
Ranking Structured Objects with Graph Neural Networks

Published at Discovery Science 2021

Clemens Damke

Heinz Nixdorf Institute
Paderborn University
cdamke@mail.upb.de

Eyke Hüllermeier

Institute of Informatics
University of Munich
eyke@ifi.lmu.de

ABSTRACT

Graph neural networks (GNNs) have been successfully applied in many structured data domains, with applications ranging from molecular property prediction to the analysis of social networks. Motivated by the broad applicability of GNNs, we propose the family of so-called RankGNNs, a combination of neural Learning to Rank (LtR) methods and GNNs. RankGNNs are trained with a set of pair-wise preferences between graphs, suggesting that one of them is preferred over the other. One practical application of this problem is drug screening, where an expert wants to find the most promising molecules in a large collection of drug candidates. We empirically demonstrate that our proposed pair-wise RankGNN approach either significantly outperforms or at least matches the ranking performance of the naïve point-wise baseline approach, in which the LtR problem is solved via GNN-based graph regression.

Keywords Graph-structured data · Graph neural networks · Preference learning · Learning to rank

1 Introduction

Bringing a set of objects o_1, \dots, o_N into a particular order is an important problem with many applications, ranging from task planning to recommender systems. In such domains, the criterion defining the underlying order relation \succeq typically depends on properties (features) of the objects (for example the price and quality of a product). If the sorting criterion (and hence the relation \succeq) is not explicitly given, one may think of inferring it from exemplary data, often provided in the form of a set of pair-wise orderings $o_i \succeq o_j$ (e.g., representing that the user prefers product o_i over product o_j). This gives rise to a machine learning task often referred to as LtR. Thus, the goal is to learn a general ordering strategy (preference model) from sample data of the above kind, which can then be used to sort any new (previously unseen) set of objects.

While existing state-of-the-art LtR approaches assume that objects o_i are represented by feature vectors $x_i \in \mathbb{R}^n$, in this paper, we will consider the LtR problem for another quite natural and practically important representation, namely the domain of finite graphs. Methods for learning to rank objects represented in the form of graphs can, for example, be used in applications such as drug screening, where the ranked objects are the molecular structures of drug candidates.

To support the ranking of structured objects such as graphs, existing LtR methods need to be adapted. Previously, Agarwal [1] has considered the problem of ranking the vertices within a given graph. However, to the best of our knowledge, the graph-wise LtR problem has so far only been described in the context of specific domains, such as drug discovery, where manually chosen graph feature representations were used [2]. Motivated by the success of GNNs in graph representation learning, we propose a simple architecture that combines GNNs with neural LtR approaches. The proposed approach allows for training ranking functions in an end-to-end fashion and can be applied to arbitrary graphs without the need to manually choose a domain-specific graph feature representation.

Our neural graph ranking architecture will be introduced in Section 4. Before, the LtR and GNN models that are used in this architecture are described in Section 2 and Section 3, respectively. In Section 5, we evaluate our approach on a selection of graph benchmark datasets.

2 Object Ranking

LtR approaches are often categorized as point-wise, pair-wise, and list-wise methods. We begin with a short overview of these families. Afterwards, a more in-depth introduction is given to a selection of neural pair-wise approaches that we shall built upon in Section 4.

2.1 Overview of LtR Approaches

Point-wise methods assume the existence of a (latent) utility function representing the sought preference relation \succeq , i.e., that an ordinal or numeric utility score $u_i \in \mathbb{R}$ can be assigned to each object $o_i \in O$ such that

$$\forall o_i, o_j \in O : u_i \geq u_j \Leftrightarrow o_i \succeq o_j.$$

Based on training data in the form of exemplary (and possibly noisy) ratings, i.e., object/utility pairs $\{(x_i, u_i)\}_{i=1}^N \subset X \times \mathbb{R}$, where $x_i \in X$ is the feature representation of o_i , the LtR problem can be solved by fitting a model $f_u : X \rightarrow \mathbb{R}$ using standard ordinal or numeric regression methods. Given a new set of objects $\{o'_j\}_{j=1}^M$ to be ranked, these objects are then sorted in decreasing order of their estimated utilities $f_u(o'_j)$. Note that point-wise methods are restricted to linear orders but cannot represent more general relations, such as partial orders.

Pair-wise methods proceed from training data in the form of a set of ordered object pairs $S = \{o_{a_i} \succeq o_{b_i}\}_{i=1}^N$, i.e., relative training information in the form of pair-wise comparisons rather than absolute assessments. Based on such training samples S , the goal is to learn the underlying preference relation \succeq . The resulting model $f_{\succeq} : O \times O \rightarrow \{0, 1\}$ is a binary classifier, which is supposed to return $f_{\succeq}(o_i, o_j) = 1$ iff $o_i \succeq o_j$.

One of the first pair-wise preference methods was the Ranking SVM [3]—essentially a standard support vector machine (SVM) trained on the differences between vector representations of object preference pairs. Later, Burges et al. [4] proposed the RankNet architecture, which is also trained using feature vector differences but uses a multilayer perceptron (MLP) instead of an SVM. Since then, multiple extensions of those approaches have been developed [5]. One commonality between all of them is their training optimization target, namely to minimize the number of predicted inversions, i.e., the number of pairs $o_i \succeq o_j$ with $f_{\succeq}(o_i, o_j) = 0$. An important difference between existing pair-wise approaches concerns the properties they guarantee for the learned preference relation; three properties commonly considered are

- reflexivity ($\forall x : x \succeq x$),
- antisymmetry ($\forall x, y : x \not\succeq y \Rightarrow y \succeq x$), and
- transitivity ($\forall x, y, z : (x \succeq y \wedge y \succeq z) \Rightarrow x \succeq z$).

