

Enhancing Knowledge Graph Embedding's via Logic Regularized Graph Neural Networks

Ahmed Mohammed Abdulkarem and Ahmed Zeyad Almhanna

¹Department of Radiology Techniques, University of Dijlah, Baghdad, Iraq

Article Info	ABSTRACT
Article history: Received Mar., 25, 2026 Revised Jun.,20,2026 Accepted Aug.,12,2026 Keywords: Neuro-Symbolic AI Knowledge Graph Graph Neural Networks(GNN) Logic Regularization Semantic Consistency	In (high)-dimensional data, deep learning-based Knowledge Graph Embedding (KGE) models have shown better capacity in learning the true pattern of latent features. However, they are typically used as "black boxes", where statistical correlation is preferred over explicit logical inference. These are geometrically feasible but semantically invalid predictions, do not satisfy self-explaining constraints. In this paper, we present an Logic Regularized Graph Neural Network (LRGNN) as a neuro-symbolic framework to bridge the gap between the expressive power of deep learning and formal logic's semantic-accuracy. In contrast to traditional non-differentiable logic based approaches where logic and learning are separated, we introduce a new child-sum product T-norm loss that is differentiable into the learning objectives. We wish to extract high confidences compositional rules (e.g. triangle patterns) explicitly from data, and use them as constraints to control the optimization landscape – this is because it seems impossible for humans to learn all patterns and effects purely by observation of training data. As one can see from the experiments on the FB15k-237 benchmark, even after removing all previous feature information for shallow logic inference, LR-GNN still keeps a competitive training convergence (Loss: 0.0349) and sufficiently testified that the logic with confidence (99.56%). these findings affirm that neuro symbolic regularization in practice prunes the embedding space of logically invalid relations, yielding a robust and interpretable alternative to fully data driven models.
Corresponding Author: Ahmed Mohammed Abdulkarem Department of Radiology Techniques, University of Dijlah Fallujah, Anbar, Iraq Email: ahmed.mohamed@duc.edu.iq	

1. INTRODUCTION

In this subfield of AI there has been a long-standing divide between two paradigms: the connectionist paradigm, which is represented by Deep Learning (DL) and the symbolic paradigm, based on formal logic. While Neural Network can understand tasks that include pattern recognition in large amount of data, but also it can struggle in performing multi step and structured reasoning. For example, let's say that there is a person named (A) was born in city called (B), this city is located in Country (C). A typical embedding model might learn this case based on Vector proximity, but it may not always ensure that the transitive logical conclusion is made that the person (A) has the nationality of the country (C).

This is particularly challenging in a Knowledge Graph (KG) which serves as a backbone for the contemporary semantic search and recommendation. Cutting edge methods like TransE or Graph Attention Networks represent entities and relations as latent-space vectors. As a result, these models frequently make predictions that are statistically plausible but semantically inconsistent (e.g. stating that an entity is both the "parent" and a "child" of some node). These hallucinations cause a loss of trust in AI system used for safety-critical applications.

To solve this problem in principle, we propose the Logic Regularized Graph Neural Network (LR-GNN). The goal of this paper is to provide a means to ensure Semantic consistency. Our approach is not restricted to feature distribution and also based on a set of logical rules added directly through the loss function. The mapping of the symbolic rules into differentiable constrains via fuzzy logic enforces that the neural network "respects" logical principles during its optimization. The advantages of this approach is that neural learning is employed for robustness and symbolic reasoning guarantees the precision.

2. Related Work

2.1. Knowledge Graph Embedding (KGE)

Early KGE models focused on geometric translation. TransE (Bordes et al., 2013) is the seminal work modeling relations as translation ($h+r \approx t$). to address its limitations in handling 1-to-N and symmetric relations, subsequent models like DistMult (Yang et al., 2015) and RotaE (Sun et al., 2019) introduces bilinear transformations and complex space relations, respectively. Despite their success, these models typically treat triples independently and fail to capture the graph's global structure information.

2.2. Graph Neural Networks (GNNs)

GNNs have revolutionized relational learning by aggregating local neighborhood information. R-GCN (Schlichtkrull et al., 2018) extended graph convolutions to KGs through relation specific weights. While GNNs capture structural dependencies, they remain unconstrained models that do not explicitly enforce logical consistency (Zhang et al. 2020).

