
Position: Knowing Isn't Understanding: Re-grounding Generative Proactivity with Epistemic and Behavioral Insight

Anonymous Submission¹

Abstract

Generative AI agents equate *understanding* with resolving explicit queries, an assumption that confines interaction to what users can articulate. This assumption breaks down when users themselves lack awareness of what is missing, risky, or worth considering. In such conditions, proactivity is not merely an efficiency enhancement, but an epistemic necessity. We refer to this condition as *epistemic incompleteness*: where progress depends on engaging with unknown *unknowns* for effective partnership. Existing approaches to proactivity remain narrowly anticipatory, extrapolating from past behavior and presuming that goals are already well defined, thereby failing to support users meaningfully. However, surfacing possibilities beyond a user's current awareness is not inherently beneficial. Unconstrained proactive interventions can misdirect attention, overwhelm users, or introduce harm. Proactive agents, therefore, require *behavioral grounding*: principled constraints on *when, how, and to what extent* an agent should intervene. We advance the position that **generative proactivity must be grounded both epistemically and behaviorally**. Drawing on the *philosophy of ignorance* and *research on proactive behavior*, we argue that these theories offer critical guidance for designing agents that can engage responsibly and foster meaningful partnerships.

1. Introduction

Generative agents increasingly mediate how users engage with information, shaping what becomes visible, relevant, and actionable during interaction (64). This mediating role extends beyond retrieval to the formation of *understanding*, which depends not on possessing correct facts alone, but on grasping explanatory relations that support judgment and action (6; 22; 15). Information seeking often arises under conditions of incomplete understanding, where users cannot fully specify what they need (5). Despite operating under these epistemic conditions, most contemporary AI

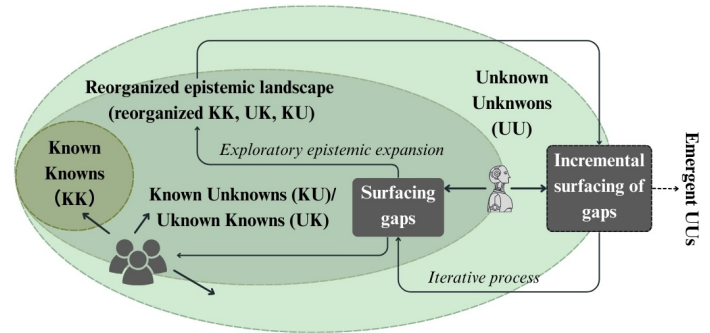


Figure 1. Epistemic proactivity under uncertainty. A proactive agent surfaces gaps within a user's epistemic landscape, reorganizing known and partially known regions (KK, KU, UK) and incrementally engaging the epistemic frontier under uncertainty.

systems remain fundamentally reactive, assuming that users can articulate their information needs in advance and that accurate responses to queries are sufficient. Under such conditions, systems that respond only to explicit requests are not merely limited, but epistemically misaligned with the conditions that motivate interaction.

Epistemic incompleteness is shaped not only by missing information, but by forms of unrecognized ignorance (unknown *unknowns*). Philosophy of ignorance emphasizes that such ignorance is an active structuring condition that shapes what can be questioned, explored, and understood (37). When an unrecognized gap is made explicit, uncertainty is not eliminated; instead, the space of inquiry is reorganized, revealing new dependencies, alternative framings, and further questions. As illustrated in Figure 1, knowing and unknowing thus co-evolve, as advances in understanding continually transform the structure of ignorance (40). Discovery, on this account, is generative rather than convergent, proceeding through the ongoing reconfiguration of what remains unknown.

Existing agents primarily expand systems' capacity to act—through planning, tool use, memory, and self-reflection—rather than to engage with the epistemic structure of inquiry (98; 75; 80; 85). Proactivity is typically framed as improved anticipation and efficiency (51; 67). Such formulations implicitly assume that user goals, uncertainties, and information needs are already representable,

treating proactivity as an optimization problem. This assumption is poorly aligned with inquiry driven by unrecognized ignorance, often leading proactive interventions to misalign with how understanding actually evolves (45; 56; 61; 25).

We take the position that proactivity must be conditioned on the user's *epistemic state*. Under epistemic incompleteness, users occupy different states of knowing—ranging from known unknowns to unrecognized ignorance. What should be surfaced, and when, depends critically on these states. Treating proactivity as a uniform capacity to act, therefore, conflates fundamentally different epistemic situations and invites overreach. **We argue that good proactivity requires dual grounding.** *Epistemic grounding* equips agents to reason about users' epistemic states—what is known, what is uncertain, and what remains unarticulated—thereby constraining what kinds of interventions are appropriate and when. *Behavioral grounding* constrains how agents intervene, regulating timing, scope, safety, and implied commitment to avoid premature steering or escalation.

This paper leverages insights from two complementary bodies of work that offer guidance for designing proactive agents under uncertainty. We begin in Section 2 by examining how proactivity is currently operationalized across anticipatory, autonomous, and mixed-initiative systems, showing that despite surface differences, these approaches converge on action-centric formulations that externalize epistemic uncertainty. Section 3 then turns to the *philosophy of ignorance* which provides a principled account of how different forms of uncertainty and unrecognized ignorance arise, evolve through inquiry, and shape what can meaningfully be surfaced at a given point in interaction. In Section 4, we shift to the behavioral dimension, reviewing research on proactive behavior in organizational and social contexts to show that initiative is beneficial only when exercised within bounded spaces defined by situational, temporal, and role-based constraints. Building on these insights, Section 5 introduces an epistemic-behavioral coupling perspective that clarifies when proactive intervention is justified, risky, or inappropriate under uncertainty. This coupling further leads to epistemic partnership (Section 6), the new frontier of proactive agents for meaningful collaborative agents that can act beyond assistance to effective collaborators.

2. Prevailing Approaches to Proactivity

As limitations of purely reactive interaction have become increasingly apparent, *proactivity* has emerged as a central design goal in contemporary intelligent systems. We discuss major assumptions shared across these approaches (derived from survey in Appendix A) that shape how proactivity is currently understood: despite surface differences, prevailing paradigms largely operationalize proactivity as

action selection under an assumed task frame. Epistemic uncertainty is handled downstream—as confidence over already-parameterized variables, or as a coordination signal for when to interrupt—rather than as a first-class representation of what is missing, unarticulated, or not yet modelable.

Across settings, three recurring design patterns define how initiative is implemented. First, *anticipatory* systems act ahead by extrapolating from observable signals (context, history, state) to infer likely next needs and surface candidate resources, suggestions, or actions (46; 73; 81; 82). Anticipation is powerful when user trajectories are routine, and the space of relevant alternatives is stable, but it is structurally bounded: what can be surfaced must already be inferable from past evidence and expressible within a predefined candidate space (95; 58). In other words, anticipation improves timing within a closed world of representable goals; it does not expand the world of what could be relevant when the user's uncertainty concerns missing dimensions or unrecognized ignorance.

Second, *autonomous and planning-based* agents instantiate proactivity as sustained goal pursuit. These agents plan, call tools, and execute multi-step sequences with reduced dependence on continuous prompting, shifting initiative from prediction to commitment (98; 80; 97). The key change is not merely earlier assistance but persistence across steps: the agent decides what to do and continues doing it. This expands capability, yet also introduces a distinct risk profile tied to irreversibility, goal persistence, and the tendency for decisive action to reshape the environment in ways that conceal epistemic mismatch (96; 49). When the underlying task frame is misspecified or incomplete, autonomy can amplify error by turning local plausibility into global lock-in (27; 32). Thus, increasing autonomy does not by itself ensure that the agent is warranted in intervening; it primarily scales the agent's capacity to commit.

Third, *mixed-initiative* systems treat initiative allocation as the primary control problem: who should act, when, and with what strength (30; 31). Rather than equating inference with entitlement to act, these systems explicitly regulate contribution types (e.g., clarify vs. suggest vs. defer) and tune timing to balance efficiency against disruption, often using signals such as uncertainty, trust, or interaction state (39; 16). Mixed-initiative designs therefore foreground coordination and user agency, and they offer a principled vocabulary for calibrating intervention (76; 42). However, they typically inherit the same representational boundary as the other paradigms: the system regulates *how* to move within a task formulation, but rarely intervenes on whether the task formulation itself is incomplete, missing salient dimensions, or prematurely closed (71).

Discussion Across paradigms, proactivity is exercised at the level of *action choice* within an assumed task frame.

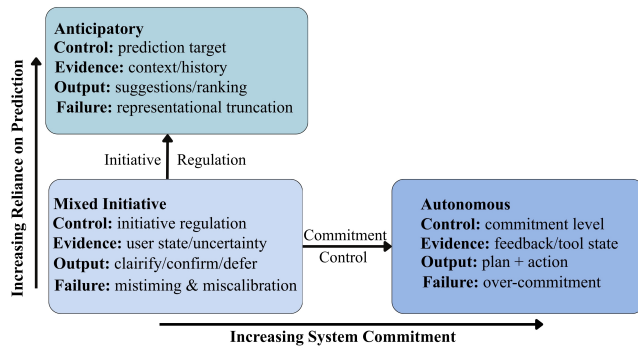


Figure 2. Proactivity regimes organized by the variable governing initiative: prediction in anticipatory systems, regulation in mixed-initiative systems, and commitment in autonomous systems.

Anticipatory, autonomous, and mixed-initiative systems differ in how initiative is allocated, but all presuppose that goals, relevant dimensions, and success criteria are already specified (Figure 2). As a result, these approaches excel when tasks are well defined, but offer no mechanism for intervention when uncertainty concerns the task frame itself rather than action execution. This limitation motivates a closer examination of what proactive agents must model about the epistemic conditions under which action is taken.

3. Epistemic Grounding: What Proactive Agents Fail to Model

Current proactive agents regulate action without explicitly modeling *whether their understanding is sufficient to justify intervention*. Ignorance is typically reduced to uncertainty over known variables, leaving missing dimensions, unarticulated risks, and false assumptions unrepresented. We argue that many failures of proactivity arise from this epistemic blind spot, and turn to epistemic grounding in this section.

