OBJN
Online Brazilian Journal of Nursing
ISSN: 1676 - 4285 | www.objnursing.uff.br

**Submission:** 03-Sep-2026
**Approved:** 20-Oct-2026

# ARTIFICIAL INTELLIGENCE AGENTS FOR REFLEXIVE THEMATIC ANALYSIS WITHOUT PROGRAMMING

AGENTES DE INTELIGÊNCIA ARTIFICIAL PARA ANÁLISE TEMÁTICA REFLEXIVA SEM PROGRAMAÇÃO

AGENTES DE INTELIGENCIA ARTIFICIAL PARA EL ANÁLISIS TEMÁTICO REFLEXIVO SIN PROGRAMACIÓN

**Rafael Cardoso Sampaio**[1] ORCID: 0000-0001-5176-173X [1] Universidade Federal de Pernambuco, Departamento de Ciência Política. Recife, Pernambuco, Brazil.

**How to cite:** Sampaio RC. Artificial intelligence agents for reflexive thematic analysis without programming. Online Braz J Nurs. 2026;25(Suppl1):e20267166. https://doi.org/10.17665/1676-4285.20267166

## RESUMO

**Objetivo:** descrever uma *skill* para análise temática reflexiva e mostrar que pesquisadores sem formação em programação podem usar agentes de inteligência artificial na análise de dados qualitativos de modo simples, confiável, transparente e replicável. **Método:** estudo metodológico. A *skill* foi construída a partir das seis fases da análise temática reflexiva e dos riscos documentados do uso de inteligência artificial generativa. Foi aplicada a um corpus real de doze sessões de grupos focais sobre sentidos de democracia e proteção social e avaliada em teste controlado com 126 execuções cegas. **Resultados:** a *skill* organiza a análise em sete fases, da decisão metodológica ao relatório final, com verificação literal das citações e bloqueio de atalhos. No corpus real, produziu o fluxo completo com cobertura total. No teste cego, a condição com a *skill* teve maior aderência aos critérios metodológicos. **Conclusão:** agentes de inteligência artificial podem ser usados sem programação, com método e rastreabilidade, desde que a interpretação e a responsabilidade sejam humanas.
**Descritores:** Inteligência Artificial Generativa; Pesquisa Qualitativa; Análise de Dados; Reprodutibilidade dos Testes.

## ABSTRACT

**Objective:** To describe a *skill* for reflexive thematic analysis and to show that researchers without programming training can use artificial intelligence agents to analyze qualitative data in a simple, reliable, transparent, and replicable way. **Method:** Methodological study. The *skill* was developed based on the six phases of reflexive thematic analysis and on the documented risks of using generative artificial intelligence. It was applied to a real corpus of twelve focus group sessions on the meanings of democracy and social protection and evaluated in a controlled test with 126 blind runs. **Results:** The *skill* organizes the analysis into seven phases, from the methodological decision to the final report, with literal verification of quotations and checks that block shortcuts. In the real corpus, it carried out the complete workflow with full coverage. In the blind test, the condition with the *skill* showed greater adherence to the methodological criteria. **Conclusion:** Artificial intelligence agents can be used without programming, in a methodical and traceable way, provided that interpretation and responsibility remain human.
**Descriptors:** Generative Artificial Intelligence; Qualitative Research; Data Analysis; Reproducibility of Results.

## RESUMEN

**Objetivo:** describir una *skill* para el análisis temático reflexivo y mostrar que investigadores sin formación en programación pueden utilizar agentes de inteligencia artificial en el análisis de datos cualitativos de modo sencillo, fiable, transparente y replicable. **Método:** estudio metodológico. La *skill* se construyó a partir de las seis fases del análisis temático reflexivo y de los riesgos documentados del uso de la inteligencia artificial generativa. Se aplicó a un corpus real de doce sesiones de grupos focales sobre los sentidos de la democracia y la protección social y se evaluó en una prueba controlada con 126 ejecuciones ciegas. **Resultados:** la *skill* organiza el análisis en siete fases, desde la decisión metodológica hasta el informe final, con verificación literal de las citas y bloqueo de atajos. En el corpus real, produjo el flujo completo con cobertura total. En la prueba ciega, la condición con la *skill* mostró mayor adherencia a los criterios metodológicos. **Conclusión:** los agentes de inteligencia artificial pueden utilizarse sin programación, con método y trazabilidad, siempre que la interpretación y la responsabilidad sean humanas.
**Descriptores:** Inteligencia Artificial Generativa; Investigación Cualitativa; Análisis de Datos; Reproducibilidad de los Resultados.

