

Human Interpretive Presence in AI-Assisted Qualitative Research: A Reproducible Workflow for Social and Emotional Learning (SEL)

*Anat Korem [0000-0001-6902-8675]^{1,2} and Sigal Chen [0000-0002-1746-6371]¹

** Equal Contribution*

¹ Levinsky-Wingate Academic College, 15 Shoshana Persitz St., Tel Aviv 69387, Israel.

² The Hebrew University of Jerusalem.

Anat.korem@l-w.ac.il ; Sigalc@l-w.ac.il

Abstract

AI and computational informatics are reshaping qualitative research, giving scholars new tools for automated coding, theme extraction, and large-scale text analysis. However, applying these tools to a domain like social and emotional learning (SEL), where meaning is bound up with belonging, vulnerability, and relationship, carries a risk: Qualitative inquiry gets flattened into mechanical data classification.

This paper introduces the principle of human interpretive presence (PHIP), a conceptual and operational framework for AI-assisted qualitative informatics. Drawing on educational epistemology and the psychology of socio-emotional experience, PHIP treats interpretation as an active, embodied, ethical encounter rather than as probabilistic pattern recognition. Case analyses involving subtext, irony, and conversational silence illustrate where standalone algorithmic text processing tends to fail when the material carries socio-emotional weight.

In keeping with AII 2026's focus on reproducibility, we translate PHIP into operational guidance. Our human-in-the-loop model is built on three pillars: contextual resonance, affective reflexivity, and ethical hermeneutics. The model also incorporates a reflexive audit trail for logging the interactions, prompts, and corrections that pass between researcher and system. Thus, the result repositions AI as an assistive collaborator rather than an independent analyst, giving the field a transparent, reproducible workflow that retains the interpretive depth that human narratives require.

Keywords: Qualitative informatics, principle of human interpretive presence (PHIP), social and emotional learning (SEL), methodological reproducibility.

1 Introduction

AI has moved quickly into qualitative research practice, introducing computational systems that can code, identify themes, summarize, and detect patterns at a scale no human team could match. That speed creates a genuine tension once these tools reach a domain like social and emotional learning (SEL). SEL is often defined as the process by which people build the knowledge, skills, and attitudes needed to form healthy identities and manage their emotions [1]. However, researchers increasingly stress that these processes are not neutral; they are relational and shaped by culture, and reading them well depends on perception, subjective experience, and intuition in ways a coding scheme cannot fully specify [2].

This tension is where this paper is positioned. Interpreting SEL data draws on lived experience of belonging, vulnerability, identity, and relationship, dimensions that resist reduction to discrete classifiable units. Therefore, the more pressing question for applied informatics is not what an automated system can do, but what human meaning is lost when delegating interpretation to automated classification.

2 Research Problem

The problem we address is methodological: How do you bring AI into a qualitative research pipeline without flattening the context-dependent meaning of socio-emotional material, and how do you hold onto interpretive depth while still meeting the field's growing demand for methodological transparency and reproducibility?

3 Aim and Contribution

This paper introduces the principle of human interpretive presence (PHIP). This principle offers an epistemic and operational framework for AI-assisted qualitative informatics that gives researchers a concrete way to keep human judgment central to the analysis while still working with computational tools at scale.

Qualitative research has never aimed only to describe what people say—its mission is to understand it. That distinction is critical. Interpretation of qualitative data is not a technical coding procedure; rather, it is a back-and-forth between a participant's narrative, its surrounding context, the researcher's theoretical lens, and the latter's own reflexive judgment. Identifying ambiguity, contradiction, silence, and emotional nuance is precisely the kind of work that standard informatics systems tend to flatten. Thus, any human-AI research design in this space has to account for that directly.

4 Conceptual-Methodological Approach

This reflexive, conceptual paper is built on the authors' extensive qualitative research experience in SEL. Rather than report findings from a single empirical study, it develops the PHIP framework through structured, critical reflection on the interpretive challenges that emerge when AI enters a qualitative computational workflow. Using case-based analysis of realistic methodological scenarios, we suggest practical guidance for bridging human thematic depth and automated data processing.

5 Theoretical Framework

Bringing AI into qualitative inquiry requires a clear theoretical line between computational data processing (e.g., pattern recognition) and the hermeneutic work SEL research actually demands. PHIP rests on two theoretical foundations: educational epistemology and the phenomenology of social-emotional experience. Together they frame PHIP as an operational structure for keeping qualitative informatics human-centered rather than as a constraint on automation.

