ARTICLE

AI-Augmented Sensemaking

From Fixed Applications to Collaborative Human-Agent Loops

By Mark Stratton

For most of software history, the application has defined the work.

Developers and product teams decide which data matters, how it will be organized, what screens users will see, which reports they can generate, and what sequence of steps constitutes the approved workflow. Users navigate menus, populate fields, interpret dashboards, and move through paths designed in advance by someone else.

That model works well when the situation is known, the process is stable, and the desired outcome can be specified ahead of time. It is the right model for processing a payment, provisioning a service, submitting an expense report, or completing a regulated transaction.

But much of the work that creates real enterprise value does not begin with a known path. It begins with ambiguity:

Something is wrong, but we do not yet know why. A customer is unhappy, but the underlying cause is unclear. The data contains a pattern, but we do not know whether it matters. A market is changing, but the implications are still forming. An engineer has an intuition, but needs evidence. A leader has to decide, but first has to understand what situation the organization is actually facing.

Traditional applications support pieces of this work, but the human usually performs the integration. We move among dashboards, documents, tickets, messages, meetings, spreadsheets, and specialists. We gather fragments, construct a working explanation, revise it as new evidence appears, and eventually decide what to do.

The applications contain the evidence. The sensemaking happens in our heads.

AI agents are beginning to change that relationship. Instead of forcing every situation through a predefined application and workflow, we can begin with an incomplete problem, an observation, a question, or an intent. An agent can then investigate the situation with us - searching data, consulting knowledge, calling tools, engaging specialist agents, presenting relevant evidence, testing possible explanations, and adapting its path as our understanding evolves.

The fundamental unit of work is no longer necessarily the screen or the workflow.

It is the collaborative loop between a person and an agent.

Sensemaking is the Missing Name for What Is Changing

Sensemaking is an established concept in organizational theory, cognitive science, information science, and human-computer interaction. It is more specific than simply "thinking" or "analyzing."

Sensemaking turns circumstances into an explicitly comprehensible situation that can become a basis for action. It is an iterative process in which people gather information, reorganize it into useful representations, develop insight, and turn that insight into a knowledge product or direct action. The process emphasizes loops: people search, organize, hypothesize, encounter evidence that does not fit, revise their representation, and continue.

Sensemaking becomes especially important when circumstances are incomplete, surprising, contradictory, or changing. It is how people answer questions such as:

- What is happening here?
- Which signals matter?
- What might explain them?
- What do we still not know?
- Which interpretation is sufficiently plausible to act upon?
- What should we do next, and what will the result teach us?

Sensemaking is related to reasoning, but it is not the same thing. Reasoning asks, "Given this representation of the situation, what follows?" Sensemaking asks, "What situation are we actually in, what matters within it, and which representation should we be using?"

Decision-making then selects an action. Agency makes the action possible. Feedback tells us whether our understanding was sound.

AI did not invent sensemaking. What AI has done is make it possible to augment this deeply human process with computational reach, persistent context, specialized knowledge, adaptive interfaces, and access to real operational systems.

That leads to a new working definition.

> AI-Augmented Sensemaking is the continuous, collaborative process through which people and AI agents jointly frame an ambiguous situation, assemble and interpret relevant evidence, generate and test explanations, decide what matters, and translate judgment into governed action - using live data, organizational knowledge, tools, memory, and feedback to adapt the path as understanding evolves.

Every part of that definition matters.

It is continuous, not a one-time prompt and answer. It is collaborative, not the wholesale delegation of judgment to a machine. It is grounded, because useful understanding depends on evidence and context. It is generative, because the path does not have to be fully defined in advance. And it is executable, because the process can move beyond insight into authorized action, observation, and learning.

Cognitive uplift is the promise. Sensemaking is the mechanism. Decision and action are the enterprise value.

The Loop Becomes the New Unit of Work

The emerging interaction can be represented simply:

Human intent -> agent interpretation -> data and tools -> evidence and hypotheses -> human judgment -> governed action -> observed outcome -> revised understanding

That is not a traditional workflow in which every branch has been anticipated. It can move forward, branch, backtrack, pause for another person, consult another agent, or return to an earlier hypothesis.

A useful design principle is:

The loop is the unit. The flow is the experience. The graph is the memory.

The loop is the recurring exchange among intention, investigation, interpretation, action, and feedback.