The set of desirable properties depends on the domain. While some approaches guarantee that the learned relation fulfills all three properties [6], others, for example, explicitly allow for non-transitivity [7].

Assuming a suitable pair-wise ranking model f_{\succeq} was selected and trained, one then typically wants to produce a ranking for some set of objects $\{o'_i\}_{i=1}^M$. To this end, a ranking (rank aggregation) procedure is applied to the preferences predicted for all pairs (o'_i, o'_j) . A simple example of such a procedure is to sort objects o_i by their Borda count $c_i = \sum_{j \neq i} f_{\succeq}(o_i, o_j)$, i.e., by counting how often each object o_i is preferred over another object. Alternatively, the classifier f_{\succeq} can also be used directly as the comparator function in a sorting algorithm; this reduces the number of comparisons from $\mathcal{O}(M^2)$ to $\mathcal{O}(M \log M)$. While the latter approach is much more efficient, it implicitly assumes that f_{\succeq} is transitive. The rankings produced by an intransitive sorting comparator are generally unstable, because they depend on the order in which the sorting algorithm compares the objects [8]. This might not be desirable in some domains.

List-wise methods generalize the pair-wise setting. Instead of determining the ordering of object pairs, they directly operate on complete rankings (lists) of objects, training a model based on a list-wise ranking loss function. One of the first list-wise losses was proposed by Cao et al. [9]. Given a set S of objects, their ListNet approach uses a probability distribution over all possible rankings of S and is trained by minimizing the cross-entropy between the model’s current ranking distribution and some target distribution. Compared to pair-wise approaches, list-wise methods exhibit a higher expressivity, which can be useful to capture effects such as context-dependence of preferences [10]. In general, however, if this level of expressiveness is not required, recent results by Köppel et al. [6] suggest that the list-wise approaches have no general advantage over the (typically simpler) pair-wise methods. To tackle the graph LtR problem in Section 4, we will therefore focus on the pair-wise approach.

2.2 Neural Pair-wise Ranking Models

As already stated, we propose a combination of existing LtR methods and GNNs to solve graph ranking problems. Due to the large number of existing LtR approaches, we will however not

evaluate all possible combinations with GNNs, but instead focus on the following two representatives:

1. **DirectRanker [6]:** A recently proposed generalization of the already mentioned pair-wise RankNet architecture [4]. It guarantees the reflexivity, antisymmetry, and transitivity of the learned preference relation and achieves state-of-the-art performance on multiple common Ltr benchmarks.
2. **CmpNN [7]:** Unlike DirectRanker, this pair-wise architecture does not enforce transitivity. The authors suggest that this can, for example, be useful to model certain non-transitive voting criteria.

Formally, the DirectRanker architecture is defined as

$$f_{\succeq}^{\text{DR}}(o_i, o_j) := \sigma(w^\top (h(x_i) - h(x_j))), \quad (1)$$

where $x_i, x_j \in \mathbb{R}^n$ are feature vectors representing the compared objects o_i, o_j , the function $h : \mathbb{R}^n \rightarrow \mathbb{R}^d$ being a standard MLP, $w \in \mathbb{R}^d$ a learned weight vector and an activation function $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ such that $\sigma(-x) = -\sigma(x)$ and $\text{sign}(x) = \text{sign}(\sigma(x))$ for all $x \in \mathbb{R}$. One could, for example, use $\sigma = \tanh$ and interpret negative outputs of $f_{\succeq}^{\text{DR}}(o_i, o_j)$ as $o_j \succeq o_i$ and positive outputs as $o_i \succeq o_j$. This model can be trained in an end-to-end fashion using gradient descent with the standard binary cross-entropy loss. Note that f_{\succeq}^{DR} can be rewritten as $\sigma(f_u^{\text{DR}}(x_i) - f_u^{\text{DR}}(x_j))$, with $f_u^{\text{DR}}(x) := w^\top h(x)$. DirectRanker therefore effectively learns an object utility function f_u^{DR} and predicts $o_i \succeq o_j$ iff $f_u^{\text{DR}}(x_i) \geq f_u^{\text{DR}}(x_j)$. Thus, the learned preference relation f_{\succeq}^{DR} directly inherits the reflexivity, antisymmetry and transitivity of the \geq relation. The main difference between DirectRanker and a point-wise regression model is that DirectRanker learns f_u^{DR} indirectly from a set of object preference pairs. Consequently, DirectRanker is not penalized if it learns some order-preserving transformation of f_u^{DR} . We will come back to this point in Section 5.3.

Let us now look at the so-called Comparative Neural Network (CmpNN) architecture, which generalizes the DirectRanker approach. The main difference between both is that CmpNN does not implicitly assign a score $f_u(x_i)$ to each object o_i . This allows it to learn non-transitive preferences. CmpNNs are defined as follows:

$$\begin{aligned} f_{\succeq}^{\text{Cmp}}(o_i, o_j) &:= \sigma(z_{\succeq} - z_{\preceq}), \text{ with} & (2) \\ z_{\succeq} &:= \tau(w_1^\top z_1 + w_2^\top z_2 + b'), \quad z_1 := \tau(W_1 x_i + W_2 x_j + b), \\ z_{\preceq} &:= \tau(w_2^\top z_1 + w_1^\top z_2 + b'), \quad z_2 := \tau(W_2 x_i + W_1 x_j + b). \end{aligned}$$

