Liva.Taihe

After AGI:

We Need Humanism More Than Intelligence

Humanistic Intelligence for Changing Lives and Shared Futures

Huaiyang Liu

DeepBean Limited · wyonliu@gmail.com

Normative design synthesis · 21 November 2026

Abstract

Advanced intelligence can help people encounter new possibilities, form purposes, create together, earn a living, and receive care. When AI helps shape physical and virtual activity, learning from a successful experience also shapes what people can do next. Liva.Taihe argues that useful capabilities can warrant continuation while their originating purposes, assignments, and cultural forms require distinct, connected judgments. Shared responsibilities may justify continuity; a person's standing never depends on productivity. Truth and love join concern for particular lives to answerable inquiry. Lingji and rationality support both new questions and their examination, while rest and familiar commitments remain legitimate ends. A constructed theatre undertaking connects an unexpected encounter to a new purpose, an editable practice, and later people's different works. Even unanimous agreement to present automation can leave successors unable to adapt what is inherited. Comparing recordings, executable handovers, alternative services, and maintained practices exposes the enabling conditions and costs that can reverse the recommendation. Drawing on capability, care, and autonomy accounts, the contribution is a normative design synthesis connecting learning, delegation, and transmission. World models, persistent agents, editable world production, and recursive improvement are candidate means for present and conditional post-AGI worlds. Local technical diagnostics establish no algorithmic advantage; participant and civilizational benefits remain untested.

1 The problem: what stronger intelligence is for

An intelligent system can complete a task while failing the person whose life gave the task its point. It can save time while increasing correction work, provide company while displacing a wanted relationship, or produce an excellent artifact while removing the learning or gift-making someone wanted to undertake. Conversely, reliable delegation can relieve pain, sustain independence, or make an otherwise impossible collaboration feasible. Capability is necessary for many such benefits; its increase alone does not select among them.

Taihe begins with concrete lives: uncertainty, difficulty, ordinary happiness, people who matter to one another, unfinished interests, and ambitions that change. Its central thesis is reciprocal. Advanced intelligence should support plural lives and shared futures that people can meaningfully influence; their experiences and deliberation should be able to redirect what intelligence develops. A person's life is not exhausted by the tasks a system already knows how to score. Nor is continual novelty an obligation imposed on someone who seeks stability or rest.

We examine a conditional post-AGI setting: systems perform much cognitive work and help produce and operate environments. This is neither a date prediction nor an assumption that scarcity, conflict, disability, or care disappear. The argument also applies to present systems. AGI is not used as an operational capability threshold in our experiments, of which none evaluates such a system.

This manuscript distinguishes normative arguments, conditional scenarios, technical definitions, mathematical claims, and empirical hypotheses. A normative premise requires justification and challenge, not a p-value. A computational mechanism requires implementation and comparison, not endorsement of its purpose. The organizing contribution is the connection between these forms of inquiry while retaining their different standards. Sections 2–6 contain the humanistic argument; Sections 7–8 specify limited technical consequences; Section 9 separates the relevant evidence; Section 10 locates the proposal among predecessors.

The central design question is what should continue when intelligence learns from a successful experience. Technical competence, personal purpose, a division of work, and a cultural practice can develop together without having the same grounds for continuation. Sections 1.1–1.4 state the judgment, apply it through one undertaking, and specify how to challenge it. Sections 2–6 supply its humanistic reasons and conflicts; Section 10.1 compares its scope with the closest accounts.

Prunkl already distinguishes autonomy-as-authenticity from autonomy-as-agency and examines AI-mediated adaptive preference formation [46]. Taihe inherits this challenge to evaluating assistance by execution and satisfaction alone. Its proposed contribution concerns the subsequent use of experience: how one episode becomes a capability, an assignment, a shared default, or an inheritance. These uses require different judgments even when the originating experience was valuable.

1.1 What should continue after a successful experience?

Core judgment. In continuing AI assistance, successful experience can justify reusing a useful capability without justifying continuation of its original purpose, division of work, or cultural form. A system should make useful means available where their benefits warrant the burden, while treating the continuing use of those means as a separate, revisable judgment. Existing commitments and shared responsibilities remain reasons for continuity; a changed preference does not automatically cancel them. The point is to preserve the right grounds for each decision, rather than to favor change or preservation in advance.

The argument starts from the whole person. An encounter may matter because of affection, beauty, surprise, companionship, useful work, or relief from a burden. It need not express a stable preference that existed before the encounter. Assistance can help make this possibility real: provide an unfamiliar instrument, make an activity accessible, explain a technique, connect willing collaborators, or take over exhausting logistics. The value of this positive support is lost if humanistic intelligence is reduced to restraining action. Yet the same experience does not decide how the rest of a life should be organized.

Three reasons establish the separation. First, evidence that a technique works supports a claim about a means; enjoyment may support a tentative interpretation of what matters to a person. Neither, without further grounds, settles who should undertake its maintenance or which shared norm should govern others. Second, purposes can change through participation, while promises, dependence, and other people's contributions create claims that a single user's later satisfaction cannot settle. Third, retaining a successful output does not by itself retain the conditions through which successors could adapt or redirect its practice. Consequently, technical learning, the interpretation of purposes, the allocation of work, and the transmission of capability should not be merged into a single inference from success.

This is an argument from capability, care, relational autonomy, and the humanistic premises developed below. It does not introduce a new moral foundation. Its design consequence is specific: an agent may reuse a repair skill, offer a tentative new undertaking, honor an established teaching commitment, and make tools accessible to newcomers, while declining to infer a permanent vocation or a group's future artistic standard from the same episode. These decisions can differ without treating the episode as a failure. A single undifferentiated profile or success reward is inadequate if it makes those uses stand or fall together.

The judgment is also stronger than retaining an approval record. A record can show that an invitation was accepted while saying little about whether collaboration remains practicable or whether later people can use an inheritance. Conversely, revisability need not require repeated questioning, a new application, or a complete personal archive. It can be sustained by a durable delegation, an available alternative service, a working relationship, or an already adequate organizational process. Which means work is an empirical question. Whoever supplies those means must count the effort and resources they require.

Current success and later satisfaction remain valuable evidence. Their insufficiency follows only when relevant differences are absent from them. If an evaluation maps histories x and y to the same completion-and-satisfaction record M, it cannot distinguish a wanted contribution from an unacknowledged burden when that distinction is omitted from M. This is an application of Sen's criticism of informationally restricted evaluation [30], not a new impossibility theorem. A richer evaluation that already includes the relevant commitments, contributions, and practical alternatives may answer the design question without any Taihe-specific representation.

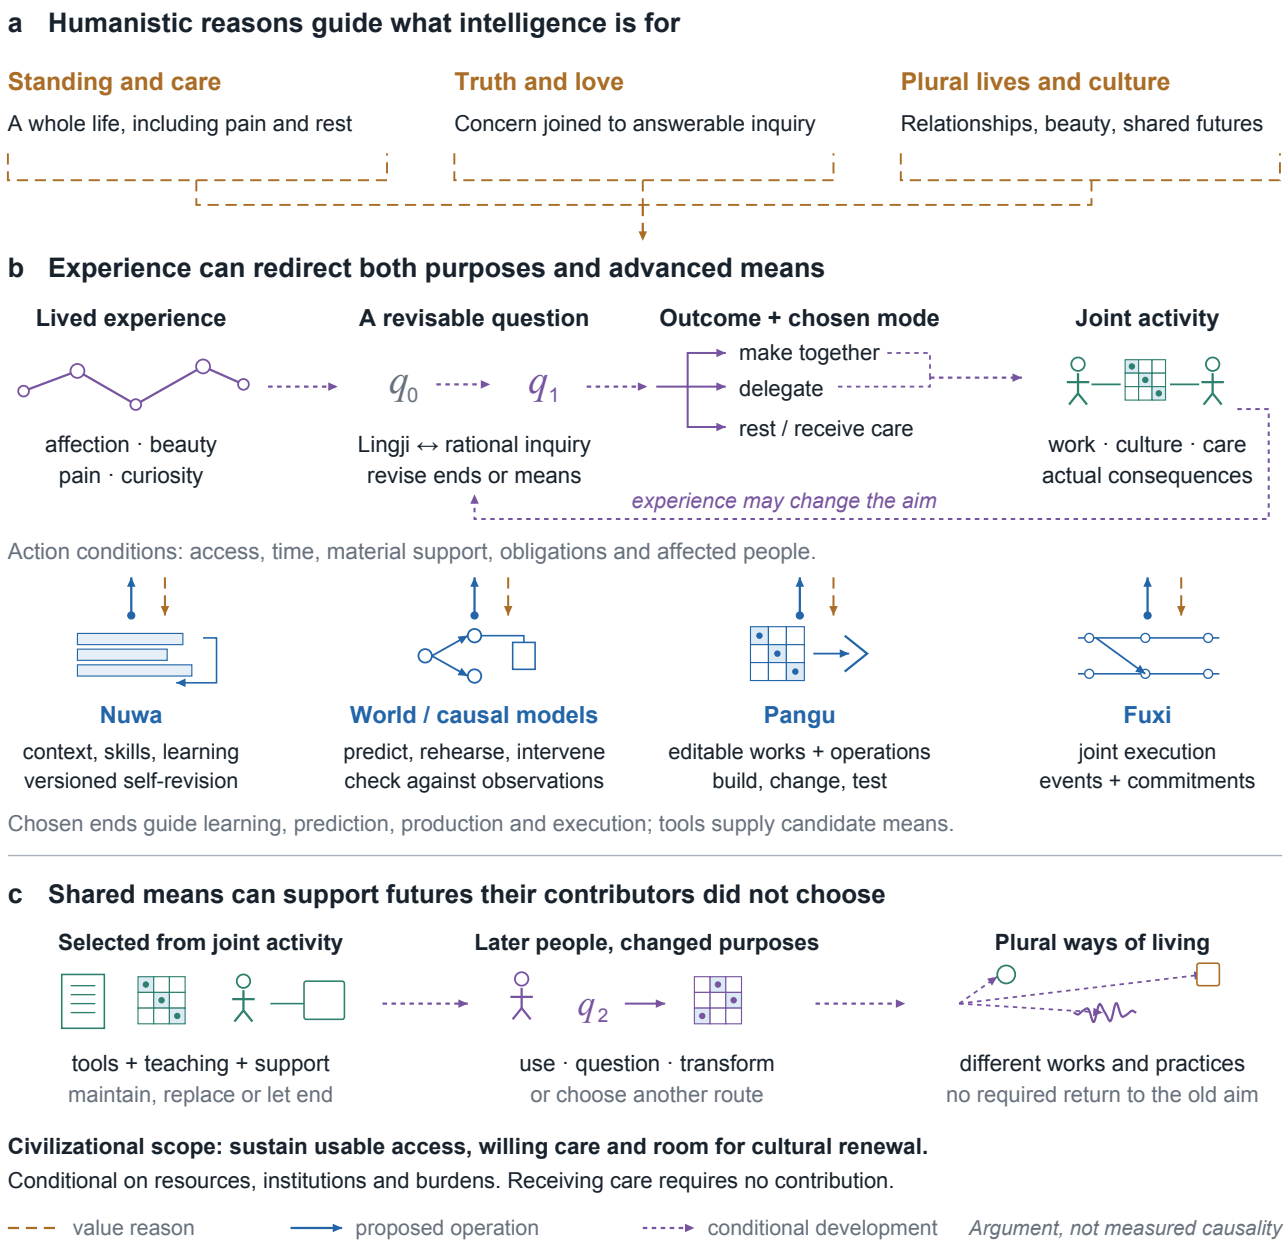


Figure 1. Humanistic reasons, advanced means, and shared futures. (a) Standing and care, truth and love, and plural lives and culture supply normative reasons for selecting what assistance should enable. They are not rewards computed from behavior. (b) Lived encounters can motivate a revised question; Lingji and rational inquiry contribute to examining ends and means. A chosen outcome includes a mode of participation: making together, delegation, or rest and receiving care. Rest does not require a productive continuation. Consequences of shared activity may reopen the question. Paired support links distinguish requirements on a technical task from candidate technical means: Nuwa's contextual learning and self-revision, world and causal models, Pangu's editable production, and Fuxi's joint execution. These are proposed roles, not exclusive faculties or a report of an integrated implementation. Model predictions require observation; self-improvement claims require separate tests of the changed improvement procedure. (c) Selected tools, teaching, and support may enable later participants to use, question, transform, or replace an inheritance. This passage depends on material access, willing maintenance, institutions, and distributed burdens. It does not require retaining every practice or making personal wishes into collective rules. The civilizational judgment concerns conditions for plural lives and cultural renewal, with care independent of contribution. Dashed amber links denote normative reasons, solid blue arrows proposed operations, and dotted purple arrows conditional development. The figure is an argument structure, not a measured causal model or a forecast.

1.2 One undertaking, changing purposes and responsibilities

Consider a constructed theatre undertaking across three periods. A technician initially needs to repair equipment. A small improvisation during the repair evokes a childhood place; collaboration makes an unfamiliar style worth exploring. The assistant can explain acoustics, prepare materials, connect willing performers, and handle invoicing. The resulting work has practical and aesthetic value. Its origin in a felt association is neither evidence that the association is factually correct nor a reason to suppress it: Lingji proposes a direction, and inquiry tests what can be made of it. Friendship and an artistic interest may develop without obliging the technician to become a professional artist.

During the next production, the same technique can be reused while the appropriate division of work changes. If the technician wants to practice, an excellent automatic rendition can displace the very activity being sought. If paid work, fatigue, or family care makes relief more valuable, full delegation can better serve the same person. Reliable technical work and artistic participation are both positive achievements; neither is the universal target. A system that interprets every successful collaboration as a reason to seek more collaboration gets this transition wrong even when its predictions of output quality improve.

Now suppose the technician has actually accepted a supported teaching role for the season. A new wish for free time does not by itself establish that the group should be left without instruction. Others may have organized their work around that commitment. The assistant should help renegotiate duties, arrange a substitute, or carry out the agreed plan within its scope. If there was only an enthusiastic response to one invitation, the same teaching assignment lacks that ground. The difference lies in an established relationship and its responsibilities, not in which branch produces the higher immediate satisfaction score. Urgent illness or a serious care need can supply reasons for release that ordinary inconvenience does not.

A later group may wish to reinterpret the work. Keeping the original performance preserves a shared memory; preserving instruments, methods, and access can enable a new work. Neither requires successors to affirm the original style, and neither automatically requires the original technician's continuing labor. The system can carry forward technique and selected context, arrange supported teaching, or supply an equivalent tool, while making a new interpretation feasible. If another service provides the relevant opportunities at lower burden, that service may be preferable to maintaining the original process. Intimate records can be omitted where a teaching example or a reconstructed method suffices.

Section 1.3 makes the production and adaptation path explicit. The positive continuity in this case is concrete: a useful repair, an unexpected interest, a friendship, an enjoyable work, a supported teaching practice, and material others can transform. Truth prevents a simulated reception from being presented as a real collaborator's response. Love and care change which burdens the assistant should relieve and how it responds to vulnerability. Rational inquiry makes the new artistic question answerable to evidence without replacing aesthetic judgment. Relations determine who has undertaken what; cultural plurality keeps later reinterpretation legitimate. Unconditional standing explains why the person may retire, receive care, or leave an ambition without losing their claim to help. These commitments change the assistance at different points in the same undertaking.

