

Responsible NLP Checklist

Paper title: *SemVink: Advancing VLMs' Semantic Understanding of Optical Illusions via Visual Global Thinking*

Authors: *Sifan Li, Yujun Cai, Yiwei Wang*

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?

Section 4.5 (Discussion).

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

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

Section 4.

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

Author information block (front page). We provide the open code repository link in the author block of the paper, where the license (CC BY 4.0) is clearly specified. The models used (e.g., O3, Gemini, Claude, Qwen2-VL, DeepSeek-VL2) are publicly accessible through official providers.

- 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)?

Section 3.1 (Data Construction) and Section 4.1 (Experimental Setup). We used only officially released APIs or open-weight models (O3, Gemini, Claude, Qwen, DeepSeek, etc.) within their intended research or evaluation context. For our created artifact (HC-Bench), we specify that its intended use is as a research benchmark for hidden-content recognition, which is consistent with the research access conditions of the underlying generative models (Stable Diffusion, ControlNet).

- 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?

Section 3.1 (Data Construction). HC-Bench contains only synthetic images generated with Stable Diffusion + ControlNet, consisting of hidden texts and objects. It does not include real people, personally identifying information, or offensive content. Therefore, no anonymization or filtering was required.

- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
Section 3.1, Table 1.
- 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 3.1 (Data Construction) and Section 4.1 (Experimental Setup).
- 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?
Section 3.1.1 (Implementation Details) and Section 4.1 (Experimental Setup).
- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?
Section 3.3 (SemVink Method) and Section 4.1 (Experimental Setup).
- 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 (Experiments).
- 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?
Section 3.1 (Data Construction) for Stable Diffusion + ControlNet details, and Section 3.3 (SemVink Method) for preprocessing parameters (downsampling resolution, brightness/contrast adjustments).
- 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?
(left blank)