5.1 Educational Epistemology and Situated Judgment

Educational epistemology holds that knowledge in pedagogical settings is co-constructed through situated judgment rather than simply discovering it [3]. An SEL researcher is an active participant in meaning-making, and their subjective sensitivity serves as an epistemic asset rather than a bias to be engineered away. That interpretive work depends on navigating ambiguity, contradiction, and the nuances of silence. These variables are complex and non-linear, thus remaining out of reach for deterministic algorithmic logic [4].

5.2 The Phenomenology and Psychology of Relational Space

Educational psychology points to a similar predicament. Social and emotional phenomena (e.g., identity negotiation, belonging, vulnerability) are deeply contextual [5]. SEL research supports this empirically, showing that social-emotional perception is embedded in cultural, linguistic, and relational context, meaning that these experiences need to be understood within the environment that produced them [6].

Much of the qualitative methodology literature locates emotional nuance specifically in the “relational space” between participant and researcher that computational tools are not built to inhabit [7]. For qualitative informatics to stay true to its purpose (i.e., capturing the complexity of human experience), it must resist the flattening effect that accompanies full automation, where layered narratives collapse into rigid data points.

5.3 Ethical Reflexivity in Human-AI Collaboration

Ethical reflexivity sits at the center of what “human touch” means in qualitative work [8]. It involves continually asking how one’s own perspective shapes the reading of someone else’s lived experience. Ethical reflexivity is a sustained, watchful kind of engagement, well beyond what pattern-recognition machine learning can offer. Thus, in an AI-assisted workflow, the researcher remains the final arbiter of meaning. This, then, becomes a structured human-in-the-loop model in which AI handles data organization and scaling, while the researcher keeps the job of building epistemic understanding [9].

Under PHIP, the researcher becomes the critical bridge between raw computational output and the context-bound reality of SEL. Treating interpretation as both a moral and intellectual responsibility lets AI-assisted research honor the dignity and complexity of the people it studies [10].

6 Illustrative Examples and Case Analysis

The need for PHIP becomes clear when we consider how automated systems handle subtext, pragmatics, and non-verbal cues. Language models are good at token frequency and statistical pattern-matching, but they tend to fail badly at what might be called the rhetoric of the unsaid.

Three scenarios below illustrate how standalone AI breaks down in SEL informatics, followed by a worked case.

6.1 The Weight of Silence

During a focus group on school bullying, one adolescent pauses for ten seconds after being asked how she feels about belonging. A standard NLP pipeline, tuned for semantic density or clean transcription, is likely to treat that silence as dead air, a gap in the data to be discarded. A human researcher will read it differently—as a communicative act in its own right, perhaps trauma, perhaps resistance, perhaps hesitation. Thus, this researcher will treat this information as a genuine data point rather than noise.

6.2 Irony and Social Masking

Teenagers often use irony to mask social anxiety. “Oh sure, being ignored is just great” is a real example. A sentiment-analysis tool working off lexical weights would likely score that as positive, because of the word “great.” A human reader will catch the gap between the literal words and the actual emotional state immediately, and read the irony correctly.

6.3 Emotional Ambiguity

A phrase like “we are so close” can describe healthy intimacy in one narrative and suffocating codependency in another. Language models built on standardized semantic embeddings tend to collapse both readings into the same generic vector space. A human researcher, grounded in the fuller context, can distinguish between the two, thus preventing a complex emotional life from being reduced to a static label.

6.4 Case Study Analysis

This statement came from an adolescent in an SEL study: “Sometimes I join my friends only because I don’t want them to think I’m strange.” An automated classifier would likely classify the statement under a generic label, such as “socializing” or “peer interaction,” because of the surface reference to joining friends. However, the real weight of the statement sits in the subtext: forced conformity, fear of being seen as different. Without a human-in-the-loop to recalibrate what the AI extracts, the computational output misses the vulnerability entirely, which summarizes the argument for keeping human interpretive presence in the pipeline.

