



HAL
open science

Knowledge Enhanced Graph Neural Networks

Luisa Werner, Nabil Layaïda, Pierre Genevès, Sarah Chlyah

► **To cite this version:**

Luisa Werner, Nabil Layaïda, Pierre Genevès, Sarah Chlyah. Knowledge Enhanced Graph Neural Networks. 2023. hal-04041691v1

HAL Id: hal-04041691

<https://inria.hal.science/hal-04041691v1>

Preprint submitted on 27 Mar 2023 (v1), last revised 30 Aug 2023 (v3)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Knowledge Enhanced Graph Neural Networks

Luisa Werner^{1,2}, Nabil Layaïda¹, Pierre Genevès¹ and Sarah Chlyah¹

¹*Institut National de Recherche en Sciences et Technologies du Numérique (INRIA), 655 Av. de l'Europe, 38330 Montbonnot-Saint-Martin, France*

²*Université Grenoble Alpes (UGA), 621 Av. Centrale, 38400 Saint-Martin-d'Hères, France*

Abstract

Graph data is omnipresent and has a large variety of applications such as natural science, social networks or semantic web. Though rich in information, graphs are often noisy and incomplete. Therefore, graph completion tasks such as node classification or link prediction have gained attention. On the one hand, neural methods such as graph neural networks have proven to be robust tools for learning rich representations of noisy graphs. On the other hand, symbolic methods enable exact reasoning on graphs. We propose KeGNN, a neuro-symbolic framework for learning on graph data that combines both paradigms and allows for the integration of prior knowledge into a graph neural network model. In essence, KeGNN consists of a graph neural network as a base on which knowledge enhancement layers are stacked with the objective of refining predictions with respect to prior knowledge. We instantiate KeGNN in conjunction with two standard graph neural networks: Graph Convolutional Networks and Graph Attention Networks, and evaluate KeGNN on multiple benchmark datasets for node classification.

Keywords

neuro-symbolic integration, graph neural networks, relational learning, knowledge graphs, fuzzy logic

1. Introduction

Graphs are ubiquitous across diverse real-world applications such as e-commerce [1], natural science [2] or social networks [3]. Graphs connect nodes by edges and allow to enrich them with features. This makes them a versatile and powerful data structure that encodes relational information. As graphs are often derived from noisy data, incompleteness and errors are common issues. Consequently, graph completion tasks such as node classification or link prediction have become increasingly important. These tasks are approached from different directions. In the field of deep learning, research on graph neural networks (GNNs) has gained momentum. Numerous models have been proposed for various graph topologies and applications [4] [5] [6] [7]. The key strength of GNNs is to find meaningful representations of noisy data, that can be used for various prediction tasks [8]. Despite this advantage, as a subcategory of deep learning methods, GNNs are criticized for their limited interpretability and large data consumption [9]. Alongside, the research field of symbolic AI addresses the above-mentioned tasks. In symbolic AI, solutions are found by performing logic-like reasoning steps that are exact, interpretable and data-

NeSy'23: 17th International Workshop on Neural-symbolic Learning and Reasoning, Siena, Italy 3rd-5th July 2023

✉ luisa.werner@inria.fr (L. Werner); nabil.layaïda@inria.fr (N. Layaïda); pierre.geneves@inria.fr (P. Genevès); sarah.chlyah@inria.fr (S. Chlyah)

🆔 0000-0002-1431-6490 (L. Werner); 0000-0001-8472-9365 (N. Layaïda); 0000-0001-7676-2755 (P. Genevès); 0009-0004-1769-5109 (S. Chlyah)



© 2023 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

 CEUR Workshop Proceedings (CEUR-WS.org)

efficient [10] [11] [12]. For large graphs, however, symbolic methods are often computationally expensive or even infeasible. Since techniques from deep learning and from symbolic AI have complementary pros and cons, the field of neuro-symbolic integration aims to combine both paradigms. Neuro-symbolic AI not only paves the way towards the application of AI to learning with limited data, but also allows for jointly using symbolic information (in the form of logical rules) and sub-symbolic information (in the form of real-valued data). This leads on the one hand to more versatile models. On the other hand, it helps to overcome the blackbox nature of deep learning and to improve interpretability through symbolic representations. [13] [14] [9].

In this paper, we present the neuro-symbolic approach **Knowledge enhanced Graph Neural Networks (KeGNN)** to conduct node classification given real-valued graph data and a set of prior knowledge. In KeGNN, differentiable knowledge enhancement layers are stacked on top of a GNN and adjust its predictions in order to increase the satisfaction of a set of prior knowledge. Additionally to the parameters of the GNN, the knowledge enhancement layers contain learnable clause weights that reflect the impact of the prior knowledge. Both components form an end-to-end differentiable model.

In the multifaceted neuro-symbolic field, KeGNN can be placed in the category of knowledge-guided learning [15], where the focus lies on learning in the presence of additional supervision induced by prior knowledge. Among these approaches, KeGNN belongs to model-based approaches where prior knowledge is an integral part of the model in the form of knowledge enhancement layers. KeGNN can be seen as an extension to knowledge enhanced neural networks (KENN) [16], which stack knowledge enhancement layers onto a multi-layer perceptron (MLP). However, an MLP is not powerful enough to incorporate graph structure into the representations. Thus, relational information can only be introduced by binary predicates in the symbolic part. In contrast, KeGNN is based on GNNs that can handle graph structure, which makes both the neural and the symbolic component sufficiently powerful to exploit graph structure. In this work, we instantiate KeGNN in conjunction with two popular GNNs: Graph Attention Networks [17] and Graph Convolutional Networks [18]. We apply KeGNN to the benchmark datasets for node classification Cora, Citeseer, PubMed [19] and Flickr [20].

2. Method: Knowledge Enhanced Graph Neural Networks

KeGNN is a neuro-symbolic approach that can be applied to node classification tasks with the capacity of handling graph structure at the base neural network level. The model takes two types of input: (1) real-valued graph data and (2) prior knowledge expressed in first-order logic.

