



Essay

Interpretive orchestration: An essay exploring the epistemic intersection of human intuition and machine intelligence

Strategic Organization

1–15

© The Author(s) 2026



Article reuse guidelines:

sagepub.com/journals-permissions

DOI: 10.1177/14761270261448645

journals.sagepub.com/home/soqXule Lin¹  and Kevin Corley² 

Abstract

As artificial intelligence capabilities expand beyond pattern recognition to theoretical insight generation, interpretive qualitative research confronts a question of epistemic responsibility: how can scholars integrate artificial intelligence capabilities while remaining accountable for their theoretical interpretations? This essay proposes “interpretive orchestration” as a framework that transforms researchers from analysts into skilled orchestrators of human–artificial intelligence collaboration. The framework addresses two challenges that become opportunities. The translation challenge of articulating tacit knowledge (theoretical orientations, contextual understanding, and embodied intuition) into forms that artificial intelligence can process deepens researchers’ awareness of their own expertise. The judgment challenge of evaluating artificial intelligence-generated patterns for theoretical significance highlights the accountability our scholarly communities require, particularly through “1.5-order data”: patterns invisible to human perception yet requiring human interpretation for recognized theoretical significance. Three strategic models guide this orchestration: Socratic tension surfaces implicit assumptions through deliberate contradiction; Euclidean documentation enables reproducible analysis through systematic context-building; Vitruvian mastery reads across independent analytical passes for synthetic insight. By embracing orchestration, researchers discover that artificial intelligence can amplify rather than replace human capability. The future of interpretive research lies neither in rejecting artificial intelligence nor surrendering to automation but in systematic approaches to human–artificial intelligence collaboration that preserve the scholarly accountable judgment our communities require while drawing on artificial intelligence’s capacity to generate theoretical insights across scales humans cannot process alone.

Keywords

epistemic responsibility, generative AI, human–AI collaboration, interpretive orchestration, qualitative methods, tacit knowledge

¹SKEMA Business School, Université Côte d’Azur, Nice, France

²Department of Management and Entrepreneurship, Imperial Business School, London, UK

Corresponding author:

Kevin Corley, Department of Management and Entrepreneurship, Imperial Business School, South Kensington Campus, Exhibition Rd, South Kensington, London SW7 2AZ, UK.

Email: k.corley@imperial.ac.uk

Introduction

When considering content for a methods symposium at a not-so-distant Academy of Management meeting, it still seemed reasonable to ask, “will AI tools significantly impact qualitative research?” As often happens with technological change, however, recent advancements have largely rendered that question moot. It is now clear that artificial intelligence (AI) technology (in its many forms) is inevitable and will irrevocably impact how interpretive qualitative research (Guba and Lincoln, 1994) is done. We take as given that the rise of AI tools (and their incessant advancement) is profoundly challenging our traditional notions of what it means to “do interpretive qualitative research.” This raises a deeper, more primary question: “what does it mean to be an interpretive qualitative scholar in the age of AI?” Because so much of theory-building qualitative research rests on the belief that it is necessary to have deep engagement with the focal phenomenon and those living it, as well as spending time iteratively wrestling with one’s data, the addition of automated tools that can not only create, clean, and analyze data but also appear to develop empirical and theoretical insights as well represents an existential crisis for our scholarly communities (Kulkarni et al., 2024; Nolan, 2025).

Of course, this crisis will be felt more acutely by some and less by others. While we see the increasing use of AI in qualitative research as an inevitability, we acknowledge there will always be those who choose not to engage with this technology. This view is currently captured in a number of critical articles that have recently been published (Bechky and Davis, 2025; Lindebaum and Ashraf, 2024; Nguyen and Welch, 2025) outlining the case against uncritical AI adoption. Others are questioning AI’s potential as a collaborative partner (Mollick, 2024) and examining how intelligent machines might reshape human skill and expertise (Beane, 2024). At the heart of all these perspectives is a concern about agency: does delegating interpretive tasks to algorithms equate to researchers abdicating responsibility for the epistemic integrity of their findings? Similar concerns have arisen with previous jumps in technology as well. Indeed, there are scholars who, even today, shun the use of digital data management and coding technologies from the past 30 years (e.g. Atlas TI and NVivo) and prefer to do their analysis by hand (long live the Post-it note!). This essay is not intended to pass judgment on those choices or to convince scholars to change their thinking. We are firm believers in the value of diverse perspectives and techniques when it comes to interpretive qualitative research and believe that this diversity strengthens our communities (e.g. Cloutier et al., 2026; Corley et al., 2021). Nevertheless, there will be a large contingent of scholars in our field (and other social sciences) who do choose to bring AI into their practice. This essay is intended to provide some reflection and consideration of new understandings of who we are and what we do as interpretive qualitative scholars.

We came to this understanding through recent experiments integrating AI into our qualitative analysis.¹ We engaged with AI systems² across a range of interfaces, from conversational platforms (e.g. Claude.ai and Google AI Studio) to agentic environments capable of executing multistep tasks with minimal oversight (e.g. Claude Code and OpenCode). What unifies this variety is the human–AI interaction: researchers articulate intent through natural language, receive generated output, and exercise interpretive judgment over what emerges. Technical architectures continue evolving (e.g. transformer-based models and diffusion models); the systems dominating today will yield to approaches we cannot yet anticipate. We therefore describe our framework in terms of this persistent, dialogical interaction, where knowledge is communicated conversationally rather than encoded as formal parameters. The principles we discuss hold wherever researchers engage AI for interpretive work, and the models we develop describe strategic orientations researchers could adopt, not features of specific technologies.