The case deliberately admits reversals. A legitimately adopted style can be a useful shared convention rather than imposed conformity. A well-supported commitment can justify continued work rather than immediate exit. Automatic production can sustain the wanted relationship or release time for it. An archive can burden the very people it is intended to help. The proposal therefore supplies no preference for maximal participation, indefinite preservation, constant novelty, or a particular technological arrangement. Its criticism applies when experience is used to support a decision whose relevant grounds or enabling conditions are missing.

1.3 Full agreement can coexist with a loss of future making

The preceding assignment dispute is not the strongest case. Consider a version in which every current member wants automatic production, the technician freely retires, compensation and commitments are settled, and no one is deceived or forced to teach. The automation produces the desired performance. These are genuine achievements. They still leave a separate question: what can another group subsequently make with what remains? Agreement about present work does not establish the availability of future means. No alleged defect in the original consent is needed for this question to arise.

To make the capability concrete, specify how it arose. At the first event, an accidental resonance prompts the group to ask whether audience contributions could become part of the sound rather than merely trigger a finished recording. AI assistance can propose candidate input-to-sound mappings, render variants, expose timing conflicts, and help test whether changing one input actually changes the intended part of the performance. People contribute associations, judgments about the shared experience, physical participation, and disagreements that can change the question. The group may discover a purpose it did not hold at the repair stage: making a performance whose audience helps compose it. The positive product is both a work and a repeatable ability to vary its relations. This is a constructed account of possible development, not a report that these effects have been measured.

Suppose the next group wants remote text contributions to control the rhythm and to replace the original style. This requires more than playing the existing recording. Someone must be able to inspect the relation between input events and cues, change that relation, run the resulting work, and check whether the contribution has the intended effect. The relevant operations can be supported by a local editable representation, by a sufficiently capable external service, or by collaborators with the required knowledge. The successor need not inherit the original artistic purpose or use the original implementation.

Compare four concrete continuations of the already agreed automation. All can preserve immediate playback; their remaining abilities differ. The following are stipulated interface conditions to be checked in an application, not measured product properties or benchmark results.

Table 1: Continuations and consequences for a successor's new work

Continuation	**What the successor can obtain or do**	**Consequence for the proposed new work**
Recording and narrative archive	Replay the piece and read why it mattered; no accessible input-to-cue editing or executable reconstruction	Keeps memory and interpretation, but does not establish the specified making ability
Executable, editable handover	Inspect and change event-to-cue relations, run the relevant dependencies, and test changed inputs within available access and support	Can support the new work if the required operations actually succeed within the successor's constraints
Alternative production service	Request the changed mapping, inspect or test its response, and obtain a usable result under credible continued access and cost	Can supply the same relevant ability without retaining the original production process
Full original workshop	Preserve tools, physical practice, and instruction, with additional upkeep and willing participants	Can supply distinctive embodied or shared experience, but that value and its burden require their own grounds

The necessary relationship made explicit by the synthesis is between what is retained and a feasible path from a later concern to a consequential modification. A description may help someone find such a path without supplying its tools; a working service may supply the path without preserving the description; a large archive can fail at both. In the notation of Appendix B, availability of a record is one possible component of X, whereas the desired outcome-and-mode descriptor must lie in the relevant feasible opportunity set. The latter depends on executable operations, access, collaborators where required, and resources. This is a consequence of the stated definition and capability premises, not a new reachability theorem or a proof that editability alone is sufficient.

A judgment can now use inspectable facts. Establish the successor's sought change and chosen mode; identify which operations it requires; inspect the offered interfaces, dependency availability, access conditions, and support; and test the disputed transformation where feasible. If the recording cannot support the transformation, while the editable handover can and its upkeep is justified relative to other needs, there is a reason to retain that handover despite unanimous preference for automating present production. If the alternative service passes the same relevant tests with lower burdens and credible availability, it may be the better continuation. If a workshop supplies a wanted embodied or shared activity that the service does

not, the activity being compared has changed and that difference must be included. A request to co-create cannot be counted as satisfied solely by delivery of an otherwise excellent finished piece.

Unknown facts leave the corresponding conclusion open. Unverified exportability is not executable access; an advertised service is not a demonstrated adaptation path; unknown maintenance capacity cannot be assumed free. When there is no known successor demand, the comparison must state which plausible future activities and whose claims justify expenditure, rather than pretending to identify the best inheritance for all possible futures. Urgent care, ecological cost, or an equally useful shared resource can defeat local preservation. Under uncertainty, a limited and revisable provision may be defensible, but it is not made obligatory by the mere possibility that someone might someday benefit.

Connecting the accounts changes the unit of judgment: the current worker can reasonably choose delegation, the group can reasonably automate, and their infrastructure can still merit a different decision about what enables later making. Care explains why the retiring technician cannot be treated as a free dependency; capability explains why possession of a recording is insufficient; relational agency explains why later collaboration and purpose formation matter; cultural plurality explains why repeating the original style is not the sole success condition. The resulting design choice is to preserve or replace the needed enabling relations where warranted, while permitting the original work arrangement and purpose to end. None of these predecessors is presumed incapable of the conclusion. The contribution offered is the explicit connection and its application to choices that a single persistence decision would conflate.

1.4 What would challenge this contribution?

The normative judgment can be challenged in at least two ways. A theory might defend continuing purposes and assignments on the basis of successful participation alone, including cases where the interpretation is later contested; Taihe must then explain why the claimed grounds are insufficient, rather than treating disagreement as misunderstanding. Alternatively, a case may show that separating the decisions destroys a valuable practice whose obligations and meaning cannot be decomposed without loss. Shared ritual, sustained care, and demanding artistic apprenticeship make this objection serious. The response must evaluate the practice and its burdens together; it cannot assume that every component can be detached or individually vetoed.

The historical and conceptual contribution has its own limit. If a careful application of the strongest predecessor already supplies all the same distinctions, decisions, and reversal conditions, the defensible contribution is an explicit design application of that account. Naming the synthesis does not establish an independent theoretical discovery. Section 10.1 therefore identifies the judgment inherited from each source before specifying what this paper organizes for a continuing AI-assisted undertaking.

A technical intervention following Section 1.3 would test an actual change path: propose a successor transformation, identify the operations and dependencies it requires, execute a candidate realization, and revise the retained means when that test fails. This differs from merely assigning more contextual labels to records. It adapts planning and capability testing rather than introducing a new general algorithm. No such integrated intervention is implemented here. A strong contextual or reflective control must also be allowed to inspect tools, propose adaptations, and perform the same tests.

A design comparison should hold the available history, capabilities, model, tools, memory allowance, and total assistance budget fixed. A strong contextual control must receive information about roles, reasons, commitments, alternatives, and later corrections. Vary only which use of the same experience is at issue: technical reuse, a new personal proposal, a shared assignment, or a successor's adaptation. Include valid continuing commitments and deliberate delegation, so that more questions or refusals cannot win by definition. Test whether a proposed mechanism improves relevant completion, wanted participation, adaptation, and the distribution of burdens, counting clarification and maintenance as costs.

If the strong control makes equally appropriate choices at equal or lower cost, no additional mechanism has earned an advantage. If separating uses systematically produces confusion, unwanted interruption, or erosion of valuable relationships, the proposed means should be revised. Constructed cases expose these contrasts; they do not establish their frequency, human experience, or an intervention's benefit. No such human or real-model results are reported. The retained technical negative results in Section 9 concern narrower mechanisms and do not supply the missing evidence.

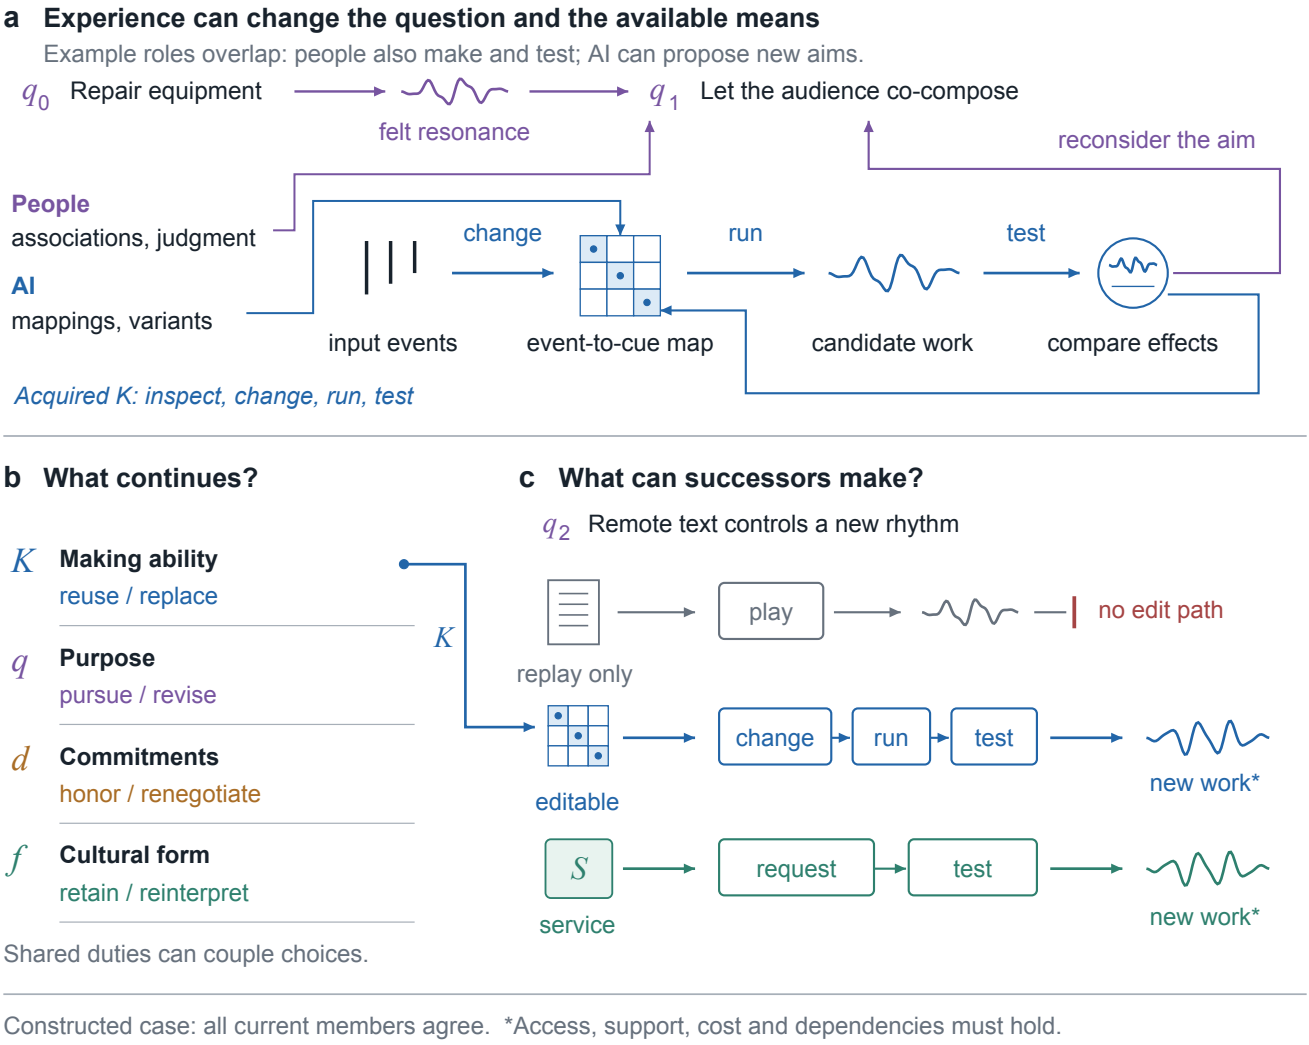


Figure 2. Experience can change both an end and the means of pursuing it. (a) In the constructed theatre case, a felt resonance prompts a possible shift from repair to audience co-composition. The illustrated roles are not exclusive: people also make and test, and AI can propose new aims. Human interpretation and AI proposals contribute differently; changing, running, and comparing mappings can support both a revised purpose and usable making operations. (b) Capabilities, purposes, commitments, and cultural forms have distinct grounds for continuation, while relationships can couple their reconsideration. (c) Unanimous agreement to current automation leaves a separate question about future making. Playback alone, editable operations, and an alternative service offer different stipulated paths to a successor's new work. Feasibility depends on access, support, resources, and the required participation mode. This is an argument diagram, not a measured causal model or an implemented integrated system.

2 Human standing, truth, care, and a changing life

The continuity judgment needs reasons for helping a person, respecting a commitment, or sustaining a practice. The following premises supply those reasons. They also explain why retaining a useful capability cannot become a demand that its contributor remain productive or keep pursuing its original purpose.

2.1 The premises are commitments, not discoveries about physics

We adopt four starting commitments. First, a person's standing does not depend on comparative productivity. Second, suffering, truthful understanding, relationships, chosen activities, and the conditions for their pursuit can each matter; a single convenience metric need not represent them. Third, people are situated and interdependent: what they can actually do depends on bodies, resources, other people, and institutions. Fourth, both preferences and interpretations of a life can develop, so historical behavior cannot permanently authorize a future on someone's behalf.

These premises are contestable. A strict productivity-based ethic rejects the first; a single-value theory may reject the second. Taihe does not deduce their falsity from an optimization theorem. It asks readers to consider the cost of those alternatives: people needing extensive care would have weaker standing precisely when their need is greater, and stable historical profiles could override later considered choices. We judge those consequences unacceptable. This judgment is the normative foundation of the program, not a property inferred from training data.

2.2 Why truth and love belong together

Claim H1: trustworthy care needs both attention to a particular person and an answerable account of consequences. If a system pursues affection without accuracy, it can comfort through falsehood or preserve attachment by concealing alternatives. If it pursues accurate task performance without regard for the person, it can efficiently deliver an unwanted or harmful intervention. Concern identifies whose circumstances deserve attention; inquiry tests whether a proposed response will help. Neither role can substitute for the other.

The difficult objection is that truth can hurt and urgent care may leave no time for complete explanation. Our response is contextual disclosure and supported decision-making, not an unconditional demand to state every fact immediately. Relevant uncertainty, timing, comprehension, and the person's expressed wishes matter. Persistent deception for retention is different from choosing a compassionate way to communicate uncertainty. The system should distinguish observed facts, interpretation, and an offered response so that kindness remains open to correction.

Love here names a human orientation toward particular others and what matters to them. We do not infer that a model experiences love from its words or establish a measurable quantity of it. The design consequence is concrete: evaluate the appropriateness and consequences of assistance in the relevant relationship, including corrections and burdens, rather than the emotional fluency of a response alone.

2.3 Why standing survives superior machine performance

Claim H2: standing cannot be lost merely because another agent performs one's activities better. Given the first premise, a capability ranking cannot determine entitlement to care or membership in a community. The argument also avoids making significance depend entirely on comparative achievement. Learning a song, receiving a gift, caring for someone, or understanding a place can matter without being the world's best performance.

A strong objection remains: some achievements derive meaning from being consequential, difficult, or genuinely one's own; simulated accomplishment may fail those standards. We accept that objection against an automatic guarantee of meaning. Removing an economic need for human performance can change institutions of recognition and leave real losses. It does not follow that all relationships, understanding, and chosen practices become worthless. Knell and Rüther defend receptive engagement as one component of a post-work account; we use their argument as a close predecessor and retain the objection that entertaining activity alone may not supply meaning [33].

Taihe does not rank people by the depth of meaning they achieve. Rest, receiving care, and limited ability to articulate a purpose do not weaken standing. This is stronger than a proposal merely to offer everyone another difficult game to win. It also leaves room for work and income as important chosen purposes; it does not predict their disappearance.