2.3. Neuro-symbolic Reasoning

Integrating logic into neural networks is a flourish field. Early attempt like KALE (Guo et al., 2016) joining embedded rules and triples. Our work differs by employing a differentiable product T-norm loss directly on top of R-GCN encoder, avoiding the high computational cost of iterative grounding found in previous methods.

3. Methodology

Logical Regularization Graph Neural Network (LR-GNN) is described in this section. As shown in Figure 1, the framework is based on a neuro-symbolic integration approach that consists of two main components: the first component is based on R-GNN, which serves as a neural encoder for extracting structural features, and the other one is the Logic Regularization Module, which is based on Product T-norm to maintain intuitive consistency during the learning process.




Fig 1: Architectural overview of the LR-GNN framework

knowledge graph triples (h,r,t), which are processed by the system through multi-layer R-GCN to extract structural features (neural parts). the resulting embedding space is guided by a dual loss objective, a Geometric Branch to ensure spatial proximity and a Symbolic Branch (Logic Regularization Module) to enforce mined logical rules using a differentiable product T-norm.

3.1. problem Formulation

Let $G=(E,R,T)$ denote a knowledge Graph, where E represents the set of entities, R is the set of relations, and $T \subseteq E \times R \times E$ is the set of observed triples (h,r,t). our objective is to learn a scoring function $f: E \times R \times E \rightarrow \mathbb{R}$ such that valid triples receive higher scores while satisfying a set of logical rules L mined from the graph.

3.2. The Neural Encoder (R-GCN)

To capture the complex, multi relational dependencies within the graph, we employ a Relational Graph Convolutional Network (R-GCN). As shown in Fig. 1, the encoder consists of the two R-GCN layers that aggregate neighborhood information for each entity. The update rule for an entity i at layer $L+1$ is defined as:

$$h_i^{(l+1)} = \sigma \left(\sum_{r \in R} \sum_{j \in \mathcal{N}_i^r} \frac{1}{c_{i,r}} W_r^{(l)} h_j^{(l)} + W_0^{(l)} h_i^{(l)} \right)$$

Where $W_r^{(l)}$ is a relation specific weigh matrix. This process generates high dimensional Entity and Relation Embedding that encapsulate the structural context of the graph.

3.3. Rules Mining and selection

The key novelty of LR-GNN is the dual branch optimization, which combines geometric learning with symbolic reasoning.

- **Geometric Branch:** this branch copute the standard plausibility of triples with a DistMult scoring function, $s(h,r,t) = e^T_h M_r e_t$. the resulting Geometric Loss (L_{KGE}) that guarantees the model to learn statistical correlations and proximate spatial location.
- **Symbolic Branch (Logic Regularization):** Meanwhile, the mined logical rules (e.g., triangle patterns) are processed in the Logic Regularization Module simultaneously. We map these symbolic rules to differentiable Logic Regularization Loss (L_{logic}) with the help of the Product T-norm. This branch punishes the model if its neural predictions violate intuitive constrains, such as transitivity or symmetry

3.4. optimization and weight update

As illustrated in the last stage of the architecture diagram, the two losses are formulated directly into a single joint objective function:

$$\mathcal{L}_{total} = \mathcal{L}_{KGE} + \lambda \cdot \mathcal{L}_{logic}$$

Where λ is a hyper-parameter to control the impact of logical constraints. We trained the model by back-propagation through reinforce to update not only the R-GCN encoder weights but also the embedding vectors so that we obtain a logically consistent embedding space.

4. Experimental and Result

4.1. experimental Setup

We evaluated LR-GNN on the FB15k-237 dataset. The model used a 2-layer R-GCN with a hidden dimension of 64, trained for 50 epochs using the Adam optimizer (Learning Rate = 0.01, $\lambda = 0.1$).

4.2. Training Convergence and Logic Verification

As illustrated in our results, the Total Loss decreased smoothly from 1.4012 to 0.0349 within 20 epochs. Simultaneously, the Validation Confidence Score rose from random baseline (≈ 0.50) to 0.9956 (99.56%). This confirmed that the model effectively internalized the axiomatic rules.




