

Knowledge Graphs and Explainable AI for Drug Repurposing on Rare Diseases

Pablo Perdomo-Quinteiro^a, Katy Wolstencroft^b, Marco Roos^c and Núria Queralt-Rosinach^{c,*}

^a *Grupo de Aplicación de Telecomunicaciones Visuales, Escuela Técnica Superior de Ingenieros de Telecomunicación, Universidad Politécnica de Madrid, Avenida Complutense 30, 28040, Madrid, Spain*

^b *Advanced Compute and Data Core, Amsterdam University Medical Centre, Meibergdreef 9, 1105 AZ, Amsterdam, The Netherlands*

^c *Department of Human Genetics, Leiden University Medical Center, Albinusdreef 2, 2333 ZA, Leiden, The Netherlands*

Abstract. Artificial Intelligence (AI)-based drug repurposing is an emerging strategy to identify drug candidates to treat rare diseases. However, cutting-edge algorithms based on Deep Learning (DL) typically don't provide a human understandable explanation supporting their predictions. This is a problem because it hampers the biologists' ability to decide which predictions are the most plausible drug candidates to test in costly lab experiments. In this study, we propose *rd-explainer* a novel AI drug repurposing method for rare diseases which obtains possible drug candidates together with human understandable explanations. The method is based on Graph Neural Network (GNN) technology and explanations were generated as semantic graphs using state-of-the-art eXplainable AI (XAI). The model learns features from current background knowledge on the target rare disease structured as a Knowledge Graph (KG), which integrates curated facts and their evidence on different biomedical entities such as symptoms, drugs, genes and ortholog genes. Our experiments demonstrate that our method has excellent performance that is superior to state-of-the-art models. We investigated the application of XAI on drug repurposing for rare diseases and we prove our method is capable of discovering plausible drug candidates based on testable explanations. The data and code are publicly available at <https://github.com/PPPerdomoQ/rare-disease-explainer>.

Keywords: Rare Disease (RD), Knowledge Graph (KG), Drug Repurposing, Graph Neural Network (GNN), Explainable AI (XAI)

1. Highlights

- We demonstrated the use of graph-based explainable AI for drug repurposing on rare diseases to accelerate sound discovery of new therapies for this underrepresented group.
- We developed *rd-explainer* for rare disease specific drug research for faster translation. It predicts drugs to treat symptoms/phenotypes, it is highly performant and novel candidates are plausible according to evidence in the scientific literature and clinical trials. Key is that it learns a GNN model that is trained on a knowledge graph built specifically for a rare disease. We provide *rd-explainer* code freely available for the community.
- *rd-explainer* is researcher-centric interpretable ML for hypothesis generation and lab-in-the-loop drug research. Explanations of predictions are semantic graphs in line with human reasoning.
- We detected an effect of knowledge graph topology on explainability. This highlights the importance of knowledge representation for the drug repurposing task.

*Corresponding author. E-mail: n.queralt_rosinach@lumc.nl.

2. Introduction

Developing new drugs can be a challenging effort that often ends with the drug not being able to launch. Recent studies have shown that around 90% of drugs fail to be approved during their clinical development [1]. This leads to a fruitless expenditure of both time and money that will yield no financial returns. The situation is even worse in the case of rare diseases, as pharmaceutical companies may consider it risky to invest large amounts of resources into developing drugs that only a small percent of the population will need. Nonetheless, in total, human beings are affected by approximately 7,000 rare diseases, of which only 5% have an effective treatment [2]; and only in Europe about 36 million people suffer from rare diseases [3].

In this scenario, drug repurposing strategies have appeared as a possible approach to solve these issues. By reusing drugs that have already been approved, companies can avoid many of the costly and time-consuming steps of clinical trials. In this context, innovative approaches to drug repurposing, such as computational strategies and AI-driven methodologies, have emerged as promising solutions to address these challenges. Graph-based drug repurposing is another noticeable strategy that has gained attention in recent years. By constructing intricate networks of molecular interactions, genes, proteins, and diseases, this approach unveils hidden relationships and connections that might otherwise go unnoticed [4].

Still, many people remain skeptical about AI-driven decisions, specially Machine Learning (ML) and Deep Learning (DL), as many of them come with no explanation that can help to understand the reason why they should be trusted (also called black-box AI). This issue is especially significant in the healthcare field, where decisions may have an important impact on people's lives. Also, giving valid explanations can help researchers to point in the right direction in the generation of hypotheses that are testable in the lab and enable a solid knowledge discovery. Furthermore, the EU General Data Protection Regulation (GDPR) is requesting the AI industry to fulfill the 'right to explanation' [5]. This 'right to explanation' implies that when a decision is significantly affected by an automated process/algorithm, the individual can demand an explanation. In recent years, many different tools have appeared to try and cover this gap in the emerging explainable AI (XAI) research area [6–8].

In this study, we explore whether AI can be used to produce both predictions and explanations in computational drug repurposing for rare diseases and, if so, how helpful can these explanations be for hypothesis generation. The main objective of this work was to develop and implement a pipeline to find marketed drugs that can be used to treat the symptoms of a rare disease. Our approach is based on cutting-edge AI algorithms used in computational drug repurposing such as graph ML using knowledge graphs (KG) and graph neural networks (GNN), and XAI methodology to provide the explanations supporting the drug predictions made by the AI model. The approach was evaluated by selecting Duchenne muscular dystrophy (DMD) as a case study, a genetic disorder that is the most common form of muscular dystrophy [9]. We demonstrate the generalizability of our approach by applying the pipeline to different rare diseases.

3. Related work

3.1. Knowledge graph-based drug repurposing

The state-of-the-art of computational drug repurposing approaches make use of graph-based structures and AI techniques to find potential drug candidates. One of the main advantages of using graph structures is that they can easily incorporate information from different sources. This is especially important in the domain of rare diseases, where information is distributed and often scarce. The ability to integrate as much relevant data as possible can confer a significant advantage. An example of this would be the recent study of Al Saleem et al. [10], where a knowledge graph was used to discover drug candidates to treat COVID-19.

Different ML algorithms can be used to analyse knowledge graphs, including matrix factorization, random-walk approaches (node2vec [11]), geometric embeddings (DistMul [12]) and GNNs [13, 14], each one of them with its own advantages and disadvantages, see Table 1. In our study, we used a combination of random-walk approaches and GNNs as in contrast to other methods (like matrix factorization or geometric embeddings) they can easily

incorporate new information without the need of retraining the ML model. This is especially relevant in the field of drug repurposing where new information about drugs, genes and diseases is being published [15–17].

Table 1
Table showing the main advantages and disadvantages of different graph ML methods.

Method	Example	Advantages	Disadvantages
<i>Matrix factorization</i>	ADA-GRMFC [18]	Simple Model Global Information	Computationally expensive Scalability Difficult to include new values
<i>Random-walks</i>	node2vec [11]	Simple model Computationally efficient	Only local information Can't learn from features Can't capture structural information
<i>Geometric Embeddings</i>	DistMult [19]	Interpretability Good performance	Only local information Can't capture structural information
<i>GNNs</i>	GraphSAGE [20]	Global and local information Structural information	Black-box Computationally expensive

3.2. Explainable AI on graph ML

One of the graph-based methods that can provide explanations of the predictions, also called local explanations, is (Graph)LIME [6], an adaptation of the popular and more general explainability method LIME [7]. The idea behind this method is the following: when trying to get an explanation for a given prediction, (Graph)LIME performs small perturbations to the features of nodes, and sees how the predictions vary with respect to the initial prediction. The more the prediction changes, the more the model is relying on that feature to obtain its prediction. This way, explanations in this model are given in the form of a set of node features. Among its drawbacks, this method can only be used in node classification tasks. Another explainability method is CRIAGE [8] where explanations are given as a set of rules.

Finally, the method chosen in this work is GNNExplainer [21]. The insight of how this method works is the following: given an initial prediction (link prediction, node classification or graph classification) obtained through a GNN, GNNExplainer finds a subset of node features and edges that are responsible for the prediction. This subset is obtained by training an edge and node mask. This method was chosen as explanations are provided in the form of a subgraph that can be easily understandable. Additionally, it is a post-hoc XAI method, i.e., it is model-agnostic, which means that if more sophisticated GNNs are developed in the future, these new GNNs can be easily incorporated into the pipeline. These features make it a popular method in the research community [22–24]. However, a major drawback is that it lacks consistency when obtaining explanations. This means that explanations on the same prediction can significantly change if running GNNExplainer several times.

