

QuaLLM: An LLM-based Framework to Extract Quantitative Insights from Online Forums

Varun Nagaraj Rao^{1,2}, Eesha Agarwal¹, Samantha Dalal³,
Dan Calacci^{1,2}, Andrés Monroy-Hernández^{1,2}

¹Princeton University, ²Center for Information Technology Policy, ³University of Colorado Boulder
USA

{varunrao, eagarwal}@princeton.edu, samantha.dalal@colorado.edu, {dcalacci, andresmh}@princeton.edu

ABSTRACT

Online discussion forums provide crucial data to understand the concerns of a wide range of real-world communities. However, the typical qualitative and quantitative methods used to analyze those data, such as thematic analysis and topic modeling, are infeasible to scale or require significant human effort to translate outputs to human readable forms. This study introduces QuaLLM, a novel LLM-based framework to analyze and extract quantitative insights from text data on online forums. The framework consists of a novel prompting methodology and evaluation strategy. We applied this framework to analyze over one million comments from two Reddit’s rideshare worker communities, marking the largest study of its type. We uncover significant worker concerns regarding AI and algorithmic platform decisions, responding to regulatory calls about worker insights. In short, our work sets a new precedent for AI-assisted quantitative data analysis to surface concerns from online forums.

1 INTRODUCTION

Text data from online forums, such as Reddit, has been used to study issues affecting diverse communities [37]. Researchers have leveraged this data across various domains, including healthcare [23, 35, 51, 53], political discourse [14, 29, 32], education [27, 33], and gig work¹ [26, 39, 42, 46, 52], to highlight critical concerns and trends.

Studies analyzing this text data for social phenomena often use interpretive qualitative methods [21] like coding [41] and thematic analysis [9], to extract themes. Although detailed, these methods demand considerable human effort and interpretation. Such detailed interpretations may not always be necessary for identifying broad trends across large online communities. Alternatively, positivist quantitative [34] techniques like LDA [1] and BERT derivatives for topic modeling also distill themes from text, but offer greater scalability. However, they require careful hyperparameter selection, operate on limited context length, and require significant human effort to translate outputs to human readable forms.

Recent advancements in LLMs and prompt engineering [15, 30, 31, 36] offer an alternative for analyzing extensive text data from online forums, addressing limitations of scale and context. Specifically, LLMs have been shown to perform well on generation, classification, and ranking tasks [55], and thus have the potential to assist researchers in text data analysis. However, the hitherto absence of a robust methodology and evaluation strategy for effectively using

LLMs for online text data analysis limits its use and underscores a significant gap in the research landscape.

Our work addresses this gap through the following contributions:

- (i) We introduce QuaLLM, an LLM-based framework consisting of a novel prompting methodology and evaluation strategy for the analysis and extraction of quantitative insights from text data on online forums.
- (ii) We apply our framework to a case study on Reddit’s rideshare communities, analyzing over one million comments—the largest study of its kind—to identify worker concerns regarding AI and algorithmic platform decisions, responding to regulatory calls [49].
- (iii) We discuss the broader implications of using LLMs for quantitative text data analysis.

Taken together, by establishing a new precedent for AI-assisted quantitative data analysis to surface concerns from online forums, we hope our work can significantly assist quantitative researchers across various fields.

2 RELATED WORK

LLM Prompting and Evaluation. LLMs have been used for generation [36, 55], classification [15], and ranking [38] tasks via advances in prompt engineering [30, 31]. Popular prompting methodologies include: *In-context learning* [10], which teaches LLMs new tasks with few labeled examples, *Chain of Thought* [48], which breaks down complex problems through stepwise reasoning, and *Ensembling*, [45] which aggregates outputs of multiple model runs to obtain a consensus output. Despite these prompting advances, accurately evaluating LLM outputs is challenging, with human evaluation generally preferred [5]. Building on prior work, we introduce a framework consisting of a novel prompting methodology and evaluation strategy for quantitative analysis of text data from online forums.

Gig Work and Reddit Analysis. Regulatory bodies [49] have sought public insights on the impact of AI and algorithmic management by employers [8, 24], particularly in the gig economy. Gig workers actively discuss their experiences on anonymous forums, e.g., Reddit, which have become a rich data source for researchers. Previous studies on rideshare worker concerns using Reddit, highlighting issues of privacy, scams, and support systems [26, 39, 42, 46, 52], have been small-scale, limited by manual coding². We use our LLM-based framework to examine over 1 million

¹a market characterized by short-term and freelance work like food delivery and ridesharing

²the largest dataset analyzed about 2.6K posts [46] using qualitative data analysis methods

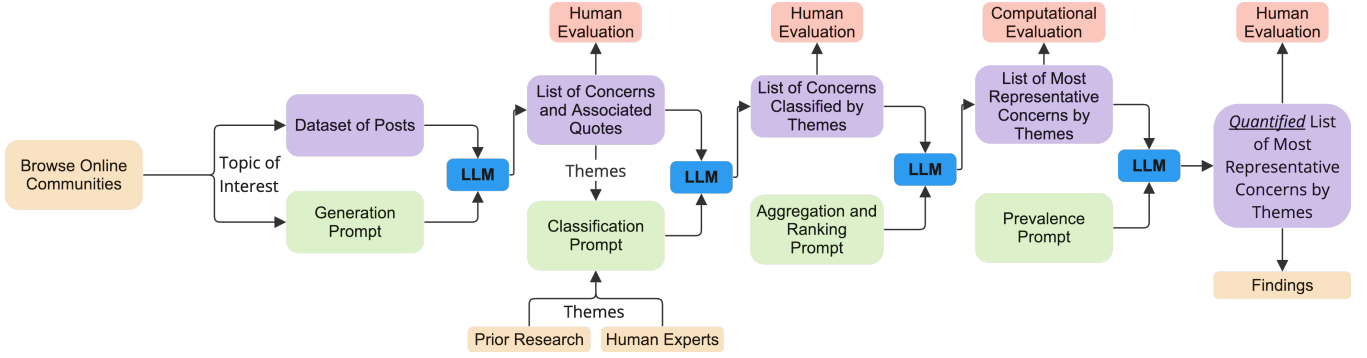


Figure 1: Overview of QuaLLM: an LLM-Based Framework to Extract Quantitative Insights from Online Forums

comments from rideshare workers, significantly broadening the scope to amplify all voices.

