

# Reconciling Methodological Paradigms: Employing Large Language Models as Novice Qualitative Research Assistants in Talent Management Research

Sreyoshi Bhaduri  
Amazon  
Arlington, Virginia, USA

Satya Kapoor  
Amazon  
Vancouver, British Columbia, Canada

Alex Gil  
Amazon  
Arlington, Virginia, USA

Anshul Mittal  
Amazon  
Arlington, Virginia, USA

Rutu Mulkar  
Amazon  
Seattle, Washington, USA

## Abstract

Qualitative data collection and analysis approaches, such as those employing interviews and focus groups, provide rich insights into customer attitudes, sentiment, and behavior. However, manually analyzing qualitative data requires extensive time and effort to identify relevant topics and thematic insights. This study proposes a novel approach to address this challenge by leveraging Retrieval Augmented Generation (RAG) based Large Language Models (LLMs) for analyzing interview transcripts. The novelty of this work lies in strategizing the research inquiry as one that is augmented by an LLM that serves as a novice research assistant. This research explores the mental model of LLMs to serve as novice qualitative research assistants for researchers in the talent management space. A RAG-based LLM approach is extended to enable topic modeling of semi-structured interview data, showcasing the versatility of these models beyond their traditional use in information retrieval and search. Our findings demonstrate that the LLM-augmented RAG approach can successfully extract topics of interest, with significant coverage compared to manually generated topics from the same dataset. This establishes the viability of employing LLMs as novice qualitative research assistants. Additionally, the study recommends that researchers leveraging such models lean heavily on quality criteria used in traditional qualitative research to ensure rigor and trustworthiness of their approach. Finally, the paper presents key recommendations for industry practitioners seeking to reconcile the use of LLMs with established qualitative research paradigms, providing a roadmap for the effective integration of these powerful, albeit novice, AI tools in the analysis of qualitative datasets within talent management research.

## Keywords

Retrieval Augmented Generation (RAG), AI in Talent Management, Qualitative Research

## 1 Introduction

Talent management researchers frequently work backwards from their customers, the employees at the organization. Understanding employee sentiment and behavior often involves conducting deep-dive interviews, explanatory in nature – e.g., demystifying the why behind customer choices, attitudes or behaviors (e.g., [22]). Talent management research, at its core, seeks to use science to equip

every employee with resources to help them best navigate their careers [40].

Consequently, qualitative research methodology plays a critical role in talent management. Many of the key considerations around employee engagement, motivation, and workforce culture involve subjective, context-dependent factors that are best explored through in-depth interviews, focus groups, and other qualitative data collection approaches. Talent management professionals often rely on rich qualitative datasets to gain deep insights into employee experiences, organizational dynamics, and the nuances of human capital. However, these qualitative paradigms can clash with the more positivist, quantitative worldview that underlies many of the analytic tools used to evaluate talent management data. Talent management researchers may find that standard statistical techniques and data visualization approaches struggle to fully capture the complexities inherent in qualitative datasets, leading to potential misinterpretations or oversimplifications of the human elements involved in managing an organization’s workforce. Navigating this tension between qualitative and quantitative approaches is an ongoing challenge for talent management professionals.

Large language models (LLMs) like BERT, GPT-3 and PaLM have demonstrated strong aptitude for summarization (e.g., [37]), classification (e.g., [31]), and information extraction (e.g., [13]) for text-based data. Consequently, LLMs are also increasingly being leveraged within talent management contexts for tasks such as interview analysis. However, language models are themselves designed primarily from a quantitative, data-driven paradigm. These models are trained on vast troves of text data using statistical machine learning techniques optimized for numerical patterns and correlations. While powerful at extracting insights from large-scale datasets, LLMs can often struggle to fully capture the nuanced, contextual nature of language [1], [14] that is critical for qualitative information sourced from interviews, focus groups, and other qualitative research methods common in talent management.

Talent management professionals must therefore continuously navigate a tension between the quantitative orientation of their analytical tools and the qualitative richness of the human dynamics they seek to understand. Bridging this gap requires innovative approaches that combine the opportunity for scale and speed offered by LLM-powered analysis augmented by borrowing evaluative nuances of traditional qualitative techniques. Talent leaders, thus,

must carefully select and configure their AI-powered tools to ensure the voices and experiences of employees are authentically represented, rather than reduced to oversimplified metrics. Mastering this balance is an ongoing challenge, but one that is critical for talent management to yield truly holistic and impactful insights.

This paper presents results from leveraging LLMs as a novice qualitative researcher to augment qualitative research workstreams, specifically for data generated through semi-structured interviews.

The purpose of this paper is two-fold – 1) provide an overview of a successful implementation of a Retrieval Augmented Generation-based model for analyzing semi-structured interviews, and more importantly, 2) enumerate pragmatic take-aways and learnings drawing from traditional qualitative research to help fellow industry practitioners in reconciling the methodological paradigms. We posit the second purpose to be valuable to the larger discussion within talent management research communities on how and where to integrate AI capabilities across different talent management workstreams.