2.1. Graph-structured Data

A Graph $\mathbf{G} = (\mathbf{N}, \mathbf{E})$ consists of set of n nodes \mathbf{N} and edges \mathbf{E} connecting two nodes $v_i, v_j \in \mathbf{N}$ with an edge $(v_i, v_j) \in \mathbf{E}$. The adjacency matrix \mathbf{A} describes the edges. The neighborhood $\mathcal{N}(v_i)$ describes the set of first-order neighbors of v_i . For an *attributed* and *labelled* graph, nodes are enriched with node features and labels. Each node v_i is attributed with a feature vector $\mathbf{x}_i \in \mathbb{R}^d$ with feature dimension d and ground truth label vector $\mathbf{y}_i \in \mathbb{R}^m$ for m classes. Graphs that consist of a single node and edge type are called *homogeneous*. In contrast, nodes and edges may have different types for *heterogeneous* graphs [5].

Example 2.1 (Citation Graph). A Citation Graph \mathbf{G}_{Cit} consists of documents that can cite each other. The documents are represented by nodes \mathbf{N}_{Cit} and citations by edges \mathbf{E}_{Cit} . Documents can be attributed with features \mathbf{X}_{Cit} that describe their content in matrix format as Word2Vec [21] vectors. Each document is labelled with a topic category that is described in \mathbf{Y}_{Cit} . Since all nodes have the same type (documents) and the graph contains only one edge type (citation), \mathbf{G}_{Cit} is homogeneous.

2.2. Prior Knowledge

The prior knowledge \mathcal{K} provided for KeGNN can be described as a set of logical formulae expressed in the function-free logical language \mathcal{L} that contains a set of constants \mathcal{C} , variables \mathcal{X} and predicates \mathcal{P} . Unary predicates express properties of variables, while binary predicates express relations. \mathcal{L} supports negation (\neg) and disjunction (\vee). A clause in the set of prior knowledge $c_i \in \mathcal{K} = \{c_0, \dots, c_{|\mathcal{K}|}\}$ can be formulated as a disjunction of literals $\bigvee_{j=1}^q l_j$ with q literals in a clause. Since the prior knowledge is general, all formulae are assumed to be universally quantified and free of constants. Clauses can be *grounded* by assigning free variables to constants. The set of all grounded clauses is $\mathcal{G}(\mathcal{K}, \mathcal{C})$.

Example 2.2 (Prior Knowledge on Citations and Topics). Let the unary predicate $\text{AI}(x)$ describe that any scientific paper x belongs to the document class AI (from the artificial intelligence domain) and let the binary predicate $\text{Cite}(x, y)$ denote a citation between two papers x and y . The clause $\forall x, y : \neg \text{AI}(x) \vee \neg \text{Cite}(x, y) \vee \text{AI}(y)$ describes that if paper x is about AI and cites paper y , paper y is also about AI.

2.3. Fuzzy Semantics

Let us consider an attributed, labelled and homogeneous graph \mathbf{G} and a set of prior knowledge \mathcal{K} as previously defined. While \mathcal{L} describes logical formulae, the neural component in KeGNN relies on continuous and differentiable representations. To interpret boolean logic in the real-valued domain, fuzzy logic [22] is used that maps binary truth values into the continuous interval $[0, 1]$. A constant is interpreted as a real-valued node feature vector $\mathbf{x} \in \mathbb{R}^d$. A unary predicate P_U is interpreted as a differentiable function $f_{P_U} : \mathbb{R}^d \mapsto [0, 1]$ that takes the vector \mathbf{x} of a constant as input and returns the truth values t of $P_U(\mathbf{x})$. T-conorm functions $\perp : [0, 1] \times [0, 1] \mapsto [0, 1]$ [23] map real-valued truth values of two literals to the truth value of their disjunction. The Gödel t-conorm function for two truth values t_1, t_2 is defined as $\perp(t_1, t_2) \mapsto \max(t_1, t_2)$. To obtain the truth value of a clause $c = l_1 \vee \dots \vee l_q$, \perp is extended to a vector t of q truth values: $\perp(t_1, t_2, \dots, t_q) = \perp(t_1, \perp(t_2 \dots \perp(t_{q-1}, t_q)))$. Fuzzy negation is defined as $t \mapsto 1 - t$ [22].

2.4. Model Architecture

The way KeGNN computes the final predictions can be divided in two stages. First, a GNN predicts the node classes given the node features and the edges. Subsequently, the knowledge enhancement layers use the predictions as truth values for the grounded unary predicates and update them with respect to the knowledge. An overview of KeGNN is given in Fig. 1.

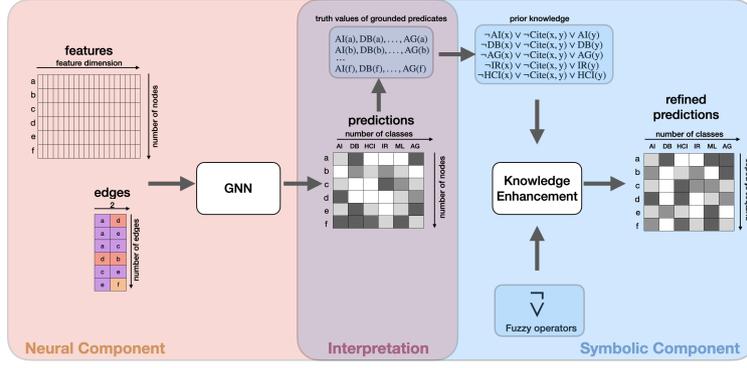


Figure (1): Overview of KeGNN.