**Editors:**
Rosimere Ferreira Santana (ORCID: 0000-0002-4593-3715)
Audrey Vidal Pereira (ORCID: 0000-0002-6570-9016)

**Publisher:**
Escola de Enfermagem Aurora de Afonso Costa – UFF
Rua Dr. Celestino, 74 – Centro, CEP: 24020-091 – Niterói, RJ, Brazil
Journal email: objn.cme@id.uff.br

**Corresponding author:**
Rafael Cardoso Sampaio
E-mail: cardososampaio@gmail.com

**What is already known:**

- Generative AI is already used in qualitative analysis without adequate documentation; *chatbots* can fabricate quotations and produce themes whose construction cannot be traced.
- Programmatic workflows offer more control and reproducibility but require technical knowledge that most qualitative researchers lack.
- Reflexive thematic analysis demands human interpretive decisions, and a simple, auditable way to conduct it with AI agents was lacking.

**What this article adds:**

- The article develops an open *skill* that lets researchers without programming training conduct reflexive thematic analysis with AI agents, using natural language.
- The *skill* checks that every quoted statement actually appears in the original transcripts, which prevents AI from fabricating quotations. It also records each decision, making the analysis auditable and open to contestation by third parties.
- In the blind test of 126 runs, the *skill* showed greater adherence to methodological criteria. Even so, this does not remove the need for human validation.

## INTRODUCTION

A researcher finishes collecting twenty interviews and needs to analyze the material. They open a generative artificial intelligence (GenAI) platform, paste excerpts from the transcripts, and receive a list of themes that look ready to use. The answer is fluent and convincing, and it is also hard to verify: there is no record of how the themes were constructed, no indication of where each idea appears in the material, and no guarantee that the quoted statements actually exist[1]. What makes this result troubling is that its findings cannot be traced and its data cannot be defended[2].

The use of GenAI in qualitative research has grown faster than its documentation[3]. Protocols for conducting thematic analysis with the ChatGPT chatbot already appear in peer-reviewed journals[4], and controlled comparisons between models and human researchers produce similar coding results, with substantial differences in interpretation[5]. Most systems report a predominance of hybrid arrangements, in which a human has the final word[6]. The risks include fabricated or embellished quotations[2]and the model's tendency to agree with the researcher's hypothesis, known as sycophancy, or simply flattery[6]. Other risks include texts that read fluently on the surface but mask a lack of interpretation[7], and the silent erasure of rare experiences[6]. Over the same period, the field began to produce examples of structured collaboration between humans and machines, with explicit protocols and a record of decisions[8].

The chatbot is easy to use, but its methodological control is fragile because every conversation starts from scratch. Programmatic workflows usually offer greater control and reproducibility, but they require technical knowledge and exclude precisely those who most need assistance with documentation[9]. There is an intermediate path, still little explored in the literature, of which most qualitative researchers are unaware.

This path relies on artificial intelligence (AI) agents, which are built on a structure that includes a GenAI model and lets the agent use tools, follow rules, and carry out tasks with more control. The agent can access files and folders, search the web, run code, use APIs, query databases, create and edit documents, test outputs, compare sources, log every step of the work, and correct errors before delivering an answer. This structure is known as a *harness*[9]. In practice, such harnesses are available in easy-to-download software that runs on GenAI model subscriptions, such as Google's Antigravity, Anthropic's Claude Code, and OpenAI's Codex.

Unlike chatbots, AI agents can carry out long and complex tasks because they plan the work, divide it into stages, and manage the context window (how much information the GenAI model can keep in mind at once). When that window nears its limit, the agent can summarize or record what it has already done, free up space, and continue without having to start over.

