Editors-in-chief, Neurosymbolic Artificial Intelligence and Editorial board feedback

... based on the reviewer's comments, your paper requires major revisions.

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

Review #1 Submitted by kaushikr@email.sc.edu

Recommendation: Minor revision

Detail Comments:

Summary of Paper: The paper presents a biologically plausible model for simulating Vector Symbolic Architectures (VSAs) at a macroscopic scale. The work builds on the Neural Engineering Framework (NEF) and introduces "relation maps" for handling symbolic processing tasks in a scalable way. It offers an interesting algorithmic approach that simulates symbolic computations using vectors in high-dimensional spaces. The authors provide an open-source implementation to benchmark the proposed approach and compare its performance against mesoscopic-level simulations.

I believe the paper addresses an important topic within the neurosymbolic AI community, specifically about biologically plausible symbolic processing.
- The proposed macroscopic simulation is a novel extension of existing VSA techniques, making it possible to efficiently simulate more complex symbolic operations. While I'm not an expert on the mathematical background that this paper requires, I found the paper to be technically sound.
- It offers a detailed explanation of how vector symbolic operations can be simulated at different scales, which was helpful for me as an unfamiliar reader.
- The introduction of "relation maps" was a uniquely interesting addition to VSAs, showing that the method can handle knowledge representations in a biologically plausible way.
- Lastly, the authors have provided an open-source implementation, which adds significant value for reproducibility and further research.

Thanks for this summary, and synthesis of the contributions, corresponding entirely to what we think have proposed.

Specific Comments: While I followed the detailed explanations to the best of my ability, perhaps certain technical explanations, particularly regarding the implementation details and the use of "relation maps" (an often overloaded term in mathematics), may be hard to follow for readers unfamiliar with VSAs or the NEF.

We have reviewed technical explanations and have specifically explicitized this particular choice of "relational map", discussing potential term overload.

Additional examples and more precise explanations could improve accessibility. I gather that this work is primarily theoretical; however; the experimental section could benefit from moving beyond benchmarking simple ontology-based tasks (e.g., pizza ontology). Expanding the validation to more real-world applications or scaling experiments would enhance the paper's contribution.

You are definitely right [if i may] and in full accordance with other reviewers, we thus have designed two additional numerical validations :
1/ We have reproduced one of the Widdows and Cohen large scale VSA storage experiment regarding a search engine that needs to be able to assess similarity between terms and documents, and
2/ We have also reused the previous data to experiment on the sequence encoding part of the Quiroz Mercado work, in both cases considering about 10^6 input tokens for 10^4 terms and 10^3 documents. We did not reproduce the whole experiments, but simply have benchmarked our implementation with such a rather large scale data set.

Review #2 Submitted by sfw5621@psu.edu

Recommendation: Major revision

Detail Comments: This paper discusses an intriguing application of Vector Symbolic Architectures (VSA), also referred to as Hyperdimensional Computing (HDC) in related literature. VSA is a family of models for representing and manipulating data in a high-dimensional space, originally proposed in cognitive psychology and neuroscience as a connectionist model for symbolic reasoning. The topic is notably promising, as it aims to advance VSA towards computational universality.

Thanks for this contextualization of the work.

However, the paper seems to overstate its contribution to the proposed implementation of VSA at a macroscopic level, due to weak empirical results and a lack of technical accuracy.

We fully understand these two feedback.

On one hand, we thus have designed two additional numerical validations :
1/ We have reproduced one of the Widdows and Cohen large scale VSA storage experiment regarding a search engine that needs to be able to assess similarity between terms and documents, and
2/ We have also reused the previous data to experiment on the sequence encoding part of the Kempitiya et al work, in both cases considering about 10^6 input tokens for 10^4 terms and 10^3 documents. We did not reproduce the whole experiments, but simply have benchmarked our implementation with such a rather large scale data set.

On the other hand, we have carefully reviewed, and when necessary rewritten, all technical statements, with their rationale, beyond the actual algebraic derivations.

Nonetheless, it provides contributions such as proposing an improved notion of belief for modal encoding within semantic vectors to implement partial knowledge, outlining basic steps of knowledge encoding such as binding and bundling, and introducing a novel data structure called associative map with the proposed macroscopic level implementation.

Thanks for this summary of the main contributions of the paper.

The paper still faces significant challenges, particularly in the experimental testing phase and in some theoretical aspects, with most of its content focusing on the macroscopic implementation of VSA. The empirical results presented are underwhelming. Although the authors argue that the main contribution of the paper is the implementation of VSA at a macroscopic level, the results primarily discuss the computation costs associated with VSA rather than practical implementation details.

This is definitely true, what we wanted to state is how much computation cost is saved simulation mesoscopic VSA at the marcoscopic level, we have now reduce this part in the experiment section, and have developed practical implementation details thanks to your input.

The preliminary experimental section is notably weak; the authors themselves admit it is far from complete. The only substantive detail is the representation of semantics in the example 'Luigi 0.5 eats thisPizza,' where 0.5 represents an introduced modality—a partial knowledge modal encoding innovation highlighted in my previous comments on paper contributions.

We also agree, the point is that we would have need a long development far beyond the scope of this paper to give enough details to be able to present a bigger experiment. We now have introduced two effective benchmarks and leaves this tiny result as an illustration.