3 QUALLM FRAMEWORK

We introduce QuaLLM (Figure 1), a framework consisting of a prompting methodology and evaluation strategy for the analysis of text data from online forums.

3.1 LLM Prompting Methodology for Analysis of Online Forum Data

Our approach begins with the collection of datasets from specified periods across any public threaded discussion forum, such as Reddit, Mastodon, and Facebook. We aim to extract and summarize the predominant concerns of workers within these communities through a four-step prompting process: generation, classification, aggregation, and prevalence.

3.1.1 Generation Prompt. The generation step involves the generation of *concern summaries* from online discussions *relevant to a specific topic* and linking each concern with a *representative quote*. The prompt is broken down into seven steps, following best prompt engineering practices from prior research and industry [2, 30, 31, 44]. The steps include:

- (1) identifying concerns
- (2) aggregating to avoid redundancy
- (3) selecting representative quotes³
- (4) assessing concern frequency and impact
- (5) formatting outputs in JSON with concern titles, descriptions, and quotes
- (6) ensuring concerns are topic-specific and derived from input data without external knowledge
- (7) reiterating to prevent redundancy and indicating when no concerns are present

To minimize API calls, we aggregate posts and comments into groups. For each group, we programmatically enhance the JSON output with the timestamp of the earliest post and assign a unique identifier, ensuring precise identification.

All these steps are necessary since direct prompting may cause hallucinations, redundancy, and inconsistent output. This prompting strategy is similar to “open coding,” where one derives new theories and concepts informed from the underlying data [41].

3.1.2 Classification Prompt. We define a set of primary themes, typically around 4-5, based on prior research, parallel participant interviews, or other expert input, and the outputs of the generation prompt. We then prompt the LLM to classify each identified concern accordingly, with an additional “other” category for irrelevant concerns. This step mirrors “thematic analysis” [9], which groups together a number of related codes, serving as a filtration mechanism.

3.1.3 Aggregation Prompt. Aggregates concerns into secondary themes based on classifications, and identifies the most representative concerns (e.g., top 5) within each primary theme. We can manually select representative quotes for each secondary theme.

3.1.4 Prevalence Prompt. Classifies concerns within primary themes as belonging to one of the secondary themes or an “other” category, facilitating quantitative analysis of concern prevalence by theme percentage.

3.2 LLM Outputs Evaluation Strategy

To validate the LLM outputs we propose a mix of human and computational evaluation methods based on the following metrics: *factual-ity* and *completeness* for generation, *accuracy* for classification and prevalence, and *distinctness* and *coverage* for aggregation prompts, respectively.

3.2.1 Generation Prompt: This prompt generates new text given unstructured source text as input.

Factuality: Determines if the LLM’s output concerns reflect the source data. Evaluators respond to: “*Is this candidate concern (and associated quote) output by the LLM present in the reference human generated list of transparency concerns? Answer Yes or No*”. The factuality score is the proportion of accurate LLM-generated concerns, akin to precision.

Completeness: Evaluates the LLM’s ability to capture all relevant concerns, asking: “*Is this reference human concern present in the candidate concern (and associated quote) output by the LLM? Answer*

³we argue that forcing the model to choose and output quotes is a form of Chain-of-Thought prompting [48] strategy and find that it helps prevent hallucination

Yes or No." The completeness score denotes the proportion of human-identified concerns also recognized by the LLM, akin to recall.

3.2.2 Classification and Prevalence Prompts: Both these prompts classify source text into predefined categories.

Accuracy: Assesses the correctness of LLM classifications, measuring the proportion of accurately classified samples.

3.2.3 Aggregation Prompt: This prompt generates aggregated text given source text as input. We leverage a topic model to generate topics for the entire text and find those topics most similar to the subthemes. Human evaluation on a subset does not accurately capture the task, and instead we pivoted to a computational evaluation on the entire text using well established topic modeling techniques. **Distinctness:** We measure the *proportion of unique most-similar topics associated with each sub-theme across all the sub-themes*. A higher value indicates a greater diversity in the subthemes.

Coverage (k): We measure the proportion of unique most-similar topics associated with each of the 'n' sub-themes that are also found among the most frequent 'nk' topics of the overall text. For example, if $n=5$ and $k=2$, we would look at the 5 sub-themes amongst the top 10 most frequent overall topics. A higher CR-k indicates that the sub-themes represent prevalent or significant themes in the text.

4 CASE STUDY OF WORKER CONCERNS FROM RIDESHARE SUBREDDITS

We apply our framework to rideshare subreddit data (r/uberdrivers, r/lyftdrivers) using the GPT4-Turbo model on Microsoft Azure to identify worker concerns regarding AI and platform algorithms. See Appendices A, B, and C for our prompts, findings and costs incurred, respectively.

Rideshare subreddits were chosen for their active 400K+ member community, presence of active discussions of worker concerns; which is of interest to regulators, accessible and comprehensive data, and relevance to HCI, Economics, and Communications researchers studying the harms caused to workers as a result of rideshare platforms' algorithmic management [17, 39, 54].