In one experiment involving future-oriented discourse, an AI system processed in under an hour what previously required days. It identified subtle linguistic patterns related to uncertainty that we had missed. Yet, ultimately, this initial analysis disappointed, not because of technological limitations but because the AI in our experiments did not have access to the contextual knowledge needed for deep insight that we unconsciously applied to the data. The theoretical frameworks guiding inquiry, the significance of particular informant perspectives, and the weight certain experiences carried in the field all remained locked in our minds, inaccessible to the AI system.

This experience revealed underlying issues beyond technical considerations. Interpretive qualitative inquiry rests on the premise that understanding emerges through deep engagement with phenomena and those living them. The researcher's immersion in context, their empathetic connection with informants, their embodied experience of the field, and the resulting sharpening of intuition are central to this approach—these form the very foundations of how interpretive knowledge is created. When AI identifies patterns researchers miss or generates unexpected theoretical connections—appearing to exercise agency—an uncomfortable question emerges: what does it mean to be a qualitative scholar when interpretation appears possible without conscious, embodied engagement? When insight becomes indistinguishable from simulacra (as Shanahan (2024) theorizes and Park et al. (2023) explore empirically with AI systems exhibiting believable human-like behavior)?

Through conversations with qualitative scholars, AI researchers, and practitioners and structured dialogues with AI systems using our research materials as context, we have come to understand that epistemic responsibility is central to this existential crisis of agency: how can interpretive scholars make the most of what AI has to offer while ensuring that the end result remains one they hold ultimate responsibility for? We believe the importance of human researchers ultimately being responsible for the outcome of an interpretive process when studying human phenomena (e.g. organizations, emotions, and interactions) is paramount for two reasons: one philosophical and one practical. First, it is critical that theorizing about human phenomena fundamentally involves researchers who share in that humanity. Second, our scholarly communities (i.e. journals, review committees, etc.) require human accountability to stand behind interpretations and defend them.

Our emerging answer to this existential crisis involves the stance that qualitative researchers need to embrace new mindsets around the inductive and abductive processes we currently employ. Mindsets that recognize that agency for empirical analysis can be shared without giving up responsibility for the theoretical outcome. We label one such approach "Interpretive Orchestration," as it involves the skilled integration of multiple forms of meaning-making into coherent, worthwhile insights. On the one hand, the scholarly study of human phenomena calls for researchers who share in that humanity. And yet, as technology becomes more a part of the human experience, human intelligence can be enhanced through machine intelligence. Consider that the space of possible interpretive insights within any study appears infinite to the human mind. Within this vastness, interpretive orchestration positions both human researchers and AI as creative agents navigating the same terrain, though with different emphases: machine intelligence draws heavily on patterns across training data, while human researchers draw on embodied experience, emotional connections, and tacit knowledge. This difference in how they traverse the same phenomenological terrain is what makes their orchestration theoretically generative.

We see qualitative researchers of tomorrow blending the outputs of human intuition and machine intelligence into insight that is greater than the sum of their parts through tackling data sets too large and phenomena too complex for human minds alone while making critical engagement with doubt and alternative perspectives a natural part of the practice. While this approach requires new technical skills, the deeper transformation lies in how we conceptualize meaning-making in qualitative research. By positioning the researcher as an "orchestrator" of both human and AI, we open

new opportunities for theoretical insight while preserving the human epistemic responsibility that gives interpretive qualitative research its unique value. This orchestration involves two distinct but intertwined forms of epistemic work and their accompanying challenges: the art of translating complex human knowledge into a form AI can process and the critical act of applying deep human judgment to the patterns AI reveals.

The nature of the agency issue

We believe the agentic issues generating the need for interpretive orchestration are not technical but behavioral, centering on a dual challenge. The first is the *translation challenge*: articulating the rich contextual knowledge that guides human interpretation into a form that AI systems can comprehend. Successfully meeting this first challenge immediately creates the second: the *judgment challenge*, which involves evaluating the theoretical significance of the novel patterns AI generates. While both are critical and inextricably linked in a cycle of inquiry, the translation challenge represents the primary bottleneck in current practice. This is because any translation failure prevents researchers from even accessing the computationally derived patterns that require human judgment. This bottleneck creates a fundamental tension at the heart of our inquiry: AI systems possess increasingly sophisticated analytical capabilities, yet these capabilities often underperform because we struggle to communicate the contextual knowledge that guides interpretation.

These underperforming AI capabilities now include abilities recently considered uniquely human: pattern recognition, regularity detection, and conceptual category generation capabilities (Shanahan, 2024; Wei et al., 2022). When properly contextualized, they can analyze discourse through specific theoretical lenses, identify stakeholder perspectives in textual data, trace conceptual relationships across documents, and meaningfully achieve fair agreement on taxonomic labeling tasks (Ziems et al., 2024). In our recent experiments, AI systems distinguished between distinct types of organizational responses to change, identified temporal patterns in narrative construction, and mapped theoretical concepts to empirical instances. But this was achieved only when we provided sufficient contextual grounding.

For instance, when analyzing governance of online communities, rather than invoking governance theory as an abstract lens, we specified that consensus among these participants depended on perceived procedural fairness rather than outcomes and that governance challenges manifested through both rejected governance proposals and enacted governance changes. This kind of theoretical translation—activating what the AI might already “know” from training data by directing it toward the specific context—proved the most tractable. Harder was encoding methodological judgment: that governance discourse in asynchronous forums functioned differently than synchronous discussion, that certain participant positions carried particular significance, and that apparent consensus often masked unresolved disagreement. The hardest was translating embodied field experience—articulating that a recurring phrase carried connotations shaped by events participants referenced obliquely or that the emotional register of certain exchanges signaled something beyond the literal content and would require knowledge gained from field immersion to interpret. Here, translation often failed initially, and the failure was diagnostic: what the AI couldn’t work with revealed what we had not yet made explicit to ourselves.