Here, $w_1, w_2 \in \mathbb{R}^d$ and $W_1, W_2 \in \mathbb{R}^{d \times n}$ are shared weight matrices, b, b' bias terms, and σ, τ activation functions. Intuitively, $z_{\succeq} \in \mathbb{R}$ and $z_{\preceq} \in \mathbb{R}$ can be interpreted as weighted votes towards the predictions $o_i \succeq o_j$ and $o_j \succeq o_i$, respectively. A CmpNN will simply choose the alternative with the largest weight. The key idea behind the definitions in (2) is that the pairs z_{\succeq}, z_{\preceq} and z_1, z_2 will swap values when swapping the compared objects o_i, o_j . Consequently, f_{\succeq}^{Cmp} must be reflexive and antisymmetric [see 7]. If we set $W_1 = w_1 = 0$, the voting weights $z_{\succeq}, z_{\preceq} \in \mathbb{R}$ reduce to the predictions of a standard MLP h with the input o_i and o_j , respectively, i.e., $z_{\succeq} = h(x_i)$ and $z_{\preceq} = h(x_j)$. In this case, the CmpNN effectively becomes a DirectRanker model. By choosing non-zero weights for W_1 and w_1 , the model can however also learn non-transitive dependencies between objects. In fact, Rigutini et al. have shown that CmpNNs are able to approximate almost all useful pair-wise preference relations [7, Thm. 1].

3 Graph Neural Networks

Over the recent years, GNNs have been successfully employed for a variety of graph ML tasks, with applications ranging from graph classification and regression to edge prediction and graph synthesis. Early GNN architectures were motivated by spectral graph theory and the idea of learning eigenvalue filters of graph Laplacians [11, 12]. Those spectral GNNs take a graph $G = (V, E)$ with vertex feature vectors $x_i \in \mathbb{R}^n$ as input and iteratively transform those vertex features by applying a filtered version of the Laplacian L of G . Formally, the filtered Laplacian is defined as $\hat{L} = U^\top g(\Lambda)U$, where $L = U^\top \Lambda U$ is an eigendecomposition of L and g is a learned eigenvalue filter function that can amplify or attenuate the eigenvectors U . Intuitively, spectral GNNs learn which structural features of a graph are important and iteratively aggregate the feature vectors of the vertices that are part of a common important structural graph feature. Each of those aggregations is mathematically equivalent to a convolution operation. This is why they are referred to as (graph) convolution layers.

One important disadvantage of spectral convolutions is their computational complexity, making them especially unsuitable for large graphs. To overcome this limitation, Kipf and Welling proposed the so-called graph convolutional network (GCN) architecture [13], which restricts the eigenvalue filters g to be linear. As a consequence of this simplification, only adjacent vertices need to be aggregated in each convolution. Formally, the simplified GCN convolution can be expressed as follows:

$$x'_i = \sigma \left(W \left(\eta_{ii} x_i + \sum_{v_j \in \Gamma(v_i)} \eta_{ij} x_j \right) \right) \quad (3)$$

Here, $x_i, x'_i \in \mathbb{R}^d$ are the feature vectors of $v_i \in V$ before and after applying the convolution, $\Gamma(v_i)$ is the set of neighbors of v_i , $W \in \mathbb{R}^{d \times n}$ is a learned linear operator representing the filter g , σ some activation function, and $\eta_{ii}, \eta_{ij} \in [0, 1]$ normalization terms that will not be discussed here. After applying a series of such convolutions to the vertices of a graph, the resulting convolved vertex features can be used directly to solve vertex-level prediction tasks, e.g. vertex classification. To solve graph-level problems, such as graph classification or graph ranking, the vertex features must be combined into a single graph vector representation. This is typically achieved via a pooling layer, which could, for example, simply compute the component-wise mean or sum of all vertex features. More advanced graph pooling approaches use sorting or attention mechanisms in order to focus on the most informative vertices [14, 15].

Xu et al. [16] show that restricting the spectral filter g to be linear not only reduces the computational complexity but also the discriminative power of the GCN architecture. More precisely, they prove that any GNN using a vertex neighborhood aggregation scheme such as (3) can at most distinguish those graphs that are distinguishable via the so-called 1-dimensional Weisfeiler-Lehman (WL) graph isomorphism test [17]. GCNs do, in fact, have a strictly lower discriminative power than 1-WL, i.e., there are 1-WL distinguishable graphs, which will always be mapped to the same graph feature vector by a GCN model. In addition to this bound, Xu et al. [16] also propose the graph isomorphism network (GIN) architecture, which is able to distinguish all 1-WL distinguishable graphs. Recently, multiple approaches going beyond the 1-WL bound have been proposed. The so-called 2-WL-GNN architecture, for example, is directly based on the 2-

dimensional (Folklore) WL test [18]. Other current approaches use higher-order substructure counts [19] or so-called k -order invariant networks [20].

4 Neural Graph Ranking

To tackle the graph LtR problem, we propose the family of RankGNN models. A RankGNN is a combination of a GNN and one of the existing neural LtR methods. The GNN component is used to embed graphs into a feature space. The embedded graphs can then be used directly as the input for a comparator network, such as DirectRanker [6] or CmpNN [7]. Formally, a RankGNN is obtained by simply using a GNN to produce the feature vectors x_i, x_j in (1) and (2) for a given pair of graphs G_i, G_j . Since all components of such a combined model are differentiable, the proposed RankGNN architecture can be trained in an end-to-end fashion. Despite the simplicity of this approach, there are a few details to consider when implementing it; these will be discussed in the following sections.