2.4.1. Neural Component

The role of the neural component is to exploit feature information in the graph structure. The key strength of a GNN is to enrich node representations with graph structure by nesting multiple message passing layers [8]. Per layer, the representations of neighboring nodes are aggregated and combined to obtain updated representations. A message passing layer k is written as:

$$v_{(i,k+1)} = \text{combine}(v_{(i,k)}, \text{aggregate}(\{v_{(j,k)} | v_j \in \mathcal{N}(v_i)\})) \quad (1)$$

The layers contain learnable parameters that can be optimized with backpropagation. In KeGNN, the GNN implements the function f_{PV} (see section 2.3) that predicts truth values for grounded unary predicates. In this work, we consider two well-known GNNs as components for KeGNN: Graph Convolutional Networks (GCN) [18] and Graph Attention Networks (GAT) [17]. While GCN considers the graph structure as given, GAT allows for assessing the importance of the neighbors with attention weights α_{ij} between node v_i and node v_j . In case of multi-head attention, the attention weights are calculated multiple times and concatenated [24].

2.4.2. Symbolic Component

To refine the predictions of the GNN, one or more knowledge enhancement layers are stacked on top of it with the goal to increase the satisfaction of the prior knowledge. The layers consist of *clause enhancers* for each clause in the prior knowledge that implement a so-called differentiable *boost function* [15]. It returns adjustments for the GNN's predictions with respect to a clause.

Let us consider the grounded clause $c = l_1 \vee \dots \vee l_q$ and the vector $t = (t_1, \dots, t_q)$ containing the truth values of the grounded literals. A boost functions $\phi : \mathbb{R}^q \mapsto \mathbb{R}^q$ provides changes to the q initial truth values t_i predicted by the GNN. The changes are computed with the objective to increase the value of the t-conorm $\perp(t)$, which measures the satisfaction of the clause. The following function [15] is used as a differentiable approximate for the Gödel t-conorm boost function:

$$\phi_{w_c}^s(\mathbf{z})_i = w_c \cdot \text{softmax}(\mathbf{z})_i = w_c \cdot \frac{e^{z_i}}{\sum_{j=1}^n e^{z_j}} \quad (2)$$

To ensure that the refined predictions remain in the interval $[0, 1]$, the boost function is applied to the preactivations $\mathbf{z} = \sigma^{-1}(t)$ of the GNN. After the enhancement, the activation function σ is applied to obtain the final predictions.

A *clause weight* w_c is assigned to each clause in the prior knowledge. $\phi_{w_c}^s$ provides changes to the predictions that are proportional to w_c . Hence, w_c determines the magnitude of the change and reflects the strength of the clause. It can be set to zero, which makes KeGNN robust to wrong knowledge. It can also be fixed to enforce the impact of a clause. Various groundings of a clause are stored in a matrix where columns represent predicates and rows represent constants. Once instantiated, a clause enhancer operates on several groundings of a clause.

For graph data, the groundings for unary predicates cannot be assumed to be independent since nodes are linked by an edge. While the groundings of unary predicates can be represented as a matrix that contains the objects as rows and grounded predicates as columns, the keys of the binary predicates are two-dimensional. Consequently, binary predicates are represented as a matrix \mathbf{M} that has as many rows as pairs of objects. To enhance clauses that contain both binary and unary predicates, all grounded predicates are joined into \mathbf{M} on which the clause enhancer can operate. Hence, unary predicates $P(x)$ are *binarized* to $P^X(x, y)$ and $P^Y(x, y)$. The boost function (Eq. 2) is then applied to \mathbf{M} in order to calculate the adjustments $\delta\mathbf{M}$. Given the adjustments $\delta\mathbf{M}$, a group-by layer collects the changes from the binarized predicates that refer to the same grounded literal. Finally, the changes by all clause enhancers are summed, added to \mathbf{z} and put to the activation function σ to obtain the updated predictions \mathbf{y}' :

$$\mathbf{y}' = \sigma\left(\mathbf{z} + \sum_{c \in \mathcal{G}(\mathcal{K}, \mathcal{C})} \phi_c(\mathbf{z})\right) \quad (3)$$

Example 2.3 (Enhancement of unary and binary predicates). Let us consider the clause $c: \neg\text{AI}(x) \vee \neg\text{Cite}(x, y) \vee \text{AI}(y)$ and the grounding $c[x/i, y/j]$. The unary predicate $\text{AI}(x)$ is binarized to $\text{AI}^X(x, y)$ and $\text{AI}^Y(x, y)$ so that the preactivations $\mathbf{z}_{\text{AI}(i)}$ and $\mathbf{z}_{\text{AI}(j)}$ can be put in the same row but in separate columns. In this case, the adjustment for $\mathbf{z}_{\text{AI}(i)}$ of the clause enhancer is given by the following term, where $e^{\mathbf{z}_{\text{Cite}(i,j)}}$ represents the preactivation of the grounded binary predicate $\text{Cite}(x, y)$:

$$\phi_s^{w_c}(\mathbf{z}_{\text{AI}(i)}) = w_c \cdot \frac{e^{\mathbf{z}_{\text{AI}(i)}}}{e^{\mathbf{z}_{\text{AI}(i)}} + e^{\mathbf{z}_{\text{AI}(j)}} + e^{\mathbf{z}_{\text{Cite}(i,j)}}}$$

3. Experimental Evaluation

To evaluate the performance of KeGNN, we apply it to the datasets Citeseer, Cora, PubMed and Flickr that are common benchmarks for node classification. In the following, KeGNN is called KeGCN and KeGAT when instantiated to a GCN or a GAT, respectively. As additional baseline, we consider KeMLP, where knowledge enhancement layers are nested on an MLP as proposed in [16]. Further, the standalone neural models MLP, GCN and GAT are used as baselines. While Citeseer, Cora and PubMed are citation graphs that encode citations between scientific papers (as in Example 2.2), Flickr [20] contains images and shared properties between them. All datasets can be modelled as homogeneous, labelled, attributed, unweighted graphs as