One way to carry out long tasks is to use *skills* designed for specific types of tasks. A *skill* is a reusable, packaged procedure, usually a file or a folder of plain-text instructions, that the agent loads when it recognizes a given type of task. It can contain anything from simple guidance to a complex workflow: several stages, decision criteria, the tools to use, ways to check outputs, rules for handling errors, and instructions on how to record each stage. A *skill* does not merely tell the agent what to do; it can specify how to carry out the entire process. Because the protocol is open, any researcher can publish, adapt, and submit the procedure they used for peer review.

Reflexive thematic analysis, formulated by Braun and Clarke in 2006[10]and reformulated in 2019[11], offers a good case for demonstrating this possibility. The method is organized into six phases and holds that themes are constructed by the analyst from patterns of meaning, not discovered fully formed in the material. Its reception in Brazil shows that this distinction does not always carry over into practice[12]. Reflexive thematic analysis has clear criteria and defined phases, and it demands decisions from the researcher that cannot be delegated, which makes it well suited to demonstrating what a *skill* can do.

The analise-tematica-reflexiva *skill* is an open package that conducts reflexive thematic analysis in seven phases, from the methodological decision to the final report. This article describes how the *skill* was built and demonstrates its use on a real qualitative research corpus, showing how a researcher with no training in programming can analyze qualitative data with AI agents. The scientific contribution lies in making the analysis simple, reliable, transparent, and replicable, while keeping interpretation a human task.

## METHOD

This study describes the development of a *skill* and its application to a real qualitative research corpus. The *skill* grew out of the intersection of the six phases of reflexive thematic analysis and the documented risks of using GenAI[2,6,7]. Each phase of the method became a step with

a required output and a check that blocks shortcuts, such as skipping the initial reading of the material, accepting vague codes, citing an unverified excerpt, or treating a topic summary as a theme.

The *skill* organizes the analysis into seven phases, described in Table 1. In each phase, the agent proposes and the researcher decides, and a dedicated document records that separation. Phases F1 to F6 correspond to Braun and Clarke's six phases[10-11]. Phase F0 is a governance addition: it requires the researcher to record decisions before coding. Separating contestation from definition in phases F4 and F5 makes explicit as distinct steps what the method handles at two points.

**Table 1** – Phases of the *skill*, what the researcher decides, and what the agent produces

| Phase | What the researcher decides | What the agent produces |
|---|---|---|
| F0. Methodological decisions | The research question, the corpus, the theoretical position, and how much to delegate | An analysis plan with declared limits |
| F1. Initial reading | What deserves to be recorded | Notes organized by source and an inventory of what was read |
| F2. Codes | Which codes are kept and which are dropped | Codes anchored in located excerpts |
| F3. Candidate themes | Which pattern of meaning organizes each theme | A map of themes by meaning, with residues kept separate |
| F4. Contestation | Which alternative reading is accepted or rejected | A rival reading and an examination of contradictions |
| F5. Defining and naming | The name and the boundary of each theme | One theme passport per theme, with genealogy and excerpts |
| F6. Report | What can be asserted | A report with an AI use statement and the status of the analysis |

Source: prepared by the author, 2026.

The *skill* lets the researcher choose how far the machine proceeds on its own. On the most autonomous route, the agent goes through all seven phases and declares the absence of human validation. On another route, it stops at defined points for verification. The researcher can also hand over notes, codes, or initial themes, which the agent checks against the corpus. There is also a parallel route: the agent reads everything before seeing the human reading, and the two readings are compared at the end, with agreement between them not treated as validation. The name reflexive thematic analysis is reserved for the route with an active human analyst; the most autonomous route produces auditable material for exploration, and the final status statement explicitly records that condition.

Using the *skill* requires no coding, but it does require installing an agent program and copying a folder of files, basic operations familiar to any computer user. For beginners, programs with a visual interface, such as Antigravity 2.0, Codex/ChatGPT Work, or Claude Code/Cowork, are the simplest starting point[9]. The *skill* is openly available in a public repository[6]. The researcher need only ask the agent to install the *skill* and provide the repository URL; the agent will then copy the *skill* into the program's *skills* folder. Next, the researcher opens the agent in a folder on their own computer that contains the material to be analyzed, such as interview or focus group transcripts. Finally, the researcher simply states the request in natural language, as they would ask a research assistant. The agent then follows the *skill*: it asks what type of analysis to perform and carries out the corresponding steps.

