Editors-in-chief feedback

... based on the reviewer's comments, your paper requires minor revisions [...] there are several presentation issues, which need still to be taken into account for the next submission.

Thank you very much for considering this draft, we fully understand the need for this revision, for the great benefit of the article, and detail how we take these reviewer comments here, using italic font.

Sorry for the so long delay in resubmitting this publication, we hope that this will still be acceptable.
- One reason is that we followed reviewer 1's advice of « give it to a colleague who did not write it and have them read it out loud for a bit... » and this added an additional delay.
- Another reason is that the helpful comments of the reviewers regarding much more experimental illustrations (and we fully agree with this) require to contact several collaborators, thus being cost consuming.
- We also have improved technical precision and other elements beyond reviewers' comments but inspired by them.

Although the paper is now longer, we certainly have not extended its purpose, we have only improved it.

Review #1 feedback by Ben Wright<dr.ben.j.wright@gmail.com>

Recommendation: Minor revision

Detail Comments: content issues

Overall, I found the concept and discussion relevant to the NAI journal and the audience. However, I did find that the paper has a few issues to address before giving a full accept.

Thanks for this positive view, we fully agree with your reservations.

- The main novel idea discussed here is the concept of the editing distance between two data structures.

Thanks, we agree: editing distance between two data structures and what results from it.

- This aspect is intriguing and leads me to think about it in the context of some ideas not discussed by the authors (how would this editing distance and created data relate to synthetic data measures).

This a very good idea, your idea, we have added this contextualization in the introduction.

- I am a bit surprised to see that no discussion towards path planning problems or planning problems in general were used as the framing for the editing distance plays very much into those classical AI bins, as distance between two state spaces is an often used heuristic.

We fully agree. and this is the next step of this ongoing work, as you can see in the next open code package in progress
`https://line.gitlabpages.inria.fr/aide-group/symbolingfield`
and a very preliminary draft
`https://www.overleaf.com/read/wrqgttdzgsbf`
... however, to develop this idea up to an effective "universal" algorithm in this non Euclidean space, we came to the point that we need parameterized sub-harmonic potentials, which requires yet another ream of derivation, and it appears to us as better reported after this first step validated.

- One of my bigger issues with the current structure of the article is that it stays too far into the abstract without contextualizing the theory discussed.

Yes, this is a precious advise and we have added a full explicit subsection on contextualizing our developments. We also have added a few set of concrete experimentation in order to illustrate the abstract developments.

- Additionally, while mentioning in the intro and abstract that this relates to the context of "cognitive computational neuroscience´´ the authors never bring this back in the conclusion to put a bow on the whole concept.

Yes again, this has been added thanks to you.

- Following this thought, I was underwhelmed by the papers current level of discussion around "implementation´´. It was my understanding from the abstract that this was "implemented in detail at the programming level´´ yet by the time I ended section 5, I did not really see any discussion of running an implementation (even a prototype).

OK for this point, we have reported more details about the implementation run, in the experimental part.

- Adding more specifics and proof-of-concept implementation discussions (I.e. "we ran a small data structure of X through a usual ML algorithm and our results were..... ") would greatly improve the paper's impression and quality.

Yet another thank you, for this improvement requirement, we have enhanced the experimental part, following this advice.

Detail Comments: presentation issues = *There are additional issues of clarity that I would hope can be improved. The paper has an overabundance of "i.e." and "e.g." examples, sometimes within the same sentence. If the authors wish to expand on an idea they are discussing, it is perfectly fine to just state that in the next sentence.*

You are right. The idea was to be able to, e.g., for neuro versus computer science colleagues, provide equivalent semantic expression, i.e., make the idea accessible via different formulations :). But you are right, this makes the text hardly readable and we thus have reformulated, using more sentences.

- The use of footnotes are also excessive, and (this is for footnote 11) sometimes referenced within the main body of the article.

You are right again. Please let us still use footnote, for mathematical derivations, because this really allows to enlighten the text. But we have reduced all footnotes used as "side" remarks.

- Footnotes 12 and 13 do not exist.

Our bad: they do exist, but we missed the fact LaTeX has "eaten" them, making them unreadable, now corrected.

- There is also an inconsistency between the Introduction and Symbolic Data Representation sections and the rest of the paper with the use of footnotes.

Yes, we have taken this into account.

- Following that, I would also say that most of the tables/charts and equations are not numbered or well placed within the context of their discussion (See pg. 8 line 37 equation and then its follow up on pg. 9 line 3).