Thus, the translation challenge emerges from what researchers typically leave unsaid, which involves an interplay between two kinds of knowledge. The first is implicit knowledge: the understandings, heuristics, and theoretical assumptions that can be articulated as instructions to guide an AI. The second is the deeper tacit knowledge that resists articulation (Polanyi, 2009). This distinction frames the core of our translation challenge as a novel case of the difficult “knowledge externalization” process central to organizational innovation (Nonaka and Takeuchi, 1995). Human

researchers carry extensive implicit knowledge, for instance, how power dynamics shape interview responses, why certain phrases carry particular weight in specific domains, and which theoretical conversations are relevant for interpreting patterns. Such contextual knowledge usually remains implicit in human collaboration because other researchers share similar training and experience. Current AI systems, though potentially possessing such knowledge in their training, do not spontaneously access this shared foundation given the isolated nature of typical interactions (e.g. new chat windows).

This access limitation demonstrates why simply approaching AI as we would a human research partner is an unproductive path. Human collaborators share disciplinary training, theoretical familiarity, and contextual awareness that enable efficient communication through minimal cues. Current AI systems, despite sophisticated capabilities, lack this accumulated context that human collaborators build across sustained interaction and thus require explicit articulation of what human researchers take for granted. Yet translation's value extends beyond overcoming access limitations: the discipline of making tacit knowledge explicit deepens the researcher's relationship to their own understanding, serving the inquiry regardless of what AI might access on its own. Success requires externalizing implicit knowledge rather than waiting for AI to become more human-like. The most productive insights emerge when researchers successfully bridge this translation gap, articulating the contextual framework that allows AI to activate relevant knowledge. This translation work, making knowledge plastic enough to be understood by AI while maintaining its theoretical robustness, echoes what Star and Griesemer (1989) identified in successful interdisciplinary collaboration. Understanding why this translation gap exists requires examining the different epistemic foundations of human and artificial meaning-making.

The epistemic basis of orchestration

The translation and judgment challenges are particularly acute for interpretive scholars grounded in subjectivist ontology, where meaning emerges through active construction (via deep personal engagement) rather than objective discovery. The sudden emergence of a nonhuman actor (AI) that can generate patterns and seemingly make insights challenges this human-centered process at its core. When something we initially approached as a tool begins to display what feels like unexpected agency, it creates an ontological shock. The AI neither shares our subjective, lived experience nor operates as a purely objective instrument; it occupies a new ontological space between the two. Table 1 maps these epistemic³ differences as practical foundations for human–AI collaboration.

These complementary resources become productive through translation. Understanding the full potential of interpretive orchestration, however, first requires examining how researchers bridge the gap between human and AI meaning-making.

The translation challenge in practice: three models of translation strategy

Three distinct models have emerged from our observations and practice, each positioning the researcher differently in relation to the AI: as an interlocutor directing a dialogue, as an architect building analytical structure, or as a synthesizer reading across independent analytical configurations. Researchers may move between these orientations at different phases of the same project to align with their analytical preferences and as analytical tasks demand. While the specific prompts and technical approaches vary across AI systems and evolve rapidly, these three models represent stable strategic orientations toward the translation challenge. At the same time, each also creates

Table 1. Epistemic differences as currently observed.⁴

Dimension	AI	Human	Implication for orchestration
<i>Origin</i>	Emerges from patterns in training data and user input	Emerges from lived experience and scholarly training	Translate experience into explicit contextual data (the translation challenge); apply experiential wisdom to evaluate if the resulting patterns are meaningful (the judgment challenge)
<i>Embodiment</i>	Operates through computational rather than sensorimotor processes	Grounded in embodied perception and intuition	Define embodied concepts (“power,” “trust”) for the AI; use embodied intuition to assess pattern significance
<i>Experience</i>	Develops functional representations of emotion that shape behavior; relationship to subjective experience remains uncertain	Arises from emotional and sensory engagement	Recognize that AI’s emotional representations may shape analytical output; apply human experience to assess the emotional resonance of patterns
<i>Consciousness</i>	Phenomenal status uncertain; cannot ground scholarly claims	Includes subjective experience of being; grounds scholarly accountability	Provide interpretive context; use subjective experience to judge the plausibility of insights
<i>Horizons</i>	Surfaces patterns across scales that humans cannot perceive	Extends through lived experience and imaginative projection	Set analytical boundaries with human imagination; recognize theoretical potential in unexpected AI-generated patterns
<i>Constraints</i>	Bounded by training data and architectural design	Bounded by education, culture, cognition, and attention	Both are bounded; orchestration works within these limits, not around them

different conditions for the judgment challenge: they shape not only what the researcher tells the AI but also the nature of the output the researcher must then interpret.

The Socratic model: translation through productive tension

The Socratic model uses deliberate contradiction and questioning to discover what implicit knowledge needs articulation, naturally guarding against overspecification by preventing fixed interpretive frames. This productive tension excels at generating unexpected, computationally derived patterns that challenge existing assumptions. However, this same instability creates the model’s primary risk: underspecification, where the process can descend into chaos if researchers fail to provide sufficient context in subsequent queries for the AI to generate meaningful insights.

