

Working Understanding

Intent, mutual adaptation, and the system that builds the system.

Christian Blank | Working edition
Working draft · v0.1 · 2026-10-13

This is my working investigation into a recurring problem: people can have a strong sense of what they want, what good looks like, or what will fail, while finding it expensive to make that understanding actionable for someone else. The same difficulty appears when working with AI. A request can be clear as language and still leave the intended scale of change, quality bar, assumptions, and freedom to act unresolved.

The proposed direction is an environment in which useful understanding accumulates through work. People and AI should become better at interpreting, questioning, contributing, and acting together, while the people using the system develop capabilities they can take beyond it.

This paper extends the **System of Work** through two companion investigations. **Building Through Emergence** establishes the continuous practice of observing, building, operating, and learning. **Adaptive Collaboration** explores shared reality, attention, and multiple projections in community coordination. This paper examines how participants develop enough shared meaning and mutual calibration to use that infrastructure well. It does not replace the thesis, the earlier papers, or dated field evidence.[1, §§1–2, 23; 2, §§1–2]

Working premise

A substantial part of collaboration overhead may come from repeatedly reconstructing understanding that the working environment fails to preserve. The opportunity is to make that understanding cumulative, inspectable, and revisable—without requiring identical minds, exhaustive documentation, or surrender of human agency.

Status	Seed draft · Living document · Proposals remain open to correction
Scope	Problem model, product hypothesis, development practice, and research direction
Starting context	My own building and collaboration work; practitioner reports in the discussion that produced this paper; the two companion working papers
Evidence boundary	No new user study, measured product result, or validated architecture is reported here
Operating rule	Every increment should create useful value now and make the next increment easier to understand or build

1. Purpose and North Star

This paper preserves the reasoning behind a platform that helps people and collectives organize, act, learn, and adapt. An intent-development tool is the first candidate application. The larger ambition is a system whose use improves both the work and its enabling conditions.

The human objective includes meaningful contribution, agency, belonging, and a better experience of everyday life. Recovered attention can go into craft, judgment, relationships, play, or rest. Greater output is one possible benefit, not the definition of a person's success.

North Star

Enable different people and AI to turn evolving intentions into useful shared outcomes, with less avoidable translation, coordination, and repair—and with growing human capability and the capacity to change.

A specification editor, prompt library, or agent platform is a possible intervention. The persistent problem is how different participants can contribute effectively and remain coherent as their goals, abilities, relationships, and environment evolve.

The contribution is to connect participation costs, working understanding, its preservation across artifacts and actors, and building the enabling system through its own use. This is a working synthesis, not a single explanation for every organizational problem.

2. Starting Observations and the Boundary of the Claims

The discussion began with independently useful tools that accelerate further work, produce public learning, and grow into a coherent platform. Two problem framings followed.

The first was an **attention and coordination tax**: decisions, context switching, synchronization, repeated explanation, artifact maintenance, and learning loops all consume human capacity. The second asked why so much of that work is necessary. One proposed explanation was that collaborators lack a sufficiently accurate, shared understanding of each other and of the work.

The starting evidence is situated. I described my own experience and conversations with capable, mission-driven people who notice dependencies and possible improvements but still find communication, decomposition, or repeated coordination expensive. These are practitioner reports, not a representative sample or a clinical finding. The artist, product thinker, operator, and systems thinker examples below are illustrations used to inspect the model; they are not fixed personality categories or measured populations.

The companion *Adaptive Collaboration* paper supplies a more concrete adjacent case: events and decisions originate in conversation, while their durable representations live elsewhere, leaving participants to maintain the connection. Its shared-substrate proposal remains a hypothesis.[2, §§3–6]

The claims have different statuses: human agency is a design commitment; the problem stack is an interpretive model; accumulating working understanding is a product hypothesis; evaluation is future work. These distinctions follow *Building Through Emergence*. [1, §§3, 15]

3. Attention Is Finite; Not All Effort Is Waste

“Tax” is useful language for a cost that repeatedly stands between an intention and a worthwhile outcome. It becomes misleading when it classifies all effort as something to eliminate.

A conversation can resolve a real disagreement, build a relationship, or compensate for missing information. A manual task can be unnecessary repetition or deliberate practice. A review can exercise important judgment or reread unchanged material because the system cannot show what changed.

Expenditure	What makes it worthwhile?
Meaningful contribution	It creates something the person or group values.
Capability investment	It develops skill, trust, infrastructure, or understanding that is worth the cost.
Necessary coordination	It resolves a real dependency, uncertainty, commitment, or conflict.
Avoidable overhead	It reconstructs, duplicates, or repairs something the environment could reasonably preserve or make easier.

