Implications of Topological Imbalance for Representation Learning on Biomedical Knowledge Graphs

Stephen Bonner[1]*, **Ufuk Kirik**[1]*, **Ola Engkvist**[2], **Jian Tang**[3,4,5], **Ian P Barrett**[1]

[1]Data Sciences and Quantitative Biology, Discovery Sciences, R&D, AstraZeneca, Cambridge, UK
[2]Molecular AI, Discovery Sciences, R&D, AstraZeneca, Gothenburg, Sweden
[3]HEC Montreal, Canada
[4]Mila - Quebec AI Institute, Montreal, Canada
[5]Canadian Institute for Advanced Research (CIFAR), Canada

Abstract

Adoption of recently developed methods from machine learning has given rise to creation of drug-discovery knowledge graphs (KG) that utilize the interconnected nature of the domain. Graph-based modelling of the data, combined with KG embedding (KGE) methods, are promising as they provide a more intuitive representation and are suitable for inference tasks such as predicting missing links. One common application is to produce ranked lists of genes for a given disease, where the rank is based on the perceived likelihood of association between the gene and the disease. It is thus critical that these predictions are not only pertinent but also biologically meaningful. However, KGs can be biased either directly due to the underlying data sources that are integrated or due to modeling choices in the construction of the graph, one consequence of which is that certain entities can get topologically overrepresented. We demonstrate the effect of these inherent structural imbalances, resulting in densely-connected entities being highly ranked no matter the context. We provide support for this observation across different datasets, models as well as predictive tasks. Further, we present various graph perturbation experiments which yield more support to the observation that KGE models can be more influenced by the frequency of entities rather than any biological information encoded within the relations. Our results highlight the importance of data modeling choices, and emphasizes the need for practitioners to be mindful of these issues when interpreting model outputs and during KG composition.

1 Introduction

Drug discovery and development is a highly complex, lengthy and multi-disciplinary endeavour [11], associated with high costs and attrition, particularly due to efficacy failure [29] which is in turn underpinned by the hypothesis of which drug target(s) is thought to play a key role in a given disease. It is a process entailing many different critical decision points, for example which protein target(s) to focus on, what type of pharmacological agent to use to modulate a target, and which molecule to take into clinical trials. The biomedical domain is characterised by a high rate of technological evolution [16, 38], multitudes of low and high dimensional data types [13, 17], and a

---

*Authors contributed equally to this work

high volume and rate of semi-/unstructured information emerging constantly via scientific literature, conferences and patents.

These characteristics of the domain have led to exploration of Knowledge Graphs (KG) to improve representation of information, as well as delivery of predictions and hypotheses to improve the quality of the decisions made [5, 15]. In a drug discovery KG, nodes often represent entities such as genes, diseases or drugs, whilst the relations between them capture their interactions with each other. Graph representation of these relationships not only allows for a more intuitive navigation of the complex domain for exploratory tasks, but is also particularly suitable for inferential analysis as Knowledge Graph Embedding (KGE) models can be used to derive meaningful, lower-dimensional representations of these biological entities and relationships. These representations in the embedding space can then be used for predicting missing links between these entities. In the case of target prediction, such approaches would produce a ranked list of genes, ordered by the score from the embedding model, predicted to be associated with a given disease. The top k elements from this list would then be inspected via subject matter experts (SME) to be triaged and further refined into a smaller list for a more manageable target validation process [32]. For example, performing drug target identification has been addressed as link prediction between gene and disease entities using the ComplEx model [47] on a drug discovery graph [32].

In the drug discovery domain, working with KGs has the promise of countering cognitive biases by leveraging massive amounts of information, well-beyond what is humanly feasible to comprehend, to inform a given hypothesis or decision through a computationally sound process that is agnostic to *a priori* experience. Challenges remain however, e.g. dealing effectively with the structural biases in the information sources [31] and the relative lack of negative data for validation in particular. To the best of our knowledge, the precise effect of the structural imbalance inherent in biomedical KGs, on predictive tasks using KGE models has not been explored thoroughly in the literature.

In this study, we aim to fill in this gap by providing a comprehensive analysis of how structural imbalance in the underlying data in a KG affects the scores for predicted relationships. We demonstrate that the KGE methods tend to overrate the highly-connected entities, and that this result is reproducible in different datasets (i.e. graphs), for different embedding models, across different diseases and different predictive tasks. We go on to study the extent of this imbalance by further investigating the type of relationships the entities have, by introducing various types of perturbations to the graph to assess the stability of the rankings, as well as presenting a thorough case study in target discovery task set on the unperturbed graph. Seen altogether, our results would indicate that the total volume of connections an entity has within the graph seemingly matters more than any biological information encoded within. We thus highlight the need for careful consideration of graph composition in relation to the analytical methods employed for a given task.

Example code for the experiments presented in thus study is available online.[2] Results are presented on only publicly available datasets.

2 Background & Related Work

2.1 Background

Knowledge Graphs. A knowledge graph is a way to capture and represent, often widely scoped, knowledge within a domain, based on graph theory principles. Semantically, a KG is a heterogeneous, multi-relation and directed graph, connecting a set of entities $\mathcal{E}$ to each other by a set of relationships $\mathcal{R}$, defined as $\mathcal{K} \subseteq \mathcal{E} \times \mathcal{R} \times \mathcal{E}$ [53]. KGs are often considered as a series of triples $(h, r, t) \in \mathcal{K}$, where $h, t \in \mathcal{E}$ are the head and tail entities connected via the relationship $r \in \mathcal{R}$. A hypothetical triple from a drug discovery KG could be $(Gene_i, associates, Disease_j)$, where the entities $Gene_i$ and $Disease_j$ are connected via the relationship $associates$. Given that a KG is still a graph, its topological structure can be measured using a well studied range of techniques [30]. A common measure to consider at the level of individual entities is that of *degree*, which is the number of relationships of any type a given entity has to all others (possibly including itself) within the graph.

---

[2]https://github.com/AstraZeneca/biomedical-kg-topological-imbalance

In many real-world KGs, the set of triples can be noisy and incomplete [1]. Thus numerous techniques attempt to impute missing information based on the existing data in $\mathcal{K}$ through multi-relation link prediction [41]. Such techniques consider the partial triple $(Gene_i, associates, ?)$ and attempt to predict the correct tail entity, i.e. try to answer the question *'with which disease(s) is $Gene_i$ associated?'* or be given $(?, associates, Disease_j)$ and attempt to predict the correct head entity, i.e. *'which gene(s) are associated with $Disease_j$?'*

Knowledge Graph Embeddings. In this work we consider KG embedding models [21, 49], which aim to learn low-dimensional representations of all entities and relations. These embeddings are combined in various ways to produce a score of how likely a given triple is to be true, with a larger score typically implying a more plausible triple [1]. More concretely, a model $f : \mathcal{E} \times \mathcal{R} \times \mathcal{E} \to \mathbb{R}$, calculates a scalar value $s_{(h,r,t)}$ representing the plausibility for each potential triple $(h, r, t) \in \mathcal{K}$. For KGE approaches, f is a learned model which operates only with the embeddings of the elements in the triples, $f(\mathbf{h}, \mathbf{r}, \mathbf{t})$, where $\mathbf{h}, \mathbf{t} \in \mathbb{R}^m$ and $\mathbf{r} \in \mathbb{R}^n$. The values of m and n represent the dimension of the entity and relation embedding respectively. It should be noted that the values of the scaler scores $s_{(h,r,t)}$ are not directly meaningful by themselves and are generally used to rank true triples above a sampled set of negative ones. Furthermore, the scale of score values varies according to the function used by the particular KGE model [1], meaning comparing scores from different models side by side is meaningless.

Given a trained KGE model, it is possible to pass it an incomplete triple $(h, r, ?)$ as described earlier and it will predict which entity is most likely to complete it. In practice, this is achieved by enumerating over all entities in the graph and the model scoring the triple with each entity taking the place of the missing element: $\mathcal{S} = \{f(h, r, t') \mid t' \in \mathcal{E}\}$. This list of scores $\mathcal{S}$ is then sorted and a user can interpret it as a ranked list from most to least likely. Many of the common metrics, such as Hits@k or Mean Reciprocal Rank (MRR) for assessing the predictive performance of KGE models are based on this ranked output, with the hope that the model ranks true, but previously unknown triples, above meaningless ones [1].

One of the most attractive properties of KGE methods is that it is possible to query the same trained model for link prediction tasks no matter the exact context as the model is trained to reconstruct the entire graph, agnostic of a particular type of inference. In other words, it is possible to reuse the same model across different tasks, whether predicting the most likely genes to associate with a given disease, or the most likely drugs to interact with a given gene.

Drug Discovery and Knowledge Graphs. Drug discovery is driven by the need to improve the standard of care for a particular disease, either by providing a more efficacious treatment or one that has fewer, less severe side effects. A crucial step early in the process is target discovery, that is to correctly identify molecular entities that are associated with the disease or symptom in question, for which a drug can be designed and developed. These entities can be genes, proteins or metabolites, which play an important role in disease causation or progression and alternatively play an important role in managing the symptoms of the disease (e.g. while steroid nose sprays like mometasone do not affect the direct cause of allergy, they help alleviate symptoms like running nose). With the cost of drug discovery estimated at hundreds of millions of dollars [51], it is crucial to find the right target [29].

One use case for KG predictions is for them to be integrated into a complex drug discovery pipeline leading to resource intensive experiments being performed to validate the predictions [32]. In the case of target prediction, such approaches would produce a ranked list of genes, ordered by the score from the model, predicted to be associated with a given disease. The top k elements from this list would then be inspected via subject matter experts (SME) to be triaged and further refined into a smaller list for a more manageable target validation process. In one study, the top 600 genes predicted to be associated with rheumatoid arthritis were reduced down to 55 by SMEs to be taken forward for experimental validation [32]. It is therefore imperative that predictions are made using pertinent biological information so that viable targets can be discovered.

