

## Responsible NLP Checklist

Paper title: *How to Contextualize Empirical Data for Risk Analysis with LLMs: A Case Study of Power Outages*

Authors: *Haiyun Huang, Yukun Li, Marco A Pretell, Jacob Naroian, Ebadah Khan, Liping Liu*

How to read the checklist symbols:

- the authors responded 'yes'
- the authors responded 'no'
- N/A* the authors indicated that the question does not apply to their work
- the authors did not respond to the checkbox question

For background on the checklist and guidance provided to the authors, see the [Responsible NLP Checklist](#) page at ACL Rolling Review.

---

### A. Questions mandatory for all submissions.

#### A1. Did you describe the limitations of your work?

*This paper has a Limitations section.*

#### A2. Did you discuss any potential risks of your work?

*The research focuses on statistical risk analysis of power outages using public historical data. It does not involve sensitive personal information, social engineering, or the generation of harmful content that could pose significant safety or security risks.*

### B. Did you use or create scientific artifacts? (e.g. code, datasets, models)

#### B1. Did you cite the creators of artifacts you used?

*We cite the creators of the power outage data (mass.gov) and the weather models (NOAA) in Section 3.2, Appendix A.0.1, and Appendix A.0.2.*

#### B2. Did you discuss the license or terms for use and/or distribution of any artifacts?

*The data sources and their public availability for research purposes are discussed in Section 3.2 and Appendix A.0.3.*

#### B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?

*Our use of power outage and weather data is consistent with their intended purpose of public safety analysis, as detailed in Section 3.2 and Appendix A.*

#### B4. Did you discuss the steps taken to check whether the data that was collected/used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect/anonymize it?

*The datasets consist of timestamps, geographic coordinates, and weather conditions. No personally identifying information or offensive content was present in the raw data or included in our study.*

#### B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?

*The Responsible NLP Checklist used at ACL Rolling Review is adopted from NAACL 2022, with the addition of ACL 2023 question on AI writing assistance and further refinements based on ARR practice.*

*Detailed documentation of data attributes, including weather variables and geospatial features, is provided in Appendix A.*

- B6. Did you report relevant statistics like the number of examples, details of train/test/dev splits, etc. for the data that you used/created?

*We report data statistics, including the 2018-2023 time range and sampling details, in Section 3.2, Section 4, and the relative risk tables in Appendix B.1.*

**C. Did you run computational experiments?**

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?

*We specify the models used (Llama3-8b, Qwen-2.5-7B, and GPT-4o) and the experimental setup in Section 4.*

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?

*The experimental setup, including the data splitting and prompting strategy, is discussed in Section 4.*

- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?

*We report the area under the ROC curve (AUC) and the F1 score for the outage class as evaluation metrics in Section 4 and Table 3.*

- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation, such as NLTK, SpaCy, ROUGE, etc.), did you report the implementation, model, and parameter settings used?

*We specify the use of Llama3-8b-Instruct, Qwen-2.5-7B, GPT-4o, and GBDT models in Section 4 and Table 3.*

**D. Did you use human annotators (e.g., crowdworkers) or research with human subjects?**

- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?

*(left blank)*

- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?

*(left blank)*

- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating (e.g., did your instructions explain how the data would be used)?

*(left blank)*

- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?

*(left blank)*

- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?

*(left blank)*

**E. Did you use AI assistants (e.g., ChatGPT, Copilot) in your research, coding, or writing?**

- E1. If you used AI assistants, did you include information about their use?

*We used ChatGPT for language refinement and programming assistance, as noted in the Acknowledgments section (after Section 7).*