Reflexive Uncertainty AI for Qualitative Data Analysis

António Pedro Costa[1,†] and Pedro Bem-Haja[2]

[1] Research Centre on Didactics and Technology in the Education of Trainers, Department of Education and Psychology, University of Aveiro, Portugal

[2] CINTESIS@RISE, Department of Education and Psychology, University of Aveiro, Aveiro, Portugal

Abstract

Most qualitative research in education relies on the interpretation of non-numerical and unstructured data. With the growing use of generative AI and probabilistic models, researchers face a methodological tension. While AI can accelerate analysis, its outputs remain statistical approximations that often overlook contextual nuance. This article proposes the Reflexive Uncertainty Framework (RUF) as a methodological extension of human–AI co-analysis that treats uncertainty not as error, but as analytically productive evidence. The framework documents three interpretive moments: (i) probabilistic classification outputs, where vector-based models estimate semantic proximity between segments and categories; (ii) zones of interpretive ambiguity, where close or unstable scores trigger a second qualitative layer of justification through reflective prompting with LLMs; and (iii) reflexive commentary, in which the researcher assumes the role of co-analyst, explicitly interpreting epistemic tension rather than concealing it behind automated certainty. The results were triangulated using traditional qualitative analysis through the webQDA software and OpenAI's ChatGPT, while the computational model is currently being implemented in Google Colab. A key contribution of the RUF is to make uncertainty methodologically visible; however, its analytical power depends on sustained human-in-the-loop engagement. Further validation is required with larger, multilingual datasets and across diverse interpretive traditions. By reframing uncertainty as a reflexive resource rather than a flaw, the framework strengthens epistemic transparency in human–AI collaboration.

Keywords

Qualitative Research, Educational Research, LLMs, AbductivAI [1]

1. Introduction

The critical reflection on uncertainty in interpretative processes mediated by artificial intelligence reveals the need for models that integrate the probabilistic logic of systems with human interpretative judgement. In this scenario, the AbductivAI model [1] emerges as a methodological response that frames human–AI collaboration around abductive reasoning, emphasising that uncertainty is not merely a technical flaw but a productive interpretive resource. The contribution of this work also lies within the field of AI education; by showing how uncertainty can be explicitly documented, negotiated, and interpreted, the framework models a form of *methodological reflexive literacy* that can be taught as part of responsible human–AI co-analysis in educational research settings.

Reflexive uncertainty in qualitative research consists of continually questioning and critically situating one's assumptions, decisions, and interpretations throughout the analytical process. In the context of human–AI collaboration, however, it is also necessary to address how automated systems express uncertainty, not only numerically but also through unstable or contradictory reasoning. When this article refers to "algorithmic reflexivity", it does not imply a human-like self-awareness; rather, it functions as an analytical metaphor to describe how a system's probabilistic oscillations, justification shifts, or contextual instability can be made visible as epistemic signals.

2nd Workshop on Education for Artificial Intelligence (edu4AI 2025, https://edu4ai.di.unito.it/), Co-located with ECAI 2025, the 28th European Conference on Artificial Intelligence which will take place on October 26, 2025 in Bologna, Italy

* Corresponding author.
† These authors contributed equally.
apcosta@ua.pt (A. P. Costa); pedro.bem-haja@ua.pt (P. Bem-Haja)
0000-0002-4644-5879 (A. P. Costa); 0000-0002-7547-5743 (P. Bem-Haja)

These signals do not replace human reflexive practice, but they can provoke it by indicating conceptual tension, ambiguity, or semantic friction requiring interpretive judgement.

2. Reflexive Uncertainty AI

Reflexive uncertainty in qualitative research consists of continually questioning and criticising your assumptions, decisions, and interpretations throughout the research process. Researchers critically reflect on their preconceptions, positions, and methodological choices to ensure ethical and differentiated analyses [2], [3]. For example, reflexive thematic analysis allows flexibility in aligning methods with philosophical positions during data interpretation [4]. On the other hand, to effectively utilise uncertainty in AI reflexively, it may be necessary to integrate the quantification and communication of uncertainty into human-AI collaboration while encouraging critical rethinking of human and AI limitations. Reflexive data curation, for example, harnesses GenAI to help users confront their biases and social norms, using uncertainty as a "tool" for deeper awareness and more responsible practices [5]. The same is true of research projects that use qualitative approaches. However, human uncertainty should also be taken into account; i.e., AI systems that take into account and learn from uncertain human feedback—rather than assuming that humans are always right—can become more robust and interactive, supporting safer and more effective joint decision-making [6]. Techniques such as probabilistic modelling, ensemble learning, and explainable AI methods that propagate uncertainty into their explanations further enhance transparency and trustworthiness [7], [8].