In the 'Binding Magnitude Verification' section, there is a reference to a non-existent Table 4, which undermines the credibility of the results discussed.

Sorry for the caveat, the table was there but pushed at the end of the paper and the numbering was noty easily readabl, after LaTeX compilation, now corrrected.

In the "Knowledge Structure Encoding" section, Table 1 describes dictionaries/maps as lacking enumeration capabilities. This is confusing because, typically, maps or dictionaries do allow for the enumeration of keys, values, or key-value pairs, contrary to what the authors have described.

This is formally exact, but usually VSA implementations of thus do not, we have rewritten the section to avoid this confusing formulation, and clearly make the difference between VSA data structure and programming dat satructure.

Figure 5, found in the "Implementation at the Macroscopic Scale" section, presents a mystery. The figure's basic declaration does not clearly convey the operational details or behaviors associated with adding, updating, or removing symbols as described in the caption. It merely shows the data structure's format without any operational logic or methods that manipulate the map's contents.

This is true. More than that, presenting it as a "figure" was inappropriate, since the figure only states that we use a symbol-index to Symbol map. All this has been redesigned, and the operational methods that manipulate the map's contents are now described in the text, to avoid the confusion.

Furthermore, in the section on "Binding Canonical Representation," the author claims that applying the r1 and r2 operators recursively guarantees no residual bindings/unbindings. This statement lacks a formal proof or detailed example, which is necessary to verify its accuracy and applicability.

Indeed, we have now provide the evidence of this mechanism.

Lastly, I recommend that the authors check for typos in this paper, e.g., on page 9: 'guaranty' should be 'guarantee.' Although this shall not impact the final decision, it will need to be corrected.

We did, both by our best careful reading and using a grammatical tool.

In conclusion, while the paper sets forth a promising framework for advancing VSA in universal semantic computation, it falls short in providing empirical support and clarity in its theoretical assertions. Future work should focus on enhancing the experimental designs, clarifying theoretical explanations, and providing proofs for the claims made.

Thanks a lot of for this encouraging conclusion, and precious guidance for resubmitting a major revision.

Review #3 Submitted by alessandro.oltramari_960

Recommendation: Major revision

Detail Comments: This paper presents an innovative application of Vector Symbolic Architectures (VSA), also known as Hyperdimensional Computing (HDC)—computational frameworks designed for representing and manipulating data in high-dimensional spaces. Rooted in cognitive psychology and neuroscience as a connectionist model for symbolic reasoning, the study extends this paradigm by introducing a biologically plausible model for macroscopic-scale simulation of VSAs.

Thanks for this contextualization of the work.

Building on the Neural Engineering Framework (NEF), the authors propose "relation maps" to efficiently and scalably handle symbolic processing tasks. This novel algorithmic approach would not only facilitate symbolic computations using high-dimensional vectors but would also bridge the gap between theoretical constructs and practical applications.

This work represents an original contribution that leverages concepts and explanatory theories from cognitive psychology to narrow the gap between neural processing and symbolic reasoning.

Thanks for this summary of the main contributions of the paper.

However, the study reveals two interrelated limitations that warrant further exploration and significant modifications:

We fully understand and agree with the limitations pointed out.

1) The empirical example presented, which is based on a small ontology model, is constrained in scope. It remains unclear how the proposed methods would scale and apply to more complex, real-world scenarios.

We fully agree with this requirement, and as also pointed out by other reviewers, we thus have designed two additional numerical validations :
1/ We have reproduced one of the Widdows and Cohen large scale VSA storage experiment regarding a search engine that needs to be able to assess similarity between terms and documents, and
2/ We have also reused the previous data to experiment on the sequence encoding part of the Quiroz Mercado al work, in both cases considering about 10^6 input tokens for 10^4 terms and 10^3 documents. We did not reproduce the whole experiments, but simply have benchmarked our implementation with such a rather large scale data set.

2) Symbolic Representation and Complexity: In section 3.0.2, the authors assert that triple statements encapsulate the essence of "symbolic representation." This claim is debatable, as triples (e.g., subject-predicate-object structures) represent only one form of knowledge representation. More expressive frameworks, such as logical representations, frame-based semantics, and hyper-graphs, offer richer and more nuanced models. It remains an open question how the proposed approach could extend to accommodate these more sophisticated symbolic formalisms and how associative memory mechanisms would handle such levels of complexity in symbolic structures.

This is absolutely right. Although we mentioned «... basic idea of symbolic representation» more than «essence» the word «idea» is a clumsy choice. Your comment is even more interesting: We have to briefly, but significantly discussed the link with other representations. Logical representations has been extensively discussed ans related to ontology, so that the literature can be easily quoted. The link between frame-based semantics has also been studied, and we have to make it explicit, now done. The hyper-graph representation through there is a canonical representation as a bipartite graph requires more specific VSA representations (we make a preliminary proposal in the discussion) but mainly state that it is beyond the limit this work. This is issue is thus addressed twice in the paper, when introducing the symbolic representation choice and when discussing its *limit.*

Addressing these limitations could significantly enhance the study's applicability and theoretical robustness, particularly in advancing the understanding of how symbolic reasoning can emerge from neural-like computations.

Thanks a lot of for this encouraging conclusion, and precious guidance for resubmitting a major revision.