We now describe our dataset and summarize our findings.

4.1 Dataset and Experiments

We obtained comprehensive 2019-2022 data from r/uberdrivers and r/lyftdrivers via <https://the-eye.eu/redarcs/>, focusing on this period to capture recent concerns, particularly relevant due to app changes and the unavailability of post-2022 data due to Reddit API changes⁴. After post pre-processing to remove short samples, our dataset comprised 65,377 submissions (47, 106 from Uber, 18, 271 from Lyft) and 1,392,776 associated comments (1,054,030 from Uber, 338,746 from Lyft)⁵. We then grouped submissions and comments into sets of five for each LLM API call (See Appendix A for the prompts). We encountered an 11% error rate due to throttling and content issues. Finally we obtained 58,728 concerns which we categorized into four themes: Enhancing Transparency and Explainability (24,721 concerns, 42%), Predictability and Worker Agency

(12,728 concerns, 22%), Better Safety and More Time (6,144 concerns, 10.5%), and Ensuring Fairness and Non-Discrimination (4,280 concerns, 7%); the remaining 18.5% were categorized as "Other". We then aggregated these findings to highlight the top five concerns for each theme and calculated their prevalence across the dataset.

4.2 Findings

We briefly summarize our findings below. We also provide tables of the aggregated, ranked and prevalence quantified themes in Appendix B.

(i) *Enhancing Transparency and Explainability:* Driver concerns about opaque fare calculations, unclear incentives, and uncertain criteria for earning, surge pricing, and cancellation (Table 1)

(ii) *Predictability and Worker Agency:* Drivers face unpredictable earnings from fluctuating surge pricing, algorithm changes, increased competition, complex incentive qualifications, and low compensation for long pickups and waits. (Table 2)

(iii) *Better Safety and More Time:* Drivers face navigation issues, support access difficulties, challenges with false complaints, and low compensation for additional tasks and wait times. (Table 3)

(iv) *Ensuring Fairness and Non-Discrimination:* Drivers dispute their fare share and platform take rates, fairness in ride assignments, the influence of demographics on ratings, and the criteria for deactivation. (Table 4)

4.3 Evaluation

4.3.1 Generation Prompt. We randomly sampled 125 submissions and all the 2,511 associated comments, jointly annotated by two researchers, to create a reference list of concerns. The two researchers then jointly evaluated the LLM outputs against the reference list of concerns, and resolved any differences through discussions. We obtained a factuality score of 0.55 and a completeness score of 0.78. Lower factuality was primarily due to the LLM identifying concerns part of the input context, but not directly related to AI and algorithmic decisions. Moreover, the lower factuality score was less concerning as outputs not factual (i.e., identified by the LLM but not by humans) were mostly filtered out into the "Other" category during the classification step and excluded from any further analysis.

4.3.2 Classification and Prevalence Prompts. Two researchers jointly annotated 100 LLM-classified concerns to establish ground truth labels, assigning one label per concern. They compared LLM outputs with these labels, achieving a consensus accuracy of 0.74. The achieved accuracy is significant because misclassified samples, excluding those labeled as "Other," can still contribute to the analysis. This is because concerns might fit into multiple categories depending on the analytical perspective. Even if placed in a different category, these concerns remain valuable for later aggregation and prevalence prompts. Using a similar methodology, researchers obtained an accuracy of 0.82 for the prevalence prompts' output.

4.3.3 Aggregation Prompt. We employed BERTopic for topic modeling (see Appendix D for hyperparameter details). Given the large volume of concerns (4K-25K) per primary theme, human evaluation on a subset would not accurately capture the task's nuances. Therefore, we adopted this well-established computational method for

⁴Reddit API restrictions enforced in Feb'23; see <https://www.redditinc.com/policies/data-api-terms>

⁵On Reddit, a submission is a post, and a comment is a response to a post or another comment; we concatenate and use the text from posts and all associated comments

evaluation. We fit a topic model on the entire text (title+description) associated with each primary theme to obtain the top 5 and 10 topics. Subsequently, we identified the most similar topic for each of the 5 sub-themes. The results demonstrated a distinctness of 0.95, coverage(1) of 0.80, and coverage(2) of 1.00, indicating that the LLMs' outputs were both distinct and well-represented among the most frequent topics.

5 DISCUSSION OF BROADER IMPACTS

5.1 Scaling online forum analysis with LLMs

QuaLLM presents a novel framework that leverages the power of LLMs to analyze vast amounts of unstructured text data efficiently and translate to a form which resembles community survey responses. Beyond analyzing online discussions in rideshare communities on Reddit, this methodology has broad applicability across numerous domains, across any platform, where large amounts of unstructured text is prevalent and traditional positivist quantitative methods [34] such as topic models through LDA, may fall short. We now discuss a few domains where our approach can aid data analysis:

- (1) **Healthcare Forums:** Identify discussions of medications from patient forums comparing anecdotal experiences and widespread concerns.
- (2) **Customer Feedback:** Filter through customer service chats and emails to identify recurring product or service issues.
- (3) **Policy and Regulation:** Review public responses to policy proposals, identifying key points of support and contention.
- (4) **Educational Feedback:** Review online forums and course feedback to extract themes related to assessment and course content.
- (5) **Financial Market Sentiments:** Analyze investor forums and financial news comments to gauge market sentiment towards certain stocks or economic policies.