## 2 Quantitative and Qualitative Paradigms

Quantitative and qualitative research represent two fundamental paradigms or philosophical frameworks that guide research strategies, methods, analysis, and use of results [38]. While both methodological approaches seek to rigorously study research problems, they are based on distinct assumptions and procedures adapted to investigating particular types of questions and drawing different conclusions. Quantitative research is based on the assumptions of positivism, the philosophical tradition premised on the application of natural science methods to the study of social reality and beyond [6]. Quantitative researchers believe that objective facts and truths about human behavior and society can be measured and quantified numerically. Quantitative methods such as surveys, structured observations, and experiments aim to test hypotheses derived from theories by examining relationships between precisely measured variables statistically analyzed using large sample sizes [10]. These methods seek to minimize subjectivity and generalize findings to a population. In contrast, qualitative research aligns with interpretivist and constructivist philosophical traditions by embracing subjectivity and focused meaning-making by and with research participants [11].

Qualitative researchers often use an inductive approach aimed at discovering and understanding processes, experiences, and worldviews by collecting non-numerical data through methods like in-depth interviews, ethnographic fieldwork, and document analysis. Findings derive from themes that emerge openly from the data rather than testing predetermined hypotheses. Samples tend to be small and purposely selected to illuminate a phenomenon in depth and detail. The aim is particularization rather than generalization, with a priority on ecological validity and multiple realities situated in time, place, culture, and context.

While debates once positioned these paradigms in opposition, contemporary mixed methods research leverages the complementary strengths of quantitative and qualitative approaches [18]. Mixed methods investigations integrate quantitative and qualitative data collection and analysis within a single program of inquiry by combining these approaches in creative ways to deepen understanding

[8] [9] [17]. This reconciliation of methodological perspectives offers opportunities to generate more robust, contextualized insights to address complex research problems. The use of large language models (LLMs) as novice qualitative research assistants, as explored in this paper, can be considered an exercise in mixed methods research design.

Prior to LLMs, in previous work, Natural Language Processing based modeling of qualitative data from social science contexts, have also been used as "novice insight" augmented by the more expert contextualization provided by human researchers (e.g., [3], [4]). Popular traditional topic modeling techniques (e.g. Latent Dirichlet Allocation), however, suffer from several limitations (e.g. specifying number of clusters) when compared to existing deep learning-based methods. They also often fail to capture the contextual nuances and ambiguities inherent in natural language, as they rely heavily on predefined rules and patterns [12] [32]. This can make it challenging to handle the complexities and variations present in real-world text data, and may require domain-specific knowledge or fine-tuning to achieve acceptable performance [21]. Recent advancements in LLMs, such as BERT and GPT, have largely overcome these limitations by leveraging deep neural networks to learn rich, contextual representations from large amounts of text data [35] [12]. These powerful models can capture subtle semantic and pragmatic features of language, and demonstrate strong generalization capabilities through transfer learning [5] [32].

Further, in traditional qualitative research, thematic analysis is the process of gathering themes across topics from qualitative data, such as interview data, through iteratively analyzing the dataset for topics of interest [8]. Inductive coding and deductive coding are two approaches to analyzing data from semi-structured interviews. Inductive coding involves starting with raw data and gradually developing codes and categories based on patterns and topics that emerge from the data as the researcher manually interacts with it [29] [34]. This approach is bottom-up, where the data drives the development of codes and theories [16]. Deductive coding, on the other hand, involves starting with preconceived codes or theories and applying them to the data [30]. This approach is top-down, where existing theories or frameworks guide the coding process [26]. Researchers in industry typically work backwards from research question of interest. Most of the research questions in industry driving qualitative data collection are also explanatory (i.e., tend to explain the quantitative findings such as low customer satisfaction, low product adoption numbers), rather than exploratory (i.e., ethnography of a community of interest or a phenomenon) and as a result deductive approaches are often more popular than inductive coding.

Ultimately, by augmenting traditional deep-dive qualitative analysis with the time and resource efficient pattern recognition and text processing capabilities of LLMs, researchers can integrate quantitative and qualitative techniques to enhance the speed, depth, and rigor of their investigations. This mental model of a novice-LLM approach holds promise for bridging the divide between positivist and interpretive paradigms, ultimately working towards a more comprehensive understanding of the phenomenon under study.