The flow is how that process feels to the person. The experience should remain coherent even when the investigation changes direction.

The graph is the accumulating structure underneath: the evidence reviewed, connections identified, hypotheses considered, questions asked, decisions made, actions taken, and results observed. That graph becomes organizational memory. It allows a future person or agent to understand not only what was decided, but why.

Current agent architectures already reflect this shift. Some approaches distinguish predefined workflows from agents that dynamically direct their own processes and tool use. Others describe agent execution as a loop that continues until an exit condition is met, including multi-agent systems in which an orchestrator coordinates specialized agents or agents hand work to one another.

A multi-agent system can therefore act like an on-demand cognitive team. One agent may understand the customer. Another may specialize in network telemetry. Another may search product documentation and errata. Another may challenge the leading hypothesis. Another may prepare and verify an operational change. An orchestrating agent can combine their work into one coherent interaction for the human.

The point is not to maximize the number of agents. Many situations will be handled best by one capable agent with several tools. Multiple agents become useful when meaningful specialization, parallel investigation, independent verification, or separation of authority improves the outcome.

The larger shift is that the person no longer has to coordinate every system and specialist manually. The agentic environment can assemble the right cognitive and operational capabilities around the situation as it unfolds.

From Predefined Interfaces to In-the-Moment Software

The traditional application asks the user to find the right screen.

An AI-augmented sensemaking environment asks the user what they are trying to understand.

That changes the role of the interface. Instead of always presenting the same dashboard, form, or report, the system can present the representation most useful to the current moment: a timeline, topology, comparison, map, anomaly plot, decision matrix, evidence table, simulation, approval control, or combination of these.

The interface becomes a temporary instrument constructed around the question at hand.

Emerging work on generative and agent-driven interfaces points in this direction: systems can dynamically create or assemble interactive experiences in response to a user's prompt rather than selecting only from static, predefined interfaces. Research prototypes are also exploring bidirectional loops in which the agent changes the interface as the task evolves, while the user's manipulation of that interface supplies additional context back to the agent.

This allows a professional to follow intuition without leaving the flow of work.

An engineer might say:

"The aggregate numbers look normal. Show me only the devices that roamed within thirty seconds of a service interruption."

Then:

"Separate those by device generation and access-point association."

Then:

"Overlay channel utilization and recent configuration changes."

No product team had to predict those exact questions and build a permanent screen for each one. The agent can generate the queries, analysis, visual representation, and next investigative options in the moment.

This is what makes the paradigm creative rather than merely automated. The human can pursue a hunch, challenge an explanation, change the frame, or introduce context that was not apparent in the data. The agent can respond immediately by changing what it retrieves, how it organizes the evidence, and what it presents next.

Data Becomes the Shared Reality

A foundation model by itself is not an enterprise sensemaking system.

The model provides general cognitive capability. The enterprise data substrate provides situational awareness. Knowledge gives the agent domain context. Tools give it agency. Governance defines its authority.

A useful formulation is:

Model = cognitive capacity

Data and knowledge = situational awareness

Tools = operational agency

Governance = authority and accountability

For AI-augmented sensemaking, the required context may include structured operational data, documents, product knowledge, event streams, historical incidents, customer records, business definitions, organizational policies, system topology, prior decisions, and the outcomes of earlier actions.

The quality of the sensemaking will depend on more than the quantity of that data. It will depend on whether the data is timely, joinable, correctly defined, traceable to its source, permissioned appropriately, and expressed with enough semantic context for the agent to interpret it accurately.

Emerging integration standards point to the infrastructure required here: AI applications need reliable ways to access external data sources, resources, tools, and workflows. Those tools may query databases, call APIs, modify files, or trigger other operational logic.

This means the enterprise data and knowledge foundation becomes more important in the agentic era, not less.

Poorly connected data does not merely produce a lower-quality answer. It can distort the agent's representation of the situation. Missing history can make an old problem appear new. Conflicting definitions can create a false trend. Incomplete topology can produce an incorrect root-cause hypothesis. Stale documentation can lead the system toward an obsolete action.

An articulate agent without trusted enterprise context is still detached from operational reality.

A Network Engineer Troubleshooting Home Wi-Fi

Consider a domain I know well: a network engineer troubleshooting a subscriber's home Wi-Fi service.

The subscriber reports:

"The internet drops upstairs almost every evening."

