AMBIENT POINT CORP.

QUANTITATIVE AI · ADVISOR · BUILDER · OPERATOR

—— POSITION PAPER

Coordination Is the *Bottleneck*

A position paper on human–AI collaboration and decision coordination

AUTHOR	Ariel Rodriguez, Founder & CEO
PUBLISHED BY	Ambient Point Corp.
DATE	October 2026
VERSION	1.0

© 2026 Ambient Point Corp. All rights reserved.

SECTION 01

Position in brief

The value of AI in knowledge work is now limited less by what models can do than by how well people, AI systems and decisions are coordinated. Organizations that treat AI as plug-in automation will get more output and less value. Those that design the coordination layer on purpose, and invest in the human judgment it depends on, will capture the gains.

Coordination operates on two levels. **Human–AI coordination** is how people and AI systems share understanding, divide work and repair mistakes. **Decision coordination** is how AI-driven predictions and automation are carried through to consequential choices: with stated objectives and constraints, expert judgment made explicit and quantified, and a clear owner for every decision. Capability has outgrown both. Models now forecast, score and value at a scale and cost that were out of reach a few years ago, and agentic systems increasingly act on what they predict. Yet in most organizations the decisions those outputs should inform are still made the old way. **Prediction and automation have outgrown decision-making.** Closing that gap calls for quantitative solutions: models that encode how experts decide, measured by the quality of the decisions they support rather than the accuracy of their forecasts.

This paper draws on seven works that approach the problem from different levels. Klein et al. set out what any joint human–machine activity requires. Vishwarupe et al. examine grounding and repair in LLM work, and Pavlović offers constructivist principles for collaboration. Rousse and Dreyfus explain how expertise forms, and Rousse and Scapini show how AI is reshaping it. Two empirical studies test the claims: a meta-analysis of 106 human–AI experiments (Vaccaro et al.) and a field experiment with nearly 800 professionals at Procter & Gamble (Dell'Acqua et al.). Research in decision science and operations research grounds the quantitative argument (Agrawal, Gans & Goldfarb; Bertsimas & Kallus; Elmachtoub & Grigas), and research on expert judgment shows why encoding it pays (Grove et al.; Kahneman, Sibony & Sunstein). Practitioner experience, including a white paper from the Human Machine Collaboration Institute (HMCI), illustrates how the problems show up inside teams.

We argue for seven principles:

- **Treat coordination as infrastructure**, designed and owned at the level of decisions as well as conversations, not left to individual improvisation.
- **Quantify the decision**, encoding expert judgment into models built and measured for the decisions they serve.
- **Design for grounding and shared repair**, so the human is not the only one catching misunderstandings.
- **Gate AI work at the points where judgment matters**, with explicit human sign-off.
- **Keep a human answerable** for every AI-assisted output and decision that leaves the team.

- **Match structure to the task**: tight scope for routine work, open dialogue for complex problems.
- **Rebuild the path to expertise**, because tomorrow's orchestrators are today's novices.

The stakes are practical. Without these principles, teams produce polished work that is subtly wrong, forecasts and dashboards accumulate without changing any decision, senior staff become quality bottlenecks, and the entry-level roles that build judgment quietly disappear.

SECTION 02

The problem: capability has outrun coordination

AI systems have moved from answering prompts to acting. They plan, keep memory, call tools and produce finished work products across multi-step tasks. The line between AI work and human work blurs into a single workflow, and someone has to manage it.

Two decades ago, Klein et al. identified what that management requires. Any joint activity between people and automated agents depends on four things: an agreement to work together, mutual predictability, mutual influence and maintained common ground. They warned that more capable, more autonomous agents can become harder to predict and direct. Today's AI makes that warning concrete.