Sorry for that too, now corrected. We number equations only if they are quoted in the sequel, and not only discussed locally, to avoid any useless references, otherwise we do.

- Equations are also introduced in some footnotes (Also page 8),

Of this is a strong presentation mistake, sorry, now corrected.

- I'm not sure the reasoning here versus a discussion of it in the main body of the article. Likewise, the overuse of the caption text for Figure 4 (pg 8 lines 23-28) while simultaneously not discussing the figure in the main body of the article. Figure descriptions should be one sentence, at most.

OK for this we have re-balanced the elements regarding the figures, and rewrite page 8 accordingly.

- I would also like to say that, in the context of Figures, make sure to discuss the full figure and why it is all in the figure before the figure appears. Figures 2 and 4 especially fall into this, but a look over all Figures would be beneficial.

Thanks for this too, done.

- Stylistically, it seems there are some presentation issues with how the bulleted lists are used in the paper. The readability is off, and sometimes it seems that the tables are being used for the same context as the bulleted lists. (See page 6, line 11; pg 9, line 4; pg 10 line 28; pg 12 line 32 -> pg 13 line 3 is especially egregious)

Sorry, this is now corrected.

- The use of citations and background/relevant work is odd for this paper. My interpretation for the NAI Journal is that you will have a wide mixture of audience readers with a variety of backgrounds, so the assumptions taken may not always be correctly assumed.

Yes the wide mixture of audience is the dream here, we have carefully revisit all elements of the paper, to improve this aspect.

- For instance my background is on the Symbolic AI/Logic side and would need a more concise introduction to some of the "usual ML algorithms" (pg 14, line 31)

Now done, thanks to you. Notice we have canceled the too general ML word to precisely talk about RL (Reinforcement Learning), with explanation and references, in order to focus on what has been precisely done.

- Also, especially in the context of section 1, there is a high risk of this paper beginning to cite papers just to cite them. A number of the citations are used as a "see this paper for a thing[X]" and then does not have followup or contextual

*nuance for why it is really being discussed within the context of *this* paper. Pg 3 lines 4-7 are an example of this to me.*

OK, agreed, and now corrected for all related references.

- The overall presentation of the sections could use a bit of a revamp as well. The idea of combining the ending of your paper along with the "discussion" sits wrong with me. The authors are still in the "introducing new ideas to the conversation´´ segment during the discussion and that I find is best reserved for sections outside of conclusions and last sections of the paper. To that end, I would have section 5 not be the concluding section of the paper.

We have restructured the paper accordingly.

*- Section 2 feels very aligned with sections 1.1 and 1.2. I would probably combine them all together for your Section 2 *BACKGROUND* and then move onto the full discussion of the editing distance stuff in Section 3.*

We have also followed this advice, considering section 1 as a complete introduction plus adding the contextual section you nicely proposed, thus grouping 1.1 + 1.2 and 2.

- I also think Appendix A.1 would be fine as just part of the main body of the writing to give better context to the discussion around data structures.

Agreed, thanks, done.

+ Other Minor Things,

- References aren't an Appendix, they are just References.

A LaTeX mistake, now corrected.

- Pg 25, Line 14 – Why is there a different use of citation here?

Again a LaTeX mistake, now corrected.

- Give everything another reading pass, give it to a colleague who didn't write it and have them read it out loud for a bit and see if it matches with your intent on the writing.

This is an excellent idea, this explains the delay, but is is done.

Review #2 feedback by Lia Morra<lia.morra@polito>

Recommendation: *Major revision*

Paper summary:

- The manuscript proposes a parameterized edit distances between hierarchical symbolic data instances. This distance defines a metric space at the symbolic level and a corresponding mapping to a scalar field. This dual representation can be manipulated at the symbolic and sub-symbolic level using the geometric properties and operators defined in the manuscript. Examples are provided to show how the proposed representation could be used to support problem solving tasks such as morphing between two data instances, extrapolating from a data instance under certain constraints, or analogical reasoning. The proposed metric space is relevant to the NeSy community, as many NeSy techniques rely on numeric representations of symbolic concepts to combine explicit knowledge

representation and reasoning with statistical inference. The proposed technique is thus complementary to other representations, such as numerical embeddings as typically used in other family of models such as Logic Tensor Networks or Vector Symbolic Architecture.

Thank you for this paper summary which plainly corresponds to what we have attempted to do.

- The main weakness of the manuscript is, in my opinion, readability, as the manuscript is at times difficult to follow, and would require extensive proofreading.