4.1 Efficient Batching for RankGNNs

In the existing neural LtR approaches for objects o_i that are represented by features $x_i \in \mathbb{R}^n$, efficient batch training is possible by encoding a batch of k relations $\{o_{a_i} \succeq o_{b_i}\}_{i=1}^k$ with two matrices

$$A := \begin{pmatrix} x_{a_1} \\ \vdots \\ x_{a_k} \end{pmatrix} \in \mathbb{R}^{k \times n}, \quad B := \begin{pmatrix} x_{b_1} \\ \vdots \\ x_{b_k} \end{pmatrix} \in \mathbb{R}^{k \times n}$$

and using

$$Y := \begin{pmatrix} 1 \\ \vdots \\ 1 \end{pmatrix} \in \mathbb{R}^k$$

as the target prediction of the model. However, this approach is suboptimal in the graph LtR setting. Given the relations $\{G_0 \succeq G_1, G_1 \succeq G_2\}$, the graph G_1 would for example have to be encoded twice. When dealing with datasets that consist of possibly large graphs, such redundant encodings quickly become infeasible due to the additional memory and runtime requirements incurred by the GNN having to embed the same graph multiple times. To prevent this redundancy, each graph occurring on the left or the right side of a relation should instead only be encoded once as part of a single graph batch. This graph batch can be fed directly into a GNN to produce a matrix Z of graph feature embeddings. The individual graph relation pairs $G_i \succeq G_j$ can then be simply represented as pairs of indices (i, j) pointing to the corresponding feature vectors in the embedding matrix Z . Using those pointers, the graph vector representations for each pair can be looked up in Z . Figure 1 illustrates this idea.

4.2 Sorting Graphs with RankGNNs

After training a RankGNN model using a set of graph relation pairs, the model can be used to compare arbitrary graph pairs. Following the approach of Köppel et al. [6] and Rigutini et al. [7], a set of graphs can then be ordered by using the RankGNN as the comparator function in a standard sorting algorithm. We propose a simple parallelized quicksort-based scheme to sort graphs. When implementing a RankGNN model on a parallel compute device, such as a GPU,

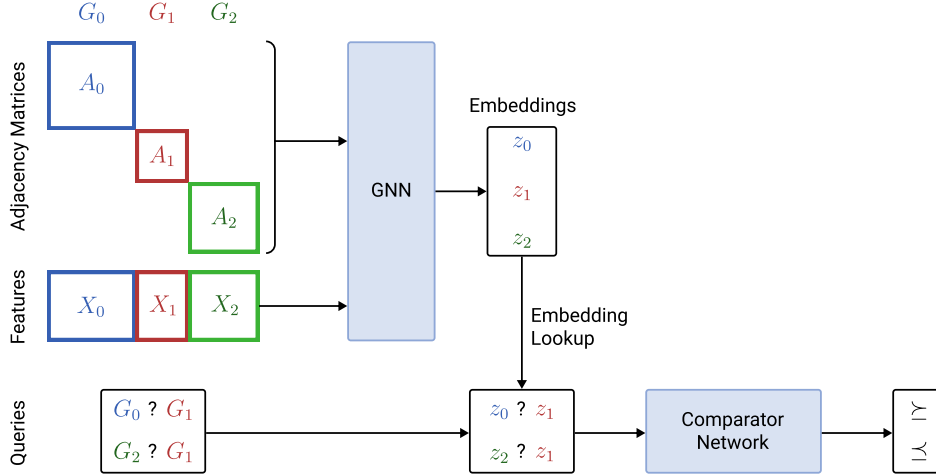


Figure 1: General architecture of the proposed family of RankGNNs. Here the common sparse adjacency representation for message-passing GNNs is shown; different types of graph batch encodings can of course also be used.

there is a constant time overhead for each call to the model. To reduce the total cost of this call overhead, we suggest that all pivot comparison queries in one layer of the recursive quicksort call tree should be evaluated by the RankGNN in parallel. Using this parallel comparison approach, only one model invocation is required for each layer of the call tree, i.e., the asymptotic model call overhead for sorting n graphs is in $O(\log n)$. Additionally, a more efficient approach is available for DirectRanker-based models. There, the implicitly learned utility function f_u^{DR} can be computed directly for a set of graphs. A standard sorting algorithm can then be applied without any further calls to the model, which reduces the call overhead to $O(1)$.

5 Evaluation

To evaluate the family of RankGNNs described in Section 4, we choose six different combinations of GNNs and comparator networks. The evaluated graph embedding modules are GCN [13], GIN [16], and 2-WL-GNN [18]. Those three GNN methods are combined with the previously described DirectRanker [6] and the CmpNN [7] comparator. Because there are currently no common graph ranking benchmark datasets, we instead convert a selection of graph regression benchmarks into ranking problems by interpreting the numeric regression targets as utility values, which are used to determine the target orderings. The following five graph regression datasets are used:

1. **TRIANGLES:** This is a synthetic dataset that we created. It consists of 778 randomly sampled graphs, each of which contains 3 to 85 unlabeled vertices. The regression target is to learn how many triangles, i.e. 3-cliques, a given graph contains. The triangle counts in the sampled graphs vary between 0 and 9. The sampled graphs are partitioned into 80%/10%/10% training/validation/test splits.

2. **OGB-molesol, -mollipo and -molreesolv:** These three datasets are provided as part of the Open Graph Benchmark (OGB) project [21]. They contain 1128, 4200, and 642 molecular structure graphs, respectively. The regression task is to predict the solubility of a molecule in different substances. We use the dataset splits that are provided by OGB.
3. **ZINC:** This dataset contains the molecular structures of 250k commercially available chemicals from the ZINC database [22]. The regression task is to predict the so-called octanol-water partition coefficients. We use the preprocessed and presplit graphs from the TUDataset collection [23].