The same gap appears one level up, between what AI predicts and what organizations decide. Agrawal, Gans and Goldfarb describe AI as a sharp drop in the cost of prediction. Prediction, though, is only one input to a decision; the other is judgment about which outcomes matter and what each kind of error costs. As prediction gets cheaper, judgment becomes more valuable, but only if it is connected to the predictions. Operations research makes the point precisely. Elmachtoub and Grigas show that models trained to minimize prediction error, without regard for the downstream decision, can be more accurate and still lead to worse decisions. Bertsimas and Kallus show that the value of data lies in its *prescriptive* content, not its predictive content alone. Forecasts that stop at a dashboard are capability without coordination.

The best recent evidence shows both the promise and the condition attached to it. In the Procter & Gamble field experiment, individuals working with AI matched the performance of two-person teams without it, and teams with AI were the most likely to produce top-tier solutions. Yet across 106 experiments, human–AI combinations on average did worse than the better of the human or the AI working alone, and they lost most on decision tasks. The difference lies in how the work and the decisions are set up, not in the model.

EVIDENCE	WHAT IT SHOWS	SOURCE
Across 106 experiments, human–AI teams beat humans alone but, on average, underperformed the better of the two working alone; they gained on creation tasks and lost on decision tasks	Adding a human to AI, or AI to a human, does not guarantee a better result, least of all in decisions	Vaccaro et al. (2024)
At Procter & Gamble, individuals with AI matched teams without AI; teams with AI were most likely to reach top-10% quality, and AI users produced balanced solutions across R&D and commercial perspectives	Well-designed human–AI work can replicate and extend the benefits of teamwork	Dell'Acqua et al. (2026)
Consultants using AI did 12.2% more tasks, 25.1% faster and at higher quality inside AI's competence, but were 19 points less likely to be correct outside it	Gains depend on knowing where the AI's frontier lies	Dell'Acqua et al. (2023)
Prediction models trained for accuracy alone ignore how their outputs feed an optimization; training against decision error improves results, especially when the model is misspecified	Quantitative models must be built for the decision they serve	Elmachtoub & Grigas (2022)
Across 136 studies, mechanical prediction matched or outperformed clinical expert judgment in the large majority of comparisons	Encoded judgment is more consistent than unaided intuition	Grove et al. (2000)
Strong AI models perform worse when they must collaborate with each other than when working alone	Collaboration is a distinct capability, not a by-product of intelligence	Davidson et al. (2025)
41% of workers reported receiving AI "workslop", costing about two hours of rework each time	Unchecked output shifts effort downstream to colleagues	Niederhoffer et al. (2025)

Practitioner reports describe the same pattern from inside organizations. HMCI reports that after an agentic rollout, output tripled in the first week while reviewers flagged about 70% of first-pass drafts as polished but misaligned. Individual gains did not add up to team gains, because each person reinvented context and handoffs went undocumented. These are one organization's self-reported figures, but they match the research. Where coordination is designed, AI adds value. Where it is assumed, AI multiplies output and moves the cost of checking it onto someone else, usually the most experienced person in the room. The same holds for decisions: where the path from prediction to action is designed, quantitative models add value; where it is assumed, they add dashboards.

—— SECTION 03

Why coordination breaks down

Four mechanisms explain most failures. Each works at a different level: the conversation, the team, the decision and the profession.

1. Weak grounding leaves the human to do all the repair

People coordinate by building common ground: shared assumptions, goals and awareness of where a task stands. Klein et al. list it as one of the four conditions of joint activity. Common ground is not a given. It is built step by step through clarifying, acknowledging and fixing misunderstandings.

Vishwarupe et al. find that most LLM work sits in what they call *weak collaboration*. The exchange is iterative and looks like teamwork, but the human must spot what went wrong, rebuild the missing assumptions and steer the system back. The appearance of partnership outpaces the interaction's capacity to ground. Research on LLM dialogue supports this: models tend to presume common ground rather than build it, and preference training reduces clarifying behavior further (Shaikh et al., 2024).

The consequence is a hidden tax. Every AI-assisted task carries repair work that falls almost entirely on the person, and it rarely shows up in productivity figures.