4. Methods

4.1. rd-explainer method overview

rd-explainer is the drug repurposing method we developed for rare diseases and its pipeline is illustrated in Figure 1. *rd-explainer* has three modules: the Knowledge Graph Construction module constructs a KG for the specific rare disease and drug repurposing task, the Prediction module trains a GNN model and predicts drug candidates for the rare disease symptoms, and the Explainer module computes the most important semantic subgraphs that explain the connection between the predicted drug and the symptom. Firstly, information related to the disease is gathered from different data sources: Monarch Initiative knowledge base [25] for disease pathology, and DrugCentral [26] and Therapeutic Target Database [27] for disease druggability. This information is then preprocessed and captured as a knowledge graph. Next, for each node in the graph a feature vector is obtained that will be used as input for

the GNN model. This is done by making use of a method known as edge2vec [28] to consider the different edge semantics in the KG for node embedding learning. We used the version extracted from GitHub (accessed in 2021)¹. The following step is to build and train the GNN model, which is done using the GraphSAGE framework for graph representation learning [20]. Next, it is the link prediction for each drug-symptom node embeddings pair by using the dot product as scoring function. Finally, we produced prediction explanations as semantic graphs using GNNExplainer [21], a recent and, to our knowledge, one of the first XAI methods for obtaining explanations from GNN predictions. GraphSAGE and GNNExplainer were implemented using using Pytorch Geometric version 2.0.4. The code is freely accessible with an open license at <https://github.com/PPerdomoQ/rare-disease-explainer>.

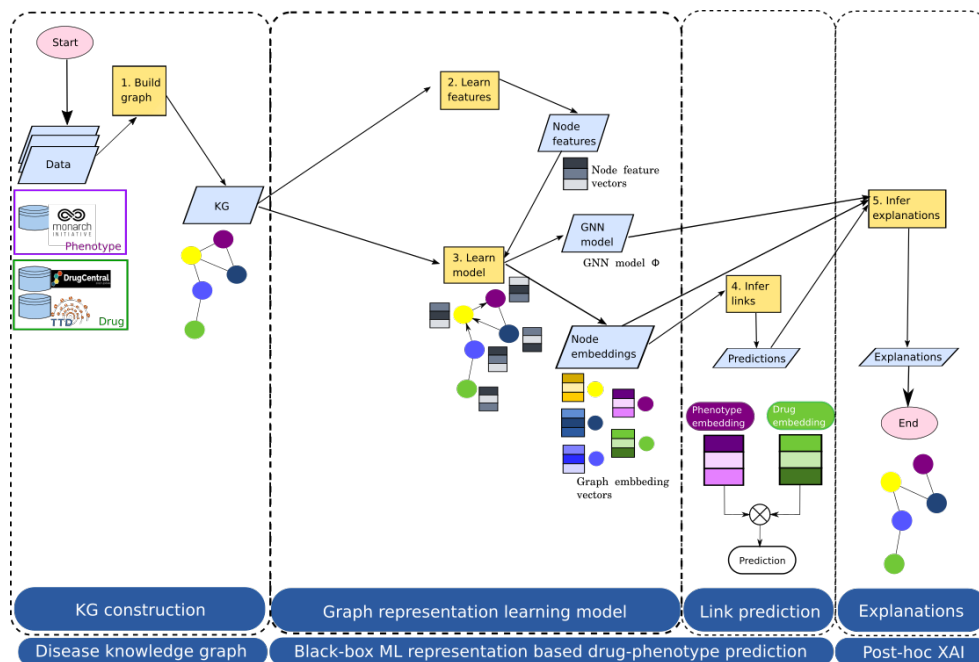


Fig. 1. *rd-explainer* drug repurposing method pipeline developed in this work.

4.2. Rare disease-specific drug repurposing knowledge graphs

4.2.1. Data sources

Data was obtained from three different sources: Monarch [25] (accessed in 2021), DrugCentral [26] (2021 version) and Therapeutic Target Database (TTD) [27] (November 8th, 2021 version). Monarch is a knowledge base built on semantic principles, unifying gene, variant, genotype, phenotype, and disease data across different species. Its primary aim is to establish links between genes and phenotypes, thereby facilitating computational exploration of human disease biology. Monarch was chosen as it contains curated information across different species. This way, because rare diseases are often less studied than common diseases, incorporating information from other species can maximize the amount of knowledge in the graph. However, Monarch does not specialize in drug information.

Drug information was incorporated from DrugCentral (drug-target information) and from Therapeutic Target Database (drug-disease information). DrugCentral is a comprehensive online database that provides information about approved drugs, active ingredients and other pharmaceutical products. One of its major features is that it is open source and its data is freely available to anyone. For this project, we made use only of the drug-target information (as it is the main piece of information that is not present in Monarch) downloaded as a tsv file from their site [27]². Similarly, TTD is a database that specializes in drugs and their respective therapeutic targets. Once

¹<https://github.com/RoyZhengGao/edge2vec>

²DrugCentral, *Download site*, accessed March 2022, https://unmtid-dbs.net/download/DrugCentral/2021_09_01/drug.target.interaction.tsv.gz

more, this database is freely accessible and its information can be easily downloaded in csv format (in this project, we just made use of the drug-disease information [27]³, once again because it is the information that is missing in Monarch).

4.2.2. Knowledge graph construction

To extract information from Monarch, the BioKnowledge Reviewer [29] tool version 1.0 was used. This tool was originally created to collect knowledge from several sources and create a knowledge graph that could be later used for hypothesis generation. It works by using several seeds (node identifiers (IDs)) as input to query the Monarch API and constructing the graph based on the neighborhood of those seeds. After introducing the seeds in the BioKnowledge Reviewer pipeline, the final output is the rare disease research question specific knowledge graph structured in two dataframes (stored as csv files). One of them contains a list of nodes with their respective name, IDs, semantic entity type, synonyms and description. The second file contains the list of edges, again containing the IDs of the entities participating in each link and other edge information such as type of edge, supporting evidence and reference date. Monarch was our main source of information, and so it served as a starting point to create the rest of the graph. This way, data from other data sources was modified to fit Monarch's standards. Finally, the graphs were constructed using the networkx Python library [30] version 2.3.6. With this library the dataframes extracted using BioKnowledge Reviewer were converted into a *Graph* object. For a more detailed understanding of the knowledge graph generation process, please consult the associated code available at: <https://github.com/PPperdomoQ/rare-disease-explainer>.

We integrated data in two different knowledge graphs to perform the experiments. Each one of them was constructed using different (number of) node seeds to extract information from Monarch. The first one (KG A) only uses two seeds: DMD seed (HGNC:2928), corresponding to the human gene that causes the disease; and DMD seed (MONDO:0010679), corresponding to the disease itself. The second graph (KG B), extends KG A by including as seeds all phenotypes of the rare disease (in total, 27 more seeds). The seeds used for the construction of each graph can be found in Tables S1 and S2. The idea of creating two different graphs is to find out if the performance of the model and the quality of the explanations increases by incorporating more (phenotypic) information.

4.3. ML model and XAI

4.3.1. Node features

At this point, none of the nodes have any specific node features. It is possible to run a GNN relying only on graph information, i.e., network topology (this is done, for example, by using the node degree as graph feature); nonetheless, this resulted in a poor performance (results not shown). To increase the efficiency of the model, edge2vec was used to produce a specific embedding for each node that captures information about its neighborhood. edge2vec [28] is a tool that generates node embeddings based on the node neighborhood and types of edges connecting each node. After executing edge2vec, each node was given a unique feature vector.

4.3.2. Data splitting

As any other machine learning task, data needs to be splitted into training, (validation) and test sets. However, when tackling a link prediction task, there are different ways to perform this split. In link or edge prediction tasks, edges can be divided into two groups: message passing edges and supervision edges. Message passing edges are the ones that will be used by our GNN to obtain the embeddings, while supervision edges are the ones that will be used to test the performance of our model [31, 32]. Additionally, when creating the supervision edges it is necessary to include negative examples by applying negative sampling. These negative sample edges are edges that are not present in our original graph, i.e., entities that it is known are not linked or there is no known link between them, and the idea is that the neural network is able to learn to distinguish true or positive edges from false or negative edges. In general, one negative edge is created for each true edge [31, 32].

In this work, the method that was selected was the all-graph transductive split [31, 32]. When applying this method the division is done in the following way: in the training dataset the supervision edges and the message passing edges are the same; in the validation dataset the message passing edges are the training edges (message and

³Therapeutic Target Database, *Download site*, accessed March 2022, https://idrblab.net/ttd/sites/default/files/ttd_database/P1-05-Drug_diseases.txt

supervision) and the supervision edges are different from the training supervision edges; finally, the test set message passing edges are formed by the validation edges and supervision edges that are different from the training and validation supervision edges.

This method is one of the standard settings when performing link prediction tasks, as the whole graph can be seen in all dataset splits [31]. The proportion used were 80% of edges used for training set, 10% for validation set and 10% for test set. The training set will be used to train the model, the validation set to select the best hyperparameters, and the test set to obtain the global performance of the model.

4.3.3. GNN model

We first utilized a GNN algorithm to learn vector representation embeddings for nodes in our knowledge graphs. Then, we applied these node embeddings for drug-phenotype link prediction. The GNN algorithm that we used in this work is called GraphSAGE [20]. GraphSAGE performs inductive graph representation learning by leveraging rich node attribute information. The main advantage that was brought by GraphSAGE is its scalability: instead of working with full batches (the whole graph is seen during the training) it works with mini-batches. Each mini-batch is a subset of computational graphs (a computational graph is the individual GNN that is built for each node) of N nodes. By applying this technique, the GNN can better manage larger graphs. The GraphSAGE model was created using the DeepSNAP library [32] version 0.2.1 to obtain the predictions. The hyperparameter optimization was performed using RayTune [33] version 1.12.1, as it is a model-agnostic library that allows to run multiple trials in parallel, reducing the training time. The list of hyperparameters that were needed to be tuned and the optimal values can be found in Table S5. In total, 30 models were created (each of them containing a random selection of parameters).

