

Responses misunderstandings paper v2

🕒 Date Created @October 9, 2024 1:15 PM

Summary

We thank the reviewer and the meta-reviewer for carefully going through the new version of the paper and ultimately suggesting that the use of tweets is perhaps more problematic than beneficial. Although we believe there is value in having some of these arguments in a position paper of this sort, ultimately, we agree it makes the article too informal in terms of its scientific arguments. Therefore, we followed the suggestion of the meta-reviewer and all references to tweets are dropped.

Other than that, there is no mention of tweets, nor are there any direct arguments against polemic positions held by AI scientists (eg Geoff Hinton). In some cases, however, we have included statements from published works, but often a remark about the general arc of the field — similar to a quotation from Judea Pearl. Based on the feedback in the review, and the non-issue with the Pearl quote, we believe this should be OK and was suggested anyway.

We hope that with these changes, the article is satisfactory. Here is a list of all the changes we have made. After that, we will once again go over the points raised in the meta-review and reiterate how we have addressed them or whether the points have been addressed.

Changes to the article

Here is the list of changes we have made in a sequential manner, with references to sections in the old version. As suggested in the title, we dropped "misunderstandings in social media." In the abstract, we dropped the phrase "circulated in social media." In the preface, we removed the point that disposition reflects objections raised by Hinton and others. At the end of section 1 in the last paragraph, we once again dropped the mention of "circulated in social media." Section 2 is now removed; however, some relevant paragraphs have been slightly moved to other sections, such as the last paragraph.

In section 3 (of the old version), which talks about logic being old-fashioned, we introduced new references and phrasing that suggest machine learning is the only viable approach to building AI systems to make the same point. In section 3.1, we removed the mention and discussion of Hinton. Nonetheless, we expanded a bit on the theory of everything approach in science by referencing an article by a leading neural network expert. At the end of that section, we dropped the mention of Geoff Hinton. In section 3.2, we tried to articulate the point about the dichotomy using references from new articles and talked about the subtlety of seeing logic only for discrete systems and machine learning for continuous data and explained the nuance here.

In section 3.5 of the old version, we fixed the interpretation of disjunction as max in fuzzy logic. In section 3.6 on monotonicity, we dropped the Neal tweet and tried to contrast how exceptions are handled in machine learning. In section 3.8.1 on flippant reasoning, we added new references on how continuing explorations of the logical reasoning abilities of large language models are being explored in many papers right now. We also included a statement from Henry Kautz from published work, similar to the Pearl quote.

We removed the tweet in relation to the intentional stance and added new references on natural language semantics and the interpretation of connectives concerning natural language.

In section 4.2 of high-level knowledge of the old version, we removed the tweet by Marcus and instead included discussions on his book and other related

explorations.

In section 4.3, we removed the tweet on the symbolic lingua franca as it didn't add much to the discussion. Likewise, figure 8 and figure 9 on circuits were removed and instead the references to the scientific papers themselves were added.

Finally, section 5.3, which was the more polemic part at the end of the conclusions, was rephrased to make a broader point that one could celebrate the success of models like AlphaGo as well as large language models, but still recognize the need for formal reasoning and correct and verifiable answers. We wrap up with that point along with a quote from an article by Hector Levesque on the notion of similar bulletism and why that's problematic in AI.

Below is a list of points from the meta-review. We scratch out the bullet points when the points have been addressed. Most of the points don't need follow-up responses because we have addressed them, or explained above. However, if follow-ups are necessary, we have provided a short explanation.

Points extracted from the meta-review

Main issues with the current version:

- ~~Extensive use of Tweets as the main narrative and source of views from the machine learning community~~
- ~~Inappropriate use of tweets to reference scientific opinions~~
- ~~Presence of ad hominem arguments~~

Suggestions for improvement:

- ✓ ~~Change the narrative to no longer be based on tweets~~
- ✓ ~~Remove all needlessly inflammatory sentences~~
- ✓ ~~Find substantial versions of arguments against logic and counter them, instead of arguing over Tweets~~

- ✓ Remove all ad hominem arguments and restrict the discussion to scientific arguments

Conditions for acceptance:

- ✓ Change the narrative to no longer be based on tweets
- ✓ Remove all needlessly inflammatory sentences

Specific suggestions for revision:

- ✓ Cut the second line of the title
- ✓ Remove all references to Tweets and rephrase them as scientific arguments from the machine learning community
- ✓ Remove "strangely" (last line of page 4)
- ✓ Remove ad hominem arguments from the first few paragraphs on page 4
- ✓ Clarify the separation between the role of logic inside an AI architecture and its use as a meta-theoretic tool (e.g., on p.6 and in 4.1)

- **We decided to explain this in Section 4.1 and expand it in a new paragraph at the end. This makes sense because in subsequent subsections, these points are discussed at length.**

- ✓ Clarify how the argument in section 3.4 implies the claim in the title

- **The title is perhaps misleading. Therefore, we have changed the title.**

- ✓ Correct the definition of fuzzy logic operators (min,max) for conjunction and disjunction on p12

- ✓ Remove sections 3.8.1 and 3.8.2

- **revised 3.8.1 by removing tweets, and discarded 3.8.2**

- ✓ Remove the last sentence of 3.8.3

- ✓ Update or remove reference to "recent" work on p19 (currently 12 years old)

- ✓ Remove "it seems irresponsible..." on p20

- **revised**

- ✓ Revise or remove "in a rough sense... is not that different..." on p22

- **revised**

- ✓ ~~Remove 2nd sentence in section 5 on p24~~
- ✓ ~~Remove first sentence of second paragraph in section 5 on p24~~
- ✓ ~~Remove section 5.3 on p25~~



We have retained the non-polemic components of this section and expanded it to include a discussion on why the field might be interested in exploring different types of architectures beyond traditional classical machine learning. We hope this section now reads well.

- ✓ ~~Consider removing the final sentence of 5.2~~

General revisions:

- ✓ ~~Sanitize the paper from social media talk, ad hominem arguments, and indignant tone of voice~~
- ✓ ~~Focus on solid arguments, explanations, and examples to merit publication as a survey discussion~~

Note: This is not an exhaustive list. The author should review the entire paper for any additional areas that need revision to meet the reviewers' concerns.



Hopefully, with the changes we have made, the reviewer is satisfied that we have addressed all the changes in line with the intent of the review.