2. Fluency hides misalignment

AI output looks finished before it is right. Well-formatted, confident prose reads as correct, and people check it less. The result has been labeled *workslop*: output that looks polished but lacks substance and shifts the real work onto colleagues (Niederhoffer et al., 2025). Practitioners describe the same effect, and link it to automation bias, to the tendency to trust what is easy to read, and to hallucination.

Pavlović describes the same risk at the individual level as *metacognitive laziness*: handing the monitoring and evaluation of one's own thinking to the AI. The opposite failure also exists. After visible errors, teams may abandon AI even where it helps. Vaccaro et al. show why calibration matters: human–AI teams lost ground on decision tasks, and did worst when the AI alone was more accurate than the human. Both over- and under-reliance are failures of calibration, not of the technology.

3. Predictions stop short of decisions

A forecast, score or valuation is not a decision. Turning one into an action requires an objective, constraints, a view of what each kind of error costs, and someone accountable for acting. In most

organizations that last mile is left to individuals with spreadsheets, or to automation configured once and rarely revisited, so decision quality depends on whoever closes the gap that day.

The expertise needed to close it often stays tacit. Experienced practitioners know which cases matter and why, but unaided judgment is also noisy: Kahneman, Sibony and Sunstein document how different experts, or the same expert on different days, reach different conclusions on identical cases, and a meta-analysis of 136 studies found that mechanical prediction matched or outperformed expert judgment in most comparisons (Grove et al., 2000). When judgment is neither encoded nor connected to the models, prediction and expertise run side by side and never meet. Vaccaro et al.'s finding that human–AI combinations lose most on decision tasks is the clearest sign of this gap.

4. The path to expertise is eroding

Judging AI output requires expertise, and expertise takes years to form. Rousse and Dreyfus show that people progress from following rules to acting on intuition only through involved, emotionally invested experience with many real cases. When situations become novel, even experts must fall back on slower, rule-like reasoning.

Rousse and Scapini argue that AI disrupts this ladder from both ends. It gives novices an elevator to expert-looking output without the learning. It also makes experts' familiar situations novel, so their intuition no longer applies directly. The Procter & Gamble experiment shows the elevator at work: employees new to product development performed at levels comparable to teams that included experienced members once they had AI. Meanwhile early data suggest AI is cutting entry-level hiring more than senior roles (Brynjolfsson, Chandar & Chen, 2026). If the bottom rungs disappear, so does the pipeline of people who can judge AI's work.

—— SECTION 04

Our position: seven principles for coordinated human–AI work

Principle 1 — Treat coordination as infrastructure

Coordination should be designed, documented and owned like any other operating system of the business. Klein et al. make the case in general terms: joint activity needs an explicit basic compact and continuous effort to stay predictable and directable. Rousse and Scapini make it from the enterprise side: AI projects tend to fail when treated as plug-in automation rather than as a new domain that requires human expertise. The Procter & Gamble study adds that AI changes how expertise flows across functions, which means team structures need to be redesigned, not just equipped.

That infrastructure extends to decisions. For each consequential decision, it should specify a stated objective, constraints, a named owner and a defined point where model output becomes action. In practice this means shared context that every person and agent can draw on, decision logs, and explicit handoffs. Practitioners describe documentation as infrastructure rather than overhead: model cards, datasheets and team decision records turn private know-how into something a team can reuse and check.

Principle 2 — Quantify the decision

Many of the decisions where AI matters most — pricing, valuation, credit, inventory, intervention — are quantitative at heart, and they deserve quantitative solutions rather than generic AI tools. That means encoding how experienced practitioners actually judge and decide into models: forecasts, valuations, optimization objectives and guardrails that make AI-driven judgment legible, testable and improvable. Encoded judgment is applied consistently, which addresses the noise that unaided expert judgment carries (Kahneman et al., 2021; Grove et al., 2000).