2.2 Previous Works

It has long been observed that graphs representing real-world data possess a long or heavy-tailed distribution of degree values [4]. This means that the majority of vertices have a limited number of connections to others, with a small number of highly connected hub-vertices being present. This

inherent connectivity imbalance present in KG data can pose challenges when using it for input to machine learning models, however to what extent it has an impact has not yet been fully explored.

The issue of biased training data has become an area of great interest within the KG field. Typically, the bias of concern is that of attributes associated with the entities in the graph, for example gender, age or race [48, 3, 8]. However, recent work [28] has shown how popularity bias is present in three frequently used non-biomedical KGs: FB15K [7], WN18 [7] and YAGO3-10 [27]. The work argues that traditional methods for addressing data imbalance, such as up-weighting or over-sampling minority classes, do not map well onto KG data where we operate in the unit of triples. By definition triples contain two linked entities, each of which might have a very different level of connectivity within the graph, so how to re-weight the training process of KG-specific models is unclear [28]. Instead, the authors propose alterations to the Hits@k and MRR ranking metrics to account for entity and relation popularity when scoring model performance.

The issue of non-uniform graph connectivity (typically in homogenous graphs) has begun to be studied in parallel by the field of Graph Neural Networks (GNN), where researchers have shown that models learn low-quality representations, and thus making more incorrect predictions for low-degree vertices [26, 25, 45]. This has also been explored in the context of homogenous graph representation learning [3] and for random walks [23, 37].

Other work has argued that degree can be a strong feature for predicting an edge between two entities in biomedical graphs, but this can cause issues when the graph representation does not truly reflect the real underlying connections between the entities [54]. The work derives a prior via a process of graph perturbations to measure the likelihood of an edge being present based solely on degree. However the impact of connectivity in relation to KGE models is not considered.

3 Experimental Setup

3.1 Dataset

Throughout this work we primarily make use of the publicly available Hetionet biomedical knowledge graph [19]. Hetionet was originally created as part of a project focusing on drug repurposing and integrates 29 public data sources, however its use has been explored in other areas such as target prediction [6]. Hetionet contains information on diseases, human protein-coding genes and compounds, among others, which are represented as entities within the graph and total over 47K. These are linked via 24 different relationship types capturing 2.2M interactions between them. An overview of Hetionet is presented in Figure 1, which shows the overall degree distribution (Figure 1a), as well as for just the gene entities (Figure 1b), both of which demonstrate a heavy-tailed distribution. The frequency of the entities (Figure 1c) and relations (Figure 1d) shows an imbalance, with certain ones being more prevalent than others[3].

As we focus on target discovery, that being the prediction of links between gene and disease entities, we make use of the $Disease \xrightarrow{associates} Gene$ (DaG) edges from Hetionet. These edges are extracted from data sources such as DISEASES [35], GWAS Catalog [10] and DisGeNET [34], all of which capture known gene-disease associations and come from a variety of sources, ranging from expert curation to text mining. Additionally, Hetionet also contains two other edge types between gene and disease entities: $Disease \xrightarrow{upregulates} Gene$ and $Disease \xrightarrow{downregulates} Gene$. This is important to consider as it could be trivial for a model to learn that if a gene-disease pair are linked via an up- or down-regulates edge, then they should also be linked via an associates edge, as it would be a common case logically.

3.2 Model

Many knowledge graph embedding models have been introduced, with the primary differentiator between them being how they score the plausibility of a given triple. Throughout this present study we make use of the TransE model which was one of the earliest approaches, and often still one of the most capable. TransE uses the notion of translational distance to learn embeddings such that relations are used to translate between entities in latent space [7]. In the TransE score function, the

---

[3]A detailed explanation of entity and relation types can be found in the original publication [19]

(a) Degree Distribution (b) Degree Distribution (Gene) (c) Entity Frequency
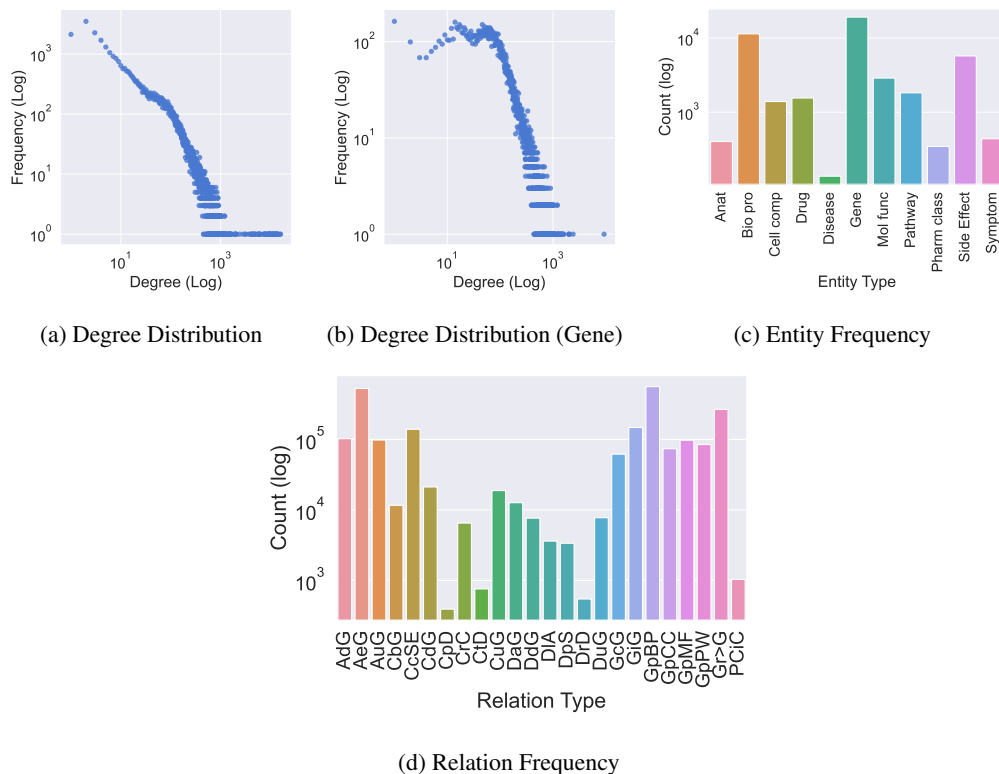

(d) Relation Frequency

Figure 1: An overview of the Hetionet dataset.

relation embedding is added to the head entity such that the result lies close to the tail embedding:

$$f(h, r, t) = -||\mathbf{h} + \mathbf{r} - \mathbf{t}||_F,$$

where F can be the l1 or l2 norm. One can interpret this score as simply the distance between the head and tail entity in the high-dimensional embedding space. Since its introduction, numerous other approaches have been proposed in the literature which build on or alter the TransE approach including TransH [50], ComplEx [47], RotatE [43] & DistMult [52], all with different strengths and weaknesses regarding relation types that can be captured [40]. However TransE has proven to still be highly competitive, when tuned appropriately, and can outperform more recent approaches [6].

Considering its comparable predictive performance and relatively modest computational requirements, the majority of our results refer to TransE, as the experiments we demonstrate in this study required a large number of training regimens. Nevertheless, we train and evaluate the performance of the above-mentioned newer models in the gene-disease prediction setting, and indeed observe the same trends regardless of the model in use (See Figure 6). Furthermore, all these models were trained using optimal hyperparameters previously established [6] on the same dataset.

3.3 Case Study Diseases

Not all entities are created equally and in our context this refers to the connectivity of diseases which naturally varies due to how well-studied a particular disease is. In order to assess how the ranking of DaG predictions vary due to the degree of the disease, we choose diseases on both ends of the connectivity scale. An overview of these five diseases and how they are connected within the graph is presented in Table 1.

On the upper end, we have Breast Cancer which is the most highly-connected disease in Hetionet, with a degree of 1159, out of which 540 are of type DaG. However, considering the heterogeneous nature of breast cancers and the diversity of both clinical and molecular subtypes [44] we also wanted to include other highly-connected diseases, with relatively more homogeneous aetiologies in the analysis. For this purpose, we chose melanoma and Parkinson's disease (PD) as representatives

5

of highly-studied diseases. Melanoma is a common type of cancer with relatively few subtypes and many known driver mutations [36]. Furthermore, it is the 4th best connected disease in the Hetionet graph, with a degree of 944 and 930 distinct neighbours. Parkinson's, on the other hand, is a complex neurodegenerative disease that has several monogenic variants, that is mutations in single genes causing the disease, as well as various identified genetic risk factors [18]. In Hetionet, Parkinson's disease has a degree of 795, and is the 8th most connected disease.

For low connectivity diseases, we chose Fuchs Endothelial Dystrophy (FED) and Fallopian Tube Cancer (FTC). Fuchs endothelial dystrophy, is a relatively rare disease with only 22 connections in the graph, all distinct neighbors. Besides being the 4th least connected disease in the graph, we found this disease to be of interest in terms of high level of characterisation and several known genetic factors as well as affected bioprocesses [14]. Fallopian tube cancer is a fairly rare type of cancer, with an incidence rate of about 0.36-0.41 per 100000 women [42]. Given its rarity, there is little in literature about this disease, especially in comparison to closely related ovarian and peritoneal carcinomas. In Hetionet, FTC has a degree of 29; connecting to 10 diseases, 8 symptom and anatomy nodes each, and most importantly to 3 genes.

Seen together these diseases provide a wide range of pathologies, both common and rare. Additionally, these diseases have high tissue-specificity and, in most cases, relatively low diversity in cell types and subtypes. Even in the case of FED and FTC, there are known genes with causative and/or risk-affecting mutations. These properties can be used in order to sanity check the predicted gene-disease associations to some extent.

