Grounding vs. Compositionality: On the Non-Complementarity of Reasoning in Neuro-Symbolic Systems

Extended Abstract

Mahnoor Shahid
Universität Duisburg-Essen
Germany
mahnoor.shahid@uni-due.de

Hannes Rothe
Universität Duisburg-Essen
Germany
hannes.rothe@uni-due.de

ABSTRACT

Neuro-symbolic AI is often motivated by the belief that once perceptual inputs are grounded into symbols, systematic compositional reasoning will follow. We test this assumption directly and find it to be false in general: strong grounding does not reliably translate into robust compositional generalization under distribution shift. We introduce the *Iterative Logic Tensor Network (iLTN)*, an extension of Logic Tensor Networks that learns an explicit multi-step inference procedure via iterative belief refinement under first-order constraints. Using a controlled taxonomy of compositional shifts—*entity*, *relational*, and *rule* composition—we compare grounding-only, reasoning-only, and joint training regimes on synthetic visual logic puzzles that disentangle perception from inference. Across all shift types, grounding-only models degrade sharply, while iLTN substantially improves zero-shot robustness by optimizing reasoning as a first-class capability. Our results show that grounding and reasoning are not automatically complementary, and suggest that agentic and multi-agent systems operating under changing entities, rules, and constraints should explicitly learn iterative reasoning rather than assume it emerges from grounded representations.

KEYWORDS

Neuro-symbolic AI; symbol grounding; compositional generalization; differentiable logic; multi-step reasoning

ACM Reference Format:
Mahnoor Shahid and Hannes Rothe. 2026. Grounding vs. Compositionality: On the Non-Complementarity of Reasoning in Neuro-Symbolic Systems: Extended Abstract. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), Paphos, Cyprus, May 25 – 29, 2026*, IFAAMAS, 3 pages. https://doi.org/10.65109/SDCB9121

1 INTRODUCTION

Modern neural networks can achieve impressive in-distribution performance yet remain brittle when generalization requires recombining familiar parts into novel structures [6, 10, 14]. This limitation is especially problematic for autonomous agents and multi-agent systems, where policies must generalize to new compositions of entities, relations, and constraints (e.g., new teammates, new task templates, new coordination rules). Neuro-symbolic AI seeks to address this by coupling neural perception with symbolic structure and logic-like inference [3, 5, 8, 21].

A common intuition is that *grounding is the main bottleneck*: once perceptual inputs are mapped into appropriate symbols, systematic reasoning should follow. Yet longstanding critiques emphasize that systematicity and compositionality are not guaranteed by connectionist representations alone [6, 7, 17]. More recently, diagnostic benchmarks for compositional reasoning (e.g., grounded language and visual reasoning) have revived the question of whether architectures learn *reasoning* or simply exploit correlations [11, 18, 25].

Claim (tested here). Grounding and compositional reasoning are *distinct capabilities*; improving grounding alone is insufficient for robust compositional generalization under structured shifts.

2 SETTING AND TAXONOMY OF COMPOSITIONAL SHIFTS

We study synthetic *visual logic puzzles* where an image I depicts an object-centric world state and must be solved subject to a knowledge base K expressed as first-order logic constraints. The task is to infer a consistent assignment (or set of assignments) that satisfies K. This setting is aligned with diagnostic reasoning paradigms such as visual question answering [1] and CLEVR-like controlled reasoning [11], while enabling clean separation of perception and inference.

To probe generalization, we adopt a decomposition of compositionality into controlled axes [10, 14]:

- **Entity composition:** test-time includes novel entities/symbols (new constants), while constraint schemas remain fixed.
- **Relational composition:** test-time introduces new axiom families/relations absent during training.
- **Rule composition:** test-time requires longer or qualitatively different multi-step inference chains than those needed in training.

This taxonomy is designed to distinguish "having the right symbols" from "using them systematically" [6, 17].

3 THE *i*LTN FRAMEWORK

Background: LTNs. Logic Tensor Networks combine differentiable predicate models with fuzzy logic semantics to integrate data and knowledge [2, 20]. LTNs are effective for enforcing constraints and performing soft satisfiability, but typical usage emphasizes one-shot optimization of a logical loss rather than *learning an explicit multi-step deductive procedure.*

iLTN: iterative refinement for learned multi-step reasoning. We introduce *Iterative Logic Tensor Networks (iLTN)*, which operationalize reasoning as an iterative refinement of a belief state $P^{(t)}$ over

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. https://doi.org/10.65109/SDCB9121

candidate symbolic assignments (e.g., per-cell distributions). Each iteration performs:

(1) **Constraint-driven update:** compute a differentiable satisfiability objective induced by K (as in LTN-style semantics) and update $P^{(t)} \rightarrow P^{(t+1)}$ via a learned or gradient-based step.
(2) **Differentiable discretization:** apply a relaxation that encourages increasingly crisp hypotheses while keeping gradients informative, similar in spirit to differentiable constraint solvers and grounding techniques [22, 26].
(3) **Adaptive unrolling:** allow the model to unroll more steps for harder instances, connecting to broader step-wise verification and iterative reasoning ideas [9, 13].

