Neurosymbolic AI on Knowledge Graphs: Characterization, Payoff, and Pitfalls

Lauren Nicole DeLong[1,*], Jacques D. Fleuriot[1]

[1]Artificial Intelligence and its Applications Institute, School of Informatics, University of Edinburgh

1. Introduction

Knowledge graphs (KGs) are becoming an increasingly popular way to represent information [1, 2]. Recently, several studies have developed neurosymbolic (NeSy) approaches for reasoning over KGs. NeSy artificial intelligence (AI) describes the combination of symbolic AI, which often includes logic and rule-based approaches, with neural networks and deep learning [3, 4]. In this presentation, we will review and discuss our recently accepted article within *IEEE Transactions on Neural Networks and Learning Systems.* Our article surveys NeSy approaches for reasoning over KGs, proposes a taxonomy by which such approaches can be categorized, discusses five unique traits which characterize them, and suggests several prospective research directions for which they can be used [5].

2. Background

A *graph* or *network* structure G is composed of a finite set of nodes, V, and a finite set of edges, E, connecting node pairs into *triples.* A *knowledge graph* (KG) uses such a graph to represent relations between real-world entities [2]. *KG completion* (KGC) denotes a series of methods which can be used to refine the graphs and uncover novel information.

Some methods for KGC utilize a set of rules (*e.g.*, from ontologies [6] or rule-mining approaches [7]), that can be used for logical inference. While these approaches tend to be inherently interpretable, they sometimes suffer scalability issues on large KGs [8, 9]. In contrast, methods which generate *KG embeddings* (KGE) typically scale well to large datasets. A KGE is a numerical vector representation of the KG constituents such that *proximity* in the embedding space approximates *similarity* in the original KG [10, 11]. However, KGE methods also have limitations: they tend to lack interpretability, and they require relatively large amounts of labeled data to capture relationships [12, 13]. To bridge this dichotomy, NeSy approaches combine aspects from both types to mitigate their respective weaknesses.

18th International conference on Neural-Symbolic Learning and Reasoning, June 09–12, 2024, Barcelona, Spain
*Corresponding author.
L.N.DELONG@sms.ed.ac.uk (L. N. DeLong); jacques.fleuriot@ed.ac.uk (J. D. Fleuriot)
https://laurendelong21.github.io/ (L. N. DeLong); https://homepages.inf.ed.ac.uk/jdf/ (J. D. Fleuriot)
0000-0001-6618-3370 (L. N. DeLong); 0000-0002-6867-9836 (J. D. Fleuriot)

3. Our Proposal: A New Characterization of NeSy Approaches

We classify the approaches into three major categories: (A) logically-informed embedding approaches, which use symbolic inference and deep learning sequentially, (B) approaches which learn embeddings with logical constraints, and (C) those which learn logical rules for reasoning. From this, we identify five *critical characteristics* of NeSy approaches:

1. Interpretability: Due to the incorporation of logical rules, several NeSy approaches are inherently interpretable.
2. Guided Training: Some NeSy approaches can integrate ontological or expert-defined knowledge into an otherwise data-driven approach [14, 15]. This guides learning toward more domain-congruous patterns.
3. Underrepresented Types: Some datasets may have fewer instances of certain types or classes. Many KGE approaches struggle to capture such underrepresented patterns [12, 13]. Through rule-based inference, some NeSy approaches might address this.
4. Heterogeneous Aggregation: Heterogeneous KGs typically comprise nodes and edges of varying types [16]. The aggregation of features between dissimilar types is a challenge, but some NeSy methods can formally define these relationships through the implementation of logical rules.
5. Long-range Dependencies: Several GNN methods suffer from local receptive fields [17]. In contrast, path-based NeSy methods account for *long-range dependencies* (*i.e.*, relationships between nodes that are several *hops* apart [18, 19]) through chained rules.

Through these, we explore how each category of approaches has unique capabilities and prospective directions for various facets of research.

4. Payoff and Pitfalls: Applications Beyond Benchmarks

Despite the potential these NeSy methods hold, most of the surveyed approaches were only tested upon KGs such as YAGO [20] and DBpedia [21], which comprise *general* knowledge. While this is ideal for benchmarking, few studies demonstrate their methods within a specific domain. In our recent investigations, we used some of our surveyed approaches within biomedical contexts. Specifically, we used two approaches for (1) drug mechanism-of-action discovery, and (2) few-shot prediction of rare side effects. In doing so, we tested how reliably these approaches operate in a real-world domain. However, we also discovered that their distinguishing characteristics, such as Interpretability, sometimes reveal limitations, like reasoning shortcuts [22].

5. Summary and Conclusions

We propose a novel taxonomy by which NeSy methods for reasoning over KGs can be grouped. Using this taxonomy alongside five unique characteristics, we discuss the capabilities and prospective directions toward which these approaches may prove useful. Finally, we present two cases in which we used some of the surveyed approaches on biomedical applications. Through these examples, we highlight the ways in which the interpretability of NeSy AI, typically considered a benefit, can help reveal its own limitations.