	MLP	KeMLP	GCN	KeGCN	GAT	KeGAT
Cora	0.7098 (0.0080)	0.8072 (0.0193)	0.8538 (0.0057)	0.8587 (0.0057)	0.8517 (0.0068)	0.8498 (0.0066)
CiteSeer	0.7278 (0.0081)	0.7529 (0.0067)	0.748 (0.0102)	0.7506 (0.0096)	0.7718 (0.0072)	0.7734 (0.0073)
PubMed	0.8844 (0.0057)	0.8931 (0.0048)	0.8855 (0.0062)	0.8840 (0.0087)	0.8769 (0.0040)	0.8686 (0.0081)
Flickr	0.4656 (0.0018)	0.4659 (0.0012)	0.5007 (0.0063)	0.4974 (0.0180)	0.4970 (0.0124)	0.4920 (0.0189)

Table (1)

Average accuracy on the test set over 50 runs (10 for Flickr). The standard deviations are reported in brackets.

defined in Section 2.1. The set of prior logic for the knowledge enhancement layers is given explicitly. If two nodes are linked and if the first node belongs to a specific class, the second node will also belong to that class. In the context of citation graphs, this implies that two documents that cite each other are referring to the same topic, while for Flickr, linked images share the same properties. According to this pattern, for all datasets a clause is instantiated for each class C_i : $\forall xy : \neg C_i(x) \vee \neg \text{Link}(x, y) \vee C_i(y)$. More details on the experiments are given in appendix A. The implementation of KeGNN and the experiments are publicly available¹.

3.1. Results

To compare the performance of all models, we examine the average test accuracy over 50 runs (10 for Flickr) for the knowledge enhanced models KeMLP, KeGCN, KeGAT and the standalone base models MLP, GCN, GAT on the named datasets. The results are given in Tab. 1 and visualized in Fig. 4 (see appendix A.3). For Cora and Citeseer, KeMLP leads to a significant improvement over MLP (p-value of one-sided t-test $\ll 0.05$). In contrast, no significant advantage of KeGCN or KeGAT in comparison to the standalone base model is observed. Nevertheless, all GNN-based models are significantly superior to KeMLP for Cora. This includes not only KeGCN and KeGAT, but also the GNN baselines. For Citeseer, KeGAT and GAT both outperform KeMLP. In the case of PubMed, only a significant improvement of KeMLP over MLP can be observed, while the GNN-based models and their enhanced versions do not provide any positive effect. For Flickr, no significant improvement between the base model and the respective knowledge enhanced model can be observed. Nevertheless, all GNN-based models outperform KeMLP, reporting significantly higher mean test accuracies for KeGAT, GAT, GCN and KeGCN.

It turns out that the performance gap between MLP and KeMLP is larger than for KeGNN in comparison to the standalone GNN. This can be attributed to the fact that MLP is a simple model that misses information on the graph structure and thus benefits from the prior knowledge provided by the knowledge enhancer. On the contrary, standalone GNNs can process graph structure by using message passing techniques to transmit learned node representations between neighbors. The prior knowledge introduced in the knowledge enhancer is simple: it encodes that two neighbors are likely to be of the same class. An explanation of the small performance gap is that the GNNs might have learned to propagate the relevant information across neighbors.

¹<https://gitlab.inria.fr/tyrex/kegnn>

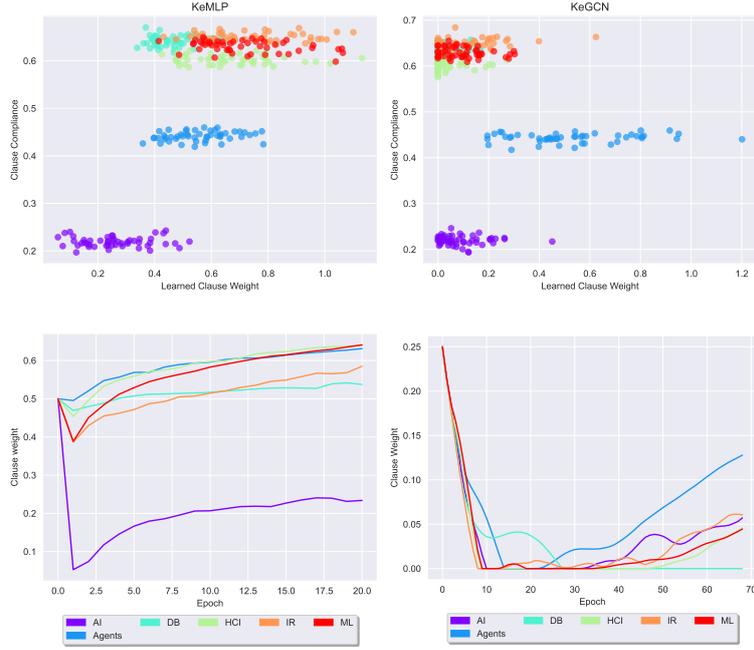


Figure (2): Comparison of KeMLP and KeGCN. Above: Clause weights of last epoch during training vs. clause compliance. Below: Clause weight evolution during training.

In other words, the introduced knowledge is redundant.

In addition, it is found that a standalone GNN can yield equal or better performance compared to a KeMLP in some cases. As GNNs are more efficient in terms of runtimes in compared to the more complex knowledge enhanced models, equal performance could be achieved faster with a pure neural model (see runtime comparison in Tab. 3 in A.3).

3.1.1. Clause Weight Analysis

The clause weights obtained during training provide insights on the extent of adjustment applied by a clause. As a reference, the *clause compliance* (see Appendix B) [15] measures the satisfaction of a clause in the training set before optimization. Fig. 2 displays the learned clause weights for KeGCN and KeMLP versus the clause compliance (above) and the clause weight evolution during training (below). For KeMLP, a positive correlation between the learned clause weights and the clause compliance on the training set is observed. This indicates that satisfied clauses in the training set have a higher impact on the adjustments. In addition, the clause weights corresponding to clauses with low compliance values make smaller modifications to the initial predictions. Consequently, clauses that are rarely satisfied learn lower clause weights during the training process. In the case of KeGCN, the clause weights are predominantly set to low values close to zero. This is in accordance with the absence of a significant performance gap between GCN and KeGCN. Since the GCN in itself already leads to valid classifications, smaller adjustments are required by the clause enhancers.