In the traditional application paradigm

A support ticket is opened. The engineer or support technician follows a troubleshooting tree. They inspect the subscriber record, gateway status, WAN or optical health, Wi-Fi controller, device inventory, historical alarms, previous tickets, and product documentation - often across different applications.

A dashboard may show average signal strength. Another system may contain channel-utilization history. A ticketing system may mention that the gateway was moved several months ago. Product documentation may contain a known issue affecting a particular firmware and client-device combination.

Each application does what it was designed to do. But the engineer has to integrate the evidence mentally.

The prescribed workflow may instruct the engineer to reboot the gateway, run a speed test, or replace equipment. It cannot easily anticipate every combination of home layout, device behavior, congestion, software version, subscriber history, and time-dependent RF conditions.

In an AI-augmented sensemaking paradigm

The engineer begins with intent:

"Help me determine why this subscriber experiences upstairs Wi-Fi interruptions in the evening. Minimize disruption and do not make a service-affecting change without my approval."

An orchestrating agent assembles the relevant context. It retrieves the subscriber's service profile, home gateway and extender topology, firmware versions, WAN health, seven days of Wi-Fi telemetry, client associations, roaming events, retransmissions, signal levels, channel utilization, configuration history, prior incidents, product documentation, and applicable operational policies.

Specialist agents examine different aspects of the situation:

- A network-access agent determines whether the underlying broadband connection is stable.
- A Wi-Fi agent analyzes coverage, interference, congestion, retries, and roaming behavior.
- A device agent looks for patterns associated with particular client types.
- A product-knowledge agent checks firmware notes and known issues.
- A verification agent challenges the leading explanation and searches for counterevidence.

The system does not overwhelm the engineer with every metric. It generates an in-the-moment evidence view showing the three most plausible explanations, the evidence for and against each, the remaining uncertainty, and the next discriminating question.

The initial evidence suggests that the WAN connection is healthy. Upstairs devices have marginal signal levels throughout the day, but interruptions become more frequent when evening channel utilization increases. One older client also appears to remain attached to the downstairs access point longer than expected instead of roaming.

The engineer follows an intuition:

"Show me the ten minutes before every interruption and separate the events by access point, client type, and roam attempt."

The agent regenerates the analysis. It finds that not every weak-signal event causes a reported interruption. The failures cluster when poor upstairs coverage, high channel utilization, and delayed roaming occur together.

The engineer now has a more useful causal explanation than "weak Wi-Fi" or "internet instability." The agent proposes two reversible actions and explains the expected effect and risk of each. With the engineer's approval, it applies a bounded configuration change through the network controller, updates the ticket, generates a subscriber-facing explanation, and begins monitoring the affected devices.

The next day, it compares the post-change results with the prior baseline. If retries and interruptions decline, it records the outcome and strengthens the working explanation. If they do not, the agent reopens the investigation rather than declaring premature success.

No application designer had to pre-author that exact investigative path.

The engineer remained responsible for intent, interpretation, risk, and approval. The agents supplied reach, correlation, domain specialization, evidence presentation, execution, and verification. The systems of record and network platforms remained essential, but they became tools within a larger collaborative cognitive process.

That is AI-Augmented Sensemaking.

Why This Goes Beyond Agentic Workflow

Agentic workflows are an important step, but they are not the full paradigm.

A workflow assumes that the organization substantially understands the process. Even when an agent dynamically chooses tools or delegates subtasks, the overall objective may still be the execution of a known business sequence.

Sensemaking begins where the correct sequence is not yet known.

The purpose of the loop is not merely to complete the workflow. It is to determine which workflow - or combination of investigation, judgment, collaboration, and action - the situation requires.

That distinction matters:

- Automation executes a known process.

- Reasoning works within a representation of the problem.
- Agentic execution dynamically performs work toward a goal.
- Sensemaking constructs and revises the representation that determines what the goal and next action should be.

Deterministic workflows will not disappear. Stable, repetitive, regulated, and safety-critical operations will continue to require explicit rules, controlled states, validation, and conventional software. In many cases, the agent will compose those deterministic capabilities rather than replace them.

The change is that static workflow will no longer be the dominant organizing principle for ambiguous knowledge work. It becomes one resource available inside a more adaptive sensemaking loop.