The design question is not simply how to reduce effort. It is **which effort serves the intended outcome, who bears it, and whether the benefit justifies its continuing cost.**

That accounting must span the whole collaboration. A five-minute request that creates an hour of interpretation for its recipient has displaced effort. An AI-generated artifact that saves writing time but expands review and maintenance has changed where the work occurs. A team supported by one person’s constant intervention may look inexpensive only because that person’s contribution is uncounted.

Time also has several forms: elapsed time, active attention, the delay imposed on others, recovery after interruption, and future upkeep. These should not be collapsed into a single productivity score. A faster result that leaves people exhausted or unable to understand their commitments may be a worse result.

The working objective is **less avoidable expenditure for a useful, trustworthy outcome, while preserving the effort people choose for learning, care, and craft.**

4. The Problem Stack

The following stack is a diagnostic map. Its layers interact; they are not a rollout plan or an assertion that every case has the same cause.

Layer	The recurring burden	What may be missing
Usable capacity	Starting, resuming, remembering, and staying oriented	Accessible context, stable cues, and support matched to the person and moment
Intent and priority	Choosing what matters and revisiting unsettled decisions	Purpose, tradeoffs, constraints, and an adequate stopping condition
Shared reality	Searching, reconciling versions, and reconstructing history	Reliable state, provenance, and distinctions between evidence and interpretation
Expression	Translating an idea into another person's or tool's terms	Shared examples, concepts, quality criteria, and scope
Execution	Repeating instructions, handing off, and integrating results	Sufficient context, appropriate capabilities, and completion criteria
Coordination	Waiting, negotiating repeatedly, and seeking permission	Clear dependencies, commitments, responsibility, and authority
Verification	Checking everything and repairing misunderstood work	Evidence, meaningful change summaries, and proportionate review
Projections	Reformatting and maintaining several accounts of the same work	Durable identity and relationships across representations
Learning and change	Mining data, repeating experiments, and repairing stale assumptions	Feedback connected to consequential decisions and affected artifacts
Participation	Proving legitimacy, hiding difficulty, and absorbing invisible work	Support, fair incentives, trust, and legitimate ways to influence the system

A complaint such as “too many meetings” does not identify the mechanism. Missing information, unclear authority, conflicting interests, and a genuinely difficult shared decision can all produce meetings. Removing the meeting without identifying its function may simply relocate the burden.

Proposed root causes

Important relationships are held primarily in people. The reason for a decision, the context behind a request, and the consequence of a change may exist only in someone's memory. Other participants repeatedly have to recover them.

The environment assumes a standard participant. It expects similar vocabulary, attention, knowledge, preferred granularity, and tolerance for uncertainty. People then do the translation the environment has omitted.

Responsibility, authority, and capacity are disconnected. Someone sees a need but cannot act; someone owns an outcome but lacks resources; someone can authorize a decision but lacks the context to judge it.

Costs are displaced across boundaries. Local measures conceal interpretation, support, coordination, or maintenance performed elsewhere. Incentives can reinforce the displacement.

Change propagates through hidden coupling. A small correction requires many independent artifacts and expectations to be reconciled. Adding participants or output increases that maintenance surface.

Improvement is a separate job. Turning a recurring difficulty into a better system requires reconstructing the evidence, explaining it, obtaining agreement, and updating everything afterward. The people closest to the problem may have the least remaining capacity to do that work.

These causes can reinforce each other. Dependence on a few attentive people leaves them less time to improve the environment, making the next change depend on them again. Better information is only part of the response: incentives, legitimate authority, resources, and human relationships remain consequential.

5. Different Minds, Learned Capabilities

The initial conversation used AuDHD and twice-exceptionality as entry points. That lens helped identify high costs from particular kinds of friction, but it should not become the boundary of the problem or a theory of who is capable.

The broader design hypothesis is to learn from people for whom a mismatch is especially visible, then test which improvements transfer. Benefit for one group does not establish benefit for everyone. W3C's cognitive accessibility guidance provides relevant practical grounding: processes should avoid unnecessary reliance on memory, support reorientation, and make help and feedback accessible.[5, objectives 5–7]

The appropriate design unit is the **person, current state, task, environment, and relationship**. A diagnostic label does not determine the right interface. Nor does high reasoning ability establish that someone needs little support.

Consider several ways a person might approach a problem. An artist may recognize an aesthetic mismatch before being able to instruct someone how to correct it. A systems thinker may anticipate dependencies and failure modes while leaving the implementation sequence implicit. A product practitioner may have a clear sense of the desired experience. An operator may recognize the practical consequence of a decision that appears sound in a document.