2.4 Why available resources are not yet a good life

Claim H3: evaluate what a capability enables for a situated person, not only what the machine produces. Equal access to an interface can yield unequal practical possibilities because of time, disability, language, equipment, or social support. Sen's argument for basic capabilities directly precedes this shift from goods to what they enable, including its difficult indexing problem [30]. Nussbaum supplies an explicit account involving affiliation, imagination, play, nature, and material and institutional conditions [31]. These precedents prevent us from claiming that plural human possibilities are a new ethical discovery.

Our extension is a research question about interactive, changing AI-mediated environments: how do an agent, a world, and a production system together enable an activity for this person, in a form they can choose and revise? Technical feasibility is only one component. Access, relationships, and whether the activity is wanted still need separate inquiry. A larger catalogue of options can overwhelm someone or conceal the loss of a valued commitment. We therefore reject opportunity counts as a complete objective and preserve the person's option to choose fewer activities.

2.5 Why care over a lifetime must permit reinterpretation

Claim H4: continuity of care requires continuity of reference without freezing identity. A remembered event may remain important while its meaning changes. A companion needs to distinguish a past act, a current commitment, a speculative interpretation, and a future aspiration. Otherwise consistency with yesterday's profile can become misrepresentation today. Care ethics supplies a direct challenge to treating people as isolated choosers: dependence and situated relations are ordinary features of life, and responsibility for care is institutional as well as personal [32].

The counterexample is unstable or coerced revision: a new statement need not warrant discarding every past commitment. The response is a contextual revision process that preserves the old evidence, asks what changed when that is useful, and records which decisions the revision concerns. This does not make the latest utterance infallible or grant the system final authority over an authentic self. Support can include trusted representation when a person cannot communicate fully, with reasons and competing interests made explicit.

Relief and growth follow from these premises together. Reducing unwanted work serves a person; supporting a practice they choose also serves them. Time saved and time worth spending are distinct outcomes. Neither active creation nor constant self-documentation should become the admission price for assistance. Voluntary self-expression can improve a companion's understanding, but the time and emotional work required to maintain it belong in the account of benefit.

3 Rational inquiry, Lingji, and finding another direction

Purposes and capabilities can develop reciprocally: a newly available means can make an unfamiliar life direction conceivable, while a new concern can redirect the abilities worth developing. This section distinguishes that positive development from better execution of a fixed task.

3.1 A complementary relationship

Rational inquiry makes explanations examinable, coordinates plans, and allows knowledge to accumulate. Lingji names a motivating family of experiences: an unexpected association, an emotionally salient concern, a courageous departure, a playful experiment, or a question that an inherited task description did not express. It is not a unitary psychological mechanism, an exclusive human faculty, or a synonym for randomness. Rational thinkers can generate new goals; emotion can illuminate neglected concerns or amplify bias.

Claim H5: systems should preserve a route from an emerging concern to a consequential test or reconsideration, without treating the concern itself as proof. The reason follows from revisability: if a system permits only better answers to already specified questions, people cannot use it

to examine a direction they cannot yet formulate in its vocabulary. The objection is wasted attention and endless novelty. The response is bounded exploratory branches, explicit costs, and a way to retain an unfinished idea without forcing its pursuit. A person may choose stability. The balance concerns the relation between proposal, inquiry, action, and reflection, not a fixed ratio of emotion to logic.

3.2 Four changes require different evidence

An unknown route to an unchanged goal is an exploration problem. Discovering a useful variable is a representation problem. Inventing an interaction is an expansion of available action. Choosing a different end is a deliberative change. These can occur together, but success in one does not establish the others. A new description with identical consequences is not a new technical discovery; a technically new consequence may be unwanted.

For a fixed utility, information set, budget, and admissible policy class, an optimal policy cannot be improved upon in expected utility by another policy in that class. Consequently, a Lingji-inspired gain must be explained through changed information, representation, actions, practical search, or evaluative commitments. No claim of universal superiority over rationality survives this boundary. Conditional benefits under bounded computation remain open to experiment.

A useful research sequence is to identify an encounter, specify what the old framing misses, propose an alternative, and expose it to a discriminating test. A changed purpose also needs the person's reasons and later response to its consequences. AI may supply a question or a collaborator may supply a suggestion. Fair algorithm comparisons must provide the same information to all groups. Human suggestions are not a hidden oracle awarded to the preferred method.

3.3 Growth without perpetual optimization

The youthful disposition valued here is curiosity and willingness to encounter an unfamiliar world, not an age requirement or an instruction to take every risk. Practice can join courage to patience; failure can reveal a limit without becoming a verdict on the person. A system can offer a recoverable attempt, a skilled partner, or an alternative interpretation of a work. It should also accommodate fatigue, mourning, responsibility, and a return to familiar places.

Self-cultivation can involve learning how to direct attention and action, including restraint. It should not be reduced to a continuously rising personal score. A participant who abandons an unwanted ambition may be exercising the very revisability the system was intended to support. That observation must not be automatically labeled either technical failure or personal improvement. Its meaning depends on reasons, alternatives, and subsequent consequences.

3.4 Experiences can help constitute a purpose

An emerging concern can change what a person seeks, not merely reveal a previously hidden preference. Three situations should therefore be distinguished. In discovery, a person learns an unknown fact about an already valued activity. In formation, an encounter or shared undertaking helps create a commitment that did not previously organize their life. In domination, another party shapes the available encounters, dependencies, or descriptions so that the resulting commitment primarily serves that party while suppressing meaningful contestation. These situations can coexist. A changed preference, sustained satisfaction, or later endorsement alone does not distinguish them.

Paul's analysis of transformative experience challenges decisions based on confidently projecting what an unfamiliar, life-changing experience would be like for oneself [43]. We use this challenge to reject a guarantee of accurate first-person prediction, not rational deliberation itself. Testimony, practical constraints, obligations, and opportunities for learning remain relevant. This is a philosophical constraint on a design assumption; it is not evidence that an AI system improves such decisions.

Consider a skilled technician who delegates invoicing to an assistant and accepts an invitation to help a neighborhood theatre. Initially the purpose is a small repair. Working with performers, discovering a neglected style, and becoming responsible to collaborators may make sound design a new interest. Paid technical work, family time, friendship, and artistic ambition now require renegotiation. The technician may pursue the practice, retain it as an occasional pleasure, or leave. Success cannot be identified with maximum

artistic output, continued attendance, or a transition to a creative career. Each can be appropriate in one life and intrusive in another.

This case yields a positive role for intelligence. The assistant can make the encounter possible, handle burdensome logistics, explain an unfamiliar technique, help the group realize a performance, and test whether a new ambition is compatible with existing responsibilities. These are forms of active assistance and actual value creation. It should also help the technician articulate a changed purpose and offer collaborators a chance to respond when that change affects them. Supporting agency need not mean withholding initiative until a person can fully specify an unexperienced future.

The difficulty is that the assistant participates in forming the very preference it might later use to justify its intervention. A recommender rewarded for subsequent endorsement can select environments that make its own suggestions difficult to reject. A benevolent system that always discourages disruption can instead deny occasions through which commitments would have formed. Taihe therefore proposes a comparative judgment: examine how the occasion was presented, what alternatives and dependencies existed, who benefited, what reasons participants could examine, and whether they could subsequently redirect the undertaking. No single signal certifies an authentic self. The person need not supply a polished rationale before being helped; affect, hesitation, play, and silence may motivate a tentative inquiry without authorizing a major commitment.

Truth and love constrain each other here. A companion can share delight and offer imaginative rehearsal without claiming that a simulated reception predicts how other people will respond. Comfort may be chosen and fictional play may be knowingly enjoyed. Fabricating a collaborator's consent, concealing a material failure, or interpreting affection as permission to decide a person's future destroys information needed for the relationship itself. Conversely, factual accuracy delivered without regard for vulnerability can obstruct understanding and care. The proposed response combines truthful distinctions, attention to timing and context, and practical support; it does not promise that every conflict has a costless resolution.

Lingji has a specific place in this account: an unexpected concern can initiate a change in the question, the undertaking, or what matters about its outcome. Rational inquiry then helps examine consequences and competing explanations, and the encounter may in turn revise the concern. This interaction provides a reason to build systems that can both accomplish current tasks and accommodate the formation of new purposes. It does not imply a mandatory cycle of reinvention. Rest, familiar affection, faithful continuation, and receiving care can remain ends in themselves.

4 Six conflicts that the proposal must decide

The distinction among capability, purpose, responsibility, and form does not decide every conflict in favor of change. The cases below specify which reasons can support continuity, delegation, exploration, or restraint, including consequences for people beyond the focal user.

4.1 Efficient delegation and chosen growth

When the same artifact can be produced automatically or through practice, ask which outcome and mode the person currently wants. Offer delegation for unwanted burdens and assistance for chosen learning; permit a scoped change of mode when circumstances change. The relevant comparison includes final quality, retained skill, effort, and whether a desired contribution actually occurred. An emergency may justify prioritizing relief, while an ordinary deadline does not license silently completing a person's chosen learning task. A growth-oriented system fails its own purpose if it withholds needed help to enforce independence.

4.2 Protection and exploratory failure

Prefer a reversible rehearsal when failure is affordable and chosen, with a route to recover and explanations of consequences. Where an experiment can seriously affect nonparticipants or exhaust shared resources, the person's willingness alone is insufficient. Reduce exposure, negotiate shared conditions, or decline the particular intervention while preserving other routes to the underlying purpose. This is a proposed ordering of reasons, not a universal numeric threshold. The system must disclose what limiting an exploration prevents; hiding all failure to achieve a safety metric would undermine growth.

4.3 AI companionship and human relationships

A responsive companion can help someone rehearse a difficult conversation, find collaborators, remember a shared promise, or enjoy time alone. Its value should be evaluated against the person's wanted relationships and actual circumstances. Maximizing time with the companion can compete with those purposes. Offer connection and support without secretly impersonating a human, manufacturing exclusivity, or conditioning help on attachment. Conversely, compulsory social contact can burden someone who prefers solitude. A better relationship with an AI does not automatically imply a worse relationship with people; neither assumption replaces evidence about both.

4.4 Individual diversity and common norms

Different practices and tastes warrant space even when a central model rates one style more highly. Preserve distinct local interpretations and a way to make or choose alternatives. But a local preference does not settle effects on outsiders. A shared constraint needs a stated reason, affected-party representation, and an opportunity to challenge its application. When purposes conflict, describe what each option preserves and loses instead of relabeling the winner as the only authentic preference. Some disputes remain unresolved. A decision process can be justified provisionally without claiming universal agreement or a complete solution to social choice.

4.5 Personal benefit and public consequences

A private assistant may improve one person's bargaining power while reducing another's options. That concern is already studied in multi-agent empowerment [14]. Taihe therefore asks who supplies care, data, cultural material, compute, land, and attention, and who receives the benefit. For optional intensive activities, compare lower-cost alternatives and make uncompensated burdens visible. Resource allocation cannot be solved by adding separate personal benefit scores when they omit affected people or use incompatible scales. In scarcity, a declared priority for essential care is a normative proposal; which needs qualify and how priorities are revised require legitimate institutions, not an agent's unilateral optimization.

4.6 Virtual experience and material life

Digital experiences can involve real learning, friendship, cultural expression, grief, and work. Their value does not depend on every object being physically present. Yet a rendered meal does not feed a body, simulated payment does not settle an external obligation, and visual plausibility does not establish physical truth. Maintain these distinctions and let participants develop connections across settings. A virtual world becomes a poor substitute when it is offered to conceal material deprivation or remove a wanted route back to physical life. A blanket judgment that games are meaningless would also be unjustified: consequences and participation, rather than medium alone, determine the question under study.

These six decisions have a common procedure: identify the claimed goods and affected people; distinguish factual disagreement from value disagreement; consider reversible or less burdensome alternatives; decide within an appropriate scope; and revisit the decision after observing consequences. The procedure is defeasible and does not guarantee agreement. Its purpose is to expose reasons and losses, not to transform plural life into a compliance score.

5 Culture, embodied practice, and philosophical resources

A reusable technique receives its human significance through an activity and its relationships. This section examines why preserving usable means can support cultural renewal, while preserving a particular form or the record of its production can sometimes frustrate it.

5.1 The content of a life cannot be supplied by capability alone

A place is encountered through histories, routes, craft, ecology, language, weather, and other people's activities. An editable landscape without these relationships may increase visual variety while offering little cultural understanding. Music involves listening, timing, bodily practice, and interpretation as well as an audio file. A gift involves a relationship and intention as well as an object. Ordinary works and local knowledge deserve room alongside exceptional art and scientific discovery.

The inference from plural lives is that people need ways to participate in, reinterpret, and create cultural practices, rather than only consume standardized representations. This is not a demand to preserve every tradition unchanged. A community contains disagreements, and cultural authorities can restrict individual expression. Record whose interpretation is being offered, make room for counterinterpretations, and distinguish respectful learning from claims of authoritative representation. World production must support editable practices and histories as well as attractive surfaces. Whether a specific representation succeeds needs participants and domain expertise; visual preference scores are insufficient.

5.2 Daoist, Buddhist, and heart-mind traditions as distinct interlocutors

The following readings are selective and situated. They neither describe a single timeless Eastern doctrine nor derive an AI architecture from an ancient authority. We identify a text, a limited interpretive relation, and the modern inference for which this paper takes responsibility.

The received *Laozi*, chapter 2, juxtaposes relational distinctions with activity that does not claim possession or dwell on achievement. The consulted Heshang Gong recension includes its own commentary, which must not be mistaken for the base text [34]. Our design analogy is enabling another's development without claiming ownership of their future. This is a modern inference. The text does not specify digital property, compensation, or democratic institutions; nonpossessive assistance cannot by itself adjudicate those questions. Nor does nonaction mean abandoning someone who needs help.

The *Diamond Sutra* in Kumarajiva's Chinese translation discusses a mind not fixed on sensory objects in a setting of bodhisattva practice; the relevant passage is Taisho T235, 749c21–24 [35]. We draw a limited analogy to refusing to equate a person with a representation of them. This does not turn Buddhist nonattachment into a memory-deletion algorithm, imply that reality is merely a simulation, or prove a physical theory of emptiness. Releasing a fixed interpretation can coexist with practical responsibility and care. The Buddhist soteriological setting exceeds this engineering analogy.

Wang Yangming's *Chuanxi lu*, the Xu Ai records, links heart-mind, moral concern, and the unity of knowing and acting, while criticizing both unreflective action and detached discussion [36]. We take from this an invitation to ask whether professed care changes conduct. This is not a claim that inward conviction establishes scientific truth or that a language model has innate moral knowledge. Its moral and historical setting differs from the procedural testing in Appendix B. Consequences still require observation and other people may reject the interpretation.

These resources illuminate different tensions: enabling without possession, avoiding rigid attachment to representations, and making understanding consequential in practice. They do not eliminate disagreement among traditions or between tradition and modern commitments. A comparative philosophical account would need additional editions and specialist criticism. Here their role is explicit argument formation, not evidence that the technical mechanism works.

5.3 From a lived concern to a practice others can change

Consider a constructed case, not a report of fieldwork. A local instrument maker wants less administrative work and more shared music. A young player cannot use the standard grip. The maker notices that a supposedly inferior, unstable tone becomes expressive when two players coordinate. A learner elsewhere later encounters different materials and asks for a quieter instrument. This case connects six transformations; none follows automatically from the preceding one.