4.3.4. Drug-phenotype link predictions

The GNN model generates embeddings for individual nodes within the graph as its final output. By applying the dot product between distinct node pairs and applying a sigmoid function, we obtain a value that shows the likelihood of a link existing between those nodes. Consequently, we obtain dot products between each drug and every phenotype in the graph, and rank them in descending order. The top-ranked dot products are considered the most promising targets. Links that were already present in the graph were removed from the ranking.

4.3.5. Graph-based prediction explanations

We applied GNNExplainer to generate explanations for every drug-phenotype prediction. To do so, we adapted the pipeline code (from Pytorch geometric version 2.0.9) to generate explanations for the link prediction task, which was not implemented in authors' version [21] (see pseudocode in Algorithm 1 in the Supplementary material). However, this XAI algorithm has a problem of robustness in the explanations it produces [34] and, additionally, it may yield disconnected graphs affecting to the interpretability of explanations by domain-users. To solve this issue, we developed the following procedure. First, we make the assumption that a complete explanation is one that connects the two targeted nodes. If drug A can treat phenotype B, there must be some common pathway that allows A to interact with B. This way, the procedure starts by running GNNExplainer for several iterations. In each iteration, networkx is used to check if, in the subgraph generated by GNNExplainer, a path exists between both nodes. If no path is found, it continues with the next iteration; if it does exist, it stops iterating and that subgraph is considered to be the final explanation. If no subgraph is found that satisfies the 'pathway' condition, the last subgraph is returned as a possible explanation.

In total 7 phenotypes were selected to evaluate the explanations (Muscular Dystrophy (HP:0003560), Respiratory Insufficiency (HP:0002093), Arrhythmia (HP:0011675), Congestive Heart Failure (HP:0001635), Dilated Cardiomyopathy (HP:0001644), Cognitive Impairment (HP:0100543) and Progressive Muscle Weakness (HP:0003323)). These phenotypes were selected to cover all the main areas that are affected by the disease (muscular, respiratory, cardiac and intellectual symptoms). For each prediction obtained in these phenotypes (three drug predictions per phenotype), an explanation was obtained. This process was done for the predictions coming from KG A and for those coming from KG B. This makes a total of 42 explanations (21 for each graph).

Regarding the parameters of GNNExplainer, because the graphs are highly connected, explanations were generated by using the 1-hop neighborhood around the graph. Using a higher k-hop neighborhood is not recommended

as the amount of nodes in the subgraph increases exponentially which can make it difficult to understand the explanation. This happens because both graphs are scale-free graphs, and thus, by increasing the number of hops there is a higher chance that a 'hub-node' is hit, and the number of nodes escalates exponentially (see Section 5.1 in the results).

Additionally, the maximum size of the explanations was set to 15 (this means that no more than 15 edges will be part of the explanation). This way, we will avoid obtaining too complex explanations with many edges that might be impossible to comprehend by researchers. This was done by selecting the edges whose contribution values are among the 15th highest values.

Finally, the maximum number of iterations was set to 10. In other words, if after 10 iterations GNNExplainer has not found an explanation that connects the drug candidate with the targeted phenotype it will conclude that no 'complete' explanation was found, and the last explanation produced by GNNExplainer will be the one that will serve as final answer. This parameter can be increased or reduced depending on the expectations of the researcher. A large number of iterations increases the chances of finding a complete explanation at the cost of more computational time. On the contrary, reducing the number of iterations reduces the computational time, which can be useful if a researcher wants to obtain explanations for a large number of predictions.

4.4. Evaluation and metrics

4.4.1. Evaluation of GNN model

Data. We used both graphs KG A and B. Data was splitted into three sets: training set, validation set and test set. **Baselines.** Our baselines include edge2vec [28], GraphSAGE [20], ComplEX [35], DistMult [19], and TransE [36]. **Evaluation metrics.** The Area Under the Precision-Recall curve (AUPRC) was used to validate and test the performance of the model, as it has been shown that it leads to better precision when evaluating link prediction [37]. Additionally, we also computed the Area Under the ROC curve (AUROC), Precision, Recall and the F1-Score metrics - the harmonic mean of precision and recall.

Other evaluations were developed to further assess the performance of the model. These evaluations include the testing of different negative sampling sizes ($n = 1, 5, 10$ and 20) to determine the importance of keeping the data balanced. Additionally, both a regular 10-fold cross validation and a biased 7-fold cross validation were performed. The biased cross validation consists of the following: in each fold 4 phenotypes were removed from the training set, and it was observed how well the model was able to predict the links of the removed phenotypes.

4.4.2. Evaluation of explanations

The evaluation of explanations was done manually, following a two-step process. Firstly, they were classified as complete or incomplete explanations based on the appearance of connection between the drug and the phenotype. We developed a function to visualize the explanations as semantic graphs (see section 9.5 in the supplementary material for further details). This way, if the explanation contains a link between the drug and the phenotype it is considered to be a complete explanation. These explanations are considered the ones that are truly useful as they are the ones that can be easily understood and interpreted. On the other hand, explanations where there is no link between drug and phenotype (where there are two separate clusters) or where only one of the target elements (either the drug or the phenotype) is missing, are considered incomplete explanations. Several illustrative examples are provided in the supplementary material (See section 9.6).

During the second step, we evaluated explanations using an objective and a subjective approach. First, complete explanations were reviewed and a manual search was performed to check whether the explanation proposed by the model had been already described in the literature (objective evaluation). This process was only performed for those predictions that have supporting evidence in the literature and that were classified as complete explanations. The examination of the literature was performed using PubMed and Google Scholar during the first half of 2022. Finally, each explanation was evaluated with our own biological knowledge (subjective evaluation).

5. Results

5.1. Rare disease KG topology and representation for drug repurposing

We generated two different drug repurposing knowledge graphs for the Duchenne muscular dystrophy rare disease. KG A contains 10786 nodes, 93905 directed edges. The average node degree of the graph ($\frac{2 \times \text{number of edges}}{\text{number of nodes}}$) is 10.83, being the node with the highest degree, the human DMD gene, with a total degree of 1683. The diameter of the graph was 6, meaning that the longest shortest path between two nodes is 6 (in other words, one can travel from one node to another in 6 steps or fewer). The final feature that was obtained is the clustering coefficient, which measures the extent to which a graph is clustered together. In a complete graph (where all nodes are connected to all nodes) this clustering coefficient is equal to 1, while in a tree-like graph this coefficient is equal to 0. In KG A this clustering coefficient is equal to 0.33. A summary of the features can be found in Table 2.

In the case of KG B (built from 29 nodes: KG A seeds extended by 27 phenotypes of DMD), the total number of nodes is 83665, with a total of 1984774 directed edges. The average degree in this case is of 34.43, being the node with the highest degree the physiological process 'Protein Binding' with a total degree of 4817. The diameter of the graph is of 7, which shows one of the features of scale-free networks: despite increasing the number of nodes 8 times and the number of edges 20 times, the diameter of graph B only increased one unit with respect to graph A. In this case, the clustering coefficient is equal to 0.48, showing that KG B is more clustered. Table 2 shows a summary of the features of both graphs.

Table 2
Table showing features of KG A and B.

Property	KG A	KG B
<i>Number of Nodes</i>	10786	83665
<i>Number of Directed Edges</i>	93855	1984774
<i>Number of Undirected Edges</i>	58435	1440418
<i>Average Degree</i>	10.83	34.43
<i>Highest Degree</i>	1683	4817
<i>Diameter</i>	6	7
<i>Average Clustering Coefficient</i>	0.33	0.48
<i>Number of drugs</i>	337	1565
<i>Number of diseases</i>	5419	25636
<i>Number of drug-disease pairs</i>	86	599

The schema of the knowledge graph, which is the same for KG A and KG B, can be seen in Figure 2 and shows how the 8 different node types interact with each other. The schema contains 24 and 29 different edge types for KG A and KG B respectively, which are not included in this figure for clarity, but are listed in the Supplementary material S3 and S4.




Fig. 2. Schema of the knowledge graph. Node types are: drugs or chemical compounds (DRUG), genes (GENE), symptoms/phenotypes or diseases (DISO), gene variants (VARI), genotypes (GENO), gene orthologs (ORTHO), anatomical structures (ANAT), and biological processes (PHYS).

5.2. GNN model performance for rare disease specific drug repurposing

In total, two GNNs were used, one trained on KG A and one trained on KG B. The hyperparameter optimization was developed using RayTune and the optimal values can be found in Table S5. These hyperparameters were obtained by training several GNN models (Random Search) on graph A; and were later used to train a GNN model on graph B.

Table 3
Precision, Recall and F1-Score obtained on each dataset, trained on each graph.