These examples distinguish **recognizing, imagining, explaining, decomposing, executing, and evaluating**. The capabilities can develop separately and combine differently across domains. Being “visionary” and being good at decomposition are not opposites. A large conceptual leap can be well founded or mistaken; a precise sequence can serve a good objective or optimize the wrong one.

One question for the proposed system is whether examples and comparison can help someone express a distinction they can already recognize. Another is whether explanation and practice can help them develop the capability itself. Assistance should respond to demonstrated needs rather than assigning permanent types.

The ambition is not to make everyone think like the original builder. It is to make differences productive without requiring every participant to understand the entire system.

6. The Abstraction-Distance Problem

Consider the request: “**Make onboarding simpler.**”

Interpretation	A plausible response
Local adjustment	Remove a field, clarify wording, or reorder screens.
Structural change	Eliminate duplicate information entry or combine disconnected workflows.
Conceptual change	Reconsider whether onboarding is a one-time form or continuing support for participation.

The same sentence can reasonably invite different interventions. Knowing the words does not settle which interpretation the speaker intends.

This paper calls the mismatch **abstraction distance**: a difference between the level at which a contribution is intended and the level at which it is interpreted. It is a provisional term, not a validated measurement. It names a practical question: *Are we changing the detail, the approach, or the model underneath it?*

Scale of change is only one dimension. A collaborator also needs to distinguish exploration from execution, a quality example from a binding requirement, a preference from a constraint, and permission to suggest from permission to act. A conceptual request can authorize a discussion without authorizing a rewrite. A detailed request can still leave substantial implementation freedom.

The person making the request may not yet know which level is appropriate. The goal is therefore not to decode an already complete intention perfectly. The interaction can help the person discover and revise the intention.

An important boundary follows: inference about meaning is not authority. Confidence that a collaborator would probably approve an action does not create permission to take it.

7. Working Understanding

Working understanding is the provisional, task-relevant understanding participants rely on to interpret contributions, make decisions, and act together. It includes what they are trying to accomplish, what good looks like, what they assume, what is unresolved, and what each participant can decide.

It does not require identical mental models. It does require enough common ground for the current purpose. Clark and Brennan describe communication as a collaborative process in which participants establish evidence of sufficient understanding, rather than merely transmit utterances. Their account grounds this distinction; it does not validate the platform proposed here.[3]

The design question is how to preserve useful understanding without claiming to model a whole mind.

Concern	What may need to become understandable
The person	Relevant knowledge, demonstrated capabilities, needs, and stated preferences
The domain	Concepts, causal assumptions, examples, and recognized failure modes
The present intent	Desired outcome, scale of intervention, tradeoffs, and stopping condition
The situation	Current state, available resources, constraints, and uncertainty
The relationship	Expectations, delegation boundaries, ways of questioning, and repair
Development	What the person wants to learn and which assistance remains useful

This is more than presentation personalization. Choosing a shorter answer may improve an interface. Learning that “simpler” refers to eliminating responsibility for synchronization changes the interpretation of the work.

Mutual adaptation

The proposed adaptation is reciprocal. The system should become better at collaborating with a person; the person should become better at expressing, evaluating, and revising intent. A correction can improve the artifact and teach a reusable distinction.

An inferred preference must remain a hypothesis. “For this task, you preferred a conceptual comparison before implementation” is a more defensible record than “you always think at a high level.” Current explicit instructions should take precedence over a historical inference. Context-specific learning should not silently become a global rule.

Understanding includes knowing where not to infer. Communicate uncertainty, invite correction, and keep disagreement visible. Predictable agreement is not the same as productive collaboration.

Calibration principle

Learn enough about the person and the work to reduce repeated interpretation, while making consequential assumptions easy to inspect, correct, or discard.

8. The Burden of Seeing the Whole

In the originating discussion, I described an “8K view”: noticing many relationships, gaps, and possible improvements at once. This is a first-person description, not a rank of intelligence or a claim about a diagnostic group.

Several kinds of gap can remain active in that experience. Reality differs from its documentation. Behavior differs from a commitment. The current system does not meet a need. A better possibility is visible. Responsibility exists without corresponding authority.

Those are different conditions. A broken commitment needs a different response from an attractive possibility. An accepted tradeoff should not continually reappear as though it were an unresolved mistake.

The proposed organizational problem is that **noticing can become an implicit obligation to carry**. The observer remembers the issue, explains it, watches for consequences, and intervenes when others lose the context. If that effort remains invisible, the group can mistake constant compensation for a well-functioning system.

