

# Knowledge Enhanced Graph Neural Networks for Graph Completion

Luisa Werner<sup>1,2</sup>, Nabil Layaïda<sup>2</sup>, Pierre Genevès<sup>1,3</sup> and Sarah Chlyah<sup>1</sup>

<sup>1</sup>Institut National de Recherche en Sciences et Technologies du Numérique (INRIA)

<sup>2</sup>Université Grenoble Alpes (UGA)

<sup>3</sup>Centre National de la Recherche Scientifique (CNRS)

luisa.werner@inria.fr, nabil.layaida@inria.fr, pierre.geneves@inria.fr, sarah.chlyah@inria.fr

## Abstract

Graph data is omnipresent and has a wide variety of applications, such as in natural science, social networks, or the semantic web. However, while being rich in information, graphs are often noisy and incomplete. As a result, graph completion tasks, such as node classification or link prediction, have gained attention. On 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 Knowledge Enhanced Graph Neural Networks (KeGNN), a neurosymbolic framework for graph completion that combines both paradigms as it allows for the integration of prior knowledge into a graph neural network model. Essentially, KeGNN consists of a graph neural network as a base upon which knowledge enhancement layers are stacked with the goal of refining predictions with respect to prior knowledge. We instantiate KeGNN in conjunction with two state of the art graph neural networks, Graph Convolutional Networks and Graph Attention Networks, and evaluate KeGNN on multiple benchmark datasets for node classification.

## 1 Introduction

Graphs are ubiquitous across diverse real-world applications such as e-commerce [Liu *et al.*, 2021], natural science [Sanchez-Gonzalez *et al.*, 2018] or social networks [Wu *et al.*, 2020]. 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 [Ma and Tang, 2021] [Wu *et al.*, 2021] [Duan *et al.*, 2022]. The key strength of GNNs is to find meaningful representations of noisy data, that can be used for prediction

tasks [Wu *et al.*, 2022]. Despite this advantage, as a subcategory of deep learning methods, GNNs are criticized for their limited interpretability and large data consumption [Susskind *et al.*, 2021]. 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-efficient. 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 AI 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 helps to overcome the blackbox nature of deep learning methods and to improve interpretability through symbolic representations [Susskind *et al.*, 2021].

In this paper, we present the neuro-symbolic approach Knowledge enhanced Graph Neural Networks (KeGNN) to conduct node classification given graph data and a set of prior knowledge. In KeGNN, 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. In addition to the parameters of the GNN, the knowledge enhancement layers contain learnable clause weights that reflect the impact of the prior knowledge on the predictions. Both components form an end-to-end differentiable model. KeGNN can be seen as an extension to knowledge enhanced neural networks (KENN) [Daniele and Serafini, 2022], 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 of KENN. In contrast, KeGNN is based on GNNs that process the graph structure, which makes both the neural and symbolic components sufficiently powerful to exploit the graph structure. In this work, we instantiate KeGNN in conjunction with two popular GNNs: Graph Attention Networks [Veličković *et al.*, 2018] and Graph Convolutional Networks [Kipf and Welling, 2017]. We apply KeGNN to the benchmark datasets for node classification Cora, Citeseer, PubMed [Yang *et al.*, 2016] and Flickr [Zeng *et al.*, 2020].

## 2 Method: KeGNN

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 a set of  $n$  nodes  $\mathbf{N}$  and a set of  $k$  edges  $\mathbf{E}$  where each edge of the form  $(v_i, v_j)$  connects two nodes  $v_i \in \mathbf{N}$  and  $v_j \in \mathbf{N}$ . 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 features and labels. Each node has a feature vector  $\mathbf{x} \in \mathbb{R}^d$  of dimension  $d$  and a label vector  $\mathbf{y} \in \mathbb{R}^m$ . The label vector  $\mathbf{y}$  contains one-hot encoded ground truth labels for  $m$  classes. In matrix notation, the features and labels of the entire graph are described as  $\mathbf{X} \in \mathbb{R}^{n \times d}$  and  $\mathbf{Y} \in \mathbb{R}^{n \times m}$ . A graph is *typed* if the type functions  $f_{\mathbf{E}}$  and  $f_{\mathbf{N}}$  assign edge types and node types to the edges and nodes, respectively. A graph with constant type functions (that assign the same edge and node type to all edges and nodes) is called *homogeneous*, whereas for *heterogeneous* graphs, nodes and edges may have different types [Ma and Tang, 2021].

**Example 2.1.** A Citation Graph  $\mathbf{G}_{\text{Cit}}$  consists of documents and citations. Fig. 1 shows an extract of the Citeseer citation graph that is used as example to guide through the method section. The documents are represented by nodes  $\mathbf{N}_{\text{Cit}}$  and citations by edges  $\mathbf{E}_{\text{Cit}}$ . Documents can be attributed with features  $\mathbf{X}_{\text{Cit}}$  that describe their content as Word2Vec [Adebumi et al., 2020] vectors. Each node is labelled with one of six topic categories  $\{\text{AI}, \text{DB}, \text{HCI}, \text{IR}, \text{ML}, \text{AG}\}$ <sup>1</sup> that are encoded in  $\mathbf{Y}_{\text{Cit}}$ . Since all nodes (documents) and edges (citations) have the same type,  $\mathbf{G}_{\text{Cit}}$  is homogeneous.

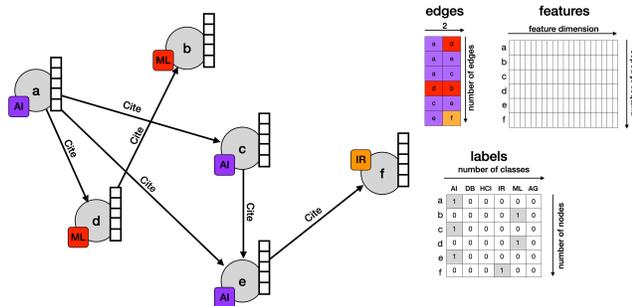


Figure (1) Example extract of the Citeseer citation graph.

### 2.2 Prior Knowledge

The prior knowledge  $\mathcal{K}$  provided to KeGNN can be described as a set of  $\ell$  logical clauses expressed in the logical language

<sup>1</sup>The classes are abbreviations for the categories *Artificial Intelligence*, *Databases*, *Human-Computer Interaction*, *Information Retrieval*, *Machine Learning* and *Agents*.

$\mathcal{L}$  that is defined as sets of constants  $\mathcal{C}$ , variables  $\mathcal{X}$  and predicates  $\mathcal{P}$ . Predicates have an arity  $r$  of one (unary) or two (binary):  $\mathcal{P} = \mathcal{P}_U \cup \mathcal{P}_B$ . Predicates of arity  $r > 2$  are not considered in this work. Unary predicates express properties, whereas binary predicates express relations.  $\mathcal{L}$  supports the operators negation ( $\neg$ ) and disjunction ( $\vee$ ). Each clause  $\varphi \in \mathcal{K} = \{\varphi_1, \dots, \varphi_\ell\}$  can be formulated as a disjunction of (possibly negated) atoms  $\bigvee_{j=1}^q o_j$  with  $q$  atoms  $\{o_1, \dots, o_q\}$ . Since the prior knowledge is general, all clauses are assumed to be universally quantified. Clauses can be *grounded* by assigning constants to the free variables. A grounded clause is denoted as  $\varphi[x_1, x_2, \dots | c_1, c_2, \dots]$  with variables  $x_i \in \mathcal{X}$  and constants  $c_i \in \mathcal{C}$ . The set of all grounded clauses in a graph is  $\mathcal{G}(\mathcal{K}, \mathcal{C})$ .

**Example 2.2.** The graph  $\mathbf{G}_{\text{Cit}}$  in Fig. 1 can be expressed in  $\mathcal{L}$ . Nodes are represented by a set of constants  $\mathcal{C} = \{a, b, \dots, f\}$ . Node labels are expressed as a set of unary predicates  $\mathcal{P}_U = \{\text{AI}, \text{DB}, \dots, \text{AG}\}$  and edges as a set of binary predicates  $\mathcal{P}_B = \{\text{Cite}\}$ .  $\mathcal{L}$  has a set of variables  $\mathcal{X} = \{x, y\}$ . The atom  $\text{AI}(x)$  expresses the membership of  $x$  to the class AI and  $\text{Cite}(x, y)$  expresses the existence of a citation between  $x$  and  $y$ . A set of prior knowledge  $\mathcal{K}$  can be written as  $\ell = 6$  disjunctive clauses in  $\mathcal{L}$ . Here, the assumption is denoted that two papers that cite each other have the same document class:

$$\begin{aligned} & \forall xy \neg \text{AI}(x) \vee \neg \text{Cite}(x, y) \vee \text{AI}(y) \\ & \forall xy \neg \text{DB}(x) \vee \neg \text{Cite}(x, y) \vee \text{DB}(y) \end{aligned}$$

...