Dataset	Precision		Recall		F1-Score	
	KG A	KG B	KG A	KG B	KG A	KG B
Training	0.93	0.96	0.96	0.93	0.95	0.95
Validation	0.93	0.96	0.93	0.92	0.93	0.94
Test	0.93	0.96	0.93	0.92	0.93	0.94

To measure link prediction performance, the scores obtained were Precision, Recall and the F1-Score, and can be found in Table 3 (the threshold used was 0.8). We found that both models (the one trained with KG A and the one trained with KG B) yield to high performance (F1-Score = 0.93 and 0.94 in KG A and B, respectively). To visualise the performance of the link prediction task, the ROC curve of KG A and KG B obtained on the test set can be found in Figure 3 and Figure 4, respectively.




Fig. 3. AUROC on the test dataset using KG A.




Fig. 4. AUROC on the test dataset using KG B.

5.3. Evaluating *rd-explainer* with state-of-the-art methods

Firstly, we evaluated our GNN model applying different strategies and compared its performance to the state-of-the-art graph embeddings used in drug repurposing methods. Then, we evaluated our approach based on its ability to predict drugs that are already reported in the literature for a new symptom or phenotype.

We performed a regular 10-fold cross-validation and a biased 7-fold cross-validation evaluation in KG A. The regular 10-fold cross-validation obtained an average AUPRC of 0.98 and an average AUROC of 0.98. For the biased 7-fold cross-validation, in each fold 4 symptoms (along with the edges connected to those symptoms) were removed from the training set. Then the performance of the model was tested on the removed symptoms. In this case, the average AUPRC was 0.75 and the AUROC was 0.8.

The performance of the pipeline was evaluated for a different number of negative edges. This evaluation was only performed in KG A due to the large increase in the number of edges in the evaluation tests (and the consequential increase in the computational time). The results can be seen in Table 4. It is seen that as the number of negative edges increases, the PR curve is affected while the ROC curve remains mostly intact, a result that has previously been reported [38].

Table 4

Performance as the number of negative edge samples increases. This results were obtained using KG A.

<i>Number of negative edges</i>	Precision	Recall	F1-Score	AUROC	AUPRC
<i>1</i>	0.95	0.95	0.95	0.99	0.99
<i>5</i>	0.94	0.90	0.92	0.94	0.94
<i>10</i>	0.93	0.91	0.86	0.86	0.85
<i>20</i>	0.93	0.57	0.62	0.82	0.70

Finally, the performance of rd-explainer (tested in KG A) was also compared to other state-of-the-art methods, including edge2vec, GraphSAGE, ComplEX, DistMult, and TransE. Our results can be seen in Table 5 and they revealed that rd-explainer outperformed all the other methods based on the different evaluation metrics measured.

Table 5

Prediction performance metrics comparing rd-explainer with other state-of-the-art graph embedding methods including edge2vec, GraphSAGE, ComplEX, DistMult and TransE. The best results are **highlighted**. In the headings, **P** stands for *Precision*, **R** for *Recall*, and **F1** for *F1-Score*.

<i>Method</i>	P	R	F1	AUROC	AUPRC
<i>edge2vec</i>	0.90	0.90	0.90	0.98	0.97
<i>GraphSage</i>	0.71	0.65	0.62	0.64	0.87
<i>ComplEX</i>	0.84	0.76	0.74	0.95	0.99
<i>DistMult</i>	0.93	0.93	0.92	0.95	0.98
<i>TransE</i>	0.88	0.87	0.87	0.95	0.95
<i>rd-explainer</i>	0.95	0.95	0.95	0.99	0.99

5.4. Drug predictions validation based on the scientific literature

We also evaluated the prediction performance based on the capacity of our method to discover marketed drugs already reported being used for a new phenotype. First, we listed for each of the 7 selected phenotypes the three drugs with the highest scores. Because the objective is to find new indications for drugs; if any of the reported drugs already appears in the graph as a treatment for the targeted symptom, this drug will be skipped and the next one with the highest score will be selected. For example, if aprindine is selected as the drug with the highest score to treat arrhythmia, but the relation 'aprindine is a substance that treats arrhythmia' is already present in our graph, aprindine won't be reported as a possible drug candidate.

For each possible drug candidate, a literature search was carried out to find preliminary evidence if that drug had already been used to treat the symptom. If the drug was contraindicated to treat the symptom (or if it could cause the symptom) it was also annotated. Results regarding each drug candidate obtained using KG A can be found in Table S9. Additionally, Table 6 summarizes the amount of drugs (in percentage) that contained supporting evidence, contraindication evidence or no evidence at all. We found that only a fifth of the drug candidates had supporting evidence in the literature, and that the vast majority of the candidates (65.43%) did not have any evidence at all. There is a small percentage of them that are actually contraindicated to treat the targeted symptom/phenotype. Finally, the amount of supporting/contraindicating evidence can be found summarized in Table S7.

Table 6

Percentage of drugs containing supporting evidence, contraindication evidence or no evidence at all for both Graph A and B.

<i>Property</i>	KG A	KG B
<i>Supporting Evidence</i>	20.99 %	27.16 %
<i>Contraindication Evidence</i>	13.58 %	14.82 %
<i>No Evidence</i>	65.43 %	58.02 %

The same approach was followed in the case of KG B. Information regarding the drug candidates for each symptom (as well as the supporting evidence) can be found also in Table S10. Additionally, the percentage of drugs with supporting evidence, contraindication evidence or no evidence at all can be seen in Table 6. In this case, the number of drug candidates with evidence has increased with respect to the drug candidates obtained with KG A (27% in B vs 21% in A), and the number of drug candidates with no evidence has been reduced (58% in B vs 65% in A). The number of drug candidates with contraindications remains almost the same (13% in A vs 14% in B).

5.5. Evaluating drug repurposing explanations as semantic graphs

Evaluating an explanation is a tough task and many different benchmarks are recently appearing to evaluate them [39]. In this work, we followed two different approaches to evaluate the explanations: a more subjective one, where the explanation was evaluated with our own biological knowledge; and a more objective one, where a manual literature search and curation was performed to check if the suggested explanation has already been reported. We selected 7 phenotypes (muscular dystrophy, respiratory insufficiency, arrhythmia, dilated cardiomyopathy, congestive heart failure, progressive muscle weakness and cognitive impairment) and their top 3 predictions, then explanations were produced from the models trained on both KGs. The selection of these phenotypes aimed to cover the diverse systems affected by the disease. Each explanation was analyzed and, if possible, compared to the one that was found in the literature.

Explanations were classified into complete and incomplete explanations. Complete explanations are those that show a connection (path) between the drug candidate and the targeted symptom/phenotype (Figure S3). They are considered complete as they allow for an easy human-understandable interpretation. On the other hand, incomplete explanations are those where the explanation is composed of two separated clusters (one for the drug and one for the phenotype) (Figure S4) or by a unique cluster where either the drug or the phenotype is missing (Figure S5).

Table 7
Number and percentage of complete and incomplete explanations in each evidence type.

	Complete Explanations	Percentage Complete Explanations	Incomplete Explanations	Percentage Incomplete Explanations
<i>Supporting Evidence</i>	13	68 %	6	32 %
<i>Contraindication Evidence</i>	3	30 %	7	70 %
<i>No Evidence</i>	5	38 %	8	62 %
<i>Total</i>	21	50 %	21	50 %

The global analysis of the completeness of explanations generated can be seen in Table 7 (amount of complete and incomplete explanations in each type of supporting evidence) and Table 8 (amount of supporting evidence in each type of explanations). This analysis was performed taking into account the explanations from both graphs. As it can be seen in Table 7, in total the same number of complete and incomplete explanations was obtained (21 each). However, when looking at each category separately, it is seen that when there is evidence GNNExplainer tends to produce complete explanations (68 %), and conversely when there is no supporting evidence or when the drug is contraindicated the resulting explanation is usually incomplete (62 % and 70 %, respectively). As it can be seen in Table 8, when a complete explanation is created, almost 2/3 of the time the explanation contains supporting evidence (62 %); while when the explanation is incomplete, only 1/4 of the times it contains supporting evidence (28 %).

Table 8

Number and percentage of explanations with no evidence, with supporting evidence and with contraindications in each type of explanation.

	Supporting Evidence	Percentage with Evidence	Contraindication Evidence	Percentage with Contraindications	No Evidence	Percentage No Evidence
<i>Complete Explanations</i>	13	62 %	3	14 %	5	24 %
<i>Incomplete Explanations</i>	6	28%	7	33 %	8	38%

An additional analysis was performed, this time considering each graph separately. This can be seen in Table 9 and Table S7. There is a clear difference between the explanations obtained in graph A and B. Firstly, KG A explanations are more likely to be complete (72 % in A vs 28 % in B), while KG B produces more incomplete explanations (72 % in B vs 28 % in A) (Table 9).

Table 9

Number and percentage of complete and incomplete explanations in each evidence type and in each graph.