A useful observation should be able to become shared evidence, an owned action, an explicit tradeoff, a question, or a deliberately deferred possibility. Its status should be dependable enough that its observer no longer has to keep rediscovering it.

Recording an issue alone may not produce relief. Trustworthy ownership, visible follow-through, and realistic expectations matter. Whether the proposed environment reduces this carrying burden is a question to test.

The intended outcome is to be able to see the whole without having to personally carry the whole.

9. A Loop for Developing Understanding

The first interaction model should allow a person to start with what they have: a sentence, example, sketch, contradiction, draft, or concrete requirement. It should not require a finished objective before helping them form one.

A candidate loop is:

Contribute → interpret → make the interpretation inspectable → clarify or compare → act → observe → update.

The steps can be compressed, repeated, or skipped when they add no value. An inspectable interpretation might be a concise restatement, a small example, or a prototype. A question is useful when its answer changes a consequential choice. Generating alternatives is useful when comparison reveals more than another abstract explanation would.

When a result is rejected, the system should help locate the difference. Was the execution weak? Was the requirement misread? Was the scope too local or too broad? Was a constraint invented? Did the user discover that the initial objective was wrong?

Those are different updates. Rewriting the output without preserving the relevant distinction can make the next interaction repeat the same mistake. Conversely, storing every correction as a permanent personal preference can make learning brittle.

This loop is bidirectional across levels. Strategy can guide a detail; a detail can expose a flawed strategy. Coherence belongs in the relationships among the current commitments, not in a mandatory top-down sequence of thought.

The aim of “getting it right the first time” is therefore less avoidable misunderstanding, not the elimination of discovery. Some iterations are precisely how the right objective becomes known.

10. Shared Reality, Many Projections

Adaptive Collaboration proposes one shared semantic substrate with different views and an attention layer. Its principle is that participants should not have to synchronize separate representations of the same event or decision by hand.[2, §§6–8, 11–12]

Working understanding adds a requirement: preserve not only an object’s current fields, but the relevant meaning, scope, and evidence behind them.

A specification can describe intended behavior while an observation records different actual behavior. A hypothesis can explain the difference while an authorized decision establishes a temporary response. These records can coexist coherently because their statuses differ.

“One source of truth” is therefore used here to mean **a dependable, scoped account of claims, decisions, commitments, and state**. It is not a claim that stored information is infallible, that everyone agrees, or that everything must inhabit one database. This is a clarification of the shared-reality requirement, not a settled architecture.

Projections must preserve meaning

A conversation, product brief, technical specification, task view, agent instruction, and public article may express different aspects of the same work. Each should expose what its audience needs while retaining a route to the relevant source and rationale.

A projection must not silently remove a material uncertainty, promote a suggestion into a decision, or obscure who has authority. When someone introduces new meaning through a projection, that contribution needs a route back to the shared account. Otherwise, tailored views become another source of divergence.

This does not require collecting everything. *Building Through Emergence* calls for making knowledge explicit when another actor needs it for a consequential decision. The practical question here is the smallest durable context that lets the work continue responsibly.[1, §13]

Personal, organizational, and shared context

The proposed platform should distinguish what belongs to a person, what belongs to a collective, and what participants have chosen to share. General working preferences may remain useful across contexts; another organization’s proprietary information does not become portable merely because it influenced an interaction.

The design commitment is user-governed adaptation with explicit scope. People need ways to inspect and revise what the system assumes about them. Useful coordination should not require unrestricted access to private thought or a permanent behavioral profile.

11. AI Delegation and the Human Capability Test

Automating an activity changes the allocation of work. The relevant boundary runs from intent to a usable result, including context preparation, supervision, review, correction, integration, and maintenance.

Bainbridge's account of industrial automation shows why the work left with the human deserves explicit attention: automation can leave difficult monitoring and abnormal-condition responsibilities rather than simply removing work. Applying that concern to AI collaboration is an analogy, not a measured result for this product.[4]

The proposed system should therefore optimize for useful completed outcomes with proportionate human involvement, not for the number of generated artifacts or agent actions. More output can enlarge the review queue at the very point constrained by human attention.

Meaningful delegation needs an understandable objective, relevant context, a boundary of authority, a way to recognize completion, and a route for uncertainty or failure to return. The required rigor should match the consequences and reversibility of the task.

The person should become more capable

A distinctive commitment is that the tool should teach transferable judgment through useful work. A user might learn to distinguish objectives from mechanisms, assumptions from evidence, examples from constraints, or acceptance criteria from a list of implementation steps.

Success can include a person eventually producing equally effective work in Markdown, Word, or Google Docs. Continued product value must then come from useful collaboration and infrastructure, rather than withholding competence or context.

