

AAAI ICWSM Paper Checklist

Written by the AAAI ICWSM 2024 Organizing Committee

pc.chairs@icwsm.org

Abstract

This document offers a Paper Checklist to be appended at the end of all submissions to, at a minimum, the September 2023 and January 2024 rounds of the AAAI ICWSM conference.

Overview

This document offers a checklist to append at the end of a AAAI ICWSM 2024 submission. The paper checklist has been adapted from the NeurIPS 2023 guidelines (NeurIPS 2021), the Natural Language Processing (NLP) reviewing checklist compiled by Benotti et al. (2023), and the consensus-based transparency checklist (Aczel et al. 2020). The checklist follows the references. While addressing these questions in the body of their manuscript, authors can explore the discussions provided in prior work (NeurIPS 2021; Aczel et al. 2020; Benotti et al. 2023; Ashurst et al. 2020; Gebru et al. 2021) as a starting point. The ethics reading list¹ compiled by Benotti et al. (2023) provides examples of papers discussing ethical considerations in NLP research.

Detailed Instructions

Please do not modify the questions and only use the provided macros for your answers. In your paper, please delete all text in the **Overview** and **Detailed Instructions** sections, as well as all subsection headers, keeping only the **Checklist** section heading above along with the questions/answers below.

For each question, change the default **Answer** to **Yes**, and **No**, **because**, or **NA**, when the question seems inappropriate for your research study. You are strongly encouraged to include a **justification to your answer**, either by referencing the appropriate section of your paper or providing a brief inline description. Within the Checklist section too, you may supplement your answers with a brief discussion that expands on answers to the checklist where necessary. For example:

- Did you include the license to the code and datasets? **Yes**, see the **Methods** and the **Appendix**.
- Did you include the license to the code and datasets? **No**, because the code and the data are proprietary.
- Did you include the license to the code and datasets? **NA**

Copyright © 2022, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹<https://github.com/acl-org/ethics-reading-list>

Paper Checklist to be included in your paper

1. For most authors...

- (a) Would answering this research question advance science without violating social contracts, such as violating privacy norms, perpetuating unfair profiling, exacerbating the socio-economic divide, or implying disrespect to societies or cultures? **Answer**
- (b) Do your main claims in the abstract and introduction accurately reflect the paper’s contributions and scope? **Answer**
- (c) Do you clarify how the proposed methodological approach is appropriate for the claims made? **Answer**
- (d) Do you clarify what are possible artifacts in the data used, given population-specific distributions? **Answer**
- (e) Did you describe the limitations of your work? **Answer**
- (f) Did you discuss any potential negative societal impacts of your work? **Answer**
- (g) Did you discuss any potential misuse of your work? **Answer**
- (h) Did you describe steps taken to prevent or mitigate potential negative outcomes of the research, such as data and model documentation, data anonymization, responsible release, access control, and the reproducibility of findings? **Answer**
- (i) Have you read the ethics review guidelines and ensured that your paper conforms to them? **Answer**

2. Additionally, if your study involves hypotheses testing...

- (a) Did you clearly state the assumptions underlying all theoretical results? **Answer**
- (b) Have you provided justifications for all theoretical results? **Answer**
- (c) Did you discuss competing hypotheses or theories that might challenge or complement your theoretical results? **Answer**
- (d) Have you considered alternative mechanisms or explanations that might account for the same outcomes observed in your study? **Answer**
- (e) Did you address potential biases or limitations in your theoretical framework? **Answer**
- (f) Have you related your theoretical results to the existing literature in social science? **Answer**

- (g) Did you discuss the implications of your theoretical results for policy, practice, or further research in the social science domain? **Answer**
3. Additionally, if you are including theoretical proofs...
- (a) Did you state the full set of assumptions of all theoretical results? **Answer**
- (b) Did you include complete proofs of all theoretical results? **Answer**
4. Additionally, if you ran machine learning experiments...
- (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? **Answer**
- (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? **Answer**
- (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? **Answer**
- (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? **Answer**
- (e) Do you justify how the proposed evaluation is sufficient and appropriate to the claims made? **Answer**
- (f) Do you discuss what is “the cost“ of misclassification and fault (in)tolerance? **Answer**
5. Additionally, if you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
- (a) If your work uses existing assets, did you cite the creators? **Answer**
- (b) Did you mention the license of the assets? **Answer**
- (c) Did you include any new assets in the supplemental material or as a URL? **Answer**
- (d) Did you discuss whether and how consent was obtained from people whose data you’re using/curating? **Answer**
- (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? **Answer**
- (f) If you are curating or releasing new datasets, did you discuss how you intend to make your datasets FAIR (see FORCE11 (2020))? **Answer**
- (g) If you are curating or releasing new datasets, did you create a Datasheet for the Dataset (see Gebru et al. (2021))? **Answer**
6. Additionally, if you used crowdsourcing or conducted research with human subjects...
- (a) Did you include the full text of instructions given to participants and screenshots? **Answer**
- (b) Did you describe any potential participant risks, with mentions of Institutional Review Board (IRB) approvals? **Answer**
- (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? **Answer**

- (d) Did you discuss how data is stored, shared, and de-identified? **Answer**

References

- Aczel, B.; Szaszi, B.; Sarafoglou, A.; Kekecs, Z.; Kucharský, Š.; Benjamin, D.; Chambers, C. D.; Fisher, A.; Gelman, A.; Gernsbacher, M. A.; et al. 2020. A consensus-based transparency checklist. *Nature human behaviour*, 4(1): 4–6.
- Ashurst, C.; Anderljung, M.; Prunkl, C.; Leike, J.; Gal, Y.; Shevlane, T.; and Dafoe, A. 2020. A guide to writing the NeurIPS impact statement. *Centre for the Governance of AI*. URL: <https://perma.cc/B5R8-2B9V>.
- Benotti, L.; Fort, K.; Kan, M.-Y.; and Tsvetkov, Y. 2023. Understanding Ethics in NLP Authoring and Reviewing. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics: Tutorial Abstracts*, 19–24.
- FORCE11. 2020. The FAIR Data principles. <https://force11.org/info/the-fair-data-principles/>.
- Gebru, T.; Morgenstern, J.; Vecchione, B.; Vaughan, J. W.; Wallach, H.; Iii, H. D.; and Crawford, K. 2021. Datasheets for datasets. *Communications of the ACM*, 64(12): 86–92.
- NeurIPS. 2021. NeurIPS 2021 Paper Checklist Guidelines. <https://neurips.cc/Conferences/2021/PaperInformation/PaperChecklist>.