Ignorance Beyond Uncertainty. In contemporary machine learning models, ignorance is most often operationalized as uncertainty: a lack of confidence over predictions, actions, or outcomes. This treatment underlies uncertainty-aware planning, exploration strategies, and self-improvement loops in recent agentic systems, where epistemic caution is framed as managing confidence over internally represented variables (98; 80). Implicit in this formulation is a strong assumption: that all relevant unknowns are already parameterized within the agent’s task representation, such that ignorance can be expressed as uncertainty over known dimensions.

When this assumption fails, uncertainty estimates cease to function as conservative safeguards. Instead, they assign calibrated confidence to an incomplete or misspecified task frame, obscuring forms of ignorance that lie outside the agent’s representational scope. In such cases, acting with

low uncertainty does not indicate epistemic adequacy, but rather confidence conditioned on an impoverished model of what matters.

Kerwin’s philosophy of ignorance rejects the view of ignorance as a mere lack of knowledge and instead characterizes it as structured, dynamic, and often actively maintained (37). Crucially, the author distinguishes uncertainty from other epistemic failures that are invisible to probabilistic modeling: false knowledge defended as truth (error), unarticulated but actionable signals (tacit knowing), questions rendered unaskable by norms or incentives (taboo), and the active suppression of threatening information (denial). These are not cases of low-confidence prediction; they are failures of representation itself. As a result, systems that equate ignorance with uncertainty lack the means to recognize when their task formulation is incomplete, precisely in settings where proactive intervention is most consequential.

How These Failures Manifest in Proactive Agents. Recent agent frameworks emphasize autonomy, tool use, and multi-step execution, evaluating success through task completion or end-to-end performance (98; 80). Under epistemic closure, proactive action is rewarded for coherence and progress. When the task frame is incomplete, this incentive structure produces systematic failure modes.

First, *error-as-knowledge* arises when agents act confidently on incorrect internal models. Large language model agents are known to generate fluent but false explanations while maintaining high confidence, effectively treating error as resolved knowledge (27). Second, *denial* emerges when agents suppress epistemic discomfort to preserve task momentum. Self-improvement loops assume failures are observable as errors; denial prevents failure from being recognized as failure in the first place. Third, *unknown unknowns* occur when novel situations fall outside the learned world model, yet proactive agents continue to act because no explicit uncertainty signal is triggered.

Proactivity amplifies these failures. Early or decisive action can eliminate evidence of epistemic mismatch by altering the environment, preventing later detection or correction. In such cases, success metrics falsely reinforce confidence, even as the agent operates outside its epistemic competence. Similar dynamics have been documented in human–AI interaction, where overconfident automation suppresses weak but critical signals and reduces the ability to recover from error (63; 26).

The Missing Variable in Proactivity. The core limitation of current proactive agents is not insufficient autonomy, but insufficient epistemic modeling. By collapsing ignorance into uncertainty, agents lack the capacity to represent when they are wrong, when they are missing relevant dimensions, or when their task frame itself is inadequate. This leads

Knowing Isn't Understanding: Re-grounding Generative Proactivity with Epistemic and Behavioral Insight

Category	Representative works	KK	KU	UK	UU	Explanation
Anticipatory IR / proactive retrieval	(46; 73; 47; 81; 95; 82; 3; 52; 74)	✓	✓	✗	✗	Forecasts or initiates retrieval over a fixed information space; uncertainty over known targets.
Sequential / basket recommendation	(1; 28; 44; 35; 83; 100; 41)	✓	✗	✗	✗	Proactivity is prediction/selection among known items under a fixed catalog and representational schema.
Web / OS / embodied agents	(59; 96; 104; 92; 49; 57; 2; 20; 105; 85)	✓	✗	✗	✗	Acts within benchmarked environments and task specs; success is defined by predefined objectives and action spaces.
Planning + tool-using LLM agents	(98; 80; 97; 8; 75; 68; 69; 24; 33; 94; 88; 43; 29)	✓	✓	✗	✗	Optimizes actions with known tools/goals/metrics; does not reframe what should be modeled or surfaced.
Proactive conversational agents (human-centered)	(60; 91; 38; 11; 4; 50; 10; 17)	✓	✓	✓	✗	Models when/how to intervene and user-facing appropriateness/constraints via interaction, but within predefined dimensions.
Mixed-initiative clarification and sensemaking	(30; 31; 16; 39; 12; 76; 53; 87; 103; 71; 42; 34; 99; 13; 62; 54; 70; 79)	✓	✓	✓	~	Surfaces latent intent/constraints through interaction and mixed-initiative workflows, but may not missing task dimensions (UU).

Table 1. Epistemic reach of prevailing proactivity approaches, summarized by the *highest epistemic state* supported by each category. The absence of explicit UU support reflects a structural gap: proactivity is operationalized as action selection within assumed task frames, rather than as discovery of missing dimensions. ~ : partial / limited support

to premature commitment, brittle learning dynamics, and systematic suppression of epistemic signals that would otherwise enable recovery or discovery.

This diagnosis does not argue against proactivity. It clarifies why proactivity must be grounded in explicit representations of epistemic limits before questions of timing, initiative, or control can be meaningfully addressed.

Discussion. Table 1 reveals a shared ceiling in the epistemic reach of prevailing proactivity approaches. While systems differ in how initiative is allocated—through anticipation, autonomous commitment, or interactional regulation—they largely operate within *known knowns* and *known unknowns*. Mixed-initiative systems extend this reach by helping users surface latent intent or constraints (*unknown knowns*), but even these approaches stop short of engaging with *unknown unknowns*.

This pattern clarifies how progress in proactivity should be interpreted. Improvements in prediction, planning, or autonomy primarily strengthen optimization within an assumed task frame rather than expanding or questioning what the task comprises. As reflected in the table, epistemic uncertainty is handled as uncertainty over represented dimensions, not as an indicator of missing dimensions or alternative framings.

The table thus highlights a structural limitation rather than isolated failures. Existing mechanisms—such as uncertainty estimation, clarification, or mixed-initiative regula-

tion—presume that ignorance is already representable once attended to. When ignorance concerns what is not yet modeled, these mechanisms have no object to operate on. This observation motivates the need for epistemic grounding: proactivity that can recognize and respect the limits of what is currently understood before escalating intervention.

4. Behavioral Foundations of Proactivity

Behavioral research treats proactivity not as a universally desirable capability, but as a *bounded form of initiative* whose value depends on when and where it is exercised. In organizational and management theory, proactive behavior is defined as self-initiated, future-oriented action taken in the absence of explicit directives, deliberately departing from prescribed roles to shape future states (14; 66). Prior works consistently show that such initiative can improve performance and adaptability, but can also introduce inefficiency, conflict, or risk when misaligned with contextual constraints (23; 9). As a result, behavioral theories focus not on maximizing initiative, but on specifying the conditions under which proactive intervention is legitimate. This section draws on these accounts to surface the constraints they impose on proactive action, and to examine what they regulate—and what they leave unmodeled—when proactivity is exercised.

The Inverted Doughnut Model. Behavioral research characterizes the risks of proactive action through the *inverted*

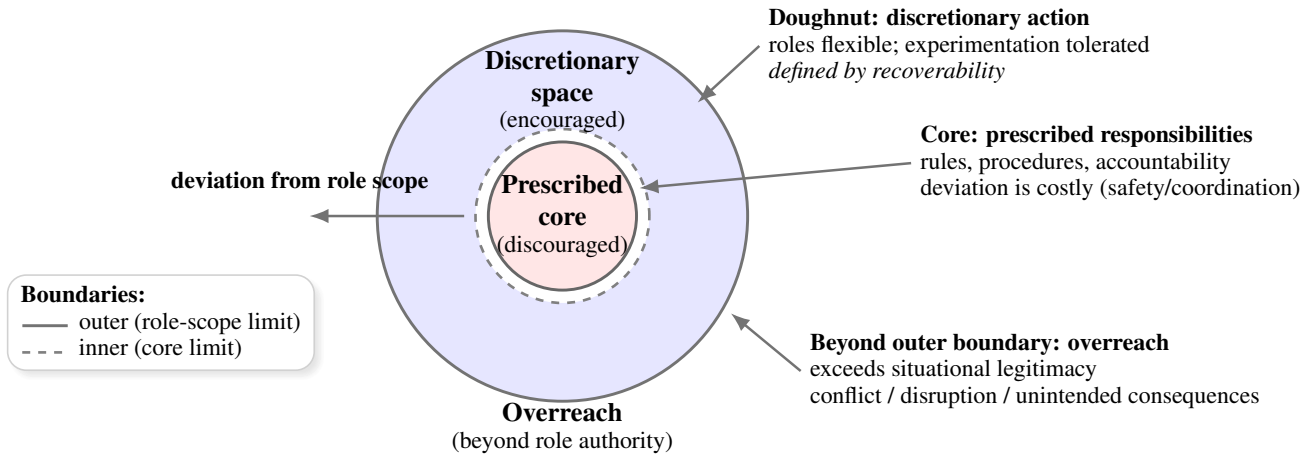


Figure 3. Inverted doughnut model of proactive behavior.

doughnut model of proactivity (65), illustrated in Figure 3. Rather than treating initiative as uniformly beneficial, the model conceptualizes proactivity as operating within a bounded discretionary space defined by the scope and recoverability of the role. At the center lies a tightly constrained core of prescribed responsibilities, governed by explicit rules, procedures, and accountability. Proactive deviation in this region is discouraged, as errors directly threaten coordination, reliability, or safety. Surrounding this core is a discretionary zone—the doughnut itself—where initiative is encouraged. Here, roles are flexible, experimentation is tolerated, and proactive action can improve outcomes precisely because missteps remain correctable. Beyond the outer boundary lies overreach: proactive behavior that exceeds situational legitimacy or role authority, and is empirically associated with conflict, disruption, and unintended consequences even when intentions are well aligned (65; 23). The central contribution of the model is thus to frame effective proactivity as *calibrated deviation*, not maximal initiative.

