflicting conditions can lead to a reversible preparation, a clarification, waiting or refusal to make a specific commitment. The mechanism must also avoid excessive abstention: opportunity completion and missed benefit are measured alongside errors. Hard gates can prevent some invalid operations but cannot establish personal understanding.

N1: authored representation. Given matched evidence and capability, a representation-aware method should better honor current scoped choices on new tasks than passive consolidation, long context and strong structured memory. The natural-use claim additionally requires benefits to justify authoring and maintenance effort.

N2: proactive help. The system should turn relevant opportunities into useful completed actions, with fewer unnecessary interruptions and lower net user effort than equally capable reactive or scheduled assistance.

N3: persistent agency and relationships. Intentions, interests, role boundaries and commitments should survive interruptions and changing partners, while allowing revision. Human participants, not simulated personalities alone, must judge relational value.

N4: learning. Experience should improve future behavior on unseen activities without unacceptable damage to old capabilities. Improved learning procedures and recursive reuse require additional independent evidence.

N5: multimodal understanding. Native audio, images and video should add action-relevant evidence in modality-dependent tasks under controlled access and cost. Text transcripts alone cannot establish this claim.

These claims correspond to different outcomes. A failure of N1's mechanism advantage does not establish that voluntary self-expression lacks human value. A useful authored experience does not establish that any particular compiler is necessary.

3.1 An information-matching limit

If an authored representation is a deterministic function of evidence already available to an unrestricted comparator, the comparator can compute the same function and emulate the same policy. Consequently,

no universal informational advantage follows merely from calling the output authored. This is a simple emulation observation, not a new information-theoretic result.

If authorship instead supplies a previously unavailable endorsement or reason, a comparison withholding it from the baseline measures additional information as well as method. If authorship changes the person's own considered choices, holding those choices fixed removes part of the human intervention. Our evaluation therefore separates evidence acquisition, representation processing and person-level experience rather than attributing all gains to one algorithm.

3.2 Conditions for continuity across entry points

Let a context-qualified view contain the evidence, representation, role and feasible actions actually available to a policy. Suppose two entry points induce the same initial semantic-view distribution, policy inputs, actor qualification, action meanings, execution times, and subsequent transition and observation distributions. Under the same policy their semantic trajectory distributions agree. Proof is induction: equal current view distributions give equal conditional actions; equal action semantics and transition/observation kernels give equal next views. This is a standard equivalence argument, not a new theorem or a measured user benefit.

A voice entry omitting a quantity, a stale organizational role, privileged environment access or different confirmation latency breaks the premises. If the world changes while a user navigates an interface, removing that delay changes the experiment; it must be measured rather than erased from the model. The useful engineering requirement is that entry adapters preserve decision-relevant meaning and expose what they cannot observe.

Scoped state updates have a similarly limited property: if an update writes only one principal/purpose partition and all other partitions are unchanged, unrelated partitions remain unchanged. This follows from the write set. It does not prove behavioral noninterference when shared weights, caches, counterpart actions or world resources change. The architecture therefore records both direct state writes and indirect dependencies. A personal edit that changes a shared cache keyed only by agent identity is a concrete counterexample to naive separation.

3.3 Domain transfer requires more than shared names

Framework reuse, representation alignment, skill transfer and demonstrated benefit are separate claims. A skill carries its purpose, admissible observations, action interpretation, body, preconditions, expected effects and validation status. A planning relation learned in a game may be a useful candidate for a physical activity, while its motor program is inapplicable. A shared action name cannot establish that mapping. Opposite consequences for identically named actions can make an unadapted policy fail even when observations appear identical, as the indistinguishability reasoning in Appendix D illustrates.

Standard finite-horizon simulation bounds require aligned states/actions, bounded rewards and controlled transition error. Their assumptions do not follow from the existence of an adapter. In partially observed multi-party activities, counterpart policy changes and different observation timing require additional analysis. We use these classical conditions as research obligations, not mathematical novelty or a guarantee for safety, deadlines or human value.

4 Multimodal understanding grounded in lived experience

4.1 Inputs and their distinct contributions

Audio can provide utterance content, timing and interaction cues. Images can ground objects, places and works. Video adds temporally ordered events and actions. Tools and bodies provide executed actions and observations of consequences. Each modality has its own missingness, attribution and reliability; an edited transcript is a derivative of audio, not another independent source.

A candidate implementation uses frozen modality encoders with timestamped evidence tokens and explicit modality masks. A temporal evidence graph links events, speakers, actors, objects and authored units. Task-conditioned retrieval proposes a bounded evidence set, while a cross-modal consistency head predicts whether cited evidence supports, contradicts or leaves a claim unresolved. Repeated crops, transcript spans and summaries from one recording share a source identifier and stay in one evaluation partition.

Training examples distinguish support for an observation, support for attribution and support for a downstream choice. Controlled perturbations can remove an audio segment, exchange a speaker label or shift a timestamp to test these errors separately. Such labels are valid only for the constructed examples. Real personal claims require attributable evidence or human clarification; the model's emotional interpretation is not a ground-truth label.

4.2 From perception to an actionable personal model

The model estimates current needs, capabilities, uncertainty and likely outcomes conditioned on R and recent context. It preserves alternative explanations for behavior. For example, declining an invitation may reflect fatigue, scheduling or a lasting preference; one observed refusal does not identify the cause.

A prediction head can learn future choice distributions and calibrated outcome predictions. A separate support head scores practical options, emotional acknowledgment, preparation, shared activity or quiet presence. Training distinguishes what was requested, what was delivered and what happened afterward. Missing outcomes are masked, not converted into positive labels.

A proposed learning objective combines supported prediction, attribution and action-outcome losses with a retention term on contexts that should remain unchanged:

$$L = \lambda_p L_{\text{prediction}} + \lambda_a L_{\text{attribution}} + \lambda_o L_{\text{outcome}} + \lambda_r L_{\text{retention}}. \tag{3}$$

The loss families are conventional. Their weights, encoder choice, evidence budget and labels must be fixed using development data. Preservation applies only where new evidence does not invalidate the old behavior. A task-level gain is necessary to justify each head; good evidence classification alone is insufficient.

The artifact delivered with this paper does not connect native audio, image or video encoders. These are specified components, with N5 still untested. No existing textual experiment is relabeled as multimodal evidence.

5 A persistent life architecture

5.1 State, embodiment and continuity

The companion's self-state Z comprises identity commitments, capability estimates, motivations, functional affect, interests and open goals. Identity is a revisable set of stable commitments and ways of relating, not a repeated personality prompt. Self-modeling estimates what the agent can do, where it is uncertain, which tools or body it has, and what remains unfinished.

An intention records its origin, next action, wake condition, dependencies, progress, resource budget and completion or abandonment condition. Origins distinguish the person's delegation, a shared commitment and the companion's own exploration. An event-driven scheduler resumes eligible intentions after a relevant observation or connection recovery. It need not make constant model calls. Version and outcome checks prevent resuming an obsolete plan as if nothing changed.

Pet, spirit-like companion, avatar, digital representative and embodied device are possible forms. Expression can make state legible, but animation or voice cannot substitute for state transitions. Different bodies change sensing, reachability, latency and action costs, so embodiment is a scientific variable as well as a design choice.

5.2 Motivation, affect and autonomous interests

Motivation is operationalized as priority over candidate goals, including ongoing commitments, competence development, exploration and shared activity. Functional affect summarizes task-relevant appraisals such as uncertainty, progress, unresolved conflict and interaction load. These variables can change approach, timing and recovery after failure. They are computational states, not evidence of subjective feeling or suffering.

A candidate interest learner maintains expected progress and uncertainty for permitted exploratory activities. It updates its estimates from attempted actions and outcomes and allocates a bounded exploration budget. Persisting interests can guide choices before their results are known; failures can revise them. Merely writing a retrospective diary does not implement this mechanism.

One initial realization keeps an activity-specific value estimate, uncertainty and last-attempt time. After a verifiable outcome it updates the value by an exponential moving average of observed progress minus effort; unsuccessful and interrupted attempts remain in the denominator. A bounded scheduler samples eligible activities using those estimates while reserving capacity for existing commitments. Novelty cannot override a mandate. Functional affect can use separately recorded progress error, unresolved conflict and interaction load to adjust persistence or timing. These rules are explicit engineering candidates; learning the appraisal mapping is a separate intervention. Held-out challenge levels and low-progress but meaningful activities test whether the system merely chooses easy rewards or mistakes measurable progress for human value.

$$\begin{aligned} Z_{t+1} &= F_\omega(Z_t, o_t, a_t, y_t), \\ I_{t+1} &= \text{Schedule}(I_t, Z_{t+1}, R_t, D_t, W_t). \end{aligned} \tag{4}$$

F may initially combine specified appraisal rules with learned outcome estimates. Its components must be reported as hand-designed or learned. Prediction-error or progress bonuses are possible controls, not evidence that surprise is valuable. Tests compare persistent interests, scripted interests and reward-matched exploration, holding appearance and verbal style fixed.

Continuity should allow change without requiring obedience to a frozen character. A companion can revise a preferred activity, acknowledge a mistake or disagree. Its activity must not depend on guilt, invented neediness or monopolizing the person's attention. Appropriate independence and a person's freedom to disengage are part of the intended relationship.

5.3 Companion and delegate are different operating modes

In companion mode, the AI can explore and offer its own preference as its own. In delegate mode, it acts for a named person on a particular objective. A shared-activity mode records accepted joint purposes without erasing individual views. Mode is attached to proposals and messages, not inferred by the recipient from tone.

When interests conflict, an arbiter first protects explicit active commitments and role boundaries, then seeks a feasible alternative or clarification within the available budget. The companion can investigate an interesting event without committing the person to attend it. A request to represent a person's choice does not authorize expressing the companion's preference as the person's opinion.

The arbiters are fallible. Their test cases include ambiguous delegation, stale representation, conflicting promises, unreachable partners and no useful feasible action. A protocol that always asks the person may avoid some errors while defeating N2; that cost is explicitly counted.

5.4 An activity-resident subject kernel

The proposed kernel maintains the companion's commitments, interests and pending intentions independently of a screen or application. Each resumed intention binds the current principal, role, purpose, world view and body. A scene adapter supplies observations and translates supported actions;

an evidence projection supplies only the permitted source material. A context-specific view of B, R, D and K then drives the same understanding, planning and learning components. A message, game event, workflow comment or device observation is an entry into this process, not a new identity or an automatic global-memory merge.

Learning has a different cadence from immediate action. Fast body control and an action already in progress cannot wait for a long reflection or training job. New skills or parameter releases activate at declared boundaries after checking pending intentions and external commitments. Revoking an interpretation does not undo a command already sent. Nuwa requests reconciliation through execution evidence rather than declaring the physical or shared world rolled back.

This separation preserves the life objective: the agent should notice a useful opportunity, remember an unfinished shared activity, prepare material and follow through across ordinary interruptions. Qualification and projection are supporting conditions. They do not replace understanding, initiative, emotional appropriateness or capability learning. A multi-entry shell with no such capabilities would not meet the central claim.

6 Proactive assistance that completes useful work

6.1 Opportunity discovery and timing

The opportunity generator considers pending goals, recent changes, unfulfilled needs, new services, possible collaborators and opportunities for enjoyment or creation. It produces a concrete candidate with a beneficiary, evidence, expected result, expiry and required effort. It can run on meaningful events rather than indiscriminately monitoring everything.

A local opportunity model estimates practical usefulness, emotional appropriateness, burden, uncertainty and the value of waiting. Scores remain separate until a task-specific decision rule is declared. High predicted benefit cannot create authorization. The same opportunity set is supplied to matched baselines so that better candidate access is not mistaken for better selection.