The atoms are grounded by replacing the variables  $x$  and  $y$  with the constants  $\{a, b, \dots, f\}$  to obtain sets of unary groundings  $\{\text{AI}(a), \text{ML}(b), \dots, \text{IR}(f)\}$  and binary groundings  $\{\text{Cite}(a, d), \text{Cite}(a, e), \dots, \text{Cite}(a, f)\}$ . Assuming a closed world and exclusive classes, other facts could be derived, such as  $\{\neg \text{DB}(a), \neg \text{IR}(a), \dots, \neg \text{Cite}(a, b)\}$ . For the sake of simplicity, these are omitted here.

### 2.3 Node Classification

Node classification is a subtask of knowledge graph completion on a graph  $\mathbf{G}$  with the objective to assign classes to nodes where they are unknown. This task is accomplished given node features  $\mathbf{X}$ , edges  $\mathbf{E}$  and some prior knowledge  $\mathcal{K}$  encoded as a set of clauses in  $\mathcal{L}$ . A predictive model is trained on a subset of the graph  $\mathbf{G}_{\text{train}}$  with ground truth labels  $\mathbf{Y}_{\text{train}}$  and validated on a test set  $\mathbf{G}_{\text{test}}$  for which the ground truth labels are compared to the predictions in order to assess the predictive performance. Node classification can be studied in a *transductive* or *inductive* setting. In a transductive setting, the entire graph is available for training, but the true labels of the test nodes are masked. In an inductive setting, only the nodes in the training set and the edges connecting them are available, making it more challenging to classify unseen nodes.

### 2.4 Fuzzy Semantics

Let us consider an attributed and labelled graph  $\mathbf{G}$  and a set of prior knowledge  $\mathcal{K}$ . While  $\mathcal{K}$  can be defined in the logic language  $\mathcal{L}$ , the neural component in KeGNN relies on continuous and differentiable representations. To interpret Boolean

164 logic in the real-valued domain, KeGNN uses fuzzy logic  
 165 [Zadeh, 1988], which maps Boolean truth values to the con-  
 166 tinuous interval  $[0, 1] \subset \mathbb{R}$ . A constant in  $\mathcal{C}$  is interpreted as  
 167 a real-valued feature vector  $\mathbf{x} \in \mathbb{R}^d$ . A predicate in  $\mathcal{P}$  with  
 168 arity  $r$  is interpreted as a function  $f_{\mathcal{P}} : \mathbb{R}^{r \times d} \mapsto [0, 1]$  that  
 169 takes  $r$  feature vectors as input and returns a truth value.

**Example 2.3.** In the example, a unary predicate  $P_U \in \mathcal{P}_U = \{\text{AI}, \text{DB}, \dots\}$  is interpreted as a function  $f_{P_U} : \mathbb{R}^d \mapsto [0, 1]$  that takes a feature vector  $\mathbf{x}$  and returns a truth value indicating whether the node belongs to the class encoded as  $P_U$ . The binary predicate  $\text{Cite} \in \mathcal{P}_B$  is interpreted as the function

$$f_{\text{Cite}}(v_i, v_j) = \begin{cases} 1, & \text{if } (v_i, v_j) \in \mathbf{E}_{\text{Cit}} \\ 0, & \text{else.} \end{cases}$$

170  $f_{\text{Cite}}$  returns 1 (true) if there is an edge between two nodes  $v_i$   
 171 and  $v_j$  in  $\mathbf{G}_{\text{Cit}}$  and 0 otherwise.

T-conorm functions  $\perp : [0, 1] \times [0, 1] \mapsto [0, 1]$  [Klement *et al.*, 2013] take real-valued truth values of two literals<sup>2</sup> and define the truth value of their disjunction. The Gödel t-conorm function for two truth values  $\mathbf{t}_i, \mathbf{t}_j$  is defined as

$$\perp(\mathbf{t}_i, \mathbf{t}_j) \mapsto \max(\mathbf{t}_i, \mathbf{t}_j).$$

172 To obtain the truth value of a clause  $\varphi : o_1 \vee \dots \vee o_q$ ,  
 173 the function  $\perp$  is extended to a vector  $\mathbf{t}$  of  $q$  truth values:  
 174  $\perp(\mathbf{t}_1, \mathbf{t}_2, \dots, \mathbf{t}_q) = \perp(\mathbf{t}_1, \perp(\mathbf{t}_2, \dots, \perp(\mathbf{t}_{q-1}, \mathbf{t}_q)))$ . Fuzzy nega-  
 175 tion over truth values is defined as  $\mathbf{t} \mapsto 1 - \mathbf{t}$  [Zadeh, 1988].

176 **Example 2.4.** Given the clause  $\varphi_{\text{AI}} : \forall xy \neg \text{AI}(x) \vee$   
 177  $\neg \text{Cite}(x, y) \vee \text{AI}(y)$  and its grounding  $\varphi_{\text{AI}}[x, y|a, b] :$   
 178  $\text{AI}(a) \vee \neg \text{Cite}(a, b) \vee \text{AI}(b)$  to the constants  $a$  and  $b$  and truth  
 179 values for the grounded predicates  $\text{AI}(a) = \mathbf{t}_1$ ,  $\text{AI}(b) = \mathbf{t}_2$   
 180 and  $\text{Cite}(a, b) = \mathbf{t}_3$ , the truth value of  $\varphi_{\text{AI}}[x, y|a, b]$  is  
 181  $\max\{\max\{(1 - \mathbf{t}_1), (1 - \mathbf{t}_3)\}, \mathbf{t}_2\}$ .

## 182 2.5 Model Architecture

183 The way KeGNN computes the final predictions can be di-  
 184 vided in two stages. First, a GNN predicts the node classes  
 185 given the features and the edges. Subsequently, the knowl-  
 186 edge enhancement layers use the predictions as truth values  
 187 for the grounded unary predicates and update them with re-  
 188 spect to the knowledge. An overview of KeGNN is given in  
 189 Fig. 2.

### 190 Neural Component

The role of the GNN in 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 struc-  
 ture by nesting  $k$  message passing layers [Wu *et al.*, 2022].  
 Per layer, the representations of neighboring nodes are aggre-  
 gated and combined to obtain updated representations. The  
 node representation  $v_i^{k+1}$  in the  $k$ -th message passing layer is

$$v_i^{k+1} = \text{combine}(v_i^k, \text{aggregate}(\{v_j^k | v_j^k \in \mathcal{N}(v_i)\})).$$

191 The layers contain learnable parameters that are optimized  
 192 with backpropagation. In this work, we consider two well-  
 193 known GNNs as components for KeGNN: Graph Convolu-  
 194 tional Networks (GCN) [Kipf and Welling, 2017] and Graph

<sup>2</sup>A literal is a (possibly negated) grounded atom, e.g.  $\text{AI}(a)$

Attention Networks (GAT) [Veličković *et al.*, 2018]. While  
 GCN considers the graph structure as given, GAT allows  
 for assessing the importance of the neighbors with attention  
 weights  $\alpha_{ij}$  between node  $v_i$  and node  $v_j$ . In case of multi-  
 head attention, the attention weights are calculated multiple  
 times and concatenated which allows for capturing different  
 aspects of the input data. In KeGNN, the GNN implements  
 the functions  $f_{P_U}$  (see Section 2.4). In other words, the pre-  
 dictions are used as truth values for the grounded unary predi-  
 cates in the symbolic component.