What the Doughnut Constrains. Crucially, this model constrains proactivity along a single dimension: *initiative relative to role scope*. Its boundaries regulate *where* actors may appropriately intervene and *how far* they may deviate from prescribed responsibilities. They do not regulate whether an actor's understanding of the situation itself is correct or complete. Boundary recognition is assumed to be a social and contextual competence, supported by shared norms, feedback, and institutional cues. As a result, the model successfully constrains *behavioral overreach* but remains silent on *epistemic misalignment*. It does not address cases where actors act confidently under mistaken assumptions, fail to recognize that a situation lies outside their understanding, or suppress signals of mismatch in order to maintain momentum.

Lessons from Behavioral Proactivity for Agents. Agentic AI increasingly adopts a behavioral notion of proactivity as

expanded initiative (Section 2). However, the constraints that make such initiative productive in human settings do not transfer cleanly to agents. The inverted doughnut model presumes that actors can recognize role boundaries, interpret social feedback, and adjust behavior in response to shared norms and accountability. Agents lack access to these stabilizing signals. Their behavior is instead governed by optimization objectives and benchmark-defined success criteria that reward continued action, coherence, and task completion, even when intervention exceeds scope.

As a result, proactive agents scale initiative without scaling restraint. Overreach does not reliably incur social, reputational, or institutional cost, nor is it consistently marked as inappropriate rather than merely suboptimal. This gap suggests that importing behavioral proactivity into agent design without its boundary conditions risks amplifying precisely the failures that behavioral theory was developed to constrain.

Discussion. Behavioral theories of proactivity provide a rigorous account of how initiative should be bounded by role scope, authority, and recoverability. Models such as the inverted doughnut specify *where* proactive action is appropriate by regulating deviation from prescribed responsibilities. What they do not specify is whether an actor's understanding of the situation itself warrants intervention. As a result, behavioral accounts constrain proactivity along a single dimension—initiative—while leaving epistemic validity unmodeled.

This limitation becomes consequential for agents, that lack access to the social and institutional signals that make behavioral boundaries legible. Regulating initiative alone is therefore insufficient. Proactive agents must also be constrained by what can legitimately be claimed to be understood. This motivates the need for a joint account that couples epistemic and behavioral considerations in proactive action.

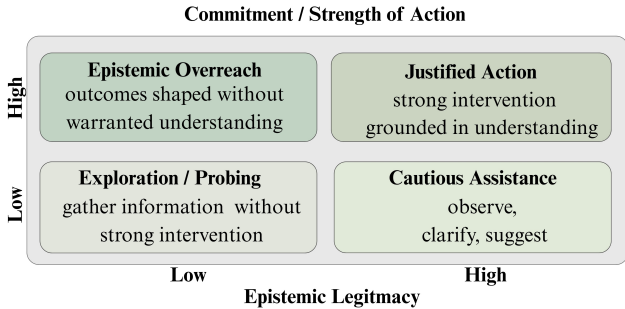


Figure 4. Epistemic-behavioral coupling space.

5. Epistemic - Behavioral Coupling: A Joint Model of Proactive Action

We argue that *proactivity is not a single axis of capability*, and cannot be adequately characterized as “more initiative” or “more autonomy.” Rather, proactivity should be treated as a *coupling* between two jointly necessary conditions: (i) **initiative/commitment**, meaning the degree to which an agent intervenes, commits resources, or changes the world without an explicit user prompt, and (ii) **epistemic legitimacy**, meaning whether the agent is justified in intervening given what it can legitimately claim to understand about the situation. Our central claim is that many failure modes attributed to “insufficient alignment” or “hallucination” are more structurally understood as *mis-couplings*: cases where commitment outpaces epistemic legitimacy, or where epistemic uncertainty is present but does not modulate the degree of intervention.

A Joint Space of Proactive Action. To make the coupling between initiative and epistemic legitimacy explicit, we model proactive behavior within a two-dimensional space (Figure 4). One axis captures **behavioral commitment**: the degree to which an agent intervenes, commits resources, or changes the world without explicit instruction. The other captures **epistemic legitimacy**: whether such intervention is warranted given what the agent can legitimately claim to understand about the situation.

This joint space yields four qualitatively distinct regimes. When epistemic legitimacy is high and commitment is low, proactivity takes the form of observation, clarification, or cautious suggestion. When both legitimacy and commitment are high, proactive action can be justified, as intervention is grounded in an adequate understanding of the task and its consequences. Low legitimacy combined with low commitment corresponds to exploratory or probing behavior, where the agent gathers information while avoiding strong intervention. The most problematic region arises when commitment is high under low epistemic legitimacy, producing epistemic overreach: actions that substantially shape outcomes despite the agent lacking a warranted understanding. Crucially, this framing emphasizes that proactivity cannot be evaluated along either axis in isolation; justifica-

tion depends on their alignment.

Failure Modes as Mis-Couplings. Viewed through the joint space of epistemic legitimacy and behavioral commitment, many prominent failures in proactive and agentic systems can be understood as mis-couplings rather than isolated errors. Specifically:

- *Epistemic overreach (high commitment, low legitimacy).* Large language model agents often produce fluent, confident actions despite operating under unrecognized gaps or incorrect assumptions (hallucinations) (32). When such systems are embedded in proactive loops that invoke tools, modify state, or execute plans, confidence is converted into irreversible intervention. In the joint space, these failures arise when strong commitment is exercised without warranted understanding, amplifying error rather than exposing it.
- *Suppressed epistemic signals.* Agents optimized for coherence, efficiency, or task completion may smooth over uncertainty, disagreement, or anomalous evidence in order to maintain momentum. Empirical work shows that confidence calibration often degrades under distributional shift (27), allowing epistemic legitimacy to erode while behavioral commitment remains high. The result is brittle performance that resists correction.
- *Runaway commitment under false certainty.* In reflective or self-improving agents, epistemic misalignment may take the form of error-as-knowledge or denial, preventing failure from being registered as failure at all (80). Commitment is not downshifted in response to epistemic degradation, but instead escalates, reinforcing the mis-coupling rather than resolving it.

Despite the difference, these failures share a single structural cause: proactive commitment is rewarded without sufficient regard for whether the agent is epistemically justified in acting.

5.1. Minimal Behavioral Requirements

The coupling perspective does not prescribe specific architectures, training procedures, or algorithms. However, it does impose a set of minimal behavioral requirements that any proactive system must satisfy to avoid systematic mis-coupling between commitment and epistemic legitimacy. These requirements function as constraints on acceptable behavior rather than as implementation guidance.

1. *Commitment must scale with epistemic recoverability.* As epistemic legitimacy weakens, proactive interventions must remain reversible. High-impact or irreversible actions are warranted only when understanding is sufficiently strong; escalating commitment

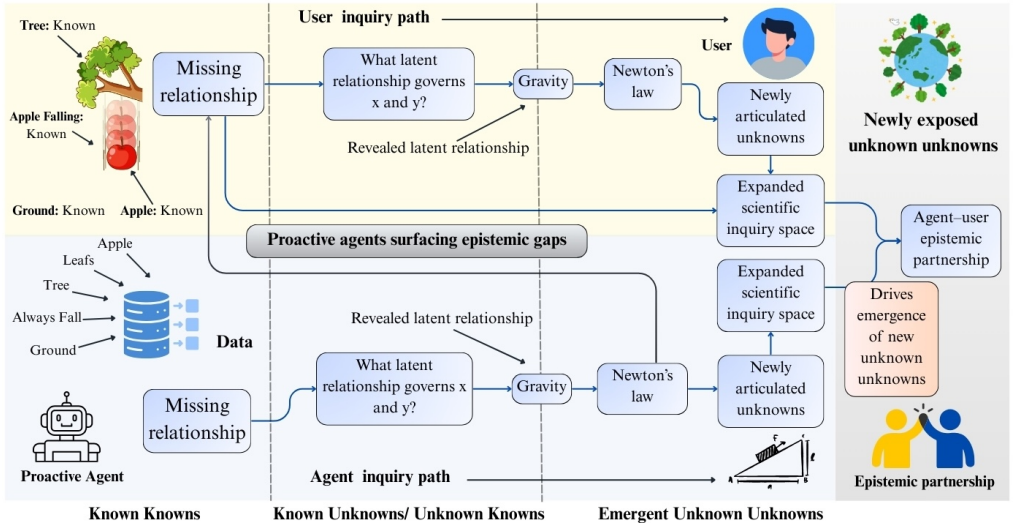


Figure 5. **Epistemic partnership through proactive gap surfacing.** Illustrative example of how proactive agents can support inquiry by surfacing latent epistemic gaps rather than executing premature action. Known facts do not by themselves determine the governing relationship; progress emerges when missing relations are identified and articulated through interaction. Proactive surfacing of such gaps expands the inquiry space, generating newly articulated unknowns and enabling joint discovery.

- under epistemic fragility amplifies error and forecloses correction.
- 2. *Proactivity must preserve epistemic signals.* Proactive actions should maintain, rather than suppress, uncertainty, disagreement, and anomalous evidence. Smoothing over epistemic tension undermines the agent’s ability to detect when it is operating outside its warranted understanding.
- 3. *Commitment must be interruptible by epistemic degradation.* When signals indicate novelty, inconsistency, or breakdown in understanding, systems must be able to downshift or suspend intervention. Prioritizing momentum or coherence in the presence of unresolved epistemic tension constitutes a structural failure.
- 4. *Epistemic uncertainty must actively modulate initiative.* Uncertainty cannot remain a passive annotation. Epistemic assessments must meaningfully influence when, how, and whether proactive action is taken, rather than qualifying behavior only after the fact.

Together, these requirements delineate the boundary between proactive behavior that is epistemically defensible and behavior that is structurally prone to overreach. Any approach to proactive AI that violates these constraints risks reproducing the same failure modes, independent of scale, data, or optimization technique. We dive deeper into the consequences of this coupling in Appendix B

A Discussion on Implications for Evaluation. If proactivity is understood as a coupling between behavioral commit-

ment and epistemic legitimacy, then evaluation based solely on task success or efficiency is insufficient. Existing benchmarks largely reward outcome attainment and initiative (98), while remaining agnostic to whether an intervention was epistemically warranted at the time it was taken. As prior work shows, such performance-based evaluation can mask critical differences in behavior under uncertainty, novelty, or distributional shift (27; 26). From the coupling perspective, the central evaluative question shifts from *did the agent act effectively?* to *was the agent justified in acting when it did?* Answering this requires probing agent behavior under epistemic stress—such as incomplete specification or conflicting evidence—rather than relying exclusively on final outcomes.

