

Responsible NLP Checklist

Paper title: *InsAT: Instance-aware Semantic Alignment and Transfer from Human Object Keypoints for Zero-to-Few-shot Action Understanding*

Authors: *Kazuki Tsutsukawa*

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?

We address this point at the Ethical Considerations section (after Limitations).

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?

All pose data consists of anonymized keypoints, and textual descriptions were generated from class labels and, in some cases, from publicly available images.

- 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?

See Section 4.1 and Appendix (B) for dataset statistics and train/test splits, and Section 3.3 and Appendix (E.1) for details on the generated textual descriptions.

C. Did you run computational experiments?

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

In Section 4 (Experiments), we describe the experimental setup, and provide implementation details and hyperparameters in the Appendix (C).

- 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 per-split results for zero-shot experiments without aggregated statistics, and the mean over five runs for zero-to-few-shot adaptation. For comparisons with RGB-based VLMs, we follow prior work and report results averaged over three runs.

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.

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.?

N/A

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)?

N/A

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)?

N/A

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

N/A

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?

AI assistants were used only for translation and language polishing. Their use is not described in the paper, as they were not involved in the research methodology or experimental content.