### Symbolic Component

To refine the predictions of the GNN, one or more knowl-  
 edge enhancement layers are stacked onto the GNN to update  
 its predictions  $\mathbf{Y}$  to  $\mathbf{Y}'$ . The goal is to increase the satisfac-  
 tion of the prior knowledge. The predictions  $\mathbf{Y}$  of the GNN  
 serve as input to the symbolic component where they are in-  
 terpreted as fuzzy truth values for the unary grounded predi-  
 cates  $\mathbf{U} := \mathbf{Y}$  with  $\mathbf{U} \in \mathbb{R}^{n \times m}$ . Fuzzy truth values for the  
 groundings of binary predicates are encoded as a matrix  $\mathbf{B}$   
 where each row represents an edge  $(v_i, v_j)$  and each column  
 represents an edge type  $e$ . In the context of node classifica-  
 tion, the GNN returns only predictions for the node classes,  
 while the edges are assumed to be given. A binary grounded  
 predicate is therefore set to truth value 1 (true) if an edge be-  
 tween two nodes  $v_i$  and  $v_j$  exists:

$$\mathbf{B}_{[(v_i, v_j), e]} = \begin{cases} 1, & \text{if } (v_i, v_j) \text{ of type } e \in \mathbf{E} \\ 0, & \text{else.} \end{cases}$$

**Example 2.5.** In case of the beforementioned citation graph  
 of Fig. 1,  $\mathbf{U}$  and  $\mathbf{B}$  are defined as:

$$\mathbf{U} := \begin{bmatrix} \text{AI}(a) & \dots & \text{AG}(a) \\ \text{AI}(b) & \dots & \text{AG}(b) \\ \vdots & & \vdots \\ \text{AI}(f) & \dots & \text{AG}(f) \end{bmatrix} \quad \mathbf{B} := \begin{bmatrix} \text{Cite}(a, d) \\ \text{Cite}(a, e) \\ \text{Cite}(a, c) \\ \vdots \\ \text{Cite}(c, e) \\ \text{Cite}(e, f) \end{bmatrix}$$

To enhance the satisfaction of clauses that contain both  
 unary and binary predicates, their groundings are joined into  
 one matrix  $\mathbf{M} \in \mathbb{R}^{k \times P}$  with  $P = 2 \cdot |\mathcal{P}_U| + |\mathcal{P}_B|$ .  $\mathbf{M}$  is  
 computed by joining  $\mathbf{U}$  and  $\mathbf{B}$  so that each row of  $\mathbf{M}$  repre-  
 sents an edge  $(v_i, v_j)$ . As a result,  $\mathbf{M}$  contains all required  
 grounded unary predicates for  $v_i$  and  $v_j$ .

**Example 2.6.** For the example citation graph, we obtain  $\mathbf{M}$   
 as follows:

$$\mathbf{M} = \begin{array}{c} \begin{array}{ccc} \text{unary grounded predicates for } x & \text{unary grounded predicates for } y & \text{binary groundings} \end{array} \\ \left[ \begin{array}{ccc|ccc|c} \text{AI}(a) & \text{DB}(a) & \dots & \text{AG}(a) & \text{AI}(d) & \text{DB}(d) & \dots & \text{AG}(d) & \text{Cite}(a, d) \\ \text{AI}(a) & \text{DB}(a) & \dots & \text{AG}(a) & \text{AI}(e) & \text{DB}(e) & \dots & \text{AG}(e) & \text{Cite}(a, e) \\ \text{AI}(a) & \text{DB}(a) & \dots & \text{AG}(a) & \text{AI}(c) & \text{DB}(c) & \dots & \text{AG}(c) & \text{Cite}(a, c) \\ \vdots & & & \vdots & \vdots & & & \vdots & \vdots \\ \text{AI}(c) & \text{DB}(c) & \dots & \text{AG}(c) & \text{AI}(e) & \text{DB}(e) & \dots & \text{AG}(e) & \text{Cite}(c, e) \\ \text{AI}(e) & \text{DB}(e) & \dots & \text{AG}(e) & \text{AI}(f) & \text{DB}(f) & \dots & \text{AG}(f) & \text{Cite}(e, f) \end{array} \right] \end{array}$$

For each clause  $\varphi \in \mathcal{K}$ , a *clause enhancer* is instantiated.  
 Its aim is to compute updates  $\delta \mathbf{M}_{\varphi}$  for the groundings in  $\mathbf{M}$

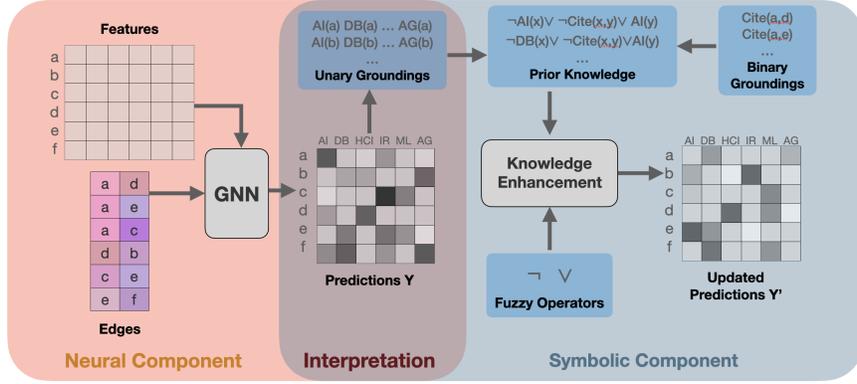


Figure (2) Overview of KeGNN.

that increase the satisfaction of  $\varphi$ . First, fuzzy negation is applied to the columns of  $\mathbf{M}$  that correspond to negated atoms in  $\varphi$ . Then  $\delta\mathbf{M}_\varphi$  is computed by a *t-conorm boost function*  $\phi$  [Daniele and Serafini, 2020].  $\phi : [0, 1]^q \mapsto [0, 1]^q$  takes  $q$  truth values and returns changes to those truth values such that  $\perp(\mathbf{t}) \leq \perp(\mathbf{t} + \phi(\mathbf{t}))$ . [Daniele and Serafini, 2020] propose the following differentiable t-conorm boost function

$$\phi_{w_\varphi}(\mathbf{t})_i = w_\varphi \cdot \frac{e^{t_i}}{\sum_{j=1}^q e^{t_j}}.$$

214 The boost function  $\phi_{w_\varphi}$  employs a clause weight  $w_\varphi$  as a  
 215 learnable parameter so that the updates for the groundings are  
 216 proportional to  $w_\varphi$ . Therefore,  $w_\varphi$  determines the magnitude  
 217 of the update and thus reflects the impact of a clause. The  
 218 changes to atoms that do not appear in a clause are set to zero.  
 219 The boost function is applied row-wise to  $\mathbf{M}$  as illustrated in  
 220 the following example.

**Example 2.7.** Given the clause  $\varphi_{AI} : \forall xy \neg AI(x) \vee \neg Cit(x, y) \vee AI(y)$  with the clause weight  $w_{AI}$ , the changes for this clause are  $\delta\mathbf{M}_{\varphi_{AI}} =$

$$w_{AI} \cdot \begin{bmatrix} \delta_{\neg AI^x(a)} & 0 & \dots & \delta_{AI^y(c)} & 0 & \dots & \delta_{\neg Cit(a,c)} \\ \delta_{\neg AI^x(a)} & 0 & \dots & \delta_{AI^y(e)} & 0 & \dots & \delta_{\neg Cit(e,a)} \\ \delta_{\neg AI^x(a)} & 0 & \dots & \delta_{AI^y(d)} & 0 & \dots & \delta_{\neg Cit(c,d)} \\ \vdots & & & \vdots & \vdots & & \\ \delta_{\neg AI^x(e)} & 0 & \dots & \delta_{AI^y(f)} & 0 & \dots & \delta_{\neg Cit(e,f)} \end{bmatrix}$$

The values of  $\delta\mathbf{M}_{\varphi_{AI}}$  are calculated by  $\phi_{w_{AI}}$ , for example:

$$\delta_{\neg AI^x(a)} = \phi_{w_{AI}}(\mathbf{z})_a = -\frac{e^{-\mathbf{z}_{AI(a)}}}{e^{-\mathbf{z}_{AI(a)}} + e^{-\mathbf{z}_{Cit(a,c)}} + e^{\mathbf{z}_{AI(c)}}$$