4. Limitations

In this work, we focused on homogeneous graphs. However, in reality graphs are often heterogeneous with multiple node and edge types [4]. Adaptations are necessary on both the neural and the symbolic side to apply KeGNN to heterogeneous graphs. The restriction to homogeneous graphs also limits the scope for formulating complex prior knowledge. Eventually, the datasets used in this work and the set of prior knowledge are too simple for KeGNN to exploit its potential and lead to significant improvements. The extension of KeGNN to more generic data structures such as incomplete and heterogeneous knowledge graphs in conjunction with more complex prior knowledge is an ongoing work.

A further limitation of KeGNN is scalability. With increasing depth of stacked knowledge enhancement layers, the affected node neighborhood grows exponentially, which can lead to significant memory overhead. This problem is referred as neighborhood explosion [7] and is particularly problematic in the context of training on memory-constrained GPUs. This affects both the GNN and the knowledge enhancement layers that encode binary knowledge. Methods from scalable graph learning [25] [20] [26] represent potential solutions for the neighborhood explosion problem in KeGNN.

Furthermore, limitations appear in the context of link prediction with KeGNN. In link prediction tasks, the knowledge enhancement layers must obtain truth values from the neural component not only for unary predicates but also for binary predicates. At present, KeGNN can handle clauses containing binary predicates, but their truth values are initialized with artificial predictions, where a high value encodes the presence of an edge. This also limits KeGNN to datasets where the graph structure is assumed to be complete and true. For link prediction, a neural component is required that predicts truth values for binary predicates.

5. Conclusion and Future Work

In this work, we introduced KeGNN, a neuro-symbolic model that integrates GNNs with symbolic knowledge enhancement layers into an end-to-end differentiable model. This allows the use of prior knowledge to improve node classification by exploiting the expressive representations of a GNN. Experimental studies show that the inclusion of prior knowledge has the potential to improve simple neural models (as observed in the case of MLP). However, the knowledge enhancement of GNNs is harder to achieve on the underlying and limited benchmarks. Nevertheless, this work is a step towards a holistic neuro-symbolic method on incomplete and noisy semantic data, such as knowledge graphs. Semantically richer datasets in conjunction with more expressive prior knowledge should be investigated to fully explore the promising benefits of KeGNN. A future work in this direction is the application of KeGNN to the completion of noisy, feature-enriched knowledge graphs such as Wikialumni [27] based on YAGO4 [28] and Wikidata [29] with more complex prior knowledge in form of ontologies [30].

Acknowledgments

This work has been partially supported by MIAI @ Grenoble Alpes (ANR-19-P3IA-0003).

References

- [1] W. Liu, Y. Zhang, J. Wang, Y. He, J. Caverlee, P. Chan, D. Yeung, P.-A. Heng, Item relationship graph neural networks for e-commerce, *IEEE Transactions on Neural Networks and Learning Systems* PP (2021) 1–15. doi:10.1109/TNNLS.2021.3060872.
- [2] A. Sanchez-Gonzalez, N. Heess, J. T. Springenberg, J. Merel, M. Riedmiller, R. Hadsell, P. Battaglia, Graph networks as learnable physics engines for inference and control, in: J. Dy, A. Krause (Eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, PMLR, 2018, pp. 4470–4479. URL: <https://proceedings.mlr.press/v80/sanchez-gonzalez18a.html>.
- [3] Y. Wu, D. Lian, Y. Xu, L. Wu, E. Chen, Graph convolutional networks with markov random field reasoning for social spammer detection, *Proceedings of the AAAI Conference on Artificial Intelligence* 34 (2020) 1054–1061. URL: <https://ojs.aaai.org/index.php/AAAI/article/view/5455>. doi:10.1609/aaai.v34i01.5455.
- [4] X. Yang, M. Yan, S. Pan, X. Ye, D. Fan, Simple and efficient heterogeneous graph neural network, 2022. URL: <https://arxiv.org/abs/2207.02547>. doi:10.48550/ARXIV.2207.02547.
- [5] Y. Ma, J. Tang, *Deep Learning on Graphs*, Cambridge University Press, 2021.
- [6] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, P. S. Yu, A comprehensive survey on graph neural networks, *IEEE Transactions on Neural Networks and Learning Systems* 32 (2021) 4–24. URL: <https://doi.org/10.1109/TNNLS.2020.2978386>. doi:10.1109/tnnls.2020.2978386.
- [7] K. Duan, Z. Liu, P. Wang, W. Zheng, K. Zhou, T. Chen, X. Hu, Z. Wang, A comprehensive study on large-scale graph training: Benchmarking and rethinking, 2022. URL: <https://arxiv.org/abs/2210.07494>. doi:10.48550/ARXIV.2210.07494.
- [8] L. Wu, P. Cui, J. Pei, L. Zhao, *Graph Neural Networks: Foundations, Frontiers, and Applications*, Springer Singapore, Singapore, 2022.
- [9] Z. Susskind, B. Arden, L. K. John, P. Stockton, E. B. John, Neuro-symbolic ai: An emerging class of ai workloads and their characterization, 2021. URL: <https://arxiv.org/abs/2109.06133>. doi:10.48550/ARXIV.2109.06133.
- [10] H. Wu, Z. Wang, K. Wang, Y.-D. Shen, Learning Typed Rules over Knowledge Graphs, in: *Proceedings of the 19th International Conference on Principles of Knowledge Representation and Reasoning*, 2022, pp. 494–503. URL: <https://doi.org/10.24963/kr.2022/51>. doi:10.24963/kr.2022/51.
- [11] P. Körner, M. Leuschel, J. Barbosa, V. S. Costa, V. Dahl, M. V. Hermenegildo, J. F. Morales, J. Wielemaker, D. Diaz, S. Abreu, G. Ciatto, Fifty years of prolog and beyond, 2022. URL: <https://arxiv.org/abs/2201.10816>. doi:10.48550/ARXIV.2201.10816.
- [12] G. Marra, S. Dumančić, R. Manhaeve, L. De Raedt, From statistical relational to neural symbolic artificial intelligence: a survey, 2021. URL: <https://arxiv.org/abs/2108.11451>. doi:10.48550/ARXIV.2108.11451.
- [13] M. Garnelo, M. Shanahan, Reconciling deep learning with symbolic artificial intelligence: representing objects and relations, *Current Opinion in Behavioral Sciences* 29 (2019) 17–23. URL: <https://www.sciencedirect.com/science/article/pii/S2352154618301943>. doi:https://doi.org/10.1016/j.cobeha.2018.12.010, *artificial Intelligence*.
- [14] G. Marra, M. Diligenti, F. Giannini, M. Gori, M. Maggini, *Relational neural machines*, 2020.