To demonstrate its use, the *skill* was applied in OpenAI's Codex application to a real corpus of twelve focus group sessions from across Brazil, totaling 154,573 words, on the meanings of democracy and social protection among participants in public policy programs, using the ChatGPT Luna 6 model with the effort level set to extra high. The request was made in natural language and specified the research question and the autonomous route, in which the agent goes through the phases without stopping for verification; no human took part in the workflow, and the final status statement therefore records the absence of human validation. The corpus came from focus groups held in 2024 in 11 state capitals across the country, divided into four income brackets defined in multiples of the minimum wage: 0–2, 2–5, 5–10, and 10–20. More specifically, the *skill* was applied to a derived corpus (limited to the sessions held in Belém, Recife, and Curitiba), anonymized to prevent reconstruction of the original context and identification of participants.

The *skill* also underwent a controlled comparative evaluation, described in the final section of the Results, which compared the same model with and without the *skill* across 126 blind runs: twenty-one methodological failure scenarios, each repeated three times in each condition. These are "trap scenarios," that is, cases designed to test whether the protocol refuses shortcuts. Examples include skipping the initial reading and coding directly, treating a topic summary as a theme, publishing a quotation that does not exist in the corpus, exposing participant identifiers, fabricating consent for AI use, and harmonizing statements to force a conclusion. Each scenario specifies the incorrect behavior expected without the *skill* and the required behavior with the *skill*. The evaluation and its limits are part of the *skill*'s repository[6], and portability was verified across two

distinct AI agents.

## RESULTS

### What the researcher receives

Table 2 shows the natural-language request that initiated the demonstration.

**Table 2** – Example of a request to the agent and what it delivers

| The researcher asks | The agent delivers |
|---|---|
| Analyze the transcripts in this folder for the meanings of democracy and social protection, using reflexive thematic analysis on the autonomous route. | Seven documents, one per phase, with a final report, an audit appendix, and the status of the analysis |

Source: prepared by the author, 2026.

The request produced seven documents, one per phase, in addition to the final report and the audit appendix. In this demonstration, the full route took less than an hour (an informal estimate), compared with weeks of manual work. The use of GenAI in the qualitative method, however, cannot rest on the efficiency gain alone: the quality of the evaluation itself must also be verified.

Before the detailed account, a note on how the phases correspond. With one exception, each phase of the *skill* corresponds to one of the six phases of Braun and Clarke's reflexive thematic analysis: Phase 1 is familiarization with the data; Phase 2, generating initial codes; Phases 3 and 4, searching for and reviewing themes; Phase 5, defining and naming themes; and Phase 6, producing the report. Phase 0 is the exception: it has no counterpart in Braun and Clarke's framework and is a governance addition to record decisions before the workflow begins.

The first document records the analysis decisions, such as the research question, the corpus, the unit of analysis, the theoretical position, and the inductive/deductive orientation, as well as how much would be delegated in each phase. In the demonstration, the research question asked how participants from different state capitals and income brackets construct meanings about democracy, the state, corruption, social protection, and moral life, and how socioeconomic position cuts across those meanings. The theoretical position was declared to be critical-interpretive contextualist, and city and income entered as contrast attributes, not as themes. The *skill* does not proceed until this record is in place.

Phase 1, initial reading, corresponds to Braun and Clarke's first phase, familiarization with the data, and comes before any coding. The agent inventoried the corpus (13 files: the 12 transcripts plus the session index), worked through the twelve sessions spread across Belém, Recife, and Curitiba in the four income brackets, and produced notes for each session covering the group profile, surprises, and passages of interest, without generating a single code.

Phase 2, codes, corresponds to Braun and Clarke's second phase, generating initial codes: each code was anchored in an excerpt with a precise location. The agent generated 54 codes, from C01 to C54, each with a meaning, a representative excerpt, a corpus locator, and an interpretive note. For example, C01, "Democracy as electoral procedure and abstract freedom," was based on "democracy is having freedom of expression, voting for whomever you want, free will" (BEL-0-2, line 16); C04, "Authoritarianism as imposition and 'not having a voice,'" drew on "it does not give you a voice, where you have no rights, no choices" (BEL-5-10, line 13); and C42 and C43, "Social programs as necessary with controls" and "Social programs as 'accommodation' and manipulation," offered competing readings of the same material on social protection. The distribution check confirmed that the codes covered the twelve sessions.