A clause enhancer is instantiated for each clause  $\varphi \in \mathcal{K}$ . Each clause enhancer computes updates  $\delta\mathbf{M}_\varphi$  for a clause independently. The updates of all clause enhancers are finally added, resulting in a matrix  $\delta\mathbf{M} = \sum_{\varphi \in \mathcal{K}} \delta\mathbf{M}_\varphi$ . To apply the updates to the initial predictions,  $\delta\mathbf{M}$  has to be added to  $\mathbf{Y}$ . The updates in  $\delta\mathbf{M}$  can not directly be applied to the predictions  $\mathbf{Y}$  of the GNN. Since the unary groundings  $\mathbf{U}$  were joined with  $\mathbf{B}$ , multiple changes may be proposed for the same grounded unary atom. For example, for the grounded atom  $AI(c)$  the changes  $\delta_{\neg AI^y(c)}$  and  $\delta_{\neg AI^x(c)}$  are proposed,

since  $c$  appears in first place of edge  $(a, c)$  and in second place of edge  $(c, e)$ . Therefore, all updates for the same grounded atom are summed, reducing the size of  $\mathbf{M}$  to the size of  $\mathbf{U}$ . To ensure that the updated predictions remain truth values in the range of  $[0, 1]$ , the knowledge enhancer works with the preactivations  $\mathbf{Z}$  of the GNN and applies the activation function  $\sigma$  to the updated preactivations  $\mathbf{Z}'$  to obtain the final predictions:  $\mathbf{Y}' = \sigma(\mathbf{Z}')$ . Therefore, the knowledge enhancer transforms  $\mathbf{Z}$  to  $\mathbf{Z}'$  (with  $\mathbf{Z}, \mathbf{Z}' \in \mathbb{R}^{n \times m}$ ). Regarding the binary groundings, the values in  $\mathbf{B}$  are set to a high positive value that results in one when  $\sigma$  is applied. In the last step, the updates by the knowledge enhancer are added to the preactivations  $\mathbf{Z}$  of the GNN and passed to  $\sigma$  to obtain the updated predictions

$$\mathbf{Y}' = \sigma \left( \mathbf{Z}x + \sum_{\varphi \in \mathcal{K}} \delta\mathbf{U}_\varphi \right)$$

where  $\delta\mathbf{U}_\varphi$  is the matrix obtained by extracting the changes to the unary predicates from  $\delta\mathbf{M}_\varphi$ .

### 3 Related Work

The field of knowledge graph completion is addressed from several research directions. Symbolic methods exist that conduct link prediction given a set of prior knowledge [Dou *et al.*, 2015] [Meilicke *et al.*, 2019]. Embedding-based methods [Dai *et al.*, 2020] are mostly sub-symbolic methods to obtain node embeddings that are used for knowledge graph completion tasks. Usually, their common objective is to find similar embeddings for nodes that are located closely in the graph. The majority of these methods only encodes the graph structure, but does not consider node-specific feature information [Abboud and Ceylan, 2021]. However, KeGNN is based on GNNs that are suited for learning representations of graphs attributed with node features. It stacks additional layers that interpret the outputs of the GNN in fuzzy logic and modify them to increase the satisfiability. Therefore, it is considered a neuro-symbolic method. In the multifaceted neuro-symbolic field, KeGNN can be placed in the category of knowledge-guided learning [Daniele and Serafini, 2020], where the focus lies on learning in the presence of additional supervision introduced as prior knowledge. Within this cate-

244 gory, KeGNN belongs to the model-based approaches, where  
 245 prior knowledge in the form of knowledge enhancement lay-  
 246 ers is an integral part of the model. Beyond, loss-based meth-  
 247 ods such as logic tensor networks [Badreddine *et al.*, 2022]  
 248 exist that encode the satisfiability of prior knowledge as an  
 249 optimization objective.

250 Further, in [DeLong *et al.*, 2023] neuro-symbolic ap-  
 251 proaches dealing with graph structures are classified into  
 252 three categories. First, logically informed embedding ap-  
 253 proaches [Li *et al.*, 2023] [Jain *et al.*, 2021] use predefined  
 254 logical rules that provide knowledge to a neural system, while  
 255 both components are mostly distinct. Second, approaches for  
 256 knowledge graph embedding with logical constraints [Fatemi  
 257 *et al.*, 2019] [Guo *et al.*, 2016] use prior knowledge as con-  
 258 straints on the neural knowledge graph embedding method in  
 259 order to modify predictions or embeddings. Thirdly, neuro-  
 260 symbolic methods are used for learning rules for graph reason-  
 261 ing tasks [Hu *et al.*, 2020] [Qu *et al.*, 2021]. This allows  
 262 for rule generation or confidence scores for prior knowledge  
 263 and makes the models robust to exceptions or soft knowledge.  
 264 KeGNN best falls into the second category, since the prior  
 265 knowledge is interpreted in fuzzy logic to be integrated with  
 266 the neural model and update the GNN’s predictions. The idea  
 267 of confidence values in category three shares the common  
 268 property of relativating knowledge as with KeGNN’s clause  
 269 weights. However, even though KeGNN’s clause weights in-  
 270 troduce a notion of impact of a clause when predictions are  
 271 made, they cannot directly be interpreted as the confidence in  
 272 a rule. In the well-known *Kautz Taxonomy* [Kautz, 2022] that  
 273 classifies neuro-symbolic approaches according to the inte-  
 274 gration of neural and symbolic modules, KeGNN falls best  
 275 into the category *Neuro[Symbolic]* (Type 6) of fully-  
 276 integrated neuro-symbolic systems that embed symbolic rea-  
 277 soning in a neural architecture.

## 278 4 Experimental Evaluation

279 To evaluate the performance of KeGNN, we apply it to the  
 280 datasets Citeseer, Cora, PubMed and Flickr that are com-  
 281 mon benchmarks for node classification in a transductive set-  
 282 ting. In the following, KeGNN is called KeGCN and KeGAT  
 283 when instantiated to a GCN or a GAT, respectively. As ad-  
 284 ditional baseline, we consider KeMLP, that stacks knowledge  
 285 enhancement layers onto an MLP, as proposed in [Daniele  
 286 and Serafini, 2022]. Further, the standalone neural models  
 287 MLP, GCN and GAT are used as baselines. While Citeseer,  
 288 Cora and PubMed are citation graphs that encode citations  
 289 between scientific papers (as in Example 2.2), Flickr contains  
 290 images and shared properties between them. All datasets can  
 291 be modelled as homogeneous, labelled and attributed graphs  
 292 as defined in Section 2.1. The set of prior logic for the knowl-  
 293 edge enhancement layers is given explicitly. In this work,  
 294 we encode the assumption that the existence of an edge for  
 295 a node pair points to their membership to the same class and  
 296 hence provides added value to the node classification task.  
 297 In the context of citation graphs, this implies that two docu-  
 298 ments that cite each other refer to the same topic, while for  
 299 Flickr, linked images share the same properties. Following  
 300 this pattern for all datasets, a clause  $\varphi: \forall xy : \neg \text{Cls}_i(x) \vee$

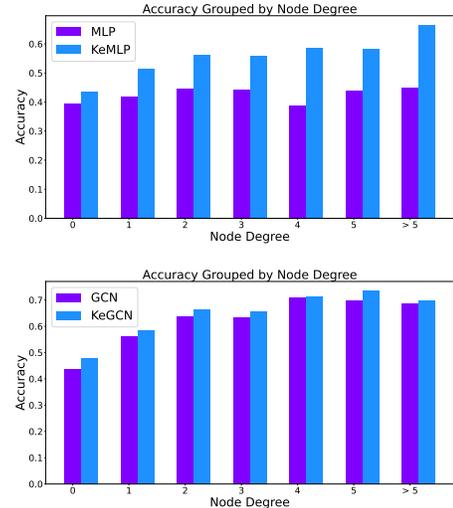


Figure (3) The accuracy grouped by the node degree for MLP vs. KeMLP (above) and GCN and KeGCN (below) on Citeseer.

–Link( $x, y$ )  $\vee$   $\text{Cls}_i(y)$  is instantiated for each node class  $\text{Cls}_i$ . 301  
 More details on the experiments are given in Appendix A. 302  
 The implementation and the experiments is publicly available 303  
 on Gitlab<sup>3</sup>. 304

### 4.1 Results 305

To compare the performance of all models, we examine the 306  
 average test accuracy over 50 runs (10 for Flickr) for the 307  
 knowledge enhanced models KeMLP, KeGCN, KeGAT and 308  
 the standalone base models MLP, GCN, GAT on the named 309  
 datasets. The results are given in Tab. 1 and visualized in 310  
 Fig. 7 (see Appendix A.3). For Cora and Citeseer, KeMLP 311  
 leads to a significant improvement over MLP (p-value of one- 312  
 sided t-test  $\ll 0.05$ ). In contrast, no significant advantage 313  
 of KeGCN or KeGAT in comparison to the standalone base 314  
 model is observed. Nevertheless, all GNN-based models are 315  
 significantly superior to KeMLP for Cora. This includes not 316  
 only KeGCN and KeGAT, but also the GNN baselines. For 317  
 Citeseer, KeGAT and GAT both outperform KeMLP. In the 318  
 case of PubMed, only a significant improvement of KeMLP 319  
 over MLP can be observed, while the GNN-based models 320  
 and their enhanced versions do not provide any positive ef- 321  
 fect. For Flickr, no significant improvement between the base 322  
 model and the respective knowledge enhanced model can be 323  
 observed. Nevertheless, all GNN-based models outperform 324  
 KeMLP, reporting significantly higher mean test accuracies 325  
 for KeGAT, GAT, GCN and KeGCN. 326

### Exploitation of the Graph Structure 327

It turns out that the performance gap between MLP and 328  
 KeMLP is larger than for KeGCN in comparison to the stan- 329  
 dalone GCN (or KeGAT vs. GAT, respectively). To explain 330  
 this observation, we examine how the graph structure affects 331  
 the prediction performance. Therefore, in Fig. 3 we analyze 332  
 the accuracy grouped by the node degree for the entire graph 333

<sup>3</sup><https://gitlab.inria.fr/tyrex/kegnn>

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

Table (1) Average test accuracy of 50 runs (10 for Flickr). The standard deviations are reported in brackets.