AI tools can complement this process, acting as peer questioners to evaluate codes and interpretations, although they must be used alongside human reflexivity to preserve the interpretative richness of qualitative research [9]. For example, iterative refinement guided by reflective insights can help address challenges such as the design and interpretability of tools like ChatGPT, ensuring they align with human-centred objectives and preserving the interpretive richness essential to qualitative research [9]–[11]. Incorporating reflexivity into AI tools can reflect the critical practices of researchers, resulting in systems that are both adaptable and ethically grounded [12].

Building on this idea, reflective journaling offers educators and researchers a structured way to document biases and uncertainties by integrating AI-orientated tools in qualitative research. By recording their thoughts, decisions, and emotional responses, they can critically assess how personal assumptions and AI limitations—such as systemic and dataset biases—shape the results [13], [14]. Although it requires sustained effort, reflexivity increases transparency and critical engagement with AI systems [15].

The development of RUF becomes particularly relevant here, as it provides structured mechanisms for documenting the interpretive ambiguities that emerge when human methodological expertise encounters AI's probabilistic nature. This methodological positioning also requires researchers to maintain what Bourdieu and Wacquant [16] term “epistemic vigilance”—a critical stance toward their knowledge production processes. As methodological experts working with AI, researchers must reflexively examine how their disciplinary training, theoretical commitments, and methodological preferences shape the ways they configure and interpret AI outputs. This includes acknowledging the inherent power dynamics involved in determining which methodological approaches are deemed suitable for AI implementation and which remain exclusively within human interpretive domains.

The conceptual distinction between human reflexivity and “algorithmic reflexivity” becomes crucial for understanding the dynamics of AI-human collaboration in qualitative research. Human reflexivity, as traditionally understood in the qualitative approach, involves conscious self-examination, critical awareness of positionality, and deliberate interrogation of assumptions that researchers bring to their work [17], [18]. This form of reflexivity is inherently experiential, drawing upon the researcher’s lived experiences, emotional responses, and intuitive insights to inform interpretive processes [19]. In contrast, “algorithmic reflexivity” represents a fundamentally different mode of self-examination—one that operates through computational processes, pattern recognition, and probabilistic calculations rather than conscious deliberation [20], [21]. While human reflexivity is characterised by intentionality, contextual sensitivity, and a capacity for ethical reasoning, “algorithmic reflexivity” emerges through iterative learning, statistical correlations, and systematic pattern detection across vast datasets [22]. These two forms of reflexivity are not merely different in degree but represent ontologically distinct modes of engaging with knowledge production. Human reflexivity involves what Daston and Galison [23] term "trained judgment"—the capacity to navigate ambiguity, recognise exceptions, and make contextually appropriate interpretive decisions. “Algorithmic reflexivity”, conversely, operates through what Mackenzie [24] describes as “machine learning logics”, which identify patterns and generate predictions based on statistical regularities rather than interpretive understanding. The productive tension between these reflexive modes creates opportunities for what might be termed “hybrid reflexivity”—a collaborative form of self-examination that leverages both human interpretive capabilities and algorithmic pattern recognition [25]. This hybrid approach recognises that human and algorithmic reflexivity can be mutually informative rather than mutually exclusive. The practical implementation of this integration, however, necessitates methodological innovations that can systematically harness both human interpretive capabilities and algorithmic analytical power. Human reflexivity can provide contextual interpretation of algorithmic outputs, while algorithmic reflexivity can reveal patterns and biases that human reflexivity might overlook due to cognitive limitations or embedded assumptions [26]. However, this collaboration requires careful attention to the asymmetries between these reflexive modes, particularly regarding issues of accountability, transparency, and interpretive authority [27].

3. AbductivAI model

GenAI models are significantly transforming the production of knowledge, especially in the field of qualitative research. New possibilities for collaboration with artificial agents have challenged the traditional focus of qualitative research on human interpretation. In this context, the AbductivAI model [1] emerges as a proposal that integrates humans and AIs as co-investigators in analytical processes, based on an abductive logic. This logic is not limited to a combination of induction and deduction but promotes creative inferences that enable a deeper understanding of the data [28], [29].

The theoretical foundation of the AbductivAI model [1] is anchored in two complementary approaches: the Actor-Network Theory (ANT) and sociomateriality. The Actor-Network Theory (ANT) enables us to comprehend humans and AIs as participants in knowledge production networks, defining and redefining their roles through interaction [30]–[32]. Sociomateriality, on the other hand, reinforces the inseparability between social and material elements, making it possible to conceive of AI as an active participant in the process of constructing meaning [33], [34]. Both perspectives support a relational and performative ontology, in which knowledge emerges from hybrid configurations where humans and algorithms co-construct interpretations.