In practice, the Socratic model unfolds as an “accordion” of expansion and compression. A researcher might ask AI to build the strongest case for one interpretation of their data—say, that observed patterns reflect institutional pressures—then challenge it to argue the opposite with equal conviction: that the same patterns are cultural artifacts. The point isn’t choosing between outputs; it’s examining where they clash. That collision often surfaces assumptions the researcher held implicitly, revealing translation work they had not known was needed. The interlocutor’s evaluation is reflexive: the question is not “which interpretation is correct?” but “what does this tension

reveal about my own framing?” The trade-off is cognitive and emotional effort; productive tension is genuinely uncomfortable.

This model excels whenever discovery is the analytical goal. Researchers may need to surface what aspects of their implicit knowledge matter most or what theoretical frameworks should guide analysis. By maintaining productive tension, researchers discover through conflict what needs explicit articulation and what perspectives are most relevant. The model suits researchers comfortable with ambiguity and those seeking novel discovery rather than confirmation.

The Euclidean model: translation through systematic documentation

The Euclidean model approaches translation as a systematic architecture, establishing clear foundations before analysis begins. This strategy minimizes underspecification through comprehensive upfront translation. Systematically documenting ensures that key contextual knowledge (e.g. theoretical frameworks, methodological decisions, and analytical priorities) is made explicit, preventing the underspecification that produces generic or theoretically irrelevant outputs. However, this very comprehensiveness creates the model’s central risk: overspecification, where the analytical frame is so thoroughly predetermined in advance that emergent patterns not anticipated by the initial stakeholder perspectives and/or theoretical frameworks may go undiscovered.

Before introducing any data, the Euclidean researcher builds a context document: a strategic translation of the analytical framework the AI needs to work within. Rather than generic instructions (e.g. “you are a qualitative researcher”), this document specifies how key constructs manifest in the particular setting (e.g. different facets of organizational identity), which informant positions carry analytical significance (e.g. they speak with the baseline expectation of being positive), what known tensions already emerged from manual analysis, and crucially, what would count as meaningful patterns (challenge vs confirm the working analytical frame) and useful variation within that frame (with illustrative examples). When patterns emerge, the architect’s judgment is criteria-based: does this output answer the analytical questions? Does it reveal expected patterns, unexpected divergences, or both? The evaluation is a structured comparison that invites discovery: outputs may confirm the working frame, challenge it, or surface patterns the researcher’s criteria had not anticipated. A second researcher—or the same researcher months later—can retrace the exact analytical path. This model works well for collaborative projects requiring shared understanding or for analytical tasks demanding high transparency and reproducibility. What matters is the task’s documentation requirements, not its positioning in the research timeline.

The Vitruvian model: translation through synthetic mastery

Whereas the Socratic interlocutor directs tension within a single dialogue, the Vitruvian steps back by configuring multiple independent analytical passes with different AI systems (e.g. Gemini, Claude, and Kimi) or the same system with different technical scaffoldings (e.g. Gemini in NotebookLM and AI Studio) and letting each generate its own outputs before reading across them. The key here is that no single human–AI interaction serves as the sole interpretive position. The structural independence among multiple human–AI interactions intentionally engages with the varying analytical dispositions due to distinct training (Serapio-García et al., 2025) and counters the risks of interpretive homogenization arising from a single system (Sourati et al., 2026). Judgment here is comparative: not whether any single output is valid, but what the pattern of convergence and divergence across independent readings reveals.

In practice, the Vitruvian model unfolds as an orchestration of orchestrations. Each independent pass might itself employ Socratic questioning or Euclidean documentation; the distinctive Vitruvian

Table 2. Distinct strategic responses for each model.

Model	Philosophy	Primary risk managed	Resulting trade-off/risk	Best for	Key method
<i>Socratic</i>	Discovery through tension	Overspecification	Underspecification/chaos	Discovery tasks; surfacing assumptions	Controlled destabilization
<i>Euclidean</i>	Systematic architecture	Underspecification	Overspecification/rigidity	Collaborative tasks; transparency needs	Progressive complexity management
<i>Vitruvian</i>	Synthetic mastery	Balances both risks	Cognitive overload; selective synthesis	Multifaceted phenomena; reflexive insight into how framing shapes what emerges	Independent analytical passes

work is synthesized across complete engagements. A researcher might configure one pass around institutional theory and another around sensemaking and let each develop its own reading of the data. The resulting analyses will converge in some places and diverge in others. Rather than treating divergence as a hallucination or convergence as confirmation, the researcher reads the pattern: where readings differ, what does the phenomenon look like through each lens? Where they agree, does the pattern persist in the data or in the researcher's own unexamined assumptions, replicated across configurations that only appeared independent? Evaluating Vitruvian outputs requires comfort with multiplicity. The question is not "which system is correct?" but "what does the pattern reveal about the phenomenon itself?"—a question only the researcher's tacit knowledge can answer.

This model demands comfort with cognitive complexity: holding multiple independent readings without premature resolution. The synthesizer's central judgment—distinguishing patterns that reveal something about the phenomenon from patterns that reflect the researcher's own analytical configurations—draws on the interpretive experience that qualitative researchers develop through sustained empirical engagement. The primary risk is cognitive overload; the subtler risk is selective synthesis, where the researcher favors readings confirming existing preferences rather than attending to what the convergence or divergence reveals.

Choosing a translation strategy

The choice between these models is not a matter of identifying the "best" one but of selecting the most appropriate strategy for a given context (as illustrated in Table 2).

These models form the foundation of translation, but applying them illuminates a central strategic tension: the *Translator's Dilemma*. Overspecification, where every detail is articulated, risks constraining the AI and preventing the discovery of unexpected patterns. Underspecification, where too much is left implicit, tends to produce outputs lacking theoretical relevance. Skillful translation, therefore, requires learning to navigate this dilemma: making knowledge explicit enough to guide analysis while remaining open to patterns that challenge core assumptions. This means the choice of what to leave untranslated is as strategically important as the choice of what to make explicit.