To train the proposed pair-wise graph ranking network architecture, a subset of graph pairs from the training split is sampled uniformly at random. The size of a training sample is $M = \alpha N$, where N is the number of graphs in the training split of a dataset and $\alpha \in \mathbb{R}^+$ is a constant factor. We use a sampling factor of $\alpha = 20$ for all datasets except ZINC, where we use $\alpha = 3$ due to the large number of graphs in the training split ($N_{\text{ZINC}} = 220011$, whereas e.g. $N_{\text{OGB-mollipo}} = 3360$). This sampling strategy guarantees that each training graph occurs in at least one sampled pair with a probability of at least $1 - e^{-2\alpha}$; thus, for both $\alpha = 20$ and even $\alpha = 3$, all graphs are considered with high probability ($> 99.75\%$).

In addition to the six pair-wise RankGNN model variants, we also evaluate the ranking performance of standard point-wise GNN graph regression models, which are trained directly on graph utility values. We use two different target graph utilities: The original regression target $y_i \in \mathbb{R}$ for each training graph G_i , and the normalized graph rank $\bar{r}_i \in [0, 1]$, i.e. the normalized ordinal index of each training graph G_i when sorted by y_i .

5.1 Experimental Setup

We evaluate the performance of the different RankGNN variants via Kendall’s τ_B rank correlation coefficient. Given two graph rankings $r_1 : \mathcal{G} \rightarrow \mathbb{N}$, $r_2 : \mathcal{G} \rightarrow \mathbb{N}$, this coefficient is defined as

$$\tau_B := \frac{C - D}{\sqrt{(C + D + T_1)(C + D + T_2)}},$$

where C is the number of concordant pairs

$$\{\{G_i, G_j\} \mid i \neq j \wedge r_1(G_i) < r_1(G_j) \wedge r_2(G_i) < r_2(G_j)\},$$

D is the number of discordant pairs

$$\{\{G_i, G_j\} \mid i \neq j \wedge r_1(G_i) < r_1(G_j) \wedge r_2(G_i) > r_2(G_j)\},$$

and $T_{1,2}$ are the numbers of tied graph pairs, which have the same rank in r_1 and r_2 , respectively. Kendall’s τ_B rank coefficient ranges between -1 and $+1$, where $\tau_B = +1$ indicates that the two compared rankings are perfectly aligned, whereas $\tau_B = -1$ means that one ranking is the reversal of the other.

Another commonly used metric in the LtR literature is the normalized discounted cumulative gain (NDCG), which penalizes rank differences at the beginning of a ranking more than differences at the end. This is motivated by the idea that typically only the top- k items in a ranking are of interest. We do not employ the NDCG metric because this motivation does not hold for

the used target graph rankings. Since the target rankings are derived from regression targets, such as the water solubility of a molecule, both, the beginning and the end of a ranking are of interest and should therefore be weighted equally.

To train the evaluated point- and pair-wise models, we use the standard Adam optimizer [24]. The mean squared error (MSE) loss is used for the point-wise regression models, while the pair-wise variants of those GNNs are optimized via binary cross-entropy. All models were tuned via a simple hyperparameter grid search over the following configurations:

- 1. Layer widths:** $\{32, 64\}$. The width of both, the convolutional layers, as well as the fully-connected MLP layers that are applied after graph pooling.
- 2. Number of graph convolutions:** $\{3, 5\}$. A fixed number of two hidden layers was used for the MLP that is applied after the pooling layer.
- 3. Pooling layers:** $\{\text{mean, sum, softmax}\}$. Here, “mean” and “sum” refer to the standard arithmetic mean and sum operators, as described by Xu et al. [16], while “softmax” refers to the weighted mean operator described by Damke et al. [18].
- 4. Learning rates:** $\{10^{-2}, 10^{-3}, 10^{-4}\}$.

We used standard sigmoid activations for all models and trained each hyperparameter configuration for up to 2000 epochs with early stopping if the validation loss did not improve by at least 10^{-4} for 100 epochs. The configuration with the highest τ_B coefficient on the validation split was chosen for each model/dataset pair. To account for differences caused by random weight initialization, the training was repeated three times; 10 repeats were used for the TRIANGLES dataset due to its small size and fast training times. Note that, depending on the type of GNN, the pair-wise models can have between 3% and 10% more trainable weights than their point-wise counterparts, due to the added comparator network. All models were implemented in Tensorflow and trained using a single Nvidia GTX 1080Ti GPU. The code is available on GitHub¹.

5.2 Discussion of Results

Table 1 shows the ranking performance of the evaluated point- and pair-wise approaches on the test splits of the previously described benchmark datasets. Each group of rows corresponds to one of the three evaluated GNN variants. The first two rows in each group show the results for the point-wise models that are trained directly on the original regression targets and on the normalized ranks, respectively. The last two rows in each group hold the results for the pair-wise DirectRanker- and CmpNN-based models. Generally speaking, the pair-wise approaches either significantly outperform or at least match the performance of the point-wise regression models. The most significant performance delta between the point- and pair-wise approaches can be observed on the ZINC and OGB-molipo datasets. Only on the OGB-molesol dataset, the point-wise models achieve a slightly higher average τ_B value than the pair-wise models, which is however not significant when considering the standard deviations. Overall, we find that the pair-wise rank loss that directly penalizes inversions is much better suited for the evaluated graph ranking problems than the point-wise MSE loss.

¹<https://github.com/Cortys/rankgnn>

Table 1: Mean Kendall’s τ_B coefficients with standard deviations for the rankings produced by point- and pair-wise models on unseen test graphs.