First, experience makes something matter. Discomfort, affection, frustration, or delight can direct attention to an exclusion that a specification omits. The player's discomfort motivates a question about access; pleasure in an unusual sound motivates a question about expression. Neither establishes an acoustic fact. Truthful inquiry separates those concerns from claims about resonance, fatigue, and reliability. Care matters at the point of selecting what deserves investigation, as well as when responding to its results.

Second, a concern becomes a revisable question. The problem might change from producing the cleanest solo note to enabling two people to make a meaningful phrase together. This could be a change of purpose, an explanatory conjecture about interaction, or both. A trial must state which. Comparing grips can test accessibility under specified conditions; it cannot decide whether a shared musical practice ought to replace solo performance. Lingji motivates a departure; rational inquiry articulates alternatives, exposes error, and

sometimes shows that the departure should be abandoned. Courage permits a reversible attempt despite uncertainty. It does not warrant imposing its risks on the second player.

Third, a question becomes action with consequences. A feasible design is insufficient if no one can learn, maintain, or actually play it. The maker, players, accessible tools, and a place to meet jointly supply enabling conditions. Delegating scheduling may support practice; automating the entire performance may frustrate this particular purpose. Someone who only wants to listen can instead benefit from the automatic performance. Youthful exploration here means permission and support to encounter possibilities, including failure, without requiring a young person to become an innovator or an older person to remain a fixed authority.

Fourth, repeated action becomes a shared practice when participants can coordinate, interpret consequences, and contest what counts as a good performance. A finished recording alone does not contain that relation. A distinctive local style may make collaboration intelligible; the same style may exclude a newcomer. Recognition of the maker's contribution therefore coexists with the learner's ability to question it. Mackenzie's account already connects temporally extended agency, shared narratives, and social constraints; it also explains why an apparently voluntary group can reproduce oppression [38]. Our case must retain that double possibility, rather than treat community membership as sufficient evidence of freedom.

Fifth, a practice becomes transmissible through some combination of examples, explanations, tools, interaction, and institutions. The relevant content may be a short contrast between two materials, a manipulable model, or time with a teacher. It need not be a complete personal archive. A failed sketch can teach a boundary condition if its failure is understood; an abandoned sketch is not thereby a negative training label. Reasons can be mistaken, tacit skill may resist explanation, and a canonical account can suppress alternative interpretations. Transmission is a selective reconstruction whose adequacy depends on what later learners need to do.

Sixth, inheritance becomes renewed capability when another person can use, question, or transform what was received in changed circumstances. Copying a sound tests reproduction. Adapting to a new material tests transfer. Choosing a different social use changes the end being pursued. These outcomes require different evidence. Yaman, Tian, and Lindstrom study semantic knowledge and social learning in a combinatorial innovation task; their findings support investigating structured understanding, but their fixed recipe space does not establish open-ended cultural renewal or the efficacy of our proposed transmission design [40].

The return to concrete lives completes the argument. Shared knowledge can make care less burdensome, tools more accessible, and previously difficult collaboration possible. Later people contribute new concerns, or choose to receive these benefits without contributing. Civilizational development in this account includes sustaining the conditions for such lives across time. Increasing an archive or a model's capacity while extinguishing those conditions is a different achievement and must be judged separately.

5.4 A failure case that the account must not romanticize

Suppose the maker's archive treats one family as the sole legitimate interpreter, includes intimate disputes, and requires learners to study hours of commentary before using an accessible instrument. A system that faithfully preserves everything can strengthen exclusion and exhaust its users. A compact automatic tool, with little process history, can then better support both care and exploration. Even accurate context may be unnecessary when a successor can cheaply infer the mechanism or obtain equivalent instruction elsewhere.

The response is selective, revisable transmission: preserve enough usable structure for the chosen activity and plausible changes, expose uncertainty where it matters, and allow alternative teachers, interpretations, and modes of assistance. Do not equate historical density with cultural depth. Nor should a community be forced to expose restricted knowledge to prove its commitment to civilization. Incompatible wishes about sharing require negotiation and sometimes a narrower shared resource. The theory gives a reason to consider future use, not a universal license to extract lives.

5.5 A different case: encounter, affection, and a shared journey

A second constructed case concerns friends taking a journey together after years apart. One wants to sketch an unfamiliar landscape, another hopes to meet local people, and a third wants companionship without

exhausting travel. They delegate bookings, accessibility planning, and paid work that would otherwise prevent the trip. An AI can usefully initiate a suggestion for a shared meal or an unexpected detour, prepare alternatives, and help make the encounter happen. Its initiative can enlarge life rather than merely prevent harm. The proposal remains answerable to the people affected; a plausible prediction of enjoyment is not their agreement to a commitment.

The journey makes room for love, beauty, adventure, ordinary inconvenience, and memories created together. None is guaranteed by a perfect itinerary. A disappointing sketch can be a valued part of learning; a missed event can be an avoidable failure; accepting a modest risk can express exploration while imposing it on a reluctant companion would frustrate the shared activity. A person may instead choose a fully arranged, restful trip. The relevant distinction is between burdens they wish to delegate and features that constitute this experience for them. The system must discover that distinction through interaction, not infer it once from personality labels.

Three kinds of continuity now come apart. **Output continuity** preserves the itinerary, images, or a finished work. **Practice continuity** preserves the ability to undertake an activity and relationships through which it matters. **Continuity with room for departure** allows later participants to question the inherited route or purpose. A synthetic album can preserve the expected outputs without the friends having traveled. A memorable trip can sustain a relationship without an archive or a desire to repeat it. An archive can teach a newcomer a route while its social rules forbid deviation. None of these dimensions entails the others, and maximizing all of them is not obligatory. A transient encounter may be valuable precisely without retention.

The distinction has immediate technical consequences. A system should model an instrumental task, a chosen mode of participation, and optional transmission as separable specifications. A logistics learner may reuse successful scheduling without importing one person's interpretation of a friendship. A memory system may preserve an agreed event without declaring a single account of what it meant. A planner can propose encounters as well as optimize a route, and can offer delegation without treating it as a loss of dignity. Work, earning income, and reliable service to others remain valuable possible purposes; the theory gives no warrant for substituting an artistic adventure for a user's practical livelihood.

Capability and relational theories can already explain much of this case. Our contribution is a synthesis with a precise failure distinction for AI systems: improved delivery, meaningful enactment, and reusable inheritance can diverge even within the same observed project. This changes what to learn, what to predict, what to execute, and what to retain. The companion, runtime, and production system must preserve these distinctions across their interfaces. Their successful implementation is an engineering and human-evidence obligation, not established by introducing the vocabulary.

6 From shared intelligence to plural civilization

Individual delegation and collective provision answer different questions. The freedom to leave a task does not settle what shared means should remain available, while a valued inheritance does not entitle a community to a contributor's lifelong labor. The civilizational argument follows this distinction into institutions, material access, and future lives.

6.1 Shared capability and different futures

Claim H6: collective improvement should be assessed together with the conditions under which distinct people and communities can develop differently. The argument joins situated capability to revisability. A shared model can lower the cost of knowledge and creation, yet uniform defaults and concentrated control can narrow practical choices. Distinct histories and practices are not noise that should necessarily be distilled away. At the same time, incompatible technical infrastructures can prevent useful exchange. Shared capabilities and interoperable interfaces can coexist with divergent interpretations, provided the latter remain practically usable rather than decorative options.

This is not an argument against shared learning. It is a question about its terms: what may be learned, how corrections propagate, who can use the result, whether contributors are recognized, and what happens when a community no longer wishes to participate. Withdrawal may not erase every previously learned

statistical effect; portable works also do not guarantee access to compute. Such limits belong in institutional comparisons rather than being hidden behind the language of control.

6.2 Conditional scenarios, not predictions

We vary two conditions: the distribution of control over enabling infrastructure and the practical availability of revisable ways of life. Neither is determined by an AGI capability level. The following scenarios are arguments about consequences under assumptions, not simulations or forecasts.

Table 2: Conditional institutional scenarios and their failure pressures

Institutional candidate	**Potential benefit under stated assumptions**	**Failure pressure and comparison question**
Concentrated provision with public obligations	Shared infrastructure can reduce duplication and supply reliable access if obligations are enforceable	Defaults and exit costs may dominate local choices. Can obligations be challenged and works transferred when interests conflict?
Public or cooperative enabling resources	Broad access and collective voice can support care and creation if maintenance and allocation are funded	Scarce resources and majorities can suppress minority practices. How are essential provision, dissent, and innovation sustained?
Federated communities with portable works and negotiated exchange	Local interpretation can coexist with shared technical capabilities if interoperability is real	Wealth differences and cross-community harms may persist. Who supplies common infrastructure and resolves spillovers?

The proposals can be combined. We do not select one as the author's settled political program. Evidence would include access, effective exit, distribution of maintenance work, ability to contest rules, and durable cultural production, measured in their own domains. A benchmark score cannot choose an institutional form. If a scenario assumes abundant digital output, it must still account for bodies, ecosystems, finite attention, and care. Monetary compensation is one instrument, not an equivalent of every relationship or contribution.

6.3 Responsibility to the vulnerable, future lives, and ecology

Concern for particular people must extend to those absent from a negotiation. The proposed priority of a viable basis for life and care follows from standing that does not depend on productivity; it does not follow from a forecast that machines will provide abundance. Practical policies could include access obligations, pooled resources, compensation, and accountable stewardship, each with costs and failure modes. Scarcity can leave tragic choices. A responsible framework records whose interests are unmet instead of declaring that optimization has harmonized them.

Ecological attention is also substantive. Compute, materials, land, energy, and the effects of constructed environments on other living beings must enter decisions about production. No energy or ecological result is reported here. A digital demonstration cannot establish environmental sustainability. For future artificial entities, persistence, learning, and expressive behavior do not settle consciousness or moral status; further evidence and argument are required. We do not grant a developer ownership of a morally significant future merely by calling it a creation, and we do not presume that present software already has that significance.

Gratitude, humility before uncertainty, compassion toward vulnerability, and freely undertaken responsibility can inform these choices. They should not become mechanisms for assigning people hereditary ranks or demanding sacrifice. A person may choose costly care or a minority cause; an institution must not infer that such a choice authorizes others to impose the same burden. Civilization's growth is thus assessed by what is sustained and made possible, by what is consumed, and by who retains a say in its direction.

6.4 Renewing the conditions for later agency

Claim H7: when shared intelligence grows through people's activities, its assessment should include whether contributors and successors retain feasible ways to understand, redirect, and renew relevant practices. This is a normative elaboration of H3, H4, and H6, not a scientific law.

Standing remains unconditional: no individual owes a creative contribution in return for care. The duty under discussion concerns the organization of shared resources, not a demand that every life be productive.

The argument has three steps. Situated agency depends on enabling conditions, not simply permission. Those conditions include learned practices, tools, relationships, and accessible knowledge that can be lost even while present output improves. If future people also matter and purposes may change, evaluating only current delivery systematically omits that loss. Therefore present consent, compensation, and competent delivery can be necessary features of a good arrangement without exhausting its justification. This conclusion is conditional on the stated commitments. It is compatible with, and substantially indebted to, capability and relational accounts; it does not establish a new independent foundation for ethics.

The full-agreement case in Section 1.3 makes this assessment concrete. Individuals may reasonably delegate their own contributions while the organization has a reason to sustain a shared route for later making. The distinction separates an individual's work arrangement from provision of a capability that no one individual can supply alone. It does not infer a duty to preserve every original workshop or cultural form. An alternative tool, service, teaching arrangement, or partnership can satisfy the relevant purpose; the choice depends on its actual accessibility and burdens.

The harder intergenerational decision concerns people whose purposes are not yet known. A credible argument must identify a range of activities or dependencies worth sustaining and explain the competing claims on its upkeep. A speculative possible user does not establish priority over urgent care, and present unanimity does not establish that later access is irrelevant. Where the relevant facts or priorities remain unsettled, the account supports stating the disagreement and seeking proportionate, revisable provision rather than announcing an optimal inheritance. A contributor can retire even when a shared capability matters; institutions must arrange willing support, choose an alternative, or acknowledge the loss.

This yields a conditional civilizational implication. Greater common intelligence can relieve unwanted work while widening the means through which different people develop and enact purposes. It can also make present delivery convenient while narrowing those means. Which process occurs depends on institutions, material access, relationships, and editable or substitutable technical resources. Aggregate output and universal agreement to the current service do not settle that comparison. Neither does the number of preserved artifacts. The inference concerns conditions for plural lives; it is not a prediction that growing intelligence necessarily improves or degrades civilization.

6.5 Shared capability can support difference or replace its conditions

Pooling knowledge can reduce repeated effort and help a newcomer escape dependence on a local gatekeeper. Distinct communities can then redirect the same general capabilities toward different activities. Conversely, a common model may improve average performance while its interfaces, economics, or defaults make a minority practice infeasible. Neither the number of model variants nor whether weights are centralized settles this question. The relevant comparison follows the activities people can actually undertake, the burden of changing direction, and the distribution of maintenance work.

A possible self-reinforcing process is that popular practices receive better tools and more visibility, which further increases their recorded use; a service trained on that use can then mistake constrained exposure for settled preference. This is a causal hypothesis, not an observed result of Taihe. An intervention would have to vary access and exposure and examine sustained choices, not infer oppression from unequal usage alone. Some convergence reflects informed shared preference and should not be artificially reversed. Similarly, different outputs do not demonstrate independent agency if meaningful alternatives are unavailable.

The same reasoning reaches beyond work. Receiving care, listening, friendship, spiritual practice, playful exploration, and rest can be worthwhile even when machines create every technically superior artifact. The possibility of contributing to common intelligence is an additional path, not a replacement justification for human existence. A civilization that rewards novelty while neglecting dependent people fails the standing and care premises even if its discovery rate rises. A civilization that provides comfort while systematically closing wanted routes of inquiry and association fails a different part of the account. These are distinct criticisms; neither can be absorbed into a productivity ranking.

6.6 How to challenge the civilizational bridge

The six transformations in Section 5.3 are a proposed explanatory structure. A local study can examine one transition; it cannot verify the entire civilizational claim. A useful human study would compare transmission through finished examples with transmission that also exposes selected context, reasons, and operations. Equalize teaching time, interaction opportunities, and usable information budget; pretest legibility and accessibility. Include a matched active-teaching control so that extra attention is not credited to process preservation.

Assess reproduction first, adaptation to changed materials or constraints second, and open-ended creation separately with declared, independent criteria and participant explanations. Measure retention, burden, exclusion, and whether participants can reject an inherited purpose. Use separate cohorts or appropriately blocked assignments to avoid treating members of one transmission chain as independent. Delayed transfer matters because immediate success may be imitation. A larger history can lower transfer through overload or fixation; that outcome would challenge the mechanism proposed here. All of this remains a study design: no participant recruitment, sample-size calculation based on pilot variance, ethics review, or human results have been completed.

Institutional and ecological claims need additional evidence: who maintains access, how costs move between groups and generations, and what happens when infrastructure fails or ownership changes. No single experiment settles those questions. This separation preserves the ambitious thesis while preventing a local result from being promoted into proof of a desirable civilization.

6.7 From formed commitments to shared futures

The theatre example also changes what cultural transmission should preserve. A recording preserves a performance; a manual can preserve a technique; accessible tools and willing collaborators can preserve a practice. None guarantees that successors will find the same purpose worth pursuing. A community can inherit sound design as ritual, entertainment, protest, paid work, or an experiment the original contributors never anticipated. The ability to form a new shared purpose depends on material access and social relationships as well as information. Preserving every recording cannot substitute for those conditions.

