

## Responsible NLP Checklist

Paper title: *iTAG: Inverse Design for Natural Text Generation with Accurate Causal Graph Annotations*

Authors: *Wenshuo Wang, Boyu Cao, Nan Zhuang, Wei Li*

How to read the checklist symbols:

- the authors responded 'yes'
- the authors responded 'no'
- 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?  
*Limitations section.*

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

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?  
*NA*

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?  
*Section 4.1.2; Appendix D.1.2D.1.3; Table 28.*

### C. Did you run computational experiments?

C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?  
*Section 4.1.1; Appendix B.3.2B.4.4; Tables 67.*

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?  
*Section 4.1.3; Section 4.2; Appendix B.4.5; Appendix C.1, C.4, C.5; Table 2 and Table 21.*

### 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.?  
*Appendix D.3.*

*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. ACL 2026 used a subset of ARR checklist form.*

- 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)?

*Sections 4.1.2, 4.1.3, and Appendix D.3; we did not report participant recruitment, compensation, or compensation adequacy because the human component of this study was limited to blinded trained-annotator labeling and human detectability judgments on de-identified/public text, and the papers claims concern annotation fidelity and benchmark transfer rather than any participant-level or demographic effect.*

- 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?  
*Sections 4.1.2, 4.1.3, and Appendix D; we did not report ethics-board approval or exemption because, as presented in the manuscript, the study is a low-risk text-only annotation/evaluation workflow over de-identified or publicly released materials, and it does not make any human-subject claim beyond the resulting labels and judgments.*

**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?  
*Acknowledgments.*