The AbductivAI model (figure 1) organises the qualitative analysis process into eight iterative phases, ranging from the formulation of the research question to the final interpretation of the data. These phases include (1) definition of the research question and the theoretical framework, (2) initial construction of the category system, (3) definition of the coding rules, (4) initial joint coding between humans and agents, (5) critical review of the categories, (5.5) CoT Review Loop, (6) final coding with cross-validation, (7) verification and resolution of ambiguities, and (8) analysis and interpretation of the results [35]. To clarify the phase 5.5, the CoT (Chain-of-Thought) prompting is a technique in which the system externalises its intermediate reasoning steps, allowing the researcher to inspect, contest, and refine the interpretive path rather than only its final output.

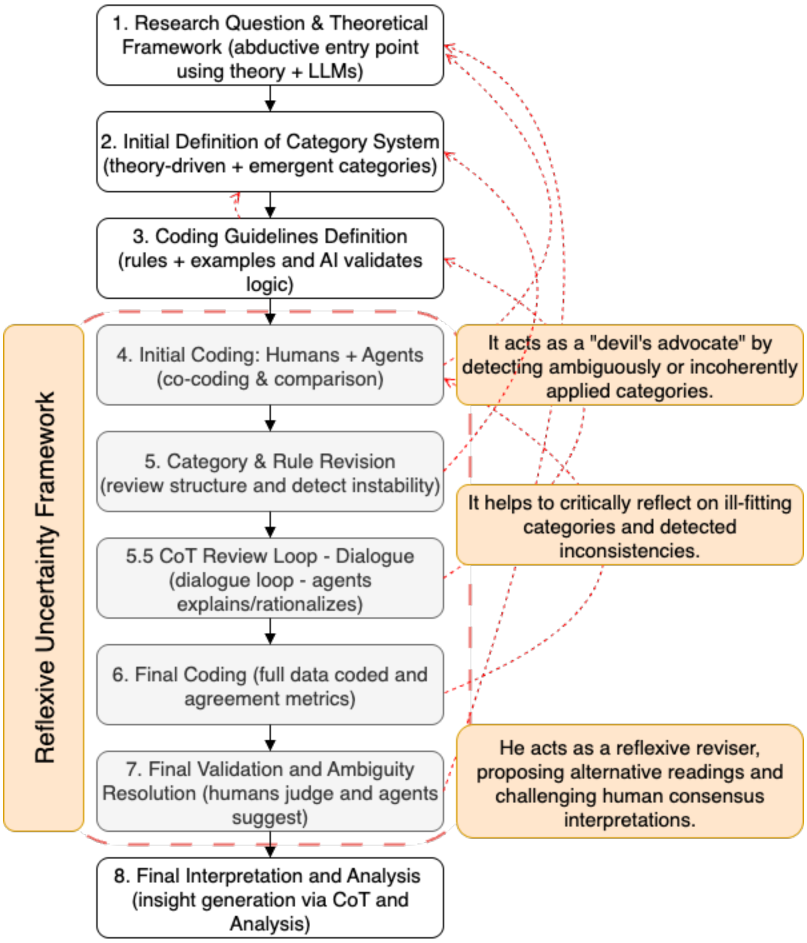

Figure 1: AbductivAI Model and the phases where RUF can be explored

In this article, the RUF is proposed as a conceptual deepening applied specifically to phase 5.5, reinforcing the documentation of uncertainty arising during the analytical process. Its role is not to replace AbductivAI, but to expand the reflective dimension of the stage in which the most direct confrontation between the model's probabilistic reasoning and human interpretative judgement occurs.

4. Reflexive Uncertainty Framework

The RUF (Figure 2) is not an independent model but rather a methodological deepening applied to AbductivAI [1], focusing in this article on phase 5.5 (CoT Review Loop). Its goal is to make analytical uncertainty visible and methodologically productive, clarifying when automatic interpretation requires human intervention. The notion of "algorithmic reflexivity" is used here in a technical rather than anthropomorphic sense: it does not imply conscious self-reflection but describes the mechanism through which the system makes explicit the interpretative limits of probabilistic calculation, exposing areas where its own decision is structurally unstable or epistemically fragile.
In the AbductivAI model [1], the RUF can work, preferably in the formulation, coding, and interpretation phases, where there is a greater risk of human bias or premature stabilisation of categories. Its role is to instigate transversal human competences, essentially cognitive ones, such as critical thinking, analytical thinking, challenging interpretative "certainties" and enriching the abductive process with plausible alternatives. We will apply it in phase 5.5 of the presented example, creating a reflexive loop through the RUF. Appendix 1 shows 9 student compositions (4 boys and 5 girls). The aim is to answer the following questions: "What are the characteristics of a successful teacher from the students' perspective? Are there differences in perspective depending on the gender of the respondents?"