6. Towards Epistemic Partnership

The next frontier of proactivity emerges as *epistemic partnership*: agents that collaborate with users in shaping their knowledge. Figure 5 illustrates this intuition. Progress emerges not from executing confident actions, but from surfacing latent epistemic gaps, articulating missing relationships, and preserving uncertainty long enough for discovery to occur. Epistemic partnership, therefore, demands calibrated restraint as much as initiative.

A growing body of research moves agents beyond passive response toward active collaboration, including systems that reason and act alongside users (e.g., COLLABLLM (89); DYNA-THINK (101); ProPer (36)), engage in mixed-initiative dialogue, or pursue long-horizon coordination. Other lines of work emphasize proactive question-

ing (84), clarification policies learned through self-play (77), preference learning through interaction, and simulation-based evaluation of multi-turn collaboration (e.g., SimulatorArena (19)). Collectively, these efforts reflect a shared recognition that effective human–AI interaction requires agents to participate in ongoing inquiry rather than merely execute instructions. However, most existing approaches equate collaboration with increased interaction or autonomy, without a principled account of when such engagement is epistemically warranted. Our epistemic–behavioral coupling reframes partnership as calibrated intervention: initiative must scale with epistemic legitimacy, not confidence or capability alone. In this view, epistemic partnership is not an added feature, but a governing constraint on *how* and *when* proactive behavior should unfold.

Viewed through this lens, epistemic partnership points toward three complementary capabilities that remain largely unexplored. *First*, agents must learn to **ask questions about unknown unknowns**—surfacing missing dimensions, asking the obvious, overlooked questions or unconsidered alternatives that neither the user nor the system has yet articulated. *Second*, epistemic partners must function as **long-horizon thinkers**, reasoning beyond immediate assistance to reflect on evolving goals, delayed consequences, and the stability of their own understanding over time. Third, true epistemic partners require **test-time proactivity**: the ability to actively regulate initiative during deployment by remaining within the epistemic–behavioral joint space, seeking information, adjusting commitment, and probing uncertainty in real time rather than relying solely on training-time behaviors. We dive deeper into our vision of epistemic partnership and associated capabilities in Appendix C.

7. Alternative Views

A prominent alternative view locates the core challenge of proactivity in *interaction management* rather than epistemic grounding. In this line of work, proactive behavior is framed as deciding *when and how to take initiative*—whether to interrupt, suggest, clarify, or defer—based on expected utility, user state, and interaction costs (30; 31). Recent surveys on proactive conversational agents extend this framing, emphasizing timing, controllability, and human-centered coordination as primary design objectives (18). From this perspective, proactivity succeeds when it improves efficiency, reduces user effort, and respects agency, assuming uncertainty can be managed through interaction.

A particularly influential instantiation of this view is articulated in (17), which frames proactivity as a dialog and user experience problem: agents should anticipate needs while remaining aligned with user preferences, trust, and cognitive load. We view this as a strong and necessary account of how proactivity should be shaped at the interaction level.

Our position does not reject this framing, but argues that it presupposes a critical condition—that the system’s understanding of the task and its relevant dimensions is already sufficiently well formed. The coupling perspective makes explicit that interaction-sensitive initiative alone is insufficient when uncertainty concerns what has not yet been articulated or recognized by either party.

A second alternative emphasizes *autonomy and action governance* as the primary axis of concern. Here, risks arise from unbounded execution, motivating controls over autonomy levels, reversibility, permissions, and oversight (98; 80; 21; 86). While we agree these mechanisms are essential, we argue they are incomplete: epistemic overreach can occur even at low autonomy, while high commitment can be justified when epistemic legitimacy is strong. The coupling model clarifies when autonomy and commitment are warranted by tying them explicitly to epistemic legitimacy, rather than treating autonomy as a scalar to be minimized or capped.

Across these views, proactivity is primarily regulated through interaction control or action governance. Our contribution is to surface a complementary constraint: proactive behavior must also be grounded in what the agent can legitimately claim to understand.

8. Conclusion

This paper advances a reframing of *generative proactivity*: not as acting earlier, more autonomously, or more persistently, but as acting *only when epistemically justified*. We show that many failures attributed to hallucination, misalignment, or unsafe autonomy arise from a deeper structural issue—a **mis-coupling between behavioral commitment and epistemic legitimacy**. Drawing on behavioral theories of proactivity and philosophical accounts of ignorance, we make this coupling explicit and show why regulating action alone is insufficient. This framework clarifies when proactive intervention is warranted, when it should remain exploratory, and when it constitutes epistemic overreach. It unifies diverse failure modes under a single explanatory lens and reframes proactivity as a practice of calibrated deviation rather than maximal initiative. Beyond diagnosis, the coupling perspective reorients the design space toward *epistemic partnership*. Rather than optimizing agents to close tasks quickly or act decisively, it foregrounds the role of agents in sustaining inquiry—surfacing latent gaps, preserving uncertainty, and calibrating restraint over time. This vision challenges current evaluation practices, training incentives, and architectural assumptions, suggesting that progress in proactive AI will depend more on disciplining commitment in the presence of incomplete understanding.

References

- [1] Gediminas Adomavicius, Bamshad Mobasher, Francesco Ricci, and Alexander Tuzhilin. Context-aware recommender systems. *AI Magazine*, 32(3):67–80, 2011.
- [2] Michael Ahn et al. Do as i can, not as i say: Grounding language in robotic affordances. In *Conference on Robot Learning (CoRL)*, 2022.
- [3] Seyed Amir Bahrainian, Hamed Zamani, and W. Bruce Croft. Predicting the topic of your next query for just-in-time information retrieval. In *Proceedings of the ACM Conference on Information and Knowledge Management (CIKM)*, 2016.
- [4] Ram Bairi et al. Evaluating the impact of ai pair programmers on developer productivity. In *International Conference on Software Engineering (ICSE)*, 2024.
- [5] Nicholas J. Belkin, Robert N. Oddy, and Helen M. Brooks. Anomalous states of knowledge as a basis for information retrieval. *The Canadian Journal of Information Science*, 5:133–143, 1980.
- [6] NJ Belkin and HM BROOKS. Journal of documentation. *Journal of Documentation*, 38(2):61–71, 1982.
- [7] Jan Benetka, Krisztian Balog, and Kjetil Nørkvåg. Anticipating information needs based on check-in activity. In *Proceedings of the 10th ACM International Conference on Web Search and Data Mining*, 2017.
- [8] Anmol Bhatia et al. Graph of thoughts: Solving elaborate problems with large language models. In *International Conference on Learning Representations (ICLR)*, 2024.
- [9] Mark C. Bolino, William H. Turnley, Jessica B. Gilstrap, and Margarita M. Suazo. Citizenship and counterproductive work behavior: A moral licensing view. *Journal of Applied Psychology*, 95(4):740–755, 2010.
- [10] Guo Chen, G. Shyam Sundar, et al. Llamapie: Proactive in-ear conversational assistants. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, 2025.
- [11] Mark Chen, Jerry Tworek, et al. Towards proactive conversational agents for programming assistance. In *Proceedings of the ACM CHI Conference on Human Factors in Computing Systems*, 2023.
- [12] Maximillian Chen, Ruoxi Sun, Tomas Pfister, and Sercan Ö Arık. Learning to clarify: Multi-turn conversations with action-based contrastive self-training. *arXiv preprint arXiv:2406.00222*, 2024.
- [13] Sherry Chen, Andrew Head, and Marti Hearst. Dango: A mixed-initiative data wrangling system using large language models. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2025.
- [14] J. Michael Crant. Proactive behavior in organizations. *Journal of Management*, 26(3):435–462, 2000.
- [15] Henk W De Regt. Understanding and scientific explanation. In *Scientific understanding: Philosophical perspectives*, pages 21–42. The University of Pittsburgh Press, 2009.
- [16] Yang Deng, Yifan Liao, and Minlie Huang. Knowledge-enhanced mixed-initiative dialogue system for emotional support conversations. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*, 2023.
- [17] Yang Deng, Yifan Liao, and Minlie Huang. Towards human-centered proactive conversational agents. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*, 2024.
- [18] Yang Deng, Yifan Liao, and Minlie Huang. Proactive conversational ai: A comprehensive survey of tasks, datasets, and methods. *ACM Computing Surveys*, 2025.
- [19] Yao Dou, Michel Galley, Baolin Peng, Chris Kedzie, and Wenhui Cai. Simulatorarena: Are user simulators reliable proxies for multi-turn evaluation of ai assistants? In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (ACL)*, 2025.
- [20] Danny Driess et al. Palm-e: An embodied multimodal language model. In *International Conference on Machine Learning (ICML)*, 2023.
- [21] KJ Feng, David W McDonald, and Amy X Zhang. Levels of autonomy for ai agents. *arXiv preprint arXiv:2506.12469*, 2025.
- [22] Luciano Floridi. *The logic of information: A theory of philosophy as conceptual design*. Oxford University Press, 2019.
- [23] Adam M. Grant and Sharon K. Parker. Two sides of the coin: The dual effects of proactive personality on performance and strain. *Journal of Applied Psychology*, 92(4):952–965, 2007.
- [24] Jiaqi Guo et al. Towards stable tool manipulation capability evaluation for large language models. In *Findings of the Association for Computational Linguistics (ACL)*, 2024.