	Evidence Type	Complete Explanations		Incomplete Explanations	
		Number	Percentage	Number	Percentage
<i>KG A</i>	Supporting Evidence	9	100%	0	0%
	Contraindication Evidence	1	17%	5	83%
	No Evidence	5	83%	1	17%
	Total	15	72%	6	28%
<i>KG B</i>	Supporting Evidence	4	40%	6	60%
	Contraindication Evidence	2	50%	2	50%
	No Evidence	0	0%	7	100%
	Total	6	28%	15	72%

An example of an explanation produced by rd-explainer can be seen in Figure 5. This explanation is classified into complete and suggests why Doxorubicin should be considered for treating respiratory insufficiency; as it is a drug that targets CHRM1 a gene that interacts with DAG1, which causes the disease. Throughout this section explanations have been classified into complete and incomplete. However, an explanation being complete does not make it a good explanation. This way, for example, an explanation of the type 'Drug A targets Gene B, Gene B interacts with Gene C, and Gene C causes Disease D' can make biological sense such as in Figure 5. On the other hand, an explanation of the type 'Drug A treats Disease B, Disease B is caused by Gene C, Gene C causes Disease D' does not make full biological sense (Drug A could treat Disease B by targeting a gene other than Gene C; this way, the same treatment could not be applied for Disease D). This is in fact what is observed in Figure S6, where disopyramide is said to treat muscular dystrophy following the next explanation: disopyramide treats urinary incontinence, affectation in DMD gene can cause urinary incontinence, and DMD gene has as phenotype muscular dystrophy. In this case, a person may have urinary incontinence for several reasons, and disopyramide may be able to treat one of them, but not necessarily the one caused by affectation in DMD gene.

The objective evaluation is undoubtedly more unbiased and equitable. Nonetheless, subjective evaluations are also significant since there are drug-phenotype interactions that are not fully understood (specially when a certain drug is producing an undesired side effect), and so they are not well established in the literature. But, analyzing the proposed explanations based on expert domain knowledge might shed light on the interaction and help to formulate a hypothesis that can be clearly designed to be tested in the wet laboratory.

After applying the objective evaluation only one explanation (levosimendan - progressive muscle weakness) was found to have supporting evidence (where levosimendan treats the disease by increasing the troponin C affinity for

calcium), and two links' explanations (doxorubicin - respiratory insufficiency and sorafenib - respiratory insufficiency) contained unclear interactions (both were of type contraindications). The results after applying this evaluation can be found summarised in the Table S6. Regarding the subjective evaluations, 17 out of 21 explanations were found to be good explanations (they were in accordance with biological reasoning) such as the one illustrated by doxorubicin - respiratory insufficiency in Figure 5; and 4 were considered bad explanations (they made no biological sense), the previously mentioned disopyramide - muscular dystrophy in Figure S6, and the explanations in Figures S7, S8 and S9.

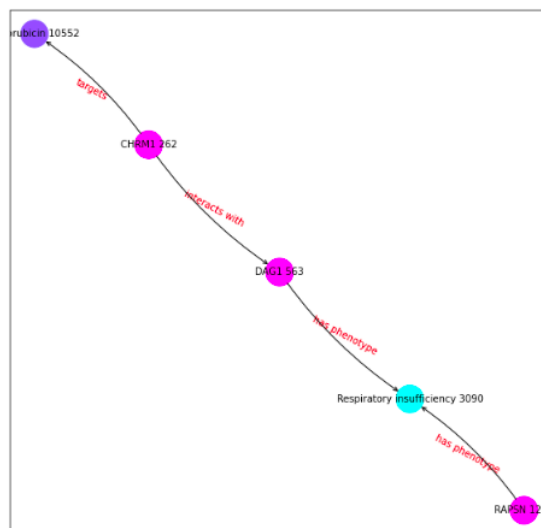


Fig. 5. Explanation of drug candidate Doxorubicin as possible treatment for Respiratory Insufficiency. Classified as complete explanation.

5.6. Generalizability of rd-explainer tested on other case studies

To show that this method can be extended to other rare diseases it was also tested in Alzheimer's Disease (AD) and Amyotrophic Lateral Sclerosis (ALS) type 1. Despite Alzheimer Disease not being a rare disease, there are different types of Alzheimer with very little prevalence. This way, for the Alzheimer's knowledge graph we used the general disease (MONDO:0004975) and all its causal genes that were present in Monarch (APP (HGNC:620), APOE (HGNC:613), PSEN1 (HGNC:9501) and PSEN2 (HGNC:9509)) as seeds. The final result would be a knowledge graph that specializes in Alzheimer diseases and that we can use to focus on the symptoms of the rare types of the disease. For the ALS type 1 knowledge graph we used the seed for the disease (MONDO:0007103) and the causal gene according to Monarch (SOD1 (HGNC:11179)). Table 10 shows the GNN performance in both diseases, showing once more a high AUROC and AUPRC for these diseases.

Table 10
Table showing different performance metrics tested in AD and ALS.

	Precision	Recall	F1-score	AUROC	AUPRC
AD	0.95	0.95	0.95	0.98	0.97
ALS	0.94	0.93	0.94	0.97	0.97

Next, the same approach that was followed for DMD was followed for both diseases: for each symptom we analysed the three drug candidates with the highest score (this drug candidates should not appear in the knowledge graph); then a literature search was performed to check if the drug candidates had been reported by the scientific

community. The complete list of phenotypes as well as the drug candidates and scores for each phenotype can be found in Tables S11 and S12. These tables also contain whether the drug candidates had supporting evidence in the literature.

Among the predictions, it is worth mentioning Pexidartinib, a drug candidate that was proposed by the model to treat memory impairment in AD and that is currently undergoing a clinical trial as a drug that could be potentially beneficial to treat the disease [40].

6. Discussion

We integrated disease-specific knowledge graphs in combination with GNN and XAI for interpretable drug repurposing. We found that state-of-the-art XAI methods based on GNNs support *in silico* predictions of candidate repurposable drugs for rare diseases by providing interpretable reasoning paths of mechanism of action. We developed *rd-explainer*, a method to perform computational drug repurposing specifically for rare diseases. It utilizes cutting-edge deep learning methods such as edge2vec and GNNs and provides drug-symptom/phenotype predictions with high performance scores, and utilizes a modified version of GNNExplainer to provide explanations as semantic graphs for the interpretability of the results. We also found that these explanations have different levels of usefulness to generate testable hypotheses: paths linking drug and phenotype nodes are more understandable versus isolated clusters since they are similar to human reasoning; adding semantics to relations adds biological meaning to help to formulate a hypothesis and design the experiment in the laboratory; and providing clear semantic graphs by removing relations that are not contributors in the learning process. We tested the generalizability of our method executing it on two additional diseases: ALS and AD. ALS type 1 was selected to test the pipeline in another monogenic disease with fewer information available. AD was selected as it is a common disease with rare subtypes that can be caused by several genes, and we wanted to test the pipeline in a polygenic and multifactorial disease. We demonstrated that our pipeline performs well on mono- and polygenic rare diseases.

rd-explainer is a researcher-centered drug repurposing method that has been demonstrated as an innovative AI based method for rare disease drug research. *rd-explainer*'s main advantage is its interpretability. The main motivation of this study was to provide explanations underlying AI predictions. *rd-explainer* provides explanations as semantic graphs, a type of explanation that resembles to human reasoning. This is in line with current research on user-centric XAI [41]. Not only does this have the high value to support rare disease researchers to formulate evidence-based hypotheses testable in the wet laboratory (and reduce cost, time and risk), but to gain new disease knowledge and speed up robust drug research. Our approach was to use state-of-the-art AI and XAI methods used in drug repurposing such as knowledge graphs to naturally represent known associations among biological entities with expressive semantics and supporting curated evidence, graph learning, and graph based XAI methods. The advance in the rare disease field is that we provide interpretable predictions thanks to a pipeline that it seamlessly integrates a graph learning model with an explainer, combining results of both model performance and explanation accuracy to mitigate the black-box problem and foster XAI adoption in the field [42]. BioKnowledge Reviewer tool provides rare disease specific knowledge graphs for disease biology data collection by means of the Monarch knowledge base API [29]. We argue that a tool or approach that can collect associations from a virtual, federated knowledge graph via APIs could extend this feature to any biomedical associations such as for drug data collection, and improve data and knowledge driven research. Another great advantage of the *rd-explainer* method is its modular implementation; this means that different parts of the workflow (data, features, GNN and explanations) can be independently modified and the pipeline can still be run. For example, if one is interested in using another node feature embedding algorithm instead of edge2vec, one can just modify that component of the pipeline and still run the rest of the workflow.

Our results showed that *rd-explainer* is a highly performant graph ML based drug repurposing method. Our method builds rare disease-specific models trained on newly generated KG for the disease of focus and enriched with data for the prediction task. In comparison with state-of-the-art AI-based drug repurposing approaches, *rd-explainer* demonstrates outstanding performance. Throughout this paper, we have compared *rd-explainer* with various AI methods that employ different techniques for their predictions, including GNNs such as GraphSAGE, random walk embeddings like edge2vec, and geometric embeddings using models like CompLEX, DistMult, and TransE.