334 for MLP vs. KeMLP and GCN vs. KeGCN<sup>4</sup>. It is observed  
335 that KeMLP performs better compared to MLP as the node  
336 degree increases. By contrast, when comparing GCN and  
337 KeGCN, for both models, the accuracy increases for nodes  
338 with a higher degree. This shows that rich graph structure  
339 is helpful for the node classification in general. Indeed, the  
340 MLP is a simple model that misses information on the graph  
341 structure and thus benefits from graph structure in the form  
342 of binary predicates contributed by KeMLP. On the contrary,  
343 standalone GNNs can process graph structure by using mes-  
344 sage passing techniques to transmit learned node representa-  
345 tions between neighbors. The prior knowledge introduced in  
346 the knowledge enhancer is simple. It encodes that two neigh-  
347 bors are likely to be of the same class. An explanation for the  
348 small difference in performance is that GNNs may be able to  
349 capture and propagate this simple knowledge across neigh-  
350 bors implicitly, using its message passing technique. In other  
351 words we observe that, in this particular case, the introduced  
352 knowledge happens to be redundant for GNNs. However, the  
353 introduced knowledge significantly improves the accuracy of  
354 MLPs. In this context, we discuss perspectives for future  
355 work in Section 5.

### 356 Robustness to wrong knowledge

357 Furthermore, a question of interest is how the knowledge en-  
358 hanced model finds a balance between knowledge and graph  
359 data in case of knowledge that is not consistent with the  
360 graph data. In other words, can the KeGNN successfully deal  
361 with nodes having mainly neighbors that belong to a differ-  
362 ent ground truth class and thus contribute misleading infor-  
363 mation to the node classification? To analyze this question,  
364 we categorize the accuracy by the proportion of misleading  
365 nodes in the neighborhood, see Fig. 4. Misleading nodes are  
366 nodes that have a different ground truth class than the node  
367 to be classified. It turns out that KeMLP is particularly helpful  
368 over MLP when the neighborhood provides the right infor-  
369 mation. However, if the neighborhood is misleading (if most or  
370 even all of the neighbors belong to a different class), an MLP  
371 that ignores the graph structure can lead to even better results.  
372 When comparing KeGCN and GCN, there is no clear differ-  
373 ence. This is expected, since both models are equally affected  
374 by misleading nodes as they utilise the graph structure. Just as  
375 a GCN, the KeGCN is not necessarily robust to wrong prior  
376 knowledge since the GCN component uses the entire neigh-  
377 borhood, including the misleading nodes. When comparing

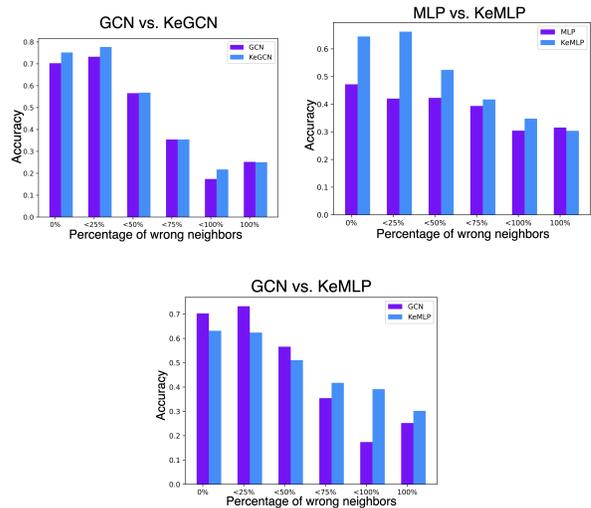


Figure (4) The accuracy grouped by the ratio of misleading first-order neighbors for GCN vs. KeGCN (left), MLP vs. KeMLP (right), GCN vs. KeMLP (below) on CiteSeer.

378 GCN to KeMLP, see plot below in Fig.4, KeMLP is more ro-  
379 bust to misleading neighbors. While GCN takes the graph  
380 structure as given and includes all neighbors equally in the  
381 embeddings by graph convolution, the clause weights in the  
382 knowledge enhancement methods provide a way to devalue  
383 knowledge. If the data frequently contradicts a clause, the  
384 model has the capacity to reduce the respective clause weight  
385 in the learning process and reduce its impact.

### 386 Clause Weight Learning

387 The clause weights learned during training provide insights  
388 on the updates made by a clause. The *clause compliance* (see

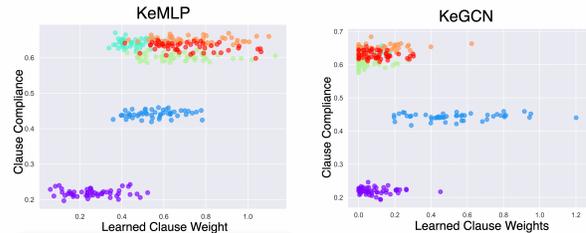


Figure (5) Learned clause weights vs. clause compliance for KeMLP (left) and KeGCN (right) on CiteSeer.

<sup>4</sup>The findings for KeGAT are in line with those for KeGCN, see Fig. 8 in Section A.3

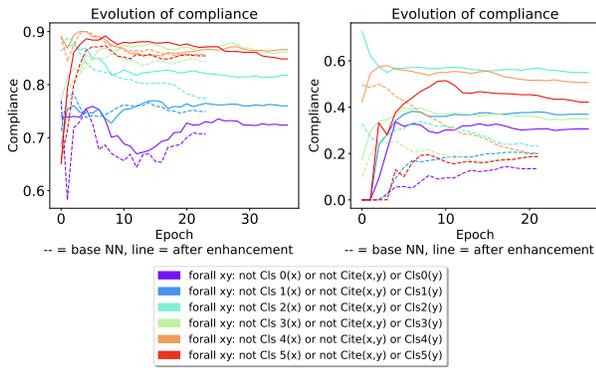


Figure (6) Clause compliance during training for GCN vs. KeGCN (left) and MLP vs. KeMLP (right) on Citeseer.

Adaptations are necessary on both the neural and the sym- 429  
 bolic side to apply KeGNN to heterogeneous graphs. The res- 430  
 triction to homogeneous graphs also limits the scope of for- 431  
 mulating complex prior knowledge. Eventually, the datasets 432  
 used in this work and the set of prior knowledge are too sim- 433  
 ple for KeGNN to exploit its potential and lead to a signifi- 434  
 cant improvement over the GNN. Experimental results show 435  
 that knowledge encoded by the symbolic component leads to 436  
 significant improvement over a model that is not capable 437  
 to capture and learn that knowledge. This indicates that for 438  
 more complex knowledge that is harder for a GNN to learn, 439  
 KeGNN has the potential to bring higher improvements. A 440  
 perspective for further work is the extension of KeGNN to 441  
 more generic data structures such as incomplete and hetero- 442  
 geneous knowledge graphs in conjunction with more com- 443  
 plex prior knowledge. 444

Another limitation of KeGNN is scalability. With an in- 445  
 creasing number of stacked knowledge enhancement layers, 446  
 the affected node neighborhood grows exponentially, which 447  
 can lead to significant memory overhead. This problem 448  
 is referred as neighborhood explosion [Duan *et al.*, 2022] 449  
 and is particularly problematic in the context of training on 450  
 memory-constrained GPUs. This affects both the GNN and 451  
 the knowledge enhancement layers that encode binary knowl- 452  
 edge. Methods from scalable graph learning [Fey *et al.*, 2021] 453  
 [Zeng *et al.*, 2020] [Hamilton *et al.*, 2017] represent po- 454  
 tential solutions for the neighborhood explosion problem in 455  
 KeGNN. 456

