

QRCA
VIEWS

TOOLBOX

Agentic AI and Qualitative Research: A Work-in-Progress View

BY USHMA KAPADIA

Marketing Content Specialist

flowres.io / MyTranscriptionPlace.com

Lewes, Delaware

ushma@flowres.io



BY KAWALPREET JUNEJA

Head, Product Development

flowres.io / MyTranscriptionPlace.com

Lewes, Delaware

kp@flowres.io



When it comes to qualitative research debriefs, trust often rides on the researcher’s skillset. A senior researcher signs off on the “story,” and the “story” stands. Sure, context and experience stay at the core of the qualitative research craft. However, as stakeholders begin to increasingly consider do-it-yourself qualitative research (i.e., conducting their own research using self-service artificial intelligence platforms), they’re likely to want to integrate evidence “in the moment.” They’re likely to seek reassurance from the workflow itself—i.e., they want only verifiable claims to be cited in the debrief, every consumer segment to be included, and ready-to-access links to consumer verbatims. This is exactly where Agentic AI can step in and deliver trust.

Taking a Step Back...What is Agentic AI, Really?

The term Agentic AI describes artificial intelligence systems or agents that plan, act, and adapt toward a goal—coordinating steps, employing tools, and self-correcting, rather than producing one-and-done outputs. Learnings from industry leaders like IBM¹ and AWS underscore the emphasis on autonomy with guardrails: AI agents pursue goals, route work to the right tools for each task, and still keep humans in charge of critical decisions. In short, the advent of Agentic AI signals a move from reactive, automated assistants to proactive, human-guided, AI-powered collaborators.

In the context of qualitative research, Agentic AI is a system that acts autonomously, takes charge of repetitive tasks, pursues a research project task step-by-step. It is not just a chatbot that answers once and stops.

“Agentic AI can gather, check, and propose—but humans must frame the question, set the evidence bar, evaluate the agent-proposed actions, and frame recommendations that address client objectives and budgets.”

Agentic AI research applications for qualitative research can span across a typical project cycle. Think of AI-agents simply as “software colleagues”—that can play various roles: a fact finder, an intern, a planner, an analyzer, a data checker, and even a draft report writer. Each does one thing, and the system records who did what, when.

For instance, Agentic AI can help qualitative researchers:

1. Conduct secondary research to support writing a proposal.
2. Generate a discussion guide.
3. Manage individual parts of a live project (e.g., drafting stakeholder invitations, scheduling interviews, managing research-participant incentives).
4. Take, collate, and circulate qualitative session notes.
5. Transcribe and translate qualitative sessions.
6. Plan analysis steps (e.g., build a draft codeframe of verbatims (list of themes), then analyze for similarities and differences by consumer segment).
7. Execute some of the analysis steps (e.g., tag quotes, identify and group themes, identify contradictions).
8. Check and interrogate the AI-generated themes (e.g., “Did we cover all consumer segments? Which claims are weak?”).
9. Provide proof of analyzed themes (e.g., link every claim or insight to the exact quotes and timestamps).
10. Create a debrief in PowerPoint format with speaker notes.

No, Agentic AI isn’t an adopt-or-die harbinger of qualitative research in the future. At this stage, it’s

better to decide whether and how to use it, once you’ve asked yourself:

- Is it addressing a pain point we (qualitative researchers) frequently face? How often do we research a brief from scratch, before writing a proposal? A commonly occurring use case is one that involves tedious, repetitive tasks, such as researching current market conditions to help define a target recruit.
- Do we really need it? For a small, tightly designed study, a strong human analyst plus an efficient, automated search of AI-generated transcripts might deliver most of the desired value.
- Will it distort what we see? When AI automatically sets up the analysis —such as selecting segments, variables, or comparison groups —it might omit something important. For example, if the Agentic AI fails to include gender as a segmentation dimension when evaluating

often, and fall back to “manual” labor (scrapping Agentic AI) when the system’s confidence seems low.

Our stance is practical: use Agentic AI where it clearly improves efficiency, quality, or speed; don’t use it where a simpler, human-led process is clearer and safer.