This does not imply that every task should become a lesson. Someone may want deep practice today and direct execution tomorrow. The system should support learning when it serves the person's goals, with scaffolding that can become lighter as capability grows.

For AI, "learning" here means an observable improvement in future collaboration through retained context, revised instructions, feedback, or another mechanism to be evaluated. It does not imply that every interaction changes model weights or that a particular implementation is available.

12. The First Candidate Application: An Intent Workspace

The first candidate is a workspace for turning evolving intentions into documents that people and AI can use. It should be valuable as a standalone tool even if the wider platform never arrives.

A person can begin from their strategy, a rough idea, an existing draft, or a specific mismatch. They can write directly, use a template, ask for an initial draft, or invite questions. The structure should develop only as it helps the work.

Proposed interaction

Capture without premature commitment. Preserve the original contribution and distinguish what the person has asserted from what the system inferred.

Clarify at the useful level. Connect the immediate request to purpose and constraints. Ask about abstraction distance, authority, or quality when ambiguity would materially change the result.

Make correction local. Let the user mark the exact passage or behavior that does not fit, add the missing distinction, and either edit directly or compare a small number of alternatives. Three alternatives are an initial interaction idea, not a principle.

Maintain coherence in both directions. A changed objective should expose affected details. A discovery in the details should be able to reopen the objective. The interface should distinguish contradictions from deliberately unresolved questions.

Produce a usable handoff. Preserve intent, relevant context, instructions, acceptance criteria, and open decisions in a form suitable for the recipient. Export should remain useful without the application.

Carry learning forward selectively. Retain a correction when it improves the artifact, the shared project context, or an explicitly scoped collaboration preference. Keep its source and limits understandable.

These behaviors form a candidate research surface. They are not claims of implemented functionality.

What the document needs to carry

An effective handoff should answer enough of the following to support the task: what is being pursued, for whom, why it matters, what must remain true, what is decided, what is uncertain, what the recipient may decide, and how the result will be judged.

The answer may include examples and counterexamples as well as prose. Not every quality criterion can be reduced to a binary test. Aesthetic, interpersonal, and exploratory work can require informed judgment, with the relevant evaluator and examples made explicit.

A coherence check should help find a mismatch, not manufacture a claim that the document is complete or correct. The user remains able to disagree with the check.

Product boundaries still to determine

The intended access includes local web use and an effective iOS experience, with eventual integration into a broader platform. The particular rendering, native packaging, synchronization, and storage mechanisms remain undecided.

Likewise, the desired AI environments include the tools people already use, such as ChatGPT/Codex, Claude/Claude Code, and separately deployed models. These are integration candidates, not interchangeable APIs or promises of working support. Capability, authorization, context handling, and deployment constraints must be evaluated during implementation.

Templates, repository layout, document schemas, and model adapters remain mechanisms. None should become sacred because it was the first approach considered.

13. Worked Example: “Make Onboarding Simpler”

The following is an illustrative walkthrough, not a report of a deployed system or measured outcome.

A small team is preparing a community application. One collaborator says, “Make onboarding simpler.” An AI proposes fewer fields and clearer labels. The requester rejects the draft: the forms are only the visible symptom.

A low-cost comparison exposes three possibilities: improve the existing form; eliminate repeated information entry; or support a person’s continuing path into participation. The requester recognizes the third as the intended direction, with a near-term constraint: the team needs a useful first improvement rather than a full redesign.

The working intent becomes: **help a new member find a suitable first contribution, understand the commitment, and ask for support without needing to understand the whole organization.**

The operator adds a practical constraint: existing leads cannot absorb a new stream of unowned requests. The designer points out that guidance should remain available after signup. The engineer identifies a small intervention using existing contribution opportunities. Each participant adds knowledge the original sentence did not contain.

The team agrees on a bounded increment. New members can see a small set of suitable opportunities, each with an owner, an understandable commitment, and a way to ask a question. The team does not yet commit to automatic matching or organization-wide workflow changes.

Acceptance criteria now concern the desired participation: can someone identify a plausible contribution, understand what accepting it entails, and reach the responsible person? The group also checks the lead’s burden. An increase in member clicks alone would not establish success.

The product explanation, implementation handoff, and evaluation note are different views of this decision. A later correction to how commitments work should expose the relevant dependent statements rather than require someone to remember every copy.

The reusable learning is scoped: in this task, “simpler” meant reducing the burden of becoming a participant, not merely reducing the number of form fields. It is not recorded as a rule that this person always wants conceptual redesign.

After real use, evidence could challenge any layer. Unowned questions might reveal an operational gap. Poor comprehension might require a local wording change. Low participation might challenge the assumed problem. The system should help the team identify which kind of correction is warranted.