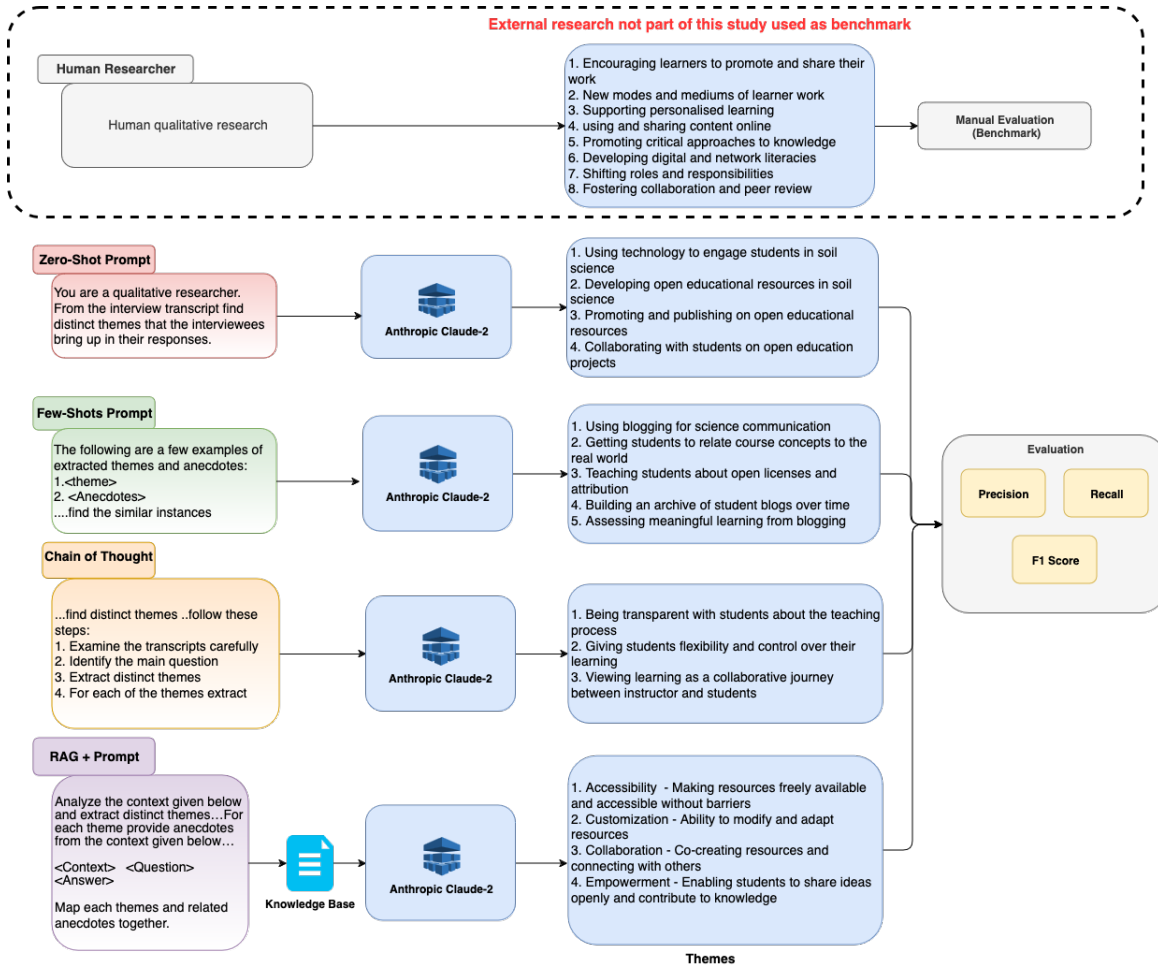


Figure 1: Comparison across prompting approaches

### 3 Dataset

We used an open-source dataset [28] to demonstrate how an LLM prompted as a novice researcher can enhance traditional qualitative deductive thematic coding. This dataset was originally collected to explore educators’ experiences implementing open educational practices [28]. The dataset contains eight transcripts each from hour-long interviews conducted with educators to understand how they are using openly accessible sources of knowledge and open-source tools. The original research involved a deep-dive qualitative analysis through using a phenomenological approach to extract topics manually from the dataset. We chose this open-source dataset for two reasons – 1) structural match to proprietary dataset, and 2) rich description and manually identified topics by an expert to serve as a gold standard to measure the efficacy of our LLM based approach. Semi-structured interviews provide critical insights through participant perspectives, making them foundational in various industry settings.

The semi-structured approach used to create this dataset is a close match to proprietary talent management data from our organization, where employees are interviewed on a particular phenomenon to get deeper understanding of their related sentiment, attitudes, and behaviors. Manually extracted topics serve as gold standard for benchmarking findings from our LLM-based approach. The paper [28] describing the dataset explains the manual process establishing how each transcript was read twice: first, for a comprehensive analysis, and subsequently, to initiate a thematic exploration. Additional reviewing continued as codes and topics emerged and intersected among the interviews. A manual qualitative coding approach was applied at each iteration to reveal themes, following constant comparison methodology [16].

We posit that our approach, as demonstrated on this sample semi-structured interview dataset, can easily extend to multiple industry settings in talent management research where researchers conduct interviews and focus groups.

**Interviewer:** Are you encouraging your students to share and engage the community or larger audiences with their work, then, sounds like.

**Interviewee:** Yeah. So, that's definitely a big part of my practice is thinking around this term in Indigenous film called screen sovereignty which I also try to think of because I do a lot of podcasting and audio work, so audio sovereignty. And that is really considering audience from the get-go. So you're building a project and thinking about relationship to community. So, yeah, that's often where we start. It's like, you are not just writing for me, your prof. Who is this project for? Who do you want it to speak to? And how can we find ways to make that viable?

**Interviewer:** How do find students respond to that? And is it untraditional from other work they might be doing?

**Interviewee:** Usually, I mean, always really well. I just finished a podcasting assignment right now, they just are turning them in as we speak and the feedback has just been wonderful. It has created amazing community in the classroom, people are really loving the work their partners are making.