It also means building and measuring models for the decision they serve. Elmachtoub and Grigas show that training against the decision error a prediction induces, rather than prediction error alone, improves decisions, especially when the model is imperfect. The practical test applies before any build: **if this model were perfect, what would we do differently tomorrow?** Once a model is running, the measure of success is decision quality — outcomes against a benchmark, override rates, rework — not accuracy alone. Encoding is not replacing: a good model handles routine cases consistently and routes novel ones to the expert, whose intuition matters most when situations depart from the normal.

Principle 3 — Design for grounding and shared repair

A working human–AI system should make its understanding visible and share the work of fixing mistakes. Vishwarupe et al. identify three mechanisms that raise an interaction's grounding capacity:

- **Scoping** — narrow the task, fix formats and break work into stages, so less must be kept in shared view at once.
- **Signaling** — have the system restate its assumptions, summarize task state and flag uncertainty.
- **Repair** — build clear paths for revision, rollback, clarification and challenge.

These mechanisms put Klein et al.'s requirements into practice. Signaling makes the system predictable; repair makes it directable. Pavlović's *distributed metacognition* adds the psychological version: the AI should prompt reflection on the hypotheses being tested, explain why it suggests something and state its confidence. In a quantitative decision system, signaling means showing the key drivers and confidence behind every recommendation, and repair means making overrides cheap and recording the reason for each one so it improves the model. The practical test for any tool or workflow is simple: **who carries the repair burden?** If the answer is always the human, the collaboration is weaker than it looks.

Principle 4 — Gate AI work where judgment matters

AI should not be used less, but more deliberately. Human judgment is most valuable at specific points: setting direction before work starts, and selecting and verifying what AI produces. The Procter & Gamble study found that AI mainly improved idea generation, while human judgment kept its value in choosing among options. Vaccaro et al. found the largest losses on decision tasks, where people either deferred to AI too readily or overrode it without reason.

A practical way to protect these points is a small number of gates with explicit human sign-off. HMCI's practitioner framework offers a workable sequence:

- **Context** — who is this for, what problem are we solving, what do we already know?
- **Strategy** — what is the goal, what does success look like, what constraints apply?
- **Copy** — is the substance accurate, aligned and valuable?
- **Final form** — is the format, structure and design ready for its audience?

The first two gates matter most because they set the direction before any AI output exists. For quantitative decisions, the same logic applies with gates suited to the work: **frame** the decision and its objective, check the **model** and whether the case falls inside its validated range, review the **recommendation** against policy and expert sense, and **commit** with a named sign-off. Research on automation supports making such checks mandatory: accountability and exposure to failure cases reduce uncritical reliance. Gates need not be fixed. Autonomy can be earned with evidence: as a decision's track record builds, gates can lighten and automation can expand.

Principle 5 — Keep a human answerable

Responsibility for AI-assisted work cannot be delegated to the AI. Rousse and Scapini argue that delivering work is a promise to the person receiving it: that it will do what was agreed, and that its author will support it, fix it and make amends if it fails. Drawing on Winograd and Flores, they treat organizations as networks of such commitments. The same applies to decisions: a valuation, a credit decision or an operational intervention needs a named person who vouches for it.

This gives "human in the loop" definite content. The human is not a rubber stamp at the end of a pipeline. They are the party who vouches for the result and answers when it goes wrong. Pavlović's *distributed agency* makes the same demand from the design side: what is automated and what is human-led should be visible, negotiated and reversible.

Principle 6 — Match structure to the task

No single interaction style fits all work. Pavlović makes the case for open, dialogical collaboration in which the AI asks questions and offers possibilities rather than final answers. She also concedes this frustrates users on simple, factual tasks. Vishwarupe et al. favor tight scoping and staged work.

These views are complementary, and the evidence shows where each applies. Human–AI teams tend to gain on open-ended creation tasks and lose on pick-an-answer decision tasks (Vaccaro et

al., 2024). In the Procter & Gamble study, AI lifted the quality of ideas while people remained essential for selection. Routine, well-defined work — including rule-bound quantitative decisions — benefits from narrow scope, fixed formats, guardrails and checkpoints. Complex, ambiguous problems benefit from dialogue that surfaces assumptions and alternative framings.