14. From Individual Support to Self-Organizing Groups

The same problem appears at different scales, but the appropriate intervention can differ substantially. A founder working with agents, a group of friends arranging a trip, a volunteer community, and a company have different stakes, permissions, and expectations. Generalization across them remains a research question.

The proposed common requirement is enough shared understanding for local action. Participants should be able to understand the purpose, current state, relevant commitments, available support, and boundaries of their decisions without continually requesting permission or reconstructing history.

This carries forward the “elastic container” from *Building Through Emergence*: preserve intent and boundaries while allowing local adaptation. Its human premise remains that people are ends, not resources to optimize for a system’s continuation.[1, §§8, 10]

Coordination and incentives belong in the model

Coordination concerns the dependencies among contributions. Governance concerns legitimate authority, decision procedures, and repair. Incentives concern what participation makes attractive, costly, recognized, or safe. These are interacting concerns, not problems that disappear when information becomes clearer.[1, §§2, 10]

A team can understand a disagreement perfectly and still have conflicting interests or insufficient resources. The platform should make that situation easier to address; it should not reinterpret dissent as a communication defect.

Human participation includes initiating, shaping, questioning, declining, teaching, caring, and changing the goal. AI can participate through clear capabilities, context, permissions, and accountability for outputs without making humans interchangeable with computational actors.

The aspiration for exceptionally effective teams should mean clarity, complementary capability, initiative, dependable support, and recovery—not permanent emergency intensity. Necessary maintenance and support still need legitimate ownership. They must not become invisible work assigned by default to whoever notices first.

A meaningful test is whether the group can sustain useful action as people join, leave, rest, or change roles. Less dependence on constant compensating intervention is a stronger signal than a brief increase in activity.

15. The Platform That Builds the Platform

The recursive ambition has three parts.

Tools improve the next tools. An intent workspace can help define another capability, preserve decisions, or prepare a collaborator's handoff. Using it exposes limitations that inform its next version.

People become more capable. Better expression, judgment, and shared understanding should improve the quality of future contributions. That is a distinct outcome from making today's artifact faster.

Experience improves the practice. Real use produces cases, questions, failures, and findings. Those can change the product, the working method, and the thesis.

The loops have to produce value before the full platform exists. Otherwise the project becomes a prerequisite for its own usefulness. *Building Through Emergence* supplies the operating discipline: a useful increment should stand on its own, remain composable, and be replaceable as understanding changes.[1, §§4–6, 16–17]

Architectural intent without premature architecture

The enduring requirement is independently useful components that can participate in a coherent whole. A module should have an understandable purpose and boundary, retain the context its operation requires, and integrate without forcing the user to maintain another incompatible account of reality.

A monorepo, multiple repositories, or a hybrid arrangement may serve that requirement. The choice should follow actual development, ownership, reuse, release, and contribution needs. Repository topology is not the thesis.

The same applies to deployment and interfaces. Shared infrastructure should be extracted where repeated use justifies it. An application should not require a complete organizational operating system to solve one person's immediate problem.

As needs become clear, participants should be able to compose or adapt applications and components to their work, including through AI-assisted contributions. Local adaptation should preserve the shared meaning and boundaries on which others depend. Extensibility earns its place by enabling useful variation, not by making every person configure their own platform.

The system must pay for itself in attention

Every abstraction, integration, evaluation routine, and generated artifact has a carrying cost. The platform must make its own cost visible and retire structures that no longer serve the work.

"Entropy" here is a practical metaphor for accumulating inconsistency, stale assumptions, unowned obligations, and maintenance. It is not a physical law used to explain human behavior. A healthy system needs pruning and recovery as well as growth.[1, §§3, 12, 17]

16. Building in Public as Cumulative Evidence

The development practice has several intended returns: useful applications, personal leverage, applied research, peer collaboration, educational material, and evidence of a cross-disciplinary mode of work. These should draw from the same underlying work rather than become unrelated production streams.

A complete increment can create immediate value, leave an inspectable decision trail, and produce a lesson worth sharing. A public account should explain the problem, the judgment exercised, what changed, and what remains uncertain. A polished interface alone may not reveal the broader systems contribution.

The primary near-term professional aim is to make that contribution legible to people who could provide an appropriate scope and environment in which to continue it. A parallel aim is an independently valuable product that could grow into a sustainable company or an open collaborative effort. The work should retain value across those possibilities.

The educational ambition is similarly incremental: articles grounded in actual decisions and outcomes can accumulate into lessons, then potentially a course or book. The learning value should not depend on reproducing one toolchain. It should help others reason, collaborate, and build more effectively in their own circumstances.