The Language Is Beginning to Converge

"AI-Augmented Sensemaking" is not yet a universally standardized category, but several closely related names are now appearing in current research.

Sensemaking AI describes AI that supports collective meaning-making in evolving human-AI networks. This perspective argues that AI does not merely supply information; it participates in the social and cognitive processes through which people interpret situations, coordinate, and act.

AI-Assisted Sensemaking describes human-led systems that use AI to identify themes, summarize evidence, and scale qualitative analysis while preserving human agency.

Collaborative Causal Sensemaking argues that current agents are too often trained as answer engines rather than partners capable of co-constructing explanations, surfacing uncertainty, adapting goals, and maintaining a shared mental model with a human expert. The shift is from oracle-like systems toward AI teammates that co-reason with people.

Other human-AI collaboration research similarly treats collaboration as an evolving sociotechnical process in which AI contributes speed, scale, and pattern recognition while humans contribute contextual judgment, interpretation, ethical responsibility, and accountability.

Adjacent terms include mixed-initiative systems, human-AI teaming, human-agent co-generation, cybernetic teammate, generative UI, and human-in-the-loop agentic systems.

No single label has won. But the direction is increasingly clear.

I prefer AI-Augmented Sensemaking because it keeps the human process at the center.

"AI-assisted" can understate how actively agents may investigate, challenge, coordinate, and act. "Sensemaking AI" can sound as though the machine itself is the principal sensemaker. "Human-in-the-loop" often reduces the person to an approval or exception-handling mechanism. "Agentic workflow" centers execution rather than the construction of understanding.

AI-Augmented Sensemaking describes a partnership in which AI materially expands human cognitive and operational capacity without relocating meaning, values, responsibility, or final accountability to the machine.

The Next Stage of Enterprise AI

Many enterprises began their generative AI journey with lists of isolated use cases: summarize a document, draft a message, answer a question, generate code, or classify a ticket.

The next step was to embed AI into workflows. Organizations began connecting models to data, applications, and tools so that AI could participate in multi-step business processes.

The agentic wave goes further by allowing systems to plan, select tools, coordinate specialized agents, and execute parts of those processes with greater independence. The direction is a shift from asking models for individual outputs toward delegating more complex, multi-step workflows.

AI-Augmented Sensemaking is the next conceptual step:

1. Use-case AI: Produce an output.
2. Workflow AI: Assist a predefined sequence of work.
3. Agentic AI: Dynamically execute and coordinate work toward a goal.
4. AI-Augmented Sensemaking: Jointly determine what is happening, what matters, what path is required, and what action should follow.

This is where many of the most valuable enterprise opportunities sit: incident diagnosis, strategic planning, product design, research, operational optimization, complex customer escalations, forecasting, risk assessment, architecture, investigation, and executive decision support.

These are not merely tasks waiting to be automated. They are situations that must be interpreted.

The enterprise design question therefore changes from:

"Where can we insert an agent into this workflow?"

to:

"Where do our people repeatedly have to assemble fragmented evidence, develop an explanation, make a judgment, coordinate expertise, and act under uncertainty - and how can agents improve that entire loop?"

That is a more powerful way to identify AI opportunities. It also moves the organization beyond disconnected bots toward domain-level sensemaking environments grounded in shared data, knowledge, tools, memory, and governance.

The Governance Challenge Is Premature Closure, Not Just Hallucination

The most commonly discussed AI risk is that a model may produce an incorrect answer.

In sensemaking, there is another risk: the system may establish the wrong frame too early.

An agent can prioritize some evidence, omit other evidence, reinforce the user's initial assumptions, or make one plausible narrative appear more certain than it is. Sensemaking-oriented AI research warns that AI can shape sensemaking trajectories, reinforce particular narratives, and marginalize alternative interpretations.

A trustworthy AI-augmented sensemaking environment should therefore be designed to:

1. Establish meaning before reasoning. Clarify the purpose, stakeholders, constraints, values, and consequences before optimizing an answer.
2. Make evidence visible. Preserve provenance, distinguish observation from inference, and allow people to inspect the basis for an explanation.
3. Maintain interpretive flexibility. Present competing hypotheses, search for counterevidence, and make uncertainty explicit.
4. Use bounded agency. Separate recommendation, simulation, drafting, approval, execution, and verification according to risk.
5. Close the loop. Observe the outcome of an action and incorporate what happened into future understanding.
6. Retain deterministic controls. Use conventional software, policy engines, permissions, validations, and audit trails where reliability requires them.