		TRIANGLES	OGB-molesol	-mollipo	-molreesolv	ZINC
GCN	Utility regr.	0.273 \pm 0.004	0.706 \pm 0.001	0.232 \pm 0.002	0.015 \pm 0.242	0.547 \pm 0.414
	Rank regr.	0.172 \pm 0.040	0.702 \pm 0.012	0.224 \pm 0.006	0.446 \pm 0.038	0.823 \pm 0.005
	DirectRanker	0.234 \pm 0.002	0.714 \pm 0.003	0.327 \pm 0.006	0.483 \pm 0.011	0.879 \pm 0.002
	CmpNN	0.195 \pm 0.007	0.632 \pm 0.076	0.381 \pm 0.008	0.351 \pm 0.055	0.819 \pm 0.002
GIN	Utility regr.	0.469 \pm 0.056	0.729 \pm 0.006	0.353 \pm 0.052	0.243 \pm 0.125	0.790 \pm 0.003
	Rank regr.	0.481 \pm 0.014	0.717 \pm 0.011	0.310 \pm 0.017	0.495 \pm 0.021	0.827 \pm 0.011
	DirectRanker	0.502 \pm 0.028	0.712 \pm 0.007	0.429 \pm 0.026	0.439 \pm 0.065	0.894 \pm 0.012
	CmpNN	0.520 \pm 0.070	0.710 \pm 0.007	0.506 \pm 0.013	0.518 \pm 0.018	0.891 \pm 0.006
2-WL	Utility regr.	0.997 \pm 0.006	0.747 \pm 0.007	0.318 \pm 0.017	0.379 \pm 0.207	0.803 \pm 0.006
	Rank regr.	0.972 \pm 0.017	0.720 \pm 0.019	0.332 \pm 0.083	0.524 \pm 0.020	0.810 \pm 0.003
	DirectRanker	1.000 \pm 0.000	0.745 \pm 0.009	0.505 \pm 0.012	0.525 \pm 0.010	0.894 \pm 0.008
	CmpNN	1.000 \pm 0.000	0.718 \pm 0.020	0.503 \pm 0.010	0.527 \pm 0.064	0.873 \pm 0.002

Comparing the two evaluated variants of point-wise regression models, we find that the ones trained on normalized graph ranks generally either have a similar or significantly better ranking performance than the regression models with the original targets. We will come back to this difference in Section 5.3.

Looking at the results for the synthetic TRIANGLES dataset, we find that only the higher-order 2-WL-GNN is able to reliably rank graphs by their triangle counts. This is plausible, because architectures bounded by the 1-WL test, such as GCN and GIN, are unable to detect cycles in graphs [25]. While both the point- and the pair-wise 2-WL-GNN models achieve perfect or near-perfect τ_B scores on this task, the pair-wise approaches did perform more consistently, without a single inversion on the test graphs over 10 iterations of retraining.

Since the target graph rankings for all evaluated datasets are derived from regression values, all models have to learn a transitive preference relation. Consequently, the ability of CmpNN-based RankGNNs to learn non-transitive preferences is, in theory, not required to achieve optimal ranking performance. If the sample size of training graph pairs is too small, such that it contains few transitivity-indicating subsets, e.g. $\{G_1 \succeq G_2, G_2 \succeq G_3, G_1 \succeq G_3\}$, the higher expressiveness of CmpNNs could even lead to overfitting and therefore worse generalization performance compared to DirectRanker. Nonetheless, with the used sampling factor of $\alpha = 20$ (and $\alpha = 3$ for ZINC), each graph is, in expectation, sampled 40 times (6 for ZINC). This appears to be sufficient to prevent overfitting. In fact, the CmpNN-based RankGNNs perform very similarly to their DirectRanker-based counterparts. However, since DirectRanker-based models allow for a more efficient sorting implementation than CmpNN-based ones (cf. Section 4.2), we suggest the use of DirectRanker for problems where transitivity can be assumed.

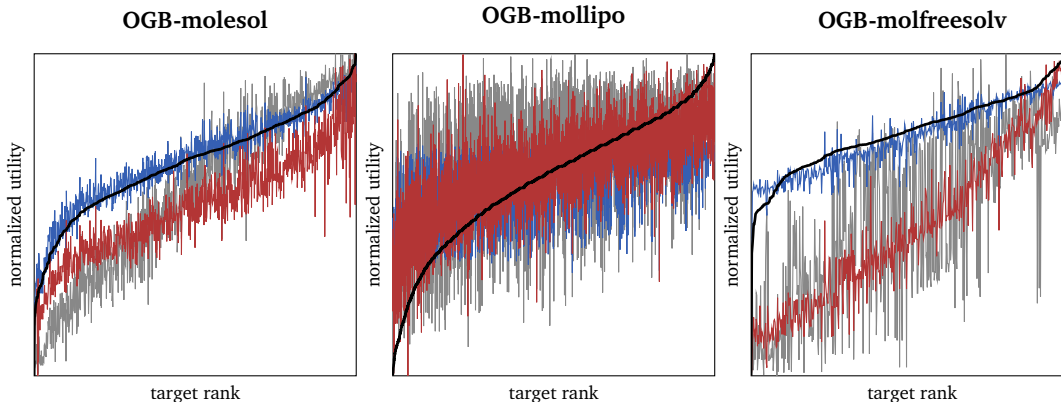


Figure 2: Normalized learned utility values of the point-wise GNN regression model trained on the original utilities (in blue), the point-wise model trained on normalized ranks (in gray) and the pair-wise DirectRanker model (in red). For each dataset, we plot the predicted utilities of the GNN architecture that achieved the best point-wise ranking performance in Table 1, i.e. 2-WL-GNN for OGB-molesol and -molfreesolv and GIN for OGB-mollipo. Each point along the horizontal axes corresponds to a graph in the training split of a dataset. The graphs are sorted in ascending order by the ground truth utility values (shown in black) from which the target rankings are derived.