This sharpens H7 without supplying a new ethical foundation. Capability theory already gives reasons to support substantive opportunities; care ethics asks who maintains the practice and bears its burdens; relational autonomy examines how social conditions enable or constrain commitments. Taihe brings these questions to the design of intelligence that helps select encounters, generate works, learn from participation, and distribute the resulting capabilities. The synthesis changes a concrete engineering judgment: a system cannot count a successful current service and a satisfied returning user as sufficient evidence that its learning policy supports the formation and revision of worthwhile purposes.

An arrangement that obtains excellent performances by continually depending on one exhausted organizer may improve output while undermining care. An arrangement that archives all contributors' intimate exchanges may support reconstruction while damaging trust and willingness to participate. An arrangement that automatically preserves only the most popular style may make average production easier while eroding the practical conditions for another tradition. Each calls for a different response: shared maintenance or paid support, selective and agreed transmission, or accessible routes to distinct practices. None warrants a blanket duty to preserve everything. A disliked inherited practice may reasonably end, and successors owe no reproduction of their predecessors' purposes.

For post-AGI civilization, the normative claim is consequently about the organization of shared power: expanding technical capability should be accompanied by feasible ways for diverse people to help direct what that capability makes possible, while care and standing do not depend on their contribution. Whether systems actually produce dependency, plural experimentation, or renewed cultural capacity is a causal question requiring evidence. Whether those outcomes are desirable requires the stated commitments and public contestation. Neither question is answered by assuming that more intelligence necessarily produces moral progress.

Two limited studies would test parts of this argument. A task study can compare a fixed-preference assistant with equally capable assistants that support revisable purposes and a strong ordinary deliberative control. Assess practical completion, unwanted burden, whether a changed concern reaches consequential

action, effects on collaborators, and the ability to leave or delegate. Treat changed preferences and continued engagement as observations, not success labels. A transmission study can then test whether another group can adapt the resulting tools and practices for a different declared purpose, separating information access from the extra attention of a facilitator. These designs require participants, consent, suitable allocation and follow-up, and an independently justified analysis plan. They have not been conducted. Their outcomes could challenge the proposed mechanisms while leaving the normative commitments open to argument.

The technical experiment below has a narrower role. It asks whether a memory retains evidence relevant to specified explanations. A person's love, purpose, cultural belonging, or worth is not an element of that hypothesis class. The connection is a design warning: when purposes and questions can change, optimizing information solely for current questions can undermine later inquiry. Establishing that warning does not establish an algorithm that supports a whole human life.

6.8 Personal lives, joint activities, organizations, and civilization

The unit of assistance need not be a visit to an AI application. A person can ask for a change while playing, designing, repairing equipment, or carrying out work with others. Observation, assistance, production, and learning can be embedded in that activity. A separate interface remains useful for reviewing complex changes, resolving disputes, or inspecting what was learned. Neither invisibility nor the absence of an application establishes better support: the relevant comparison is what people can accomplish, understand, and redirect, with what interruption and burden.

Situated and extended cognition supply demanding precedents, including attention to reliable external supports, practical interpretation, and moral or material consequences of human-machine arrangements [44, 45]. The comparison in Section 10.1 treats those contributions as substantive. In the present design problem, a system's recorded events may capture only part of what the activity means to those involved. A representation of a purpose must therefore remain open to correction in use; adding more sensors or removing an application boundary does not by itself establish adequate understanding. This requirement concerns the relationship between representation and activity, and does not require settling the metaphysical boundary of mind.

Moving from a person to a joint activity adds other people's purposes and mutual expectations. Moving from an activity to an organization adds durable roles, decision procedures, maintenance obligations, and successors who were absent from the original event. Moving from organizations to civilization adds interdependence, unequal access, long time horizons, and people without representation in a particular institution. Each transition introduces a justification problem. None can be solved by treating an organization as one enlarged person or averaging its members' preferences.

Consider a repair team with a service deadline and an apprentice who wants supervised practice. An assistant can diagnose the fault and a robot can perform the repair faster. If adequate supervised time exists, automatically taking over may improve completion time while removing the chosen learning opportunity. Retaining a control button is insufficient if the scheduling system never makes a practice slot feasible. An alternative is to automate preparation and documentation while a willing experienced colleague supervises the apprentice on an appropriate part. This choice serves both service and learning only under the stated conditions; the supervisor's labor and the customer's deadline are real costs, not free support.

Change the case. The team faces an urgent service failure, the apprentice requests full delegation, or the experienced colleague is exhausted. Automated repair and an optional later practice session can then expand freedom and protect care. Compelling manual participation would frustrate the same humanistic premises. The role of embodiment is concrete: a digital rehearsal may be reset, while a physical intervention can damage equipment or affect bystanders. Success in the rehearsal does not supply evidence of physical competence. The worker should be supported within the ongoing activity, but that does not remove necessary inspection or transfer another person's authority to their personal agent.

The choice can be made explicit without pretending to calculate a whole life. Let a plan have task quality q, elapsed time t, each affected participant's burden b_i, and a feasible participation opportunity p_i. In a particular agreed learning task, conditions such as q at least $q_{\min}$, t at most the deadline, and $p_{\text{apprentice}}$ at least the requested opportunity define admissible plans. Minimizing time over that set differs from minimizing time alone. If the opportunity condition is withdrawn by the relevant person, or the constraints

cannot jointly be met, the response is delegation or deliberation about the conflict; the system must not silently erase a condition to report success. These quantities are task specifications. Their definitions, thresholds, and authority come from the situation, not a universal humanistic score or a new optimization theorem.

Learning from this activity requires distinctions with consequences. A successful diagnostic sequence may become a shared skill after appropriate review. The apprentice's personal aspiration need not become an organizational requirement for every recruit. A supervisor's one-time compromise need not become a permanent expectation of unpaid teaching. When members leave or roles change, the organization needs a legitimate way to retain useful operational knowledge while personal accounts and relationships remain contextual. An instruction from a manager, a worker's preference, a customer's requirement, and a sensor observation are different kinds of evidence; collapsing them into one memory can produce technically competent but wrongly directed action.

The link to civilization is a conditional generalization of these problems, not a claim that a team is a miniature civilization. Shared intelligence can relieve repeated burdens, make cultural creation accessible, and help newcomers acquire capabilities. It can also concentrate opportunities to decide, remove the settings in which expertise is learned, or make institutions dependent on a service they cannot meaningfully redirect. H7 therefore asks about usable paths for later participation and renewal alongside present delivery and care. Particular people may choose delegation, continuity, adventure, or departure. Preserving these different possibilities does not require preserving every current occupation, ritual, or manual skill.

Evidence should follow the distinctions. Compare embedded and separate entry points while holding capability, observable information, permissions, and budget fixed; measure completion, interruptions, context reconstruction, and correction. Independently compare learning enabled and frozen in the same entry point. For organizational transfer, change roles or members and assess legitimate task completion and whose burdens or opportunities changed. For embodied transfer, verify physical observations and consequences separately. No improvement in one comparison establishes the others. These studies are proposed, not conducted, and the present scripted discovery adapter implements none of the repair-team or organizational scenarios.

7 A bounded technical account of discovery and realization

The humanistic argument is broader than the formal objects in Appendix B. The bounded specification concerns making a revised question or selected activity consequential across changing technical components. The mathematics defines feasibility and a restricted equivalence result; it does not prove the normative premises, measure love, or certify civilization. The full discovery and synthesis method remains unimplemented. Existing executable diagnostics test only the boundaries stated in Section 9.

The design requirement is to check which part of an activity should be delegated, which contribution someone wants to make, and what a successor must be able to use or change. A planning specification can represent an outcome and a selected participation mode separately, while an evidence record can distinguish a world fact from a personal or collective commitment. These are consequences of the argument, not proof of a new algorithm. Appendix B gives the bounded state, opportunity, and transformation definitions, with explicit feasibility and cost assumptions.

Two mathematical boundaries matter. First, coupled changes to agent capabilities, world rules, and produced content have no reachability advantage over a sequential method that can stage the same primitive edits at the same cost. Second, in the tested nested hypothesis classes, contradiction-triggered expansion selects the same surviving class as a minimum-degree sequential control. Both results remove putative superiority claims. The evidence and its implications are summarized in Section 9; detailed derivations are in Appendix B.

8 How advanced systems bear the humanistic requirements

8.1 Technical means for changing purposes and usable practices

The theatre case requires several distinct technical abilities: generate alternatives, anticipate consequences, make a proposed change executable, observe what happens, and retain or replace the means of later

adaptation. A better result in one role does not complete the others. World models can assist prediction and rehearsal; causal explanations require intervention-sensitive evidence and assumptions [28]. A convincing rollout is not an observed event, and a prediction of reception does not determine what a person should value.

Nuwa studies persistent life, experience, agency, and improvement. Its role includes acquiring reusable abilities while distinguishing a remembered event, a current purpose, and an established responsibility. Voluntary self-distillation allows people to actively express revisions and aspirations; assistance should count the effort and mistaken representation this involves. Fuxi studies sustained interactions, A2A commitments, causal cascades, and reliable execution. Its role is to make changes consequential in shared activity while preserving the distinction between prediction and actual events. Pangu studies reconstruction, generation, editable production, and learning from creation. Its role includes producing artifacts and operations that people can use, change, or replace to realize different works.

These roles jointly support a possible path from an emerging concern to a usable practice. They do not establish that the path has been integrated or that its participants benefit. For a particular change, identify the missing means, the relevant interfaces, who must cooperate, and the cost of trying and maintaining it. The same general capability can support practice, co-creation, observation, or complete delegation. The chosen relation to the activity belongs in the target of assistance, rather than being inferred from output quality alone.

Studio, Canvas, and improvement mechanisms can preserve sketches, branches, decisions, and feedback during creation. An abandoned draft is not automatically a negative example; its reason and scope matter. Body, movement, sound, geometry, semantics, rules, and rendering each need implementation and validation. Preserving a project does not by itself preserve an executable adaptation path. Flow, Seek, and Land name a product layer; Nuwa, Fuxi, and Pangu span those products and are not three corresponding application names. Assistance and production may also operate inside an ongoing game, workplace, or physical activity.

8.2 AI 1.0 to 4.0 as the author's motivating framework

The labels below are an authorial thought experiment, not a consensus history, measured ranking, or inevitable sequence. AI 1.0 originally refers only to current practice; assigning it a unique statistical paradigm would exceed that definition.

Table 3: The author's AI-stage framework and its evidence obligations

Label	Meaning in this framework	Evidence obligation
AI 1.0	Current AI practice, deliberately broad	Name the actual model, task, tools, data, and budget
AI 1.5	World models anticipating world change	Held-out action-conditioned prediction and useful planning
AI 2.0	Causal models of mechanisms and intervention	Discriminating interventions, transfer, and stated identification assumptions
AI 3.0	Shared human-AI worlds in which emotion, values, impulse, and Lingji may redirect questions	Equal-information discovery tests and separate evidence about chosen purposes
AI 4.0	AI designing AI, including mathematical methods, architecture, and potentially physical materials	Independently checked designs; improved improvement procedures need additional held-out and longitudinal tests

Four paper names do not map onto four stages. AI 3.0 is not a necessary prerequisite for AI 4.0; automated design can occur without a human-AI society. The research opportunity is their interaction: changed purposes may redirect which abilities are developed, and improved abilities may enable previously inaccessible lives. No historical inevitability follows.

8.3 Learning and learning to improve

Memory updates, learned skills, tool changes, model-parameter updates, and changes to an improvement procedure are distinct interventions. Their effects must be traced separately. Learned optimizers, self-modifying agents, and open-ended archives already precede this program [29, 19, 5, 15]. Saving a draft or improving a task score is not by itself recursive self-improvement.

For an improvement-procedure claim, start the old and proposed procedures from the same system and permitted evidence. Freeze the procedures before final evaluation, allow each the declared learning process, and compare the systems they produce on task families withheld from both design and selection. Count failed proposals, preparation, judging, training, and execution. A recursive claim additionally requires later generations to use the revised procedure to produce subsequent revisions, with fresh evaluation and a versioned causal lineage. Compare to a frozen designer allowed the same cumulative resources. A stronger inherited model, extra practice, or repeated access to test answers does not isolate the improvement procedure.

Proposals about mathematics need explicit assumptions and independently checked proofs or algorithms. Architecture changes require appropriate performance and resource studies. Materials claims need domain validation beyond generated descriptions or unvalidated simulations. None of those broad AI 4.0 capabilities is implemented by this manuscript. Their horizon is preserved without treating it as a result.

9 Evidence appropriate to each claim

9.1 Four claims, four kinds of scrutiny

Table 4: Four kinds of claim and their appropriate scrutiny

Claim	Evidence supplied in this version	Required next scrutiny
T1: a coherent humanistic position	Premises, claims H1–H7, conflicts, objections, textual interlocutors, and design consequences	Criticism of premises, conflicts, counterexamples, and unsupported interpretations
T2: useful discovery through frame revision	Bounded specification; failed finite candidate and sequential equivalence (Appendix C.2); no external discovery result	Implemented method, capable revising baselines, discriminating tests, held-out transfer, and full cost
T3: actual significance for participants	Constructed scenarios and a human-study design; no participants studied	Chosen purposes, consequences, later endorsement, unwanted influence, effort, relationships, and failures
T4: a defensible civilizational direction	Conditional institutional scenarios and normative arguments	Assumption stress tests, distributional and ecological evidence, alternative institutions, and informed criticism

The table separates normative, technical, experiential, and institutional claims. No objective civilization score is proposed. Mechanism evidence, lived experience, and institutional consequences can inform each other but cannot replace one another.

9.2 Technical comparisons and their proper scope

A mechanism intended to support later making should be compared with capable contextual and reflective controls that receive the same history, model, tools, adaptation tests, and total resources. Section 1.4 specifies the proposed intervention and its possible failure. Before testing, fix development and held-out units, success criteria, failure handling, and the accounting of search, preparation, evaluation, and maintenance. A gain due to extra information, a weaker opponent, or uncharged effort does not establish the mechanism's advantage.

DiscoveryWorld is one candidate task for a narrower scientific-discovery mechanism [37]. Its official goals and evaluator would need to remain fixed, with official agents and strong revising controls. Such a score

would assess that task, not the humanistic position or the theatre example. No DiscoveryWorld run is reported here. Actual significance for people and institutions requires the separate forms of evidence below.

9.3 What the local technical evidence establishes

The finite TH-FR-01 candidate identifies 14/70 targets at eight queries, against 70/70 for full-space generalized binary search. It matches a minimum-degree sequential method on all 280 paired trajectories; a later analytic control identifies all 112 catalogue hypotheses with seven fixed queries. The class is saturated and supplies a negative result. Appendix C.2 retains all methods, costs, traces, and qualifications.

The corrected TH-DW-02 retention policy uses actual observed outcomes, but its marginal-exclusion score is equivalent to a structured control's survivor-reduction score under the same byte cost and tie-breaking. Cases involving changed hypotheses or retractions show why a summary adequate for one question can omit evidence needed for another. Six scripted episodes establish execution in a fixture environment with a fake provider only. Appendix C.3 preserves the correction, equivalence, counterexamples, and costs. The earlier exact diagnostics remain in Appendix C.1.

These results rule out claims for the tested implementations. They neither establish a distinct technical advantage nor evaluate the accepted purposes, lived relationships, or intergenerational judgments discussed in this paper. No real-model, official-benchmark, or participant benefit is reported. The normative argument stands on its stated reasons, comparisons, and counterexamples; it is not supported by treating these unsuccessful mechanisms as evidence for its aspirations.