Principle 7 — Rebuild the path to expertise

The most valuable skill in AI-assisted work is what Rousse and Scapini call *orchestration expertise*. It means guiding AI capabilities beyond one's own, telling good output from merely plausible output, and standing behind the result. Their account of how it develops has three stages: *awakening* to the fact that the domain has changed, *calibration* through deliberate practice with an unreliable but powerful partner, and *reimagination* of the work itself.

Orchestration depends on domain expertise. The accountant in their case study could direct AI to build a system she could not have coded, but only because she could check every result against records she understood. Vaccaro et al. point the same way: human–AI teams gained when the human was the stronger performer. Organizations therefore need what Rousse and Scapini call **AI-calibrated apprenticeship**. That means protected novice roles with bounded but real assignments, routine human review, shared libraries of AI failures, peer calibration groups and explicit verification rituals. Novices should learn with AI, not around it, and learn to answer for what it helps them produce.

—— SECTION 05

Objections and responses

"Better models will close the gap on their own." Capability and collaboration are different skills. Davidson et al. found that models strong on their own degrade when they must work with a partner, and Shaikh et al. found that current training can reduce grounding behavior. Klein et al. warned two decades ago that more autonomous agents can become harder to predict and direct. Better models will help, but they will not assign responsibility, set strategy or decide what is worth doing.

"Our models are already accurate; the problem is adoption." Accuracy is not the same as decision quality. A forecast can be more accurate and still lead to a worse decision if it is not built for the choice it serves (Elmachtoub & Grigas, 2022). What looks like an adoption problem is usually a decision coordination problem: no stated objective, no owner, no defined point where the model's output becomes an action. Fix those, and adoption tends to follow.

"Checkpoints are bureaucracy that cancels the speed gain." The alternative is not speed but rework. Workslop costs its recipients about two hours per instance (Niederhoffer et al., 2025). In HMCI's own rollout, the leader became the single quality bottleneck within three weeks, and the organization reports that direct interventions dropped significantly after it introduced its gates. The

gates are cheapest at the start: an error caught at the Context checkpoint never becomes a polished draft.

"Evidence shows human–AI teams often do worse than AI alone." Vaccaro et al. do find that, on average, human–AI combinations underperform the best of the two alone. But the same study finds gains when the human is more capable than the AI at the task, and on creative rather than decision tasks. The Procter & Gamble experiment found that teams using AI produced the strongest solutions of all. The lesson is not to remove people. It is to allocate work and decisions deliberately and to invest in the expertise that makes human judgment worth adding.

"AI can now do entry-level work, so we need fewer juniors." This is the oldest irony of automation: automating routine work removes the practice that builds the skill needed for rare, critical interventions (Bainbridge, 1983). Cutting junior roles saves money now and hollows out the supply of future orchestrators. There is a better option. AI raised novice productivity by 34% in one large field study (Brynjolfsson, Li & Raymond, 2025), and let newcomers perform like experienced teams at Procter & Gamble. Used within a designed apprenticeship, AI can speed up learning rather than replace it.

"Checklists contradict expert intuition." The Dreyfus model holds that experts act on intuition only when things proceed normally. In novel situations, experts return to deliberate reasoning. AI has made many familiar situations novel, so structured checks are exactly what the model predicts experts need during the transition. As calibration builds new intuition, the checks can become lighter.

—— SECTION 06

Recommendations

Each group has a distinct part to play. The actions below follow from the seven principles.