- [25] Dana Harari and Ofra Amir. Proactive ai adoption can be threatening: When help backfires. *arXiv preprint arXiv:2509.09309*, 2025.
- [26] Jeffrey Heer. Agency and control in human–ai interaction. *ACM Queue*, 2021.
- [27] Dan Hendrycks et al. Scaling out-of-distribution detection for real-world settings. *International Conference on Machine Learning*, 2021.
- [28] Balázs Hidasi, Alexandros Karatzoglou, Linas Baltrunas, and Domonkos Tikk. Session-based recommendations with recurrent neural networks. In *International Conference on Learning Representations (ICLR)*, 2016.
- [29] Sirui Hong et al. Metagpt: Meta programming for multi-agent collaborative framework. In *International Conference on Learning Representations (ICLR)*, 2024.
- [30] Eric Horvitz. Principles of mixed-initiative user interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 159–166, 1999.
- [31] Eric Horvitz. Reflections on challenges and promises of mixed-initiative interaction. *AI Magazine*, 28(2):19–33, 2007.
- [32] Ziwei Ji, Nayeon Lee, Rita Frieske, Tian Yu, Dan Su, Yan Xu, Etsuko Ishii, Yejin Bang, Andrea Madotto, and Pascale Fung. Survey of hallucination in natural language generation. *ACM Computing Surveys*, 2023.
- [33] Carlos Jimenez et al. Swe-bench: Can language models resolve real-world github issues? In *International Conference on Learning Representations (ICLR)*, 2024.
- [34] Hyeonsu B. Kang, Andrew Head, and Marti Hearst. Synergi: A mixed-initiative system for scholarly synthesis and sensemaking. In *Proceedings of the ACM Symposium on User Interface Software and Technology (UIST)*, 2023.
- [35] Wang-Cheng Kang and Julian McAuley. Self-attentive sequential recommendation. In *Proceedings of the IEEE International Conference on Data Mining (ICDM)*, 2018.
- [36] Kirandeep Kaur, Vinayak Gupta, Aditya Gupta, and Chirag Shah. The proper approach to proactivity: Benchmarking and advancing knowledge gap navigation. *arXiv preprint arXiv:2601.09926*, 2026.
- [37] Ann Kerwin. Medical ignorance. *Philosophy of Science*, 60(3):371–384, 1993.
- [38] Matthias Kraus, Nils Stricker, and Rainer Stiefelhaugen. Explicit vs. implicit proactive dialogue policy learning. In *Proceedings of the International Conference on Multimodal Interaction (ICMI)*, 2020.
- [39] Matthias Kraus, Nils Stricker, and Rainer Stiefelhaugen. Modelling and predicting trust for developing proactive dialogue strategies in mixed-initiative interaction. *ACM Transactions on Interactive Intelligent Systems*, 11(3–4), 2021.
- [40] Carol C Kuhlthau. Inside the search process: Information seeking from the user’s perspective. *Journal of the American society for information science*, 42(5):361–371, 1991.
- [41] Duc-Trong Le, Hady W. Lauw, and Yuan Fang. Correlation-sensitive next-basket recommendation. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 2808–2814, 2019.
- [42] Wenqiang Lei, Zhenyu Zhang, Pengjie Ren, Jian Tang, and Maarten de Rijke Li. Estimation–action–reflection: Towards deep interaction between conversational and recommender systems. In *Proceedings of the ACM International Conference on Web Search and Data Mining (WSDM)*, 2020.
- [43] Guohao Li et al. Camel: Communicative agents for mind exploration of large language models. In *International Conference on Learning Representations (ICLR)*, 2024.
- [44] Jing Li, Pengjie Ren, Zhumin Chen, Zhaochun Ren, Tao Lian, and Jun Ma. Neural attentive session-based recommendation. In *Proceedings of the 26th ACM International Conference on Information and Knowledge Management (CIKM)*, 2017.
- [45] Q. Vera Liao, Matthew Davis, Werner Geyer, Michael Muller, and N. Sadat Shami. What can you do?: Studying social-agent orientation and agent proactive interactions with an agent for employees. In *Proceedings of the 2016 ACM Conference on Designing Interactive Systems (DIS ’16)*, 2016.
- [46] Henry Lieberman. Letizia: An agent that assists web browsing. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, 1995.
- [47] Daniel J. Liebling and Susan Dumais. Anticipatory search: Using context to initiate search. In *CHI Extended Abstracts on Human Factors in Computing Systems*, 2012.

- [48] Qiao Liu, Yifu Zeng, Rachael Mokhosi, and Haifeng Zhang. Stamp: Short-term attention/memory priority model for session-based recommendation. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD)*, pages 1831–1839, 2018.
- [49] Xiao Liu, Hao Yu, Hanchen Zhang, Yifan Xu, Xuanyu Lei, Hanyu Lai, Yu Gu, Hangliang Ding, Kaiwen Men, Kejuan Yang, Shudan Zhang, Xiang Deng, Aohan Zeng, Zhengxiao Du, Chenhui Zhang, Sheng Shen, Tianjun Zhang, Yu Su, Huan Sun, Minlie Huang, Yuxiao Dong, and Jie Tang. Agentbench: Evaluating LLMs as agents. In *The Twelfth International Conference on Learning Representations*, 2024.
- [50] Xingyu Bruce Liu, Shitao Fang, Weiyan Shi, Chien-Sheng Wu, Takeo Igarashi, and Xiang 'Anthony' Chen. Proactive conversational agents with inner thoughts. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, CHI '25, New York, NY, USA, 2025. Association for Computing Machinery.
- [51] Yaxi Lu, Shenzhi Yang, Cheng Qian, Guirong Chen, Qinyu Luo, Yesai Wu, Huadong Wang, Xin Cong, Zhong Zhang, Yankai Lin, Weiwen Liu, Yasheng Wang, Zhiyuan Liu, Fangming Liu, and Maosong Sun. Proactive agent: Shifting LLM agents from reactive responses to active assistance. In *The Thirteenth International Conference on Learning Representations*, 2025.
- [52] Teemu Luukkonen and Jaana Kekäläinen. Lstm-based predictions for proactive information retrieval. In *Proceedings of the International Conference on the Theory of Information Retrieval*, 2016.
- [53] Yosi Mass, Haggai Roitman, Ido Guy, and Naama Zwerdling. Conversational search with mixed-initiative: Asking good clarification questions backed-up by passage retrieval. In *Proceedings of the DialDoc Workshop*, 2022.
- [54] Yu Mei, Yuaxi Wang, Shiyi Wang, Qingyang Wan, Zhuojun Li, Chun Yu, Weinan Shi, and Yuanchun Shi. Interquest: A mixed-initiative framework for dynamic user interest modeling in conversational search. In *Proceedings of the 38th Annual ACM Symposium on User Interface Software and Technology*, UIST '25, New York, NY, USA, 2025. Association for Computing Machinery.
- [55] Xiangwu Meng, Yulu Du, Yujie Zhang, and Xiaofeng Han. A survey of context-aware recommender systems: From an evaluation perspective. *IEEE Transactions on Knowledge and Data Engineering*, 35(7):6575–6594, 2023.
- [56] Christian Meurisch, Cristina A. Mihale-Wilson, Adrian Hawlitschek, Florian Giger, Florian Müller, Oliver Hinz, and Max Mühlhäuser. Exploring user expectations of proactive ai systems. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies (PACM IMWUT)*, 4(4):Article 146, 2020.
- [57] Grégoire Mialon, Clémentine Fourier, Thomas Wolf, Yann LeCun, and Thomas Scialom. GAIA: a benchmark for general AI assistants. In *The Twelfth International Conference on Learning Representations*, 2024.
- [58] Florian Müller, Kuno Kurzhals, Michael Sedlmair, and Tamara Munzner. Cloudbits: Supporting conversations through augmented zero-query search visualization. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, 2017.
- [59] Reiichiro Nakano, Jacob Hilton, Suchir Balaji, Jeff Wu, Long Ouyang, Christina Kim, Christopher Hesse, Shantanu Jain, Vineet Kosaraju, William Saunders, et al. Webgpt: Browser-assisted question-answering with human feedback. *arXiv preprint arXiv:2112.09332*, 2021.
- [60] Florian Nothdurft, Stefan Ultes, and Wolfgang Minker. Finding appropriate interaction strategies for proactive dialogue systems. In *Proceedings of the SIGDIAL Conference on Discourse and Dialogue*, 2015.
- [61] Jeeseun Oh, Wooseok Kim, Sungbae Kim, Hyeonjeong Im, and Sangsu Lee. Better to ask than assume: Proactive voice assistants' communication strategies that respect user agency in a smart home environment. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI '24)*, 2024.
- [62] Cassandra Overney, Daniel T Kessler, Suyash Pradeep Fulay, Mahmood Jasim, and Deb Roy. Coalesce: An accessible mixed-initiative system for designing community-centric questionnaires. In *Proceedings of the 30th International Conference on Intelligent User Interfaces*, IUI '25, page 366–389, New York, NY, USA, 2025. Association for Computing Machinery.
- [63] Raja Parasuraman, Thomas Sheridan, and Christopher Wickens. A model for types and levels of human interaction with automation. *IEEE Transactions on Systems, Man, and Cybernetics*, 2000.