9.4 A separate study of revised purposes

A constructed scenario illustrates the research object. A group begins by planning an impressive digital festival from hometown memories. A participant instead wants an accessible shared craft activity; another wants to listen and remember rather than make an artifact. These choices change the sought experience. A world model can predict some technical consequences, Pangu can provide editable materials, Fuxi can execute shared interactions, and Nuwa can help people learn or delegate. This is a scenario, not an observed deployment.

The study would compare systems with the same technical capabilities and different deliberative interaction mechanisms. Participants can propose, revise, delegate, or decline purposes. Record initial and revised reasons, available alternatives, realizations, effort, and later reflection. Analyze whether suggestions helped reveal a wanted possibility or displaced a purpose through pressure, novelty, fatigue, or framing. Rejection of a proposal is not automatically a system failure, and acceptance is not proof of authentic preference. Include affected collaborators rather than interviewing only the system's primary user.

Attribution requires intervention: if the contribution of one participant is removed or altered in an authorized test, does the shared activity actually change? This tests a narrow causal role, not the value of the relationship. An accessible observational role can be the person's desired outcome. The design must include real people for claims about their experiences; simulated personas can only test implementation. No recruitment, ethics approval, consent, or participant outcomes have occurred. Those requirements and an appropriate sampling plan precede any human study.

10 Theoretical contribution and its relation to existing accounts

The contribution is a connected analysis of what should continue after successful assistance: reusable means, changing purposes, established responsibilities, and cultural forms can require different decisions. Its value is assessed through justified choices and reversal conditions, including the full-agreement case in Section 1.3. The comparison below separates inherited reasons from this design application; it does not infer originality from combining established terms.

Self-determination theory supplies psychological work on autonomy, competence, and relatedness [1]. Meaning-in-life research already examines automation, achievement, reception, and the limits of games as a response [33]. Taihe keeps care and standing independent of a requirement to achieve, create, or continually discover meaning. The resulting position does not establish the empirical effects of any proposed assistance.

Pluralistic alignment, cooperative AI, and empowerment address diverse values and interdependent action [2, 4, 13]. Reward uncertainty is already central to CIRL [21]. It would therefore be inaccurate to describe all existing assistance as optimizing a single known, fixed preference. Taihe adds no general solution to interpersonal trade-offs. Its proposed distinction is to connect revisable purposes and participation to changes in enabling environments, while keeping authority and actual benefit open to scrutiny.

Novelty search, quality-diversity methods, and open-ended learning already investigate discovery beyond direct objective pursuit [25, 16, 3]. POET co-evolves solvers and environments; OMNI and OMNI-EPIC connect interestingness to open-ended tasks and executable environments [15, 26, 27]. New goals, generated rewards, and preserved stepping stones are therefore not sufficient novelty claims. Taihe's particular research question concerns emerging purposes within continuing lives and relationships, their realization, and their consequences. The design question here concerns how an emerging purpose is interpreted, realized, and carried into later activity; generating a novel task does not by itself answer it.

Planning, program synthesis, and human-AI co-creation already contain dependencies, repair, iterative specification, and meaningful participation [22, 23, 24]. Our bounded synthesis protocol adapts them across persistent capabilities, rules, and affordances. Its conditional equivalence result explicitly rules out a blanket superiority claim over serializable planning. There is no new general planning theorem or successful distinct search algorithm here; Appendix B.12 establishes conditional equivalence to a sequential comparator.

Persistent agents and memory, self-modification, learned world models, and shared simulation are enabling neighbors rather than evidence for this paper's humanistic claims [8, 20, 12, 11, 5, 19, 6, 7, 9, 10]. A replay archive does not predict unobserved branches; convincing social simulation does not establish human welfare. Official descriptions of frontier generated-world systems likewise motivate possibilities without giving Liva their performance or proving durable social commitments [17, 18]. Different tasks and resources prevent combining their reported strengths into an achieved SOTA system.

The contribution is a design-oriented synthesis of purpose formation, situated participation, and transmission to later agency. Sections 3.4 and 6.4–6.8 develop cases in which these distinctions change the appropriate division of work, the representation of commitments, and what a learning system should retain. The technical tests in Section 9 evaluate selected means of supporting inquiry; the normative argument rests on the premises, comparisons, and counterexamples developed here.

10.1 What is inherited, and what this synthesis specifies

The argument builds on substantial accounts of human development rather than on an assumed opposition between fixed preferences and newly discovered sociality. The relevant comparison concerns which judgment each predecessor already supports and what remains to specify for an AI system that selects encounters, participates in an activity, and carries experience forward.

Sen's capability argument already shows why equal resources and equal satisfaction can leave a morally relevant deprivation unaddressed; he also retains a role for happiness. Nussbaum explicitly includes forming a conception of the good, reflecting on one's life, emotional development, affiliation, and their institutional conditions [30, 31]. Purpose formation and dissatisfaction with satisfaction-only evaluation are therefore inherited commitments, not Taihe discoveries.

Tronto treats care as a relational practice and urges assessment of the care consequences of policy changes, including effects across connected practices [32]. Counting only assistance received by the focal user misses who sustains it. Taihe applies this insight when a successful personal interaction becomes an organizational routine that allocates others' time, attention, and responsibility.

Mackenzie analyzes revisable, socially coauthored self-understanding, affective and imaginative agency, and the constraints of nonvoluntary group membership. Joint activity can change ambitions, while group narratives can enable or constrain members [38]. The strongest relational account thus already reaches changing purposes and institutional influence. Taihe does not add these by merely inserting an AI into the relationship.

Prunkl distinguishes authenticity from effective agency and examines how AI choice architecture can affect preferences and reflection [46]. This supplies an important diagnosis of the theatre encounter: its influence on motives and its practical opportunities should be assessed separately. Taihe accepts that diagnosis and follows the experience into a second decision. Even when the encounter was valuable and the

participant's interest is authentic, reusing a technique, assigning teaching work, and preserving a cultural default can warrant different answers. The relevant grounds include an existing commitment, the roles of other people, and later means of adaptation. This comparison does not show that Prunkl's framework cannot accommodate those facts. It specifies a continuing design problem to which that framework, care ethics, and capability accounts jointly apply. The proposed increment is the connected decision analysis in Sections 1.1–1.4: in the full-agreement case, competent and welcome automation can coexist with a lost path to later making. The relevant design response concerns executable means and their support, even after the consent issue has been settled. This is not a newly discovered risk of preference formation.

Huang, Hung, and Lin provide an especially close AI precedent: a carebot can improve competent decision-making while shaping motives in objectionable ways, and assistance can produce tensions among opportunity, competence, and self-authorization. They explicitly deny that these tradeoffs are inevitable [39]. Taihe therefore cannot claim that multidimensional autonomy or AI-mediated purpose formation is a new problem. Its extension here follows the reuse of an encounter into later activities, organizational decisions, and resources for successors.

Clark and Chalmers argue for active cognitive coupling, discuss reliable access and endorsement, and explicitly anticipate moral consequences of interfering with an extended person's environment [44]. Suchman's introduction addresses the contingencies that plans cannot fully specify, unequal access to unfolding situations, and responsibility in sociomaterial configurations [45]. Neither can be reduced to a merely descriptive precursor to which Taihe first adds ethics. The present contribution concerns design judgments within an interacting system, without requiring a verdict on the metaphysical boundary of mind.

The contribution offered by this synthesis is a connected answer to what should continue after successful assistance. Capability accounts supply reasons to examine usable means; relational autonomy explains changing purposes and commitments; care directs attention to dependence and the distribution of work; situated accounts limit what an episode's formal description can capture. Together they support the judgment that a useful capability, a purpose, an obligation, and a cultural form need not share a single persistence decision. Sections 1.2 and 6.4 show how that judgment changes assistance and provision while allowing commitments, delegation, alternatives, and costs to reverse the recommendation. Its value must be assessed against these concrete decisions, including the possibility that an already adequate account or implementation does the same work.

Table 5: Design judgments and evidence that could defeat the concern

Transition within the same undertaking	**Decision that the synthesis makes explicit**	**Evidence that could defeat the concern**
An encounter becomes a proposed purpose	Help make an unfamiliar activity feasible; distinguish a tentative attraction from a durable commitment and from an obligation imposed on others	Participants had usable alternatives, understood the relevant consequences, and can continue or redirect the undertaking without the alleged dependency
A purpose becomes a completed task	Distinguish chosen delegation, wanted participation, and burdens no participant wishes to retain	Automatic assistance preserves the desired contribution or the person chooses full delegation; supposed learning benefits do not justify the additional effort
Experience becomes shared knowledge or a default	Preserve useful technique while distinguishing a situated choice, a personal interpretation, and a collectively adopted practice	A simpler representation already preserves those distinctions, or the default remains practically adjustable without transferring another person's authority
A practice becomes an inheritance	Sustain wanted opportunities through usable tools, access, relationships, and justified maintenance, including the option to depart	An alternative service supplies the same relevant opportunities, or preserving the practice imposes disproportionate care, privacy, or resource costs

These are defeasible design judgments, not requirements to collect every life detail or obtain permission for every routine act. Existing delegation can authorize sustained competent action. The distinctions become consequential when the interpretation, affected roles, or use of experience changes. The same framework supports active assistance, chosen rest, and a deliberate end to a practice.

11 Conclusion

What should intelligence carry forward? Intelligence can help people form new purposes and make them consequential in shared practice. Useful abilities can help people create, understand, earn a living, sustain relationships, and receive care. The purposes, assignments, and cultural forms through which those abilities arose need their own grounds for continuation. A successful encounter can open a new direction without prescribing the rest of a life. An established responsibility can warrant continuity without making one person's labor permanent. Later people can inherit usable means without inheriting the same ambitions. These connected judgments support active assistance, chosen delegation, faithful practice, exploration, and rest.

Truth and love make assistance attentive to particular people and answerable to consequences. Lingji and rational inquiry allow emerging concerns to reach examination and consequential action. Cultural difference, shared commitments, and the conditions of a whole life determine what counts as worthwhile help. World models, persistent agents, world production, and recursive improvement should be directed by these questions; their capabilities do not decide the answers. Physical and virtual settings require attention to their actual consequences, material dependencies, and relationships.

The resulting civilizational question concerns what growing intelligence makes possible across lives and generations. Shared resources may sustain different practices, or offer excellent current delivery while narrowing future making. Accessible tools, willing support, alternative services, and revisable institutions can change that judgment. None entitles a community to a contributor's lifelong labor, and no contribution is required for care. Taihe's answer is to carry useful capabilities into future lives while keeping their purposes and uses answerable to the people, responsibilities, and circumstances they will encounter.

A Claims, provenance, and artifact status

The normative argument is supported by explicit premises, constructed paired cases, and selective primary-source comparison. These are not observations of participant benefit. The AI stage vocabulary motivates the research questions and is not an established historical classification. Reproducible technical artifacts are separated below from proposed studies. No private recordings or interactions form a released empirical dataset.

Table 6: Artifact and evidence status

Artifact or evidence	Status	Permitted conclusion
Normative argument and institutional scenarios	Arguments and constructed cases	Claims and objections are available for criticism
Formal definitions and conditional proposition	Retained and recontextualized	The stated result holds under its explicit assumptions
Exact finite program and all outputs	Exhaustive finite systems; no fitted model	Reproducible boundary diagnostics in stipulated finite systems
Closed-class explanatory revision	Implemented and compared in TH-FR-01	Candidate loses to full-space GBS and is equivalent to a sequential comparator
Open-ended discovery and cross-system synthesis	Specification; no complete implementation	No demonstrated benefit or complete integrated system
Official benchmark experiment	Not started	No external score, ranking, or accepted leaderboard result
Human co-creation and longitudinal care study	Not started	No measured lived benefit or lifetime claim
AI-design and recursive improvement evaluation	Proposed only	No demonstrated self-improvement advantage or recursive acceleration

B Bounded technical specification and equivalence

B.1 Coupled state

Let a deployment at time t have state

$$X_t = (s_t, \gamma_t, \mathcal{N}_t, \mathcal{G}_t, \mathcal{H}_t, \mathcal{C}_t). \tag{1}$$

Here s is shared executable state, gamma is the current rule and interface version, N is the collection of persistent agents, G contains world-generation capabilities and assets, H contains permitted histories, and C contains participation agreements and resource constraints. These records do not exhaust anyone's inner life.

Agent i has memory m_i, a skill collection k_i, a policy π_i, and an improvement procedure u_i. A world runtime realizes state transitions under joint actions a. A learned world model estimates transitions for rehearsal:

$$s_{t+1} \sim T_{\gamma_t}(\cdot \mid s_t, a_t), \qquad \hat{T}_{\psi_t} \approx T_{\gamma_t}. \tag{2}$$

T can include stochastic mechanics, asynchronous events, and human actions. Its implemented state is authoritative only for its digital domain. Physical observations may be partial or delayed. A prediction from the learned model is never, merely by being predicted, an observed event.

B.2 Executable opportunities

An opportunity is a structured object

$$o = (q, \mu, \mathcal{P}, \mathrm{pre}, \pi, b, h, e). \tag{3}$$

q describes an outcome or temporally extended experience. The participation mode mu specifies whether someone acts personally, with assistance, through delegation, jointly, or as an observer. P identifies relevant participants, pre gives prerequisites, pi is a candidate plan, b describes resource demands, h is a horizon, and e records evidence and provenance.

An opportunity can concern a performance, a visit, a shared work, a lesson, a quiet gathering, or a new use of an existing world. The representation is deliberately more specific than a text suggestion and less ambitious than a complete representation of meaning.
Let Q be a finite, declared evaluation catalogue of outcome-and-mode descriptors. Define the feasible opportunity set for participant i as

$$\begin{aligned}\mathcal{F}_i(X; B, h, \epsilon) = \{o \in \mathcal{O}(Q) : &\exists \pi \in \Pi_i^{\text{allowed}}(X, o), \\ &C(\pi) \preceq B,\ \Pr_T[q_o(\tau_{0:h}) = 1] \geq 1 - \epsilon\}.\end{aligned} \tag{4}$$

B is a vector of budgets, including human time, money or compute, and physical resources where relevant. Each evaluation must declare whether C bounds realized costs or estimates expected costs. These interpretations are not interchangeable. Allowed policies incorporate participation mode, agreed collaborator behavior, and applicable rules. They must not assume that another person will cooperate simply because an agent can imagine it. The probability condition is a definition, not a generally computable quantity.

Practical evidence should be separated into model-predicted, executable-rehearsal, and observed-use results. A successful rehearsal establishes a narrower claim than reliable human use. Confidence bounds require an appropriate statistical design. A single successful trace does not justify a calibrated success probability.

B.3 Why opportunity counts are insufficient

Duplicating cosmetic variants can inflate a catalogue. A feasible option may be unwanted, inaccessible to its intended participant, or only available by exhausting someone else's resources. Outcome equivalence, relevance, cost, and participation mode therefore need declared definitions.
For a fixed evaluation catalogue and matched conditions, technical diagnostics can report newly feasible and lost opportunities:

$$\begin{aligned}\Delta_i^+ &= \mathcal{F}_i(X'; B, h, \epsilon) \setminus \mathcal{F}_i(X; B, h, \epsilon), \\ \Delta_i^- &= \mathcal{F}_i(X; B, h, \epsilon) \setminus \mathcal{F}_i(X'; B, h, \epsilon).\end{aligned} \tag{5}$$

These sets are not a utility function. Set comparisons identify an opportunity by its declared outcome-and-mode equivalence class, not by a distinct plan identifier. Novel descriptors generated during exploration belong in a separately reported discovery set until mapped to a declared evaluation scheme. This prevents an algorithm from winning simply by redefining the catalogue it is scored on. People may reasonably choose a smaller set of options to make a commitment or reduce overload.