AUDIENCE	ACTION	REF.
Organizational leaders	Name an owner for the coordination layer: shared context, handoffs, decision logs and review standards	1
Organizational leaders	Inventory the highest-value recurring decisions and name an owner for each	1, 2
Organizational leaders	Start with a bounded internal AI lab where teams can test, fail and feed lessons back into operations	1, 7
Organizational leaders	Measure rework, review load and decision quality, not just output volume or model accuracy	2, 3, 4

AUDIENCE	ACTION	REF.
Organizational leaders	Keep and redesign entry-level roles as AI-calibrated apprenticeships	7
Team leads	For each model, write down the decision it serves: objective, constraints and the cost of each kind of error	2
Team leads	Adopt the Context, Strategy, Copy and Final form gates (or Frame, Model, Recommend and Commit for quantitative decisions) with explicit sign-off	4
Team leads	Assign a named human to vouch for every AI-assisted deliverable and decision that leaves the team	5
Team leads	Keep a shared library of AI failures and near misses	3, 7
System designers	Train and evaluate models against decision error, not prediction error alone	2
System designers	Build signaling in: restated assumptions, key drivers, task-state summaries and confidence flags	3
System designers	Make repair cheap: revision, rollback, clarification and contesting outputs, with override reasons captured	3
System designers	Show clearly what is automated and what is human-led, and let users take back control	5
System designers	Offer both scoped and dialogical modes, matched to task complexity	6
Practitioners	Make your decision rules, cues and exceptions explicit, and review the encoded version	2
Practitioners	Ask AI “what are the steps?” as well as “do this”, and use it to learn	7
Practitioners	Verify outputs against sources and benchmarks you already trust	4, 5
Practitioners	Join or form a peer group to share practices and failures	7
Educators and policymakers	Teach AI literacy as a mindset of active agency, not tool operation	6, 7
Educators and policymakers	Support protected learning spaces, such as sandboxes and apprenticeships, where judgment can form	7

A sensible first step for most organizations is small. Pick one recurring team deliverable or one recurring quantitative decision, add the gates, name an answerable owner and track rework and outcomes for a month. The results will show where the coordination gap is widest.

—— SECTION 07

Conclusion

The next gains from AI will not come mainly from better models. They will come from better coordination: shared understanding that is built rather than assumed, repair that is shared rather than left to the user, predictions carried through to decisions rather than stopping at a dashboard, and judgment kept at the points where it matters. They will depend on quantitative solutions that encode how experts decide, and on people who remain answerable for the work and are formed deeply enough to judge it.

Whether the expert as AI orchestrator becomes a lasting role or a passing one is not a technical question, as Rousse and Scapini conclude. It depends on choices organizations make now: to design the coordination layer, from conversations to decisions, to keep humans responsible, and to rebuild the ladder of expertise rather than let it fall. Coordination is the bottleneck, and it is within our control.

—— SOURCES

References

Primary sources

- Dell'Acqua, F., Ayoubi, C., Lifshitz, H., Sadun, R., Mollick, E., Mollick, L., Han, Y., Goldman, J., Nair, H., Taub, S., & Lakhani, K. R. (2026). The cybernetic teammate: A field experiment on generative AI and teamwork. *Organization Science*, 37(4).
- Klein, G., Woods, D. D., Bradshaw, J. M., Hoffman, R. R., & Feltovich, P. J. (2004). Ten challenges for making automation a "team player" in joint human-agent activity. *IEEE Intelligent Systems*, 19(6), 91–95.
- Pavlović, J. (2025). Constructivist psychology principles of human–AI collaboration. *Frontiers in Psychology*, 16, 1638774. doi:10.3389/fpsyg.2025.1638774
- Rousse, B. S., & Dreyfus, S. E. (2021). Revisiting the six stages of skill acquisition. In E. S. Mangiante & K. Peno (Eds.), *Teaching and Learning for Adult Skill Acquisition* (pp. 3–28). Information Age Publishing. Publisher page
- Rousse, B. S., & Scapini, M. (2026). The fate of expertise in the age of AI. *Journal of Expertise*, 9(2), 91–105. journalofexpertise.org
- Vaccaro, M., Almaatouq, A., & Malone, T. (2024). When combinations of humans and AI are useful: A systematic review and meta-analysis. *Nature Human Behaviour*, 8(12), 2293–2303.
- Vishwarupe, V., Jirotka, M., Shadbolt, N., & Flechais, I. (2026). The collaboration gap in human–AI work: Grounding and repair conditions for stable collaboration. *ECSCW 2026 – Posters and Demos*. EUSSET Digital Library