Disease	Degree	# Distinct	# DaG	Connectivity
Breast Cancer	1159	1123	540	High
Melanoma	944	930	342	High
Parkinson's Disease	795	789	143	High
Fallopian Tube Cancer	29	29	3	Low
Fuchs Endothelial Dystrophy	22	22	6	Low

Table 1: The diseases we focus on in this study and associated metrics for them from Hetionet. Where # Distinct is the number of unique entities connected to the disease.

3.4 Implementation Details

All work has been performed using the open-source PyKEEN framework [2], a python library for knowledge graph embeddings built on top of PyTorch [33]. All experiments were performed on machines with Intel(R) Xeon(R) Gold 5218 CPUs and NVIDIA(R) V100 32GB GPUs. Additionally, we kept the software environment consistent throughout all experimentation using python 3.8, CUDA 10.1, PyTorch 1.9, and PyKEEN 1.5.0. Models were trained using optimal hyper-parameters, presented in Table 2, as discovered through a detailed optimisation process of over 100 unique parameter configurations per model [6]. Additionally, all models were trained on the same fixed random split of the Hetionet dataset.

Parameter	Value By Approach				
	ComplEx	DistMult	RotatE	TransE	TransH
Embedding Dim	272	80	512	304	480
Num Epochs	700	400	500	500	800
Learning Rate	0.03	0.02	0.03	0.02	0.005
Num Negatives	91	41	41	61	1
Optimiser	Adagrad				
Inverse Relations	False				
Loss Function	Negative Sampling Self-Adversarial Loss [43]				

Table 2: Model hyperparameters and training setups.

A note on runtime: The training of knowledge graph embedding models can be costly in regards to runtime on KGs the size of Hetionet [6]. For example, training a TransE model with the hyperparameters detailed in Table 2, took six hours on the previously detailed hardware. All other models demonstrated runtimes equal to, or often, greater than this. This has particular implications for the experiments where the KG is being perturbed in some way (i.e. the edge removal, addition or rewiring experiments) as each individual data point and each experiment repeat requires a full model retraining, meaning that these results took several hundred GPU hours to compute. This is ultimately why we limited this analysis to specific diseases, rather than repeating it over all available in Hetionet.

4 Results

In this section we present the results of our experimental evaluation. Unless otherwise stated, all experiments employ the $Disease \xrightarrow{associates} Gene$ relation type from Hetionet and make use of the TransE model as described in Section 3.

4.1 Correlation with Topological Structure

We begin by studying correlation between an entity's degree and the score assigned to it by the model measuring how likely it is to be the correct one to complete the triple.

As highlighted in Section 2.1, KGE models are essentially agnostic of the data schema when completing missing edges during inference time. As such they provide a score for each and every entry regardless of their type. Figure 2 displays the score assigned by TransE[4] to each entity in the Hetionet graph when completing the triple (Disease, DaG, ?) plotted against the degree for that entity, from which two main patterns can be seen. First and foremost, in both Melanoma (Figure 2a) and FED (Figure 2b) we can see a clear separation for Gene entities from everything else in the score space. This implies that the model is assigning a higher confidence to Genes being the correct entities to complete the triple. This is encouraging since we can conclude that the model at least is able to learn that the DaG relation only ever connects a Disease and a Gene in the training dataset.

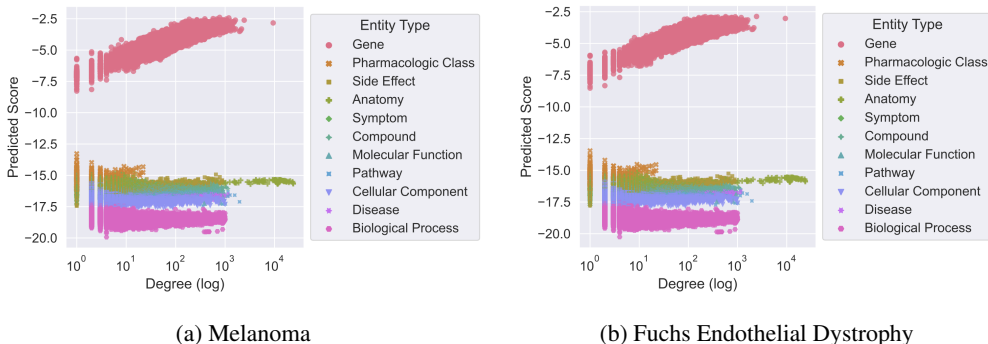(a) Melanoma   (b) Fuchs Endothelial Dystrophy

Figure 2: Predicted score for each entity using TransE when completing the partial triple (Disease, DaG, ?) versus the degree of the entity. Points are styled by entity type. The higher the score, the more likely the model considers the entity and being the correct one to complete the triple.

Having passed this first sanity check for the predictions, we thus move to focus on the relationship between the score assigned to just the Gene entities. The second pattern observed in Figure 2 was that there appears to be a correlation between degree and score. We now explore this in further detail for the various diseases explained in Section 3.3.

In Figure 3 the relationship between the score assigned to the gene entities by the TransE model and their degree is presented across both low and highly connected diseases. The points are coloured according to whether the particular entity was seen completing the triple in the train or test datasets,

---

[4]This score is the direct value taken from the TransE scoring function as detailed in Section 3.2, hence the negative value.

or whether it is novel (i.e. a DaG relationship not present in Hetionet between the given disease and that particular gene). The figures show there is a very clear relationship between the degree of the entity and the score assigned to it by the model – the higher the degree, the more likely the model considers the entity to complete the triple. This observation holds for both well-connected (Figures 3a and 3b) and sparsely connected diseases (Figures 3c and 3d). Indeed the pattern is remarkably similar, not only across all the diseases shown here in detail, but also for all 137 unique diseases that exist in Hetionet (Figure 4), suggesting that the connectivity of the gene has a large impact on whether it is likely to be predicted as associated with a disease, no matter which disease we investigate.

Another observation that can be made from Figure 3 is that many entities not seen during training have a higher score given by the model than those in the training data. This is especially true for the low connectivity diseases (Figures 3c and 3d) where the model ranks many genes with high degree over those with an actual connection in the training data. In other words, the model is more confident in a highly-connected gene than a true positive it has seen during training. While it is prudent not to draw big conclusions with so few data points in these diseases, it is nevertheless an important finding that a high-degree node can be ranked higher than a known positive. Indeed, the perceived wisdom would be that the model would over-fit to the training data and one would expect the training points to all be given high scores regardless of their degree, which does not appear to be the case.

(a) Melanoma

(b) Parkinsons

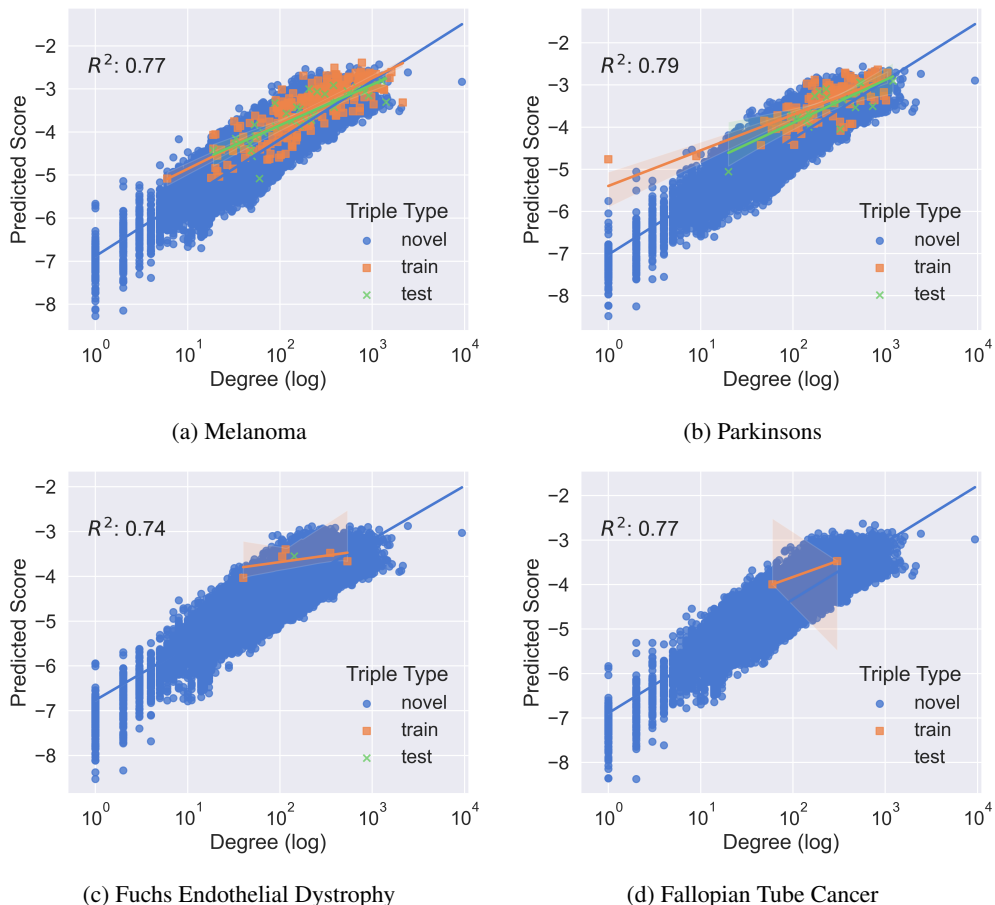(c) Fuchs Endothelial Dystrophy

(d) Fallopian Tube Cancer

Figure 3: Relationship between gene entity degree and the score assigned to it by TransE assessing the likelihood of it associating with the given disease.