Figure 2: Sample of the interview transcript

## 4 Thematic Analysis Using LLMs

In traditional, manual qualitative research, deductive thematic analysis process begins with the researcher first formulating the research questions. Then, upon collection of the data, such as interview transcripts, the researcher iterates manually through the transcripts to identify and extract themes or topics of interest. This labor-intensive process involves carefully reading through the data, taking notes, and organizing the topics iteratively into broader coherent themes that address the research questions. The researcher may go through multiple rounds of coding and analysis to refine the themes and ensure they comprehensively capture the key insights from the data. Our approach finds that LLMs can quickly uncover topics of interest from the dataset which can then be iterated upon to garner broader themes of interest across topics. Thus, for our novice-LLM led approach, we leveraged the power of Large Language Models (LLMs) as a novice research assistant in the thematic analysis process. Specifically, we used the open-source framework called Langchain to create dynamic prompt templates, such as few-shot prompts and chain of thoughts, that guided the LLM in performing topic modeling and generating insights from the interview transcripts. We then opted to use Anthropic's Claude2 model to execute these prompts and extract the relevant themes.

To initiate the analysis, we first selected a main research question and corresponding sub-questions from our dataset [28]. We then fed these research questions, along with the interview transcripts, into the LLM-powered Langchain framework. The model was able to quickly identify and summarize the key topics, and iteratively, themes emerging from the data. This approach provided a quick yet relatively comprehensive analysis that would have taken a human researcher significant time and effort to reproduce manually.

### 4.1 Thematic analysis enhanced through Retrieval Augmented Generation (RAG)

In our LLM based approaches, we experiment with four methods - zero-shot prompting, few-shot prompting, chain-of-thought reasoning, and Retrieval Augmented Generation based Question Answering. In zero-shot prompting we provide a single prompt to the model. In few-shot prompting, we provide a set of topics

and anecdotes to the model as examples. In the chain of thought (COT) approach, we provide a set of instructions for the model to follow. Finally, for Retrieval Augmented Generation (RAG) we provide context and questions to the model, from which it extracts information.

Zero-shot prompts are simple instructions or tasks given to an LLM that have not been specifically trained on that task. It serves as a baseline because it demonstrates the model's fundamental ability to understand and respond to prompts based solely on its pre-training [19]. In few-shot prompting, a small set of examples illustrating the desired outcome are manually selected and provided to the LLM. These examples allow the model to understand the tasks at hand and generate similar results [5]. Chain-of-thought prompting provides a set of intermediate steps to guide the LLM to mimic human-like reasoning. This significantly improves the capability of the LLM to understand complex reasoning and generate better topics [36]. Retrieval-augmented generation (RAG) combines the capabilities of an LLM with a retrieval system to source and integrate additional information into its responses [23]. This effort provides contextually richer and ultimately more accurate outputs. We do this by providing all the interview transcripts to the LLM as a custom knowledge base. Two considerations helped the RAG approach outperform the other approaches:

**4.1.1 Focused Analysis:** In our approach, LLM searches the knowledge base to find and retrieve parts of documents that are most relevant to the question in the query. This narrows the focus to the most relevant information and ensures attention to critical topics and nuances.

**4.1.2 Context Dilution/Managing Information Overload:** Using all transcripts as input in a single instance creates information overload scenarios, ultimately leading to dilution of important topics or nuances. If the dataset is too large or complex, LLM might lose track of what's most relevant to specific query, leading hallucinations. Hallucinations or inaccuracies within this context refers to instances where the model generates information which is not grounded in input data. In our approach, the use of RAG mitigates some of the hallucination by anchoring LLM responses relevant

information, and providing a form of contextual validation for the output.

## 5 Findings

In the paper describing the dataset leveraged for this work, the authors collected and conducted a manual analysis [28]. Their research led to identification of significant, recurring topics within the interviews. Our evaluation strategy uses these manually generated topics from the paper’s work as gold standard to compare against topics generated by the LLMs-based approach. We use Precision (Equation 1), Recall (Equation 2), and F1-score (Equation 3) to benchmark topics generated by our LLM-augmented qualitative research approach against the topics generated by the human researcher.

$$R_{BERT} = \frac{1}{|x|} \sum_{x_j \in x} \max_{\hat{x}_i \in \hat{x}} x_i^T \hat{x}_j \quad (1)$$

$$P_{BERT} = \frac{1}{|\hat{x}|} \sum_{\hat{x}_i \in \hat{x}} \max_{x_j \in x} \hat{x}_i^T x_j \quad (2)$$

$$F_{BERT} = 2 \cdot \frac{P_{BERT} \cdot R_{BERT}}{P_{BERT} + R_{BERT}} \quad (3)$$

$$\text{Cosine Similarity} = \frac{x_i^T \hat{x}_j}{\|x_i\| \|\hat{x}_j\|} \quad (4)$$

These metrics are the current evaluation standard for classification models, but they can be adapted for text generation tasks [39]. Precision and Recall measure the proportion of correctly identified positive cases. In the context of our experiment, every word from predicted text gets matched to a word in the referenced text to compute recall. This process is inverted to then compute precision. The precision and recall values are then combined to compute an F1 score. These metrics use cosine similarity (Equation 4) in which each predicted word is paired with its closest corresponding word from the reference text with the aim of maximizing the similarity score.