This architecture makes the *reasoning trajectory* a trainable object, rather than assuming compositional behavior emerges from grounded representations alone.

4 EVALUATION: GROUNDING-ONLY VS. REASONING-ONLY VS. FULL ILTN

To isolate contributions, we compare three conditions:

- **Grounding-only baseline:** end-to-end perception-to-solution training without explicit iterative reasoning supervision (a common "implicit reasoning" regime).
- **Reasoning-only ablation:** provides pre-grounded symbolic inputs and trains only the iterative reasoning mechanism (tests inference without perception).
- **Full iLTN:** trains perception and iterative reasoning jointly, optimizing both grounding and multi-step deductive behavior.

Our design parallels "disentangling reasoning from perception" methodologies used in neuro-symbolic VQA [25] and program/logic-guided vision systems [15, 16], while focusing specifically on compositional shifts.

5 FINDINGS

Across entity, relational, and rule composition regimes, we observe a consistent pattern:

Grounding-only fails under compositional shifts. Even when the model achieves strong in-distribution accuracy, performance deteriorates sharply on compositional splits, consistent with prior concerns that correlation-based solutions do not guarantee systematicity [6, 14, 24].

Explicit iterative reasoning improves robustness. The full iLTN exhibits markedly stronger zero-shot robustness under all three shift types, indicating that learning an explicit multi-step refinement procedure supports generalization beyond what grounding alone provides. This aligns with the broader perspective that compositional generalization often benefits from inductive biases or objectives that reflect compositional structure [4, 12, 18].

End-to-end training can benefit reasoning itself. In several cases, the full iLTN compares favorably to the reasoning-only ablation, suggesting that training reasoning under perceptual uncertainty can regularize inference and improve tolerance to imperfect representations, echoing themes in grounded world model learning [19] and grounded specification learning [23].

6 IMPLICATIONS FOR AAMAS

For autonomous agents and multi-agent systems, compositional generalization is often the rule rather than the exception: new teammates, new task graphs, new constraint sets, and new negotiation structures arise frequently. Our results caution against a design pattern that treats symbol grounding as a sufficient "unlock" for robust reasoning. Instead, reasoning should be treated as a *first-class, optimizable capability*, with explicit objectives and architectures that support multi-step inference [5, 8, 21]. This is directly relevant to agent planning under constraints, coalition formation with changing rules, and verification-heavy coordination settings where step-wise deduction matters.

7 CONCLUSION

We present evidence that grounding and compositional reasoning are not automatically complementary: solving grounding does not reliably yield systematic compositional generalization. By extending LTNs with explicit iterative refinement, iLTN learns multi-step deductive behavior that substantially improves robustness under entity, relational, and rule composition shifts. The key takeaway for agentic and multi-agent applications is practical: to obtain reliable out-of-distribution coordination and constraint satisfaction, reasoning must be explicitly optimized, not assumed to emerge from grounded symbols.