7 Core Conceptual Shifts in AI-Assisted Qualitative Informatics

Bringing applied intelligence into qualitative inquiry sets off a bigger epistemological shift than initially evident. Whereas AI is built for pattern recognition and statistical similarity, qualitative researchers are after contextual meaning-making, a clearly different task. Much of the current debate treats AI mainly as a tool to cut human bias and boost objectivity. This real contribution, however, carries a quiet assumption worth challenging—that human interpretation is mostly a limitation to correct for, rather than a primary source of understanding in its own right.

We would argue the opposite. Human interpretation is not noise that needs algorithmic cleanup; it is a necessary condition for understanding social and emotional phenomena that rarely sort neatly into predefined categories. In SEL, for instance, silence often carries more meaning than explicit language, contradictions mark developmental process rather than analytical error, and ambiguity is itself a finding rather than a flaw. Moving from manual to automated analysis brings three shifts worth naming directly.

7.1 From Interpretation to Categorization.

AI's reliance on automated classification and clustering underscores the risk that contextual nuance, contradiction, and complexity will be flattened into fixed categorical labels by systems ill-equipped to accommodate them. Returning to the case above ("Sometimes I join my friends..."), a standalone AI system might correctly tag it under peer relationships, social participation, or fear of rejection. A human researcher's interpretation will reveal a much richer layer underneath—vulnerability, identity negotiation, social anxiety, the tension between authenticity and belonging. While the AI classification is not wrong, mechanical categorization alone cannot hold the layered structure of the experience.

7.2 From Dialogue to Computation. Qualitative knowledge has traditionally emerged from an ongoing, iterative dialogue between a participant's narrative, the researcher, theory, and context. Incorporating AI in the analysis changes that; AI mediates the process through computational representations (vector embeddings, algorithmic weights) that reorganize how meaning gets produced, trading dialogical reflexivity for deterministic calculation.

7.3 The Reconfiguration of the Researcher's Role. Researchers risk drifting from active interpreters to mere supervisors of an automated pipeline. Thus, keeping human interpretive engagement present throughout the process is critical to maintaining the research's epistemological validity.

8 The Principle of Human Interpretive Presence (PHIP)

Rather than framing AI and qualitative inquiry as a binary—human interpretation versus machine interpretation—we shift to ask what conditions make meaningful qualitative understanding possible. PHIP formalizes this process, providing a core conceptual framework for AI-assisted qualitative research.

PHIP claims that interpreting complex social and emotional material requires more than accurate coding or sophisticated pattern recognition; it needs the researcher's ongoing epistemic, relational, and ethical engagement across the whole analytic pipeline.

Human interpretive presence refers to the human capacity to jointly hold contextual understanding, emotional nuance, linguistic ambiguity, reflexive judgment, ethical responsibility, and theoretical reasoning while deciphering meaning out of someone's lived experience. Indeed, this is a condition for carrying out this kind of research responsibly.

SEL illustrates this condition. We cannot understand belonging, empathy, trust, and identity apart from the relational dynamics in which they transpire. Their meaning rests as much on hesitation, silence, and contradiction as on anything stated explicitly. From that angle, AI should not be replacing qualitative interpretation. AI works best as a computational partner that handles scaling and organization, while the harder job of building meaning remains in human hands.

9 The Core Pillars of PHIP: Formalizing the Framework

PHIP holds that the epistemic quality of qualitative inquiry in SEL is directly tied to how present the researcher stays within the data pipeline. Thus, interpretation is an active, embodied, ethical encounter between researcher and participant, not simply mechanical theme extraction or algorithmic labeling. Whereas AI systems work through probabilistic pattern recognition and vector proximity, PHIP-driven qualitative informatics rests on three interlocking pillars (see Fig. 1).

9.1 Contextual Resonance

This pillar recognizes that emotional expressions (e.g., silence, irony) are not isolated data points; they sit inside specific temporal, social, and cultural structures. The human researcher functions as a contextual anchor here, grounding what would otherwise be a fleeting textual token in the participant's broader biography and circumstances and going beyond what context-window optimization in current large language models (LLMs) can do.

9.2 Affective Reflexivity

PHIP asks researchers to notice and examine their own emotional response to the data as part of the analysis. In SEL work, that inter-subjective space—the resonance a researcher feels while sitting with a narrative of vulnerability—often reveals a deeper understanding than the transcript alone would give. As AI systems have no subjective experience to draw on here, this piece of the work stays with the human researcher.