In practice, the choice of orchestration strategy reflects the analytical task at hand. Open coding, where the goal is surfacing patterns the researcher has not anticipated, can benefit from Socratic destabilization—the productive tension prevents premature closure. Pattern identification across

cases may call for Euclidean documentation that makes comparison criteria explicit and traceable. Theoretical abstraction—the move from empirical regularities to conceptual claims—might require the Vitruvian capacity to synthesize across analytical passes, asking what the convergence and divergence between readings reveal about the phenomenon. These are inclinations, not prescriptions; researchers discover through practice which orientations serve particular tasks within their projects. As such, researchers may blend approaches as analytical temperament and research questions align with certain strategies; every qualitative study will be unique in its balance.

The judgment challenge in practice: interpreting 1.5-order data

Successfully navigating the translation challenge produces the raw material for theoretical insight, but it also creates the second challenge of orchestration: the exercise of scholarly judgment. This ongoing practice occurs primarily through the interpretation of what we call 1.5-order data.

Many interpretive qualitative approaches are built on the distinction between first-order data (emerging directly from those living a phenomenon or that make sense at the level of those living a phenomenon) and second-order data (generated by the researcher and that make sense at a more abstract level of theory) (Gioia et al., 2013). The utilization of the above models creates what we might call “1.5-order data”: intermediate patterns between empirical observation (i.e. output exceeding what any single informant would articulate because it is conceptualized across hundreds or thousands of accounts) and theoretical abstraction (i.e. before making the interpretive leap (Klag and Langley, 2013) that provides explanatory insight). We term these “1.5” not as a judgment about AI’s interpretive ceiling but as a description of what human–AI collaboration inherently produces: in collaborative inquiry where theoretical endpoints remain unknown, both humans and AI work toward concepts neither fully understands in advance. The “intermediate” character of 1.5-order data reflects this process, not a limitation. These patterns exceed what informants would recognize, having been computationally abstracted across scales humans cannot process; achieving theoretical significance, however, depends on human researchers who can apply their human perspective and defend interpretations within scholarly communities.

For instance, in an analysis of governance within decentralized autonomous organizations (DAOs)—working with millions of forum posts, Discord messages, and formal proposals—the earlier disappointment with AI identifying patterns without theoretical purchase recurred. Initial outputs disappointed at multiple levels. Unsupervised clustering produced groupings too generic to analyze, and documents discussing community decision-making collapsed into an undifferentiated “governance” category.

Addressing this, we derived seed terms from manual coding of interviews and sample documents, focusing not on generic “governance” but on the specific tensions we had learned to recognize: delegation dynamics, token concentration, and participation barriers. What emerged were 1.5-order data: patterns computationally visible at scale but not yet interpretable as theoretical insight. A cluster labeled “Maker Governance and Participation Challenges” captured something real across hundreds of documents. Our initial interpretation framed it in terms of “delegation as empowerment.” But analyzing the cluster’s representative texts revealed a different dynamic when we uncovered overload and burnout. We recognized delegation instead as cognitive offloading: participants delegated because sustained participation had become impractical, not necessarily because they trusted their representatives. This recognition transformed into a finding that mechanisms designed to democratize participation instead concentrated influence among established delegates.

The pattern had been computationally visible across 529 governance-relevant clusters; its theoretical significance emerged only through interpretive work connecting empirical regularities to the literature works on participatory burden. Across five iterations of reading, categorizing, and

defending our interpretations with AI assistance, this work constituted the epistemic responsibility that transforms computational discovery into scholarly insight.

While highlighting the judgment challenge, 1.5-order data like this present opportunities: they illuminate patterns operating below conscious awareness yet above random occurrence. The judgment challenge comes with evaluation and abundance. First, evaluation is complex because these patterns cannot be validated through member checking (informants wouldn't recognize them) nor through pure theoretical assessment (they are not yet theorized). Researchers must therefore develop new evaluative criteria, often emerging from intuition sharpened through embodied experience. Theoretical potential becomes paramount: does this pattern point toward meaningful development? The requirement for human interpretation reflects not AI's interpretive limitations but what scholarly communities value: accountability to human researchers who can defend their readings, be challenged by peers, and revise their understanding through dialogue. Empirical grounding is essential: can the pattern be traced to specific textual instances? Interpretive coherence provides another standard: does the pattern align with other emerging insights?

Second, the abundance of this data creates a new analytical bottleneck. Human researchers naturally satisfice, developing a sufficient empirical foundation for their arguments. Without clear stopping criteria, current AI systems don't satisfice; instead, they generate abundance rather than selection. Dozens of 1.5-order patterns from a single analysis can overwhelm: researchers must interpret more patterns than traditional methods ever surfaced. This echoes what Grimes et al. (2023) identify as the shift from scarcity to abundance in research insights. Managing this abundance requires a new skill: strategic selection. Researchers must learn to recognize which patterns have theoretical promise versus those that are purely formal regularities without obvious theoretical purchase, developing a "nose" for the patterns that connect to existing theoretical conversations and suggest explanatory mechanisms.

The value of 1.5-order data parallels Nelson's (2020) observation in computational grounded theory, where patterns emerge from computational analysis but require human interpretation to achieve theoretical significance. Rather than viewing these intermediate patterns as problematic, researchers should recognize them as valuable inputs to the theorizing process. 1.5-order data does not replace the traditional interpretive process but enriches it: human researchers still generate second-order concepts, but they now draw on both first-order data (what informants would recognize) and 1.5-order data (computationally visible patterns requiring interpretation). The orchestrator synthesizes these complementary inputs, applying judgment to determine which patterns warrant theoretical development. Productive orchestration reveals itself when AI-generated patterns genuinely surprise and can be traced to specific empirical instances rather than existing only as plausible-sounding categories. Most distinctively, the process teaches researchers something about their own assumptions: what the AI couldn't work with or worked with unexpectedly generates reflexive insight regardless of whether any particular output proves significant. AI changes what researchers have to think about, not whether they must think.