REFERENCES

[1] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C. Lawrence Zitnick, and Devi Parikh. 2015. VQA: Visual Question Answering. *Proceedings of the IEEE International Conference on Computer Vision (ICCV)* (2015), 2425–2433.
[2] Samy Badreddine, Artur d'Avila Garcez, Luciano Serafini, and Michael Spranger. 2022. Logic Tensor Networks. *Artificial Intelligence* 303 (2022), 103649.
[3] Bikram Pratim Bhuyan, Amar Ramdane-Cherif, Ravi Tomar, and TP Singh. 2024. Neuro-symbolic Artificial Intelligence: A Survey. *Neural Computing and Applications* (2024).
[4] Henry Conklin, Bailin Wang, Kenny Smith, and Ivan Titov. 2021. Meta-learning to Compositionally Generalize. arXiv:2106.04252
[5] Jonathan Feldstein, Paulius Dilkas, Vaishak Belle, and Efthymia Tsamoura. 2024. Mapping the Neuro-Symbolic AI Landscape by Architectures: A Handbook on Augmenting Deep Learning Through Symbolic Reasoning. arXiv:2410.22077
[6] Jerry A. Fodor and Zenon W. Pylyshyn. 1988. Connectionism and Cognitive Architecture: A Critical Analysis. *Cognition* 28, 1–2 (1988), 3–71.
[7] Steven M. Frankland and Joshua D. Greene. 2020. Concepts and Compositionality: In Search of the Brain's Language of Thought. *Annual Review of Psychology* 71 (2020), 273–303.
[8] Artur d'Avila Garcez and Luis C. Lamb. 2023. Neurosymbolic AI: The 3rd Wave. *Artificial Intelligence Review* 56, 11 (2023), 12387–12406.
[9] Drew A. Hudson and Christopher D. Manning. 2018. Compositional Attention Networks for Machine Reasoning. *International Conference on Learning Representations (ICLR)* (2018).
[10] Dieuwke Hupkes, Verna Dankers, Mathijs Mul, and Elia Bruni. 2020. Compositionality Decomposed: How Do Neural Networks Generalise? *Journal of Artificial Intelligence Research* 67 (2020), 757–795.
[11] Justin Johnson, Bharath Hariharan, Laurens Van Der Maaten, Li Fei-Fei, C. Lawrence Zitnick, and Ross Girshick. 2017. CLEVR: A Diagnostic Dataset for Compositional Language and Elementary Visual Reasoning. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2017), 2901–2910.
[12] Brenden M. Lake and Marco Baroni. 2018. Generalization Without Systematicity: On the Compositional Skills of Sequence-to-Sequence Recurrent Networks. *International Conference on Machine Learning (ICML)* (2018), 2873–2882.

[13] Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. 2023. Let's Verify Step by Step. *International Conference on Learning Representations (ICLR)* (2023).

[14] Baihan Lin, Djallel Bouneffouf, and Irina Rish. 2023. A Survey on Compositional Generalization in Applications. arXiv:2302.01067

[15] Jiayuan Mao, Chuang Gan, Pushmeet Kohli, Joshua B. Tenenbaum, and Jiajun Wu. 2019. The Neuro-Symbolic Concept Learner: Interpreting Scenes, Words, and Sentences from Natural Supervision. arXiv:1904.12584

[16] Jiayuan Mao, Xiuming Zhang, Yikai Li, William T. Freeman, Joshua B. Tenenbaum, and Jiajun Wu. 2019. Program-Guided Image Manipulators. *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)* (2019), 4030–4039.

[17] Gary F. Marcus. 2003. *The Algebraic Mind: Integrating Connectionism and Cognitive Science.* MIT Press.

[18] Laura Ruis, Jacob Andreas, Marco Baroni, Diane Bouchacourt, and Brenden M. Lake. 2020. A Benchmark for Systematic Generalization in Grounded Language Understanding. *Advances in Neural Information Processing Systems (NeurIPS)* 33 (2020), 11139–11150.

[19] Atharva Sehgal, Arya Grayeli, Jennifer J. Sun, and Swarat Chaudhuri. 2024. Neurosymbolic Grounding for Compositional World Models. *International Conference on Learning Representations (ICLR)* (2024).

[20] Luciano Serafini and Artur d'Avila Garcez. 2016. Logic Tensor Networks: Deep Learning and Logical Reasoning from Data and Knowledge. arXiv:1606.04422

[21] Amit Sheth, Kaushik Roy, and Manas Gaur. 2023. Neurosymbolic Artificial Intelligence (Why, What, and How). *IEEE Intelligent Systems* 38, 3 (2023), 56–62.

[22] Sever Topan, David Rolnick, and Xujie Si. 2021. Techniques for Symbol Grounding with SATNet. *Advances in Neural Information Processing Systems (NeurIPS)* 34 (2021), 20733–20744.

[23] Elena Umili, Roberto Capobianco, and Giuseppe De Giacomo. 2023. Grounding LTLf Specifications in Image Sequences. *International Conference on Principles of Knowledge Representation and Reasoning (KR)* 19 (2023), 668–678.

[24] Tim Woydt, Moritz Willig, Antonia Wüst, Lukas Helff, Wolfgang Stammer, Constantin A. Rothkopf, and Kristian Kersting. 2025. Fodor and Pylyshyn's Legacy—Still No Human-like Systematic Compositionality in Neural Networks. arXiv:2506.01820

[25] Kexin Yi, Jiajun Wu, Chuang Gan, Antonio Torralba, Pushmeet Kohli, and Joshua Tenenbaum. 2018. Neural-Symbolic VQA: Disentangling Reasoning from Vision and Language Understanding. *Advances in Neural Information Processing Systems (NeurIPS)* 31 (2018).

[26] Eric Zhan, Jennifer J. Sun, Ann Kennedy, Yisong Yue, and Swarat Chaudhuri. 2021. Unsupervised Learning of Neurosymbolic Encoders. arXiv:2107.13132