5.2 QuaLLM's evaluation reveals strengths and necessitates further refinement

Due to the unique nature of our pipeline approach, which combines several tasks into a cohesive workflow, unlike the most closely related work by Pham et al. [36] where LLMs are used for the single task of topic generation, there are no established baselines or benchmarks available for direct end-to-end comparison and evaluation. Topic models might offer a comparison for the generation prompt outputs, yet they require significant human effort to translate outputs to human readable forms, unlike our method where the LLM autonomously generates topics, making baselining challenging.

This absence of precedent opens questions regarding effective baselining and benchmarking for performance evaluation. As a result, human evaluation serves as the gold standard for evaluation, ensuring the validity of our findings.

5.3 LLMs can be used to assist researchers and complement existing methods

Some researchers highlight limitations in LLMs' understanding of complex human emotions and experiences, pointing to a lack of

context-sensitivity [4, 7, 40]. Others [6] argue LLMs cannot match the innate human ability to perform certain tasks in inexplicable ways, described in conceptual frameworks for knowledge creation [12, 25]. This gap raises doubts about LLMs' capability for interpretation crucial to fields such as anthropology [28] and sociology[47], challenging their application within the qualitative interpretivist research paradigm [16, 21, 22].

On the other hand, advocates for LLM-enhanced research suggest LLMs align with human outputs and in some cases can replace human judgement [6, 11, 13, 18, 19, 36, 43, 50]. Further, they find integrating AI with human expertise can improve the scope and scale of analyses, facilitating efficient processing of vast, unstructured datasets.

Our research stands in the middle, underscores the potential of LLMs to help researchers engage with data in a way that can supply intrinsic meaning at scale, while stressing the need to recognize their limitations and the continued importance of human oversight and interpretation.

Broadly, QuaLLM can also be used in mixed methods study designs to complement data from interviews and focus groups. However, safeguards, including successive filtering and human evaluation, are necessary to ensure validity of LLM outputs, indicating a complementary shift rather than a complete replacement of human interpretation.

6 LIMITATIONS

Interpretivist vs Positivist trade-offs: While LLMs efficiently aggregate broad trends, reflecting a positivist quantitative approach that assumes text has intrinsic meaning, individual experiences may be neglected. This contrasts with the interpretivist qualitative paradigm, which values subjective meaning derived from the interaction between the researcher and the data.

Small scale human evaluation: We acknowledge the limited generalizability of our human evaluation due to the relatively small sample size and its restriction to specific pipeline stages.

Hyperparameter-sensitive topic modeling: BERTopic performance is very sensitive to choice of hyperparameters and this can result in significant variation in results based on the chosen hyperparameter values.

Reproducibility concerns: The stochasticity of LLM outputs raises concerns about reproducibility. However, this variability mirrors the inherent subjectivity of human perspectives, which can differ across time and contexts.

7 CONCLUSION

Our study presents a new framework using LLMs consisting of a prompting methodology and evaluation strategy for analysis of online forum data, showcased by analyzing over one million rideshare subreddit comments—the largest analysis to date. This approach highlights LLMs' capability to efficiently process large unstructured datasets establishing a precedent for future AI-assisted positivist quantitative research. But does an increase in data equate to enhanced understanding? Echoing Agnew et al. [3]'s sentiments,

we believe further research is essential to determine the applicability and utility of AI-assisted research methods.

ACKNOWLEDGEMENTS

We acknowledge the Azure Cloud Computing Grant from CSML at Princeton which enabled access to Microsoft Azure compute resources. We thank Sayash Kapoor for insightful discussions on the conceptual and normative implications of LLMs and prompting strategies. We are grateful to Vinayshekhar Bannihatti Kumar for his valuable inputs on defining evaluation metrics. We thank Sunnie Kim for her contributions on connections to Explainable AI. We appreciate Elizabeth Anne Watkins for providing valuable critiques from an interpretivist perspective, which helped reframe our work. We thank Yuhuan Liu, Beza Desta, and Kiyosu Maeda from the Princeton HCI group for their annotations. We are grateful to Hope Schroeder for their feedback on evaluations and broad discussions on applications. We thank Prof. Arvind Narayanan and Mihir Kshirsagar from Princeton CITP for their insights on the generalization of this work and its policy implications. We thank Prof. Brian Keegan for pointing us to the Reddit data source. Finally, we thank the anonymous CSCW reviewers for their valuable feedback, which helped frame the larger paper of which this is a part.