Decision science and judgment

- Agrawal, A., Gans, J., & Goldfarb, A. (2018). *Prediction machines: The simple economics of artificial intelligence*. Harvard Business Review Press.
- Bertsimas, D., & Kallus, N. (2020). From predictive to prescriptive analytics. *Management Science*, 66(3), 1025–1044.
- Elmachtoub, A. N., & Grigas, P. (2022). Smart "predict, then optimize". *Management Science*, 68(1), 9–26.
- Grove, W. M., Zald, D. H., Lebow, B. S., Snitz, B. E., & Nelson, C. (2000). Clinical versus mechanical prediction: A meta-analysis. *Psychological Assessment*, 12(1), 19–30.
- Kahneman, D., Sibony, O., & Sunstein, C. R. (2021). *Noise: A flaw in human judgment*. Little, Brown Spark.

Supporting evidence

- Bainbridge, L. (1983). Ironies of automation. *Automatica*, 19(6), 775–779.
- Brynjolfsson, E., Chandar, B., & Chen, R. (2026, revised August 12). Canaries in the coal mine? Six facts about the recent employment effects of AI. Stanford Digital Economy Lab.

- Brynjolfsson, E., Li, D., & Raymond, L. (2025). Generative AI at work. *Quarterly Journal of Economics*, 140(2), 889–942.
- Davidson, T. R., Fourney, A., Amershi, S., West, R., Horvitz, E., & Kamar, E. (2025). The collaboration gap. arXiv:2511.02687.
- Dell'Acqua, F., McFowland, E., Mollick, E., et al. (2023). Navigating the jagged technological frontier. HBS Working Paper 24-013.
- Niederhoffer, K., Kellerman, G. R., Lee, A., Liebscher, A., Rapuano, K., & Hancock, J. T. (2025, September 22). AI-generated "workslop" is destroying productivity. *Harvard Business Review*.
- Shaikh, O., Gligorić, K., Khetan, A., Gerstgrasser, M., Yang, D., & Jurafsky, D. (2024). Grounding gaps in language model generations. *NAACL 2024*, 6279–6296.
- St. Lawrence, S. (2026, April). *Managing Human and AI Teams* (White Paper No. 01). Human Machine Collaboration Institute. hmci.ai

AMBIENT POINT CORP.

—— ABOUT AMBIENT POINT

Models are not the bottleneck. *Decisions are.*

Ambient Point puts quantitative AI to work — as advisor, builder and operator. We help organizations decide where analytics and AI belong in their most consequential decisions, build the models and systems that put them there, and run companies where people, AI agents and data operate as one team. With domain experts as partners, we encode their judgment into quantitative systems they can trust.

Our model has three parts. **Encode** how practitioners work and decide into quantitative models. **Optimize** how models, optimization engines, people and AI agents divide work, hand off, escalate and verify. **Operate** the ventures that put this to work.

OFFICE

CIRA Centre, 2929 Arch Street, Suite 1700, Philadelphia, PA 19104

EMAIL · TELEPHONE · WEB

info@ambientpointcorp.com · (484) 734-0668 · ambientpointcorp.com

SUGGESTED CITATION

Rodriguez, A. (2026). *Coordination is the bottleneck: A position paper on human–AI collaboration and decision coordination* (Version 1.0). Ambient Point Corp.

This paper is provided for general information and discussion. It does not constitute professional, legal or investment advice. Figures attributed to third-party sources are reported as published by those sources; readers should consult the original works cited. © 2026 Ambient Point Corp. All rights reserved.