Time and sequence matter

A platform, a useful standalone app, a compelling public artifact, adoption, and a professional opportunity mature on different schedules. The development sequence must produce benefit early enough to matter. Completeness of the platform is not automatically the highest-value next step.

Each increment need not produce software, a paper, a lesson, and a career artifact simultaneously. Requiring that would create the overhead the system is intended to reduce. Reuse the evidence where useful; publish only what earns the reader's attention.

"Building in the open" also remains a scoped choice. It does not imply unrestricted release of all source code, conversations, or organizational knowledge. Public claims should be separated from private context, and licensing and contribution boundaries should be explicit when chosen.

The existing method offers an appropriate contribution standard: return context, intervention, expected outcome, observed outcome, surprise, and the resulting model update. Disagreement grounded in experience is useful input.[1, §18]

17. Evaluation: What Would Make the Hypothesis Earn Its Place?

The proposed benefit is cumulative collaboration quality, not simply a more pleasing first response. Evaluation should include repeated use, re-entry after interruption, changing participants, and the cost carried by recipients and maintainers.

Hypothesis	Supporting signal to seek	Result that would challenge it
H1. Calibration reduces repeated misunderstanding.	Fewer recurring scope and meaning corrections on comparable tasks, without hiding assumptions.	Additional memory produces overconfident or stale interpretations.
H2. Local comparison improves intent formation.	Users resolve an important distinction more easily with examples or alternatives.	Alternatives create extra review, anchor the user, or exclude the useful possibility.
H3. Shared context reduces total handoff cost.	Recipients act with less clarification and rework at comparable quality.	Sender savings are offset by recipient interpretation or maintenance.
H4. Different projections improve participation.	People with different needs can contribute without losing material meaning.	Tailored views conceal disagreement, consequences, or important information.
H5. Assistance develops transferable capability.	Users perform unfamiliar or unaided tasks more effectively after use.	Improvement is limited to the tool, or assistance weakens independent judgment.
H6. Recursive use creates practical leverage.	The tool helps complete subsequent work with a manageable support burden.	Building and maintaining the platform consumes more capacity than it releases.

These are proposed tests, not results. No effect size, adoption outcome, or retention claim is established in this edition.

A practical first comparison

Start with a recurring task that already matters, such as producing a handoff another person must act on. Compare the proposed interaction with a credible baseline: a direct conversation, an ordinary document with a template, or the person's existing assistant workflow. The baseline should include reasonable context, not be made artificially weak.

Record time spent preparing, clarifying, reviewing, correcting, and integrating, alongside the usefulness and correctness of the result. Ask the recipient whether they can act. Capture the failures and workarounds that simple timing misses.

Repeated matched tasks can help separate a useful interaction from novelty or task familiarity. A later study may require controlled conditions and independent assessment. Record the tool and model versions used so a changed implementation is not confused with a changed principle.

Begin with the author's work for rapid learning, then include people who do not share the author's vocabulary, preferences, or systems perspective. Designing for a sensitive initial group does not establish universal benefit.

Human outcomes need direct attention: do people understand their choices, retain control, feel able to contribute, and gain useful capability? Engagement and throughput are insufficient substitutes. Measurement itself should remain proportionate to the decision it informs.

18. Failure Modes and Revision Conditions

Over-modeling the person. Historical behavior becomes a fixed identity. Revise toward explicit, scoped, revisable assumptions; discard inferences that do not improve collaboration.

Sycophancy presented as understanding. The system mirrors the person's language while failing to challenge a mistaken premise. Evaluate whether disagreement and counterexamples improve decisions, not whether every response feels affirming.

Premature interpretation. The system decides what an incomplete contribution means too early. Preserve the source and make meaningful alternatives available before a consequential commitment.

False coherence. A clean narrative hides uncertainty or legitimate disagreement. Keep claim status and decision authority distinct from stylistic consistency.

Bureaucracy in the name of context. Every contribution requires tagging, classification, or explanation. Remove structure whose upkeep exceeds its demonstrated benefit.

Automation that relocates burden. Generated work expands supervision or leaves difficult recovery to an unprepared human. Reconsider the delegation boundary using the full task cost.

A platform-shaped delay. Shared infrastructure postpones useful applications, while public writing substitutes for contact with real use. Return to an independently valuable intervention.

Hidden dependency. The product appears effective because the founder, coordinator, or expert quietly repairs it. Include their effort and test what happens when they step away.