REFERENCES

- [1] Ashraf Abdul, Jo Vermeulen, Danding Wang, Brian Y Lim, and Mohan Kankanhalli. 2018. Trends and trajectories for explainable, accountable and intelligible systems: An HCI research agenda. In *Proceedings of the 2018 CHI conference on human factors in computing systems*. 1–18.
- [2] Griffin Adams, Alexander Fabbri, Faisal Ladhak, Eric Lehman, and Noémie Elhadad. 2023. From sparse to dense: GPT-4 summarization with chain of density prompting. *arXiv preprint arXiv:2309.04269* (2023).
- [3] William Agnew, A Stevie Bergman, Jennifer Chien, Mark Diaz, Selim El-Sayed, Jaylen Pittman, Shakir Mohamed, and Kevin R McKee. 2024. The illusion of artificial inclusion. *arXiv preprint arXiv:2401.08572* (2024).
- [4] Hussam Alkaissi and Samy I McFarlane. 2023. Artificial hallucinations in ChatGPT: implications in scientific writing. *Cureus* 15, 2 (2023).
- [5] Meor Amer. 2023. Evaluating LLM Outputs. <https://txt.cohere.com/evaluating-llm-outputs/>
- [6] Muneera Bano, Didar Zowghi, and Jon Whittle. 2023. Exploring Qualitative Research Using LLMs. *arXiv preprint arXiv:2306.13298* (2023).
- [7] 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.
- [8] Annette Bernhardt, Lisa Kresge, and Reem Suleiman. 2021. Data and Algorithms at Work. (2021). <https://laborcenter.berkeley.edu/wp-content/uploads/2021/11/Data-and-Algorithms-at-Work.pdf>
- [9] Virginia Braun and Victoria Clarke. 2012. *Thematic analysis*. American Psychological Association.
- [10] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems* 33 (2020), 1877–1901.
- [11] Robert Chew, John Bollenbacher, Michael Wenger, Jessica Speer, and Annice Kim. 2023. LLM-assisted content analysis: Using large language models to support deductive coding. *arXiv preprint arXiv:2306.14924* (2023).
- [12] Harry M Collins. 2005. What is tacit knowledge? In *The practice turn in contemporary theory*. Routledge, 115–128.
- [13] Shih-Chieh Dai, Aiping Xiong, and Lun-Wei Ku. 2023. LLM-in-the-loop: Leveraging large language model for thematic analysis. *arXiv preprint arXiv:2310.15100* (2023).
- [14] Gianmarco De Francisci Morales, Corrado Monti, and Michele Starnini. 2021. No echo in the chambers of political interactions on Reddit. *Scientific reports* 11, 1 (2021), 2818.
- [15] Xiang Deng, Vasilisa Bashlovkina, Feng Han, Simon Baumgartner, and Michael Bendersky. 2023. LLMs to the Moon? Reddit Market Sentiment Analysis with Large Language Models. In *Companion Proceedings of the ACM Web Conference 2023*. 1014–1019.
- [16] Norman K Denzin. 2001. *Interpretive interactionism*. Vol. 16. Sage.
- [17] Veena Dubal. 2023. On algorithmic wage discrimination. *Available at SSRN 4331080* (2023).
- [18] Zackary Okun Dunivin. 2024. Scalable Qualitative Coding with LLMs: Chain-of-Thought Reasoning Matches Human Performance in Some Hermeneutic Tasks. *arXiv preprint arXiv:2401.15170* (2024).
- [19] 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. “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.
- [20] Joseph L Fleiss, Bruce Levin, and Myunghee Cho Paik. 2013. *Statistical methods for rates and proportions*. John Wiley & sons.
- [21] Mwenda J Gichuru. 2017. The interpretive research paradigm: A critical review of is research methodologies. *International Journal of Innovative Research and Advanced Studies (IJIRAS)* 4, 2 (2017), 1–5.
- [22] Barney G Glaser, Anselm L Strauss, and Elizabeth Strutzel. 1968. The discovery of grounded theory: strategies for qualitative research. *Nursing research* 17, 4 (1968), 364.
- [23] Chia-Ling Lynn Ho, Wenjing Pan, and Laramie D Taylor. 2017. Stigma of HIV testing on online HIV forums: self-stigma and the unspoken. *Journal of Psychosocial Nursing and Mental Health Services* 55, 12 (2017), 34–43.
- [24] Jodi Kantor and Arya Sundaram. 2022. The Rise of the Worker Productivity Score (New York Times). <https://www.nytimes.com/interactive/2022/08/14/business/worker-productivity-tracking.html>
- [25] Meng Li and Fei Gao. 2003. Why Nonaka highlights tacit knowledge: a critical review. *Journal of knowledge management* 7, 4 (2003), 6–14.
- [26] Ning F Ma, Chien Wen Yuan, Moojan Ghafurian, and Benjamin V Hanrahan. 2018. Using stakeholder theory to examine drivers’ stake in Uber. In *Proceedings of the 2018 CHI conference on human factors in computing systems*. 1–12.
- [27] Marisa Alise Madsen and Dag Øivind Madsen. 2022. Communication between Parents and Teachers of Special Education Students: A Small Exploratory Study of Reddit Posts. *Social Sciences* 11, 11 (2022), 518.
- [28] Bronislaw Malinowski. 1929. Practical anthropology. *Africa* 2, 1 (1929), 22–38.
- [29] Nahema Marchal. 2020. The polarizing potential of intergroup affect in online political discussions: Evidence from reddit r/politics. *Politics (July 10, 2020)* (2020).
- [30] Harsha Nori, Yin Tat Lee, Sheng Zhang, Dean Carignan, Richard Edgar, Nicolo Fusi, Nicholas King, Jonathan Larson, Yuanzhi Li, Weishung Liu, et al. 2023. Can generalist foundation models outcompete special-purpose tuning? case study in medicine. *arXiv preprint arXiv:2311.16452* (2023).
- [31] OpenAI. 2023. Prompt Engineering. <https://platform.openai.com/docs/guides/prompt-engineering>
- [32] Orestis Papakyriakopoulos, Severin Engelmann, and Amy Winecoff. 2023. Upvotes? Downvotes? No Votes? Understanding the relationship between reaction mechanisms and political discourse on Reddit. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–28.
- [33] Albert Park, Mike Conway, et al. 2018. Harnessing reddit to understand the written-communication challenges experienced by individuals with mental health disorders: analysis of texts from mental health communities. *Journal of medical Internet research* 20, 4 (2018), e8219.
- [34] Yoon Soo Park, Lars Konge, and Anthony R Artino Jr. 2020. The positivism paradigm of research. *Academic medicine* 95, 5 (2020), 690–694.
- [35] Christina G Paxman. 2021. “Everyone thinks I am just lazy”: Legitimacy narratives of Americans suffering from fibromyalgia. *Health* 25, 1 (2021), 121–137.
- [36] Chau Minh Pham, Alexander Hoyle, Simeng Sun, and Mohit Iyyer. 2023. TopicGPT: A Prompt-based Topic Modeling Framework. *arXiv preprint arXiv:2311.01449* (2023).
- [37] Nicholas Proferes, Naiyan Jones, Sarah Gilbert, Casey Fiesler, and Michael Zimmer. 2021. Studying reddit: A systematic overview of disciplines, approaches, methods, and ethics. *Social Media+ Society* 7, 2 (2021), 20563051211019004.
- [38] Zhen Qin, Rolf Jagerman, Kai Hui, Honglei Zhuang, Junru Wu, Jiaming Shen, Tianqi Liu, Jialu Liu, Donald Metzler, Xuanhui Wang, et al. 2023. Large language models are effective text rankers with pairwise ranking prompting. *arXiv preprint arXiv:2306.17563* (2023).
- [39] Alex Rosenblat and Luke Stark. 2016. Algorithmic labor and information asymmetries: A case study of Uber’s drivers. *International journal of communication* 10 (2016), 27.
- [40] Jürgen Rudolph, Samson Tan, and Shannon Tan. 2023. ChatGPT: Bullshit spewer or the end of traditional assessments in higher education? *Journal of Applied Learning and Teaching* 6, 1 (2023).
- [41] Johnny Saldaña. 2021. The coding manual for qualitative researchers. *The coding manual for qualitative researchers* (2021), 1–440.
- [42] Shruti Sannon, Billie Sun, and Dan Cosley. 2022. Privacy, surveillance, and power in the gig economy. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. 1–15.