The objective is not an AI system that always appears confident. It is a collaborative environment that helps people reach better-grounded understanding and act with appropriate confidence.

A New Working Normal

AI-Augmented Sensemaking is not simply another product feature or interface style. It is an emerging operating paradigm for knowledge work.

Applications will remain. Systems of record will remain. Databases, APIs, workflow engines, dashboards, and deterministic services will remain. But increasingly, they will become the underlying capabilities through which agents investigate and act.

The visible working surface will become more fluid.

We will express intent, explore evidence, ask follow-up questions, challenge explanations, summon specialist capabilities, generate temporary interfaces, authorize actions, and observe outcomes - all within a continuous collaborative flow.

The system will not merely retrieve what someone previously decided to place on a dashboard. It will help determine what is relevant now.

It will not merely route us through the approved workflow. It will help determine what path the present situation requires.

It will not merely provide an answer. It will help us develop an understanding, make a judgment, act, and learn.

This will gradually become as ordinary as search, dashboards, and workflow software are today. The organizations that move first will not simply deploy more agents. They will redesign the human intent-to-outcome loop around a new division of cognitive and operational labor.

Humans will contribute purpose, meaning, intuition, lived context, values, judgment, and accountability.

Agents will contribute reach, speed, synthesis, memory, specialized reasoning, tool use, coordination, and bounded action.

Together, they will make sense of situations that neither traditional applications nor isolated AI models could adequately address alone.

AI has made cognition interactive, connected, and increasingly executable. The next enterprise paradigm is not merely agentic automation. It is AI-Augmented Sensemaking.

The future of work is not humans operating inside applications. It is humans and agents making sense of the world together - and changing it, carefully, one loop at a time.

Further Reading & Sources

AI-Augmented Sensemaking draws on foundational and emerging work in organizational sensemaking, human-AI collaboration, agent design, generative interfaces, and enterprise AI.

1. Weick, Karl E., Kathleen M. Sutcliffe, and David Obstfeld. "Organizing and the Process of Sensemaking." Organization Science 16, no. 4 (2005): 409-421.
2. Pirolli, Peter, and Stuart Card. "The Sensemaking Process and Leverage Points for Analyst Technology as Identified Through Cognitive Task Analysis." International Conference on Intelligence Analysis (2005).
3. Anthropic. "Building Effective AI Agents." December 19, 2024.
4. OpenAI. "A Practical Guide to Building AI Agents."
5. Google Research. "Generative UI: A Rich, Custom, Visual Interactive User Experience for Any Prompt." November 18, 2025.
6. Google Developers Blog. "Introducing A2UI: An Open Project for Agent-Driven Interfaces." December 15, 2025.
7. Xu, Yuan, et al. "DuetUI: A Bidirectional Context Loop for Human-Agent Co-Generation of Task-Oriented Interfaces." arXiv:2509.13444 (2025).
8. Model Context Protocol. "What Is the Model Context Protocol?"
9. Comes, Tina. "Sensemaking AI: Introducing a Research and Design Agenda for Human-AI Networks." EPJ Data Science 15 (2026).
10. MIT Center for Constructive Communication. "AI-Assisted Sensemaking."
11. Jain, Raunak, and Mudita Khurana. "Collaborative Causal Sensemaking: Closing the Complementarity Gap in Human-AI Decision Support." arXiv:2512.07801 (2025).
12. Hao, Xinyue, Emrah Demir, and Daniel Eyers. "Beyond Human-in-the-Loop: Sensemaking Between Artificial Intelligence and Human Intelligence Collaboration." Sustainable Futures 10 (2025): 101152.
13. Ye, Runlong, et al. "ScholarMate: A Mixed-Initiative Tool for Qualitative Knowledge Work and Information Sensemaking." arXiv:2504.14406 (2025).
14. Microsoft Research. "Magentic-UI: Towards Human-in-the-Loop Agentic Systems." 2025.
15. Dell'Acqua, Fabrizio, et al. "The Cybernetic Teammate: A Field Experiment on Generative AI Reshaping Teamwork and Expertise." Harvard Business School Working Paper 25-043 (2025).
16. OpenAI. "The State of Enterprise AI." December 17, 2025.