- URL: <https://arxiv.org/abs/2002.02193>. doi:10.48550/ARXIV.2002.02193.
- [15] A. Daniele, L. Serafini, Neural networks enhancement with logical knowledge, 2020. URL: <https://arxiv.org/abs/2009.06087>. doi:10.48550/ARXIV.2009.06087.
- [16] A. Daniele, L. Serafini, Knowledge enhanced neural networks for relational domains, 2022. URL: <https://arxiv.org/abs/2205.15762>. doi:10.48550/ARXIV.2205.15762.
- [17] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Liò, Y. Bengio, Graph attention networks, 2017. URL: <https://arxiv.org/abs/1710.10903>. doi:10.48550/ARXIV.1710.10903.
- [18] T. N. Kipf, M. Welling, Semi-supervised classification with graph convolutional networks, 2016. URL: <https://arxiv.org/abs/1609.02907>. doi:10.48550/ARXIV.1609.02907.
- [19] Z. Yang, W. W. Cohen, R. Salakhutdinov, Revisiting semi-supervised learning with graph embeddings, 2016. URL: <https://arxiv.org/abs/1603.08861>. doi:10.48550/ARXIV.1603.08861.
- [20] H. Zeng, H. Zhou, A. Srivastava, R. Kannan, V. Prasanna, Graphsaint: Graph sampling based inductive learning method, 2019. URL: <https://arxiv.org/abs/1907.04931>. doi:10.48550/ARXIV.1907.04931.
- [21] T. P. Adewumi, F. Liwicki, M. Liwicki, Word2vec: Optimal hyper-parameters and their impact on nlp downstream tasks, 2020. URL: <https://arxiv.org/abs/2003.11645>. doi:10.48550/ARXIV.2003.11645.
- [22] L. Zadeh, Fuzzy logic, *Computer* 21 (1988) 83–93. doi:10.1109/2.53.
- [23] E. Klement, R. Mesiar, E. Pap, *Triangular Norms*, Trends in Logic, Springer Netherlands, 2013. URL: <https://books.google.fr/books?id=HXzvCAAQBAJ>.
- [24] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, I. Polosukhin, Attention is all you need, 2017. URL: <https://arxiv.org/abs/1706.03762>. doi:10.48550/ARXIV.1706.03762.
- [25] M. Fey, J. E. Lenssen, F. Weichert, J. Leskovec, Gnnautoscale: Scalable and expressive graph neural networks via historical embeddings, 2021. URL: <https://arxiv.org/abs/2106.05609>. doi:10.48550/ARXIV.2106.05609.
- [26] W. L. Hamilton, R. Ying, J. Leskovec, Inductive representation learning on large graphs, 2017. URL: <https://arxiv.org/abs/1706.02216>. doi:10.48550/ARXIV.1706.02216.
- [27] R. Abboud, I. I. Ceylan, Node classification meets link prediction on knowledge graphs, 2021. URL: <https://arxiv.org/abs/2106.07297>. doi:10.48550/ARXIV.2106.07297.
- [28] T. P. Tanon, G. Weikum, F. M. Suchanek, Yago 4: A reason-able knowledge base, *The Semantic Web* 12123 (2020) 583 – 596.
- [29] Wikidata, <https://www.wikidata.org>, ????. Accessed: 2023-03-20.
- [30] Web ontology language (owl), <https://www.w3.org/OWL/>, ????. Accessed: 2023-03-20.
- [31] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Köpf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, S. Chintala, Pytorch: An imperative style, high-performance deep learning library, 2019. URL: <https://arxiv.org/abs/1912.01703>. doi:10.48550/ARXIV.1912.01703.
- [32] M. Fey, J. E. Lenssen, Fast graph representation learning with pytorch geometric, 2019. URL: <https://arxiv.org/abs/1903.02428>. doi:10.48550/ARXIV.1903.02428.
- [33] L. Biewald, Experiment tracking with weights and biases, 2020. URL: <https://www.wandb.com/>, software available from wandb.com.

- [34] J. Chen, T. Ma, C. Xiao, Fastgcn: Fast learning with graph convolutional networks via importance sampling, 2018. URL: <https://arxiv.org/abs/1801.10247>. doi:10.48550/ARXIV.1801.10247.
- [35] S. Ioffe, C. Szegedy, Batch normalization: Accelerating deep network training by reducing internal covariate shift, 2015. URL: <https://arxiv.org/abs/1502.03167>. doi:10.48550/ARXIV.1502.03167.
- [36] D. P. Kingma, J. Ba, Adam: A method for stochastic optimization, 2014. URL: <https://arxiv.org/abs/1412.6980>. doi:10.48550/ARXIV.1412.6980.