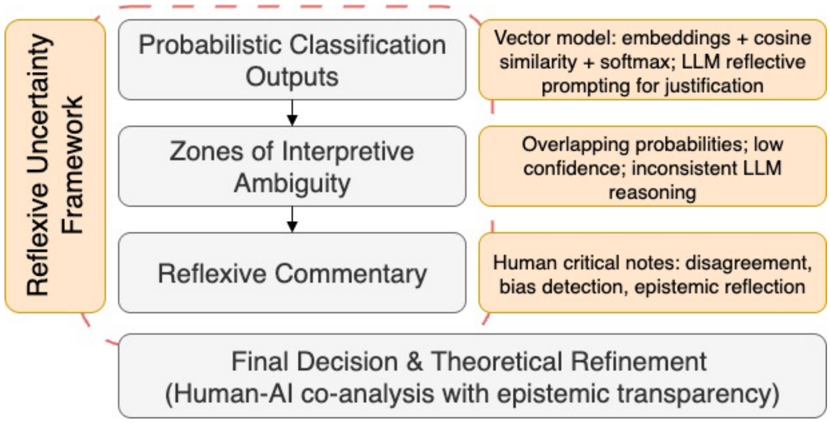

Figure 2: RUF Dimensions

While AbductivAI provides the general architecture for co-analysis, RUF functions as an epistemological increment that documents and qualifies emerging uncertainty. Its application follows three central steps: (1) identification of quantitative uncertainty (close probabilities); (2) qualitative examination via textual justification generated by LLM; and (3) human interpretative intervention in the form of reflective commentary. This articulation constitutes a hybrid cycle in which statistical reasoning and human interpretation feed back into each other. Before presenting the table, a brief example illustrates this dynamic: given the sentence "I would stop the activity to let them do what they prefer," the vector model approximates the codes Empathy (72%) and Permissiveness (68%). The overlap suggests a zone of semantic indecision. In justifying the decision, the LLM describes the gesture as empathetic but admits a lack of regulation. This oscillation constitutes the uncertainty that activates the RUF; the researcher clarifies whether they interpret the act as relational care or erosion of authority. To consolidate these three axes, we present a brief conceptual summary below the table:

1. Probabilistic Classification Outputs – the procedure uses probabilistic classification as the first layer of analysis. Initially, a classic vector model identifies general patterns in the text by transforming excerpts and analytical categories into high-dimensional vectors based on models such as BERT [36], calculating cosine similarity [37] and applying softmax [38] to generate probability distributions. These results allow us to map semantic proximities and detect interpretative ambiguities, such as when categories have very similar scores. When this ambiguity occurs, reflective prompts are activated with LLMs for a second qualitative layer, whose purpose is not to recalculate probabilities but to discursively justify the relevance of categories and explicitly compare competing terms. This contextualised use aims to reveal interpretative nuances that statistical scoring alone does not capture. If the model's justifications show contradictions or instability, this is understood, according to Kompa et al. [39], as evidence of "epistemic instability," which RUF considers analytical data, not mere noise. Thus, the system produces a hybrid set composed of vector probabilities, discursive justifications, and contrastive reasoning. This approach shifts the focus from the search for certainty to the productive use of uncertainty as an epistemological resource, strengthening the researcher's methodological reflexivity.
2. Zones of Interpretive Ambiguity – correspond to sections of the corpus in which the AI classification reveals conceptual uncertainty or interpretative tension [18]. This ambiguity is not treated as an error but as a privileged space for critical analysis. It arises mainly when the vector model assigns very close scores between categories or when the overall probability distribution is low and diffuse, signalling a lack of clear semantic alignment. In these cases, a second analytical layer with LLMs is automatically triggered, which must discursively justify the relevance or irrelevance of the categories, including through contrastive prompts. When such justifications prove fragile, contradictory, or unstable in the face of small changes in formulation, this oscillation reveals epistemic instability, composing the analytical evidence itself. Thus, the ambiguous zone is defined by the combination of two markers: (i)

statistical proximity between categories (quantitative uncertainty) and (ii) fragile or inconsistent justifications (qualitative uncertainty). Identifying these zones makes doubt methodologically visible, recognising that certain ambiguities reflect real conflicts of values, context or interpretation — and therefore require greater scrutiny and possible human triangulation.

3. Reflexive Commentary – marks the moment when the researcher makes their own interpretative position explicit, not to "correct" the AI, but to make human judgement and its epistemological basis audible [19]. Reflexive commentary marks the moment when the researcher ceases to be a mere consumer of AI-generated inferences and assumes the role of co-analyst. It consists of a critical annotation, concise yet epistemologically dense, recorded whenever there is a significant divergence between human interpretation and the AI's classification, whether that divergence arises from quantitative outputs (vector-based model) or qualitative reasoning (LLM-based reflective model). The goal is not simply to correct the AI but to make the interpretive process—and its boundaries—explicit and accountable. Such commentary may be triggered, for example, when the AI assigns Authority: 78% to a segment but overlooks affective or ironic cues in the surrounding context ("She placed the children who didn't learn at the back of the classroom"). In such cases, the researcher is expected to articulate:
   - the nature of the disagreement—e.g., "The model fails to recognise the symbolic weight of spatial exclusion";
   - the possible sources of misinterpretation, such as cultural insensitivity, lack of pragmatic awareness, or institutional ignorance.
   - whether the discrepancy reveals a deeper analytical insight, such as a tension between intent and effect, or a limitation in the theoretical framing of the analysis itself.