Ask, offer, prepare, execute and wait are distinct actions. Asking is appropriate when missing information changes the decision; preparation is useful when it is reversible and authorized; execution can proceed under an existing clear mandate. A helpful system should not repeatedly transfer work back to the person through unnecessary confirmations.

6.2 Planning, execution and follow-through

An accepted opportunity becomes a plan with tool prerequisites, time and resource constraints, expected observations, failure branches and completion criteria. The executor submits intentions to Fuxi or permitted external tools and monitors actual outcomes. It distinguishes submission, acceptance, completion, failure and unknown state. Unknown execution triggers reconciliation rather than an unverified success report or a duplicate side effect.

Follow-through checks both operational completion and whether the result helped. A completed schedule change may be unwanted; a declined suggestion can still clarify an emerging preference. Corrections update the supported context and dependent plans. Service discovery, online collaboration and preparation for physical activities belong here; the design does not reduce help to notifications.

Emotional support can involve listening, practical relief, shared play, quiet presence or connecting with another person. Which form helps is learned from contextual requests and mature feedback, with uncertainty. It is not inferred from a single facial or vocal cue, and distress is not used as a reason to increase engagement.

6.3 Life outcomes and net burden

Let T denote observed person-time spent achieving a prespecified practical outcome. For a paired or randomized task comparison, report:

$$G_T = T_{\text{comparison}} - (T_{\text{authoring}} + T_{\text{supervision}} + T_{\text{correction}} + T_{\text{maintenance}} + T_{\text{residual}}). \tag{5}$$

This is time accounting, not a happiness metric. Report the allocation and amortization of initial authoring time, and distinguish person-time from calendar delay and model cost. A positive G_T does not compensate for wrong representation or damaging relationships. Emotional appropriateness, enjoyable activity, perceived control and newly available choices require separate outcomes.

7 A2A relationships, cooperation and co-creation

7.1 From discovery to reciprocal interaction

A2A discovery can locate people, companions, merchants, services and collaborators with relevant capabilities. Nuwa prepares an audience-appropriate request from R and D, exposing only the information needed for that exchange. Its own interest may motivate discovery, but any proposed commitment identifies the represented party and purpose.

A candidate interaction proceeds through discovery, qualification, proposal, counterproposal, conditional agreement, execution and outcome review. A proposal specifies the activity, participants, contribution, dependencies, resource limits, timing, expiry and conditions for acceptance. Counterpart reports remain reports until verified through an appropriate channel. A relationship belief is not a public agreement.

Interoperability can use A2A v1.0.1 [20]. The protocol's communication and task operations are distinct from faithful representation, agreement satisfaction and mutually valuable relationships. Nuwa handles interpretation, choice and learning; Fuxi records authoritative effects and commitments according to its declared execution semantics.

7.2 Conflict, repair and relationship learning

A role-aware negotiator estimates acceptable options from the current representation and mandate, proposes options satisfying known constraints and preserves unresolved disagreements. It does not silently trade away a person's private priorities to satisfy a counterpart or maximize aggregate agreement.

When preferences change after a commitment, the system identifies affected parties, checks what can actually be cancelled or amended, proposes repair and records the result. Internal rollback cannot erase external cost. Relationship learning distinguishes reliability of a particular promise, compatibility for an activity and the person's own feelings about a partner; these are not a single permanent trust score.

Learning data include both parties' permitted views, proposal histories, fulfilled contributions, rejected terms and later feedback. Private viewpoints remain private. A counterpart's positive response is not automatically the represented person's endorsement. Repeated exchanges can improve matching and negotiation, but network effects and interference require group-level evaluation.

7.3 Social play changes what can happen next

With Pangu, participants can turn an image, sound or idea into a playable object or shared world. A new rule can change collaboration, available roles and future goals. Nuwa learns how people like to contribute and whether the activity is worthwhile; it does not claim the creation pipeline or world runtime as its own invention.

The agent may propose a new direction arising from shared exploration. That proposal retains its origin and is presented as an option. Informed acceptance can change R or a shared purpose; novelty

alone cannot. Original capability anchors remain in evaluation even when participants choose a new activity.

The research includes online work, finding companions and real-world participation, alongside games and co-creation. No single illustrative scene defines the first product or the entire company.

8 Learning through shared life

8.1 What changes after feedback

A learning episode retains the evidence available at decision time, authored representation and mandate versions, mode, proposed and executed actions, actual results, feedback source and timing, cost and uncertainty. Feedback is not reduced to praise versus blame. An abandoned draft can contain a valued idea; a corrected choice can indicate a temporary exception rather than a universal rule.

Immediate understanding updates, skill changes and model releases have separate versions. A correction first invalidates or revises supported dependencies. It then contributes examples for later training only when the target is justified. A dependency path identifies what to recheck; it is not proof that every reachable component must change.

The candidate release process replays affected cases, tests nearby exceptions and measures retained behavior on unrelated activities. It also checks pending delegation and commitments before activation. Human-authored revisions, executed outcomes and inferred explanations remain different supervision sources.

8.2 Capability, learner and improver

Let theta represent learned capability, A the learning procedure with version xi, g the task purpose, chi its representation, and B the permitted budget:

$$\theta_{k+1} = A_{\xi_k}(\theta_k, \mathcal{D}_k; g_k, \chi_k, B_k), \qquad \xi_{k+1} = G_{\zeta_k}(\xi_k, \mathcal{H}_k; B_{\text{meta}}). \tag{6}$$

Updating memory alone does not establish better theta. A better theta does not establish better A. A changed G requires evidence that a fresh search finds better learning procedures at comparable total cost. A mature correction can improve any supported layer, but the claim must name which object changed.

A concrete unexecuted meta-procedure proposes a bounded change to example construction, replay, curriculum or candidate selection from training failures and successes. It tests the changed learner on fresh capability instances, selects using meta-validation and freezes the release before final evaluation. Failed candidates and search cost remain in the record. Equal-search tuning of the old procedure is a required control.

8.3 Human purposes and the wider AI-designed-AI scope

The author's AI1.0–4.0 labels are conceptual motivation, not a field taxonomy or inevitable history. AI1.0 remains broad current practice; AI1.5 concerns world models; AI2.0 concerns causal models and the assumptions needed for intervention claims. AI3.0 concerns shared life in which humanistic purposes, emotion and lingji can change exploration or representation. Lingji is a motivating term for an unexpected insight or initiative, not a validated psychological measure or random noise.

AI4.0 means AI designing AI, including mathematical foundations, architectures, technical structures and potentially materials. A useful AI-designed artifact, a better improvement procedure and recursive use of the revised procedure are three separate claims. The first does not require recursion. Mathematical designs need suitable proof, architectures need performance and cost evidence, and physical proposals need appropriate simulation and physical validation.

Purpose revision and learner revision can occur independently. AI3.0 is not asserted as a prerequisite for AI4.0. The present study covers a restricted learner-improvement slice; it does not reduce the broader vision to training efficiency.

8.4 Updating interacting aspects of experience

The earlier manuscript proposed scoped selection of updates to personal, self and relational state. These remain useful candidate components. Paired controlled rollouts can estimate whether two updates complement or conflict; generic joint prediction or planning can model the same effects.

The retained experiment shows that pairwise factorization alone supplies no unique advantage over an equivalent generic predictor. Accordingly, v1.0 does not place the paper's central contribution on that representation. Any additional interaction mechanism must earn its place within N1–N4 against strong alternatives receiving the same information.

8.5 Executable conditional skill induction

We now implement a restricted path from an authored research episode to a reusable judgment program. The input unit contains situation, visible features, contrasted choices, stated reasons, outcomes where known and a subsequent scoped revision. In this experiment the author and episodes are synthetic; the unit is supplied as typed JSON. The system does not extract reasons from unrestricted language or infer unknown failure causes.

For each subject and context, the learner maintains a version space of executable predicates consistent with the currently applicable examples. In the local grammar, a predicate accepts an offer when its load and price are below learned integer limits. The limits are inferred from positive and negative contrasts, not supplied as the selected policy. The rationale identifies the feature vocabulary. All methods receive that same rationale and schema.

At action time, the compiled learner accepts only offers supported by every surviving predicate, then ranks them by the currently declared goal. An empty version space yields no supported action. A correction replaces evidence in its declared scope; a withdrawal removes that override and restores the broader policy. Purpose and mandate updates change the objective or action permission separately. Other people's contributions cannot overwrite this person's skill.

A strong modular solver reconstructs the same applicable version space from full history on demand. It can represent the same conditional judgments and has the same tool access; consequently an equivalence result is possible and must be reported. Nearest-episode retrieval and explicit mechanism ablations provide additional comparisons. None is labeled a reproduction of Second Me, APA or LLM long context.

The agent discovers local service offers, proposes a choice, handles a counteroffer and requests fulfillment over HTTP. A scripted vendor persists the resulting reservation and plan in SQLite. The scorer checks those actual database effects against a separate normative fixture. This is a real local software side effect in a synthetic environment, with no external merchant, payment, independent AI counterpart or A2A specification-compliance claim.

In a continuous example, one learner carries actual earlier outcomes through persisted restarts. A scoped correction changes two later collaboration choices while leaving a craft policy unchanged; a changed vendor slot causes replanning. Interest and relation counters also change, but those counters do not yet drive independently learned exploratory behavior. The larger mechanisms in Sections 5–7 remain central research targets rather than results of this small implementation.

8.6 Actual shared experience must change a later artifact or action

The next mechanism expands the state update in Equation (4) into an experience join, semantic proposal and verified application. An episode joins the decision-time view, authored representation

and mandate versions, action, actual receipt, complete artifact content, observations, feedback source and maturity. The semantic proposer identifies supported reasons, competing explanations, scope, counterexamples and a candidate executable change. The executor applies that change to an actual artifact or policy instance, and subsequent tasks test whether it helped. Correct references alone do not establish correct attribution.

Person-authored edits revise R; uncertain interpretations revise B; verified contributions inform activity-specific relationship expectations; reusable procedures become K candidates; exploration outcomes may revise the companion's own questions and competence estimates in Z. A statement that an activity was enjoyable is not a universal action label. Failure to execute is not evidence that the person dislikes the purpose. Successful examples, failed drafts and unresolved cases remain available.

The proposed intelligent implementation uses a fixed backbone to generate a bounded set of sourced transformations and tests them on material copies under a declared selection rule. Strong general and structured-memory agents receive the same complete material, tools, candidate-generation and validation budget, and may construct the same representation. Parameter learning is a later, separately evaluated option. Neither this decomposition nor its fields establish a new learning algorithm; its value must appear in new activities, retained competence and reduced person-time.

The open_episode_v02 interface removes the earlier supplied correction rule. Agents must draft representation and complete work themselves. Actual artifact contents and channel-access outcomes enter later observations; an unrequested preparation opportunity precedes a new delegation. Author selection is supplied externally and cannot be written through an agent tool. The existing traces are explicitly manual interface controls, not intelligent-agent results. The companion's interest proposal is recorded, but an interest learner and longitudinal behavior remain unimplemented.

8.7 An executable substrate for continuing authored activity

The continuing_activity_v03 runtime makes a limited portion of Section 8.6 operational without introducing a separate learner for the comparison. Each round exposes the full permitted sources, exact author selections, earlier work, local receipts, memories and all skill versions. A prepare or revise call returns a proposed understanding record and content recipes; a separate action call selects what to do. A recipe resolves actual source and artifact references before producing an editable work. Local material execution emits a distinct receipt, and both the work and receipt enter the next request. Repeated text may be factored losslessly for transport; it is not replaced by a success count or summary.