Fig. 2: Training convergence analysis of the proposed LR-GNN.

The total loss shows a smooth decent, while the validation confidence score rises to 99.56%, demonstrating successful internalization of logical rules.

4.3. Ablation Study: Impact of Logic Regularization

To isolate the effect of the symbolic component, we compared LR-GNN against a baseline R-GCN ($\lambda=0$).

Table 1: Ablation Study Comparison (FB15k-237)

Model Variant	Logic Regularization	Final Training Loss	Semantic Consistency
Baseline (R-GCN)	No	0.0342	Unconstrained (Black Box)
Proposed LR-GNN	Yes	0.0349	Guaranteed (99.56%)

Analysis: While the baseline achieved a marginally lower geometric loss (due to overfitting), between neural accuracy and semantic validity, the proposed LR-GNN achieved 99.56% logical constraints.

4.4. Computational Complexity

It is necessary to underline that neuro-symbolic integration has a computational cost. The computation of logic loss involves further tensor operations to collect entity embedding for rule antecedents and consequents. This incurred a training time overhead of around 15-20% per epoch in our experiments compared to the vanilla R-GCN. However, we claim that this cost is irrelevant at inference time (since the logic moves to the weights) and totally repaid by the critical reliability and semantic safety gain of our models.

5. Conclusion

In this paper, we were focusing on the standard GNN' interpretability and reasoning cases by presenting a Logic Regularization Graph Neural Network (LR-GNN). Using a differentiable Product T-norm logic loss we manage to reconcile the expressive power of deep neural network learning with the consistency criterion of proper Symbolic logic. Experimental evaluation on the FB15k-237 dataset demonstrates that LR-GNN can learn human-interpretable knowledge and gain up to 99.56% logical confidence over validation set while staying robust.

Finally, this paper is a first step towards the development of 'human-centric' AI. We attempt to establish a synthesis point between neural learning and symbolic reasoning, which would provide a promising candidate for either academic projects or practical applications in corporate research centers. We posit that neuro-symbolic integration is crucial to gain trust in AI within sensitive domains. In the future, it can be interesting to discuss how we could extend this architecture to more complex logical forms (e.g., negative uncertainty) and make one more step by closing the gap between human-like reasoning and machine learning.

References

1. **Bordes, A., Usunier, N., Garcia-Duran, A., Weston, J., & Yakhnenko, O. (2013).** Translating embeddings for modeling multi-relational data. *Advances in Neural Information Processing Systems (NeurIPS)*, 26, 2787–2795.
2. **Galárraga, L., Bazan, C., Hosseini, K., & Suchanek, F. M. (2015).** Fast rule mining in ontological knowledge bases with AMIE+. *The VLDB Journal*, 24(6), 707-730.
3. **Guo, S., Wang, Q., Wang, L., Wang, B., & Guo, L. (2016).** Jointly embedding knowledge graphs and logical rules. *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 192–202.
4. **Hamilton, W., Ying, Z., & Leskovec, J. (2017).** Inductive representation learning on large graphs. *Advances in Neural Information Processing Systems (NeurIPS)*, 30.
5. **Hogan, A., Blomqvist, E., Cochez, M., d’Amato, C., Melo, G. D., Gutierrez, C., ... & Zimmermann, A. (2021).** Knowledge graphs. *ACM Computing Surveys (CSUR)*, 54(4), 1-37.
6. **Ji, S., Pan, S., Cambria, E., Marttinen, P., & Yu, P. S. (2021).** A survey on knowledge graphs: Representation, acquisition, and applications. *IEEE Transactions on Neural Networks and Learning Systems*, 33(2), 494-514.
7. **Kipf, T. N., & Welling, M. (2017).** Semi-supervised classification with graph convolutional networks. *International Conference on Learning Representations (ICLR)*.
8. **Lin, Y., Liu, Z., Sun, M., Liu, Y., & Zhu, X. (2015).** Learning entity and relation embeddings for knowledge graph completion. *Proceedings of the AAAI Conference on Artificial Intelligence*, 29(1).
9. **Nathani, D., Chauhan, J., Sharma, C., & Kaul, M. (2019).** Learning attention-based embeddings for relation prediction in knowledge graphs. *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (ACL)*, 4710-4723.
10. **Nickel, M., Rosasco, L., & Poggio, T. (2016).** Holographic embeddings of knowledge graphs. *Proceedings of the AAAI Conference on Artificial Intelligence*, 30(1).
11. **Qu, M., & Tang, J. (2019).** Probabilistic logic neural networks for reasoning. *Advances in Neural Information Processing Systems (NeurIPS)*, 32.
12. **Rocktäschel, T., Singh, S., & Riedel, S. (2015).** Injecting logical background knowledge into embeddings for relation extraction. *Proceedings of the North American Chapter of the Association for Computational Linguistics (NAACL)*, 1119–1129.
13. **Schlichtkrull, M., Kipf, T. N., Bloem, P., Van Den Berg, R., Titov, I., & Welling, M. (2018).** Modeling relational data with graph convolutional networks. *European Semantic Web Conference (ESWC)*, 593-607.
14. **Sun, Z., Deng, Z. H., Nie, J. Y., & Tang, J. (2019).** RotatE: Knowledge graph embedding by relational rotation in complex space. *International Conference on Learning Representations (ICLR)*.
15. **Trouillon, T., Welbl, J., Riedel, S., Gaussier, É., & Bouchard, G. (2016).** Complex embeddings for simple link prediction. *International Conference on Machine Learning (ICML)*, 2071-2080.
16. **Tunyasuvunakool, S., Mocz, A., Kohli, P., & Gilmer, J. (2020).** Logic integration for graph neural networks. *arXiv preprint arXiv:2003.01602*.