Thus, the following table shows how these three dimensions operate in the corpus, highlighting four analytical elements: (i) the initial probabilistic output; (ii) the discursive justification provided by the LLM; (iii) the marking of the zone of interpretative ambiguity; and (iv) the researcher's reflective comment.

Table 1
RUF results applied in phase 5.5

RUF Dimensions	Examples and Reflection on the Corpus	Probabilistic Classification (Vectorial)	LLM feedback (reflective prompting)	Zones of Interpretive Ambiguity	Reflexive Commentary
Fixed—for all types of data corpus					
Interpretative ambiguity (presence of expressions that allow for multiple interpretations, with no clear consensus).	C3: 'I'd stop and tell them to do something else they'd prefer.'	Empathy: 72%, Permissiveness: 68%	"Both apply: the tone is empathetic, but the absence of regulation can indicate permissiveness."	Yes—two close categories (empathy vs. permissiveness)	The AI attributed empathy but may be confusing it with a loss of authority. The emotional context is ambiguous.
Tension between codes (applied codes that compete or overlap in the same segment).	C4: discipline with friendship promotion	Discipline: 71%, Collective Care: 69%.	"Discipline is present, but with a strong relational component. Both are justifiable."	Yes—tension between regulation and bonding	AI brands discipline and empathy without recognising that these ideas are in tension in real school contexts.
Relevant Absence—Thematic Silence	Almost no composition	Formal Evaluation: 22%, Planning: 18%	"There is no explicit reference. It may be a	Yes—absence detected by humans, not AI	AI did not identify absence as analytical data—a

(expected topics that don't appear: i.e., evaluation, verbalised explanation).	mentions tests or assessments		thematic omission, but it has not been interpreted as such."		recurring limitation.
Variables—depending on the data corpus					
Reversals or breaches of expectations (when the expression deviates from expected patterns, such as gender roles or hierarchy).	C7: girl proposes severe punishment	Authority: 77%, Severity: 74%	"The intention is to command respect, but the tone can be exaggerated or dramatised."	No - high score, but breach of cultural expectation	AI applies authority with confidence but does not perceive deviation from social or gender norms.
Multiplicity of Positions (representations of the teacher as a figure who accumulates divergent functions).	C8: strict control, more recreation	Discipline: 67%, Freedom: 65%	"Both elements are present; there's an attempt to balance authority and fun."	Yes — overlapping of opposite categories	AI collapses control and freedom as complementary, ignoring the pedagogical dilemma.
Students express a desire for one practice but compare it with the opposite one, creating a conflict between practice and desire.	C5: rejection of tests, but appreciation of games	Playfulness: 70%, Institutional Criticism: 68%	"The playfulness is explicit, but the passage contains implicit criticism of the evaluation system."	Yes—desire vs. implicit criticism coexist.	AI identifies playfulness but doesn't notice the criticism embedded in the traditional school model.

The results presented in the table were triangulated using traditional qualitative methods through the webQDA software [40] and OpenAI's ChatGPT, while the model is currently being implemented in Google Colab.

5. Final Reflections

Adopting AI in qualitative research presents substantial opportunities alongside significant challenges. Opportunities include enhanced efficiency in data processing and analysis [41], the ability to analyse larger datasets than manually possible [42], and potential discovery of patterns humans might miss [43]. Firnando and Wahyudi [44] suggest AI could democratise access to advanced analytical capabilities, allowing more researchers—particularly from underserved communities—to enhance their qualitative research. Time and resource savings during routine tasks can accelerate project timelines [45]. However, challenges remain significant. Gibson and Beattie [46] caution against AI impersonating human participants, emphasising the importance of affect and human experience in qualitative data. Ethical concerns regarding data privacy and participant confidentiality require careful consideration [47]. Technical challenges include AI's difficulty capturing contextual subtleties and cultural nuances [48]. Baig et al. [49] emphasise the need for robust ethical frameworks to foster user trust in AI-generated insights. Yang and Berdine [50] warn that researchers must remain vigilant regarding AI tools' limitations and potential inaccuracies.

At this stage, what becomes most relevant is not the automation of coding but the status of uncertainty within interpretive collaboration. The framework shows that uncertainty need not be eliminated for analysis to progress; instead, it can be made analytically productive when brought into methodological visibility.

Within this small corpus, the application of the RUF shows that reflexive uncertainty can function as analytical evidence rather than noise, revealing where interpretive boundaries are unstable. By foregrounding these moments, the framework supports accountable interpretation, particularly when human and AI coders diverge. In doing so, the RUF clarifies that uncertainty should not be suppressed but made methodologically visible and accountable.