A Simple Analysis Workflow: One You Can Instantly Visualize

In a traditional qualitative project, analysis follows a linear, manual path:

- Build a codeframe of verbatims
- Manually code transcripts and notes
- Identify patterns and themes
- Draft the insights and build the final report

With an Agentic AI workflow, this becomes a coordinated, step-by-step sequence handled by specialized agents:

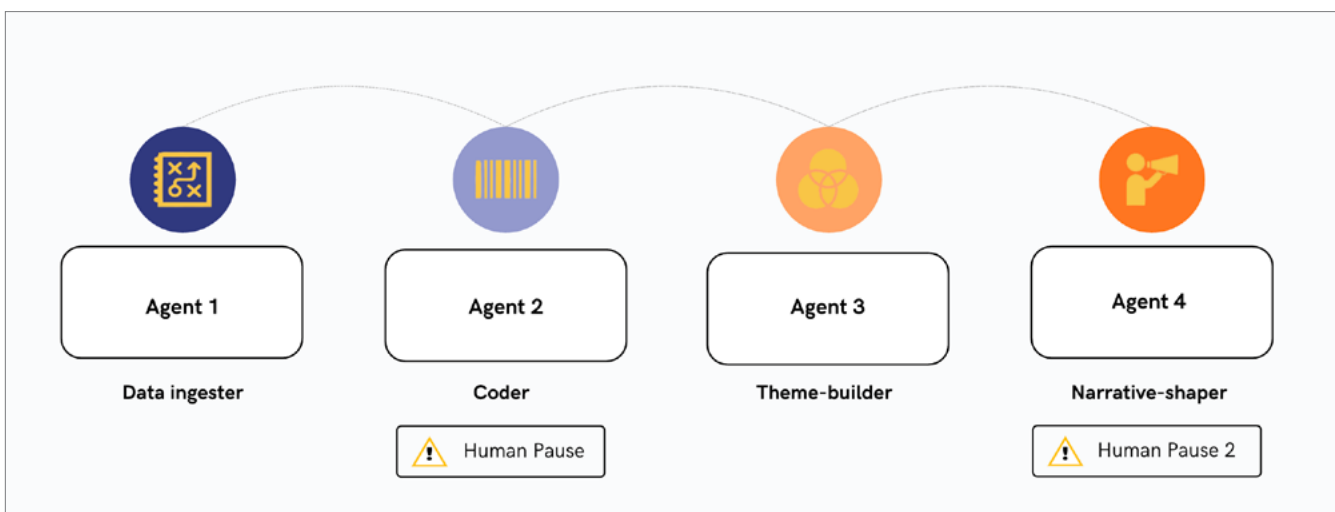


Figure 1. An example of an Agentic AI workflow applied to a typical qualitative research study.

reactions to a new product concept, and instead analyzes only by age, the resulting insights will be incomplete or misleading.

- Is it creating value (by saving time, money, or effort)? Running long, unattended Agentic AI sessions can waste research time and budget. Keep sessions short, review them

The researcher remains in control at two key checkpoints or pause points across the Agentic AI workflow (see Figure 1).

1. **After fieldwork:** approving or adjusting how the AI structures the data and codes

2. **Before final delivery:** reviewing, correcting, and shaping the storyline and visuals

What Changes in Practice and What Doesn't Change

First, time is reallocated, not merely reduced. Agentic AI compresses mechanical steps (e.g., first-pass verbatim coding, insight identification and retrieval, evidence compilation) so researchers can spend more time on judgment, i.e., reframing research questions, testing possible explanations, and confidently navigating stakeholder queries.

Second, insights become verifiable. Instead of a “trust me!” approach, every claim or insight can include its supporting verbatims: what data (verbatim) was consulted, which alternative demographic cuts (geographies, age-group, gender, etc.) were tested? Human pause points can be strategically placed, where humans must review before the Agentic AI proceeds. This is particularly useful for studies on sensitive categories and with large sample sizes.

Third, the profession stays humble. Gartner's² recent warning about “agent-washing” (mislabeling macros or other tools as “Agentic”) and the high failure rate of early pilot projects is a healthy reminder (of what can go wrong and what to avoid). Value is derived when we design for clarity, controllability, and measurement; not because we simply rename traditional automated macros as “Agents.” In other words, we only call something “Agentic” when it truly plans steps, takes actions, can explain itself under scrutiny; otherwise, it's just automation with a new label.

“Agentic AI can make the analysis process more transparent—showing which quotes or data points informed each theme—and give clients a clearer view into how conclusions were reached.”

Where Agentic AI in Qualitative Research is Headed

Our stance: the relevant debate isn't “Agentic AI in qualitative research—go or no-go?” Instead, it's “Agentic AI in qualitative research—how?” The need is to design Agentic AI systems that respect the craft of qualitative research and reserve human effort for the parts that can make-or-break

business outcomes based on research findings. Agentic AI can gather, check, and propose—but humans must frame the question, set the evidence bar, evaluate the agent-proposed actions, and frame recommendations that address client objectives and budgets. Done poorly, agentic tooling will simply generate bad stories faster. Done well, it will mean that researchers spend less time proving claims or insights and more time interpreting results.