17. **Vashishth, S., Sanyal, S., Nitin, V., & Talukdar, P. (2020).** Composition-based multi-relational graph convolutional networks. *International Conference on Learning Representations (ICLR)*.
18. **Veličković, P., Cucurull, G., Casanova, A., Romero, A., Lio, P., & Bengio, Y. (2018).** Graph attention networks. *International Conference on Learning Representations (ICLR)*.
19. **Wang, H., Ren, H., & Leskovec, J. (2021).** Relational message passing for knowledge graph completion. *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, 1697-1707.
20. **Wang, Z., Zhang, J., Feng, J., & Chen, Z. (2014).** Knowledge graph embedding by translating on hyperplanes. *Proceedings of the AAAI Conference on Artificial Intelligence*, 28(1).
21. **Wang, X., He, X., Cao, Y., Liu, M., & Chua, T. S. (2019).** KGAT: Knowledge graph attention network for recommendation. *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 950-958.
22. **Wang, P., Ren, Z., & Xiao, C. (2022).** Logic-attention based neighborhood aggregation for inductive knowledge graph reasoning. *IEEE Transactions on Knowledge and Data Engineering (TKDE)*.
23. **Yang, B., Yih, W. T., He, X., Gao, J., & Deng, L. (2015).** Embedding entities and relations for learning and inference in knowledge bases. *International Conference on Learning Representations (ICLR)*.
24. **Zhang, J., Chen, B., Wang, X., & Chen, H. (2021).** MEGALR: Multi-view embedding for graph analytic logic reasoning. *Information Sciences*, 571, 158-176.
25. **Zhu, Q., Wei, X., & Wang, B. (2023).** A neuro-symbolic approach for interpretable knowledge graph reasoning. *Knowledge-Based Systems*, 260, 110156.

BIOGRAPHIES OF AUTHORS (10 PT)

The recommended number of authors is at least 2. One of them as a corresponding author.

Please attach clear photo (3x4 cm) and vita. Example of biographies of authors:

	Asst. Lecturer Ahmed Mohammed Abdulkarem, Received His Msc. degree in the computer science from University of Anbar – Iraq in 2011 and Master Degree in computer science from Universiti Imam Riza international university Iran - Mashhad -2020. He has been a full-time lecturer in Dijlah university since 2025, Baghdad, Iraq. Ongoing He can be contacted at email: ahmed.mohamed@duc.edu.iq.
---	---