Additionally, we wanted to explore how this apparent degree bias impacts the commonly used metrics for ranking entities. We demonstrate this by measuring the Mean Rank (MR) and Mean Reciprocal Rank (MRR), as well as Hits@k, with k=1 and k=10, of the tail entities for all DaG triples present in a holdout testset over 10 different dataset splits, as shown in Table 3. Results are sepa-

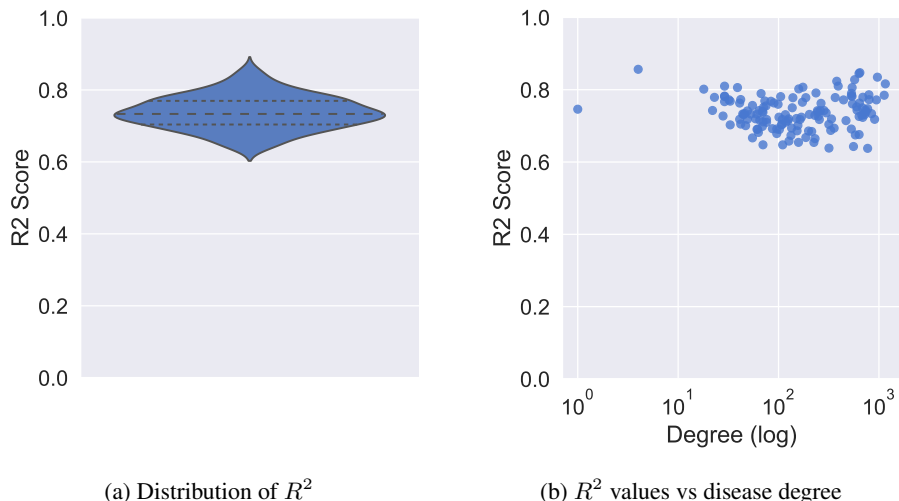(a) Distribution of R^2 (b) R^2 values vs disease degree

Figure 4: Correlation between the gene degree and score is consistent across all diseases in Hetionet as shown here. All diseases in Hetionet show high correlation coefficients between gene degree vs score, while the disease degree seems not to influence the results.

rated into high-degree (genes with degree > 1000) and low-degree (genes with degree < 200). The table shows that the model indeed demonstrates significantly lower performance, often by an order of magnitude, when evaluating on low-degree ground truth data.

Degree	Metric			
	MR $\downarrow$	MRR $\uparrow$	Hits@1 $\uparrow$	Hits@10 $\uparrow$
Low	4353.7±198.8	0.006±0.001	0.001±0.001	0.012±0.003
High	300.85 ±96.6	0.053±0.018	0.010±0.017	0.118±0.053

Table 3: Ranking metrics over high and low degree DaG edges from the holdout testsets. Arrows indicate if a higher or lower value is best for a given metric.

Correlation Across Diseases Figure 5 shows how the R-Squared value for the correlation between gene degree and predicted score across all 137 diseases in Hetionet, where the diseases have been categorised into families. The figure highlights how the strong correlation is present across the whole range of diseases in the dataset.

Different Models. Figure 6 highlights that this relationship between degree and score is not just a TransE specific issue but is prevalent across many knowledge graph embedding models used in the literature (including TransH [50], ComplEx [47], RotatE [43] & DistMult [52] which were trained using optimised hyper-parameters for Hetionet as detailed previously [6]). The figure also displays the Hits@10 score on a hold-out testset comprising 10% of the triples for all models. There appears to be little direct relationship between the Hits@10 score and the R-squared value, although the worse performing model ComplEx does have the weakest correlation between degree and the score assigned to the gene entities.

Different Datasets. Further, Figure 7 shows this issue not to be a dataset specific one as the same pattern can be found on two other biomedical knowledge graphs: the Drug Repurposing Knowledge Graph (DRKG) [20] and OpenBioLink [9]. Here the TransE model is used to predict the genes most likely to be associated with breast cancer. Figure 7 shows that a correlation between predicted score and degree is indeed present in both datasets, although OpenBioLink has a weaker correlation. Although we only present results for breast cancer for these datasets, we observed similar patterns across all of our chosen diseases.
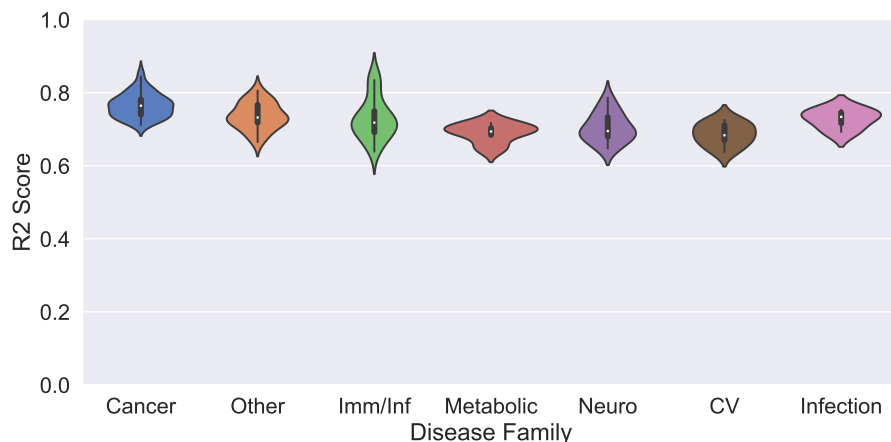
9

Figure 5: R-Squared value of the relationship between gene degree and score across all 137 diseases in Hetionet categorised into disease families. Disease families are Cancer, Immune/Inflammatory (Imm/Inf), Metabolic, Neurological (Neuro), Cardio-Vascular (CV), Infection and Other.

(a) TransH (0.11)   (b) ComplEx (0.07)   (c) RotatE (0.27)   (d) DistMult (0.14)

Figure 6: Relationship between gene degree versus predicted score for association with Breast Cancer. Numbers detail the Hits@10 score for that model on a holdout testset. Note that the change scale for the score is due to the different objective function used by the models.



(a) DRKG   (b) OpenBioLink

Figure 7: Relationship between the predicted score for gene association with breast cancer and degree across other biomedical knowledge graphs.

Different Tasks. We show that the issue is not specific to the task of gene-disease association prediction, but is replicated across many drug discovery tasks. Figure 8 highlights the relationship between score and entity degree for three other tasks, namely: Drug Target Interaction (DTI), Protein-Protein Interaction (PPI) and Drug Repurposing (DR). For the DTI task (Figures 8a-8c), we use the $Compound \xrightarrow{binds} Gene$ (CbG) edge types to rank genes by how likely they are to bind to the drugs Sunitinib, Auranofin and Doxorubicin. These compounds are selected in order to showcase different aspects of the dataset. Sunitinib and Doxorubicin are both chemotheraphy agents against cancer; the former is a receptor tyrosine kinase while the latter is an anthracycline topoisomerase inhibitor. Both drugs are well studied and thus well connected in the graph, Sunitinib has the most

10

CbG connections, while Doxorubicin is extremely well connected both in terms of overall degree, and specifically connections to genes. On the opposing end of the connectivity scale Auranofin has fewer overall connections in the graph, and very few CbG connections thus provides a good example when there is little information to train on.

In the PPI task (Figures 8d-8f), the $Gene \xrightarrow{interacts} Gene$ (GiG) edge types are used and genes are ranked based on likelihood to interact with the query genes MAPK1, PCNA and UBC. These are all fairly well studied genes, despite their apparent differences in terms of degree in Hetionet. An overwhelming majority of the connectivity of UBC is to other genes, specifically of the type GiG we consider here. PCNA on the other hand has again majority of its connections to other genes, however these are of a different relationship type. Lastly, MAPK1 has a much more balanced connectivity profile, in terms of diversity to other entities in the graph. Thus it provides a good test case to see if connectivity to other node types make any tangible contribution to the predictive performance in GiG relationships.

Finally, in the DR task (Figures 8g-8i), the $Compound \xrightarrow{treats} Disease$ edges are utilised to rank compounds by how likely the model considerers them to treat the diseases Rheumatoid arthritis, Fuchs endothelial dystrophy and Hypertension. This choice of diseases is motivated by two well-connected diseases of different aetiologies, both with various compounds for treatment options currently used in the clinic. The third disease here is a weakly connected one, without any compound connectivity and surgery as a primary means of treatment. Thus it provides a negative control for the predictive performance for the model. Last but not least, all three diseases are chosen to be distinct from various forms of cancer, as these disease dominate the overall connectivity, as well as to compounds in particular. Many cancer drugs are used for various indications and indeed a model could simply do well by extrapolating a well-connected compound to a well-connected form of cancer.

Across all tasks and example cases, we observe a strong correlation between score and degree highlighting that this is not a task-specific issue and is prevalent across many drug discovery problems.

4.2   Relations with Disease

We now consider whether the number of connections a gene has to disease entities within the training graph can affect ranking.

Figure 9 shows this analysis where melanoma is the query disease in question. Intuitively, one would expect a gene that has multiple connections to diseases to have a higher likelihood of being associated with another disease. This follows from the fact that in order to be associated with any disease in the first place, the gene would need to be characterized and often these associations are experimentally validated. Therefore, it is more likely that a well-characterized gene is associated with another disease, than a completely unknown gene is to be linked to a disease for the first time.

As expected, we found that the genes that have links to disease entities tend to have a higher predicted score (and thus higher rank) than those that have no disease links, as shown in Figure 9a. As mentioned above, gene-disease annotations are not uniformly distributed over genes. In other words, some genes have many disease annotations while others have few or none. Considering this, we wanted to check whether or not we can see a difference between high-ranking and low-ranking genes in terms of total number of connections to other diseases (Figure 9c).