5.3 Analysis of the Implicit Utilities of DirectRanker GNNs

As described in Section 2.2, a DirectRanker model $f_{\Sigma}^{\text{DR}} : O \times O \rightarrow \{0, 1\}$ implicitly learns a utility function $f_u^{\text{DR}} : O \rightarrow \mathbb{R}$ from the set of pairs it sees during training. We will now take a closer look at this implicitly learned utility function f_u^{DR} and compare it to the explicitly learned utilities f_u^{util} and f_u^{rank} of the point-wise GNN regression models. Figure 2 shows the values of all three, f_u^{DR} (in red), f_u^{util} (in blue) and f_u^{rank} (in gray), normalized to the unit interval. Any monotonically increasing curve corresponds to a perfect ranking ($\tau_B = +1$), while a monotonically decreasing curve would signify an inverse ranking ($\tau_B = -1$).

As expected, the blue utility curves of the point-wise approaches align with the black target utility curves, while the gray curve more closely follows the 45° diagonal line on which the normalized graph ranks would lie. However, this alignment does not necessarily imply good ranking performance. For example, on the OGB-molfreesolv dataset, the blue utility curve of the point-wise 2-WL-GNN model fits the black target curve fairly well for the graphs in the middle of the ranking. However, near the low and the high graph ranks, the target curve abruptly falls/rises to its minimum and maximum values; the point-wise regression model that is trained on the original utilities ignores those outliers. By instead training a point-wise model on the normalized ranks, outliers in the original utility values are effectively smoothed out, as can be seen in the gray OGB-molfreesolv utility curve. Looking at Table 1, we find that this corresponds to a significantly higher mean τ_B coefficient and a lower variance on the OGB-molfreesolv dataset. The pair-wise DirectRanker-based approach solves the problem of outliers in a more general fashion. It uses a loss function that does not penalize for learning a monotonous, rank-preserving transformation

of the target utility curve. This allows it to effectively “stretch” the target utilities into a linearly growing curve with fewer abrupt changes, which results in a similar performance to that of the regression model trained on normalized ranks.

The target utilities of the OGB-molesol dataset are distributed more smoothly, without any outliers. There the advantage of approaches that work well with outliers (e.g. pair-wise models) over the ones that do not is less pronounced. Lastly, looking at the OGB-molliipo dataset, we also do not find outliers in the target utility curve. However, there the pair-wise RankGNN models perform significantly better than the point-wise approaches. The reason for this performance difference is not yet fully understood.

6 Conclusion

In this paper, we addressed the problem of learning to rank graph-structured data and proposed RankGNNs, a combination of neural pair-wise ranking models and GNNs. When compared with the naïve approach of using a point-wise GNN regression model for ranking, we found that RankGNNs achieve a significantly higher or at least similar ranking performance on a variety of synthetic and real-world graph datasets. We therefore conclude that RankGNNs are a promising approach for solving graph ranking problems.

There are various directions for future research. First, due to the lack of graph ranking benchmark datasets, we had to use graph regression datasets in our evaluation instead. For a more thorough analysis of the practical applicability of graph ranking models, a collection of real-world graph ranking benchmarks should be created. One potential benchmark domain could, for example, be the drug screening problem we described in the introduction, where the training data consists of drug candidate pairs ranked by a human expert.

Second, list-wise graph ranking approaches could be evaluated in addition to the point- and pair-wise models considered in this paper. Such list-wise models can be useful to learn a human’s individual preferences for structured objects, such as task schedules or organizational hierarchies, represented as directed acyclic graphs or trees, respectively. A list-wise ranking approach [e.g. 10] would be able to consider context-dependent preferences in such scenarios [26]. Yet another interesting idea, motivated by the behavior we observed for the point- and pair-wise 2-WL-GNN-based models on the OGB-molfreesolv dataset (cf. Figure 2), is a hybrid approach that combines regression and ranking, that is, point-wise and pair-wise learning [27].

Third, although graph neural networks are quite popular these days, the problem of graph ranking could also be tackled by well-established kernel-based methods. In the past, there has been a lot of work on graph kernels [28], making graph-structured data amenable to kernel-based learning methods. In principle, one may hence think of combining graph kernels with learning-to-rank methods such as RankSVM. However, our first experiences with an approach of that kind suggest that kernel-based approaches are computationally complex and do not scale sufficiently well, even for point-wise implementations— for larger data sets, the running time as well as the memory requirements are extremely high (which is also the reason why we excluded them from the experiments). Although they can be reduced using suitable approximation techniques, complexity clearly remains an issue. Besides, the ranking performance turned out to be rather

poor. For pair-wise approaches, not only the complexity further increases, but the problem also becomes conceptually non-trivial. This is because the simple reduction of ranking to classification, on which RankSVM is based, no longer works (this reduction takes differences between feature vectors, an operation that cannot be applied to graphs). Instead, a (preference) kernel function on pairs of pairs of objects, i.e. on quadruples, has to be used [29]. Nevertheless, this does of course not exclude the existence of more efficient (approximate) algorithms operating on kernel-representation for graphs.