9.3 Ethical Hermeneutics

Interpreting someone's emotions is a normative act, not a neutral classification task. Calling an utterance "resilience" instead of "compliance" carries real moral weight. PHIP holds the researcher as the final arbiter of meaning, accountable for what the findings imply — something an agentless computational tool structurally cannot be.

None of this amounts to rejecting AI in qualitative informatics. PHIP recalibrates the AI system's role, from independent "analyst" down to assistive computational collaborator. Under this architecture, AI handles structural complexity and macro-level pattern detection; the human researcher does the interpretive work.




Fig. 1. The Core Pillars of the Principle of Human Interpretive Presence (PHIP) Supporting Epistemic Validity in AI-Assisted Qualitative Research. *Contextual Resonance* grounds interpretation in broader social and contextual structures; *Affective Reflexivity* treats the researcher’s emotional response as a gateway to deeper understanding; *Ethical Hermeneutics* positions the human researcher as the ultimate ethical and interpretive arbiter.

10 Practical Implications for AI-Assisted Qualitative Informatics

Moving toward AI-augmented qualitative inquiry demands rethinking how the research process operates routinely. PHIP does not mean rejecting digital tools; it means re-engineering how researcher and system interact, across three practical areas.

10.1 Establishing a Reflexive Audit Trail. Meeting the reproducibility standard required for applied informatics means researchers working with AI need a structured, transparent audit trail (Table 1) — logging AI’s output (automated theme extraction, vector clustering) as well as documenting the specific points where the researcher accepted, rejected, or modified an AI suggestion, and why [3]. That transparency keeps the human interpretive loop visible and accountable rather than invisible inside a black box [4].

Table 1. PHIP Reflexive Audit Trail: A One-Page Template for Reproducible AI-Assisted Qualitative Analysis

Audit element	What should be documented	Responsibility
---------------	---------------------------	----------------

Data segment ID	Anonymized identifier of the excerpt or unit being analyzed	Researcher
Analytic task	Coding, categorization, theme generation, comparison, or summarization	Researcher
AI system/model	Tool and model/version used, when available	Documented
Prompt version	Exact prompt and sequential version (e.g., P1, P2, P3)	Documented
AI output	Original AI-generated code, theme, interpretation, or summary	AI-assisted
Human-led decision	Accept, modify, reject, merge, or rename the AI suggestion	Human-led
Interpretive rationale	Contextual, theoretical, affective, or ethical justification	Human-led
PHIP dimension	Contextual Resonance, Affective Reflexivity, and/or Ethical Hermeneutics	Human-led
Revision	Changes made following human review	Human-led
Final interpretation	Final researcher-approved code, theme, or interpretation	Human-led
Archive record	Date, researcher initials, prompt/output version, file name, and storage location	Researcher

This template (Table 1) supports methodological reproducibility by preserving a transparent, versioned record of the pathway from the initial AI output to the final human-approved interpretation. Human-led decisions include accepting, rejecting, modifying, naming, and contextually or ethically interpreting codes and themes.

10.2 Developing Interpretive Literacy in AI Prompting

Training the next generation of scholars needs to place “interpretive literacy” on equal footing with prompt engineering. Besides learning how to prompt an LLM well, researchers also need to know where algorithmic sentiment analysis breaks down, especially with non-linear SEL phenomena like emotional regulation or identity development [7]. Practically, this distinction means treating AI as a junior partner for lower-level work (e.g., brute-force coding, structural sorting, macro-pattern detection) while reserving higher-order hermeneutic work for the human scholar [9].

10.3 Algorithmic Governance and Ethical Gatekeeping

Bringing AI into qualitative data work also means taking ethical responsibility seriously. Foundation models are trained on large, homogenized datasets, which means they carry statistical biases that often miss the specific, idiosyncratic texture of marginalized or under-represented populations [8]. Under PHIP, the researcher acts as an ethical gatekeeper, ensuring that automated theme generation does not flatten or erase participants' authentic voices and cultural specificity [10].

11 Conclusion

AI is genuinely one of the more significant methodological developments in contemporary qualitative informatics. Its ability to organize, classify, and process huge volumes of unstructured text creates real opportunities for scaling qualitative work. However, the value of qualitative inquiry has always been in making sense of the depth of human experience, not throughput or efficiency [11]. SEL research is a useful reminder of that value: understanding emotion, relationship, vulnerability, and identity takes more than algorithmic sophistication. It takes researchers who stay interpretively present throughout the process.