Furthermore, limitations appear in the context of link pre- 457  
 diction with KeGNN. For link prediction, a neural component 458  
 is required that predicts fuzzy truth values for binary predi- 459  
 cates. At present, KeGNN can handle clauses containing bi- 460  
 nary predicates, but their truth values are initialized with ar- 461  
 tificial predictions, where a high value encodes the presence 462  
 of an edge. This limits the application of KeGNN to datasets 463  
 for which the graph structure is complete and known a priori. 464

## 6 Conclusion 465

In this work, we introduced KeGNN, a neuro-symbolic model 466  
 that integrates GNNs with symbolic knowledge enhancement 467  
 layers to create a fully differentiable end-to-end model. This 468  
 allows the use of prior knowledge to improve node classi- 469  
 fication while exploiting the expressive representations of a 470  
 GNN. Experimental studies show that the inclusion of prior 471  
 knowledge has the potential to improve simple neural models 472  
 (as observed in the case of MLP). However, the knowledge 473  
 enhancement of GNNs is harder to achieve on the underly- 474  
 ing and limited benchmarks for which the injection of sim- 475  
 ple knowledge concerning local neighborhood is redundant 476  
 with the representations that GNNs are able to learn. Never- 477  
 theless, KeGNN has not only the potential to improve graph 478  
 completion tasks from a performance perspective, but also to 479  
 increase interpretability through clause weights. This work is 480  
 a step towards a holistic neuro-symbolic method on incom- 481  
 plete and noisy semantic data, such as knowledge graphs. 482

## 5 Limitations and Perspectives 424

The method of KeGNN is limited in some aspects, which we 425  
 present in this section. In this work, we focused on homoge- 426  
 neous graphs. In reality, however, graphs are often heteroge- 427  
 neous with multiple node and edge types [Yang *et al.*, 2022]. 428

Appendix B) [Daniele and Serafini, 2020] measures how well 389  
 the prior knowledge is satisfied in a graph. It can be calcu- 390  
 lated on the ground truth classes or the predicted classes. As 391  
 a reference, we measure the clause compliance based on the 392  
 ground truth labels in the training set. Fig. 5 displays the 393  
 learned clause weights for KeGCN and KeMLP versus the 394  
 clause compliance. For KeMLP, a positive correlation be- 395  
 tween the learned clause weights and the clause compliance 396  
 on the training set is observed. This indicates that higher 397  
 clause weights are learned for clauses that are satisfied in 398  
 the training set. Consequently, these clauses have a higher 399  
 impact on the updates of the predictions. In addition, the 400  
 clause weights corresponding to clauses with low compliance 401  
 values make smaller updates to the initial predictions. Ac- 402  
 cordingly, clauses that are rarely satisfied learn lower clause 403  
 weights during the training process. In the case of KeGCN, the 404  
 clause weights are predominantly set to values close to zero. This 405  
 is in accordance with the absence of a significant performance 406  
 gap between GCN and KeGCN. Since the GCN itself already 407  
 leads to valid classifications, smaller updates are required 408  
 by the clause enhancers. 409

Furthermore, we analyse how the compliance evolves dur- 410  
 ing training to investigate whether the models learn predic- 411  
 tions that increase the satisfaction of the prior knowledge. 412  
 Fig. 6 plots the evolution of the clause compliance for the 413  
 six clauses for GCN vs. KeGCN and MLP vs. KeMLP. It 414  
 is observed that GCN and KeGCN yield similar results as 415  
 the evolution of the compliance during training for both mod- 416  
 els is mostly aligned. For MLP vs. KeMLP the clause compli- 417  
 ance of the prediction of the MLP converges to lower values 418  
 for all classes than the clause compliance obtained with the 419  
 KeMLP. This gives evidence that the knowledge enhancement 420  
 layer actually improves the satisfiability of the prior knowl- 421  
 edge. As already observed, the GCN is also able to implicitly 422  
 satisfy the prior knowledge even though it is not explicitly 423  
 defined.

## 483 A Experiment Details

### 484 A.1 Implementation

485 The implementation of KeGNN and the described experi-  
 486 ments is publicly available on GitLab<sup>5</sup>. The code is based on  
 487 PyTorch [Paszke *et al.*, 2019] and the graph learning library  
 488 PyTorch Geometric [Fey and Lenssen, 2019]. The Weights  
 489 & Biases tracking tool [Biewald, 2020] is used to monitor the  
 490 experiments. All experiments are conducted on a machine  
 491 running an Ubuntu 20.4 equipped with an Intel(R) Xeon(R)  
 492 Silver 4114 CPU 2.20GHz processor, 192G of RAM and one  
 493 GPU Nvidia Quadro P5000.

### 494 A.2 Datasets

495 Tab. 2 gives an overview of the named datasets in this work.  
 496 The datasets are publicly available on the dataset collection<sup>6</sup>  
 497 of PyTorch Geometric [Fey and Lenssen, 2019]. For the split  
 498 into train, valid and test set, we take the predefined splits in  
 499 [Chen *et al.*, 2018] for the citation graphs and in [Zeng *et al.*,  
 500 2020] for Flickr. Word2Vec vectors [Adewumi *et al.*, 2020]  
 501 are used as node features for the citation graphs and image  
 502 data for Flickr. Fig. 1 visualizes the graph structure of the  
 503 underlying datasets in this work as a homogeneous, attributed  
 504 and labelled graph (on the example of Citeseer).

### 505 A.3 Results

506 The average test accuracies obtained for the node classifica-  
 507 tion experiments on Cora, Citeseer, PubMed and Flickr over  
 508 all tested models are visualized in Fig. 7. The average run-  
 509 times per epoch on the Citeseer dataset are compared for all  
 510 models in Tab 3. The runtimes were calculated for models  
 511 with three hidden layers and three knowledge enhancement  
 512 layers in full-batch training. It can be noted that the knowl-  
 513 edge enhancement layers lead to increased epoch times since  
 514 the model complexity is higher.

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.

515 Fig. 8 shows the accuracy grouped by node degree for GAT  
 516 vs. KeGAT.

<sup>5</sup><https://gitlab.inria.fr/tyrex/keggn>

<sup>6</sup><https://pytorch-geometric.readthedocs.io/en/latest/modules/datasets.html>

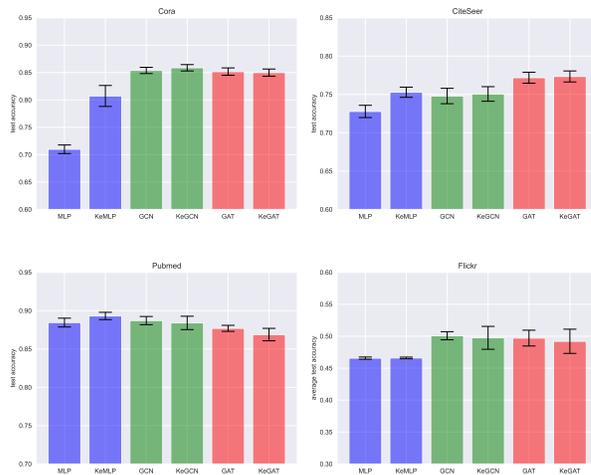


Figure (7) Average test accuracies over 50 runs for Cora, Citeseer and PubMed and 10 runs on Flickr. Error bars denote standard deviation.

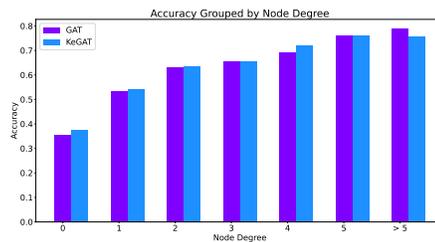


Figure (8) The accuracy grouped by the node degree for GAT vs. KeGAT on Citeseer.

### 517 A.4 Hyperparameter Tuning

518 KeGNN contains a set of hyperparameters. Batch normal-  
 519 ization [Ioffe and Szegedy, 2015] is applied after each hid-  
 520 den layer of the GNN. The Adam optimizer [Kingma and Ba,  
 521 2015] is used as optimizer for all models. Concerning the hy-  
 522 perparameters specific to the knowledge enhancement layers,  
 523 the initialization of the preactivations of the binary predicates  
 524 (which are assumed to be known) is taken as a hyperparam-  
 525 eter. They are set to a high positive value for edges that are  
 526 known to exist and correspond to the grounding of the bi-  
 527 nary predicate. Furthermore, different initializations of clause  
 528 weights and constraints on them are tested. Moreover, the  
 529 number of stacked knowledge enhancement layers is a hyper-  
 530 parameter. We further allow the model to randomly neglect  
 531 a proportion of edges by setting an edges drop rate param-  
 532 eter. Further, we test whether the normalization of the edges  
 533 with the diagonal matrix  $\tilde{\mathbf{D}} = \sum_j \tilde{\mathbf{A}}_{i,j}$  (with  $\tilde{\mathbf{A}} = \mathbf{A} + \mathbf{I}$ ) is  
 534 helpful.