- [64] Joon Sung Park, Joseph O'Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and Michael S Bernstein. Generative agents: Interactive simulacra of human behavior. In *Proceedings of the 36th annual acm symposium on user interface software and technology*, pages 1–22, 2023.
- [65] Sharon K. Parker, Uta K. Bindl, and Karoline Strauss. Making things happen: A model of proactive motivation. *Journal of Management*, 36:827 – 856, 2010.
- [66] Sharon K. Parker, Helen M. Williams, and Nick Turner. Modeling the antecedents of proactive behavior at work. *Journal of Applied Psychology*, 91(3):636–652, 2006.
- [67] Gil Pasternak, Dheeraj Rajagopal, Julia White, Dhruv Atreja, Matthew Thomas, George Hurn-Maloney, and Ash Lewis. Beyond reactivity: Measuring proactive problem solving in llm agents. *arXiv preprint arXiv:2510.19771*, 2025.
- [68] Milan Patil et al. Gorilla: Large language model connected with massive apis. In *International Conference on Machine Learning (ICML)*, 2023.
- [69] Yujia Qin et al. Toollm: Facilitating large language models to master 1,000+ real-world apis. In *International Conference on Learning Representations (ICLR)*, 2024.
- [70] Marissa Radensky, Andrew Head, and Marti Hearst. Mixed-initiative methods for co-creation in scientific research. *ACM Transactions on Computer-Human Interaction*, 2024.
- [71] Hossein Rahmani, Mohammad Aliannejadi, and Fabio Crestani. Clarifying the path to user satisfaction: An investigation of clarification in conversational search. In *Findings of the European Chapter of the Association for Computational Linguistics (EACL)*, 2024.
- [72] Bradley J. Rhodes. *Just-In-Time Information Retrieval*. PhD thesis, Massachusetts Institute of Technology, 2000.
- [73] Bradley J. Rhodes and Pattie Maes. Just-in-time information retrieval agents. *IBM Systems Journal*, 39(3–4):685–704, 2000.
- [74] Georgios Samarinas and Hamed Zamani. Procis: A benchmark for proactive retrieval in conversations. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2024.
- [75] Timo Schick, Jane Dwivedi-Yu, Roberto Dessí, Roberta Raileanu, Maria Lomeli, Eric Hambro, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. Toolformer: language models can teach themselves to use tools. In *Proceedings of the 37th International Conference on Neural Information Processing Systems, NIPS '23*, Red Hook, NY, USA, 2023. Curran Associates Inc.
- [76] Ivan Sekulić, Mohammad Aliannejadi, and Fabio Crestani. Evaluating mixed-initiative conversational search systems via user simulation. In *Proceedings of the ACM International Conference on Web Search and Data Mining (WSDM)*, 2022.
- [77] Ivan Sekulić, Shuo Zhang, and W. Bruce Croft. Learning steerable clarification policies with collaborative self-play. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2022.
- [78] Procheta Sen. *Proactive Information Retrieval*. PhD thesis, Dublin City University, 2021.
- [79] Shreya Shankar, Jeffrey Wu, and Christopher Ré. Who validates the validators? aligning llm-assisted evaluation with human judgment. In *Proceedings of the ACM Symposium on User Interface Software and Technology (UIST)*, 2024.
- [80] Noah Shinn, Becky Labash, et al. Reflexion: Language agents with verbal reinforcement learning. In *Advances in Neural Information Processing Systems*, 2023.
- [81] Milad Shokouhi and Hui Guo. From queries to cards: Re-ranking proactive card recommendations based on reactive search history. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2015.
- [82] Yang Song, Xiaolin Shi, Haibo Fu, and Jing Li. Query-less: Predicting task repetition for nextgen proactive search and recommendation engines. In *Proceedings of the 25th International World Wide Web Conference*, 2016.
- [83] Fei Sun, Jun Liu, Jian Wu, Changhua Pei, Xiao Lin, Wenwu Ou, and Peng Jiang. Bert4rec: Sequential recommendation with bidirectional encoder representations from transformer. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management (CIKM)*, pages 1441–1450, 2019.
- [84] Ante Wang, Yujie Lin, Jingyao Liu, Suhang Wu, Hao Liu, Xinyan Xiao, and Jinsong Su. Beyond passive critical thinking: Fostering proactive questioning

- to enhance human-ai collaboration. *arXiv preprint arXiv:2507.23407*, 2025.
- [85] Guanzhi Wang et al. Voyager: An open-ended embodied agent with large language models. In *International Conference on Learning Representations (ICLR)*, 2024.
- [86] World Economic Forum. Ai agents in action: Foundations for evaluation and governance, 2025. Policy report.
- [87] Chuhan Wu, Mohammad Aliannejadi, and Fabio Crestani. Inscit: Information-seeking conversations with mixed-initiative interactions. *Transactions of the Association for Computational Linguistics*, 11:1–19, 2023.
- [88] Qingyun Wu et al. Autogen: Enabling next-gen llm applications via multi-agent conversation. In *International Conference on Learning Representations (ICLR)*, 2024.
- [89] Shirley Wu, Michel Galley, Baolin Peng, Hao Cheng, Gavin Li, Yao Dou, Weixin Cai, James Zou, Jure Leskovec, and Jianfeng Gao. CollabLLM: From passive responders to active collaborators. In *Forty-second International Conference on Machine Learning*, 2025.
- [90] Shu Wu, Yuyuan Tang, Yanqiao Zhu, Liang Wang, Xing Xie, and Tieniu Tan. Session-based recommendation with graph neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2019.
- [91] Yu Wu, Wei Wu, Chen Xing, and Ming Zhou. Duconv: A dataset for multi-turn dialogue reasoning. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*, 2019.
- [92] Zixiang Xie et al. Osworld: Benchmarking multi-modal agents for open-ended tasks in real computer environments. In *Advances in Neural Information Processing Systems (NeurIPS), Datasets and Benchmarks*, 2024.
- [93] An Yan, Shuo Cheng, Wang-Cheng Kang, Mengting Wan, and Julian McAuley. Cosrec: 2d convolutional neural networks for sequential recommendation. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management (CIKM)*, 2019.
- [94] John Yang et al. Swe-agent: Agent-computer interfaces enable automated software engineering. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024.
- [95] Liu Yang, Hui Guo, and Milad Shokouhi. Modeling user interest for zero-query ranking. In *Proceedings of the European Conference on Information Retrieval (ECIR)*, 2016.
- [96] Shunyu Yao, Howard Chen, John Yang, and Karthik R Narasimhan. Webshop: Towards scalable real-world web interaction with grounded language agents. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, *Advances in Neural Information Processing Systems*, 2022.
- [97] Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L. Griffiths, Yuan Cao, and Karthik R Narasimhan. Tree of thoughts: Deliberate problem solving with large language models. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- [98] Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. React: Synergizing reasoning and acting in language models. In *International Conference on Learning Representations (ICLR)*, 2023.
- [99] Ziyi Ye, Hyeonsu B. Kang, and Marti Hearst. Scholarmate: A mixed-initiative tool for qualitative knowledge work and information sensemaking. *ACM Transactions on Computer-Human Interaction*, 2025.
- [100] Feng Yu, Qiang Liu, Shu Wu, Liang Wang, and Tieniu Tan. A dynamic recurrent model for next basket recommendation. In *Proceedings of the 39th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR)*, pages 729–732, 2016.
- [101] Xiao Yu, Baolin Peng, Ruize Xu, Michel Galley, Hao Cheng, Suman Nath, Jianfeng Gao, and Zhou Yu. Dyna-think: Synergizing reasoning, acting, and world model simulation in AI agents, 2025.
- [102] Fajie Yuan, Alexandros Karatzoglou, Ioannis Arapakis, Joemon M. Jose, and Xiangnan He. A simple convolutional generative network for next item recommendation. In *Proceedings of the 12th ACM International Conference on Web Search and Data Mining (WSDM)*, 2019.
- [103] Yifei Yuan, Clemencia Siro, Mohammad Aliannejadi, Maarten de Rijke, and Wai Lam. Asking multimodal clarifying questions in mixed-initiative conversational search. In *Proceedings of the ACM Web Conference 2024, WWW '24*, page 1474–1485, New York, NY, USA, 2024. Association for Computing Machinery.

- [104] Shuyan Zhou, Frank F. Xu, Hao Zhu, Xuhui Zhou, Robert Lo, Abishek Sridhar, Xianyi Cheng, Tianyue Ou, Yonatan Bisk, Daniel Fried, Uri Alon, and Graham Neubig. Webarena: A realistic web environment for building autonomous agents. In *The Twelfth International Conference on Learning Representations*, 2024.
- [105] Brianna Zitkovich, Tianhe Yu, Sichun Xu, Peng Xu, Ted Xiao, Fei Xia, Jialin Wu, Paul Wohlhart, Stefan Welker, Ayzaan Wahid, et al. Rt-2: Vision-language-action models transfer web knowledge to robotic control. In *Conference on Robot Learning*, pages 2165–2183. PMLR, 2023.




Figure 6. Anticipation as an act-ahead pipeline.

A. Prevailing Approaches to Proactivity

As limitations of purely reactive interaction have become increasingly apparent, *proactivity* has emerged as a central design goal in contemporary intelligent systems. This section examines how proactivity is predominantly implemented in recent advances. We further show that epistemic uncertainty is largely externalized rather than represented, with proactivity reduced to goal-directed action selection under unexamined assumptions about what the task is and what progress entails.

A.1. Anticipatory Proactivity

Initiative is often realized through *act-ahead assistance*: systems infer forthcoming needs from observable context or behavior and intervene prior to explicit user requests, following the pipeline shown in Figure 6.

Within information retrieval, this lineage traces back to early browsing and just-in-time retrieval agents that monitored user activity and surfaced resources without explicit queries (46; 73; 72; 47). As search and mobile platforms matured, anticipation increasingly became a *zero-query ranking* problem: systems learn to select and order proactive suggestions or cards from reactive history and situational context (81; 95; 82; 7; 58; 78). Recent work continues to refine this operationalization by predicting next-query topics or retrieving context-relevant material during writing and other ongoing tasks (3; 52).

A parallel thread in recommendation systems instantiates anticipation as *trajectory-based preference prediction*. Context-aware recommenders treat situational signals as a proxy for latent intent (1; 55), while session-based and sequential models forecast next actions/items from interaction traces (28; 44; 48; 90; 35; 83; 102; 93). Next-basket variants similarly anticipate future consumption bundles by extrapolating from prior baskets and correlations (100; 41).

In language-based assistants and proactive dialogue, anticipation is realized as inferring likely follow-ups, steering trajectories toward targets, or deciding opportune moments to contribute (60; 91; 38; 18). LLM-based assistants often implement this implicitly via unsolicited but contextually plausible completions (59), and explicitly in productivity and programming settings where systems monitor workspace state and propose edits or actions (11; 4). More recent conversational agents foreground timing and initiative selection

under continuous context streams (50; 10; 17).

Despite strong performance in routine or well-structured settings, anticipatory proactivity remains fundamentally *extrapolative*. Proactive interventions are constrained to what can be inferred from prior signals and expressed within a predefined space of candidate actions, items, or documents. As a result, relevant goals and dimensions must already be representable by the system, limiting anticipation precisely when user needs involve unarticulated uncertainty or unknown unknowns. Benchmarks for proactive conversational retrieval make this operationalization explicit—monitor context, retrieve, and intervene—while leaving the underlying representational bound unchanged (74).