1 By combining random walk models (edge2vec) with GNNs (GraphSAGE), rd-explainer achieves superior results 1
2 in the link prediction task. Notably, edge2vec outperforms GraphSAGE, suggesting that the exceptional perfor- 2
3 mance of rd-explainer is primarily attributed to the random walk model, with the GNN providing an additional 3
4 performance boost. This level of performance rivals other models developed for drug repurposing, such as deepDR 4
5 (AUROC = 0.908) [43]. Although there are benchmarks and frameworks to evaluate the performance of GNNs [44– 5
6 49] to the best of our knowledge there is no a standard for drug repurposing, and makes it challenging to directly 6
7 compare rd-explainer to other methods due to one of its key features: the creation of high-quality disease specific 7
8 knowledge graphs. These knowledge graphs are enriched with data from a wide array of sources including domain 8
9 expert knowledge via the seed nodes, and curated known relations among genes, anatomical structures, biological 9
10 processes and diseases not only from humans, but also importantly numerous other species to fill the lack of 10
11 molecular knowledge. This comprehensive approach significantly boosts the graph’s richness and diversity, making 11
12 it a valuable resource for tackling rare diseases, which often suffer from limited research attention. By maximizing 12
13 the information available, rd-explainer enhances our ability to identify potential treatments for these understudied 13
14 conditions and, ultimately, enable more effective and faster translation. Conversely, Huang et al. recently proposed 14
15 a clinician-centered drug repurposing foundation model pre-trained on a medical KG composed of 17.000 diseases 15
16 and transfer learning by disease mechanism similarity [50]. It would be interesting to combine both approaches and 16
17 investigate the effect of extending our KGs with similar disease networks from well-known diseases. 17

18 Our new predictions are valid drug candidates since they are consistent with recent findings in the literature. 18
19 We demonstrated that rd-explainer can provide new interesting drug-phenotype predictions. For instance, Sunitinib, 19
20 one of the drugs that appear to be a good candidate to treat the symptoms of the disease according to both models 20
21 (using KG A and KG B), has been considered as a good drug candidate to treat DMD and in 2019 appeared to 21
22 be in preclinical trials [51]. This drug belongs to the group of tyrosine kinase inhibitors, and many other drugs 22
23 that belong to this category have been proposed by our model (Fedratinib, Sorafenib, Bosutinib, Ruxolitinib and 23
24 Midostaurin). Similarly, Mezlocillin, an antibiotic used to treat gram-negative bacterial infections, has also been 24
25 proposed by our model; while Gentamicin, another gram-negative antibiotic, was in 2019 in clinical trials to treat 25
26 DMD [51]. This way, despite not producing drugs candidates that are undergoing a clinical trial or treating the 26
27 disease, it produces drug candidates that participate in similar biological processes (i.e., tyrosine kinases inhibitor, 27
28 gram negative antibiotics) 28

29 Importantly, explanations for hypothesis generation may enable to move towards *lab-in-the-loop* framework. 29
30 With respect to the interpretability and utility of explanations, one of the 21 examined explanations was supported 30
31 by evidence in the literature. Nonetheless, this does not mean that the explanations are useless. A good example of 31
32 this would be the explanation for the Methylprednisolone-Muscular Dystrophy link (Figure S10). The explanation 32
33 is simple: ‘Methylprednisolone treats DMD, DMD has Muscular Dystrophy as phenotype; thus methylprednisolone 33
34 can treat Muscular Dystrophy’. In this case the explanation does not contain supporting evidence but the explanation 34
35 still makes sense. In the literature, methylprednisolone is said to be a good candidate to treat muscular dystrophies 35
36 because it interacts with the glucocorticoid receptor and this leads to the activation of anti-inflammatory signaling 36
37 and the inhibition of proinflammatory signaling [52]. The explanation proposed by rd-explainer doesn’t provide the 37
38 underlying causative mechanism that relates methylprednisolone and muscular dystrophy, but a researcher can still 38
39 be able to see that muscular dystrophies and methylprednisolone are interrelated. This illustrates how even though 39
40 an explanation may lack comprehensive supporting evidence, it can still provide valuable directional cues for further 40
41 more precise investigation. Another important aspect is that rare disease findings in the lab can be introduced back in 41
42 the knowledge graph to update and improve the disease specific AI model for continual learning and enabling precise 42
43 experimental design. Besides, this synergy fosters collaboration between computational and wet lab researchers to 43
44 increase efficiency for disease specific drug research [29]. 44

45 Finally, we found that knowledge graph topology has an impact on explainability. It was also seen that KG A usu- 45
46 ally produces more complete explanations, while in KG B incomplete explanations appear to be more numerous. 46
47 This could happen due to the difference in the graph structure itself: graph A has a smaller clustering coefficient than 47
48 graph B (see Section 5.1), which leads to more edges being present in the subgraphs produced by GNNExplainer. 48
49 This way, because the 15th edges with the highest scores are selected, it is more likely to find a path between drug 49
50 and phenotype in KG A than in B. Another interesting difference is that explanations generated with KG A tend 50
51 to have a higher ‘sensitivity’, while explanations generated with KG B tend to have a higher ‘specificity’. When an 51

1 incomplete explanation is produced using graph A it is very unlikely that the explanation will contain supporting 1
2 evidence (0 explanations were found to have evidence if the explanation was incomplete in KG A). Similarly, when 2
3 a complete explanation is produced in KG B, it is very likely that the explanations have supporting evidence or 3
4 contraindication evidence (67% of complete explanations had supporting evidence and 33% of complete explanations 4
5 had contraindication evidence). For this reason, if one remains skeptical about the explanations themselves, 5
6 this quality of the explanations might be used as filter/validation. For example, if an incomplete explanation is ob- 6
7 tained with KG A, it is unlikely that it is trustworthy (none of the incomplete explanations had supporting evidence). 7
8 Similarly, if a complete explanation is obtained using KG B, it is likely that there is some interaction between the 8
9 drug and the phenotype (all of the complete explanations generated with graph B had either supporting or con- 9
10 traindication evidence). Our findings are aligned with recent studies where the influence of clustering coefficient 10
11 and topology has been observed in embedding-based predictions [53, 54], here we extend these observations to its 11
12 impact on graph-based explanations. 12

13 *Limitations and future directions* 13

14 An important limitation of this study is that we only utilize one XAI method, which is not model agnostic. 14
15 XAI is a hot research topic in the AI field, where new and more sophisticated methods are frequently published 15
16 [55]. It would be good to extend our study to other XAI types to check how applicable they are given the unique 16
17 characteristics of rare diseases, including limited annotated data, lack of knowledge of pertinent entity relations, 17
18 and lack of a gold drug-phenotype standard. Another important limitation is the lack of standard benchmarking and 18
19 metrics to systematically evaluate explainers and explanations. Currently, there are some initial efforts going in this 19
20 direction [34, 56–60], but there is still a lack of a common standard [61]. The known reproducibility issue of our 20
21 explainer [34] that may imply that the explanations are different each time it is used, may reduce the confidence and 21
22 reliance on the explanations. We did several experiments to try and bring consistency to explanations; for example, 22
23 executing GNNExplainer several times and using the mean mask as the final mask or increasing the number of 23
24 epochs. However, this still did not solve the issue. This experience makes us strongly recommend to work on the 24
25 standard evaluation of explanations by the XAI community to foster trust on the application of AI in bioinformatics 25
26 and biomedicine. Additionally, many times the explanation would consist in a subgraph where the two targeted 26
27 nodes would be disconnected from each other, which might bring confusion and could be seen as a ‘bad’ explanation. 27
28 Therefore, work towards methods that prioritize or focus on providing just connecting paths such as metapath based 28
29 ones [62–66] and on improving path visualisation for user interpretation [67–69] is arguably recommended. Finally, 29
30 while we focused primarily on integrating a graph ML model with an explainer, a clear line of research will be 30
31 to work on interpretability and reproducibility of explanations in the context of the drug repurposing task. The 31
32 reproducibility/ inconsistency could be affected by the size and complexity of our data. This inconsistency could make 32
33 the users of this pipeline skeptical about its explanations and for this reason more investigation should be done in 33
34 this element of the pipeline to make it a more robust model. To improve this, ontologies could be incorporated into 34
35 the knowledge graph to increase the quality and interpretability of our data. Ontologies help to standardize data 35
36 into the shared meaning by a community enhancing thus interpretability by domain users. Importantly, the formal 36
37 description of knowledge embedded in ontologies can be leveraged for data consistency checking, and for inference 37
38 to add implicit knowledge into the graph [70]. Nonetheless, knowledge graph and ontology changes pose a great 38
39 interoperability challenge to the community to keep up downstream bioinformatics and data science workflows and 39
40 analyses [71, 72]. Finally, it would make our work more ‘FAIR’ [73], i.e., not only understandable by humans, but 40
41 also by machines, by providing our drug repurposing for DMD KG from a FAIR data point [74], and rd-explainer 41
42 from workflowHub [75]. 42
43 43
44 44
45 45