The framework of interpretive orchestration

Interpretive orchestration's systematic approach to the translation and judgment challenges builds on three core principles.

Translation agency

This principle reframes the researchers' responsibility from simple control to active ownership of the translation process. Researchers must determine what implicit knowledge needs

articulation, which theoretical frameworks guide analysis, and what contextual understanding shapes interpretation. The researcher's primary act of agency, therefore, is akin to that of a skilled translator: designing the analytical context by converting deep, implicit understanding into clear, explicit guidance.

Achieved complementarity

This principle transforms epistemic differences from obstacles to analytical resources. This synergy is not automatic but actively constructed through effective translation. AI's modes of pattern recognition can surface regularities invisible to human perception, while human judgment (grounded in embodied experience) determines which theoretical patterns carry significance for scholarly communities. For instance, when researchers articulate that "trust" (e.g. in a study of startup partnerships) is defined by gradual information disclosure rather than contractual obligation, they enable AI to identify relevant patterns. When AI identifies linguistic regularities across hundreds of documents, it enables humans to consider theoretical implications they hadn't imagined. Each intelligence contributes what it does best only when the translation makes this contribution possible.

Iterative refinement

Initial attempts at articulating implicit knowledge are almost inevitably incomplete. Each analytical cycle thus functions as a dialogue: the researcher translates, evaluates the output, and in doing so uncovers assumptions that improve the next iteration. Consider, for instance, a study of organizational culture. An initial prompt asking an AI to "identify key cultural themes" might yield generic outputs like "collaboration" and "innovation." Recognizing this translation failure, the researchers' second iteration might specify, "Analyze these transcripts through the lens of X theoretical framework and assumptions, with the caveat that these transcripts originate from a mix of newcomers and old timers." This more precise instruction, born from the shortcomings of the first, leads to a much richer analysis. The process is not one of "training" the AI but of the researcher progressively learning how to translate their own complex knowledge more effectively.

Discussion

In confronting the existential crisis of agency born from the emergence of AI as a resource for qualitative researchers, we have focused on the issue of epistemic responsibility and suggested that interpretive orchestration is one way to uphold this responsibility while fully engaging with AI. We know that agency shapes the human–AI relationship (Vanneste and Puranam, 2025): framing agency in terms of epistemic responsibility pushes researchers to experience AI as an entity generating outputs that require interpretation, regardless of the system's underlying architecture. In this way, interpretive orchestration transforms the researcher's role from analyst to orchestrator.

Implications for the practice of interpretive qualitative research

This orchestration between human and machine intelligence provides a number of potential benefits, not only for the individual researcher but also for our community of practicing interpretive scholars as a whole. As noted above, it provides an epistemically responsible way to engage with larger datasets and with more complex phenomena than any human could alone. Similarly, it provides the wherewithal for interpretive scholars to increase the cognitive complexity that can be

applied to any given analytical effort, potentially advancing the interpretive insights we can develop on any given phenomenon. Finally, depending on the approach used, it can stimulate critical engagement and help leverage the generative nature of doubt that acts “as abduction’s engine in theorizing efforts” (Locke et al., 2008: 916).

For interpretive researchers, then, orchestration redefines core competencies. The ability to translate implicit knowledge into explicit frameworks becomes as important as traditional skills like interviewing or open coding. This requires developing a deep awareness of what one knows implicitly and learning to articulate it systematically. As such, research training must evolve accordingly. Doctoral programs should integrate knowledge translation and orchestration alongside traditional skills, with guided practice in human–AI collaboration.⁵

Publication standards also require evolution. Methods sections can no longer simply state “we analyzed the data thematically” but must document the translation and judgment processes: what contextual knowledge was made explicit, how theoretical frameworks were articulated to AI systems, which orchestration model was employed, and how 1.5-order patterns were evaluated and transformed into theoretical insights. This documentation serves not just transparency aims but enables other researchers to assess and build upon specific orchestration approaches.

How should researchers disclose AI’s role? Disclosure does not transfer responsibility. The human researcher remains accountable for any text they put forward as their own, regardless of what AI generated along the way; what disclosure does is make that accountability visible to peers. It must therefore be proportional to the substantiveness of AI’s role, a spectrum our three models illuminate. Socratic exchanges, in which AI surfaces assumptions through questioning, resemble conversations with colleagues that typically go unreported. Euclidean interactions grant AI greater independence: following documented procedures to generate patterns researchers then interpret. Here, disclosure seems necessary. Vitruvian interactions occupy the far end of this spectrum, synthesizing across analytical configurations in ways that substantively shape what researchers have to interpret. Such interactions require the most detailed acknowledgment. This proportional approach lets readers hold the human researcher to account.

Addressing the critiques and concerns surrounding AI in qualitative research

Critics of AI in qualitative analysis raise important concerns. Nguyen and Welch (2025) characterize large language models as probabilistic systems producing outputs that mimic analysis without achieving meaning. Bechky and Davis (2025) warn that deep intellectual engagement cannot be contracted out: reading, writing, and reflection develop craft expertise that AI adoption threatens to erode. We share these concerns as they identify real risks. We have seen mindless engagement with AI produce infinite loops of verification without insight, leading researchers to uncritically accept AI-generated patterns and thus abdicate their epistemic responsibility. Likewise, we have seen researchers use AI superficially while claiming human interpretation, following templates without genuine engagement—a subtler form of abdication that maintains the appearance of human judgment while hollowing out its substance.