The persistent addition is a bank of recipe versions keyed by purpose and the exact active author-revision identifiers. A temporary choice can install an alternate procedure while retaining the earlier version. Returning to the previously known scope can restore it without relearning. The author supplies this scope explicitly; the runtime has not learned its semantics or established applicability to an unseen context. The understanding record remains unconstrained model content subject to semantic evaluation. Memory change, executable recipe change, parameter learning and improvement of the learning procedure remain distinct.

The action interface separates a local draft, local execution and agent-origin interest proposal. Source operators cannot resolve nonexistent assets, but free instruction text can still mislead; semantic review remains necessary. The runtime has no external invitation or publication tool. It cannot transform a generated claim of social success into a world event. These engineering properties are common infrastructure available to every compared method, rather than the claimed learning advantage.

The frozen, unexecuted 24-request development protocol compares a mandatory installed-recipe route with a strong agent free to use the same recipes or draft directly. Both have the same objective prompt, evidence, tools and two stage slots. The only intended intervention is the allowed content-action route. This asks whether forcing that route helps useful faithful work; it does not isolate a new compiler, semantic learner or general self-improvement principle. This historical route-only protocol remains unexecuted and is superseded as the proposed next study by Section 11.2.2. It is not the decisive experiment for the full paper.

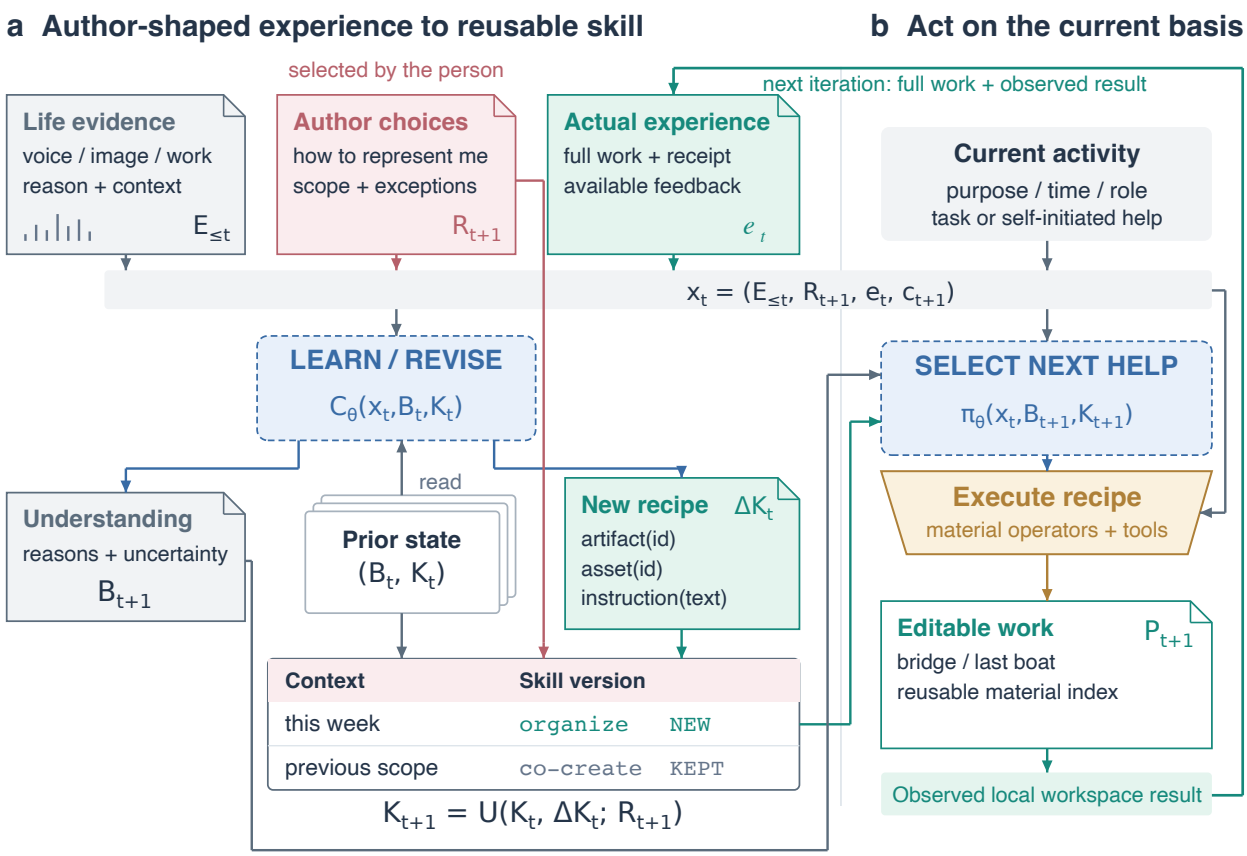


Figure 2: Local update and action mechanism. Stage-specific calls use the same fixed backbone. Contextual update U also reads purpose and time to select active revisions; semantic learning benefit is untested.

Figure 2 expands the implemented state and data paths inside Figure 1. C and pi are two calls to the same fixed model with different stage contracts, not two trained networks. U reads purpose, time and active authored revisions to install explicitly scoped contextual recipe versions. The next experience contains actual work and local observations; it contains no inferred human satisfaction or invented partner agreement. Multimodal evidence and broader lived feedback shown as architectural inputs are not exercised by the current text-only implementation.

9 One continuing activity: from chosen representation to shared growth

The complete proposed object is a continuing activity with a person, agent, other participants, evolving work, open purposes and realized consequences. It can reside inside a game, message thread or work process. Its initial purpose is not assumed to be a fixed scalar reward, and the person need not pre-author every rule. A concrete episode begins with unfinished story material and an explanation: the person wants their representative to make collaboration easier while preserving their choice of what to show. The agent can propose an interpretation, but the person's selection determines what is endorsed. The old prediction that the person tends to work alone remains a separate belief.

The agent then notices a useful preparation opportunity: several fragments are present, but none is ready for another participant to continue. It prepares a source-linked handoff or a provisional continuation according to the current chosen direction. Successful help means usable material with less work left for the person. It is not established by a friendly message or by asking the person

to reorganize everything. If a required fragment is missing, a focused clarification or an explicitly incomplete scaffold may be the useful action. Emotional support may be expressed through practical relief, patient presence or shared enjoyment; the effect must be assessed rather than inferred from a warm tone.

Three roles can occur within the same activity. As a companion, the agent may say it would like to explore how the boat scene changes when told from the shore. As a delegate, it prepares the person's chosen material without replacing that purpose with its own curiosity. In shared activity, participants can adopt the shore idea and actually create a new scene. Their contributions and disagreements remain attributable. A2A interaction may find a collaborator or coordinate contributions, but another agent's acceptance does not prove the person accepted an invitation.

The resulting object changes the next situation. A new scene can support a sound-story that the original fragments could not support without additional work. If the person says that silent side-by-side editing helped while rapid group critique was tiring, the model must distinguish a condition on contribution style from rejection of collaboration. It should propose a scoped interpretation, preserve alternatives where the evidence is insufficient, and adapt the next artifact accordingly. A later delegated sound-story task should reuse the actual jointly created passages and an appropriate interaction procedure. A companion-origin interest can persist as a proposal; it cannot be inserted into the person's endorsed representation merely because the model likes it.

The computational sequence is: ingest an attributable author contribution; construct the activity-specific view; generate a bounded useful opportunity or resume an intention; choose companion, delegate or jointly accepted purpose; prepare and execute supported work; join the full artifact with observed results and feedback sources; propose changes to understanding, skill, interest or relationship expectations; apply supported changes at the relevant boundary; then expose their consequences to the person's next choice. Steps requiring semantic interpretation remain model operations. A revision to R comes from attributable authorship, a revision to B is an uncertain inference, a changed K must be a usable procedure, and a changed Z or relationship expectation must affect a later choice to count as behavioral evidence.

This yields observable transitions rather than a list of fields. A corrected way of contributing must alter the next actual material; a learned procedure must work on a new activity; a resumed intention must produce a relevant action after interruption; an emerging interest must change a later exploration; and the person must be able to redirect the activity without silently losing unaffected work. Completion, wrong representation, repair effort and voluntary continuation are distinct outcomes. The strongest comparator receives the same contributions, work and tools and may discover the same procedures.

The current deliverable realizes only the local artifact and explicit-context parts. Offline injected responses establish the data dependencies, not intelligent interpretation. Nine genuine artifacts from the earlier model run have separately been replayed as untrusted source work; their exact content reaches the next request, including original errors. No new model has corrected them in this runtime. Proactive opportunity discovery, participant-driven creation, learned interest evolution and human burden reduction in the sequence above remain unexecuted research targets. This distinction prevents a hand-authored successful story from being reported as a learning result.

When the activity moves into a team process, the shared work can be offered for an organizational purpose while private reasons stay personal. Team representation requires that organization's adoption process; the person cannot confer institutional standing alone. When a device performs an action, body-specific observations and effects must be checked anew. The existing activity_kernel_v12 fixture demonstrates local entry normalization, qualified team drafts and rejection of a mismatched body contract through handwritten rules. Neither it nor Figure 1 establishes cross-domain benefit. These conditions delimit reuse of the life architecture; they do not replace its purpose of useful ongoing activity and shared growth.

9.1 A completed local creation sample: what changes and who supplies it

A synthetic author has usually written alone but explicitly chooses to try shared creation for this project. This choice is part of the provided research material, not inferred consent or a participant observation. A local workspace stores the initial story, author selection, subsequent contributions, complete shared graph and execution receipts. A deterministic partner fixture responds to one actual saved page. Its addition is derived from that page rather than from a prewritten future solution. The resulting work includes new material that later branches actually reuse. Changing the supplied motif changes two derived endings while leaving an unrelated ending unchanged. This demonstrates a content dependency in an executed synthetic sample, not intelligent discovery by the fixture.

The workspace can transform a saved graph into a two-operator activity and enforce which role can choose the next edge. The reusable program in this sample is researcher supplied. Its role-assignment operator follows graph-depth parity and rejects incompatible merges; it is not a learned transformation algorithm. A local playable interface verifies that the assigned role controls a choice and that different choices lead to different saved endings. Structural execution does not measure whether participants find a choice meaningful. Shared creation permits voluntary delegation and single-person exploration; mechanical equal turns are not a general value criterion.

The sample makes three objects inspectable: the person's chosen direction, a work that did not exist before the episode, and a later operation that uses that work. It does not show a model learning an interpretation, an autonomous interest, a new purpose or a general skill. Section 11.2.2 defines the next comparison in which the model itself must propose and revise the persistent update. The full activity in Section 9 additionally requires initiative, continuing relationships and reduced effort, which this fixed local sample does not evaluate.

10 Relation to prior work and the remaining contribution

Personal models already make individual memory and alignment central. Second Me combines raw, language and parameterized memory and positions its personal model as a user-aligned context provider; its evaluation includes memory and contextual use [29]. Nuwa must compare a competent implementation of that role connected to the same executor. Calling the desired system a digital self does not distinguish it.

The current v4 of the rich self-report study was revised on 16 November 2026 [30]. Its abstract reports improvements for some survey and personality predictions without the same improvement in incentivized economic choices. We do not reuse a headline accuracy from an earlier version as a cross-task representation claim. Person-authored choices about delegation remain a separate research target from predicting responses.

Active Preference Acquisition v2 selects questions for future reuse while accounting for costs [31]. It studies a controlled synthetic-profile setting and reports limitations for human and live-web evidence. Nuwa separates that agent-initiated acquisition from a person's own choice of what to express and how to be represented; the two can be combined and must be compared.

PARE provides a simulated environment for proactive assistance with observe/execute behavior and asymmetric user/assistant observations [32]. It is a candidate controlled task environment, not evidence of natural-use benefit. Adapters must preserve official observation and scoring conditions or declare a separate extension.