A failed test need not invalidate the whole direction. Change the lowest layer that explains the evidence: local implementation, workflow, decision, hypothesis, or underlying commitment. Keep the reason for that revision connected to affected work.^[1, §24]

19. Current Conclusions and Open Questions

The strongest current conclusion is a reframing: **the problem is not exhausted by producing better instructions or storing more information. It includes developing and maintaining enough understanding for different participants to act together.**

The resulting design commitments are to preserve human agency, make important assumptions and authority inspectable, support different modes of contribution, evaluate total collaborative cost, and deliver value before requiring a complete platform.

The working hypotheses are that calibration can reduce repeated misunderstanding; that linked context and projections can reduce repair; that assistance can increase transferable capability; and that recursive use can make the system easier to develop. They remain to be tested.

Several questions are deliberately open. What is the smallest context worth retaining? When is a question better than an example or a reversible action? Which learning belongs to a person, a task, or a relationship? How should contributors correct the system's model of them? Which distinctions continue to matter across changes in models and tools? When does a standalone capability justify shared platform infrastructure?

The next useful move is a bounded application of the intent workspace to real work, with a credible comparison and explicit revision conditions. Its first proof should be a useful outcome and less unnecessary burden—not a larger architecture.

Foundation

Build conditions in which people can understand one another, exercise judgment, and contribute without repeatedly reconstructing the whole system. Let useful understanding accumulate. Let the people become more capable. Let the tools remain replaceable. Use what the work teaches to improve the system that enables the work.

20. Provenance, Related Work, and Evolution

This edition synthesizes the October 2026 working discussion that produced it: the platform-building intent; the attention and coordination tax model; the neurodiversity entry point and subsequent broadening; abstraction distance; mutual adaptation; and the proposed intent workspace. It preserves those as reported experience, interpretations, design commitments, and hypotheses rather than promoting them into measured findings.

The principal source documents are the supplied working editions of *Building Through Emergence* and *Adaptive Collaboration*, both dated 2026-09-20. The relationship to the System of Work thesis follows those editions: a developing body of practice and investigation, open to testing and revision, distinct from the dated field evidence. The earlier papers are the source of that lineage; this edition does not claim to restate the entire thesis.

The selected external references below situate mechanisms already invoked in the discussion. They do not constitute a systematic literature review or validate the proposed platform. Broader claims about expertise, diagnosis, team performance, belonging, or retention require additional evidence and are not established by this paper.

Source lineage

Source	Contribution to this edition
Originating discussion, October 2026	Practitioner problem statements; intended product and professional outcomes; the shift from cognitive overhead to mutual understanding
Building Through Emergence	Continuous practice; claim hygiene; the elastic container; useful increments; knowledge in the runtime; human ends; revision discipline
Adaptive Collaboration	Shared reality; attention; emergent structure; multiple projections; re-entry; the situated community coordination case
Selected related work	Grounding in communication; the human responsibilities left by automation; cognitive accessibility and reorientation

References

[1] **Christian Blank.** *Building Through Emergence: A working blueprint for observing, building, operating, and learning together while preserving human dignity and the capacity to change.* Working edition, 2026-09-20. References above identify sections of the supplied 17-page edition. [Source PDF](#).

[2] **Christian Blank.** *Adaptive Collaboration: A working investigation into attention, shared reality, and structures that emerge from how people actually collaborate.* Working edition, 2026-09-20. References above identify sections of the supplied 14-page edition. [Source PDF](#).

[3] **Herbert H. Clark and Susan E. Brennan.** "Grounding in Communication." In *Perspectives on Socially Shared Cognition*, edited by Lauren B. Resnick, John M. Levine, and Stephanie D. Teasley, 127–149. American Psychological Association, 1991. Used for the distinction between transmitting an utterance and collaboratively establishing sufficient understanding. [Paper](#).

[4] **Lisanne Bainbridge.** "Ironies of Automation." *Automatica* 19, no. 6 (1983): 775–779. Used for the need to examine monitoring, recovery, and other responsibilities remaining with the human. Its original setting is industrial automation. [Paper](#).

[5] **World Wide Web Consortium.** *Making Content Usable for People with Cognitive and Learning Disabilities.* W3C Working Group Note. Consulted during preparation, 2026-10-13; particularly the objectives addressing focus, memory, and help. Used as accessibility design guidance, not as evidence that one interface benefits every population. [Guidance](#).

Evolution record

v0.1 · 2026-10-13. Initial synthesis. Establishes working understanding as the central proposed mechanism; connects it to the problem stack, the first application, the public development practice, and six testable hypotheses. No product validation is claimed.

Future editions should record the observation that prompted a change, the layer revised, affected claims or artifacts, and what would now count as counter-evidence. The goal is cumulative understanding rather than cumulative text.