In Table 1, the performance of various LLM prompting techniques including Chain of Thought, Few Shot, Zero Shot and RAG, are compared across different embedding models (Distillbert-base-uncased, Bert-base-uncased, and Roberta-large). This comparison aims to evaluate the robustness and effectiveness of these prompting techniques. Our results indicate that while each prompting technique shows varying level of precision, recall and F1-score, RAG consistently outperform the others on all three metrics, achieving highest performance across all models.

## 6 Learnings

Treating large language models (LLMs) as novice research assistants during thematic analysis offered valuable insights for our research. By framing the LLM as a novice collaborator with little knowledge or insight of the context, prompts can be crafted to better guide the model and leverage its capabilities. Used prudently, similar novice LLM-augmented approaches can significantly increase time and resource efficiency compared to traditional qualitative coding methods in talent management research. The following sections explore

some of our key learnings that may benefit other researchers considering designing LLMs as novice researchers to optimize thematic analysis.

### 6.1 Approaching LLMs as Novice Research Assistants can help prepare better prompts

A novice is a person who, “has no experience with the situations in which they are expected to perform tasks” [2]. The novice is thus at a basic proficiency level for skill acquisition, with limited information and prior experience related to a task at hand [27]. For large qualitative datasets analyzed using LLMs we propose that a novice-led approach to analysis is a good fit. In our approach the human behaves as an expert prompting the novice LLM to provide insights related to topics of interest. We found this framework as a helpful mental model to ground the primary researcher prompting the LLM as they iteratively uncover insights from the dataset.

### 6.2 Used prudently, LLMs can help increase time effectiveness and resource efficiency

LLMs have advanced the field of natural language processing with their ability to understand and generate responses that closely mimic human language [33]. The strengths of LLMs extend beyond metrics, these models are adept at processing vast amounts of text rapidly, demonstrating a level of topic modeling that can mimic human analysis. Manual topic modeling is human labor intensive and time inefficient [7]. LLMs also enhance efficiency by streamlining the processing of large datasets, allowing for the extraction of topics from qualitative data more quickly. Improvisations of these model using techniques like few-shot and zero-shot learning capabilities further reduce the need for expensive data labeling and annotations. In a nutshell, LLMs boost speed, reduce human effort, scale to massive datasets, and lower labeling costs. However, human expertise is still essential for judgment, validation and end-to-end framework design.

### 6.3 LLM augmented approaches offer significant increase in ease and enhanced context compared to traditional NLP approaches.

Using a RAG approach towards an LLM-augmented qualitative research analyzing semi-structure interviews shows great promise compared to natural language processing methods like Latent Dirichlet allocation (LDA). Currently, there are no widely accepted methods for comparing the two approaches as there is no bridge to compare keywords to themes, except from a human-evaluator ease of interpretability standpoint. We performed topic modeling analysis on the same dataset with the broader aim of finding themes. Manually comparing both approaches, each researcher of this workflow independently found that any of the approaches using an LLM yielded much greater context and consequently, better interpretability than the traditional LDA approach. This is likely because, with LDA, the model outputs a list of words and probability for each topic. With these words, the researcher would then have to manually define the topic. While this approach increases researcher flexibility, it remains time and resource consuming. In contrast,

<b>Distillbert-base-uncased</b>	Precision	Recall	F1-Score
Chain of Thought	67%	62%	64%
Few Shot	72%	67%	70%
Zero Shot	68%	66%	67%
<b>RAG</b>	<b>79%</b>	<b>80%</b>	<b>79%</b>

<b>Bert-base-uncased</b>	Precision	Recall	F1-Score
Chain of Thought	56%	48%	52%
Few Shot	64%	56%	60%
Zero Shot	59%	55%	57%
<b>RAG</b>	<b>70%</b>	<b>70%</b>	<b>70%</b>

<b>Roberta-large</b>	Precision	Recall	F1-Score
Chain of Thought	89%	85%	87%
Few Shot	90%	87%	88%
Zero Shot	89%	86%	88%
<b>RAG</b>	<b>92%</b>	<b>91%</b>	<b>91%</b>

**Table 1: Comparison of Results across the LLM Enhanced Thematic Analysis Strategies Employed**

Example: Keywords from LDA Topic One

Students  
 Course  
 Develop  
 People  
 Institution  
 Project  
 Science  
 Discipline  
 Material  
 Start

Example: Output from LLM approach

Collaboration: Co-creating resources and connecting with others

*Corresponding Anecdote: You can also in your teaching have students connect with people outside the course in various ways. Like, maybe some people outside the course are commenting on blogs and student are getting in a conversation around that.*

**Table 2: Example Output from LLM approach compared to Keywords from LDA Topic One**

with the LLM approach, the output is richer in context of what particular topics mean. For example, our LDA model yielded 5 topics (see: Appendix A Figure 3). The first 10 words for topic 1 can also be seen in Table 2. Putting these words together into a comprehensive

theme can be challenging without more context. However, an LLM is able to generate context grounded in the participant's voice for researchers to work with. An example of an extracted theme and

its corresponding anecdote using an LLM can also be seen in Table 2, above.