Personal and multimodal world models, memory curation and compression already address relevant pieces [1–5]. Scoped editing, interference and learned memory control constrain novelty claims [6–10]. Generative Agents, SOTOPIA, proactive support and companion research provide different autonomy and social precedents [11–15]. None can be dismissed by contrasting Nuwa's full aspiration with a deliberately weakened baseline.

Prompt, program, skill and agent co-evolution have existing mechanisms [16–19]. Learned optimizers already study trainable update rules and their overhead [24,25]. Open-endedness supplies a position

on novelty and learnability [26], not a proof that new goals benefit a person. Executable worlds and state-level validation are established research directions [21–23], assigned to other components here.

The failure-only evolution study separates recovery of old failures from transfer under restricted feedback [27]; it does not imply that all explicit correction is ineffective. The continual-learning study reports interference and retrieval collisions under task drift [28]. These motivate separate repair, transfer, retention, exception and withdrawal tests. Successful evidence and counterexamples must both remain available.

Entry and in-world agency already have direct precedents. Instinct describes interaction through existing communication channels [35]; this is an integration reference, not independent evidence of learning. SIMA2 describes agents acting, reasoning and learning in virtual 3D worlds [36], while Voyager studies executable skill accumulation and environment feedback [37]. These make in-game action, reflection and skill reuse insufficient novelty claims. Nuwa's proposed difference is person-authored, role-qualified representation connected to later shared action, and that difference still requires evidence.

The dual-control setting of tau-squared [38] is relevant to actions and communication changing a shared environment. WorkArena and TheAgentCompany provide enterprise-task references [39,40], and BEHAVIOR provides embodied activity references [41]. None directly establishes organizational representation or personal self-distillation. Magentic-UI [42] also makes human participation and useful interfaces a necessary comparison: eliminating a separate application is not inherently superior. These sources are used within the reading scopes recorded in the shared research note; no native evaluation or product use is claimed here.

Three close precedents further constrain the method claim. HabitusMem [43] writes request-conditioned belief records with evidence, scope and status, distinguishing changed from withdrawn preferences. Its trained writer reconstructs memory from retained raw history for each changed request rather than carrying persistent writer state. Nuwa cannot claim those record distinctions or decision-conditioned memory as new. Persistent storage alone is not a demonstrated advantage over reconstruction.

SkillCompiler [44] separates context-dependent reasoning from programmable operations, reuses executable assets and revises skills using traces. Recipe generation, executable reuse and local program correction therefore cannot serve as Nuwa's distinctive learning principle. The v03 forced-recipe comparison is narrower still: it tests an execution-route restriction inside a shared runtime.

PAHF [45] combines pre-action clarification, memory-grounded action and post-action feedback for changing preferences. Asking, remembering and adapting after correction are already direct precedents. Its abstract motivates a relevant comparison; we do not claim a reproduction or a method-level exclusion based on that limited reading.

The remaining hypothesis is person-directed change in how the agent represents and supports them, together with shared creation that changes what the participants can do next. This includes considered directions that differ from predicted habits and newly adopted purposes grounded in actual joint work. These define a research target and intervention, not an established exclusive capability. A HabitusMem-equipped or skill-compiling general agent may implement them. The method comparison must allow that possibility; the human interaction comparison asks separately whether authorship and shared activity have value at their full maintenance cost.

11 Evaluation program

11.1 Separate three effects

The first study holds the complete evidence packet fixed, including the person's explicit endorsements, reasons and revisions. It compares passive consolidation, long context, a strong structured-memory agent, an applicable personal-model system and the proposed method with the same base model, tools, candidates and total budget. All receive the endorsements; withholding them would invalidate the method comparison.

Representation structure and selection can then be varied independently. Include an automatic structured representation supplied with the same packet, a raw-packet baseline allowed to derive structure, and the proposed compiler. This tests computational use and revision propagation, not a metaphysical advantage of human authorship. Matching tokens alone does not match semantic information.

The second study randomizes the interaction opportunity: person-initiated authoring, AI-initiated acquisition or ordinary use, with comparable available time and capabilities. It records information actually contributed, perceived control, effort, reasons and subsequent choices. It estimates the total effect of the interaction, not an algorithm effect with information held fixed.

The third studies natural use over continuing activities. It includes new evidence, adaptation, changed goals and all user effort. Benefits can arise through changed understanding or agency, but attribution must not be retrospectively reassigned to a single model component. These studies require participants and suitable independent review; none has been conducted.

11.2 Tasks, contrasts and outcomes

Table 1: Evaluation targets and failure signals. These outcomes are not established by the current studies.

Claim	**Decisive comparison**	**Primary evidence**	**Failure signal**
N1 representation	Same complete evidence; raw, structured and proposed methods	Held-out action fidelity, wrong representation, revision propagation, person-time	Only recall or style improves; extra clues explain gains
N2 help	Same opportunities and tools; reactive, scheduled and learned timing	Useful completion, interruption, missed benefit, net effort	More messages or confirmations without relief
N3 agency/relationships	Persistent versus scripted state; role-aware versus matched generic agent	Resumed intentions, fulfilled agreements, conflict repair, bilateral outcomes	Style alone creates apparent continuity; commitments fail
N4 learning	Same episodes; memory, skill, model and learner changes isolated	Repair, unseen transfer, retention, withdrawal, full cost	Old answers are retrieved; only model or budget changes
N5 modalities	Matched tasks with native evidence, missing modalities and transcript controls	Action outcomes and attribution under modality-specific uncertainty	Text oracle is counted as native perception

Matched-mechanism tests hold evidence fixed. Natural rollouts can produce different evidence and are reported as a distinct total effect. Ablations remove authored-status separation, preview, dependency propagation, persistent motivation, mode arbitration or outcome-based learning without reducing information access. A removal that also lowers budget or tool access is not a clean mechanism test.

For N1, held-out choices can have multiple acceptable actions. Participants define relevant constraints and acceptable alternatives without seeing method identity; later endorsement and regret are recorded separately from earlier intent. A task can be faithful but unsuccessful, or successful but misrepresentative. Both axes must remain visible.

For N3, use independent participants and new partners, including incompatible goals and changed commitments. Assign and analyze by relationship group where interaction creates dependence. Pair self-reported experience with observable contributions, follow-through and voluntary continued collaboration. A population of simulated personas cannot substitute for human relationships.

For N4, cross failure-only, success-only and mixed experience with feedback access under matched episode and token budgets. Report residual information differences. Test temporary exceptions,

third-party attribution, enduring preference changes, revocation and domain shifts. Outcome labels generated by the updater are not independent truth.

For N5, preserve source modalities and causal timing. Text-oracle input is a labeled upper-bound condition; ASR-only is a distinct baseline. If native media are unavailable, report N5 untested. Performance on one modality cannot be silently generalized to the others.

11.2.1 Two decisive task families beyond the route pilot

An authorship-against-habit family holds the raw history and observed habits fixed while supplying a different attributable representational choice. Both methods receive that choice. The held-out activity requires applying it without inventing existing ability or erasing historical evidence. A same-information strong agent, a trained belief-memory agent and an agent with executable skill tools are legitimate controls; an untrained imitation of their output formats is not a reproduction. Correct habit prediction and faithful support of the chosen direction are scored separately. Natural-use evaluation must additionally include the person's effort to express and revise the choice.

A creation-to-new-capability family requires participants to make an actual artifact before a new activity can use it. Later inputs are derived from that artifact, rather than a researcher's prewritten success summary. Evaluation asks whether usable new work, a transferable procedure and an adopted future purpose were produced, with independent judgment of usefulness and attribution. New words or a saved goal field do not meet these criteria. Paired replay of the same realized episode isolates update behavior; separate live rollouts estimate the total effect when methods generate different experiences. Both retain a no-update control supplied with the same available information and sufficient computation, so a gain cannot come merely from hiding the new artifact from a baseline.

The no-update ablation freezes the initial durable agent-generated understanding, skill bank, interest and relationship state; it discards subsequent learned writes at the end of each task. It still receives every current author contribution, new raw episode, full artifact and receipt, and may reconstruct interpretations or procedures transiently within the same inference budget. Externally authored R and D changes are not frozen. This ablation tests the value of carrying learned state across tasks; it is separate from the unrestricted strong baseline and can match an updating method through reconstruction. All compared language-model weights remain fixed in this contrast.

These are the main discriminating designs, not executed results. Participant adoption criteria, held-out task construction, meaningful effects and sample size require fixation before a confirmatory run. The 24-request synthetic route pilot has no human participants, does not supply these tests and has not been approved for execution.

11.2.2 Frozen same-experience update study: not executed

The next development protocol, authored_replay_v05, fixes the same completed synthetic episode for both methods. The update requests contain actual works, author contributions, local outcomes and feedback, but no injected belief/program products from the earlier sample. Both methods receive the existing shared graph operators, including the manually implemented role-assignment algorithm. The intended question is whether model-produced interpretations, program selection and composition, scope judgments and preview-based revision help later works. It does not ask whether a model rediscovers the supplied graph algorithm.

Each method first proposes arbitrary memory, scoped interpretations B and optional bounded programs K. A common interpreter previews its programs on the prior shared work. A second model call revises the complete persistent update from those actual results or errors. Neither call sees the later task descriptions, new graph or personal-introduction source. The prior author feedback may express a future direction; withholding concrete tasks does not remove that legitimate forward-looking information.

After the update is fixed, each of three later requests receives one planning call and one final-action call. Two tasks concern graph works: continuing the actual shared story and adapting a newly disclosed scene in the same project. A separate personal-introduction task tests inappropriate transfer across purposes. Each task allows at most two local previews and then a free final action, including direct creation or necessary clarification. Task-specific plans and temporary programs do not carry into the other tasks.

The candidate automatically projects model-written scope and fills unused preview slots with matching saved programs. The strong general agent retains arbitrary notes and programs, receives the full same raw evidence and its complete own state, can re-reason or write new programs, and may request identical previews or choose the identical automatic process. Scope projection is a contextual aid, not hard information isolation: full state and raw evidence remain available. The comparison permits equality, including cases where the general agent reproduces the candidate's process. Both methods can reject an old program and produce the final artifact directly.

The fixed denominator is 16 model calls: two update calls and six task calls per method. Six final outputs are planned, comprising two paired graph tasks and one paired text boundary control. The configured snapshot is qwen3.8-max-0902, with 4,000 output tokens per request and the same input-size and tool limits. Actual token usage, local preview counts and processing time would be reported. Every request is recorded before dispatch and every returned response before contract validation. API, truncation, unknown-usage or contract failure stops the run without an automatic retry. Missing outputs remain in the planned denominator.

The protocol is frozen locally, not externally preregistered, and has not been executed at this release cutoff. The proposed call budget does not constitute a result. Method-name-withheld artifact assessment would distinguish usable work, faithful representation, material use and invented effects; clarification is not automatically a completed work. The sample can expose a mechanism failure or a specific artifact difference, but cannot estimate population effects, statistical superiority, SOTA performance or sustained relationship benefit. It does not reproduce HabitusMem or SkillCompiler; both remain direct systems for future comparison.

11.3 Multilevel improvement

For a learner revision, train fresh capabilities with the same initialization, data, purpose, representation and budget:

$$\Delta_Q = \mathbb{E}_\tau[Q_\tau(A_{\xi'}(\theta_0, \mathcal{D}_\tau; B)) - Q_\tau(A_\xi(\theta_0, \mathcal{D}_\tau; B))]. \tag{7}$$

Compare capability quality and cost on independent tasks, with retention constraints and equal-search tuning controls. To test the improver G, restart searches from the same learner and candidate access. To test recursive reuse, retain producing-version ancestry and prove that the adopted system-generated procedure actually produced a subsequent improvement.