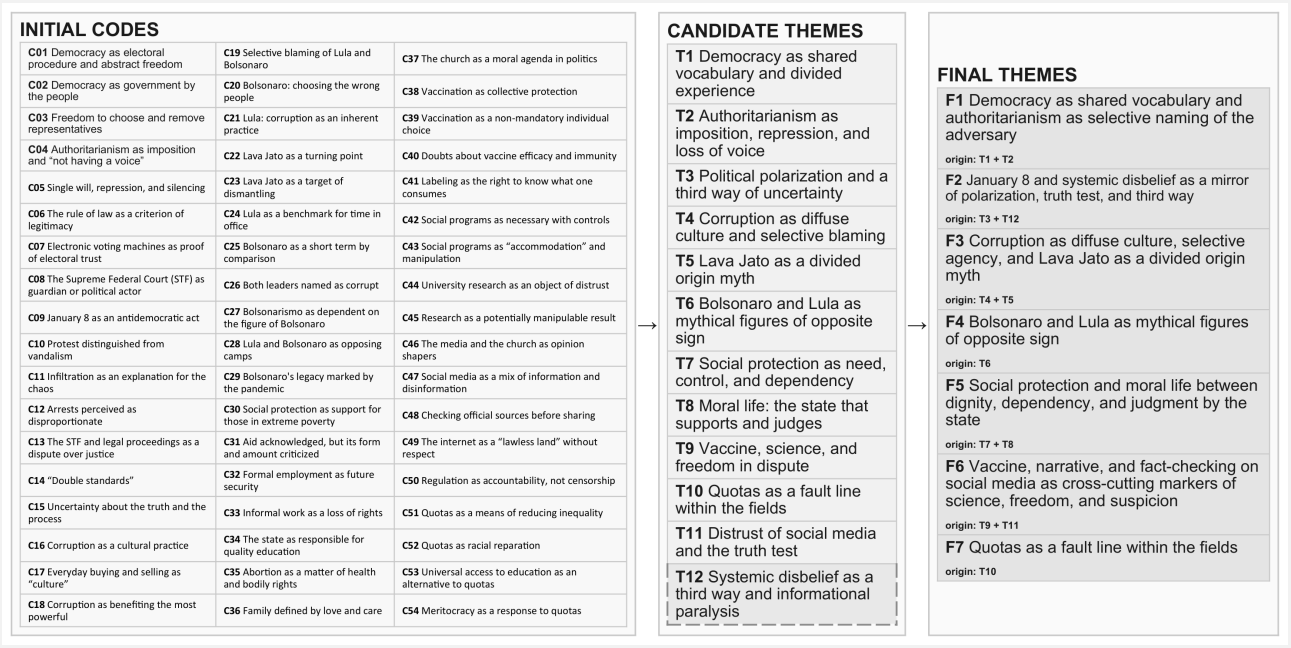
Phase 3, candidate themes, corresponds to Braun and Clarke's initial search for themes. In this phase, the agent built candidate themes from patterns of shared meaning, not by grouping the material by subject matter. Each candidate theme had to carry an organizing concept and an interpretive claim, never a domain label; income and city were rejected as themes because they function as cross-cutting dimensions, not as organizing keys. The map of 54 codes yielded 11 candidates, from T1 to T11. The first, "Democracy as shared vocabulary and divided experience," argued that almost everyone defines democracy by the same formula (voting, freedom, choice) but that its practical meaning diverges according to experiences of electoral defeat, institutional trust, and income. Whatever did not fit a candidate theme went into a provisional folder of residues rather than being discarded.

Phase 4, contestation, did not accept the thematic map at face value. The integrating agent required a plausible alternative reading and a test of contradictions and residues. The rival reading took the form of a hypothesis: the master organizer of meanings is neither income nor city but the electoral position declared in the paired sessions, so that income and city would remain the context in which material experience is produced. The rival reading did not overturn any theme; that was not its role. Its role was to force scrutiny, and it prompted several changes: T3 was revised to incorporate the third way ("I don't know if it was true"); T9 and T10 were strengthened with cases of dissent; T11 was split because it mixed two distinct movements, that of building a narrative and that of distrusting fact-checking; and a new theme emerged, T12, "Systemic disbelief as a third way and informational paralysis."