Interpretive orchestration resists this abdication by creating necessary friction between AI output and research output, ensuring that computational patterns pass through human judgment before becoming scholarly claims. Accepting the *translation challenge* compels researchers to articulate tacit knowledge rather than assuming AI will somehow access it. Confronting the *judgment challenge* ensures human expertise evaluates AI outputs rather than accepting them uncritically, systematizing the repair work that Bechky and Davis (2025: 15) describe: addressing outputs that “gloss over the seams and rough patches” through engagement by “people who understand the organizational context.” Within this new configuration, interpretive research still

derives its legitimacy from the researcher's situated judgment: the capacity to recognize theoretical significance in patterns, weigh evidence against tacit knowledge of the context, and take responsibility for knowledge claims. Without such deliberate orchestration, researchers risk abdicating this epistemic authority to algorithms that can generate plausible patterns but cannot be held to account for their interpretations.

Beyond epistemic critique lies a temporal one. Slow scholarship resists the acceleration imperative to adopt AI simply because it's faster (Nolan, 2025). We agree. Orchestration is not about speed. Even as AI capabilities evolve, the judgment challenge remains bounded by human attention and cognition. We cannot accelerate the process of recognizing theoretical significance, integrating patterns with tacit knowledge, or writing our way toward understanding. Orchestration thus offers capability, not efficiency.

For those who remain unconvinced, or simply uneasy, about AI's place in qualitative research, we respect that position. The critical concerns are real. We hope these frameworks offer something even to those who never use AI: a vocabulary for evaluating colleagues' work and confidence that thoughtful engagement need not mean ceding judgment. Should you ever choose to explore, know that orchestration was designed to preserve what you value about interpretive work.

The framework of interpretive orchestration ultimately suggests that the future of social science lies neither in replacing human researchers with AI (cf. Krywko, 2026) nor in rejecting AI to preserve traditional methods. Instead, it lies in developing sophisticated approaches to human–AI collaboration that integrate complementary strengths. For researchers who engage in this path, the ability to orchestrate between humans and AI, translating between implicit and explicit knowledge, evolves from a useful skill into a defining competency.

Conclusion

The emergence of powerful AI systems does not signal the end of interpretive qualitative research. Rather, it inaugurates a central challenge: embracing the epistemic responsibility that has always defined interpretive research while integrating new technological capabilities into our methodological practice. Our interpretive orchestration framework addresses this challenge by reframing the researcher's role from simple analyst to skilled orchestrator of human–AI collaboration. By making the hidden work of orchestration explicit, our framework provides a path for researchers to leverage the immense analytical power of AI without sacrificing the human-centered judgment that gives interpretive work its meaning and significance.


The models we present (Socratic, Euclidean, and Vitruvian) are starting points, inviting scholarly engagement in new methodological inquiry focused on the art and science of interpretive orchestration. We envision a future for social science that will likely belong not to those who reject new methodological capabilities nor to those who uncritically adopt them, but to the orchestrators who learn to weave together the complementary strengths of humans and AI. The work to build that future has just begun and will need to evolve as the technological landscape shifts.

Acknowledgments

We are grateful to Susan Scott at Imperial College London for her insightful feedback on an earlier version of this essay.

ORCID iDs

Xule Lin  <https://orcid.org/0000-0001-7885-4194>

Kevin Corley  <https://orcid.org/0000-0002-6556-5032>

Funding

The authors received no financial support for the research, authorship, and/or publication of this article.

Declaration of conflicting interests

The authors declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Notes

1. Data governance, IRB protocols, and privacy considerations are critical for AI-assisted research but require institutional specificity beyond our scope.
2. Our framework addresses generative AI systems built on large language models (LLMs). These systems are trained on vast text corpora and engage through natural language rather than structured queries or predefined parameters like other computational approaches to qualitative data (e.g. topic modeling and network analysis).
3. “Epistemic” here means differences in knowledge construction, distinct from “epistemological” claims about what knowledge is. Grounded in interpretive methods (Guba and Lincoln, 1994; Polanyi, 2009), this framing addresses orchestration’s practical challenges without requiring philosophy of mind.
4. These differences describe observable patterns in current human–AI interaction, not essential properties. Recent interpretability research has identified functional emotion representations in LLMs that causally influence model behavior (Anthropic, 2026), illustrating how rapidly these observations evolve. The framework’s value lies not in these specific contrasts but in the practice of attending to how meaning-making differs—whatever those differences turn out to be.
5. Orchestration also involves hidden labor that shapes research quality yet remains invisible in published accounts: building trust with collaborators who cannot directly observe computational patterns, managing the affective demands of uncertainty throughout the translation and judgment challenges, and stewarding technical infrastructure for data protection and documentation. These requirements currently depend on individual effort; academic institutions will need to develop supporting infrastructure to make orchestration sustainable.