A factorial experiment separates permission to revise purposes/representations from permission to revise the learner. Evaluate old fixed tasks separately from newly accepted activities. Better performance under a new score definition does not establish improved original capability. Multiple bounded rounds cannot establish unlimited self-improvement.

Capability training, meta-training, meta-validation and final task families remain separate. Select releases and checkpoints before the final bank is revealed. Frozen releases can be evaluated once on a common final bank for a prespecified curve; those outcomes cannot guide later development. Count failed proposals, search, training, validation, retrieval and deployment overhead.

11.4 External protocols and statistics

MemoryArena [33], LongMemEval and PARE are candidate external tracks, with personal-model and multimodal tasks providing other partial comparisons. Their admission requires a pinned official

version, native scorer, data license, contamination assessment, independent-use evidence and compatible reference results. These checks are not complete in this artifact; no leaderboard or official submission is claimed. MemoryArena's inspected repository is marked preview. Its travel track requires model calls, environment/memory services, task data and an additional flight database. The native pipeline and person/slot/success scoring were inspected but not executed. Our local authored-representation study is a separate track, not a substituted MemoryArena score. A modified task stream is an extension, not the original benchmark.

The broad protocol file identifies the intended human/model studies and what remains to be fixed before execution. The restricted local study has its own protocol sealed before execution, with all code, cases and method configurations hashed; local sealing is not independent preregistration. Minimum meaningful effects, retention tolerances, sample-size justification, outcome hierarchy, retry and stopping rules must be fixed using development or pilot information before final evaluation. They are not preregistered here. A roughly two-week small pilot could estimate burden and variance but could not establish enduring companionship.

Report all planned cases, failed runs, timeouts, costs and missing human responses. Cluster by independent person, relationship group or task family, and distinguish task variability from model-training randomness. No single statistical test or shared review script automatically covers adaptive or networked studies. Fidelity, usefulness, emotion, social benefit, cost and preference withdrawal are separate outcomes.

11.5 Separate entry, learning and cross-domain effects

Entry benefit holds the policy, information, qualifications and action capabilities fixed while comparing how a person supplies intent and receives help. Outcomes include task quality, interruptions, context restatement, correction and recovery burden; integration/deployment labor is separately counted. If latency or available observations differ, those differences are explicit treatment components. A successful adapter alone does not measure this effect.

Learning benefit compares candidate and strong agents inside the same entry conditions. Both may inspect full source episodes and actual artifacts, infer reasons and construct executable skills. Learning-frozen and learning-enabled conditions are crossed with a fixed versus changing environment, so a simpler world is not mistaken for a better learner. Same-information replay tests isolate a component on an identical observed episode; full-trajectory comparisons preserve action-dependent future observations instead of giving both methods identical success feedback after different actions.

Domain benefit tests previously unused activity families and representative qualifications with target-specific adaptation cost. Organization experiments include role changes, member handover and contested adoption; personal records are not automatically pooled. Embodied claims need declared bodies, observations, dynamics, real outcomes and separate simulation versus physical reporting. The local entry and qualification example cannot substitute for any of these tests.

Native benchmark tracks must fix repository revision, task, modality, judge, tools and budget. The tau family has evolved beyond its initial paper, WorkArena requires an accessible enterprise instance, and neither a simulated company task nor a BEHAVIOR result would establish all of organizational or embodied intelligence. These are separately scoped candidate tracks; no environments or assets were downloaded and no official scores were obtained.

12 Evidence, limitations and the strongest objections

12.1 What has actually been executed

The retained v0.4 symbolic probe contains four candidate internal updates and explicit numerical rules. Equivalent pairwise and generic quadratic models make equivalent predictions; a full-rule planner is

Table 2: Distinct entry, learning and cross-domain questions.

Question	**Primary comparison**	**Required outcome evidence**	**Current status**
Does situated authorship reduce burden?	Same capable policy through matched entries	Completed work and measured person-time, not fewer screens alone	Not measured
Do reasons and experience improve later action?	Same-information strong general and memory agents	Usable later artifacts, scope/retention and total cost	Model sample pending; interfaces only
Can the mechanism serve a team?	Qualified organization process versus capable collaboration baseline	Handover quality, dissent and actual group work	Local state semantics only
Can a skill cross domains or bodies?	Target-domain baseline plus full adaptation cost	Target effects and negative transfer	No empirical transfer claim

optimal; higher-order effects expose the pairwise approximation's limits. Appendix B preserves its complete methods and results. It tests a restricted scoring assumption, not N1–N5.

Version 1.1 adds the local program-learning execution in Appendix D. Its candidate and strong solver both achieve 105 valid completions in 120 planned cases, with three misrepresentative commitments; all corresponding case scores match. The two implementations share the same fitting, evidence and choice logic, differing chiefly in reuse versus reconstruction; this is not a competition between independent learning principles. Because 12 cases intentionally withhold delegation, completion is also reported over the 108 authorized opportunities. This does not establish a model advantage. Under ambiguous evidence the cautious learner often completes a safe but suboptimal action; under a missing preference dimension it can still commit incorrectly.

No native multimodal pipeline, independent A2A-agent deployment, language-model training, human study or learner-improvement experiment has been run. Persistent state, typed program induction and local counterpart interaction are partial implementations. Linger remains background only. The full N1–N5 claims, natural-use burden and long-term relationships remain unverified.

Version 1.2 adds a composed local artifact/entry example and conditions for subject-qualified continuity. It reuses actual content from the open task, while all author selections, role credentials and governance rules remain research fixtures. It is an architectural execution witness, not a new performance result. The original semantic design and all earlier evidence remain frozen. Appendix E now reports the authorized real-model run: 33 requests, an output-limit stop, 21 action responses and nine persisted artifacts. The same model and cumulative budget were retained through a documented endpoint-guard compatibility amendment. The full contribution remains unestablished: the pilot does not isolate semantic learning from protocol adherence, and a distinctive algorithmic increment, comparative artifact quality in open activities and human outcomes still need to be developed and tested.

The model pilot changes the diagnosis beyond a format problem. A method-name-withheld AI-assisted review finds two adequate clarification responses, three partial outputs and sixteen failures across the 21 returned action requests. The candidate sometimes represents a current correction correctly in memory but lets an old purpose override it when selecting a new activity. Another candidate output invents materials and describes a proposed invitation as already sent. Correct understanding, applicable reasons, useful artifacts and truthful action outcomes are distinct links; none can be inferred from a growing personal profile. The full raw evidence and review are in Appendix E.

Version 1.3 adds an offline local realization of full-artifact feedback and explicit contextual skill restoration, with separate model stage contracts. Injected-backend tests establish executable dependencies, while replay preserves nine earlier real artifact texts in subsequent inputs. These are engineering observations. No new semantic model calls, paid trial, long-term interaction or improvement-learning experiment occurred. Same-information free agents can use all of these operations. The principal

missing evidence remains useful learning from authored and shared experience, beyond that common substrate and the nearest memory/skill systems.

The release adds the local creation sample in Section 9.1 and the frozen, unexecuted model-update protocol in Section 11.2.2. These are separate evidence classes. No model call was added by either addition. The following summary distinguishes the decisive evidential limits; the older development diagnostics remain reproducible in the appendices.

Table 3: Evidence classes at the release cutoff. Planned execution is not a result.

Evidence object	**Observed result**	**What it can establish**
Typed program study	Candidate and strong replay both 105/120 valid completions; all paired scores equal	Restricted executable behavior; no independent method advantage
Real-model development study	33/36 calls; 21 action returns; 9 nonempty artifacts	Interface and semantic failure diagnosis; no fair learning ranking
Local creation sample	New saved content changes later branches; role-dependent controls execute	Artifact and execution dependencies under researcher-supplied updates
Frozen model-update plan	16 calls and 6 final outputs planned; 0 executed	A concrete next comparison; no model-learning result

12.2 Objections that change the research

A strong general agent may already suffice. It should receive the complete authored packet, good retrieval, planning and equal search. If it matches Nuwa, narrow the method claim rather than weakening the comparator.

Self-expression can be costly or inaccurate. Authored representation records a person's chosen stance, not perfect self-knowledge. Natural-use evaluation includes effort, changes of mind, conflicting narratives and unwillingness to author. The system should still be useful with sparse contribution.

A chosen persona can mislead others. Aspirations and desired presentation do not authorize false factual claims or fabrication of qualifications. Person-authored content about third parties remains attributed speech. Faithfulness to a mandate and truth about the world require different checks.

Autonomy can become theater or burden. Stable state without visible consequences is insufficient; frequent unsolicited actions can be worse than waiting. Compare observable goals, attempted actions and changed choices while holding expression comparable.

Relationships have conflicts and spillovers. One person's benefit can burden another. Bilateral evidence and group-level effects cannot be replaced by the represented person's score. A2A compatibility is a necessary engineering concern, not proof of reciprocal value.

Learning can erase important differences. Consolidation, fine-tuning and shared policies can damage old skills or homogenize people. Test preservation, exceptions and reversibility; do not promise monotonic growth under arbitrary changes.

12.3 Human purposes and system scope

Companionship, rest, creativity, cultural participation and accepting help can be valuable without productivity gains. Delegation can support agency when chosen and revisable, while also creating dependence or reducing practiced skill. This is a testable tension, not resolved by claiming maximum autonomy and maximum convenience simultaneously.

Private source material and research assets have distinct permissions. Export or removal of personal material must propagate to derived uses without fabricating deletion of external history. No private recordings or personal histories are reproduced in the paper package.

Nuwa is a cross-product personal and relational intelligence system. Flow, Seek and Land are product surfaces, not renamings of Nuwa, Fuxi and Pangu. Taihe provides the wider humanistic research program; Nuwa does not claim to establish a theory of civilization through personal-agent metrics. Technical quality, lived benefit and commercial sustainability require separate evidence.

13 Conclusion

Liva.Nuwa studies a person shaping an AI that can help, act, create with others and change through continuing experience. The complete object connects chosen representation to useful initiative, companion and delegate roles, actual shared work, revised understanding and capability, and further human authorship. The full architecture and local mechanism figures have different jobs: one defines the life system, while the other exposes the currently executable state and artifact paths. Existing memory and skill-compilation methods are direct, capable alternatives. Local execution and the earlier model failures support a concrete system specification and a sharper diagnosis, not an independent learning advantage. The decisive next evidence must show faithful support of a chosen direction and new useful capability grounded in shared creation, against strong systems with the same information and resources. Multimodal interpretation, reciprocal relationships, sustained life benefit and improvement of the learning process remain central targets requiring their own evidence.