46 **7. Conclusion** 46

47 We present the application of explainable AI on state-of-the-art computational drug repurposing for rare diseases. 47
48 Our knowledge graph based deep learning method provides human understandable explanations for the phenotype- 48
49 drug link prediction and we demonstrated that graph XAI can be applied to rare diseases. The *rd-explainer* method 49
50 provides an innovative approach that can maximize the available disease-specific knowledge and generate valuable 50
51 51

1 predictions with its explanations. Our model has proven to obtain high evaluation scores, providing drug candidates 1
2 that are often supported by evidence. The key contribution of our study is that our pipeline gives possible explana- 2
3 tions in the form of semantic graphs that may help rare disease researchers to make informed decisions to exper- 3
4 imentally validate deep learning model predictions. However, we detected that data topology affects explanations, 4
5 highlighting the importance of investigating further how best represent graphical knowledge for model performance 5
6 and explanation accuracy. *rd-explainer* can be extended to other rare diseases and provide computer-aided guidance 6
7 for biologists and accelerate translational research. Finally, future studies should advance our understanding of the 7
8 necessary standard mechanism to evaluate explainability to foster adoption from domain experts and to mitigate 8
9 the black-box problem of trust on AI, especially for biomedicine where decisions can have an important impact on 9
10 people's lives. 10
11

12 **8. Code availability** 12

13 The code to run the repurposing pipeline is available at <https://github.com/PPerdomoQ/rare-disease-explainer>. 13
14 14
15 15
16 16
17 17
18 18
19 19
20 20
21 21
22 22
23 23
24 24
25 25
26 26
27 27
28 28
29 29
30 30
31 31
32 32
33 33
34 34
35 35
36 36
37 37
38 38
39 39
40 40
41 41
42 42
43 43
44 44
45 45
46 46
47 47
48 48
49 49
50 50
51 51

References

- [1] Duxin Sun, Wei Gao, Hongxiang Hu, and Simon Zhou. Why 90 *Acta Pharmaceutica Sinica. B*, 12:3049, 7 2022. ISSN 22113843. . URL [/pmc/articles/PMC9293739/](https://www.ncbi.nlm.nih.gov/pmc/articles/PMC9293739/).
- [2] Melissa Haendel, Nicole Vasilevsky, Deepak Unni, Cristian Bologna, Nomi Harris, Heidi Rehm, Ada Hamosh, Gareth Baynam, Tudor Groza, Julie McMurry, Hugh Dawkins, Ana Rath, Courtney Thaxon, Giovanni Bocci, Marcin P. Joachimiak, Sebastian Köhler, Peter N. Robinson, Chris Mungall, and Tudor I. Oprea. How many rare diseases are there? *19(2):77–78*. ISSN 1474-1776. . URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7771654/>.
- [3] Rare diseases. . URL https://ec.europa.eu/health/non-communicable-diseases/steering-group/rare-diseases_en.
- [4] Emre Guney, Jörg Menche, Marc Vidal, and Albert-László Barabási. Network-based in silico drug efficacy screening. *7(1):10331*. ISSN 2041-1723. . URL <https://www.nature.com/articles/ncomms10331>. Number: 1 Publisher: Nature Publishing Group.
- [5] Bryce Goodman and Seth Flaxman. European Union Regulations on Algorithmic Decision-Making and a 'Right to Explanation'. *AI Magazine*, 38(3):50–57, sep 2017. ISSN 07384602. . URL <https://api.semanticscholar.org/CorpusID:7373959#id=name=S2CID>.
- [6] Qiang Huang, Makoto Yamada, Yuan Tian, Dinesh Singh, Dawei Yin, and Yi Chang. GraphLIME: Local interpretable model explanations for graph neural networks. (arXiv:2001.06216). . URL <http://arxiv.org/abs/2001.06216>. type: article.
- [7] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. "why should i trust you?": Explaining the predictions of any classifier. (arXiv:1602.04938). . URL <http://arxiv.org/abs/1602.04938>. type: article.
- [8] Pouya Pezeshkpour, Yifan Tian, and Sameer Singh. Investigating robustness and interpretability of link prediction via adversarial modifications. (arXiv:1905.00563). . URL <http://arxiv.org/abs/1905.00563>. type: article.
- [9] Cristina Al-Khalili Szgyarto and Pietro Spitali. Biomarkers of duchenne muscular dystrophy: current findings. *Degenerative Neurological and Neuromuscular Disease*, 8:1–13, 1 2018. . URL <https://www.dovepress.com/biomarkers-of-duchenne-muscular-dystrophy-current-findings-peer-reviewed-fulltext-article-DNND>.
- [10] Jacob Al-Saleem, Roger Granet, Srinivasan Ramakrishnan, Natalie A. Ciancetta, Catherine Saveson, Chris Gessner, and Qiongqiong Zhou. Knowledge graph-based approaches to drug repurposing for COVID-19. page acs.jcim.1c00642. ISSN 1549-9596. . URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC8340579/>.
- [11] Aditya Grover and Jure Leskovec. node2vec: Scalable feature learning for networks. (arXiv:1607.00653). URL <http://arxiv.org/abs/1607.00653>. type: article.
- [12] Bishan Yang, Wen-tau Yih, Xiaodong He, Jianfeng Gao, and Li Deng. Embedding entities and relations for learning and inference in knowledge bases. (arXiv:1412.6575). . URL <http://arxiv.org/abs/1412.6575>. type: article.
- [13] Xiang Yue, Zhen Wang, Jingong Huang, Srinivasan Parthasarathy, Soheil Moosavinasab, Yungui Huang, Simon M. Lin, Wen Zhang, Ping Zhang, and Huan Sun. Graph embedding on biomedical networks: methods, applications and evaluations. *Bioinformatics*, 36:1241–1251, 2 2020. ISSN 1367-4803. . URL <https://dx.doi.org/10.1093/bioinformatics/btz718>.
- [14] Ilaria Ferrari, Giacomo Frisoni, Paolo Italiani, Gianluca Moro, and Claudio Sartori. Comprehensive analysis of knowledge graph embedding techniques benchmarked on link prediction. *Electronics 2022, Vol. 11, Page 3866*, 11:3866, 11 2022. ISSN 2079-9292. . URL <https://www.mdpi.com/2079-9292/11/23/3866/html>.
- [15] Shaghayegh Sadeghi, Jianguo Lu, and Alioune Ngom. An integrative heterogeneous graph neural network-based method for multi-labeled drug repurposing. *Frontiers in Pharmacology*, 13:908549, 7 2022. ISSN 16639812. .
- [16] Yuchen Zhang, Xiujian Lei, Yi Pan, and Fang Xiang Wu. Drug repositioning with graphsage and clustering constraints based on drug and disease networks. *Frontiers in Pharmacology*, 13:872785, 5 2022. ISSN 16639812. .
- [17] Kanglin Hsieh, Yinyin Wang, Luyao Chen, Zhongming Zhao, Sean Savitz, Xiaoqian Jiang, Jing Tang, and Yejin Kim. Drug repurposing for covid-19 using graph neural network and harmonizing multiple evidence. *Scientific Reports 2021 11:1*, 11:1–13, 11 2021. ISSN 2045-2322. . URL <https://www.nature.com/articles/s41598-021-02353-5>.
- [18] Junjun Zhang and Minzhu Xie. Graph regularized non-negative matrix factorization with prior knowledge consistency constraint for drug–target interactions prediction. *BMC Bioinformatics*, 23, 12 2022. ISSN 14712105. . URL [/pmc/articles/PMC9798666/](https://www.ncbi.nlm.nih.gov/pmc/articles/PMC9798666/).
- [19] Bishan Yang, Wen tau Yih, Xiaodong He, Jianfeng Gao, and Li Deng. Embedding entities and relations for learning and inference in knowledge bases. *3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings*, 12 2014. URL <https://arxiv.org/abs/1412.6575v4>.
- [20] William L. Hamilton, Rex Ying, and Jure Leskovec. Inductive Representation Learning on Large Graphs. *Advances in Neural Information Processing Systems*, 2017–December:1025–1035, jun 2017. ISSN 10495258. . URL <https://arxiv.org/abs/1706.02216v4>.
- [21] Rex Ying, Dylan Bourgeois, Jiaxuan You, Marinka Zitnik, and Jure Leskovec. GNNExplainer: Generating explanations for graph neural networks. (arXiv:1903.03894). . URL <http://arxiv.org/abs/1903.03894>. type: article.
- [22] Bastian Pfeifer, Anna Saranti, and Andreas Holzinger. Gnn-subnet: disease subnetwork detection with explainable graph neural networks. *Bioinformatics*, 38:ii120–ii126, 9 2022. ISSN 1367-4803. . URL <https://dx.doi.org/10.1093/bioinformatics/btac478>.
- [23] Yu Sun, Xiang Xu, Lin Lin, Kang Xu, Yang Zheng, Chao Ren, Huan Tao, Xu Wang, Huan Zhao, Weiwei Tu, Xuemei Bai, Junting Wang, Qiya Huang, Yaru Li, Hebing Chen, Hao Li, and Xiaochen Bo. A graph neural network-based interpretable framework reveals a novel dna fragility-associated chromatin structural unit. *Genome Biology*, 24, 12 2023. ISSN 1474760X. . URL [/pmc/articles/PMC10124043/](https://www.ncbi.nlm.nih.gov/pmc/articles/PMC10124043/).
- [24] So Yeon Kim. Personalized explanations for early diagnosis of alzheimer's disease using explainable graph neural networks with population graphs. *Bioengineering*, 10, 6 2023. ISSN 23065354. . URL [/pmc/articles/PMC10295378/](https://www.ncbi.nlm.nih.gov/pmc/articles/PMC10295378/).