Overall, we observed there is a connection between the number of edges to other diseases and the ranking, such that genes with a higher number of connections to other diseases, tend to have higher scores. However, not all high ranking genes had a large number of edges, which brings us to the question of whether it is the proportion of the edges of a gene to disease nodes that play a significant role. This did not turn out to be the case, as shown in Figure 9d there is little relationship between the proportion of edges a gene has to disease entities and the score assigned to it. However, ultimately it is hard to disentangle the impact of disease connections versus the total number, as highly connected genes are also more likely to be connected to diseases.

4.3   Trivial Relations

We now consider whether the presence of other relation types between a given gene and disease pair in the training graph affects score.
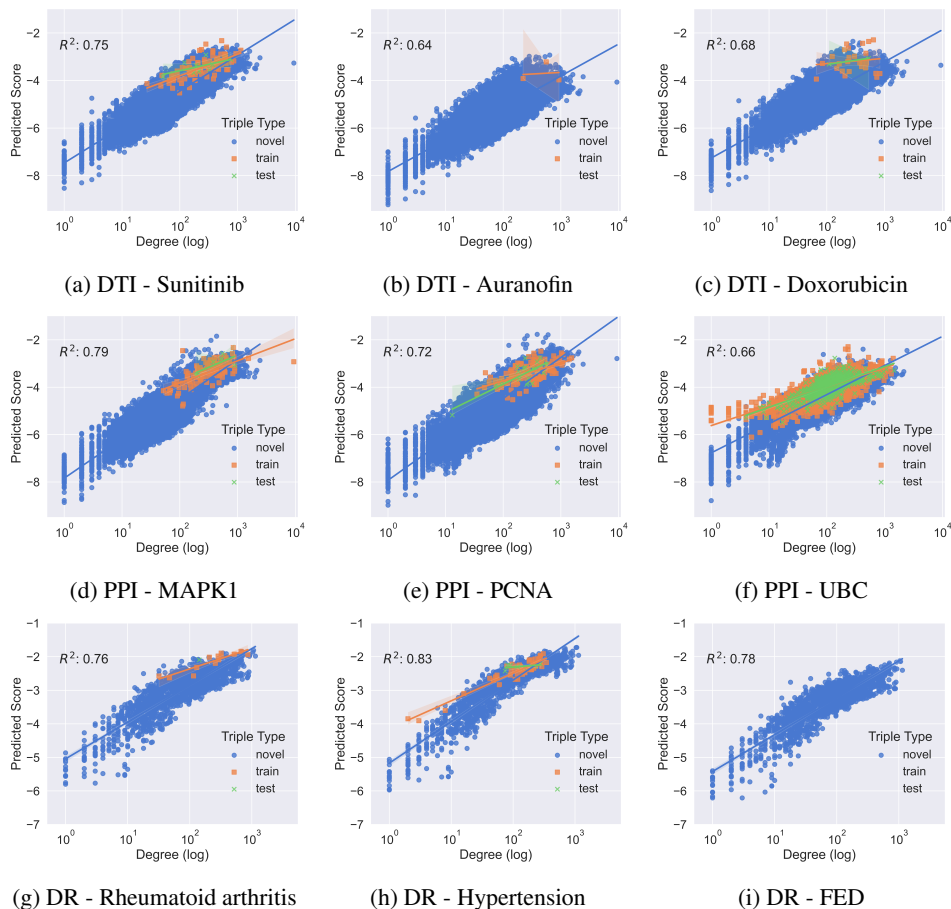
Figure 8: Relationship between the prediced score and degree for other drug discovery tasks; in a-c we query drug-target interactions for three compounds, d-f protein-protein interaction partners for three genes are queried, and in g-i we query for compounds for treatment of three diseases.

This is a subtle but important aspect, as multiple relationship types between entities in biological KGs like Hetionet typically allow for information leakage in the form of trivialized predictions. The presence of trivial relations, such as inverse ones, is known to result in over-optimistic models, which although appear as though they are demonstrating good predictive performance, in reality have exploited trivial patterns in the data [46, 12]. In Hetionet besides the DaG relation type we have used so far, there are also two other relationship types that are allowed between disease and gene entities, specifically $Disease \xrightarrow{upregulates} Gene$ and $Disease \xrightarrow{downregulates} Gene$. Figure 10 shows the relationship between score and degree, as before, however here we stratified the genes by whether they have another type of relationship to the queried disease, for melanoma and Parkinson's Disease.

Curiously, while the pattern of high correlation between degree and score is maintained for both groups, there is little difference between them in terms of ranking. In other words, the existence of another relation type between the same disease and gene does not appear to affect the associates relation. Especially interesting are the genes at the lower end of the degree distribution, where there are genes whose only edge is to the disease of interest which are still ranked lowly. We could not replicate this with the other two, sparsely-connected, diseases since they did not have multiple relation types between them and genes.
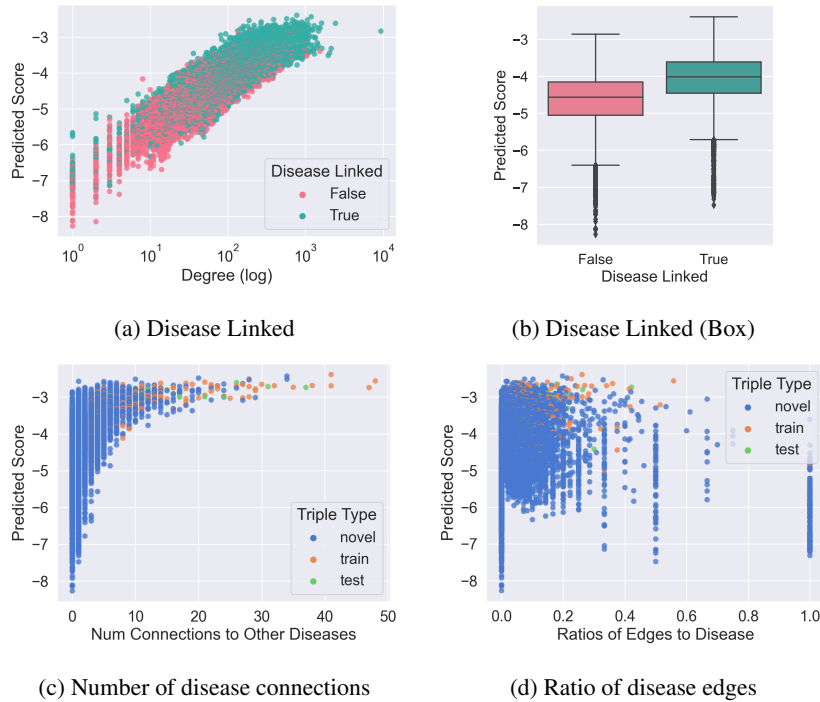
(a) Disease Linked

(b) Disease Linked (Box)

(c) Number of disease connections

(d) Ratio of disease edges

Figure 9: The relationship between the prediced score for Melanoma and various measures of connectivity between genes and the disease entities.
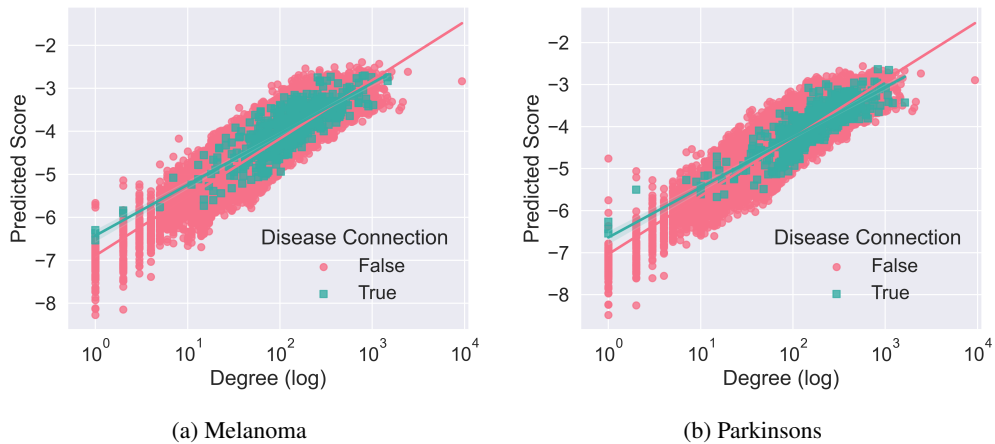

(a) Melanoma

(b) Parkinsons

Figure 10: Relationship between predicted gene-disease association score and entity degree. Points are coloured to indicate whether the gene has an edge of a different type to the disease of interest.

4.4 Graph Perturbations

In this section we aim to understand whether we can artificially influence the rank assigned to a gene by a model via a perturbing process whereby edges are removed or added, thus altering the graph topology. Any change in score relative to a disease could then be measured by retraining the model on this new graph. For all experiments presented in this section, 10 repeats are performed over different sets of random edges, with results being presented as the mean with 95% confidence intervals.

Edge Removal. In this experiment, a random selection edges are incrementally removed from the highest ranked novel gene for a given disease and any change in rank is measured. Two strategies are used for this removal process:

- *Disease:* edges are removed randomly from the gene to only other disease entities.
- *Random:* edges are removed randomly from the gene to any other entity.