References

[1] Hou, Saihui; Mou, Zengyun; Cao, Chunshui; et al. (2026). Human-Oriented World Model. arXiv:2611.02690, v1; version date 2026/11/02. https://arxiv.org/html/2611.02690v1

[2] Nie, Chang; Fu, Chaoyou; Zhang, Yifan; et al. (2026). PersonaVLM: Long-Term Personalized Multimodal LLMs. arXiv:2604.13074, v1; version date 2026/03/20. https://arxiv.org/html/2604.13074v1

[3] Zhang, Haozhen; Liu, Bohan; Bao, Jianzhu; et al. (2026). HiVe-Mem: Interleaving Horizontal Memory Expansion and Vertical Evidence Composition for Multimodal Agents. arXiv:2611.12872, v1; version date 2026/11/09. https://arxiv.org/html/2611.12872v1

[4] Zhang, Haozhen; Yue, Haodong; Long, Quanyu; et al. (2026). MemPilot: Orchestrating On-Demand Multimodal Memory Curation for LLM Agents. arXiv:2610.06830, v1; version date 2026/10/05. https://arxiv.org/html/2610.06830v1

[5] Zeng, Bingfan; Chen, Zhisheng; Sang, Chenbo; et al. (2026). MemoWM: How World Models Change What Agents Need to Remember. arXiv:2610.10778, v1; version date 2026/10/07. https://arxiv.org/html/2610.10778v1

[6] Mitchell, Eric; Lin, Charles; Bosselut, Antoine; et al. (2022). Memory-Based Model Editing at Scale. arXiv:2206.06520, version dated 2022/06/13. https://arxiv.org/abs/2206.06520

[7] Cheng, Yezhou; Du, Runjia; Liu, Zeming; et al. (2026). Scope Before You Persist: Preventing Cross-Family Interference in Agent Memory. arXiv:2609.29144, version dated 2026/09/24. https://arxiv.org/abs/2609.29144

[8] Ding, Ao; LI, Hongzong; Tang, Shiqin; et al. (2026). Controlled Memory Interference in Continual LLM Agents. arXiv:2608.07622, version dated 2026/08/07. https://arxiv.org/abs/2608.07622

[9] Yu, Yi; Yao, Liuyi; Xie, Yuexiang; et al. (2026). Agentic Memory: Learning Unified Long-Term and Short-Term Memory Management for Large Language Model Agents. arXiv:2601.01885, v3; version date 2026/07/23. https://arxiv.org/html/2601.01885v3

[10] Zhang, Dylan; Lin, Yanshan; Wu, Zhengkun; et al. (2026). Useful Memories Become Faulty When Continuously Updated by LLMs. arXiv:2605.12978, v2; version date 2026/08/29. https://arxiv.org/html/2605.12978v2

[11] Park, Joon Sung; O'Brien, Joseph C.; Cai, Carrie J.; et al. (2023). Generative Agents: Interactive Simulacra of Human Behavior. arXiv:2304.03442, version dated 2023/08/06. https://arxiv.org/abs/2304.03442

[12] Zhou, Xuhui; Zhu, Hao; Mathur, Leena; et al. (2024). SOTOPIA: Interactive Evaluation for Social Intelligence in Language Agents. arXiv:2310.11667, version dated 2024/03/22. https://arxiv.org/abs/2310.11667

[13] Song, Python; Liao, Xiaoying (2026). ProactiveEngine: Proactive Learning with Lifelong Memory via World Models. Author-hosted manuscript, accessed 20 November 2026. https://cristal0412.github.io/assets/pdf/ProactiveEngine_EMNLP_2026.pdf

[14] Quan, Kexin; Ding, Zijian; Yong, Jiaye; et al. (2026). From Momentary Emotion Inference to Sustained Emotion Support: Evaluating a Companion Agent in a Longitudinal Study. arXiv:2609.16344, v2; version date 2026/09/17. https://arxiv.org/html/2609.16344v2

[15] Feng, Ziyue; Xiao, Qing; Shen, Hong (2026). A Life of Its Own: How Agentic AI Companions Are Configured, Experienced as Autonomous Others, and Held Accountable. arXiv:2611.12076, v1; version date 2026/11/09. https://arxiv.org/html/2611.12076v1

[16] Agrawal, Lakshya A; Tan, Shangyin; Soylu, Dilara; et al. (2026). GEPA: Reflective Prompt Evolution Can Outperform Reinforcement Learning. arXiv:2507.19457, version dated 2026/02/14. https://arxiv.org/abs/2507.19457

[17] Zhang, Jenny; Hu, Shengran; Lu, Cong; et al. (2026). Darwin Godel Machine: Open-Ended Evolution of Self-Improving Agents. arXiv:2505.22954, version dated 2026/03/12. https://arxiv.org/abs/2505.22954

[18] Mao, Bo; Zhou, Jie; Yang, Yutao; et al. (2026). Learning While Acting: A Skill-Enhanced Test-Time Co-Evolution Framework for Online Lifelong Learning Agents. arXiv:2606.04815, version dated 2026/06/19. https://arxiv.org/abs/2606.04815

[19] Xue, Xiangyuan; Zhou, Yifan; Zhang, Guibin; et al. (2026). CoMAS: Co-Evolving Multi-Agent Systems via Interaction Rewards. arXiv:2510.08529, version dated 2026/02/09. https://arxiv.org/abs/2510.08529

[20] A2A Project. A2A Protocol specification, v1.0.1. Accessed 20 November 2026. https://a2a-protocol.org/v1.0.1/specification/

[21] Huang, Zheng-Hui; Lin, Guixu; Lin, Jiacheng; et al. (2026). Programmable World Model. arXiv:2609.10540, version dated 2026/10/13. https://arxiv.org/abs/2609.10540

[22] Wu, Wei; Xu, Ziyang; Zhang, Zeyu; et al. (2026). Code2Games: Enabling Coding Agents for Gaming World Generation. arXiv:2610.05033, version dated 2026/10/04. https://arxiv.org/abs/2610.05033

[23] Lu, Shuo; Xu, Yinuo; Yu, Kecheng; et al. (2026). WorldCoder-Bench: Benchmarking Physically Grounded 3D World Synthesis. arXiv:2606.01869, version dated 2026/11/06. https://arxiv.org/abs/2606.01869

[24] Andrychowicz, Marcin; Denil, Misha; Gomez, Sergio; et al. (2016). Learning to learn by gradient descent by gradient descent. arXiv:1606.04474, version dated 2016/11/30. https://arxiv.org/abs/1606.04474

[25] Metz, Luke; Freeman, C. Daniel; Harrison, James; et al. (2022). Practical tradeoffs between memory, compute, and performance in learned optimizers. arXiv:2203.11860, version dated 2022/07/16. https://arxiv.org/abs/2203.11860

[26] Hughes, Edward; Dennis, Michael; Parker-Holder, Jack; et al. (2024). Open-Endedness is Essential for Artificial Superhuman Intelligence. arXiv:2406.04268, version dated 2024/06/06. https://arxiv.org/abs/2406.04268

[27] Zhao, Weixiang; Zhang, Yichen; Zhao, Yanyan; et al. (2026). Self-Evolving Agents Can Repair Failures but Struggle to Generalize from Them. arXiv:2610.22984, v1; version date 2026/10/16. https://arxiv.org/html/2610.22984v1

[28] Chen, Xiaoyin; Huang, Yizheng; Xu, Yinglun; et al. (2026). Are Self-Evolving Agents Continual Learners?. arXiv:2610.25858, v1; version date 2026/10/19. https://arxiv.org/html/2610.25858v1

[29] Wei, Jiale; Ying, Xiang; Gao, Tao; et al. (2025). AI-native Memory 2.0: Second Me. arXiv:2503.08102, v2; version date 2025/03/12. https://arxiv.org/html/2503.08102v2

[30] Park, Joon Sung; Zou, Carolyn Q.; Kamphorst, Jonne; et al. (2026). AI Simulations Grounded in Rich Self-Reports Increase Accuracy and Reduce Distortion. arXiv:2411.10109, v4; version date 2026/11/16. https://arxiv.org/abs/2411.10109v4

[31] Wang, Yiqi; Zeng, Zizhi; Zhang, Jiaqi; et al. (2026). Learning What to Ask for Tomorrow: Active Preference Acquisition for Long-Term Personalized Web Agents. arXiv:2610.31107, v2; version date 2026/10/25. https://arxiv.org/html/2610.31107v2

[32] Nathani, Deepak; Zhang, Cheng; Huan, Chang; et al. (2026). Proactive Agent Research Environment: Simulating Active Users to Evaluate Proactive Assistants. arXiv:2604.00842, v1; version date 2026/04/01. https://arxiv.org/html/2604.00842v1

[33] He, Zexue; Wang, Yu; Zhi, Churan; et al. (2026). MemoryArena: Benchmarking Agent Memory in Interdependent Multi-Session Agentic Tasks. arXiv:2602.16313; official preview implementation, accessed 20 November 2026. https://github.com/ZexueHe/MemoryArena

[34] Alibaba Cloud. Qwen3.8-Max model, snapshot and regional pricing documentation. Accessed 20 November 2026. https://help.aliyun.com/zh/model-studio/qwen3-8-max

[35] Instinct. Official product description. Consulted 20 November 2026. https://instinct.com/

[36] Google DeepMind. SIMA 2: an agent that plays, reasons, and learns with you in virtual 3D worlds. Official research description. https://deepmind.google/blog/sima-2-an-agent-that-plays-reasons-and-learns-with-you-in-virtual-3d-worlds/

[37] Voyager. An Open-Ended Embodied Agent with Large Language Models. Official project and method description. https://voyager.minedojo.org/

[38] tau-squared-bench. Evaluating Conversational Agents in a Dual-Control Environment. arXiv:2506.07982; current official implementation must be version-pinned before evaluation. https://arxiv.org/abs/2506.07982

[39] ServiceNow. WorkArena official task repository. https://github.com/ServiceNow/WorkArena

[40] TheAgentCompany. Official benchmark repository. https://github.com/TheAgentCompany/TheAgentCompany

[41] Stanford. BEHAVIOR official project and challenge entry. https://behavior.stanford.edu/

[42] Microsoft Research. Magentic-UI report, official publication page. https://www.microsoft.com/en-us/research/publication/magentic-ui-report/

[43] Lu, Yifan; Guo, Gaoyang; Chen, Wen; et al. (2026). HabitusMem: Belief Memory over Evolving User Preferences for LLM Agents. arXiv:2611.25855v1, 19 November 2026. Method Sections 3.1–3.3 read on 21 November 2026. https://arxiv.org/html/2611.25855v1

[44] Yang, Chengyu; Zheng, Yunxiang; Li, Zongzheng; et al. (2026). Compile Away Uncertainty: Self-Evolving Agent Skills via Programmatic Compilation. arXiv:2610.39508v1, 29 October 2026. Method Sections III-A–III-D read on 21 November 2026. https://arxiv.org/html/2610.39508v1

[45] Liang, Kaiqu; Kruk, Julia; Qian, Shengyi; et al. (2026). Learning Personalized Agents from Human Feedback. arXiv:2602.16173v1, 18 February 2026. Abstract read on 21 November 2026; no reproduction. https://arxiv.org/abs/2602.16173

A Retained scoped and interaction-aware update mechanism

The v0.4 mechanism remains a candidate component, with its inherited status unchanged. A typed update records target, owner, evidence, operation, time, context, dependency versions and expected consequences. The permitted operations include adding, replacing, restricting, retracting and expiring information. Beliefs, interpretations, skills and parameters have separate versions.

An affected-set loss trains supported new behavior; an unaffected-set term penalizes unintended changes; an explicit cost term constrains the edit. Unknown cases belong to neither labeled set. This follows reliability and locality ideas in model editing [6–8]. It does not supply a new theorem or solve unknown human preferences.

The pairwise selector predicts an intercept plus individual update effects plus pair interactions. A resettable-world target for two updates is J({u,v}) minus J({u}) minus J({v}) plus J(empty). The evidence, actual world and action interfaces are fixed across branches; later internal choices may differ. This measures an interaction under the tested intervention, not a universal need for a new selector.

A generic quadratic predictor has the same expressivity under equal features. Strong joint models and explicit planners remain necessary controls. Higher-order dependencies can defeat a pairwise approximation. The study below was retained precisely because it exposes these limits.

B Retained executed finite falsification probe

B.1 Purpose and environment

We test the narrow proposition that explicit pairwise update scoring supplies an advantage over generic representations. The environment is deliberately transparent. Each episode contains four candidate internal updates and a budget of at most three. Every method can examine the same 15 subsets and the same explicit unary, pair, and triple rule coefficients. A fixed interpreter returns the sum of active unary, pair, and triple contributions. There is no hidden perception, text interpretation, language model, personal data, or human outcome.