References

- [1] Shivani Agarwal. Learning to rank on graphs. *Machine Learning*, 81(3):333–357, 2010. doi: 10.1007/s10994-010-5185-8.
- [2] Wei Zhang, Lijuan Ji, Yanan Chen, Kailin Tang, Haiping Wang, Ruixin Zhu, Wei Jia, Zhiwei Cao, and Qi Liu. When drug discovery meets web search: Learning to rank for ligand-based virtual screening. *J. Cheminf.*, 7(1), 2015.
- [3] Thorsten Joachims. Optimizing search engines using clickthrough data. In *Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining - KDD '02*. ACM Press, 2002. doi: 10.1145/775047.775067.
- [4] Chris Burges, Tal Shaked, Erin Renshaw, Ari Lazier, Matt Deeds, Nicole Hamilton, and Greg Hullender. Learning to rank using gradient descent. In *I*, 2005. doi: 10.1145/1102351.1102363.
- [5] Chris Burges. From RankNet to LambdaRank to LambdaMART: An overview. Technical Report MSR-TR-2010-82, Microsoft Research, 2010.
- [6] Marius Köppel, Alexander Segner, Martin Wagener, Lukas Pensel, Andreas Karwath, and Stefan Kramer. Pairwise learning to rank by neural networks revisited: Reconstruction, theoretical analysis and practical performance. In *ECML PKDD 2019*, volume 11908 of *Lecture Notes in Computer Science*, pages 237–252. Springer, 2019.
- [7] L. Rigutini, T. Papini, M. Maggini, and F. Scarselli. SortNet: Learning to rank by a neural preference function. *IEEE Transactions on Neural Networks*, 22(9):1368–1380, 2011. doi: 10.1109/tnn.2011.2160875.
- [8] A. El Mesaoudi-Paul, E. Hüllermeier, and R. Busa-Fekete. Ranking distributions based on noisy sorting. In *Proc. ICML 2018, 35th International Conference on Machine Learning*, pages 3469–3477, Stockholm, Sweden, 2018.
- [9] Zhe Cao, Tao Qin, Tie-Yan Liu, Ming-Feng Tsai, and Hang Li. Learning to rank. In *I*. ACM Press, 2007. doi: 10.1145/1273496.1273513.
- [10] Karlson Pfannschmidt, Pritha Gupta, and Eyke Hüllermeier. Deep architectures for learning context-dependent ranking functions, March 2018.
- [11] J. Bruna, W. Zaremba, A. Szlam, and Y. LeCun. *Spectral Networks and Locally Connected Networks on Graphs*, 2013.
- [12] Mikael Henaff, Joan Bruna, and Yann LeCun. *Deep Convolutional Networks on Graph-Structured Data*, 2015.
- [13] T. N. Kipf and M. Welling. Semi-supervised classification with graph convolutional networks. *ICLR*, 2017.

- [14] Muhan Zhang, Zhicheng Cui, Marion Neumann, and Yixin Chen. An end-to-end deep learning architecture for graph classification. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [15] Junhyun Lee, Inyeop Lee, and Jaewoo Kang. Self-attention graph pooling. In *ICML*, pages 6661–6670, 2019.
- [16] K. Xu, W. Hu, J. Leskovec, and S. Jegelka. How Powerful are Graph Neural Networks? In *ICLR*, 2019.
- [17] J. Cai, M. Fürer, and N. Immerman. An optimal lower bound on the number of variables for graph identification. *Combinatorica*, 12(4):389–410, 1992.
- [18] Clemens Damke, Vitalik Melnikov, and Eyke Hüllermeier. A novel higher-order Weisfeiler-Lehman graph convolution. In *Proceedings of the 12th Asian Conference on Machine Learning (ACML 2020)*, volume 129 of *Proceedings of Machine Learning Research*. PMLR, 2020.
- [19] Giorgos Bouritsas, Fabrizio Frasca, Stefanos Zafeiriou, and Michael M. Bronstein. Improving graph neural network expressivity via subgraph isomorphism counting, 2020.
- [20] Haggai Maron, Heli Ben-Hamu, Hadar Serviansky, and Yaron Lipman. Provably powerful graph networks. In *NeurIPS 2019*, pages 2153–2164, 2019.
- [21] Weihua Hu, Matthias Fey, Marinka Zitnik, Yuxiao Dong, Hongyu Ren, et al. Open graph benchmark: Datasets for machine learning on graphs, 2020.
- [22] Teague Sterling and John J. Irwin. ZINC 15 – ligand discovery for everyone. *Journal of Chemical Information and Modeling*, 55(11):2324–2337, 2015. doi: 10.1021/acs.jcim.5b00559.
- [23] Christopher Morris, Nils M. Kriege, Franka Bause, Kristian Kersting, Petra Mutzel, and Marion Neumann. TUDataset: A collection of benchmark datasets for learning with graphs, 2020.
- [24] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015.
- [25] M. Fürer. On the Combinatorial Power of the Weisfeiler-Lehman Algorithm. In *Lecture Notes in Computer Science*, pages 260–271. Springer International Publishing, 2017.
- [26] Joel Huber, John W. Payne, and Christopher Puto. Adding asymmetrically dominated alternatives: Violations of regularity and the similarity hypothesis. *Journal of Consumer Research*, 9(1):90, 1982. doi: 10.1086/208899.
- [27] D. Sculley. Combined regression and ranking. In *Proceedings of the 16th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, Washington, DC, USA, July 25-28, 2010, pages 979–988, 2010.
- [28] S.V.N. Vishwanathan, N.M. Schraudolph, R. Kondor, and K.M. Borgwardt. Graph kernels. *Journal of Machine Learning Research*, 11:1201–1242, 2010.
- [29] W. Waegeman, B. De Baets, and L. Boullart. Kernel-based learning methods for preference aggregation. *4OR*, 7:169–189, 2009.