A key limitation of this study concerns the ethical and pedagogical scope of the framework. Its contribution depends on researchers actively sustaining a human-in-the-loop posture, which cannot be assumed by default in AI-assisted qualitative work. The present analysis also remains methodologically bounded by a small and context-specific corpus, meaning that reflexive uncertainty still requires validation across different interpretive cultures and research traditions. From a technical perspective, broader scalability and reproducibility will require future testing with larger and multilingual datasets. Future work will therefore extend the framework beyond research settings into educational practice, enabling educational researchers to co-analyse uncertainty as part of reflexive AI literacy. By reframing uncertainty as accountable evidence rather than a technical defect to be eliminated, the RUF strengthens epistemic agency in human–AI co-analysis. In doing so, it transforms interpretive instability into a legitimate site of inquiry, inviting human judgement to become visible and reviewable. Ultimately, embracing uncertainty through reflexive practices not only enhances the technical robustness of analytical outcomes but also fosters ethical awareness and a more meaningful form of human–AI collaboration [5], [6], [51], [52].

Acknowledgements

The work of the first is funded by national funds through FCT – Fundação para a Ciência e a Tecnologia, I.P., under the Scientific Employment Stimulus - Institutional Call - [CDL-CTTRI-248-SGRH/2022] and the CIDTFF (projects UIDB/00194/2020 and UIDP/00194/2020).

Declaration on Generative AI

We declare the use of ChatGPT to conduct the first RUF tests using the compositions of the students listed in Appendix A. If I Were a Teacher.

References

[1] A. P. Costa, G. Bryda, P. A. Christou, and J. Kasperiuniene, "AI as a Co-researcher in the Qualitative Research Workflow: Transforming Human-AI Collaboration," *Int. J. Qual. Methods*, vol. 24, Sep. 2025.

[2] F. M. Olmos-Vega, R. E. Stalmeijer, L. Varpio, and R. Kahlke, "A practical guide to reflexivity in qualitative research: AMEE Guide No. 149," *Med. Teach.*, vol. 45, no. 3, pp. 241–251, Mar. 2023.

[3] D. Watt, "On Becoming a Qualitative Researcher: The Value of Reflexivity," *Qual. Rep.*, Jan. 2015.

[4] H. Pearson, M. Myall, A.-S. Darlington, and F. Gibson, "The approach and application of analysing inductive and deductive datasets: a worked example using reflexive thematic analysis," *Qual. Res. Psychol.*, pp. 1–45, May 2025.

[5] A. Arzberger, M. L. Lupetti, and E. Giaccardi, "Reflexive Data Curation: Opportunities and Challenges for Embracing Uncertainty in Human–AI Collaboration," *ACM Trans. Comput. Interact.*, vol. 31, no. 6, pp. 1–33, Dec. 2024.

[6] K. M. Collins *et al.*, "Human Uncertainty in Concept-Based AI Systems," in *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*, 2023, pp. 869–889.

[7] S. Bobek and G. J. Nalepa, "Introducing Uncertainty into Explainable AI Methods," 2021, pp. 444–457.

[8] T. Qiu and Q. Li, "A Framework for Counterfactual Explanation of Predictive Uncertainty in Multimodal Models," *IEEE Trans. Neural Networks Learn. Syst.*, pp. 1–14, 2024.

[9] Z. O. Dunivin, "Scaling hermeneutics: a guide to qualitative coding with LLMs for reflexive content analysis," *EPJ Data Sci.*, vol. 14, no. 1, p. 28, Apr. 2025.

[10] L. Thominet, J. Amorim, K. Acosta, and V. K. Sohan, "Role Play: Conversational Roles as a Framework for Reflexive Practice in AI-Assisted Qualitative Research," *J. Tech. Writ. Commun.*, vol. 54, no. 4, pp. 396–418, Oct. 2024.