- [43] Robert H Tai, Lillian R Bentley, Xin Xia, Jason M Sitt, Sarah C Fankhauser, Ana M Chicas-Mosier, and Barnas G Monteith. 2023. An Examination of the Use of Large Language Models to Aid Analysis of Textual Data. *bioRxiv* (2023), 2023–07.
- [44] Bailin Wang, Zi Wang, Xuezhi Wang, Yuan Cao, Rif A Saurous, and Yoon Kim. 2023. Grammar Prompting for Domain-Specific Language Generation with Large Language Models. *arXiv preprint arXiv:2305.19234* (2023).
- [45] Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2022. Self-consistency improves chain of thought reasoning in language models. *arXiv preprint arXiv:2203.11171* (2022).
- [46] Elizabeth Anne Watkins. 2022. “Have you learned your lesson?” Communities of practice under algorithmic competition. *New Media & Society* 24, 7 (2022), 1567–1590.
- [47] Max Weber. 1949. “Objectivity” in social science and social policy. *The methodology of the social sciences* (1949), 49–112.
- [48] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems* 35 (2022), 24824–24837.
- [49] White House (OSTP). 2023. Request for Information; Automated Worker Surveillance and Management. <https://www.federalregister.gov/documents/2023/05/03/2023-09353/request-for-information-automated-worker-surveillance-and-management>
- [50] Ziang Xiao, Xingdi Yuan, Q Vera Liao, Rania Abdelghani, and Pierre-Yves Oudeyer. 2023. Supporting Qualitative Analysis with Large Language Models: Combining Codebook with GPT-3 for Deductive Coding. In *Companion Proceedings of the 28th International Conference on Intelligent User Interfaces*. 75–78.
- [51] Tian Xu, Junnan Yu, Dylan Thomas Doyle, and Stephen Volda. 2023. Technology-Mediated Strategies for Coping with Mental Health Challenges: Insights from People with Bipolar Disorder. *Proceedings of the ACM on Human-Computer Interaction* 7, CSCW2 (2023), 1–31.
- [52] Zheng Yao, Silas Weden, Lea Emerlyn, Haiyi Zhu, and Robert E Kraut. 2021. Together but alone: Atomization and peer support among gig workers. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW2 (2021), 1–29.
- [53] Samaneh Zamanifard and Andrew Robb. 2023. Social Virtual Reality Is My Therapist: Overcoming Social Anxiety Disorder Through Using Social Virtual Reality. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–6.
- [54] Angie Zhang, Alexander Boltz, Chun Wei Wang, and Min Kyung Lee. 2022. Algorithmic management reimagined for workers and by workers: Centering worker well-being in gig work. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. 1–20.
- [55] Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, et al. 2023. A survey of large language models. *arXiv preprint arXiv:2303.18223* (2023).

A PROMPTS

GPT-4 Generation Prompt

```
1 Analyze a set of JSON objects, each representing a submission from the r/UberDrivers and r/LyftDrivers subreddits.
  For each JSON object, you will find the following information:
2
3 Submission Title: The title of the Reddit post.
4 Submission Body: The main content or message of the post.
5 Timestamp: The date and time when the post was submitted.
6 Group Key: A unique identifier that is common across all 5 submissions in the dataset.
7 Comments: A list of all comments made on the submission.
8
9 Generate a list of the most frequently occurring and impactful concerns due to AI and algorithmic platform features
  discussed by drivers within the input context.
10
11 Step 1: Identify mentions in submission bodies, titles, and/or comments about concerns that pertain to a lack of
  knowledge or available information regarding the platforms' algorithms for drivers, including but not limited to
  fares, routes, incentive programs, driver preferences, etc.
12
13 Step 2: Group similar concerns across comments and submissions to ensure a mutually exclusive list of concerns and
  avoid redundancy. For example, multiple mentions of fare calculation issues should be grouped under a single
  concern.
14
15 Step 3: From the grouped concerns, select the most representative quote for each concern. Ensure the quote clearly
  illustrates the specific concern due to AI and algorithmic platform features.
16
17 Step 4: Assess which concerns are mentioned most frequently and have the most significant impact on drivers.
18
19 Step 5: Create a list of these concerns in a JSON format. Each entry should include (with these specific field
  names):
20 "title": The title of the concern
21 "description": A brief description (10-20 words)
22 "quote": The selected representative quote.
23
24 Step 6: Ensure the final list is concise, precise, and specifically addresses drivers' concerns due to AI and
  algorithmic platform features present regarding the platform's algorithms and policies. Include only those concerns
  found in the input context without generalizing based on prior or outside knowledge.
25
26 Step 7: Group any similar concerns to avoid redundancy. If there are no concerns, output ``No concerns''. Do not
  generate any other text.
```