A. Experiment Details

A.1. Implementation

The implementation of KeGNN and the described experiments are publicly available on GitLab². The code is based on PyTorch [31] and the graph learning library PyTorch Geometric [32]. The Weights & Biases tracking tool [33] is used to monitor the experiments. For the experiments, we use a machine running an Ubuntu 20.4 equipped with an Intel(R) Xeon(R) Silver 4114 CPU 2.20GHz processor, 192G of RAM and one GPU Nvidia Quadro P5000.

A.2. Datasets

Tab. 2 gives an overview of the named datasets in this work. The datasets are publicly available on the dataset collection³ of PyTorch Geometric [32]. For the split into train, valid and test set, we take the predefined splits in [34] for the citation graphs and in [20] for Flickr. Word2Vec vectors [21] are used as node features for the citation graphs and image data for Flickr. Fig. 3 visualizes the graph structure of the underlying datasets in this work as a homogeneous, attributed and labelled graph (on the example of Citeseer).

Name	#nodes	#edges	#features	#Classes	train/valid/test split
Citeseer	3,327	9,104	3,703	6	1817/500/1000
Cora	2,708	10,556	1,433	7	1208/500/1000
PubMed	19,717	88,648	500	3	18217/500/1000
Flickr	89,250	899,756	500	7	44624/22312/22312

Table (2)

Overview of the datasets Citeseer, Cora, PubMed and Flickr

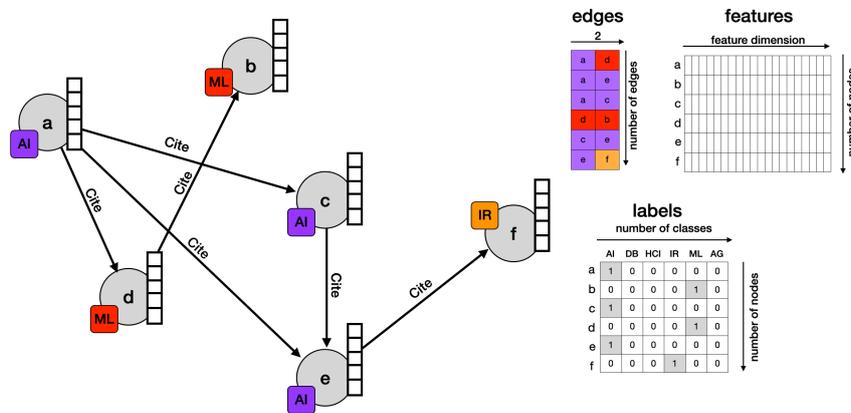


Figure (3): Visualization of the structure of the datasets used in this work with a citation graph as example.

²<https://gitlab.inria.fr/tyrex/kegnn>

³<https://pytorch-geometric.readthedocs.io/en/latest/modules/datasets.html>

A.3. Results

The average test accuracies obtained for the node classification experiments on Cora, Citeseer, PubMed and Flickr over all tested models are visualized in Fig. 4. The average runtimes per epoch on the Citeseer dataset are compared for all models in Tab 3. For runtime comparison, the runtimes were calculated for models with three hidden layers and three knowledge enhancement layers in full-batch training. It can be noted that the knowledge enhancement layers lead to increased epoch times since the model complexity is higher.

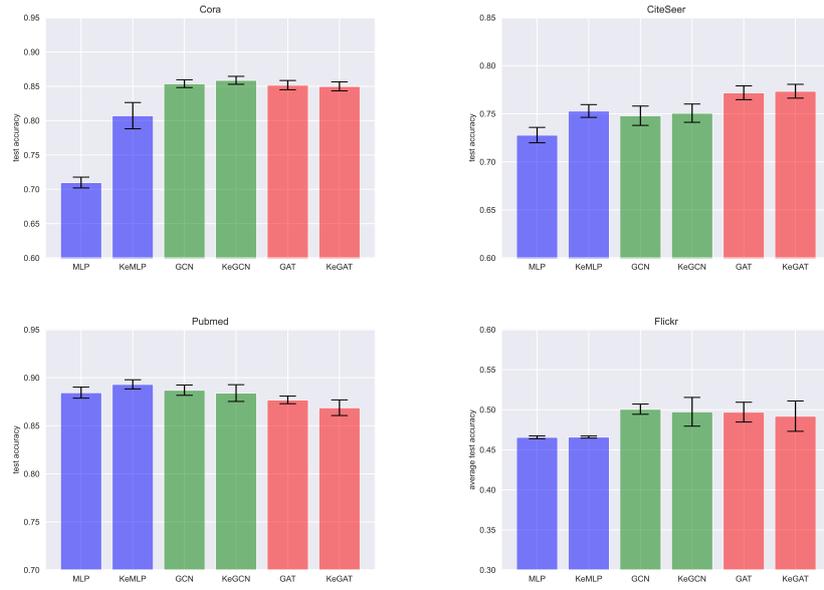


Figure (4): Average test accuracies over 50 runs for Cora, Citeseer and PubMed and 10 runs on Flickr. Error bars denote standard deviation.

Model	Avg Epoch Time
MLP	0.02684
GCN	0.03109
GAT	0.06228
KeMLP	0.04304
KeGCN	0.03747
KeGAT	0.08384

Table (3)

Comparison of the average epoch times for all models on the Citeseer dataset.

A.4. Hyperparameter Tuning

KeGNN contains a set of hyperparameters that affect the optimization and the model architecture. Batch normalization [35] is applied after each hidden layer of the GNN. Adam [36] is used as optimizer for all models. Concerning the hyperparameters specific to the knowledge

enhancement layers, the initialization of the preactivations of the binary predicates (which are assumed to be known) is taken as a hyperparameter. They are set to a high positive value for edges that are known to exist and correspond to the grounding of the binary predicate. Furthermore, different initializations of clause weights and constraints on them are tested. Moreover, the number of stacked knowledge enhancement layers is a hyperparameter. We further allow the model to randomly neglect a proportion of edges by setting an edges drop rate parameter. Further, we test whether the normalization of the edges with the diagonal matrix $\tilde{\mathbf{D}} = \sum_j \tilde{\mathbf{A}}_{i,j}$ (with $\tilde{\mathbf{A}} = \mathbf{A} + \mathbf{I}$) is helpful.