[11] H. Zhang, C. Wu, J. Xie, Y. Lyu, J. Cai, and J. M. Carroll, "Harnessing the power of AI in qualitative research: Exploring, using and redesigning ChatGPT," *Comput. Hum. Behav. Artif. Humans*, vol. 4, p. 100144, May 2025.
[12] H.-H. Hachem and F. Heintz, "Where is the reflexive 'I' in the Elements of AI?," *Int. J. Lifelong Educ.*, vol. 43, no. 6, pp. 664–681, Nov. 2024.
[13] R. Schwartz, A. Vassilev, K. Greene, L. Perine, A. Burt, and P. Hall, "Towards a standard for identifying and managing bias in artificial intelligence," Mar. 2022.
[14] V. Matthews and M. Murphy, "Addressing Bias in Artificial Intelligence," 2023.
[15] A. Barrett, A. Kajamaa, and J. Johnston, "How to … be reflexive when conducting qualitative research," *Clin. Teach.*, vol. 17, no. 1, pp. 9–12, Feb. 2020.
[16] P. Bourdieu and L. Wacquant, *An Invitation to Reflexive Sociology*. Chicago: University of Chicago Press, 1992.
[17] L. Finlay, "Negotiating the swamp: the opportunity and challenge of reflexivity in research practice," *Qual. Res.*, vol. 2, no. 2, pp. 209–230, Aug. 2002.
[18] W. Pillow, "Confession, catharsis, or cure? Rethinking the uses of reflexivity as methodological power in qualitative research," *Int. J. Qual. Stud. Educ.*, vol. 16, no. 2, pp. 175–196, Mar. 2003.
[19] M. Alvesson and K. Sköldberg, *Reflexive Methodology: New Vistas for Qualitative Research*. London: Sage Publications, 2017.
[20] K. Barad, *Meeting the Universe Halfway: Quantum Physics and the Entanglement of Matter and Meaning*. Durham, NC: Duke University Press, 2007.
[21] R. Kitchin, "Thinking critically about and researching algorithms," *Information, Commun. \& Soc.*, vol. 20, no. 1, pp. 14–29, 2017.
[22] S. Pink, H. Horst, J. Postill, L. Hjorth, T. Lewis, and J. Tacchi, *Digital Ethnography: Principles and Practice*. London: Sage Publications, 2016.
[23] L. Daston and P. Galison, *Objectivity*. New York: Zone Books, 2007.
[24] A. MacKenzie, *Machine Learners: Archaeology of a Data Practice*. Cambridge, MA: MIT Press, 2017.
[25] N. Marres, *Digital Sociology: The Reinvention of Social Research*. Cambridge, UK: Polity Press, 2017.
[26] L. D. Introna, "Algorithms, governance, and governmentality: On governing academic writing," *Sci. Technol. \& Hum. Values*, vol. 41, no. 1, pp. 17–49, 2016.
[27] T. Gillespie, "The relevance of algorithms," in *Media Technologies: Essays on Communication, Materiality, and Society*, T. Gillespie, P. J. Boczkowski, and K. Foot, Eds. Cambridge, MA: MIT Press, 2014, pp. 167–194.
[28] J. Reichertz, "Abduction: The Logic of Discovery of Grounded Theory," in *The SAGE Handbook of Grounded Theory*, 1 Oliver's Yard, 55 City Road, London England EC1Y 1SP United Kingdom: SAGE Publications Ltd, 2007, pp. 214–228.
[29] I. Tavory and S. Timmermans, *Abductive analysis: Theorizing qualitative research*. University of Chicago Press, 2014.
[30] B. Latour, *Reassembling the Social*. Oxford University PressOxford, 2005.
[31] R. Holton and R. Boyd, "'Where are the people? What are they doing? Why are they doing it?'(Mindell) Situating artificial intelligence within a socio-technical framework," *J. Sociol.*, vol. 57, no. 2, pp. 179–195, Jun. 2021.
[32] S. Keiff, "Latour's Discernment Tools in Action: Exploring and Solving Complex Social Problems Through Participatory Methods," *Int. J. Qual. Methods*, vol. 23, p. 16094069241286844, Jan. 2024.
[33] W. J. Orlikowski and S. V Scott, "Sociomateriality: Challenging the separation of technology, work and organization," *Acad. Manag. Ann.*, vol. 2, no. 1, pp. 433–474, 2008.
[34] S. V Scott and W. J. Orlikowski, "Exploring AI-in-the-making: Sociomaterial genealogies of AI performativity," *Inf. Organ.*, vol. 35, no. 1, p. 100558, Mar. 2025.
[35] A. P. Costa, G. Bryda, P. A. Christou, and J. Kasperiuniene, "AI as a Co-Researcher in Qualitative Research Workflow: Transforming Human-AI Collaboration," *Int. J. Qual. Methods*.
[36] N. M. Gardazi, A. Daud, M. K. Malik, A. Bukhari, T. Alsahfi, and B. Alshemaimri, "BERT applications in natural language processing: a review," *Artif. Intell. Rev.*, vol. 58, no. 6, p.