Figure 11 shows how the rank of the top gene for a disease is affected as different proportions of its edges are removed. As can be seen in both panels, there is a clear relationship between the removal of edges and a decrease in rank. It should noted that this result should not be considered surprising, indeed the assumption would be that as knowledge is removed from the graph, the model is able to learn less about the gene and thus is able to make less confident predictions about it. What is more interesting however is that by removing disease connections, the rank is relatively unaffected, which is somewhat counter-intuitive.
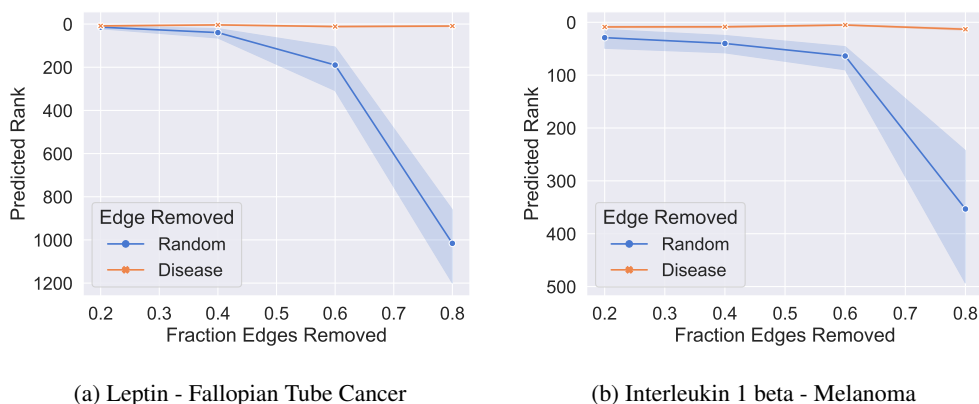

(a) Leptin - Fallopian Tube Cancer

(b) Interleukin 1 beta - Melanoma

Figure 11: Removel of edges from the top ranked novel gene for two diseases: Fallopian Tube Cancer and Melanoma

Edge Addition. In this experimental setup we add random, biologically meaningless edges to a gene to investigate how this impacts the score assigned to it by the model in relation to a set of diseases. To achieve this, we took the gene predicted to be the least likely to associate with a given disease and artificially increased its connectivity within the graph using two strategies:

- *Disease:* edges are added randomly from the gene to only other disease entities, excluding the disease we are scoring against. These edges conform to the graph schema and use the correct relation types of DaG, DuG and DdG.
- *Ant-Comp-Gene:* edges are added randomly from the gene to other gene, compound and anatomy entities. Again these edges conform to graph schema and use the correct relation types.

Figure 12a shows the results for Fuchs endothelial dystrophy, whilst Figure 12b the results for Breast Cancer. For both of these the gene to which the edges were added was T cell receptor gamma locus (TRG) as it was the least likely genes for both diseases. TRG has just a single edge in the original graph, connecting it to another gene. Overall the figures highlight the key observation that for both diseases, TRG was able to be moved from the least likely to one of the most likely gene to associate with the disease by the addition of biologically meaningless random edges. This finding provides further evidence that the predictions from KGE models can be biased by the frequency with which the entity is observed during the training process, and that this bias might outweigh any domain-specific knowledge captured in the data that may be useful to make more accurate real-world predictions.

4.5 Case Study: UBC

A consistent feature of the ranking vs degree plots is a gene to the top-right of the plot area, ranked curiously high no matter the disease in question, likely due to its particularly high degree. This gene

14

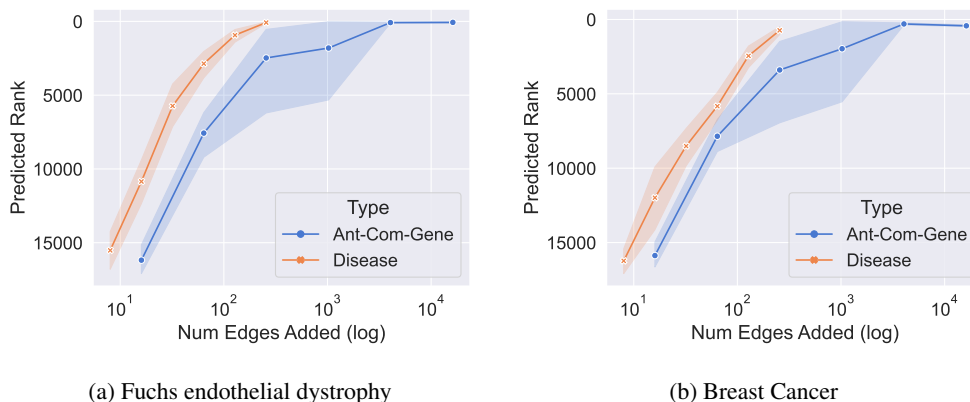(a) Fuchs endothelial dystrophy   (b) Breast Cancer

Figure 12: Addition of edges the lowest ranked gene for two diseases: Breast Cancer and Fuchs endothelial dystrophy. In both cases, the gene TRG is used as the target for the edge addition.

is *Polyubiquitin-C* (UBC), with a degree just below 10^4, nearly an order of magnitude higher than any other gene within Hetionet.

The UBC gene codes for multiple copies of the Ubiquitin protein, which in turn is a crucial component in many signalling and transport processes within the cell [22]. Specifically, ubiquitin signalling is associated with protein degradation and recycling of amino acids, which is a critical metabolic process that has important implications for normal function of a cell. With such an critical role, it is not surprising that ubiquitin, and thus indirectly ubiquitin-coding genes like UBC, are involved in many pathways.

Additionally, the overwhelming majority of the connections UBC has in Hetionet is to other genes (Table 4), likely due to UBC being directly responsible for degradation of many proteins these genes code for. Despite its predicted association with all our diseases of interest, UBC has a only single connection to a Disease and Compound entity type in the graph.

Neighbour Type	# of Conns	# of Distinct
Gene	8789	8653
BiologicalProcess	345	345
Pathway	173	173
Anatomy	55	44
CellularComponent	7	7
MolecularFunction	1	1
Compound	1	1
Disease	1	1

Table 4: Overview of UBC's connectivity within Hetionet. The disparity between number of connections and the number of distinct neighbours is due to the fact that any two entities can have multiple relationships between them.

Seen altogether, UBC provides a good example of where a highly connected entity is being over ranked as a prediction using KGE methods. UBC is ranked in the top 100 for every disease in the Hetionet dataset (See Table 5), despite having only a single relationship of the desired type (i.e. DaG), or barely any type of relationship to the entity type of interest (i.e. Disease). Furthermore, a gene that is heavily interconnected with other genes, like UBC, is likely to be considered a poor candidate disease target since any interference with its function will likely have many unintentional and undesired downstream effects, as vital metabolic and signalling processes would be altered out of the context of the disease.

As an interesting point of comparison, in the OpenBioLink dataset [9], UBC has a degree of 5937 and is only the sixth most well-connected gene, with many genes possessing a similar degree value.

Table 5 shows how in OpenBioLink, now that UBC is no longer an outlier with nearly an order of magnitude greater connections than the next gene, its predicted rank in relation to the diseases is greatly reduced.

Disease	Rank (Hetionet)	Rank (OpenBioLink)
Breast Cancer	60	232
Melanoma	82	1061
Parkinsons disease	59	1599
Fuchs Endothelial Dystrophy	23	3329
Fallopian Tube Cancer	52	2117

Table 5: The predicted rank of the gene UBC to be associated with the range of diseases. Results presented using the Hetionet and OpenBioLink graphs with the TransE model.

Graph Rewiring. Using UBC as our gene of interest, we investigated one final way of altering the graph topology: random rewiring. Here, edges from UBC to other entities in the graph are permuted by swapping the target entity via a random strategy. In this strategy, edges are rewired to the correct entity type. Further, the existing set of entities were removed from the pool of possible replacements to avoid the new edge replicating an already observed edge in the training data.

Figure 13 demonstrates the results of this rewiring process across four diseases in relation to UBC. The most apparent observation is that rewiring edges in accordance to the schema has almost no negative impact on the rank of UBC in relation to the four diseases. Indeed the rank actually increases for UBC across all diseases as we rewire edges. We hypothesise this to be caused by new edges being formed with other hub entities, thus further increasing UBC's importance within the graph. Simply put, this experiment shows we are able to go as far as completely randomising the biological information contained within UBC's edges without drastically altering how likely the model considers the gene to be associated with a disease. This is further evidence that KGE models are seemingly biased by the number of connections an entity has, rather than any domain knowledge encoded within its edges.

4.6 Discussion

The purpose of this study was not to critique a particular dataset or methodology, the trends we observed were present across multiples of both anyway, rather we wanted to take a deeper look at how topological imbalance may be affecting predictions within the drug discovery domain. We believe the evidence from this experimental evaluation suggests that degree has a strong influence on KGE models, with highly connected entities seemingly being highly ranked regardless of context. Whilst this result may not initially seem surprising, as the high-degree entities are simply observed more during the model training process, the level of correlation should nevertheless be of concern to practitioners.

Indeed the collective evidence from this study raises the prospect that domain knowledge encoded in the relationships matters less than one would expect and a highly connected, but biologically implausible, gene can be associated with a disease over one that would ultimately make a better drug target. This result also suggests that the predictions made by a KGE model, in the context of target discovery, may suffer from non-specificity – the same set of highly connected genes are highly ranked regardless of the query disease, leading to little variability in candidates.

Modeling. Many existing KGE models were designed to solve tasks in other domains, and may not have been developed against biomedical graphs which can often have a distinct topological structure, such as higher average connectivity [24]. We hope that new models will more regularly include datasets such as Hetionet during their development phase. More generally, and taking cues from the field of GNNs [25, 26, 45], new methods could be be developed which consider how best to learn meaningful representations for low-degree entities.