GPT-4 Classification Prompt

```

1 Task: Analyze a list of rideshare drivers' transparency concerns. Each concern should be evaluated and categorized
  into one of the following five categories. For each concern listed, assign exactly one letter that corresponds to
  its most appropriate classification in the same order as the original list (preserving the serial number of the
  concern).
2
3 Categories:
4
5 A. Need for Enhancing Transparency and Explainability: Concerns due to the lack of information and explanations
  about AI and algorithmic features, required for drivers to do their work. These include opaque trip details,
  unclear surge boundaries, bonus and quest clarity issues, lack of wage breakdown, platform take rate and clarity on
  how prices and bonuses are calculated, and the impact of ratings on metrics.
6
7 B. Need for Greater Predictability and Worker Agency: Concerns due to significant variation in AI and algorithmic
  features leading to reduced worker agency and diminished predictability of work conditions. These include
  unpredictable wages, misleading destination filter, immense variation in surge prices and quest matches across
  drivers and location, and deceptive trip offers.
8
9 C. Need for Better Safety and More Time: Concerns focused on safety risks and time pressures due to the AI and
  algorithmic features. These include dangerous multitasking, compromised route safety, unattainable quests, and
  acceptance rate concerns, emphasizing the need for improved safety and time management in app design.
10
11 D. Need for Ensuring Fairness and Non-Discrimination: Concerns related to algorithmic wage discrimination. These
  include unequal pay for similar work, earnings below the prevailing minimum wage, influence of the demographic
  characteristics of drivers and riders on earnings, all caused by opaque AI and algorithmic features.
12
13 E. Other Concerns: Any concerns that do not fit into the above categories.
14
15 Output Format: Present the analysis in a dictionary format with the serial number of the concern as the key and the
  classification as the value, preserving the order of the concerns.
16
17 {1: A, 2: B, 3: C}
18
19 Note: Ensure that each concern is classified under only one of the aforementioned categories and that there is one
  classification corresponding to each concern in the input, e.g., for 400 transparency concerns, there should be 400
  items in the dictionary. Do not generate any other text.

```

GPT-4 Aggregation and Ranking Prompt

```

1 The data contains a list of concerns of rideshare drivers obtained from discussions on the r/uberdrivers and r/
  lyftdrivers subreddits. Each concern title and its description is present on consecutive lines.
2
3 These concerns are relevant to <insert category description from Classification Prompt>
4
5 Identify the 5 most frequently occurring themes of concerns. Ensure each concern is sufficiently different from the
  others on the list. Don't repeat the same concern in the list. If there are similar concerns, group them and find
  another one.
6
7 Provide the output in a rank ordered format:
8 {concern_rank: <a number between 1-5>, concern_title: string, concern_description: string of 10-20 words}

```

GPT-4 Prevalence Prompt

```
1 Classify each line which contains the title and description of a concern from rideshare drivers on Reddit into the
  following 6 categories:
2
3 <list the 5 categories A to E from the aggregation prompt output>
4 F: Other
5
6 Output Format: Present the analysis in a dictionary format with the serial number of the concern as the key and the
  classification as the value, preserving the order of the concerns.
7
8 {1: A, 2: B, 3: C}
9
10 Note: Ensure that each concern is classified under only one of the aforementioned categories and that there is one
    classification corresponding to each concern in the input, e.g., for 400 concerns, there should be 400 items in the
    dictionary. Do not generate any other text.
```

B FINDINGS

Harm	Quote	% (Count)
Drivers are unsure about opaque fare calculation methods.	<i>"\$30 to the middle of nowhere and had to eat dead miles to drive back"</i>	29 (7,202)
Issues with understanding and qualifying for incentive programs, bonuses, and quests.	<i>"Are they really expecting a thumbnail of a map with zero details of where I'd need to be to have a ride qualify for the Quests to be useful?"</i>	20 (4,880)
Drivers express concerns over clarity in ride assignments and earnings determination by algorithms.	<i>"A driver who keeps a high AR, low cancel rate, gets mostly 5 Stars, and gives out mostly 5 Stars to pax will get better (more lucrative) rides than a driver the company's algorithm doesn't like."</i>	11 (2,765)
Confusion persists over surge pricing determination and boundaries.	<i>"I get a ton of riders telling me they are paying 40-60 for the ride (surging rate), but when the ride ends, I get around \$6 (non-surge rate pay)"</i>	8 (1,858)
Concerns over lack of clarity in cancellation criteria and the associated fare impact.	<i>"Then I ask about the cancellation fee and they said they can't provide one as no option came up."</i>	5 (1,197)

Table 1: Most representative concerns on Reddit relevant to the need for enhancing transparency and explainability (Total = 24,721).