This paper's core contribution is PHIP, a stabilizing epistemic principle for AI-assisted qualitative research. Rather than treating machine intelligence and human interpretation as competitors, we think the field's future depends on building responsible, transparent forms of human-AI collaboration. In keeping with the AII 2026 theme, reproducible qualitative research means systematically documenting how computational tools and human reflexive judgment actually work together. Thus, it does not mean removing the human from the loop in the name of artificial objectivity.

The framework has real limits. Maintaining a strict human-in-the-loop process reintroduces the time costs AI was partly meant to reduce, a genuine trade-off against computational efficiency. Because this is a reflexive, case-based conceptual paper, an inter-researcher pilot fell outside its scope; thus, how well the reflexive audit trail holds up across different research teams remains an open empirical question. Future work should test the framework in multi-researcher, multi-coder settings and examine its effects on interpretive consensus. The more challenging question, in the end, is not whether AI will keep getting more human-like; it is whether qualitative research retains its genuine humanity as AI becomes part of how it is conducted.

Code and Data Availability. This study is a conceptual and methodological paper and does not report on custom software or empirical datasets. Reasonable requests for additional information regarding the proposed PHIP framework and audit trail template may be addressed to the corresponding author via email [Anat.korem@l-w.ac.il].

Competing Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Jones, S. M., Doolittle, E. J.: Social and emotional learning: Introducing the issue. *The Future of Children* **27**(1), 3–11 (2017). <https://dx.doi.org/10.1353/foc.2017.0000>
2. Korem, A.: Teacher conceptual challenges in emotional education: key issues in supporting student social competence in pre-service and in-service teacher development. *International Journal of Emotional Education* **17**(2), 75–91 (2025). <https://www.um.edu.mt/library/oar/handle/123456789/140875>
3. Anito Jr, J.: Reflexive human–AI collaboration: Tracing the evolving epistemics of qualitative inquiry (2021–2025). *Social Sciences & Humanities Open* **13**, 102885 (2026). <https://doi.org/10.1016/j.ssaho.2026.102885>
4. Friese, S.: Conversational analysis with AI (CA to the power of AI): Rethinking coding in qualitative analysis (2025). <http://dx.doi.org/10.2139/ssrn.5232579>
5. Yin, Y., Su, Q., Li, S.: School belonging mediates the association between negative school climate and depressive symptoms among Chinese adolescents: a national population-based longitudinal study. *Frontiers in Psychology* **15**, 1368451 (2024). doi.org/10.3389/fpsyg.2024.1368451
6. Chen, S.: Beyond the boundaries of sharing: Parents’ perceptions on social-emotional learning (SEL) process in school. *Researchers@Parent-Education Teams* **3**, 44–85 (2025). <https://www.l-w.ac.il/articles/researchers-parent-education-teams/gilyonot-parent-education-teams/issue-3-sel-english/>
7. Meland, E. A., Brion-Meisels, G.: An integrative model for culturally sustaining SEL in the classroom. *Social and Emotional Learning: Research, Practice, and Policy* **3**, 100042 (2024) <https://doi.org/10.1016/j.sel.2024.100042>
8. Olmos-Vega, F. M., Stalmeijer, R. E., Varpio, L., Kahlke, R.: A practical guide to reflexivity in qualitative research: AMEE Guide No. 149. *Medical Teacher* **45**(3), 241–251 (2023). <https://doi.org/10.1080/0142159X.2022.2057287>
9. Pahl, A.: The AI-Reflexivity Checklist (ARC): a pre-analysis pause for LLM-assisted coding. *Qualitative Health Research* **36**(2-3), 181–190 (2026). <https://doi.org/10.1177/10497323251401503>

10. Resnik, D. B., Hosseini, M.: The ethics of using artificial intelligence in scientific research: new guidance needed for a new tool. *AI and Ethics* 5(2), 1499–1521 (2025). <https://doi.org/10.1007/s43681-024-00493-8>
11. Morrow, S. L.: Quality and trustworthiness in qualitative research in counseling psychology. *Journal of Counseling Psychology* 52, 250–260 (2005). <https://doi.org/10.1037/0022-0167.52.2.250>