## 7 Recommendations

Traditional qualitative research is evaluated based on several criteria that ensure quality and rigor of the research, both in terms of methods as well as findings. Prior research has established four criteria for increased rigor and trustworthiness of qualitative research studies around credibility, dependability, confirmability, and transferability [24]. We recommend three ways in which quality criteria from traditional qualitative research can be used by practitioners employing LLM augmented analysis of qualitative data.

### 7.1 Establishing credibility of findings by incorporating mechanism for member checks.

Member checks, i.e., the strategy of soliciting insights from research participants on research findings, are often relied on as the gold standard for increasing trustworthiness of qualitative research approaches (e.g., [29] [20]). Qualitative researchers employing LLMs can work on deepening their understanding of the research context using appropriate data-collection methods and tools that work best for particular contexts, as well as conduct adequate member checking to ensure the accuracy of findings.

### 7.2 Practicing increased researcher reflexivity.

Qualitative researchers are recommended that they acknowledge and address their own biases, thus recognizing the influence of their own experiences and opinions on the research process [15]. Similar exercises on reflectivity can also be helpful for researchers augmenting qualitative data analysis through employing LLMs. Researcher reflexivity in such instances can extend to querying the LLM to ask for rationale on why certain topics were extracted, grounding topics in anecdotes from the transcripts, and recognizing the influence the human researcher's prior knowledge and biases will have on the prompts used. Future work in extending LLMs for qualitative research should continue to draw on evaluation criteria grounded in traditional qualitative research paradigm.

### 7.3 Increasing transparency of decisions made throughout the research study.

Qualitative researchers are recommended to thoroughly document all decisions that guide their analysis process by providing thick descriptions, allowing for increased transparency. This practice enhances reliability and reproducibility of the research [24]. Qualitative researchers employing LLMs should also similarly strategize maximizing transparency through mechanisms such as documenting changes in workflow, sharing prompts, and detailing model preferences.

## 8 Closing Thoughts

The approach outlined in this paper offers a promising avenue for industry-based talent management practitioners seeking to increase the time and resource efficiency of qualitative interview data analysis. By leveraging large language models (LLMs) as novice

qualitative research assistants, organizations can potentially accelerate the coding, categorization, and thematic synthesis of rich interview data - a critical bottleneck in many talent management research initiatives.

However, as the field of LLM-assisted qualitative research matures, it will be essential to not only benchmark model performance against traditional quantitative evaluation metrics, but also consider quality criteria more prominent within the qualitative research paradigm. Factors such as credibility, transferability, dependability, and confirmability will need to be carefully evaluated as LLMs are integrated into qualitative workflows. Furthermore, the ethical use of AI assistants in sensitive domains like talent management will require close, multi-disciplinary attention to issues at the intersection of data privacy, algorithmic bias, and model transparency, for which researchers will have to be trained [25].

Future research should seek to establish guidelines and best practices for LLM-augmented qualitative analysis that uphold the rigor and trustworthiness expected within the qualitative research community. Only by doing so can talent management scholars and practitioners unlock the full potential of these powerful language models, while respecting the epistemological foundations of qualitative inquiry. As the field evolves, we believe that a judicious, ethically-grounded approach to LLM integration can yield substantial gains in research efficiency and organizational impact.

## References

- [1] Emily M Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the dangers of stochastic parrots: Can language models be too big?. In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*. 610–623.
- [2] Patricia Benner. 1982. From novice to expert. *AJN The American Journal of Nursing* 82, 3 (1982), 402–407.
- [3] Sreyoshi Bhaduri. 2018. NLP in Engineering Education-Demonstrating the use of Natural Language Processing Techniques for Use in Engineering Education Classrooms and Research. (2018).
- [4] Sreyoshi Bhaduri, Michelle Soledad, Tamoghna Roy, Homero Murzi, and Tamara Knott. 2021. A Semester Like No Other: Use of Natural Language Processing for Novice-Led Analysis on End-of-Semester Responses on Students' Experience of Changing Learning Environments Due to COVID-19. In *2021 ASEE Virtual Annual Conference Content Access*.
- [5] Tom B Brown. 2020. Language models are few-shot learners. *arXiv preprint ArXiv:2005.14165* (2020).
- [6] Alan Bryman. 2016. *Social research methods*. Oxford university press.
- [7] Victoria Clarke and Virginia Braun. 2017. Thematic analysis. *The journal of positive psychology* 12, 3 (2017), 297–298.
- [8] Elizabeth G Creamer. 2017. *An introduction to fully integrated mixed methods research*. sage publications.
- [9] Elizabeth G Creamer. 2018. Striving for methodological integrity in mixed methods research: The difference between mixed methods and mixed-up methods. , 526–530 pages.
- [10] John W Creswell and J David Creswell. 2017. *Research design: Qualitative, quantitative, and mixed methods approaches*. Sage publications.
- [11] Norman K Denzin, Yvonna S Lincoln, Michael D Giardina, and Gail S Cannella. 2023. *The Sage handbook of qualitative research*. Sage publications.
- [12] Jacob Devlin. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).
- [13] Alexander Dunn, John Dagdelen, Nicholas Walker, Sanghoon Lee, Andrew S Rosen, Gerbrand Ceder, Kristin Persson, and Anubhav Jain. 2022. Structured information extraction from complex scientific text with fine-tuned large language models. *arXiv preprint arXiv:2212.05238* (2022).
- [14] Yogesh K Dwivedi, Nir Kshetri, Laurie Hughes, Emma Louise Slade, Anand Jeyaraj, Arpan Kumar Kar, Abdullah M Baabdullah, Alex Koohang, Vishnupriya Raghavan, Manju Ahuja, et al. 2023. Opinion Paper: "So what if ChatGPT wrote it?" Multidisciplinary perspectives on opportunities, challenges and implications of generative conversational AI for research, practice and policy. *International Journal of Information Management* 71 (2023), 102642.
- [15] Linda Finlay. 2002. Negotiating the swamp: the opportunity and challenge of reflexivity in research practice. *Qualitative research* 2, 2 (2002), 209–230.