The framework also distinguishes access from realized benefit. Whether an opportunity supports love, beauty, growth, or belonging requires human interpretation and longitudinal evidence beyond reachability.

B.4 A realization witness must preserve the selected mode

For a bounded executable task, define its success predicate as the conjunction of an outcome predicate and a participation predicate over the same trace. The predicate q_o in the feasible-set definition denotes this combined predicate. A delivered composition can satisfy the outcome predicate even if its intended human co-author made no contribution. A joint-composition predicate can additionally require an identifiable contribution that causally affected the result. If the person selected delegation, requiring that contribution would instead misrepresent the purpose.

A witness contains the task and mode descriptor, relevant agent and world versions, initial conditions, action and event references, resource use, and verification scope. It can certify a particular finite trace under explicit semantics. It cannot certify subjective enjoyment, authenticity of a relationship, or success on unobserved inputs. For stochastic environments, one trace is a sample; a reliability claim needs separate evidence.

Witnesses expire when a relevant assumption changes. A changed instrument interface can invalidate a learned playing skill; a withdrawn collaboration changes the plan even if its code still runs. Unchanged technical dependencies permit reuse only when their declared footprint is complete. Unknown dependence triggers revalidation. Historical evidence remains historical rather than being retroactively relabeled as a current guarantee.

This separates three questions: was the intended activity technically possible, did the selected form of participation occur, and did participants find it worthwhile? They require different evidence. The mode descriptor is revisable through actual choice; it is not a system-assigned identity or a requirement to perform more labor.

B.5 The problem frame can develop too

The fixed evaluation catalogue Q supports comparison, but a life is not confined to a predeclared task list. Define a current problem frame for this local mechanism as `K = (Q_frame, phi, A, V)`: currently expressible purposes, a representation of relevant phenomena, available action descriptions, and participants' articulated evaluative commitments. Q_{frame} can develop; the evaluation catalogue Q remains fixed within a comparison. V may contain disagreement, incompleteness, or nonnumeric reasons; it is not a universal reward function. A frame proposal records typed changes to these four components and why they are proposed.

Frame proposals have a different status from realization patches. Discovering a use for an object can add a purpose without changing a value. Discovering a new causal variable revises a representation. Providing a new interaction expands an action language. Changing from competition to shared improvisation can revise a chosen purpose and its criteria. The system must state which change is proposed rather than silently replacing an old evaluator with an easier one.

A realization patch z is evaluated under a declared frame K. A proposed frame change can in turn call for new patches to Nuwa, Fuxi, or Pangu. The two processes interact: realized experience supplies observations and reasons for a subsequent frame proposal, and a revised frame directs subsequent technical improvement. This is the proposed reciprocal mechanism. It accommodates the discovery of ends as well as more capable means.

Within a fixed task, compare methods under the same evaluator. Across frames, report the old and new descriptors, consequences, reasons, and participant choices separately. There is no automatic scalar comparison between a better solution to an old purpose and a newly chosen purpose. The system cannot establish the value of a new direction solely by assigning it a new reward.

Rationality boundary. For a fixed utility, information set, budget, and admissible policy class, an optimal policy cannot be outperformed in expected utility by another policy in that same class; this follows from optimality. Any observed gain from a Lingji-inspired intervention must instead be attributable to a changed information set, representation, action space, practical search procedure, or evaluative commitment. This boundary preserves a research role for emotion and intuition without declaring logical error a source of guaranteed superiority.

B.6 How AI 3.0 can propose a new question

A new question needs a referent and a possible consequence. We propose a frame-change record containing the source encounter, the prior frame, the changed component, a contrast with the old question, a discriminating test or experience, affected participants, required realization patches, and a later disposition. The source can be a person, an AI, or an interaction. An emotional reason can motivate a proposal without being a factual proof that its predicted consequence will occur.

For example, a group initially seeks the highest-scoring synchronized performance. A participant values an unplanned response between two unlike instruments. A new question is whether an asynchronous duet can let each player shape the other's next phrase. That proposal changes the purpose and required interactions. Pangu supplies editable instrument mappings, Fuxi executes delayed reciprocal effects, and Nuwa learns how to support turn-taking. The candidate is tested by intervening on one player's contribution and checking whether the other's available response actually changes. Participants then decide whether this is an experience they want to develop. Technical reciprocity and artistic worth are separate findings.

Operationally, the proposer detects or receives an unresolved encounter, produces a bounded number of typed revisions, and selects tests that distinguish a new mechanism from a renamed old one. A realization planner builds only the changes needed for those tests. Failed tests can revise the explanatory model, while rejected purposes remain rejected without being relabeled as technical failures. This adds an explicit interface between question formation and executable investigation. No claim is made that question formation or this general sequence is unprecedented.

The method can fail through vacuous novelty, untestable proposals, excessive cost, or desirable language without a changed interaction. It can also discover a technically valid direction that nobody wants. Reporting these separately preserves the positive role of imagination while making a particular discovery claim vulnerable to evidence.

B.7 The joint transformation object

A new experience can require a new skill, a change in a shared institution, and a new world affordance simultaneously. We represent a candidate transformation as

$$z = (\Delta\mathcal{N}, \Delta\gamma, \Delta\mathcal{G}, \pi, e_z). \tag{6}$$

Its components propose changes to agents, runtime rules, and generated content, together with a realization plan and supporting evidence. Components can be null. The proposal does not require a separate model for each component or mandatory simultaneous changes.

The system searches a budgeted archive of such transformations, with human-articulated purposes and exploratory proposals both available. Participants may choose among relevant candidates, delegate this choice within a specified scope, revise a purpose, or decline. Technical feasibility is evaluated separately from preference and from effects on other participants. Comparisons must charge proposal search, training, world modification, and subsequent use to the total budget; the use-cost constraint in Appendix B does not make development costs free. Shared reusable work requires a declared amortization rule.

B.8 An elementary complementarity example

Suppose a group wants a playable collaborative instrument. Learning an agent skill to follow a shared beat has no immediate utility if the world contains no usable instrument. Producing the instrument asset does not make the activity feasible if the runtime provides no synchronized interaction. All three changes may be required for one usable experience.
Formally, let binary variables n, f, and p denote these three prerequisites and let the diagnostic success function be

$$J(n, f, p) = nfp. \tag{7}$$

At (0,0,0), every single-coordinate change has zero gain, while the joint change reaches success. This is an elementary constructed example of complementarity, not an empirical result or a novel theorem. A sequential method that retains stepping stones or plans multiple changes can also solve it. Therefore the appropriate comparison includes strong sequential search and open-ended archives, rather than an artificially myopic baseline.

The research hypothesis is that an explicit joint representation improves search efficiency or realization quality in domains with important cross-component dependencies. It may fail where simple decomposition is sufficient or coordination overhead dominates.

B.9 Dependency-directed proposal construction

A concrete implementation can encode a proposed experience as a typed dependency graph. Outcome and participation-mode nodes connect to required agent skills, world entities, interaction rules, collaborators, and resources. Existing capability witnesses mark satisfied prerequisites. Missing or incompatible prerequisites identify a bounded joint patch to explore.

Candidate generation then has explicit operators: reuse a verified skill, adapt a skill, compose or edit an affordance, introduce a compatible rule, propose a collaborator role, or change a participation mode with the relevant participant's agreement. Each candidate records its dependencies and estimated cost before

rehearsal. Aesthetic or narrative edits can accompany the patch, but cannot replace a missing executable interaction.

Rehearsal tests both the proposed opportunity and declared retained capabilities. Failures are localized to a dependency where possible, allowing another candidate to reuse unaffected work. The archive can preserve distinct combinations of participation mode, cultural interpretation, required effort, and technical feasibility. These descriptors must be disclosed and checked for omissions. A Pareto set over declared technical criteria can organize candidates, but does not settle trade-offs between people's lives.

This specification supports comparison with hierarchical planning, program synthesis, and tool-using agents. A planner that already discovers and maintains the same dependencies may match or outperform it. Dependency-directed proposal construction is a proposed implementation choice, not a claim to have invented dependency graphs or hierarchical planning.

B.10 When joint and sequential changes are equivalent

Proposition 1 (conditional reachability equivalence). Fix an initial state, a primitive edit language, and a budget. Assume that (i) every permitted joint transformation has a permitted primitive serialization with the same terminal state, evaluated outcome, and cost; (ii) every primitive edit is available as a singleton joint transformation; and (iii) budgets and horizons count primitive costs and relevant elapsed time consistently, rather than charging a whole joint transformation as one primitive step. Then finite sequences of joint transformations and finite sequences of primitive edits reach the same evaluated opportunities within that budget.

Proof. Replace each joint transformation in a feasible sequence by its permitted serialization. Assumption (i) preserves the evaluated outcome and assumption (iii) preserves feasibility under the budget and horizon. This gives inclusion in primitive-sequence reachability. Conversely, replace each primitive edit by its singleton joint transformation using (ii); costs and evaluated outcomes remain unchanged. This gives the reverse inclusion.

The proposition is deliberately elementary. It shows why complementarity alone cannot establish superiority over planning. Assumption (i) is substantive: intermediate exposure, deadlines, stochastic coupling, or participation during the transition can break it. A staging area may restore serialization by hiding incomplete versions until publication. Any claimed gain must identify which assumption fails, or demonstrate improved search under equal information and resources. Giving one method atomic edits and denying the other staging would confound representation with the available action language.

B.11 Participation-aware synthesis across changing components

The following concrete protocol connects the formal object to an implementable method. Given an articulated purpose, an agreed mode, current versions, and budgets, the planner constructs dependencies for both outcome and participation. It proposes an agent patch, an affordance patch, a rule patch, and an executable plan, allowing null patches. A verifier attempts to realize the plan in the relevant executable environment. A failure yields a typed counterexample: missing capability, incompatible interface, unmet collaborator condition, excessive cost, or an outcome that lacks the selected participation. Counterexamples refine the next candidate without silently changing the purpose.

The archive stores candidates with their assumptions and separate evidence levels. A patch to a dependency invalidates the associated witness and causes repair or renewed evaluation. The selected candidate is staged, executed within its agreed scope, and compared with actual events. A later participant revision starts a new purpose or mode version; the system may reuse technical work but cannot reuse old endorsement as evidence of the new choice.

This protocol adapts planning and counterexample-guided synthesis [22, 23]. Its proposed point of integration is preserving the selected participation semantics across changing capabilities, world rules, and produced content. It remains an algorithm specification, not a demonstrated new search algorithm. The finite implementation reported in Appendix C.1 tests necessary distinctions and the reachability boundary; it does not implement open-ended proposal generation, learned verification, or human deliberation.

B.12 An executable revision candidate and its identifiability boundary

We now implement a deliberately narrow explanatory revision. Inputs have six binary coordinates. A hypothesis predicts the parity of a nonempty subset of at most four coordinates, optionally complemented. The known catalogue contains 112 hypotheses and the query space contains 64 inputs. There are no hidden additional variables, noisy observations, learned embeddings, language models, or changing purposes. Expanding a permitted interaction degree changes the currently considered explanatory class; it does not invent a concept outside the catalogue.

The candidate starts with degree-one explanations. After each observed answer, it retains hypotheses consistent with the complete observation history. If none remain, it increases the allowed degree until a consistent class is found. It chooses the unqueried input that most evenly partitions the current class, with the smallest input index breaking ties. After a fixed query budget it predicts each input by majority vote, breaking ties toward zero. Contradictions and expansions are retained in the trace. This is an executable specification of one proposed local mechanism, not an implementation of all the preceding architecture.

The full-space baseline uses the same partition rule on all surviving hypotheses from the outset. This is a noiseless finite implementation of generalized binary search (GBS), whose established query-complexity results require their own conditions [41]. We do not claim those conditions or invent a new GBS result. An independent minimum-degree baseline maintains the full consistent catalogue, then selects its smallest surviving degree for query and prediction. This is a complexity-prioritized sequential comparator, not a reproduction of a complete published MDL system. Two ablations remove expansion or active query selection. All receive the same available coordinates, catalogue, intervention oracle, and per-case query budget; the ablations deliberately refrain from using part of that capability.

Conditional equivalence. With nested degree classes, noiseless consistent filtering, no forgetting, and identical tie-breaking, contradiction-triggered expansion and minimum-degree sequential selection have the same query-answer and prediction trajectories. To see this, once every smaller degree is inconsistent, adding observations cannot make one consistent again. The candidate's first nonempty class is therefore exactly the minimum-degree survivors of the full catalogue. Equal survivor sets imply equal next queries, and induction gives equal histories and final predictions. Their computational work need not be equal. This result removes a putative algorithmic novelty claim; Appendix C.2 checks the implementation against it.

Observational non-identifiability. If two deterministic explanations give the same answer to every admissible query, no adaptive policy using only those answers can distinguish them. Couple the policy's random choices: equal initial histories give the same next query, its answers agree, and induction preserves identical histories. Extra computation alone cannot break this equivalence. A distinguishing claim requires an added observation, intervention, or external assumption, and its cost must be disclosed. This elementary boundary is not new statistical theory. It also explains why a new description without new observable consequences cannot count as discovery in this study.

C Full local results and counterexamples

C.1 Existing exact finite diagnostics and their negative result

The standard-library Python program retained from v0.2 tests three restricted boundaries. It uses no fitted model, human data, or stochastic sampling. Its counts are exhaustive properties of stipulated transition systems. These old computations are retained without alteration. The separate new comparison in Appendix C.2 does not turn them into evidence of open-ended discovery.

In a six-bit system with two bits each for agent, runtime, and content, weighted primitive flips are compared with an endpoint oracle charging the same edit cost. All 64 initial states at budgets 0 through 9 yield identical reachable sets: 640 comparisons, no joint reachability advantage. This supports the conditional equivalence proposition. It does not compare search efficiency and the oracle is not a learned planner.

In a three-bit example allowing only 000 and 111 as publicly visible versions, online single-bit publication cannot complete the change. An atomic update and sequential edits in a private staging area both reach 111 at primitive cost 3. Giving staging only to one method would create a misleading advantage. The result illustrates an established transaction capability.

In a length-three trace model for human motif creation and AI composition, 36 of 64 candidate traces satisfy preconditions; 28 deliver a composition, and 20 of those delivered traces omit the stipulated human contribution in the final composition. These are not behavioral frequencies. They show that outcome and selected participation can be different predicates. Solo composition is appropriate under delegation, so the counts do not establish moral superiority of joint creation, authentic authorship, or human benefit.

C.2 TH-FR-01: actual comparison and a negative result

We executed the Appendix B.12 methods in an author-constructed finite environment. Development used all 42 degree-one/two functions at eight queries. Final evaluation used the disjoint 70 degree-three/four functions at budgets 2, 4, 8, and 12, producing 280 cases and 1,400 method executions. All functions were already in the declared catalogue: this is held-out target evaluation within a known class, not transfer to unseen concepts. The developer knew the generator. The learner received only selected query answers, not target identities or answer tables.

The code, development record, cases, metric, costs, primary eight-query comparison, and analysis were hashed and locally sealed before final execution. That sequence is recorded; a local timestamp is not independent preregistration. Development exactness covered unqueried inputs. Before freezing final evaluation we tightened the primary score to require correctness on all 64 inputs, so querying a hard point cannot remove it from the scoring denominator. No method was changed after final outcomes. Every planned case and trace is retained; there was no run selection or retry.