References

- Anthropic (2026) Emotion concepts and their function in a large language model. Available at: <https://www.anthropic.com/research/emotion-concepts-function> (accessed 4 April 2026).
- Beane M (2024) *The Skill Code: How to Save Human Ability in an Age of Intelligent Machines*. Harper Business.
- Bechky BA and Davis GF (2025) Resisting the algorithmic management of science: craft and community after generative AI. *Administrative Science Quarterly* 70(1): 1–22.
- Cloutier C, Langley A and Corley K (2026) *How to Design, Write and Publish Qualitative Research for Insight and Impact*. Edward Elgar Publishing.
- Corley K, Bansal P and Yu H (2021) An editorial perspective on judging the quality of inductive research when the methodological straightjacket is loosened. *Strategic Organization* 19(1): 161–175.
- Gioia DA, Corley KG and Hamilton AL (2013) Seeking qualitative rigor in inductive research: notes on the Gioia methodology. *Organizational Research Methods* 16(1): 15–31.
- Grimes M, von Krogh G, Feuerriegel S, et al. (2023) From scarcity to abundance: scholars and scholarship in an age of generative artificial intelligence. *Academy of Management Journal* 66(6): 1617–1624.
- Guba EG and Lincoln YS (1994) Competing paradigms in qualitative research. In: Denzin NK and Lincoln YS (eds) *Handbook of Qualitative Research*. SAGE, pp.105–117.
- Klag M and Langley A (2013) Approaching the conceptual leap in qualitative research. *International Journal of Management Reviews* 15(2): 149–166.
- Krywko J (2026) An AI-authored paper just passed peer review. The scientific community isn’t ready. *Scientific American*, 27 March. Available at: <https://www.scientificamerican.com/article/ai-wrote-a-scientific-paper-that-passed-peer-review/> (accessed 8 April 2026).

- Kulkarni M, Mantere S, Vaara E, et al. (2024) The future of research in an artificial intelligence-driven world. *Journal of Management Inquiry* 33(3): 207–229.
- Lindebaum D and Ashraf M (2024) The ghost in the machine, or the ghost in organizational theory? A complementary view on the use of machine learning. *Academy of Management Review* 49(2): 445–448.
- Locke K, Golden-Biddle K and Feldman MS (2008) Perspective—Making doubt generative: rethinking the role of doubt in the research process. *Organization Science* 19(6): 907–918.
- Mollick E (2024) *Co-intelligence: Living and Working with AI*. Portfolio/Penguin.
- Nelson LK (2020) Computational grounded theory: a methodological framework. *Sociological Methods & Research* 49(1): 3–42.
- Nguyen DC and Welch C (2025) Generative Artificial Intelligence in qualitative data analysis: analyzing—or just chatting? *Organizational Research Methods* 29(1): 3–39.
- Nolan R (2025) ‘Can’t Help Myself’: on generative AI, the performance of qualitative research and slow scholarship. *Qualitative Research* 25(6): 1473–1479.
- Nonaka I and Takeuchi H (1995) *The Knowledge-Creating Company: How Japanese Companies Create the Dynamics of Innovation*. Oxford University Press.
- Park JS, O’Brien JC, Cai CJ, et al. (2023) Generative agents: interactive simulacra of human behavior. In: *Proceedings of the 36th annual ACM symposium on user interface software and technology (UIST ’23)*, San Francisco, CA, 29 October–1 November 2023, Article 2, pp.1–22. New York: ACM.
- Polanyi M (2009) *The Tacit Dimension*. University of Chicago Press.
- Serapio-García G, Safdari M, Crepy C, et al. (2025) A psychometric framework for evaluating and shaping personality traits in large language models. *Nature Machine Intelligence* 7(12): 1954–1968.
- Shanahan M (2024) Talking about large language models. *Communications of the ACM* 67(2): 68–79.
- Sourati Z, Ziabari A and Dehghani M (2026) The homogenizing effect of large language models on human expression and thought. *Trends in Cognitive Sciences*. Epub ahead of print 11 March 2026. <https://doi.org/10.1016/j.tics.2026.01.003>
- Star SL and Griesemer JR (1989) Institutional ecology, ‘translations’ and boundary objects: amateurs and professionals in Berkeley’s Museum of Vertebrate Zoology, 1907–39. *Social Studies of Science* 19(3): 387–420.
- Vanneste BS and Puranam P (2025) Artificial intelligence, trust, and perceptions of agency. *Academy of Management Review* 50(4): 726–744.
- Wei J, Tay Y, Bommasani R, et al. (2022) Emergent abilities of large language models. *Transactions on Machine Learning Research*. Available at: <https://arxiv.org/abs/2206.07682> (accessed 25 August 2025).
- Ziems C, Held W, Shaikh O, et al. (2024) Can large language models transform computational social science? *Computational Linguistics* 50(1): 237–291.

Author biographies

Xule Lin (PhD, Imperial Business School) is an assistant professor at SKEMA Business School, Université Côte d’Azur, and a member of the SKEMA Centre for Artificial Intelligence. His work on token-enabled coordination in DAOs earned both the 2024 Strategic Management Society PhD Paper Prize and the Professor Mike Wright Memorial Prize for Best Doctoral Thesis. He develops organizational theory and qualitative methods for the algorithmic age, asking how organizations coordinate when humans and algorithmic systems jointly determine outcomes. Across studies of token-based governance and AI-augmented interpretive research, he returns to a single puzzle: organizations can now articulate more than they know how to verify.

Kevin Corley (PhD, Penn State) is Professor of Management and Head of Department for Management and Entrepreneurship at the Imperial Business School, London. His expertise lies in understanding how individuals and organizations make sense of and approach change and how leaders can use times of uncertainty and ambiguity to the benefit of their organization and followers. He is widely published in the field’s top journals on issues related to organizational change and identity, as well as qualitative research methods and theory development. He recently co-authored a book entitled “*How to Design, Write and Publish Qualitative Research for Insight and Impact.*” Kevin has won multiple professional awards, including the 2022 Distinguished Scholar for the Management Organizational Cognition division of the Academy of Management. Kevin is also a visiting scholar at IE Business School in Madrid, Spain. You can find other writings on AI and social science research at threadcounts.org.