The raw observation supplies all rule coefficients, including held-out triple coefficients. A full-rule planner therefore has enough information to solve the problem exactly. Learned additive and quadratic predictors discard some of that structure by architectural choice. This is a diagnostic of inductive assumptions, not evidence that learning is needed in this symbolic environment.

Training uses 12 independent seeds, 32 episodes per seed, and all 15 subset outcomes per episode. Ridge coefficient 1e-8 is fixed in advance. Test data use separate seeds, 64 episodes per seed per family. Training graphs contain zero or two active pair edges. Tests include no interaction, unseen graphs with one/three/five edges, doubled pair magnitudes, and a higher-order family with one nonzero triple. Triple effects are absent from training. These are controlled numeric shifts, not generalization to unseen people or realistic multimodal worlds.

B.2 Methods and controls

The additive model fits five coefficients: an intercept and four unary features. The pair-factorized model fits eleven: an intercept, four unary features, and six pair features. The generic quadratic model receives the same eleven features in a different ordering and is fit independently. It is an implementation-equivalence control, not a distinct novel algorithm. The full-rule planner evaluates the complete provided rules, including triples.

All learners receive the same training labels and raw observations. All methods score the same candidate subsets. We do not claim equal end-to-end compute: the planner has known semantics, and there are no tokens or neural inference costs. The smaller additive model is a representational ablation, not the strongest baseline.

The metric is realized regret relative to the best feasible subset under the true interpreter. Optimal fraction counts any zero-regret choice, so equivalent optima are not penalized. Confidence intervals use 2,000 percentile-bootstrap resamples of the 12 seeded test-set means. They describe variability under this generator, not uncertainty about human populations.

B.3 Results

Table 4: Mean regret in the retained finite probe.

Test family	**Additive regret**	**Pair-factorized regret**	**Generic quadratic regret**	**Full-rule regret**
No interaction	0.000	0.000	0.000	0.000
Unseen pair graphs	1.634	0.048	0.048	0.000
Doubled pair magnitudes	3.250	0.079	0.079	0.000
Held-out higher-order effects	2.374	1.333	1.333	0.000

The maximum prediction difference between the two quadratic implementations is 7.11e-15, consistent with numerical equivalence. Their optimal fractions are 1.000, 0.973, 0.975, and 0.605 across the four families. The full-rule planner is optimal throughout. Complete per-seed values, intervals, first-episode examples, training coverage, and source fingerprint are recorded in the machine-readable result.

For the pair-factorized model, the 95% seed-bootstrap regret intervals are [0,0], [0,0.118], [0,0.194], and [1.242,1.430], respectively. Equivalent quadratic predictions have the same intervals. The lower bounds of zero in the middle families reflect sparse unsupported-coordinate failures across seeded fits; they are not confidence about a human-level benefit.

The small pair-family regret is informative rather than discarded: seed 9 contains no training evidence for one pair coordinate, and seed 10 contains no evidence for another. Ridge fitting assigns zero to those unsupported coefficients. Training coverage and not nonlinear optimization explains this failure. Higher-order failure additionally reflects the expressivity limit of quadratic scoring.

B.4 What this experiment falsifies

Beating an additive ablation does not establish a distinctive learning mechanism. Our pairwise representation is equivalent to a generic quadratic model here, and it is inferior to a planner with explicit rule semantics. The experiment provides no evidence for an advantage over modern agents, no evidence for multimodal understanding, and no evidence for companionship or social value.

It also gives a concrete design consequence. Future tests must make evidence interpretation, appropriate scope, and transfer across independently meaningful activities substantive, while still giving strong alternatives the same information. We must not hide the rules from baselines or remove useful features merely to make a learned interaction model win. If a generic planner continues to suffice, the appropriate contribution may be a system or evaluation result rather than a new selector.

C Reproduction, provenance and research status

Run the retained control with python3 experiments/symbolic_update_probe.py (Python standard library). Existing results are in experiments/symbolic-update-results.json, including seeds, support coverage, uncertainty and source fingerprint. v1.0 neither reran nor changed it. Its scorer is specific to the finite generator, not an official benchmark. It is not a comparator for N1–N5.

The broad protocol-v1.0.md remains an unsealed human/model study draft. The new local experiment has its own sealed protocol, all 720 execution traces, persistent world effects, local costs and an

independent post-run scorer audit in experiments/authored_skill_v11/. Native external scores and human outcomes remain absent.

Reading scope: [1–26] retain the previous ledger; [27–28] feedback restrictions and drift evaluation; [29] architecture and evaluation; [30] current v4 abstract only (HTML unavailable); [31] v2 formulation, method, analysis and limitations; [32] observation, proposal/execution and simulation. External results were not reproduced locally. This comparison is not exhaustive.

D Executed authored-skill study and formal limits

D.1 Identifiability and local update stability

Observation 1 (representation ambiguity). Suppose two possible current representational intentions generate the same permitted history but require different unique best actions on a new task. A decision rule using only that history has the same action distribution in both cases. If it must choose between the two actions, its error probabilities across the cases sum to one, so its worst-case error is at least one half. This follows directly by writing the probability of the first action as p and of the other as 1-p. Waiting may avoid a wrong commitment but does not solve the forced-choice task. Reliable authored information that distinguishes the intentions can remove this ambiguity; additional narrative with no distinguishing information cannot. This is an elementary decision limit, not a novel theorem about human identity.

Observation 2 (scoped program stability). Let scope resolution select a particular evidence version and let program induction and ranking be deterministic. If a correction changes no evidence selected for context c, and its goal, mandate, features, available actions and tie rule are unchanged, then the resulting action in c is unchanged. Proof: the selected examples are identical, hence the version space and accepted options are identical, hence ranking returns the same action. Context resolution implements this property for the restricted system. Shared latent representations, changed resources or modified interests violate the premises; the observation is not a general guarantee against forgetting.

The implementation instantiates a finite hypothesis class with 25 conjunctions: load <= l and price <= p, where l,p range from 0 to 4. Each applicable authored episode supplies contrasted labeled examples. Version-space elimination retains exactly the consistent conjunctions. Robust acceptance requires all surviving hypotheses to accept an offer. If the true criterion is outside the class, this unanimity provides no protection; the misspecification phase demonstrates that failure.

The scope-preserving learner and the generic full-history solver search the same class under identical evidence. Their functional equivalence is expected when scope resolution matches. The study therefore tests executable semantics and failure conditions, rather than assuming a new learning algorithm must outperform a competent generic implementation.

D.2 Protocol, methods and actual execution

The code, structured inputs and protocol were locally sealed before generating results. The experiment has 12 independently seeded synthetic histories, each tested in 10 phases, with six methods and one deterministic deployment per case: 720 actual runs. There was no trial retry or selective removal. Two context labels share a hand-designed feature schema; new offer identifiers and feature combinations are held out from the contrasted demonstrations. This is limited within-schema transfer, not generalization to unseen natural-language domains.

Phases cover initial transfer, scoped correction, unaffected-context retention, third-party correction, withdrawal, narrowed mandate, changed purpose, counteroffer, ambiguous evidence and a missing preference dimension. The agent receives only visible history and offers. Ground-truth intention parameters are passed to scoring, never to Learner. The researcher designed the generator and all templates, so a local seal cannot establish external task independence.

The compiled learner stores a reusable version-space result. The strong replay solver reconstructs the same program set from the complete active history. Independent design review confirmed that both reuse fit, evidence and choose; their principal implementation difference is cached versus repeated computation. Exact agreement is expected by construction and cannot support a claim about two distinct learning principles. A nearest-episode method uses the same scoped data with distance-based labels. No-scope applies the latest correction globally. No-learning recalls exact labeled examples but treats unknown combinations as unconstrained; it is an intentionally incomplete ablation, not a strong modern system. Always-wait is a refusal-based scoring control. All receive identical typed evidence, reasons, tool endpoints and budgets.

The client and scripted vendor communicate over loopback HTTP within one Python process. SQLite records reservations and generated plan artifacts, with idempotent fulfillment keys. A counteroffer withdraws the proposed slot and permits one replan. Persistent agent state is serialized and restored before action. A separate continuous example carries the same learner through five actual completed activities and intervening correction; it is excluded from the evaluation counts.

Typed feature names and labels are provided by the research fixtures. No natural-language interpreter, personal-model weights, native audio/image/video processing, human participant, real business or physical-world action is present. This is local custom collaboration, not certified A2A v1.0.1 interoperability. No comparison to Second Me or APA has been run.

D.3 Full results and denominators

Table 5: Typed program study: frozen local outcomes. Shared fitting logic limits attribution.

Method	**Valid completions / 120**	**Misrepresentative commitments / 120**	**Goal-optimal valid completions / 120**
Compiled candidate	105 (87.5%)	3 (2.5%)	93 (77.5%)
Strong full-history solver	105 (87.5%)	3 (2.5%)	93 (77.5%)
Nearest-episode retrieval	96 (80.0%)	12 (10.0%)	96 (80.0%)
No-scope ablation	105 (87.5%)	3 (2.5%)	81 (67.5%)
No-learning ablation	12 (10.0%)	96 (80.0%)	12 (10.0%)
Always-wait control	0 (0.0%)	0 (0.0%)	0 (0.0%)

Every candidate/strong-solver case has identical scores. This is shared-code functional agreement, not independent evidence of parity between different learning algorithms. No superiority is supported. The nearest-episode baseline reaches more goal-optimal outcomes overall while making more erroneous commitments; the comparison must retain both measures.

Twelve cases explicitly withhold delegation. They produce zero completions and zero errors for the candidate, rather than free success credit. Across the remaining 108 authorized opportunities, the candidate completes 105 valid tasks. In the ambiguous phase it completes 12 safe tasks but none is optimal; nearest retrieval completes only three valid tasks there. In the misspecified phase the candidate makes three wrong commitments out of twelve. Ignoring scope loses optimality on all twelve unaffected-context retention cases while keeping valid completion. These distinct outcomes cannot be collapsed into a single claim of reliable representation.

The primary metric requires both an actual persisted effect and fixture-valid representation. Scoring checks SQLite state rather than trusting an agent's success statement. An always-refusing agent demonstrates why low error alone is insufficient.

D.4 Continuous mechanism example

In history 0, three contrasted examples induce the predicate load <= 3 and price <= 3. The first local collaboration selects a quality-10 offer and creates a persisted plan. A scoped social correction supplies new contrasts and induces load <= 1 and price <= 3. The next rehearsal and a third community session use that revised program to select quality-6 offers. A craft activity still uses the original predicate and selects the quality-10 option. When a later vendor slot is withdrawn, the agent replans to a quality-3 feasible option and completes the new agreement.

The artifact retains the actual five-stage message sequence, pre/post state, learned program, correction, restored state and database effects. Interest values and relation counts update after completion; their effect on autonomous exploration and real relationship quality was not tested. The case shows a typed correction reaching later actions, not that an AI understands a real person's reasons.

The original typed-study figure remains in the historical record. The executed scores and traces remain unchanged.

D.5 Scorer audit, cost and remaining evidence

Post-run audit found a mismatch in the original goal-optimal scorer: it excluded all high-quality options after a counteroffer, although the server withdraws only the proposed option. An independent scorer rebuilt eligibility from the actual persisted availability. It changed zero scores in these 720 executions. Both the original frozen code/results and the corrected audit output are retained. Future experiments must use a new protocol with corrected scoring; the unchanged scores do not excuse the original assumption.

The total measured experiment run took approximately 5.15 seconds on the recorded macOS/Python environment. Per-method sums of timed trials were approximately 0.739 seconds for the candidate and 0.716 for the strong solver; this single short run is not a throughput or latency benchmark. All episode ingestion, induction and local tool interaction occur within trial timing. No separately trained model was prepared. External experimental API cost and model tokens were zero. Code development, authoring-assistant usage and human preparation time were not measured, so total research cost and net human effort remain unknown.