535 To find a suitable set hyperparameters for each dataset and  
 536 model, we perform a random search with up to 800 runs and  
 537 48h time limit and choose the parameter combination which  
 538 leads to the highest accuracy on the validation set. The hyper-  
 539 parameter tuning is executed in Weights and Biases [Biewald,  
 540 2020]. The following hyperparameter values are tested:

- Adam optimizer parameters:  $\beta_1$ : 0.9,  $\beta_2$ : 0.99,  $\epsilon$ : 1e-07

517

518

519

520

521

522

523

524

525

526

527

528

529

530

531

532

533

534

535

536

537

538

539

540

541

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

- 542 • Attention heads: {1, 2, 3, 4, 6, 8, 10}
- 543 • Batch size: {128, 512, 1024, 2048, full batch}
- 544 • Binary preactivation: {0.5, 1.0, 10.0, 100.0, 500.0}
- 545 • Clause weights initialization: {0.001, 0.1, 0.25, 0.5, random
- 546 uniform distribution on [0,1]}
- 547 • Dropout rate: 0.5
- 548 • Edges drop rate: random uniform distribution [0.0, 0.9]
- 549 • Edge normalization: {true, false}
- 550 • Early stopping:  $\delta_{min}$  : 0.001, patience: {1, 10, 100}
- 551 • Hidden layer dimension: {32, 64, 128, 256}
- 552 • Learning rate: random uniform distribution [0.0001, 0.1]
- 553 • Clause weight clipping:  $w_{min}$  : 0.0,  $w_{max}$ : random uniform
- 554 distribution: [0.8, 500.0]
- 555 • Number of knowledge enhancement layers: {1, 2, 3, 4, 5, 6}
- 556 • Number of hidden layers: {2, 3, 4, 5, 6}

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

## 562 B Clause Weight Evaluation

563 The *clause compliance* [Daniele and Serafini, 2020] indicates  
564 the level of satisfaction of a clause in the data in this experi-  
565 mental setting. Given a clause  $\varphi$ , a class  $\text{Cls}_m$ , the set of  
566 training nodes  $\mathbf{V}_{\text{train}}$ , the set of nodes of the class  $\text{Cls}_m$ :  
567  $\mathbf{V}_m = \{v_i | v_i \in \mathbf{V}_{\text{train}} \wedge \text{Cls}(v_i) == m\}$ , and the neigh-  
568 borhood  $\mathcal{N}(v_i)$  of  $v_i$ , the clause compliance on graph  $\mathbf{G}$   
569 is defined as follows:

$$\text{Compliance}(\mathbf{G}, \varphi) = \frac{\sum_{v_i \in \mathbf{V}_k} \sum_{v_j \in \mathcal{N}(v)} \mathbf{1}[\text{if } v_j \in \mathbf{V}_m]}{\sum_{v_i \in \mathbf{V}_m} |\mathcal{N}(v_i)|} \quad (1)$$

570

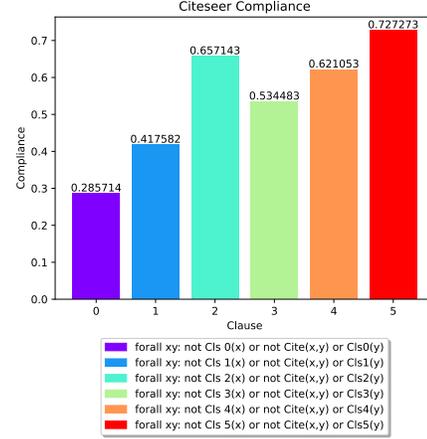


Figure (9) The clause compliance on the ground truth graph on the training set for Citeseer.

In other words, the clause compliance counts how often  
among nodes of a class the neighboring nodes are of the same  
class. [Daniele and Serafini, 2020]

571  
572  
573

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

574 **References**

- 575 [Abboud and Ceylan, 2021] Ralph Abboud and Ismail Ilkan  
576 Ceylan. Node classification meets link prediction  
577 on knowledge graphs. <https://arxiv.org/abs/2106.07297>,  
578 2021.
- 579 [Adewumi *et al.*, 2020] Tosin P. Adewumi, Foteini Liwicki,  
580 and Marcus Liwicki. Word2vec: Optimal hyper-  
581 parameters and their impact on nlp downstream tasks.  
582 <https://arxiv.org/abs/2003.11645>, 2020.
- 583 [Badreddine *et al.*, 2022] Samy Badreddine, Artur d'Avila  
584 Garcez, Luciano Serafini, and Michael Spranger. Logic  
585 tensor networks. *Artificial Intelligence*, 303:103649, feb  
586 2022.
- 587 [Biewald, 2020] Lukas Biewald. Experiment tracking with  
588 weights and biases. <https://www.wandb.com/>, 2020. Soft-  
589 ware available from wandb.com.
- 590 [Chen *et al.*, 2018] Jie Chen, Tengfei Ma, and Cao Xiao.  
591 Fastgcn: Fast learning with graph convolutional networks  
592 via importance sampling. In *ICLR (Poster)*. OpenRe-  
593 view.net, 2018.
- 594 [Dai *et al.*, 2020] Yuanfei Dai, Shiping Wang, Neal N.  
595 Xiong, and Wenzhong Guo. A survey on knowledge graph  
596 embedding: Approaches, applications and benchmarks.  
597 *Electronics*, 9(5), 2020.
- 598 [Daniele and Serafini, 2020] Alessandro Daniele and Lu-  
599 ciano Serafini. Neural networks enhancement with logical  
600 knowledge. <https://arxiv.org/abs/2009.06087>, 2020.
- 601 [Daniele and Serafini, 2022] Alessandro Daniele and Lu-  
602 ciano Serafini. Knowledge enhanced neural networks  
603 for relational domains. <https://arxiv.org/abs/2205.15762>,  
604 2022.
- [DeLong *et al.*, 2023] Lauren Nicole DeLong, Ra- 605  
mon Fernández Mir, Matthew Whyte, Zonglin 606  
Ji, and Jacques D. Fleuriot. Neurosymbolic ai 607  
for reasoning on graph structures: A survey. 608  
<https://arxiv.org/abs/2302.07200>, 2023. 609
- [Dou *et al.*, 2015] Dejing Dou, Hao Wang, and Haishan Liu. 610  
Semantic data mining: A survey of ontology-based ap- 611  
proaches. In *Proceedings of the 2015 IEEE 9th Interna- 612  
tional Conference on Semantic Computing (IEEE ICSC 613  
2015)*, pages 244–251, 2015. 614
- [Duan *et al.*, 2022] Keyu Duan, Zirui Liu, Peihao Wang, 615  
Wenqing Zheng, Kaixiong Zhou, Tianlong Chen, Xia Hu, 616  
and Zhangyang Wang. A comprehensive study on large- 617  
scale graph training: Benchmarking and rethinking. In 618  
*Thirty-sixth Conference on Neural Information Processing 619  
Systems Datasets and Benchmarks Track*, 2022. 620
- [Fatemi *et al.*, 2019] Bahare Fatemi, Siamak Ravanbakhsh, 621  
and David Poole. Improved knowledge graph embedding 622  
using background taxonomic information. *Proceedings of 623  
the AAAI Conference on Artificial Intelligence*, 33:3526– 624  
3533, 07 2019. 625
- [Fey and Lenssen, 2019] Matthias Fey and Jan E. Lenssen. 626  
Fast graph representation learning with PyTorch Geomet- 627  
ric. In *ICLR 2019 Workshop on Representation Learning 628  
on Graphs and Manifolds*, 2019. 629
- [Fey *et al.*, 2021] Matthias Fey, Jan E. Lenssen, Frank We- 630  
ichert, and Jure Leskovec. Gnnautoscale: Scalable and ex- 631  
pressive graph neural networks via historical embeddings. 632  
In Marina Meila and Tong Zhang, editors, *Proceedings 633  
of the 38th International Conference on Machine Learn- 634  
ing*, volume 139 of *Proceedings of Machine Learning Re- 635  
search*, pages 3294–3304. PMLR, 18–24 Jul 2021. 636
- [Guo *et al.*, 2016] Shu Guo, Quan Wang, Lihong Wang, Bin 637  
Wang, and Li Guo. Jointly embedding knowledge graphs 638  
and logical rules. In *Proceedings of the 2016 Conference 639  
on Empirical Methods in Natural Language Processing*, 640  
pages 192–202, Austin, Texas, November 2016. Associa- 641  
tion for Computational Linguistics. 642
- [Hamilton *et al.*, 2017] William L. Hamilton, Rex Ying, and 643  
Jure Leskovec. Inductive representation learning on large 644  
graphs. In *Proceedings of the 31st International Confer- 645  
ence on Neural Information Processing Systems, NIPS'17*, 646  
page 1025–1035, Red Hook, NY, USA, 2017. Curran As- 647  
sociates Inc. 648
- [Hu *et al.*, 2020] Yuwei Hu, Zihao Ye, Minjie Wang, Jiali 649  
Yu, Da Zheng, Mu Li, Zheng Zhang, Zhiru Zhang, and 650  
Yida Wang. Featgraph: A flexible and efficient backend 651  
for graph neural network systems. In *SC20: International 652  
Conference for High Performance Computing, Network- 653  
ing, Storage and Analysis*, pages 1–13, 2020. 654
- [Ioffe and Szegedy, 2015] Sergey Ioffe and Christian 655  
Szegedy. Batch normalization: Accelerating deep net- 656  
work training by reducing internal covariate shift. In 657  
Francis Bach and David Blei, editors, *Proceedings of the 658  
32nd International Conference on Machine Learning*, 659