Agentic AI is best understood as an operating model for effective and optimized qualitative work—not a feature to be bandied about in sales pitches. In that sense, the real opportunity for qualitative

researchers isn't to automate judgment, but to support it. Agentic AI can make the analysis process more transparent—showing which quotes or data points informed each theme—and give researchers a clearer view into how conclusions were reached. This keeps researchers focused on the higher-level interpretation and recommendations that drive client decisions.

What Does All of this Mean for Qualitative Researchers Today?

Our opening argument was that trust can no longer rest on a senior researcher's skillset or experience alone. What's required instead is a practice where the path from raw verbatims to boardroom slides is clearly evident. This level of

transparency matters because it helps clients understand how conclusions were formed, strengthens confidence in the insights, and reduces the risk of bias or misinterpretation.

So, when you finish your next qualitative research debrief, make the evidence as clickable as the headline: show the claim, show the specific quotes behind it, and be willing to look at verbatims in a different way (coded verbatims) in the room when a stakeholder asks a “What if?”

Another place to begin is even before you start coding your dataset (verbatims): use a planning Agentic AI to propose lenses; then apply a human pause point during this stage. If the Agentic AI forgets a demographic lens you know matters, you catch it early.

Third, let a reviewer-style Agentic AI do the drudgery we all avoid but clients trust: pull contradictory quotes against each major claim, check which segments were under-represented, and flag places where confidence is thin. This strengthens rigor, so the human analyst (you!) can spend attention on building defensible debriefs rather than collating raw data (verbatims).

We view Agentic AI as work-in-progress—it automates the iterative, manual effort that qualitative researchers pour into qualitative research debriefs. And it ropes in the human, at strategic stages of the qualitative process, to add value to and elevate debriefs. This is our practitioner’s view on Agentic AI and qualitative research, built from ongoing experiments and client experiences. We encourage the qualitative research community to continue refining the qualitative “playbook,” as Agentic AI technology (and guardrails) mature.

References

1. Gutowska, Anna. “What Are AI Agents?” IBM Think, IBM, accessed November 4, 2025. <https://www.ibm.com/think/topics/ai-agents>.
2. Kachwala, Zaheer (reporting). “Over 40% of agentic AI projects will be scrapped by 2027, Gartner.” *Reuters*, June 25, 2025. <https://www.reuters.com/business/over-40-agentic-ai-projects-will-be-scrapped-by-2027-gartner-says-2025-06-25/>. ■



Kawalpreet Juneja (KP) is Co-Founder and CTO at flowres.io, where he leads the technology thinking that is designed to simplify workflows. A computer science graduate, KP brings a technologist's lens to the practical challenges faced by researchers, moderators, and client observers. His writing on the flowres blog reflects his interest and experience in engineering, product design, and research operations. In particular, he enjoys proving how technology can elevate the qualitative research experience for all stakeholders involved. He is particularly passionate about speech technology, AI-led transcription, opensource tools, and the role of cloud infrastructure in making advanced research technology more accessible. KP's perspective is rooted in building scalable, useful products that make online qualitative research faster, cleaner, and easier to run.

 [Connect on LinkedIn](#)



Ushma Kapadia writes about the fast-changing world of insights; where qualitative research, ResTech, and AI collide. On the flowres blog, she writes about purpose-built LLMs, qualitative data analysis, research software, and the "Which platform should we actually use?!" questions facing qualitative research buyers. Her perspective is shaped by a career across qualitative research, social media analytics, content writing and strategy, and marketing ResTech solutions, including roles at TNS/Kantar and analytics/ResTech startups. She understands both the researcher's pressure to protect data quality and the business team's need for faster, clearer, more defensible decisions. Whether unpacking generic AI's limitations, explaining online qualitative workflows, or turning product updates into useful stories, Ushma digs into issues that are rapidly shaping how qualitative research is conducted today.

 [Connect on LinkedIn](#)



Thanks to our feature editor Ty Michael for his collaboration with the authors.

Ty Michael has managed qualitative and quantitative market research projects for over 20 years and is passionate about making sense of people, markets, and the motivations behind consumer behavior. Throughout his career, he has utilized various methodologies for his clients, including ethnographies, brand tracker surveys, focus groups, and product clinics, with a strong focus on the automotive market. What Ty enjoys most about research is finding the story in the data and translating it into insights that are meaningful and actionable for his clients. Before his current role as Account Manager at INTAGE USA, Ty held research roles at GfK, AutoPacific, and SEMA, building a career rooted in curiosity, data analysis, and practical problem-solving. He holds an MBA in Marketing from California State University, San Bernardino.



[Connect on LinkedIn](#)