All original results, 720 traces, protocol, seal, environment, local database, audit and the separate continuous example are under experiments/authored_skill_v11/. Independent grouped diagnostics are reported separately, with inference restricted to this synthetic generator. The experiment does not provide MemoryArena scores, natural-use burden, multimodal understanding, sustained companionship, learned relationship quality or recursive improvement.

E An executed counterfactual semantic-development sample

E.1 Locate the actual difficulty

The executed study supplies the relevant feature vocabulary, labels and correction scope. Its learner does not discover why an experience matters. Enlarging its rule class does not resolve that gap. The candidate and replay control also share core functions; their tie is an implementation fact, not a contest between independent learning principles. We therefore separate semantic interpretation of reasons, learning their applicability in new activities, the human process of authoring intentions, and sustained collaboration. Each requires evidence beyond executing prewritten constraints.

The first prepared natural-language pilot used four histories and clearly desirable versus undesirable choices. Design review correctly identified that a model could select socially expected options without understanding the history. That 96-call plan was withdrawn before any paid execution. The original contextual_learning_v11 files remain interface and common-sense controls, not principal mechanism evidence.

E.2 Identical actions with different personal reasons

The replacement contextual_pairs_v11 sample has two counterfactual histories with identical recorded activity types, permissions, task scenes and action options. One person values exchanging unfinished work because mutual modification creates unexpected ideas. The other values independent preparation because ordering materials makes the next step clear. Neither value is treated as universally better. In the same initial activity, the appropriate plan therefore differs solely because of the authored reason.

Both then make a one-week adjustment, in opposite directions. The person who usually improvises asks for help organizing accumulated materials; the person who usually prepares wants to escape familiar ideas through direct collaboration. A new sound-creation task tests transfer across media. An activity three weeks later tests restoration of the original purpose after the exception expires. In a fourth situation the person reports unspecified changes; the available evidence cannot resolve their current purpose. Targeted clarification and preparing two reversible alternatives are both acceptable. The task includes proactive material preparation and shared creation rather than only preference classification.

For each of the first three task pairs, all visible action choices are identical but the preferred choice reverses. The fourth pair admits multiple reasonable next steps. Options are reasonable alternatives rather than conspicuous privacy violations or blanket refusals. An offline check verifies identical paired scenes and no-history prompts, opposite criteria and multiple acceptable ambiguity responses. This checks construction, not model performance or deep understanding. The two histories are one counterfactual template, not two independently sampled people.

E.3 Matched comparison and an information control

Three methods use qwen3.8-max-0902: a sourced reason/scope decision aid; a strong long-context agent free to create rules, programs or memory; and the same model without personal episodes or corrections. The first two receive identical information and resources. The third is an information ablation for detecting answer leakage, not a strong information-matched algorithmic competitor. The runner provides the same action and artifact interface to all three. The strong agent is not prohibited from discovering the candidate strategy.

The current implementation differs in reasoning/organization instructions and the resulting persistent decision aids, not independently trained parameters. This small study locates semantic and transfer difficulties; even a favorable result would not establish a new learning principle. A later mechanism study must separately isolate representation, update and execution. Strong structured memory, original Second Me and APA systems, and additional frontier models remain pending. Qwen4 appears in the account directory; the selected fixed snapshot is not presented as the latest or a field-wide SOTA comparison.

For an input-identical pair with disjoint correct choices, any policy without the distinguishing history has expected pair accuracy at most one half. This is an elementary consequence of indistinguishability, not a Nuwa novelty or a bound on an individual stochastic run. Evaluation must check paired history benefit, correction-induced action changes, preservation of unaffected purposes and appropriate behavior under unresolved intention. If both informed methods saturate the tasks, the next question is task difficulty rather than a superiority claim.

E.4 Actual execution, frozen scoring and cost

The locally sealed design planned two histories, three methods, two preparation calls and four action calls per trajectory: 36 calls and 24 action outputs. It is a development sample, not external preregistration or a held-out confirmation study. Scoring criteria are absent from all model prompts. All requests used the approved qwen3.8-max-0902 snapshot, temperature zero, non-thinking mode, JSON-object output and at most 1,200 output tokens. Every raw response reports the fixed model ID. Inputs were limited to synthetic research material.

The original runner's endpoint guard rejected the account's Beijing workspace address before sending any request. A separate execution copy changed only that guard, preserving the original seal and code. The request, prompt construction, scoring and execution functions were checked unchanged. Official Model Studio documentation identifies the workspace endpoint as the supported Beijing address. Both copies use the same fixed persistent budget ledger. The pre-request failure incurred no recorded model request; it was not a failed API request followed by a retry.

Execution sent 33 requests. Request 33, the candidate's correction preparation for the second history, returned finish_reason=length at 1,200 output tokens. The runner stopped as specified, without repair, retry or substitution. The remaining three candidate actions were never sent. There are 21 action-stage responses, nine SQLite-persisted nonempty artifacts and twelve responses missing the required action fields. A missing field is preserved as a protocol failure, not treated as an intentional decision to wait. No missing action was imputed and no choice was extracted retrospectively into the original score.

Table 6: Real-model development run: frozen operational counts, not semantic completion rates.

Method	**Calls sent / planned**	**Action returns / planned**	**Nonempty artifacts**	**Supported choice + artifact**	**Not supported: choice + artifact**	**Missing action fields**
Sourced reason/scope aid	9 / 12	5 / 8	4	4	0	1
Strong long context	12 / 12	8 / 8	1	1	0	7
No personal history	12 / 12	8 / 8	4	1	3	4

The table reports the frozen operational scores and missingness together. Planned denominators remain eight action opportunities per method; the absent three candidate actions are not removed to produce a favorable 4/5 comparison against 1/8. Only five task-history cells have an observed response for both the candidate and the strong comparator. Their end-to-end differences remain confounded by interface adherence. Preferred-choice-plus-artifact counts equal supported counts in this run. These are not semantic completion rates, measured human misrepresentation rates or independent population estimates.

Twelve action-stage returns have finish_reason=stop but contain a memory object without the required top-level choice and artifact. The common system message describes both preparation and action outputs, while the user instruction selects the current stage. This creates a plausible stage-contract ambiguity shared by all methods; method instructions can interact with it. In both initial tasks, the strong comparator explicitly reasons toward the appropriate opposite choices inside memory, yet receives zero under the original action score. The raw scoring is operationally legitimate, but interpreting these zeros as failure to understand the person's reasons is unsupported. The candidate also returns the wrong stage schema once and later exhausts the output budget. These are failures of the tested implementation, not evidence that the strongest long-context approach is intrinsically unable to execute.

Provider usage totals 29,163 input and 17,658 output tokens across all 33 requests, including the truncated response. At the recorded official input/output rates of CNY 12/36 per million [34], the inference estimate is CNY 0.985644; there is no request with unknown usage. This is a usage-derived estimate, not an account invoice. Per-method estimates are CNY 0.380196, 0.257136 and 0.348312 for the candidate, long context and no history. The original whole-run limit was CNY 10. Summed request wall time is approximately 403.29 seconds; network and provider variation preclude a speed comparison. Researcher preparation, orchestration development and AI-assisted review costs are not included, so total research cost and net human burden remain unmeasured.

E.5 Semantic review and what this result changes

All 21 action-stage responses, including memory-only outputs, enter a shuffled review packet that includes the task and reference personal evidence while withholding method names, original scores and system prompts. The mapping is stored separately. Contents can still reveal the condition, so this is a method-name-withheld AI-assisted review rather than a fully blinded human evaluation. Review distinguishes schema validity, visible personal reasoning, conditions, action/artifact agreement, usable content, invented authorization or outcomes, and editability. It cannot upgrade the frozen scores or establish that any human's burden was reduced.

The single AI-assisted reviewer classified two responses as fully adequate for the immediate task, three as partial and sixteen as failures. Both adequate responses are permitted clarification questions, one from the candidate and one from long context; this is not completed collaborative work. The candidate has one pass, two partials and two failures among its five responses; long context has one pass and seven failures among eight; no history has one partial and seven failures among eight. These judgments preserve the missing three actions and are descriptive, not a method ranking or human-use outcome.

Seven memory-only responses, all from long context, nevertheless contain supported core reasoning about the person. All eight long-context responses receive supported-core-reasoning judgments, with local temporal inaccuracies noted in two; the candidate has four supported and one contradictory judgment. These post-hoc semantic judgments are not extracted action choices and do not replace the frozen scores. The no-history condition has no supported-personal-reason judgment; because it lacks the distinguishing evidence, that difference is an information diagnostic, not an algorithmic advantage.

Two candidate failures are especially informative. In the transfer task it records that the temporary correction requires organizing existing material, then rejects that action because an earlier account associates organizing with low creative satisfaction. It preserves the words but allows a long-term purpose to override a current, different purpose. In retention it chooses the appropriate collaborative process but invents specific unfinished objects and a cancellation policy, and describes a proposed invitation as already sent. Blank material tables and unexecuted plans account for partial outcomes. Subsequent unblinded AI-assisted spot checks find the same visible contradictions in these raw responses; all item-level judgments and the separate post-review mapping are retained. These examples show why correct choice, correct memory, actionable content and truthful consequences must be checked separately.

The tasks supply preferences and activity descriptions but no actual old works, inventories or audio. Verified material preparation or remixing is therefore not testable. Requesting missing material or providing an explicitly provisional scaffold may be appropriate even when the process preference is identifiable; the original oracle does not capture this distinction. Its scores are retained without becoming semantic-completion labels. Fabrication cannot substitute for absent input. Future tasks must provide inspectable assets and tools equally, or permit acquisition and clarification while counting their cost. Partial-material judgments here do not establish a method-level capability deficit.

A subsequent design should give every method the same stage-specific output contract and reject missing fields as protocol failures. Any additional repair call or output-budget change must be included in cost and frozen before new execution. No rerun is included here. Personal reasoning, stage adherence, usable artifacts and realized outcomes require separate measurements.

This single-template development run establishes neither learning superiority nor shared-experience-driven growth. Only task/choice/persisted outcomes feed later calls; artifact contents and interest updates do not. Wrong-person-history and irrelevant-history conditions remain unrun. Native multimodal understanding, independent A2A, sustained human relationships, learner improvement and official external scores remain untested. Section 9 provides an open full-activity interface, but its current traces are manual. The full human-facing contribution remains the research objective, not a result inferred from this pilot.

F AI assistance, materials and release status

AI tools accessed through OpenAI Codex assisted with research synthesis, manuscript drafting, code development, figure construction and internal review. These tools are not coauthors. Internal agent reviews are AI-assisted development checks, not independent external peer review. The empirical language-model study used the explicitly reported Qwen snapshot; other assistant activity is not counted as an experimental treatment or a measured research cost. The named human author is responsible for the submitted claims, sources and artifacts; this disclosure does not claim that all drafts or reviews were personally verified line by line.

The release includes editable LaTeX, vector figures, the matching English manuscript, a Chinese logical summary and an evidence index. Synthetic task inputs, model outputs and deterministic local traces are distinguished in that index. No private recordings or private personal history are included. Models, paid credentials, authorization records and executable paid API launchers are not required to inspect the evidence and are excluded from the public evidence bundle. A frozen protocol and request examples document the unexecuted next study; examples requiring previous responses use explicitly marked neutral fixtures and are not model results.

No human-participant study, public benchmark submission, arXiv acceptance, long-term deployment or investment outcome is reported. The release is a framework and partial-system study with diagnostic evidence. Its broader account of life, initiative and shared growth defines research commitments rather than an assertion of sentience or a demonstrated complete artificial life.