Evaluation. Common metrics used for evaluating KGE models such as Hits@k and MRR do not consider entity or relation frequency and thus can be biased by it [28]. As a simple mitigating step, care should be taken to ensure that low degree entities are included in test and validation sets to ascertain their impact is reflected in these metrics. Overall, global performance metrics should not

(a) Melanoma

(b) Parkinsons

(c) Fuchs Endothelial Dystrophy
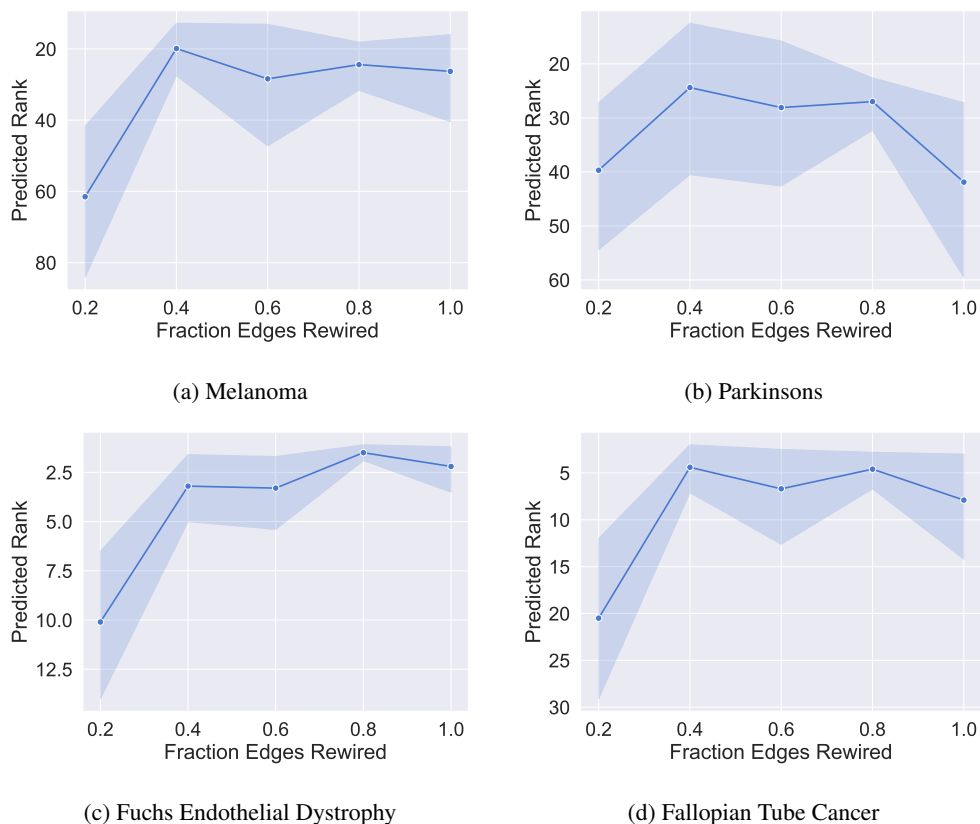
(d) Fallopian Tube Cancer

Figure 13: The change in rank assigned to the gene UBC after a varying fraction of its edges are rewired.

be the sole way of assessing model performance, rather practitioners should directly inspect the output predictions to look for any patterns contained within. For example, a highly ranked, but low degree, entity may be particularly interesting to consider.

Graph Composition. The examples demonstrated here provide a cautionary tale as relationship types seem to matter less than the overall volume of connectivity. This implies that KG practitioners should be aware of the implications of their data modelling decisions, being cognisant of any potential outliers in terms of topological structure. Thought should also be given to whether auxiliary information, which may un-proportionally increase the degree of certain entities, is really beneficial to the task at hand. This is especially important if edges are being automatically extracted, from publications for example, via NLP-based pipelines. Depending on the biological question at hand, a more focused and task-specific graph projection may lead to better overall predictions [39].

Throughout this process we have determined the following initial set of recommendations to help address topological imbalance:

- *Awareness* - We encourage practitioners to actively investigate the topologies of the graphs they are using for predictions, checking specifically for any inherent imbalance in entity or relation frequency. This also extends to awareness of which data sources and modalities have been used in graph construction, and any redundancies that may lie within.

- *Predictions* - Further, KGE users should not rely solely on performance metrics, instead the ranked list of predictions should be inspected and compared with topological features for those entities. Where metrics are required, results should be presented across different levels of connectivity.

- *Projections* - Practitioners should consider creating task-specific graph projections from larger holistic resources like Hetionet which may reduce the overall volume of connections

17

and thus imbalance. However this should be performed in a principled manner so pertinent information is not lost as the graph is reduced.

- *Edge Confidence* - One way to explore more principled projections could be to filter the graph based on confidence scores assigned to the edges. These values could be taken directly from the underlying data sources or computed where absent.

5   Conclusion

Biological knowledge graphs, especially in a drug discovery context, are heavily unbalanced in terms of degree distribution. On the one hand we observe prior knowledge bias, i.e. well-studied genes and diseases are typically have more, and often more confident, annotations. On the other hand, regardless of any *a priori* bias, biological entities are known to display a power law of connectivity, some genes have a wide range of biological functions, involved in various different diseases, and typically work in concert with a wide range of other genes. These two factors combine to create super-hub entities in the graph.

Additionally, modelling choices during the conception of the KG may lead to a very connection-dense resultant graph. For example, in the case of Hetionet, almost any two gene entities can be connected with a 1-hop path via an anatomy or cellular-component entity. This complexity is increased when one considers the multitude of ways entities can be connected to each other, depending on the data sources integrated into the graph. Automated data mining and natural text processing pipelines may further impact this by adding a significant number of relationships between the entities, often with questionable certainty, or little contextual information.

In this study, we demonstrate a reproducible tendency of KGE methods to overestimate highly-connected entities in predictive tasks. Furthermore, we show that this particular issue is independent of the choice of entities, or indeed even type of predictive task. We discuss the potential reasons for these results, as well as propose some tangible solutions that can be applied to avoid or at least ameliorate these type of problems. In the light of these findings, we believe it is prudent to pay close attention to the data modeling when creating a biological KG, application of embedding methods as well as to validate the resultant rankings for predictive tasks on these KGs. Despite the shortcomings discussed in this paper, we believe KGs and KGE methods provide valuable tools for generating meaningful representations of biological processes, reducing machine learning complexity and thus have potential for real impact in drug discovery efforts, especially when used appropriately for the underlying domain specific data. For future work, we aim to reduce the impact of connectivity imbalance through the use of a popularity bias term introduced in the model interaction function.

References

[1] Mehdi Ali, Max Berrendorf, Charles Tapley Hoyt, Laurent Vermue, Mikhail Galkin, Sahand Sharifzadeh, Asja Fischer, Volker Tresp, and Jens Lehmann. Bringing light into the dark: A large-scale evaluation of knowledge graph embedding models under a unified framework. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021.

[2] Mehdi Ali, Max Berrendorf, Charles Tapley Hoyt, Laurent Vermue, Sahand Sharifzadeh, Volker Tresp, and Jens Lehmann. Pykeen 1.0: A python library for training and evaluating knowledge graph embeddings. *Journal of Machine Learning Research*, 22(82):1–6, 2021.

[3] Mario Arduini, Lorenzo Noci, Federico Pirovano, Ce Zhang, Yash Raj Shrestha, and Bibek Paudel. Adversarial learning for debiasing knowledge graph embeddings. *arXiv preprint arXiv:2006.16309*, 2020.

[4] Albert-László Barabási and Réka Albert. Emergence of scaling in random networks. *science*, 286(5439):509–512, 1999.

[5] Stephen Bonner, Ian P Barrett, Cheng Ye, Rowan Swiers, Ola Engkvist, Andreas Bender, Charles Tapley Hoyt, and William Hamilton. A review of biomedical datasets relating to drug discovery: A knowledge graph perspective. *arXiv preprint arXiv:2102.10062*, 2021.

[6] Stephen Bonner, Ian P Barrett, Cheng Ye, Rowan Swiers, Ola Engkvist, and William L Hamilton. Understanding the performance of knowledge graph embeddings in drug discovery. *arXiv preprint arXiv:2105.10488*, 2021.

[7] Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. Translating embeddings for modeling multi-relational data. In *Advances in neural information processing systems*, pages 2787–2795, 2013.

[8] Styliani Bourli and Evaggelia Pitoura. Bias in knowledge graph embeddings. In *2020 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, pages 6–10. IEEE, 2020.

[9] Anna Breit, Simon Ott, Asan Agibetov, and Matthias Samwald. Openbiolink: A benchmarking framework for large-scale biomedical link prediction. *Bioinformatics*, 2020.

[10] Annalisa Buniello, Jacqueline A L MacArthur, Maria Cerezo, Laura W Harris, James Hayhurst, Cinzia Malangone, Aoife McMahon, Joannella Morales, Edward Mountjoy, Elliot Sollis, et al. The nhgri-ebi gwas catalog of published genome-wide association studies, targeted arrays and summary statistics 2019. *Nucleic acids research*, 47(D1):D1005–D1012, 2019.

[11] David Cook, Dearg Brown, Robert Alexander, Ruth March, Paul Morgan, Gemma Satterthwaite, and Menelas N Pangalos. Lessons learned from the fate of astrazeneca's drug pipeline: a five-dimensional framework. *Nature reviews Drug discovery*, 13(6):419–431, 2014.

[12] Tim Dettmers, Pasquale Minervini, Pontus Stenetorp, and Sebastian Riedel. Convolutional 2d knowledge graph embeddings. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.

[13] Ali Ebrahim, Elizabeth Brunk, Justin Tan, Edward J O'brien, Donghyuk Kim, Richard Szubin, Joshua A Lerman, Anna Lechner, Anand Sastry, Aarash Bordbar, et al. Multi-omic data integration enables discovery of hidden biological regularities. *Nature communications*, 7(1):1–9, 2016.

[14] Allen O. Eghrari, S. Amer Riazuddin, and John D. Gottsch. Chapter seven - fuchs corneal dystrophy. In J. Fielding Hejtmancik and John M. Nickerson, editors, *Molecular Biology of Eye Disease*, volume 134 of *Progress in Molecular Biology and Translational Science*, pages 79–97. Academic Press, 2015.