References

[1] Z. Chen, Y. Wang, B. Zhao, J. Cheng, X. Zhao, Z. Duan, Knowledge graph completion: A review, IEEE Access 8 (2020) 192435–192456.

[2] S. Ji, S. Pan, E. Cambria, P. Marttinen, S. Y. Philip, A survey on knowledge graphs: Representation, acquisition, and applications, IEEE Transactions on Neural Networks and Learning Systems 33 (2021) 494–514.

[3] A. d. Garcez, L. C. Lamb, Neurosymbolic ai: The 3 rd wave, Artificial Intelligence Review (2023) 1–20.

[4] P. Hitzler, A. Eberhart, M. Ebrahimi, M. K. Sarker, L. Zhou, Neuro-symbolic approaches in Artificial Intelligence, National Science Review 9 (2022) nwac035.

[5] L. N. DeLong, R. F. Mir, J. D. Fleuriot, Neurosymbolic AI for reasoning over knowledge graphs: A survey, IEEE Transactions on Neural Networks and Learning Systems (2024).

[6] D. Dou, H. Wang, H. Liu, Semantic data mining: A survey of ontology-based approaches, in: Proceedings of the 2015 IEEE 9th international conference on semantic computing (IEEE ICSC 2015), IEEE, 2015, pp. 244–251.

[7] L. Galárraga, C. Teflioudi, K. Hose, F. M. Suchanek, Fast rule mining in ontological knowledge bases with amie+, volume 24, Springer, 2015, pp. 707–730.

[8] D. Venugopal, V. G. Gogate, Scaling-up importance sampling for markov logic networks, Advances in Neural Information Processing Systems 27 (2014).

[9] M. Kazmi, P. Schüller, Y. Saygın, Improving scalability of inductive logic programming via pruning and best-effort optimisation, Expert Systems with Applications 87 (2017) 291–303.

[10] M. Zhang, Y. Chen, Link prediction based on graph neural networks, in: Advances in Neural Information Processing Systems, 2018, pp. 5165–5175.

[11] B. Li, D. Pi, Network representation learning: a systematic literature review, Neural Computing and Applications 32 (2020) 16647–16679.

[12] F. Zhao, T. Huang, D. Wang, Graph few-shot learning via restructuring task graph, IEEE Transactions on Neural Networks and Learning Systems (2022).

[13] J. Li, S. Feng, B. Chiu, Few-shot relation extraction with dual graph neural network interaction, IEEE Transactions on Neural Networks and Learning Systems (2023).

[14] M. Alshahrani, M. A. Khan, O. Maddouri, A. R. Kinjo, N. Queralt-Rosinach, R. Hoehndorf, Neuro-symbolic representation learning on biological knowledge graphs, Bioinformatics 33 (2017) 2723–2730.

[15] R. Wickramarachchi, C. Henson, A. Sheth, An evaluation of knowledge graph embeddings for autonomous driving data: Experience and practice, in: Proceedings of the AAAI 2020 Spring Symposium on Combining Machine Learning and Knowledge Engineering in Practice (AAAI-MAKE 2020)., 2020.

[16] Z. Shao, Y. Xu, W. Wei, F. Wang, Z. Zhang, F. Zhu, Heterogeneous graph neural network with multi-view representation learning, IEEE Transactions on Knowledge and Data Engineering 35 (2023) 11476–11488. doi:10.1109/TKDE.2022.3224193.

[17] L. Pasa, N. Navarin, W. Erb, A. Sperduti, Empowering simple graph convolutional networks, IEEE Transactions on Neural Networks and Learning Systems (2023).

[18] Y. Liu, M. Hildebrandt, M. Joblin, M. Ringsquandl, R. Raissouni, V. Tresp, Neural multi-hop reasoning with logical rules on biomedical knowledge graphs, in: European Semantic Web

Conference, Springer, 2021, pp. 375–391.
[19] F. D. Giovanni, T. K. Rusch, M. Bronstein, A. Deac, M. Lackenby, S. Mishra, P. Veličković, How does over-squashing affect the power of GNNs?, Transactions on Machine Learning Research (2024).
[20] F. M. Suchanek, G. Kasneci, G. Weikum, Yago: a core of semantic knowledge, in: Proceedings of the 16th international conference on World Wide Web, 2007, pp. 697–706.
[21] S. Auer, C. Bizer, G. Kobilarov, J. Lehmann, R. Cyganiak, Z. Ives, Dbpedia: A nucleus for a web of open data, in: The Semantic Web, Springer, 2007, pp. 722–735.
[22] E. Marconato, S. Teso, A. Vergari, A. Passerini, Not all neuro-symbolic concepts are created equal: Analysis and mitigation of reasoning shortcuts, Advances in Neural Information Processing Systems 36 (2024).

Acknowledgments

LND is funded by the University of Edinburgh (UoE) Informatics Graduate School through the Global Informatics Scholarship as well as the Artificial Intelligence and its Applications Institute. LND also thanks Enveda Biosciences for collaborating with part of her PhD work.