Nesy journal rebuttal

We thank the reviewers for their feedback. We have made several changes to the paper. In addition to adding some new references that discuss state-of-the-art developments in neurosymbolic AI -- for example, we now include discussions on the use of temporally extended reward functions in reinforcement learning and checking the reasoning of large language models using symbolic reasoners -- we have updated the abstract and the introduction. And, we have polished some of the arguments within each of the subsections. There is now a greater amount of obvious overlap between the different dimensions, and we emphasize this point.

Correspondingly, the conclusion has also been updated, and a new section have been added on how logic benefits from the machine learning literature. More generally, in some of the subsections, short introductions to some of the technical concepts are provided as a kind of schemata representing the formal objects appearing in technical papers.

Ultimately, as the meta-reviewer notes, this is not a technical paper in itself. It is something between a position paper and a survey paper, which does two things. On one hand, it discusses the breadth and diversity of solutions encompassed within symbolic logic. We believe that many of these dimensions are not obvious to people outside the logical community, and certainly, even those working within certain areas of logic might not be aware of the latest developments in statistical relational learning. On the other hand it points to common objections to using logic when building complex AI systems involving machine learning. This reflects objections raised by Geoff Hinton and echoed by others.

Therefore, this paper serves as a survey that tackles both of these aspects. As a result, we have also updated the acknowledgment section to clearly indicate the intended audience for this work.

We provide explanations for some of the objections raised below using itemized bullet points. We have also structured the paper in the updated version differently. We differentiate criticisms about the use of logic and contrast this against all the positive aspects, in which looking at the complementary properties of logic and learning, are discussed.

---

meta-review This is a type of meta-review, as there was some confusion about the type of paper. It in fact is supposed to be a position paper for the inaugural issue.

I have communicated with the reviewers, and while their reviews haven't been updated, the result of the discussion was that even as a position paper, the paper needs improvements to strengthen the qualitative arguments.

- [x] changes to the abstract to reflect better that it's a position paper.
- [x] removed the quote at the top from Michael Crichton just to avoid confusion about vague terms.

- [x] The paper has been significantly expanded with a deeper set of qualitative arguments, including boilerplate constructions of some of the technical aspects. The number of references has been bolstered, and the abstract and introduction have been slightly adapted to reflect the type of paper it is.

---

review 1:

The main objective of the paper is to conduct a critical review of the misunderstandings about the usage of logic in neurosymbolic systems.

- [x] clarified this: As such, this cannot be categorized as a research paper but as an overview paper.

The core of the paper consists of quotes from posts that appeared on social media by relatively notable researchers.

- [x] clarified and added: They do not constitute well-developed and scientific arguments.
- [x] made a statement about this: Thus, I am not convinced that reactions to them should be included in an article of a scientific (peer-reviewed) venue.
- [x] added notes: They also cannot constitute representative data for a community since the number of posts considered by the authors is too few.
- [x] hopefully changes addressed this: Additionally,
- [x] Tried to include a number of additional references that express related positions; in fact, now all of these counterexamples now include justifications from additional references: the counterarguments proposed by the author are not clearly explained, and I am not fully convinced by them. I think that the paper is more appropriate for a personal blog or an open discussion forum rather than a scientific paper.
- [x] clarified and added remark: In general, I believe that a better way to rebut misunderstandings about a scientific methodology is to produce scientific results that clarify such misunderstandings.
- [x] tried to include screenshots to twitter posts everywhere: Also, referencing social media posts in an archival paper is not a good practice as they might not be available in the future.

---

review 2:

The scope and intended goal of this paper is very relevant. But the intended target is not clear. this paper seems more a collection of blog posts / unproved statements than a scientific paper. But maybe this is not intended to be a scientific paper? This would be justified by the title. What "misunderstanding in the social media" has to do with the underlying science?

- [x] We have now added more arguments with concrete references to technical objects that specifically discuss how some of these objections to the use of logic have been overcome in recent state-of-the-art solutions in neurosymbolic AI.

What is usefulness of the notion of "thinintelligence"? Like other notions occurring in the rest of the paper, this term is associated to a lot of trivial , and sometimes ill-defined notions.

- [x] We agree that this was not properly defined and have removed this quote. Instead, the argument is more about the general need to integrate reasoning and learning.

And what about the promise of neuro-symbolic reasoning? Which promise? The author just quickly lists a bunch of citations. But could he give us a concrete example where Logic (+ML) could help solving in a problem that could NOT be solved by the SoA in ML. Something going beyond the usual arguments that we have heard many times (e.g., lack of explainability, the lack of semantics).

- [x] We have added many arguments to the updated version, including mentioning the use of model counting for exact inference, especially with hard constraints, the use of logic-based oracles to check the consistency of reasoning by large language models, and the use of temporal reward functions in reinforcement learning. Ultimately, the question of what goes beyond what is solved by state-of-the-art machine learning is a nebulous measure because many state-of-the-art machine learning papers are, in some sense, neurosymbolic even if they don't acknowledge it as much. However, we hope that by adding more instances of what is possible within logic and what it can capture, this criticism can be substantiated. We added a note on this in the updated paper.

This kind of problems apply pretty much to all sections of the paper. One last observation is on the use of the term "metalogic". From what he writes I doubt that the reader knows what a metalogic is. Or even what a metatheory is.

- [x] We are not sure what the reviewer means here, but we are simply suggesting that logic could be an underlying framework and formal structure that could contextualize and reason about machine learning outputs. In this sense, it is modeling the system at a metalinguistic level. We don't claim anything more.

This paper must be completely rewritten to be publishable. The reason why I am suggesting resubmission is the relevance and timeliness of the topic. And I do hope that the author works on this. Next time providing a better written, more convincing argumentation.

- [x] The reviewer suggested a complete rewrite, and we have attempted to update and add many details to substantiate the criticisms against the paper. Of course, it is possible that the reviewer meant something more comprehensive, including an actual rewrite of all the concepts, but we feel this would result in a completely different paper altogether. Note that the motivation here is to primarily target common criticisms against the use of logic by people from the machine learning community who don't usually peruse logical literature. As we have argued in a new section at the start called "Why Tweets," we point out that some of these opinions are informally held and never explicitly ironed out. This is precisely why we want to tackle these. Thus, the only way we see to make these arguments is to refer to these kinds of social media rants and try to make a case for why

the use of logic is powerful enough to address those concerns. We hope that with these new additions, the reviewer is slightly more satisfied with the outcome.