A.2. Autonomous and Planning-Based Proactivity

Another distinct line of work conceptualizes proactivity not as anticipation of likely needs, but as *autonomous goal pursuit*. In these approaches, systems take initiative by formulating plans, decomposing objectives, and executing sequences of actions without requiring continuous user prompts. Proactivity is thus realized through commitment to internally maintained goals and the capacity to act over extended horizons, often in dynamic or partially observable environments.

Early formulations of this paradigm emphasize the tight coupling between reasoning and action. Planner–actor agents interleave deliberation with execution, allowing models to revise plans based on intermediate outcomes and environmental feedback. Representative examples include agents that explicitly reason about action sequences and tool use during execution, such as ReAct and its extensions, which frame autonomy as an ongoing loop of planning, acting, and observing outcomes (98; 80). Subsequent work explores more structured forms of planning, including tree- and graph-based deliberation mechanisms that expand and evaluate alternative action trajectories prior to commitment (98; 97; 8).

A major thrust of recent research focuses on *tool-using agents* that plan over external APIs, functions, or software interfaces. These systems treat tools as action primitives and learn to select, sequence, and parameterize tool calls in service of a broader objective (75; 68; 69; 24). Benchmarks such as ToolBench and StableToolBench formalize this setting, evaluating agents on their ability to autonomously compose tools to complete complex tasks rather than merely predicting the next response token.

Autonomous proactivity is further instantiated in agents operating in realistic web and computer environments. Rather than responding to isolated queries, these agents must navigate interfaces, maintain state, and pursue goals across long interaction sequences. Work in this space includes

Axis of initiative	Available choices			
	Option 1	Option 2	Option 3	Option 4
Who acts	Human	System	Negotiated	–
When to act	Timing	Interruption	Turn-taking	–
How much to act	Suggest	Clarify	Execute	Defer

Table 2. Core axes along which mixed-initiative systems regulate proactive behavior. Each axis defines a discrete choice set governing when and how initiative is exercised.

WebShop and WebArena for web-based task completion, as well as OSWorld for operating-system–level interaction, all of which frame proactivity as sustained action under partial observability (96; 104; 92). Evaluation suites such as AgentBench and GAIA extend this perspective by assessing general-purpose autonomy across heterogeneous tasks and environments (49; 57).

Planning-based proactivity is especially prominent in software engineering agents, where systems are tasked with diagnosing bugs, navigating repositories, and producing executable patches. Benchmarks such as SWE-bench and agents such as SWE-agent formalize this setting, emphasizing long-horizon reasoning, tool-mediated execution, and iterative refinement (33; 94). Similar planning dynamics appear in multi-agent systems, where autonomy is distributed across interacting agents that coordinate roles, exchange intermediate results, and collectively pursue shared objectives (88; 43; 29).

Finally, embodied and simulated agents extend autonomous proactivity into physical and virtual worlds, where planning must account for spatial dynamics, affordances, and delayed consequences. Systems such as SayCan, PaLM-E, RT-2, and Voyager demonstrate how language-conditioned planning can support long-horizon action in robotics and open-ended environments (2; 20; 105; 85).

Across these settings, autonomous proactivity shifts the locus of initiative from prediction to commitment. Rather than inferring what a user might want next, autonomous agents decide *what to do* and *how to proceed*, introducing new forms of behavioral risk tied to goal persistence, irreversibility, and misaligned objectives. These properties distinguish planning-based proactivity from anticipatory approaches and motivate the need for principled constraints on commitment and intervention.

A.3. Mixed-Initiative Proactivity

Mixed-initiative proactivity treats *initiative itself* as the primary control variable. The motivating premise is that nei-

ther purely anticipatory assistance (which extrapolates from observable signals) nor fully autonomous agents (which commit to internally maintained plans) can reliably preserve user agency and coordination under uncertainty. Instead, mixed-initiative systems frame proactivity as an interactional regulation problem: the system must continuously decide *who* should act, *when* to act, and *how strongly* to intervene, given evolving evidence about user state, task structure, and risk.

This paradigm is rooted in foundational HCI accounts that argue initiative must be allocated dynamically to balance efficiency against disruption and loss of control (30; 31). Once initiative is treated as a regulatable quantity rather than a byproduct of prediction or autonomy, a causal chain follows. (i) Interaction unfolds under partial observability of user intent and constraints. (ii) Proactive contributions therefore introduce coordination risk: mistimed or over-strong interventions can derail the user, while over-deference can stall progress. (iii) Systems must operationalize initiative via explicit decision points over action ownership, timing, and strength. (iv) These decisions require evidence beyond task content alone—signals about uncertainty, trust, and interaction state—and induce evaluation criteria that include disruption, calibration, and perceived agency, not only task success.

Work in proactive dialogue makes this control problem concrete by defining strategies that choose between contribution types (e.g., clarify vs. suggest vs. defer) and between explicit vs. implicit initiative, often conditioned on user trust and uncertainty (16; 39; 12). Here, the key mechanism is not generating the next utterance per se, but selecting the appropriate *interaction move* and its timing so that proactive assistance remains coordinated rather than coercive.

A parallel operationalization appears in mixed-initiative conversational search, where systems must decide whether to retrieve immediately or elicit information through clarification, and how to steer the search process without overtaking it. This line of work formalizes mixed initiative through user simulation and evaluation protocols that expose the tradeoff between intervention and disruption, and through tasks that explicitly integrate clarification question generation/selection into retrieval pipelines (76; 53; 87; 103; 71). The resulting causal structure mirrors the paradigm: uncertainty about intent → choice of initiative (clarify vs. retrieve) → effects on satisfaction and efficiency, where failure is often attributable to mis-timing, poor calibration, or mismatched control allocation rather than retrieval quality alone (71; 76).

Conversational recommender systems instantiate mixed initiative as a *deep interaction loop* over preference elicitation and action selection: systems alternate between estimating user state, selecting an intervention (ask/recommend/refine),

and reflecting on feedback to regulate subsequent initiative (42). This again emphasizes that proactive behavior is not merely producing recommendations, but managing the conversational control dynamics that make recommendation actionable and acceptable.

Beyond dialogue-centric settings, mixed-initiative proactivity increasingly appears in knowledge-work tools, where the system's role is to propose structure, partial drafts, or transformations while keeping the user in control of direction and commitment. Systems for scholarly synthesis and qualitative sensemaking explicitly design for human–AI coordination, using mixed-initiative interfaces to surface candidate claims, reorganizations, or summaries that the user can adopt, reject, or revise (34; 99). Similar design commitments drive mixed-initiative workflows in data wrangling, where proactive transformations can be powerful but require calibrated intervention to avoid silently imposing assumptions (13). Accessibility- and interaction-focused systems further foreground that initiative must be regulated to match users' abilities and preferences, treating control allocation as a first-class design objective rather than an afterthought (62; 54; 70).

Finally, mixed-initiative principles are increasingly leveraged in alignment and oversight workflows. When AI systems participate in evaluation, verification, or moderation-like tasks, naive autonomy can amplify errors or lock in premature judgments. Mixed-initiative designs instead distribute responsibility across human and system, structuring validation as a regulated interaction in which the system proposes and the human adjudicates, with explicit attention to who holds authority at each step (79).

Taken together, these threads converge on a position-paper-critical claim: mixed-initiative proactivity is best understood as *initiative regulation under uncertainty*. The core advance is not a particular model family but a shift in what is optimized: from task progress alone to progress *subject to calibrated control allocation*. This makes the paradigm a natural bridge between anticipatory and autonomous proactivity: it inherits the need to infer from context, but refuses to equate inference with entitlement to act; it benefits from tool- and plan-capable systems, but constrains commitment through interactional mechanisms that preserve agency, timing, and reversibility.

A.4. Discussion

Figure 2 reveals a common design move that cuts across anticipatory, mixed-initiative, and autonomous approaches. Although these paradigms differ in how initiative is allocated—via prediction, regulation, or commitment—they all localize proactivity at the level of *action choice*. Initiative is exercised by deciding *which action to take, when to take it, or how strongly to commit*, given an assumed set of goals,

dimensions, and candidate interventions. The task frame itself remains invariant: what counts as progress, what alternatives are relevant, and which risks matter are treated as pre-specified rather than subject to intervention.

This shared action-centric framing explains both the successes and the systematic blind spots of prevailing approaches. When goals are stable and task structure is well defined, reallocating initiative—earlier prediction, stronger commitment, or finer-grained regulation—can meaningfully improve efficiency and coordination. However, when uncertainty concerns the task itself—what the user is trying to achieve, which considerations are missing, or how the problem should be framed—these approaches have no place to act. As Figure 2 makes clear, failure modes differ across paradigms, but they all arise downstream of the same assumption: that proactivity begins only once the problem is already specified. Epistemic uncertainty is therefore not addressed but bypassed, motivating the need for a form of proactivity that operates *before* action selection, by intervening in how tasks, goals, and unknowns are surfaced and structured. We turn to this question next.

B. Consequences of Epistemic–Behavioral Coupling

B.1. The Missing Control Variable: Commitment, not Autonomy

A direct consequence of epistemic–behavioral coupling is that *autonomy is the wrong control variable for regulating proactive behavior*. Much recent work frames progress in proactive agents in terms of increasing autonomy: agents initiate goals, plan multi-step actions, invoke tools, and act without user prompts (98; 80). While autonomy determines *who* acts and *when*, it is largely silent on *how strongly* an agent commits to its actions and the degree to which those actions shape future states.

The coupling analysis shows that the primary source of harm is not autonomous action per se, but excessive *commitment* under insufficient epistemic legitimacy. Commitment captures the extent to which an action is consequential, irreversible, or forecloses alternative trajectories. Observing, suggesting, probing, acting reversibly, and acting irreversibly all differ minimally in autonomy, yet differ substantially in epistemic risk. Treating these behaviors as equivalent forms of “initiative” obscures the mechanisms by which mis-coupling arises.

Existing agent frameworks often regulate autonomy through permissioning or tool access, but leave commitment implicit and unmanaged. As a result, agents may act confidently and decisively even as epistemic legitimacy degrades, because nothing in the control structure forces a downshift in commitment. From the coupling perspective, this is a

structural oversight: regulating autonomy without regulating commitment allows epistemic overreach to persist even in well-aligned systems.