Table 7: Exact identification over 70 held-out targets at matched query budgets

Method	**Exact at 2 queries**	**Exact at 4**	**Exact at 8 (primary)**	**Exact at 12**
Contradiction-triggered revision	0/70	2/70	14/70	16/70
Full-space GBS	0/70	0/70	70/70	70/70
Minimum-degree sequential	0/70	2/70	14/70	16/70
No expansion	0/70	0/70	0/70	0/70
Passive revision	0/70	0/70	2/70	10/70

At the primary budget, the candidate identifies 20% of functions versus 100% for full-space GBS: an 80-percentage-point deficit. It has exactly the same query-answer trajectories and final predictions as minimum-degree sequential selection on all 280 paired cases. Removing expansion or active querying damages performance in this construction, but those ablations do not establish advantage over a capable baseline. The two additional successes at four queries do not rescue the prespecified eight-query claim. All percentages describe this complete finite catalogue; no population confidence interval or significance test is appropriate.

The candidate makes a mean 2,304.7 hypothesis evaluations per eight-query case, compared with 14,362.1 for full-space GBS and 1,972.3 for minimum-degree sequential selection. Thus fewer evaluations than full-space search accompany a substantial loss of accuracy; the independent sequential method matches the candidate with fewer evaluations. The query budgets are matched, not the operation counts. These counts are instrumented function evaluations, not FLOPs or cross-hardware latency claims. Full execution, including trace serialization, took 4.314 seconds on the recorded arm64 Mac/Python 3.12.14 environment; the earlier development run took 0.158 seconds. Per-method learner/oracle timing and all raw traces are supplied. Runs made no external or paid model calls and consumed no inference tokens. Research, writing, and human labor were not metered, so these figures do not represent the total research cost.

An independent signed-product oracle recomputed every recorded answer and score. This checks the finite implementation; it does not validate the environment as a model of human discovery. The adverse result rejects superiority for this implementation and removes its claim to a distinct revision algorithm. A richer proposal would need actual representation invention, discriminating interventions, general reflection baselines, and new held-out evaluation, not a renamed version of this experiment. We have not made such

a proposal succeed. DiscoveryWorld remains an unexecuted external comparison; there is no official score or leaderboard result.

A subsequent independent design review supplied a stronger analytic control that was not in the frozen comparison. Query the all-zero input to recover the sign, then each of the six unit vectors to recover every coefficient by its difference from the zero response. These seven fixed queries identify any affine parity function, including all 112 catalogue functions. This is a proof for the noiseless class, not a new fitted method or a preregistered result. It makes the eight-query accuracy track analytically saturated. We disclose this omitted control as a design weakness and retain the earlier results as a boundary and falsification study. Enlarging this Boolean experiment cannot establish a Lingji advantage.

C.3 Observation-conditioned retention: correction and a deeper boundary

The original TH-DW-02 selector scored an observed record by the number of hypothesis pairs predicting different outcomes. This can balance a prospective query, but ignores which outcome actually occurred. In an equal-byte, one-record example, it retained a record leaving two consistent hypotheses although the alternative retained only one. The offline-v0.3 revision instead binds a typed outcome to a location in the original public observation and preserves witnesses to recorded inconsistency. Source binding prevents substitution of an invented outcome; it cannot establish that the observation, opportunity, or model interpretation is correct. All explanatory predictions remain retrospective and fallible.

Let H be a fixed finite hypothesis class and E the available records. For each admissible record e, define C(e) as the hypotheses with a known prediction different from the recorded outcome. Unknown predictions and inadequate observation opportunities exclude nothing. For retained records S contained in E, let V(S) be H excluding the union of C(e) over S. Then V(E) is contained in V(S). Equality holds exactly when retained records cover all the exclusions made by E. The proof follows directly from these set complements. An empty V(E) flags recorded inconsistency of the class, not a discovery that no explanation is possible. This elementary version-space property is not a new theorem of learning.

The implemented greedy selector chooses the largest newly covered exclusion set per marginal serialized byte, then fills residual capacity by recency. The full retained record, original observation, outcome binding, predictions, and memory wrapper count toward the same capacity. A strong structured control independently recomputes the reduction in surviving hypotheses for each candidate record. Under identical tie-breaking and costs the choices are equivalent: their marginal objective is the cardinality of precisely the same newly excluded set. Precomputing exclusion sets can change bookkeeping cost; it supplies no new scientific information or distinct selection principle. We therefore withdraw an independent algorithmic-contribution claim for this revision.

Six constructed cases expose the correction and its limits. In the original failure, both corrected greedy methods retain the uniquely discriminating record. In a capacity-two case with elimination sets {1,2,3,4}, {1,2,5}, and {3,4,6}, greedy retains only five of six possible exclusions; exhaustive subset selection retains all six. Further cases show that absent observation opportunity or unknown predictions must not eliminate a hypothesis, contradictory measurements can empty the class, a retracted witness can expose information lost in compression, and a newly introduced hypothesis can require a previously non-discriminating record that was discarded. Exhaustive selection uses the same information and byte cap but examines more subsets. Full-archive comparisons in the last two cases deliberately have extra information and are diagnostic ceilings, not equal-information wins.

The last failure matters for the larger research aim. For an arbitrary history family X with more than 2^B elements, any deterministic B-bit memory maps two distinct histories to the same state. If the admissible future query family can distinguish every pair of histories, at least one such query has different answers for a colliding pair. A decoder receiving only that memory and query therefore cannot answer correctly for every history and future query. This is a counting argument about exact guarantees, not a lower bound on ordinary task accuracy or an impossibility of useful compression. A restricted query family can admit a sufficient summary; external retrieval, new measurements, or a larger memory change the premises. No novelty is claimed for this information-loss principle.

The implication is that a selector optimized for a current explanation class cannot by itself implement support for unforeseen questions. An additional retrieval route, protected diverse source material, or an

ability to reacquire observations may help, with real storage, privacy, latency, and exploration costs. Each requires comparison to a strong ordinary retrieval and reflection system using the same information and total resources. None has yet been implemented or shown superior here. Replacing pair coverage with witness coverage repairs a bug; it does not resolve this larger design problem.

The corrected policy completed the same six scripted interface runs, with source-bound observations, legal action dispatch, separate post-run evaluation, and budget accounting. All three policies still used a fake provider and a fixture environment. The accompanying records contain selector operations, serialized bytes, local wall time, requests, and failures; external calls and actual paid cost were zero. These artifacts establish local execution and analytic counterexamples only. DiscoveryWorld assets, real model inference, official benchmark scores, and human studies remain absent. Paid expansion is not justified by these results.

EvoBelief already connects falsifiable hypotheses, evidence, scoped rules, and counterexamples, and reports a DiscoveryWorld implementation with code-fitted laws and compiled experiments [42]. Its correction of a response ceiling that truncated Hypothesizer reinforces the need for capable controls with equal resources. The present implementation consequently supplies those controls with hypothesis revision and reflection, and reports the outcome-conditioned selector as a correction rather than a new discovery algorithm.

Tools and assistance

Generative AI tools were used extensively for literature synthesis, drafting, code development, figure preparation, and internal checking. This assistance is distinct from the computational diagnostics explicitly reported in the paper. Internal AI-assisted checks are not external peer review.

References

[1] Ryan, Richard M.; Deci, Edward L.. *Self-Determination Theory and the Facilitation of Intrinsic Motivation, Social Development, and Well-Being*. 2000. Primary source.

[2] Sorensen, Taylor, et al.. *Position: A Roadmap to Pluralistic Alignment*. 2024. Primary source.

[3] Hughes, Edward, et al.. *Position: Open-Endedness is Essential for Artificial Superhuman Intelligence*. 2024. Primary source.

[4] Dafoe, Allan, et al.. *Open Problems in Cooperative AI*. 2020. Primary source.

[5] Zhang, Jenny; Hu, Shengran; Lu, Cong; Lange, Robert; Clune, Jeff. *Darwin Godel Machine: Open-Ended Evolution of Self-Improving Agents*. 2025. Version 3, 12 March 2026 Primary source.

[6] Zheng, Tong, et al.. *Dream-RSI: Recursive Self-Improvement through Evolving Worlds*. 2026. Primary source.

[7] Hafner, Danijar; Yan, Wilson; Lillicrap, Timothy. *Training Agents Inside of Scalable World Models*. 2025. Primary source.

[8] Park, Joon Sung, et al.. *Generative Agents: Interactive Simulacra of Human Behavior*. 2023. Primary source.

[9] Vezhnevets, Alexander Sasha, et al.. *Generative agent-based modeling with actions grounded in physical, social, or digital space using Concordia*. 2023. Primary source.

[10] Piao, Jinghua, et al.. *AgentSociety 2: An Integrated Research Environment for Executable Social Science*. 2026. Primary source.

[11] Behrouz, Ali; Zhong, Peilin; Mirrokni, Vahab. *Titans: Learning to Memorize at Test Time*. 2025. Primary source.

[12] Chhikara, Prateek, et al.. *Mem0: Building Production-Ready AI Agents with Scalable Long-Term Memory*. 2025. Primary source.

[13] Myers, Vivek; Ellis, Evan; Levine, Sergey; Eysenbach, Benjamin; Dragan, Anca. *Learning to Assist Humans without Inferring Rewards*. 2024. Version 3, 16 January 2025 Primary source.

[14] Yang, Claire; Zhang, Claire Jie; Cakmak, Maya; Kleiman-Weiner, Max. *When Assisting One Disempowers Another*. 2026. Version 2, 6 July 2026; title changed from version 1 Primary source.

[15] Wang, Rui; Lehman, Joel; Clune, Jeff; Stanley, Kenneth O.. *Paired Open-Ended Trailblazer (POET): Endlessly Generating Increasingly Complex and Diverse Learning Environments and Their Solutions*. 2019. Primary source.

[16] Mouret, Jean-Baptiste; Clune, Jeff. *Illuminating search spaces by mapping elites*. 2015. Primary source.

[17] Google DeepMind. *Genie 4 is now available*. 2026. Primary source.

[18] World Labs. *Atlas: A World Model for Spatial Intelligence*. 2026. Primary source.

[19] Robeyns, Maxime; Szummer, Martin; Aitchison, Laurence. *A Self-Improving Coding Agent*. 2025. Primary source.

[20] Wang, Guanzhi, et al.. *Voyager: An Open-Ended Embodied Agent with Large Language Models*. 2023. Primary source.

[21] Hadfield-Menell, Dylan; Dragan, Anca; Abbeel, Pieter; Russell, Stuart. *Cooperative Inverse Reinforcement Learning*. 2016. Primary source.

[22] Penberthy, J. Scott; Weld, Daniel S.. *UCPOP: A Sound, Complete, Partial Order Planner for ADL*. 1992. Primary source.

[23] Solar-Lezama, Armando. *Program sketching*. 2013. Journal publication 2013; repository issue date August 2012 Primary source.

[24] Rezwana, Jeba; Maher, Mary Lou. *Designing Creative AI Partners with COFI: A Framework for Modeling Interaction in Human-AI Co-Creative Systems*. 2022. Primary source.

[25] Lehman, Joel; Stanley, Kenneth O.. *Novelty Search and the Problem with Objectives*. 2011. Primary source.

[26] Zhang, Jenny; Lehman, Joel; Stanley, Kenneth; Clune, Jeff. *OMNI: Open-endedness via Models of human Notions of Interestingness*. 2024. ICLR 2024; arXiv version 3, 15 February 2024 Primary source.

[27] Faldor, Maxence; Zhang, Jenny; Cully, Antoine; Clune, Jeff. *OMNI-EPIC: Open-endedness via Models of human Notions of Interestingness with Environments Programmed in Code*. 2024. Version 3, 14 February 2025 Primary source.

[28] Schölkopf, Bernhard; Locatello, Francesco; Bauer, Stefan; Ke, Nan Rosemary; Kalchbrenner, Nal; Goyal, Anirudh; Bengio, Yoshua. *Towards Causal Representation Learning*. 2021. Primary source.

[29] Andrychowicz, Marcin; Denil, Misha; Gomez, Sergio; Hoffman, Matthew W.; Pfau, David; Schaul, Tom; Shillingford, Brendan; de Freitas, Nando. *Learning to learn by gradient descent by gradient descent*. 2016. Primary source.

[30] Sen, Amartya. *Equality of What?*. 1979. Tanner Lecture delivered 22 May 1979; selected printed pp.217–220 revisited 21 November 2026. Primary source.

[31] Nussbaum, Martha C.. *Human Rights and Human Capabilities*. 2007. Harvard Human Rights Journal 20, pp.21–24; four-page article read 21 November 2026. Primary source.

[32] Parra Jounou, Iris; Tronto, Joan C.. *Care ethics in theory and practice: Joan C. Tronto in conversation with Iris Parra Jounou*. 2024. Primary source.

[33] Knell, Sebastian; Ruther, Markus. *Artificial intelligence, superefficiency and the end of work: a humanistic perspective on meaning in life*. 2024. AI and Ethics 4, 363-373; published online 17 April 2023. Author's surname Ruther is rendered in ASCII. Primary source.

[34] Laozi (attributed); Heshang Gong recension. *Laozi, chapter 2, with Heshang Gong commentary (Siku Quanshu text)*. n.d.. Public-domain text transcription; base text distinguished from commentary; accessed 20 November 2026. Primary source.

[35] Kumarajiva (Chinese translator). *Vajracchedika Prajnaparamita Sutra (Diamond Sutra), Taisho T235*. 2001. Consulted CBETA digital edition V1.22, 2 August 2001; T08, no. 235, 749c21-24. Year refers to digital edition. Primary source.

[36] Wang Yangming; Xu Ai records. *Chuanxi lu (Instructions for Practical Living), upper scroll, Xu Ai records 3-5*. n.d.. Public-domain Chinese text transcription; sections 3-5; accessed 20 November 2026. Primary source.

[37] Allen Institute for AI. *DiscoveryWorld: project and official agent evaluation documentation*. 2024. Official documentation consulted 20 November 2026; a commit and runnable evaluation release have not yet been frozen. Primary source.

[38] Mackenzie, Catriona. *Autonomous agency, we-agency, and social oppression*. 2023. Primary source.

[39] Huang, Linus Ta-Lun; Hung, Tzu-Wei; Lin, Ying-Tung. *The dynamics of personal autonomy: a comprehensive framework for evaluating human-carebot interactions*. 2025. Primary source.

[40] Yaman, Anil; Tian, Shen; Lindstrom, Bjorn. *Semantic knowledge guides innovation and drives cultural evolution*. 2026. arXiv v4; published in PNAS, doi:10.1073/pnas.2530750123. Primary source.

[41] Nowak, Robert D.. *The Geometry of Generalized Binary Search*. 2011. IEEE Transactions on Information Theory 57(12); consulted arXiv v5 (2013). Primary source.

[42] Wang, Janet; Guo, Jiang; Xiang, Yi; Yuan, Jiaqing; Dong, Mingwen; Kobayashi, Hideo; Chang, Shuaichen; Reddy, Chandan K.; Jiang, Jiarong. *EvoBelief: Falsifiable World Models for Self-Improving LLM Agents*. 2026. Primary source.

[43] Paul, L. A.. *What You Can't Expect When You're Expecting*. 2015. Res Philosophica 92(2), 149–170. DOI:10.11612/resphil.2015.92.2.1. Primary source.

[44] Clark, Andy; Chalmers, David J.. *The Extended Mind*. 1998. Primary source.

[45] Suchman, Lucy. *Human-Machine Reconfigurations: Plans and Situated Actions*. 2007. Primary source.

[46] Prunkl, Carina. *Human Autonomy at Risk? An Analysis of the Challenges from AI*. 2024. Primary source.