[15] Thomas Gaudelet, Ben Day, Arian R Jamasb, Jyothish Soman, Cristian Regep, Gertrude Liu, Jeremy BR Hayter, Richard Vickers, Charles Roberts, Jian Tang, et al. Utilising graph machine learning within drug discovery and development. *arXiv preprint arXiv:2012.05716*, 2020.

[16] Sara Goodwin, John D McPherson, and W Richard McCombie. Coming of age: ten years of next-generation sequencing technologies. *Nature Reviews Genetics*, 17(6):333–351, 2016.

[17] Yehudit Hasin, Marcus Seldin, and Aldons Lusis. Multi-omics approaches to disease. *Genome biology*, 18(1):1–15, 2017.

[18] Dena G. Hernandez, Xylena Reed, and Andrew B. Singleton. Genetics in parkinson disease: Mendelian versus non-mendelian inheritance. *Journal of Neurochemistry*, 139(S1):59–74, 2016.

[19] Daniel Scott Himmelstein, Antoine Lizee, Christine Hessler, Leo Brueggeman, Sabrina L Chen, Dexter Hadley, Ari Green, Pouya Khankhanian, and Sergio E Baranzini. Systematic integration of biomedical knowledge prioritizes drugs for repurposing. *Elife*, 6:e26726, 2017.

[20] Vassilis N. Ioannidis, Xiang Song, Saurav Manchanda, Mufei Li, Xiaoqin Pan, Da Zheng, Xia Ning, Xiangxiang Zeng, and George Karypis. Drkg - drug repurposing knowledge graph for covid-19. `https://github.com/gnn4dr/DRKG/`, 2020.

[21] Shaoxiong Ji, Shirui Pan, Erik Cambria, Pekka Marttinen, and Philip S Yu. A survey on knowledge graphs: Representation, acquisition and applications. *arXiv preprint arXiv:2002.00388*, 2020.

[22] Yoko Kimura and Keiji Tanaka. Regulatory mechanisms involved in the control of ubiquitin homeostasis. *The Journal of Biochemistry*, 147(6):793–798, 05 2010.

[23] Sadamori Kojaku, Jisung Yoon, Isabel Constantino, and Yong-Yeol Ahn. Residual2vec: Debiasing graph embedding with random graphs. In *Thirty-Fifth Conference on Neural Information Processing Systems*, 2021.

[24] Yushan Liu, Marcel Hildebrandt, Mitchell Joblin, Martin Ringsquandl, Rime Raissouni, and Volker Tresp. Neural multi-hop reasoning with logical rules on biomedical knowledge graphs. In *European Semantic Web Conference*, pages 375–391. Springer, 2021.

[25] Zemin Liu, Trung-Kien Nguyen, and Yuan Fang. Tail-gnn: Tail-node graph neural networks. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pages 1109–1119, 2021.

[26] Zemin Liu, Wentao Zhang, Yuan Fang, Xinming Zhang, and Steven CH Hoi. Towards locality-aware meta-learning of tail node embeddings on networks. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, pages 975–984, 2020.

[27] Farzaneh Mahdisoltani, Joanna Biega, and Fabian Suchanek. Yago3: A knowledge base from multilingual wikipedias. In *7th biennial conference on innovative data systems research*. CIDR Conference, 2014.

[28] Aisha Mohamed, Shameem Parambath, Zoi Kaoudi, and Ashraf Aboulnaga. Popularity agnostic evaluation of knowledge graph embeddings. In *Conference on Uncertainty in Artificial Intelligence*, pages 1059–1068. PMLR, 2020.

[29] Paul Morgan, Dean G Brown, Simon Lennard, Mark J Anderton, J Carl Barrett, Ulf Eriksson, Mark Fidock, Bengt Hamren, Anthony Johnson, Ruth E March, et al. Impact of a five-dimensional framework on r&d productivity at astrazeneca. *Nature reviews Drug discovery*, 17(3):167, 2018.

[30] Mark Newman. *Networks*. Oxford university press, 2018.

[31] Dac-Trung Nguyen, Stephen Mathias, Cristian Bologa, Soren Brunak, Nicolas Fernandez, Anna Gaulton, Anne Hersey, Jayme Holmes, Lars Juhl Jensen, Anneli Karlsson, et al. Pharos: collating protein information to shed light on the druggable genome. *Nucleic acids research*, 45(D1):D995–D1002, 2017.

[32] Saee Paliwal, Alex de Giorgio, Daniel Neil, Jean-Baptiste Michel, and Alix MB Lacoste. Preclinical validation of therapeutic targets predicted by tensor factorization on heterogeneous graphs. *Scientific reports*, 10(1):1–19, 2020.

[33] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in Neural Information Processing Systems*, 32:8026–8037, 2019.

[34] Janet Piñero, Àlex Bravo, Núria Queralt-Rosinach, Alba Gutiérrez-Sacristán, Jordi Deu-Pons, Emilio Centeno, Javier García-García, Ferran Sanz, and Laura I Furlong. Disgenet: a comprehensive platform integrating information on human disease-associated genes and variants. *Nucleic acids research*, page gkw943, 2016.

[35] Sune Pletscher-Frankild, Albert Pallejà, Kalliopi Tsafou, Janos X Binder, and Lars Juhl Jensen. Diseases: Text mining and data integration of disease–gene associations. *Methods*, 74:83–89, 2015.

[36] Roy Rabbie, Peter Ferguson, Christian Molina-Aguilar, David J Adams, and Carla D Robles-Espinoza. Melanoma subtypes: genomic profiles, prognostic molecular markers and therapeutic possibilities. *The Journal of Pathology*, 247(5):539–551, 2019.

[37] Tahleen Rahman, Bartlomiej Surma, Michael Backes, and Yang Zhang. Fairwalk: Towards fair graph embedding. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, pages 3289–3295. International Joint Conferences on Artificial Intelligence Organization, 7 2019.

[38] F Ann Ran, Patrick D Hsu, Jason Wright, Vineeta Agarwala, David A Scott, and Feng Zhang. Genome engineering using the crispr-cas9 system. *Nature protocols*, 8(11):2281–2308, 2013.

[39] F Ratajczak, M Joblin, M Ringsquandl, and M Hildebrandt. Task-driven knowledge graph filtering improves prioritizing drugs for repurposing. *Technical Report*, 2021.

[40] Andrea Rossi, Denilson Barbosa, Donatella Firmani, Antonio Matinata, and Paolo Merialdo. Knowledge graph embedding for link prediction: A comparative analysis. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 15(2):1–49, 2021.

[41] Daniel Ruffinelli, Samuel Broscheit, and Rainer Gemulla. You can teach an old dog new tricks! on training knowledge graph embeddings. In *International Conference on Learning Representations*, 2019.

[42] Marina Stasenko, Olga Fillipova, and William P. Tew. Fallopian tube carcinoma. *Journal of Oncology Practice*, 15(7):375–382, 2019. PMID: 31283415.

[43] Zhiqing Sun, Zhi-Hong Deng, Jian-Yun Nie, and Jian Tang. Rotate: Knowledge graph embedding by relational rotation in complex space. In *International Conference on Learning Representations*, 2019.

[44] Agata Szymiczek, Amna Lone, and Mohammad R. Akbari. Molecular intrinsic versus clinical subtyping in breast cancer: A comprehensive review. *Clinical Genetics*, 99(5):613–637, 2021.

[45] Xianfeng Tang, Huaxiu Yao, Yiwei Sun, Yiqi Wang, Jiliang Tang, Charu Aggarwal, Prasenjit Mitra, and Suhang Wang. Investigating and mitigating degree-related biases in graph convoltuional networks. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, pages 1435–1444, 2020.

[46] Kristina Toutanova and Danqi Chen. Observed versus latent features for knowledge base and text inference. In *Proceedings of the 3rd workshop on continuous vector space models and their compositionality*, pages 57–66, 2015.

[47] Théo Trouillon, Johannes Welbl, Sebastian Riedel, Éric Gaussier, and Guillaume Bouchard. Complex embeddings for simple link prediction. In *International Conference on Machine Learning (ICML)*, 2016.

[48] Nan Wang, Lu Lin, Jundong Li, and Hongning Wang. Unbiased graph embedding with biased graph observations. *arXiv preprint arXiv:2110.13957*, 2021.

[49] Quan Wang, Zhendong Mao, Bin Wang, and Li Guo. Knowledge graph embedding: A survey of approaches and applications. *IEEE Transactions on Knowledge and Data Engineering*, 29(12):2724–2743, 2017.

[50] Zhen Wang, Jianwen Zhang, Jianlin Feng, and Zheng Chen. Knowledge graph embedding by translating on hyperplanes. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2014.

[51] Olivier J. Wouters, Martin McKee, and Jeroen Luyten. Estimated Research and Development Investment Needed to Bring a New Medicine to Market, 2009-2018. *JAMA*, 323(9):844–853, 03 2020.

[52] Bishan Yang, Scott Wen-tau Yih, Xiaodong He, Jianfeng Gao, and Li Deng. Embedding entities and relations for learning and inference in knowledge bases. In *Proceedings of the International Conference on Learning Representations (ICLR) 2015*, May 2015.

[53] Chuxu Zhang, Dongjin Song, Chao Huang, Ananthram Swami, and Nitesh V Chawla. Heterogeneous graph neural network. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 793–803, 2019.

[54] Michael Zietz, Daniel S Himmelstein, Kyle Kloster, Christopher Williams, Michael W Nagle, Blair D Sullivan, and Casey S Greene. The probability of edge existence due to node degree: a baseline for network-based predictions. *Manubot, Tech Rep*, 2020.