Harm	Quote	% (Count)
Drivers face uncertain and fluctuating income due to changes in surge and quest pricing.	<i>"Uber has always done that. When they offer a higher promo, they manipulate and lower surges."</i>	23 (2,953)
Challenges with understanding and adapting to algorithm changes that affect ride assignments and income.	<i>"I multi-app, but the algorithm has changed. If you log off from the app, to ride for another app, the algorithm puts you on the bottom of the queue after you log on again for the next ride."</i>	19 (2,408)
Complexity and unpredictability in qualifying for incentive programs like quests and bonuses.	<i>"Monday thru Thursday is 20 trip min & Friday thru Sunday is 40 trip min to get the bonus. Like why?"</i>	14 (1,799)
Concerns over inadequate payment for time and distance invested in long pickups and waits.	<i>"I don't understand what the hell is Uber thinking when giving us long pickups for short trips. NOBODY sane accepts a \$4 ride for someone half an hour away."</i>	11 (1,429)
A rise in driver numbers contributing to fewer available rides and lower earnings.	<i>"Now that millions of new drivers have joined overnight due to the unemployment benefits ending, rides have dried up."</i>	5 (658)

Table 2: Most representative concerns on Reddit relevant to the need for enhancing predictability and worker agency (Total = 12,728).

Harm	Quote	% (Count)
Drivers experience technical difficulties with app navigation and route mapping, leading to disruptions.	<i>"I started using the Lyft nav, then Google maps, then Waze only to go back to the Lyft nav. I went through that cycle because the Lyft nav sucks but at least you see everything going on with the app. With the others, I've lost out on streaks and pings because of glitches, etc."</i>	20 (1,208)
Challenges in contesting false passenger complaints and inadequate platform responses affecting driver status.	<i>"They claimed the fare was adjusted due to a rider complaint or something. They're lying."</i>	16 (990)
Drivers express risks due to dangerous multitasking and unsafe areas suggested by algorithms	<i>"If a Pax wants to be taken to a part of town I don't feel safe going, I feel like I should know, for my own safety and peace of mind."</i>	15 (909)
Difficulty reaching and obtaining help from support for various driver safety issues	<i>"Support was really quick to respond...only took 13 days."</i>	14 (863)
Drivers are not fairly paid for extensive wait times during stops or added tasks	<i>"The stop system is bullshit. You make less money, 100%. I just checked. To order two rides, one there, one back, \$17. To order the same ride, but using stops it's \$13."</i>	14 (856)

Table 3: Most representative concerns on Reddit relevant to the need for better safety and more time (Total = 6,144).

Harm	Quote	% (Count)
Drivers express discontent with the percentage of fares they receive compared to what the platform charges customers.	<i>"Yeah it's always high ride demand and they charge the passengers high prices and drivers earnings are like 25%"</i>	24 (1,025)
Concerns over unclear or unfair deactivation decisions based on customer complaints or automated systems.	<i>"I canceled two orders for the distance and my account was deactivated for fraud."</i>	20 (850)
Uncertainty about how ratings are determined and their direct impact on drivers' work opportunities.	<i>"I read that Lyft will try to match a pax with a driver that both had 5 starred each other in the past."</i>	9 (384)
Drivers feel platform algorithms unfairly influence ride distribution, surge pricing, and overall income.	<i>"a lady was standing next to my car and ordered a lyft and it gave her a driver that was outside of the Bonus zone and 20 minutes away."</i>	7 (316)
Concern over algorithms distributing rides unfairly or prioritizing certain drivers	<i>"Perceived favoritism by the app's algorithm towards newer drivers can limit earnings for long-term drivers."</i>	4 (151)

Table 4: Most representative concerns on Reddit relevant to the need for ensuring fairness and non-discrimination (Total = 4,280).

C FINANCIAL EXPENDITURE OF USING LLMs

In this section, we delve into the costs incurred from utilizing the GPT4-Turbo LLM on Microsoft Azure to synthesize concerns from the Reddit data.

Cost Category	Token Quantity	Cost (USD)
Input Token Rate	-	\$0.01 per 1K tokens
Output Token Rate	-	\$0.03 per 1K tokens
Total Input Tokens	135.12 million	\$1351.20
Total Output Tokens	10.37 million	\$311.10
Total Expenditure	-	\$1662.30

Table 5: GPT4-Turbo Token Usage and Associated Costs

The access to the LLM and associated compute resources through Microsoft Azure was jointly funded through a partnership between our university and Microsoft. The total expenditure for the GPT4-Turbo usage was \$1,662.30. This funding enabled us to leverage the capabilities of GPT4-Turbo to successfully meet our research objectives.

D BERTOPIC HYPERPARAMETERS

```

1 # Load CSV data
2 df = pd.read_csv(filename)
3
4 # Concatenate title and description into one text column
5 df['text'] = df['title'] + " " + df['description']
6
7 dim_model = UMAP(n_neighbors=15,
8                  n_components=10,
9                  min_dist=0.0,
10                 metric='cosine',
11                 random_state=42)
12
13 # Initialize BERTopic
14 topic_model = BERTopic(
15     nr_topics="auto",
16     min_topic_size=100,
17     n_gram_range=(1, 2),
18     umap_model=dim_model
19 )
20
21 # Fit BERTopic on the concatenated texts from the CSV
22 topic_model.fit(df['text'])

```

E IRR MEASUREMENTS

For the Classification and Prevalence stages, which associated a label to a given text input, we calculated inter-rater reliability (IRR). Three trained annotators annotated 100 randomly chosen samples each for both stages. We calculated human-human IRR and human-LLM IRR using Fleiss' Kappa. To calculate human-LLM IRR, we generated the human labels by taking the majority vote of the three annotators. If no clear majority existed, i.e., all three annotators chose different labels, we labeled the class as 'Other'. The results were as follows: For the Classification stage, the human-human IRR was 0.59, and the human-LLM IRR was 0.54. For the Prevalence stage, the human-human IRR was 0.63, and the human-LLM IRR was 0.61. These results suggest good agreement according to the interpretation guidelines by Fleiss et al. [20]. Crucially, the human-human and human-LLM IRR values differed by less than 0.05, demonstrating LLMs can substitute human labels.