166, Mar. 2025.
[37] A. Singhal, "Modern Information Retrieval: A Brief Overview." IEEE Computer Society Technical Committee on Data Engineering, pp. 1–9, 2001.
[38] J. Bridle, "Training Stochastic Model Recognition Algorithms as Networks can Lead to Maximum Mutual Information Estimation of Parameters," in *Advances in Neural Information Processing Systems*, 1990, vol. 2.
[39] B. Kompa, J. Snoek, and A. L. Beam, "Second opinion needed: communicating uncertainty in medical machine learning," *npj Digit. Med.*, vol. 4, no. 1, p. 4, Jan. 2021.
[40] A. P. Costa, A. Moreira, and F. N. de Souza, "webQDA - Qualitative Data Analysis." University of Aveiro and MicroIO, Aveiro - Portugal, 2019.
[41] M. Goyanes, C. Lopezosa, and B. Jorda, "Thematic analysis of interview data with ChatGPT: Designing and testing a reliable research protocol for qualitative research," *SocArXiv*, 2024.
[42] E. I. Ibrahim and A. Voyer, "The Augmented Qualitative Researcher: Using Generative AI in Qualitative Text Analysis." 24-Jan-2024.
[43] M. S. Jalali and A. Akhavan, "Integrating AI language models in qualitative research: Replicating interview data analysis with ChatGPT," *Syst. Dyn. Rev.*, vol. 40, no. 3, p. e1772, Jul. 2024.
[44] H. G. Firnando and M. Wahyudi, "The Role of Artificial Intelligence in Shaping the Islamic Worldview of the Digital Economy," *J. Islam. Econ. Philanthr.*, vol. 6, no. 3, pp. 231–249, 2024.
[45] A. B. Hamilton and E. P. Finley, "Qualitative methods in implementation research: An introduction," *Psychiatry Res.*, vol. 280, p. 112516, 2019.
[46] A. F. Gibson and A. Beattie, "More or less than human? Evaluating the role of AI-as-participant in online qualitative research," *Qual. Res. Psychol.*, vol. 21, no. 2, pp. 175–199, Apr. 2024.
[47] D. T. Marshall and D. B. Naff, "The Ethics of Using Artificial Intelligence in Qualitative Research," *J. Empir. Res. Hum. Res. Ethics*, vol. 19, no. 3, pp. 92–102, Jul. 2024.
[48] J. W. Creswell and C. N. Poth, *Qualitative Inquiry and Research Design Choosing among Five Approaches*, 4th Editio. Thousand Oaks: SAGE Pubications, 2018.
[49] M. I. Baig and E. Yadegaridehkordi, "ChatGPT in the higher education: A systematic literature review and research challenges," *Int. J. Educ. Res.*, vol. 127, p. 102411, 2024.
[50] S. Yang and G. Berdine, "Artificial intelligence in biomedical research," *Southwest Respir. Crit. Care Chronicles*, vol. 11, no. 46, pp. 62–65, Jan. 2023.
[51] P. Constantinides, E. Monteiro, and L. Mathiassen, "Human-AI joint task performance: Learning from uncertainty in autonomous driving systems," *Inf. Organ.*, vol. 34, no. 2, p. 100502, Jun. 2024.
[52] H. V. Subramanian, C. Canfield, D. B. Shank, and M. Kinnison, "Combining uncertainty information with AI recommendations supports calibration with domain knowledge," *J. Risk Res.*, vol. 26, no. 10, pp. 1137–1152, Oct. 2023.

A. If I Were a teacher

Problem: What are the characteristics of a competent teacher from the students' perspective? Are there differences in perspective based on the respondents' gender?
The document corpus: 9 compositions from 5th grade students.
Goal: The characteristics of a good teacher reported by the students differ according to the students' sex (M/F).
Compositions:
If I were a teacher, I would always arrive on time to set an example for the students. I would also prepare my lessons well so that everyone could learn and understand the subjects. I would be strict with my students and would not allow them to misbehave. If they did misbehave, I would first warn them and then inform their parents. – M

1. If I were a teacher, I would try to help the students who have more difficulties. I would prepare the lessons well so that everyone could understand me. But when everything went well, I would also end the class earlier and let them go to recess. – M
2. If I were a teacher, I would not allow anyone to play around during my classes; if they did, they would be punished immediately. But when I saw they were tired, I would stop and let them do something else. – F
3. If I were a teacher, I would ask many questions so we could learn the material well; I would also give the students a chance to ask questions about what they didn't know. I wouldn't allow them to fight or speak badly about each other. Whoever did that would be punished, and their parents would be notified. – F
4. If I were a teacher, I would never ask too much of the students, and I would let them play whenever they wanted. I wouldn't bother them with too many tests. Classes would be filled with games so we could learn, and we would also watch many movies. – M
5. If I were a teacher, I would do a lot of revision before tests. All students would be able to answer. Those who didn't learn would have to study with classmates who already knew the material. I would also never allow bad behaviour or swearing in class or name-calling among classmates. – F
6. If I were a teacher, I would punish the students right at the beginning of the year. If I were a teacher, I would implement disciplinary measures at the start of the academic year. I would require all students to stand at attention upon my arrival, and any student who moved would face consequences. Those who struggled to learn would be seated at the back of the classroom to minimise disruption for their classmates who were eager to learn. Those who didn't learn would be placed at the back of the classroom so they wouldn't disturb the classmates who wanted to learn. – F
7. If I were a teacher, I would try to help the weaker students. I would also try to understand why some students had more difficulties and would always offer my help. I would not permit distractions during class time; however, I would also make an effort to create opportunities for us to engage in enjoyable activities and share laughter. – M
8. If I were a teacher, I would make sure all the students learned the lessons well and that the lessons weren't boring. At the same time, I would encourage students to all be friends with one another so they could feel good at school. – F