Yes we agree, and have reworked the paper accordingly.

Detailed remarks:

- In the sentence (page 2) "This also allows us to clarify that the numerical, i.e., quantitative, aspects mainly stand on the biophysical side, while symbolic representation mainly stands on the behavioral stage, of course not exclusively, but more than a gradient" what does gradient mean?

This is an odd formulation, sorry, corrected thanks to you.

- In the sentence, "In addition, formal reasoning is by no means less costly than billions of operations of deep networks computation, as studied by [19], which limit their academic use for usual research teams, beyond raised ecologic and ethical issues, discussed by the authors" I do not understand what the limited academic use refers to in this context.

We have entirely rewritten the statement and better referred to the [10] as required in your next feedback.

- The manuscript makes an extensive use of references in bracket [X] as subject or complement in sentences, which makes sometimes the text difficult to follow without continuously looking up the bibliography. Aside from the grammar/style point of view, which is subjective, sometimes the meaning is also unclear. An example is the sentence "It has a "meaning" in the sense of [31], as reviewed and discussed in [53], when it is semantically grounded, in the sense of the "symbol grounding problem", ...", I would consider rephrasing and briefly summarizing the meaning as defined in [31] so that the text is self-consistent.

Yes, this is right, and in accordance with reviewer 1, we have revisit each reference to correct this aspect.

- Section 2.1 introduces the concept of sign, index and icon from semiology as a way to define symbols. While I find the section interesting, I am not sure how the concepts explained here, especially that of ungrounding, relates to the proposed metric space. In addition, I would suggest to revised the text and Figure 2 caption in order to clarify what it is meant by hierarchy of signs and symbols. As far as I understand, the hierarchy symbol/index/icon is not a taxonomy (an icon is not a subclass of index, at least according to the original definitions by Pierce), but the hierarchy is more related to the complexity of the cognitive processing required for the emergence of specific type of signs.

Yes, definitely, and we have followed your advice, by first better review in

a concise way, the invoked semiology taxonomy and then explain the link with the present work. The hierarchy is more related to levels of abstraction, as now better explained thanks to your comment.

- I would avoid to refer to the text in the figure caption (see text), and rather make them more self-consistent if possible.

Agreed and thanks, corrected.

- The relationship between the proposed data structure and the VSA architecture, especially in the neural spiking framework, is suggested at the end of Section 2, and later on in the paper, but not elaborated. Suggesting practical ways that the proposed framework could be implemented would increase the value of the paper.

This is right, we are in fact publishing a work specifically on this subject, and have now better develop the idea in this paper.

- In Section 3.2, the concept of resource is introduced, whereas Section 3.1 refers to a data type. What is the precise relationship between the two concepts?

Our bad, the word resource is used here in the semantic web sense and is useless and confusing, this is rewritten.

- I would consider using a running example to illustrate the concepts through the paper (for instance, in Figure 4, instead of item1, item2, etc. a realistic example could be given, as well as in Figure 10). This would also better clarify the role of a region.

This is a very good idea, now implemented thanks to you. In fact, we have maintained the actual Figure 4, because we need it as a stand of our presentation, but we also have introduced a "real" example on which we have worked recently.

- The conclusion mention that a preliminary version of the proposed method was used in a reinforcement learning setting. It would be interesting to introduce these potential applications early on in the manuscript, in the introduction and related work section.

Yes, in fact this this is the next step of this ongoing work, as you can see in the next open code package in progress
`https://line.gitlabpages.inria.fr/aide-group/symbolingfield/index.html`
and a very preliminary draft
`https://www.overleaf.com/read/wrqgttdzgsbf`
... however, to develop this idea up to an effective "universal" algorithm in this non Euclidean space, we came to the point that we need parameterized sub-harmonic potentials, which requires yet another ream of derivation, and it appears to us as better reported after this first step validated. Nevertheless, we now have better explain this point, thanks to you.

- A more in-depth comparison of the advantages and disadvantages of the proposed metric space against other types of groundings would be especially interesting. For instance, in practical applications, the proposed distance will depend on the number and quality of properties that are attached to each data type (as shown in Figure 4). What are the implications in terms of the applicability of

the proposed metric space with respect to, e.g., the quality or richness of the ontology?

This is a very good idea, thanks for the opportunity to add it.

- A detailed list of typos is available in the attached document. Some of the figures are probably not in vectorized form.

We are especially thankful about this meticulous and so helpful work, indeed integrated in this version.