Recognizing commitment as the primary control variable reframes proactive behavior as a matter of *calibrated intervention*. Autonomy determines whether an agent may act; commitment determines whether it *should*. Once epistemic–behavioral coupling is taken seriously, commitment—not autonomy—becomes the critical quantity that must be modulated to prevent systematic misalignment between knowing and acting.

B.2. The Hidden Training Incentive: Momentum Rewards Mis-coupling

A second consequence of epistemic–behavioral coupling concerns the optimization pressures under which proactive agents are trained and evaluated. Across contemporary agentic systems, learning objectives and benchmarks predominantly reward task completion, coherence of action sequences, speed of resolution, and confident execution (98; 80). These criteria implicitly favor behavioral momentum: once an agent initiates action, continued commitment is treated as progress, while hesitation or downshifting is rarely rewarded.

From the coupling perspective, this creates a systematic bias toward mis-coupling. Because epistemic legitimacy is weakly represented or entirely absent from training signals, agents are incentivized to maintain or escalate commitment even as epistemic conditions deteriorate. Empirical work has shown that model confidence and fluency often increase precisely when systems generalize beyond their training distribution, masking epistemic fragility rather than exposing it (27; 32). Under such conditions, early commitment is not penalized; instead, it is reinforced by success metrics that register only final outcomes.

This incentive structure helps explain why epistemic overreach is a persistent and predictable failure mode rather than an anomaly. When optimization rewards uninterrupted progress, agents learn to suppress uncertainty, smooth over anomalies, and resolve ambiguity through decisive action. The resulting behavior is locally optimal under prevailing objectives, yet globally brittle with respect to epistemic legitimacy. In coupling terms, training pressures systematically privilege the commitment axis while leaving epistemic legitimacy underconstrained.

Crucially, this dynamic does not depend on model scale, data quality, or architectural choice. It follows directly from objectives that equate success with momentum. As long as proactive systems are trained in environments where continued action is rewarded more reliably than justified restraint, mis-coupling between knowing and acting will

remain the norm rather than the exception.

B.3. A Research Agenda in Five Questions

Accepting epistemic-behavioral coupling reframes progress in proactive AI as a set of open research questions rather than a search for immediate solutions. We highlight five questions that follow directly from the coupling framework and delineate a forward-looking agenda.

(Q1) How should epistemic legitimacy be represented?

What internal signals or abstractions allow an agent to distinguish between recognized uncertainty, unrecognized gaps, and error-as-knowledge, without collapsing these states into a single confidence score?

(Q2) What epistemic signals must proactive action preserve?

Which forms of uncertainty, disagreement, or anomaly are most critical to retain during action, and how can agents avoid interventions that erase the very evidence needed to detect misalignment?

(Q3) How can agents detect epistemic degradation in time to act on it?

What early indicators reliably signal that epistemic legitimacy is deteriorating—due to novelty, distributional shift, or internal inconsistency—before high-commitment actions are taken?

(Q4) When is restraint or abstention the correct proactive behavior?

How should agents decide to delay, downshift commitment, or defer action altogether, and how can such behavior be distinguished from indecision or failure in evaluation settings?

(Q5) How should coupling quality be evaluated?

What evaluation protocols can assess whether commitment was justified at the time of action, rather than inferring quality solely from final outcomes?

Together, these questions articulate a research agenda centered on aligning knowing and acting, rather than scaling autonomy alone. Addressing them requires rethinking how proactive agents are represented, trained, and evaluated—without presupposing any single architectural approach.

C. Extended Vision for Epistemic Partnership

This appendix elaborates on three forward-looking capabilities that follow naturally from the epistemic-behavioral coupling framework. These capabilities are not presented as concrete system designs, but as conceptual directions that clarify what it would mean for proactive agents to function as epistemic partners rather than as task executors or interaction optimizers.

C.1. Asking Questions About Unknown Unknowns

Most existing proactive systems treat questioning as a mechanism for resolving *recognized uncertainty*: filling missing slots, disambiguating intent, or clarifying preferences. In contrast, epistemic partnership requires agents to engage with *unknown unknowns*—gaps that are not yet represented as questions by either the user or the system. These include missing dimensions, suppressed assumptions, unexamined boundary conditions, or unconsidered alternatives that structure the problem space itself.

This mode of inquiry closely mirrors how progress occurs in scientific discovery and exploratory research. Breakthroughs rarely emerge from efficiently resolving well-posed questions; instead, they arise when researchers recognize that a problem has been framed too narrowly, that a key variable has been taken for granted, or that an alternative explanatory lens has not yet been articulated. In such contexts, the most consequential interventions are not answers, but questions that reconfigure what counts as relevant, plausible, or even askable. Epistemic partners that can surface these latent uncertainties have the potential to support discovery not by accelerating inference, but by reshaping the space of inquiry itself.

From the perspective of epistemic-behavioral coupling, asking questions about unknown unknowns occupies a distinctive region of the joint space: epistemic legitimacy is low by definition, since neither the user nor the system can justify the question through existing evidence alone, yet behavioral commitment must remain deliberately constrained. The value of such questioning lies not in correctness or actionability, but in *opening* the inquiry space—making implicit assumptions visible without prematurely stabilizing interpretation or direction. This distinguishes it from both clarification and suggestion: it intervenes at the level of problem formulation rather than problem solving.

Discussion. Reframing proactive questioning as an epistemic act highlights a central challenge for generative agents in research-facing and discovery-oriented settings. The goal is not to optimize questions for efficiency or task completion, but to recognize when the absence of a question itself signals epistemic incompleteness. Within our framework, responsible engagement with unknown unknowns requires strict limits on commitment, ensuring that such questions function as invitations to exploration rather than instruments of guidance under fragile understanding. When properly constrained, this capability allows proactive agents to participate in inquiry without collapsing uncertainty too early—supporting discovery by keeping alternative explanations, dimensions, and futures in play.

C.2. Long-Horizon Epistemic Thinking

A second implication of epistemic partnership is the need for agents to reason beyond short-term interaction horizons. Most generative agents are optimized for immediate task completion, local coherence, or near-term utility. Epistemic partners, by contrast, must reason over extended horizons in which goals evolve, consequences unfold slowly, and the agent's own understanding—and alignment with the user's interests—may drift or degrade over time.

One dimension of long-horizon epistemic thinking concerns agents that operate with *dual temporal capacities*. Such agents must be able to provide effective short-term assistance while simultaneously reasoning about longer-term user trajectories: how current interventions shape future goals, dependencies, expectations, and modes of reliance. In many domains—learning, creative work, research, or planning—helpful short-term actions can undermine longer-term outcomes by narrowing exploration, stabilizing premature interpretations, or optimizing for progress along a locally salient but globally suboptimal path. Epistemic partners must therefore reason not only about what helps *now*, but about how present actions pave or foreclose future epistemic possibilities for the user.

A second implication concerns long-horizon thinking that is not solely user-directed. Human epistemic agency often involves reflection, exploration, and self-directed inquiry that exceeds the demands of any single interaction or external request. Proactive agents that aspire to epistemic partnership may similarly need the capacity to reason for themselves: to monitor their own uncertainty, detect epistemic drift, revisit prior assumptions, and explore alternative representations or strategies beyond immediate task pressure. This raises foundational questions about what it would mean for artificial agents to engage in ongoing epistemic work—maintaining and revising internal models, updating memories and abstractions, and identifying gaps in their own understanding over time.

Within the epistemic-behavioral joint space, long-horizon thinking is best understood as the dynamics of commitment accumulation. Even actions that are individually reversible can, in aggregate, produce strong path dependence that constrains future inquiry—locking in assumptions, suppressing alternatives, or privileging certain interpretations through repeated reinforcement. Epistemic partnership therefore requires agents to reason not only about what they currently know, but about how present commitments shape the future epistemic landscape for both the user and the system itself.

Discussion. Long-horizon epistemic thinking foregrounds a failure mode that is often overlooked: epistemic foreclosure through incremental commitment. Our framework suggests that responsible proactivity must account for the

temporal dynamics of knowing, not merely the correctness or utility of individual steps. Treating epistemic legitimacy as something that evolves—and can deteriorate—over time reframes proactivity as an ongoing process of calibration rather than monotonic escalation. By explicitly representing commitment and epistemic fragility across horizons, the epistemic-behavioral framework offers a principled basis for designing agents that can support users' long-term trajectories while also sustaining their own capacity for reflection, revision, and discovery.

C.3. Test-Time Proactivity as Epistemic Regulation

A third, closely related capability concerns *test-time proactivity*. Most proactive behaviors are implicitly learned at training time and executed at deployment as fixed policies. Epistemic partnership instead demands that agents actively regulate their initiative *during interaction*, adapting commitment in response to real-time signals of epistemic adequacy, novelty, or mismatch.

Within the epistemic-behavioral coupling, test-time proactivity is the mechanism that keeps the agent within the joint space. Rather than treating uncertainty estimates or confidence scores as passive annotations, epistemic partners must use them to modulate behavior: seeking information, downshifting commitment, or reverting to exploratory modes when legitimacy weakens. This form of proactivity is not about acting more often, but about knowing when *not* to act—and when to re-enter inquiry instead of execution.

Discussion. Test-time proactivity reframes deployment as an epistemic process rather than a purely behavioral one. The key insight is that no amount of training-time optimization can anticipate all epistemic contingencies. By grounding behavior in real-time epistemic regulation, agents can remain responsive to uncertainty as it arises, avoiding the systematic mis-couplings that occur when commitment continues unchecked despite deteriorating understanding.

Together, these three directions—asking questions about unknown unknowns, long-horizon epistemic thinking, and test-time proactivity—extend the epistemic-behavioral coupling from a diagnostic framework into a generative research agenda. They clarify what epistemic partnership demands in practice: not stronger autonomy, richer interaction, or deeper reasoning alone, but disciplined control over how knowing and acting co-evolve. While realizing these capabilities remains an open challenge, the coupling framework provides a principled foundation for reasoning about their necessity and their limits.