- [16] Barney G Glaser. 1965. The constant comparative method of qualitative analysis. *Social problems* 12, 4 (1965), 436–445.
- [17] Jennifer C Greene. 2008. Is mixed methods social inquiry a distinctive methodology? *Journal of mixed methods research* 2, 1 (2008), 7–22.
- [18] Elizabeth J Halcomb and Louise Hickman. 2015. Mixed methods research. (2015).
- [19] Aobo Kong, Shiwang Zhao, Hao Chen, Qicheng Li, Yong Qin, Ruiqi Sun, and Xin Zhou. 2023. Better zero-shot reasoning with role-play prompting. *arXiv preprint arXiv:2308.07702* (2023).
- [20] Mariah Kornbluh. 2015. Combatting challenges to establishing trustworthiness in qualitative research. *Qualitative research in psychology* 12, 4 (2015), 397–414.
- [21] Jieh-Sheng Lee and Jieh Hsiang. 2019. Patentbert: Patent classification with fine-tuning a pre-trained bert model. *arXiv preprint arXiv:1906.02124* (2019).
- [22] Juha Leino and Kari-Jouko Räihä. 2007. Case amazon: ratings and reviews as part of recommendations. In *Proceedings of the 2007 ACM conference on Recommender systems*. 137–140.
- [23] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems* 33 (2020), 9459–9474.
- [24] Yvonna S Lincoln and Egon G Guba. 1988. Criteria for Assessing Naturalistic Inquiries as Reports. (1988).
- [25] Tammy Mackenzie, Leslie Salgado, Sreyoshi Bhaduri, Victoria Kuketz, Solenne Savoia, and Lilianny Virguez. 2024. Beyond the Algorithm: Empowering AI Practitioners through Liberal Education. In *2024 ASEE Annual Conference & Exposition*.
- [26] Joseph A Maxwell. 2018. Collecting qualitative data: A realist approach. *The SAGE handbook of qualitative data collection* (2018), 19–32.
- [27] Devlin B Montfort, Geoffrey L Herman, Shane A Brown, Holly M Matusovich, and Ruth A Streveler. 2013. Novice-led paired thematic analysis: A method for conceptual change in engineering. In *2013 ASEE Annual Conference & Exposition*. 23–933.
- [28] Michael Paskevicius. 2018. *Exploring educators experiences implementing open educational practices*. Ph. D. Dissertation.
- [29] Michael Quinn Patton. 2014. *Qualitative research & evaluation methods: Integrating theory and practice*. Sage publications.
- [30] Noel Pearse. 2019. An illustration of deductive analysis in qualitative research. In *18th European conference on research methodology for business and management studies*. 264.
- [31] Sergio Pelaez, Gaurav Verma, Barbara Ribeiro, and Philip Shapira. 2024. Large-scale text analysis using generative language models: A case study in discovering public value expressions in AI patents. *Quantitative Science Studies* 5, 1 (2024), 153–169.
- [32] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog* 1, 8 (2019), 9.
- [33] Murray Shanahan. 2024. Talking about large language models. *Commun. ACM* 67, 2 (2024), 68–79.
- [34] Anselm Strauss and Juliet Corbin. 1998. Basics of qualitative research techniques. (1998).
- [35] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is All you Need. In *Advances in Neural Information Processing Systems*, I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Eds.), Vol. 30. Curran Associates, Inc. [https://proceedings.neurips.cc/paper\\_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf)
- [36] J Wei, X Wang, D Schuurmans, M Bosma, F Xia, and E Chi. [n. d.]. & Zhou, D.(2022). *Chain-of-thought prompting elicits reasoning in large language models* ([n. d.]), 24824–24837.
- [37] Binxia Yang, Xudong Luo, Kaili Sun, and Michael Y Luo. 2023. Recent progress on text summarisation based on bert and gpt. In *International Conference on Knowledge Science, Engineering and Management*. Springer, 225–241.
- [38] Kaya Yilmaz. 2013. Comparison of quantitative and qualitative research traditions: Epistemological, theoretical, and methodological differences. *European journal of education* 48, 2 (2013), 311–325.
- [39] Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2019. Bertscore: Evaluating text generation with bert. *arXiv preprint arXiv:1904.09675* (2019).
- [40] Wanqun Zhao. 2023. Using Science to Support and Develop Employees in the Tech Workforce—An Opportunity for Multidisciplinary Pursuits in Engineering Education. In *2023 ASEE Annual Conference & Exposition*.