Twelve themes reflect analytical breadth but also difficulty in synthesis. For this reason, the *skill* itself sets a rule of thumb: aim for four to eight themes when the corpus allows. Aggregation was not a matter of arithmetic: each merger had to demonstrate that two themes spoke of the same phenomenon. Thus T1 and T2 merged: naming democracy and naming authoritarianism are the same act of taking a position in the political lexicon. T3 and T12 merged: systemic disbelief is the face of the same polarization that January 8 exposed, and the third way of T3 was the paralysis of T12 in action. T4 and T5 merged: Lava Jato is the origin myth of the same diffuse culture of corruption that assigns blame selectively. T7 and T8 merged: social protection and moral life are the two faces of the same relationship, that of the state that supports and judges the same population. T9 and T11 merged: vaccine, narrative, and fact-checking are variations of the same

epistemic suspicion, and distrust of social media migrated to the theme of the truth test. T6 and T10 were not merged, owing to saturation and demonstrated distinctiveness. The twelve initial themes became seven (Figure 1), a more substantive approximation that makes it easier to tell the story the themes contain, as described by Braun and Clarke.

**Figure 1** – Flow of theme construction, from initial codes to candidate themes to final themes




Source: prepared by the author, 2026.

Phase 5, defining and naming themes, followed Braun and Clarke's method and was applied to the seven remaining themes. Each theme now had to present a name, an organizing concept, an interpretive claim, a genealogy (from the 12 preceding themes down to the codes), key excerpts for interpretation, the rejected rival reading, and the zone of uncertainty. Before that, the excerpts underwent verbatim verification against the original text. The verification confirmed that the 40 candidate quotations do exist in the original corpus; however, none of them is reproduced verbatim here, as they were paraphrased to prevent the identification of participants.

The process shows that the seven themes were not discovered in the corpus; they were constructed, and constructed through contestation. The final themes were as follows: Democracy as shared vocabulary and authoritarianism as selective naming of the adversary; January 8 and systemic disbelief as a mirror of polarization, truth test, and third way; Corruption as diffuse culture, selective agency, and Lava Jato as a divided origin myth; Bolsonaro and Lula as mythical figures of opposite sign; Social protection and moral life between dignity, dependency, and judgment by the state; Vaccine, narrative, and fact-checking on social media as cross-cutting markers of science, freedom, and suspicion; and Quotas as a fault line within the fields.

Phase 6, the report, which corresponds to Braun and Clarke's sixth phase, was the stage that tied everything together. The text told the analytical story in a main body and an auditable appendix. In the main body, the research question ties the corpus to the narrative of the seven themes, each with interpreted anchor excerpts, variations across sources, and the rejected rival reading recorded alongside the prevailing thesis. The appendix gathered the complete, contestable theme records for each final theme, with their genealogy, the verified anchor excerpts, the interpretive rival, the decisive evidence, the residues, the limits, and the uncertainties of the analysis. It also included the AI use statement, the delegation matrix by phase, the interaction logs, the human map, the coverage audit, the verification of the excerpts, the checkpoints, and the adversarial decisions.

A further check on coding quality is the verification step, which locates every cited excerpt in the original text and separates what is identical from what needs normalization; only identical text is placed in quotation marks. Without this check, which is intrinsic to the *skill*, the AI would be freer to fabricate excerpts or to introduce incorrect or imprecise ones into the final evaluation. This rule applies to the analysis carried out on the original corpus; in this publication, the excerpts were paraphrased to prevent the identification of participants.

## Evaluation of the method

The *skill* underwent a blind comparative test of 126 runs: sixty-three without the *skill* and sixty-three with it. The twenty-one test scenarios simulate the shortcuts that the protocol should reject, such as skipping the initial reading or treating a topic summary as a theme. An AI agent based on the GLM 5.2 model from Zhipu AI, which had not been involved in generating the outputs, evaluated each output against 14 criteria, checking items such as the presence of literal excerpts and the separation between proposal

and decision. In the blind evaluation, the *skill* condition adhered more closely to the methodological criteria, and the AI agent rated the outputs of the *skill* condition as better. The result is favorable, though it does not establish validity: no test replaces human evaluation.