To find a suitable set hyperparameters for each dataset and model, we perform a random search with up to 800 runs and 48h time limit and choose the parameter combination which leads to the highest accuracy on the validation set. The hyperparameter tuning is executed in Weights and Biases [33]. The following hyperparameter values are tested:

- Adam optimizer parameters: β_1 : 0.9, β_2 : 0.99, ϵ : 1e-07
- Attention heads: {1, 2, 3, 4, 6, 8, 10}
- Batch size: {128, 512, 1024, 2048, full batch}
- Binary preactivation value: {0.5, 1.0, 10.0, 100.0, 500.0}
- Clause weights initialization: {0.001, 0.1, 0.25, 0.5, random uniform distribution on [0,1]}
- Dropout rate: 0.5
- Edges drop rate: random uniform distribution [0.0, 0.9]
- Edge normalization: {true, false}
- Early stopping: δ_{min} : 0.001, patience: {1, 10, 100}
- Hidden layer dimension: {32, 64, 128, 256}
- Learning rate: random uniform distribution [0.0001, 0.1]
- Clause weight clipping: w_{min} : 0.0, w_{max} : random uniform distribution: [0.8, 500.0]
- Number of knowledge enhancement layers: {1, 2, 3, 4, 5, 6}
- Number of hidden layers: {2, 3, 4, 5, 6}

The obtained parameter combinations for the models KeMLP, KeGCN and KeGAT for Cora, Citeseer, PubMed and Flickr are displayed in Tab. 5 and Tab. 4. The reference models MLP, GCN and GAT are trained with the same parameter set as the respective enhanced models.

B. Clause Weight Evaluation

The *clause compliance* [15] indicates the level of satisfaction of a clause in the data in this experimental setting. Given a clause c , a class k , the set of training nodes $\mathbf{V}_{\text{train}}$, the set of nodes of the class k : $\mathbf{V}_k = \{v | v \in \mathbf{V}_{\text{train}} \wedge \text{Cls}(v) == k\}$, and the neighborhood $\mathcal{N}(v)$ of v , the clause compliance on graph \mathbf{G} is defined as follows:

$$\text{Compliance}(\mathbf{G}, c) = \frac{\sum_{v \in \mathbf{V}_k} \sum_{u \in \mathcal{N}(v)} \mathbf{1}(u \in \mathbf{V}_k)}{\sum_{v \in \mathbf{V}_k} |\mathcal{N}(v)|} \quad (4)$$

In other words, the clause compliance counts how often among nodes of a class the neighboring nodes are of the same class. [15]

Parameter	PubMed			Flickr		
	KeMLP	KeGCN	KeGAT	KeMLP	KeGCN	KeGAT
adam beta 1	0.9	0.9	0.9	0.9	0.9	0.9
ada beta 2	0.99	0.99	0.99	0.99	0.99	0.99
adam epsilon	1e-07	1e-07	1e-07	1e-07	1e-07	1e-07
attention heads	-	-	8	-	-	8
batch size	1024	full batch	1024	128	1024	2048
binary preactivation	10.0	1.0	10.0	10.0	500.0	500.0
clause weight initialization	0.001	random	0.5	0.001	0.001	0.1
dropout rate	0.5	0.5	0.5	0.5	0.5	0.5
edges drop rate	0.22	0.66	0.07	0.2	0.24	0.12
epochs	200	200	200	200	200	200
early stopping enabled	true	true	true	true	true	true
early stopping min delta	0.001	0.001	0.001	0.001	0.001	0.001
early stopping patience	100	10	10	10	10	100
hidden channels	256	256	256	32	128	64
learning rate	0.057	0.043	0.016	0.001	0.016	0.0039
max clause weight	350.0	322.0	118.0	55.0	135	113.0
min clause weight	0.0	0.0	0.0	0.0	0	0.0
normalize edges	false	false	true	true	true	false
KE layers	2	1	5	1	4	1
hidden layers	4	2	2	2	4	3
runs	50	50	50	10	10	10
seed	1234	1234	1234	1234	1234	1234

Table (4)

Hyperparameters and experiment configuration for PubMed and Flickr

Parameter	Cora			CiteSeer		
	KeMLP	KeGCN	KeGAT	KeMLP	KeGCN	KeGAT
adam beta 1	0.9	0.9	0.9	0.9	0.9	0.9
ada beta 2	0.99	0.99	0.99	0.99	0.99	0.99
adam epsilon	1e-07	1e-07	1e-07	1e-07	1e-07	1e-07
attention heads	-	-	1	-	-	3
batch size	512	512	full batch	128	full batch	1024
binary preactivation	10.0	500.0	1.0	10.0	0.5	0.5
clause weight initialization	0.5	random	0.5	0.5	0.25	0.1
dropout rate	0.5	0.5	0.5	0.5	0.5	0.5
edges drop rate	0.47	0.17	0.27	0.01	0.35	0.88
epochs	200	200	200	200	200	200
early stopping enabled	true	true	true	true	true	true
early stopping min delta	0.001	0.001	0.001	0.001	0.001	0.001
early stopping patience	1	1	10	10	10	10
hidden channels	32	256	64	256	128	32
learning rate	0.026	0.032	0.033	0.028	0.037	0.006
max clause weight	104.0	254.0	250.0	34.0	243.0	110.0
min clause weight	0.0	0.0	0.0	0.0	0.0	0.0
normalize edges	true	false	true	true	false	true
KE layers	4	2	1	1	3	2
Hidden layers	2	2	2	2	5	2
runs	50	50	50	50	50	50
seed	1234	1234	1234	1234	1234	1234

Table (5)

Hyperparameter and experiment configuration for Citeseer and Cora