for each generated topic. For instance, for Topic 1 words such as "students", "develop", "institution", "science", etc. were found important. Attempting to interpret the underlying thematic meaning of these word lists can be challenging without additional contextual information about how those words were used within the original corpus. In contrast, large language models (LLMs) have demonstrated the capability to synthesize the semantically related words and phrases into more coherent topical representations. This ability of LLMs to generate primitive yet formative contextual information threading together words and phrases of interest and thereby provide researchers with a more insightful starting point for further analysis and interpretation of the latent topics uncovered through the LDA process.

## A Results from analyzing the same dataset using an LDA Approach.

Traditional topic modeling using approaches such as Latent Dirichlet Allocation (LDA) often present the most representative words

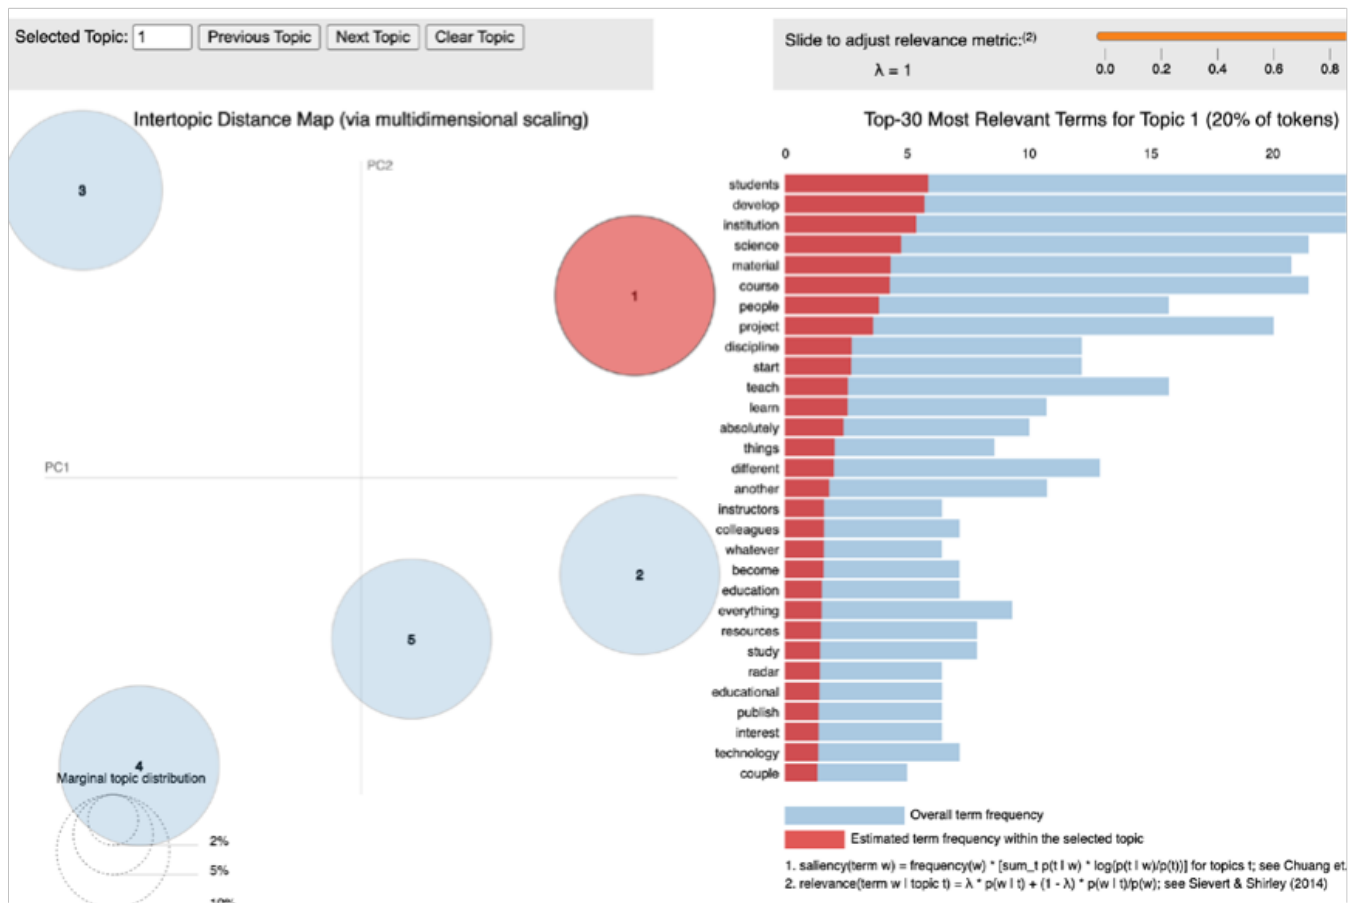


Figure 3: Topic Modeling using LDA