

Response to reviewers' comments

Dear Editor-in-Chief,

The authors would like to thank the reviewers for their valuable feedback and comments. Below is a list of changes made in response to the reviewers' suggestions. For each comment, we have provided a justification and indicate the corresponding changes in the manuscript (highlighted in yellow).

Review #1

This paper surveys neuro-symbolic AI (NSAI) architectures and examines their relationship to recent developments in generative AI. Building on Kautz's taxonomy, the authors propose a refined classification of NSAI paradigms, including an additional "fibring" architecture, and attempt to relate a range of generative AI methods to these paradigms. The paper also introduces evaluation criteria for comparing NSAI architectures and includes a 4D printing case study as an illustrative application.

Reasons to accept

The paper addresses a timely and relevant problem for the neuro-symbolic AI community: the lack of a shared architectural vocabulary and the growing need to situate NSAI relative to modern generative models. The revised manuscript is clearer about its scope as a conceptual and taxonomic survey, rather than an empirical contribution, and the architectural descriptions are more precise than in the original version.

The taxonomy-based organization provides a useful high-level structure that can help readers, especially newcomers, navigate a fragmented literature. The authors have substantially engaged with reviewer feedback and corrected several overgeneralizations in the discussion of generative AI methods.

We sincerely thank the reviewer for the thoughtful and encouraging evaluation of our manuscript. We appreciate the recognition that the paper addresses a timely and relevant issue for the neuro-symbolic AI community, and that the revised version has improved in clarity, scope, and architectural precision.

Reasons to reject

The manuscript is close to being acceptable as a survey article, but several issues should be addressed more explicitly.

While the fibring architecture is now more clearly defined (symbolic aggregation enforcing constraints across multiple neural components), its necessity as a distinct taxonomic category remains unclear. The authors should explicitly address whether fibring introduces an integration principle that cannot be subsumed under existing cooperative or nested NSAI architectures. This could be done, for example, by providing concrete examples where existing paradigms fail to capture the interaction pattern.

We thank the reviewer for this important comment. In the revised manuscript, we clarified in Section 3.5 (Fibring) that fibring is introduced as a proposed extension to Kautz's taxonomy rather than as an established consensus category. We added a paragraph explicitly distinguishing fibring from cooperative and nested architectures by emphasizing that fibring captures cases where multiple specialized neural components are coordinated through a

shared symbolic layer responsible for aggregation, conflict resolution, and cross-model constraint enforcement. We also added concrete examples, including multi-agent coordination, multi-scale design, and multi-view decision-making, to illustrate interaction patterns that are not adequately captured by a single neural-symbolic pair.

The revised manuscript appropriately clarifies that standard generative AI techniques (e.g., GANs, pre-training, fine-tuning, distillation) are not neuro-symbolic by default. However, the section still risks being interpreted as implying stronger connections than are currently established in practice. The authors should clearly label these mappings as conditional or hypothetical throughout the section, and distinguish between methods that are already instantiated in the literature and those that represent plausible future integrations.

We thank the reviewer for this helpful suggestion. In the revised manuscript, we modified the beginning of Section 4 (Leveraging NSAI in AI Technologies) and the introduction to Section 4.2 (Classification of AI Technologies within NSAI Architectures) to state explicitly that the proposed mappings are illustrative, conditional, and interpretive, rather than definitive taxonomic assignments. We also revised the relevant paragraphs on RAG, CoT prompting, fine-tuning, distillation, MoE, and multi-agent systems to clarify that these methods are not neuro-symbolic by default. Throughout the section, we now distinguish between approaches already instantiated in the literature as neuro-symbolic systems and approaches whose relation to NSAI remains partial, approximate, or prospective.

The evaluation framework is now more transparent, but the resulting tables remain based largely on expert judgment. The authors should more clearly state the limitations of these evaluations as qualitative comparisons, and avoid language that could be read as definitive rankings between architectures when empirical evidence is sparse.

We agree with the reviewer and addressed this point in the revised manuscript. In Section 5.2 (Evaluation Methodology), we added an explicit paragraph stating that the proposed ratings are a qualitative, literature-grounded synthesis rather than standardized benchmark scores obtained under a shared experimental protocol. We clarified that the yes/no judgments and the resulting High/Medium/Low summaries should be interpreted as structured comparative indicators, not as definitive rankings. In addition, we revised the wording in the Abstract, Section 5.3 (Results and Discussion), and the Conclusion to remove expressions that could imply universal superiority, replacing them with more cautious formulations.

Review #2

The authors have clarified the previous comments regarding certain NSAI architectures (e.g., CoT, RAG) not fitting into the exact categorization they attempted to map using Kautz's taxonomy. They acknowledged the concerns raised and softened their claims in the revised manuscript. I think overall they made a good attempt at organizing and mapping a broad range of NSAI architectures and proposing an extensive evaluation framework for assessing them.

We sincerely thank the reviewer for the positive assessment of our manuscript and for acknowledging the improvements made in response to the previous round of comments.