- 660 volume 37 of *Proceedings of Machine Learning Research*,  
661 pages 448–456, Lille, France, 07–09 Jul 2015. PMLR.
- 662 [Jain *et al.*, 2021] Nitisha Jain, Trung-Kien Tran, Mo-  
663 hamed H. Gad-Elrab, and Daria Stepanova. Improving  
664 knowledge graph embeddings with ontological reasoning.  
665 In *The Semantic Web – ISWC 2021: 20th International*  
666 *Semantic Web Conference, ISWC 2021, Virtual Event, Oc-*  
667 *tober 24–28, 2021, Proceedings*, page 410–426, Berlin,  
668 Heidelberg, 2021. Springer-Verlag.
- 669 [Kautz, 2022] Henry A. Kautz. The third ai summer: Aai  
670 robert s. engelmore memorial lecture. <https://onlinelibrary.wiley.com/doi/10.1002/aaai.12036>, 2022.
- 672 [Kingma and Ba, 2015] Diederik P. Kingma and Jimmy Ba.  
673 Adam: A method for stochastic optimization. In Yoshua  
674 Bengio and Yann LeCun, editors, *3rd International Con-*  
675 *ference on Learning Representations, ICLR 2015, San*  
676 *Diego, CA, USA, May 7-9, 2015, Conference Track Pro-*  
677 *ceedings*, 2015.
- 678 [Kipf and Welling, 2017] Thomas N. Kipf and Max Welling.  
679 Semi-Supervised Classification with Graph Convolutional  
680 Networks. In *Proceedings of the 5th International Confer-*  
681 *ence on Learning Representations, ICLR ’17*, 2017.
- 682 [Klement *et al.*, 2013] E.P. Klement, R. Mesiar, and E. Pap.  
683 *Triangular Norms*. Trends in Logic. Springer Netherlands,  
684 2013.
- 685 [Li *et al.*, 2023] Weidong Li, Rong Peng, and Zhi Li. Knowl-  
686 edge graph completion by jointly learning structural fea-  
687 tures and soft logical rules. *IEEE Transactions on Knowl-*  
688 *edge and Data Engineering*, 35(3):2724–2735, 2023.
- 689 [Liu *et al.*, 2021] Weiwen Liu, Yin Zhang, Jianling Wang,  
690 Yun He, James Caverlee, Patrick Chan, Daniel Yeung, and  
691 Pheng-Ann Heng. Item relationship graph neural networks  
692 for e-commerce. *IEEE Transactions on Neural Networks*  
693 *and Learning Systems*, PP:1–15, 03 2021.
- 694 [Ma and Tang, 2021] Yao Ma and Jiliang Tang. *Deep Learn-*  
695 *ing on Graphs*. Cambridge University Press, 2021.
- 696 [Meilicke *et al.*, 2019] Christian Meilicke,  
697 Melisachew Wudage Chekol, Daniel Ruffinelli, and  
698 Heiner Stuckenschmidt. Anytime bottom-up rule learning  
699 for knowledge graph completion. In *Proceedings of*  
700 *the 28th International Joint Conference on Artificial*  
701 *Intelligence, IJCAI’19*, page 3137–3143. AAAI Press,  
702 2019.
- 703 [Paszke *et al.*, 2019] Adam Paszke, Sam Gross, Francisco  
704 Massa, Adam Lerer, James Bradbury, Gregory Chanan,  
705 Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca  
706 Antiga, Alban Desmaison, Andreas Köpf, Edward Yang,  
707 Zach DeVito, Martin Raison, Alykhan Tejani, Sasank  
708 Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and  
709 Soumith Chintala. Pytorch: An imperative style, high-  
710 performance deep learning library, 2019.
- 711 [Qu *et al.*, 2021] Meng Qu, Junkun Chen, Louis-Pascal  
712 Xhonneux, Yoshua Bengio, and Jian Tang. {RNNL}ogic:  
713 Learning logic rules for reasoning on knowledge graphs.  
714 In *International Conference on Learning Representations*,  
715 2021.
- [Sanchez-Gonzalez *et al.*, 2018] Alvaro Sanchez-Gonzalez,  
Nicolas Heess, Jost Tobias Springenberg, Josh Merel,  
Martin Riedmiller, Raia Hadsell, and Peter Battaglia.  
Graph networks as learnable physics engines for inference  
and control. In Jennifer Dy and Andreas Krause, edi-  
tors, *Proceedings of the 35th International Conference on*  
*Machine Learning*, volume 80 of *Proceedings of Machine*  
*Learning Research*, pages 4470–4479. PMLR, 10–15 Jul  
2018.
- [Susskind *et al.*, 2021] Zachary Susskind, Bryce Arden,  
Lizy K. John, Patrick Stockton, and Eugene B. John.  
Neuro-symbolic AI: an emerging class of AI workloads  
and their characterization. *CoRR*, abs/2109.06133, 2021.
- [Veličković *et al.*, 2018] Petar Veličković, Guillem Cucurull,  
Arantxa Casanova, Adriana Romero, Pietro Liò, and  
Yoshua Bengio. Graph attention networks. In *Internat-*  
*ional Conference on Learning Representations*, 2018.
- [Wu *et al.*, 2020] Yongji Wu, Defu Lian, Yiheng Xu, Le Wu,  
and Enhong Chen. Graph convolutional networks with  
markov random field reasoning for social spammer detec-  
tion. *Proceedings of the AAAI Conference on Artificial*  
*Intelligence*, 34(01):1054–1061, Apr. 2020.
- [Wu *et al.*, 2021] Zonghan Wu, Shirui Pan, Fengwen Chen,  
Guodong Long, Chengqi Zhang, and Philip S. Yu. A  
comprehensive survey on graph neural networks. *IEEE*  
*Transactions on Neural Networks and Learning Systems*,  
32(1):4–24, jan 2021.
- [Wu *et al.*, 2022] Lingfei Wu, Peng Cui, Jian Pei, and Liang  
Zhao. *Graph Neural Networks: Foundations, Frontiers,*  
*and Applications*. Springer Singapore, Singapore, 2022.
- [Yang *et al.*, 2016] Zhilin Yang, William W. Cohen, and  
Ruslan Salakhutdinov. Revisiting semi-supervised learn-  
ing with graph embeddings. In *Proceedings of the 33rd*  
*International Conference on International Conference on*  
*Machine Learning - Volume 48, ICML’16*, page 40–48.  
JMLR.org, 2016.
- [Yang *et al.*, 2022] Xiaocheng Yang, Mingyu Yan, Shirui  
Pan, Xiaochun Ye, and Dongrui Fan. Simple and efficient  
heterogeneous graph neural network. <https://arxiv.org/abs/2207.02547>, 2022.
- [Zadeh, 1988] L.A. Zadeh. Fuzzy logic. *Computer*,  
21(4):83–93, 1988.
- [Zeng *et al.*, 2020] Hanqing Zeng, Hongkuan Zhou, Ajitesh  
Srivastava, Rajgopal Kannan, and Viktor Prasanna. Graph-  
saint: Graph sampling based inductive learning method.  
In *International Conference on Learning Representations*,  
2020.