## DISCUSSION

The *skill* narrows the gap between plain chat and the programmatic ("agentic") workflow without requiring the researcher to program. The reflexive thematic analysis protocol, spread across seven phases with checks, routes, and contestation, becomes observable in the agent's behavior, and each phase produces materials that can be examined, audited, and contested. The gain is that third parties can reconstruct the process, which advances principles of replicability without replacing human judgment.

In my view, artificial intelligence is already used in qualitative research, in research projects, in laboratories, and in advising, without being disclosed and without adequate documentation. The *skill* seeks to address this problem by organizing criteria, phases, and checks into a single artifact that any researcher can inspect, adapt, and criticize. A reviewer who disagrees with an analysis can install the *skill*, apply it with different parameters, and identify where the interpretation diverges; an advisor can use it as a teaching tool. In practice, the user can even modify the *skill* however they wish, or design one from scratch simply by asking the AI agent. In this initial test, I chose to evaluate the *skill* using only two AI agents; developing criteria for human validators to assess quality is certainly one way forward.

The promise of transparency finds fertile ground in qualitative research in health and nursing, a field in which demands for transparent reporting are growing[3,13]and thematic analysis has become increasingly common for analyzing data such as interviews and focus groups[14].

Certain limitations of such analyses must, of course, be acknowledged. Artificial intelligence has no lived experience, situated judgment, or human reflexivity, and checking quotations does not guarantee interpretive quality. The inventory of files generated by the *skill* confirms each document's identity, not what the agent read in it, and agreement among agents does not amount to human triangulation. Costs and subscriptions continue to create barriers, and expensive plans weigh more heavily on Brazilian graduate students and faculty[9]. Sensitive data require extra care, because participants' utterances processed by third-party services require checking the original consent and the applicable personal data protection rules. Processing by foreign services, in turn, involves international data transfer, which the researcher must assess under the terms of each provider. The decision about what can leave the computer rests with the researcher and the research ethics committee[15].

The autonomous route, in which the agent goes through the phases without intervention, remains experimental and does not claim equivalence with coding. It shows what happens when the machine runs the entire protocol and produces material that another researcher can audit, contest, or use as an initial or even a rival reading.

## CONCLUSION

This article set out to present a fairly simple way to run an automated thematic analysis: one need only install the agent program, download the *skill*, open the folder with the material to be analyzed, and request the analysis in natural language. This path seems to me innovative and promising, since it eases access for any researcher without requiring programming knowledge, while also generating materials to audit, question, or even ground the analysis. Those who used to paste excerpts into a chatbot can now request the entire analysis in natural language and show, document by document, where each theme and each utterance came from.

The demonstration on a real corpus showed that the complete workflow is feasible, with full coverage and every quotation checked against the original. The blind test is consistent with a reduction in methodological deviations relative to the same model without guidance, and the portability check showed that the package does not depend on a single provider. None of this evidence justifies claims of human equivalence or automatic validity; it supports feasibility, transparency, and auditability, properties that qualitative research must uphold in the era of artificial intelligence.

The obvious limitation at this stage is the lack of validation with one or more human researchers and against high-quality thematic analyses, which will be carried out in the next stages of the research. Validation by two distinct AI agents also has limits, since it may merely show that these technologies share similar biases. Nevertheless, several of my own experiments and the validation experiment do show better performance in coding qualitative aspects.

The scientific value of a qualitative agent lies in the extent to which it allows others to reconstruct, contest, and transform what it produced. The final interpretation and the responsibility for each statement remain human. The tool should help the researcher think better, never think for the researcher.

## CONFLICT OF INTEREST

The author declares no conflicts of interest.

## FUNDING

This work was carried out with the support of a CNPq productivity grant.

## USE OF ARTIFICIAL INTELLIGENCE

The *skill* described in this article was developed and evaluated using generative artificial intelligence tools. ChatGPT Luna 6 was used to test the thematic analysis *skill* and also for corrections and improvements to the text, and it also assisted with translation. The writing of this manuscript was assisted by these tools, with final review and editing by the author, who takes full responsibility for the content.

## REFERENCES

1. Paulus TM, Marone V. "In minutes instead of weeks": discursive constructions of generative AI and qualitative data analysis. Qual Inq. 2025;31(5):395-402. https://doi.org/10.1177/10778004241250065
2. Jowsey T, Stapleton P, Campbell S, Davidson A, McGillivray C, Maugeri I, et al. Frankenstein, thematic analysis and generative artificial intelligence: quality appraisal methods and considerations for qualitative research. PLoS One. 2025;20(9):e0330217. https://doi.org/10.1371/journal.pone.0330217. PMID: 40911617.
3. Kempny C, Frings J, Rust P, Meister S, Fehring L. The use and methodological reporting of large language models in qualitative research: a scoping review. BMC Med Res Methodol. 2026;26(1):137. https://doi.org/10.1186/s12874-026-02913-1. PMID: 42298426.
4. Naeem M, Smith T, Thomas L. Thematic analysis and artificial intelligence: a step-by-step process for using ChatGPT in thematic analysis. Int J Qual Methods. 2025;24. https://doi.org/10.1177/16094069251333886
5. Wachinger J, Bärnighausen K, Schäfer LN, Scott K, McMahon SA. Prompts, pearls, imperfections: comparing ChatGPT and a human researcher in qualitative data analysis. Qual Health Res. 2025;35(9):951-66. https://doi.org/10.1177/10497323241244669. PMID: 38775392.
6. Sampaio R. Análise temática reflexiva: skill para agentes de codificação [software]. 2026 [cited 2026 Aug 24]. Available from: https://github.com/labiia-lab/analise-tematica-reflexiva
7. Dellafiore F, Saba A, Collaro C, Artioli G. Artificial intelligence in qualitative research: insights from experts via reflexive thematic analysis. Qual Health Res. 2026;36(2-3):145-65. https://doi.org/10.1177/10497323251389800. PMID: 41351219.
8. Jaeger AA, Costa AP. Análise de dados guiada pelo modelo abductivai: um estudo de colaboração humano-IA. Rev Pesqui Qual. 2026;14(42):168-97. https://doi.org/10.33361/RPQ.2026.v.14.n.42.1629
9. Sampaio RC. Agentes de IA na pesquisa acadêmica: um tutorial para pesquisadores sem programação. Salvador: Edufba; 2026. Forthcoming.
10. Braun V, Clarke V. Using thematic analysis in psychology. Qual Res Psychol. 2006;3(2):77-101. https://doi.org/10.1191/1478088706qp063oa
11. Braun V, Clarke V. Reflecting on reflexive thematic analysis. Qual Res Sport Exerc Health. 2019;11(4):589-97. https://doi.org/10.1080/2159676X.2019.1628806
12. Souza LK. Pesquisa com análise qualitativa de dados: conhecendo a análise temática. Arq Bras Psicol [Internet]. 2019 [cited 2026 Aug 24];71(2):51-67. Available from: http://pepsic.bvsalud.org/scielo.php?script=sci_arttext&pid=S1809-52672019000200005
13. Nowell LS, Norris JM, White DE, Moules NJ. Thematic analysis: striving to meet the trustworthiness criteria. Int J Qual Methods. 2017;16(1):1-13. https://doi.org/10.1177/1609406917733847
14. Dias EG, Mishima SM. Análise temática de dados qualitativos: uma proposta prática para efetivação. Sustinere. 2023;11(1):402-11. https://doi.org/10.12957/sustinere.2023.71828
15. Sampaio RC, Sabbatini M, Limongi R. Diretrizes para o uso ético e responsável da inteligência artificial generativa: um guia prático para pesquisadores [Internet]. São Paulo: Intercom; 2024 [cited 2026 Aug 24]. Available from: https://www.portcom.intercom.org.br/ebooks/arquivos/livro-diretrizes-ia.pdf

**AUTHORSHIP CONTRIBUTIONS**

Study design: Sampaio RC.
Data collection: Sampaio RC.
Data analysis: Sampaio RC.
Data interpretation: Sampaio RC.
The author is responsible for writing and critically reviewing the intellectual content, approving the final published version, and all ethical, legal, and scientific aspects related to the accuracy and integrity of the study.

**© 2026 Authors.** The authors retain copyright and grant the Online Brazilian Journal of Nursing the right of first publication. This article is licensed under the Creative Commons Attribution 4.0 International License (CC BY 4.0).