- [25] Monarch initiative explorer, . URL <https://monarchinitiative.org/>. 1
- [26] Sorin Avram, Cristian G Bologa, Jayme Holmes, Giovanni Bocci, Thomas B Wilson, Dac-Trung Nguyen, Ramona Curpan, Liliana Halip, Alina Bora, Jeremy J Yang, Jeffrey Knockel, Suman Sirimulla, Oleg Ursu, and Tudor I Oprea. DrugCentral 2021 supports drug discovery and repositioning. 49:D1160–D1169. ISSN 0305-1048, 1362-4962. . URL <https://academic.oup.com/nar/article/49/D1/D1160/5957163>. 2
- [27] Ying Zhou, Yintao Zhang, Xichen Lian, Fengcheng Li, Chaoxin Wang, Feng Zhu, Yunqing Qiu, and Yuzong Chen. Therapeutic target database update 2022: facilitating drug discovery with enriched comparative data of targeted agents. 50:D1398–D1407. ISSN 0305-1048, 1362-4962. . URL <https://academic.oup.com/nar/article/50/D1/D1398/6413598>. 3
- [28] Zheng Gao, Gang Fu, Chunping Ouyang, Satoshi Tsutsui, Xiaozhong Liu, Jeremy Yang, Christopher Gessner, Brian Foote, David Wild, Ying Ding, and Qi Yu. edge2vec: Representation learning using edge semantics for biomedical knowledge discovery. 20(1):306. ISSN 1471-2105. . URL <https://doi.org/10.1186/s12859-019-2914-2>. 4
- [29] Núria Queralt-Rosinach, Gregory S. Stupp, Tong Shu Li, Michael Mayers, Maureen E. Hoatlin, Matthew Might, Benjamin M. Good, and Andrew I. Su. Structured reviews for data and knowledge-driven research. 2020:baaa015. ISSN 1758-0463. . 5
- [30] NetworkX — NetworkX documentation, . URL <https://networkx.org/>. 6
- [31] CS224W Machine Learning with Graph Home. URL <http://web.stanford.edu/class/cs224w/>. 7
- [32] DeepSNAP documentation — DeepSNAP 0.2.0 documentation, . URL <https://snap.stanford.edu/deepsnap/>. 8
- [33] Richard Liaw, Eric Liang, Robert Nishihara, Philipp Moritz, Joseph E. Gonzalez, and Ion Stoica. Tune: A research platform for distributed model selection and training. 7 2018. URL <http://arxiv.org/abs/1807.05118>. 9
- [34] Chirag Agarwal, Owen Queen, Himabindu Lakkaraju, and Marinka Zitnik. Evaluating explainability for graph neural networks. *Scientific Data* 2023 10:1, 10:1–18, 3 2023. ISSN 2052-4463. . URL <https://www.nature.com/articles/s41597-023-01974-x>. 10
- [35] Théo Trouillon, Johannes Welbl, Sebastian Riedel, Eric Gaussier, and Guillaume Bouchard. Complex embeddings for simple link prediction. In Maria Florina Balcan and Kilian Q. Weinberger, editors, *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 2071–2080, New York, New York, USA, 20–22 Jun 2016. PMLR. URL <https://proceedings.mlr.press/v48/trouillon16.html>. 11
- [36] 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*, volume 26. Curran Associates, Inc. URL <https://proceedings.neurips.cc/paper/2013/hash/1cecc7a77928ca8133fa24680a88d2f9-Abstract.html>. 12
- [37] Yang Yang, Ryan N. Lichtenwalter, and Nitesh V. Chawla. Evaluating link prediction methods. *Knowledge and Information Systems*, 45: 751–782, 5 2015. ISSN 02193116. . URL <https://arxiv.org/abs/1505.04094v1>. 13
- [38] Ruthwik R. Junuthula, Kevin S. Xu, and Vijay K. Devabhaktuni. Evaluating link prediction accuracy on dynamic networks with added and removed edges. *Proceedings - 2016 IEEE International Conferences on Big Data and Cloud Computing, BDCloud 2016, Social Computing and Networking, SocialCom 2016 and Sustainable Computing and Communications, SustainCom 2016*, pages 377–384, 7 2016. . URL <https://arxiv.org/abs/1607.07330v1>. 14
- [39] Aniek F. Markus, Jan A. Kors, and Peter R. Rijnbeek. The role of explainability in creating trustworthy artificial intelligence for health care: A comprehensive survey of the terminology, design choices, and evaluation strategies. *Journal of Biomedical Informatics*, 113: 103655, 1 2021. ISSN 1532-0464. . 15
- [40] Antonio Ancidoni, Ilaria Bacigalupo, Giulia Remoli, Eleonora Lacorte, Paola Piscopo, Giulia Sarti, Massimo Corbo, Nicola Vanacore, and Marco Canevelli. Anticancer drugs repurposed for alzheimer’s disease: a systematic review. *Alzheimer’s Research & Therapy*, 13, 12 2021. ISSN 17589193. . URL <https://pubmed.ncbi.nlm.nih.gov/pmc/articles/PMC8101105/>. 16
- [41] Qianwen Wang, Kexin Huang, Payal Chandak, Marinka Zitnik, and Nils Gehlenborg. Extending the nested model for user-centric xai: A design study on gnn-based drug repurposing. *IEEE Transactions on Visualization and Computer Graphics*, 29(1):1266–1276, 2023. . 17
- [42] Claudio Borile, Alan Perotti, and André Panisson. Evaluating link prediction explanations for graph neural networks. *CoRR*, abs/2308.01682, 2023. . URL <https://doi.org/10.48550/arXiv.2308.01682>. 18
- [43] Xiangxiang Zeng, Siyi Zhu, Xiangrong Liu, Yadi Zhou, Ruth Nussinov, and Feixiong Cheng. deepdr: a network-based deep learning approach to in silico drug repositioning. *Bioinformatics*, 35:5191–5198, 12 2019. ISSN 1367-4803. . URL <https://dx.doi.org/10.1093/bioinformatics/btz418>. 19
- [44] Weihua Hu, Matthias Fey, Marinka Zitnik, Yuxiao Dong, Hongyu Ren, Bowen Liu, Michele Catasta, and Jure Leskovec. Open graph benchmark: datasets for machine learning on graphs. In *Proceedings of the 34th International Conference on Neural Information Processing Systems, NIPS ’20*, Red Hook, NY, USA, 2020. Curran Associates Inc. ISBN 9781713829546. 20
- [45] Wentao Zhang, Zeang Sheng, Yuezihan Jiang, Yikuan Xia, Jun Gao, Zhi Yang, and Bin Cui. Evaluating deep graph neural networks. 2021. 21
- [46] Vijay Prakash Dwivedi, Chaitanya K. Joshi, Anh Tuan Luu, Thomas Laurent, Yoshua Bengio, and Xavier Bresson. Benchmarking graph neural networks. *Journal of Machine Learning Research*, 24(43):1–48, 2023. URL <http://jmlr.org/papers/v24/22-0567.html>. 22
- [47] Zhiyao Zhou, Sheng Zhou, Bochao Mao, Xuanyi Zhou, Jiawei Chen, Qiaoyu Tan, Daochen Zha, Yan Feng, Chun Chen, and Can Wang. Opengsl: a comprehensive benchmark for graph structure learning. In *Proceedings of the 37th International Conference on Neural Information Processing Systems, NIPS ’23*, Red Hook, NY, USA, 2024. Curran Associates Inc. 23
- [48] Juanhui Li, Harry Shomer, Haitao Mao, Shenglai Zeng, Yao Ma, Neil Shah, Jiliang Tang, and Dawei Yin. Evaluating graph neural networks for link prediction: current pitfalls and new benchmarking. In *Proceedings of the 37th International Conference on Neural Information Processing Systems, NIPS ’23*, Red Hook, NY, USA, 2024. Curran Associates Inc. 24

- [49] Xin Zheng, Miao Zhang, Chunyang Chen, Soheila Molaei, Chuan Zhou, and Shirui Pan. Gnevaluator: evaluating gnn performance on unseen graphs without labels. In *Proceedings of the 37th International Conference on Neural Information Processing Systems, NIPS '23*, Red Hook, NY, USA, 2024. Curran Associates Inc.
- [50] Kexin Huang, Payal Chandak, Qianwen Wang, Shreyas Havaldar, Akhil Vaid, Jure Leskovec, Girish N Nadkarni, Benjamin S Glicksberg, Nils Gehlenborg, and Marinka Zitnik. A foundation model for clinician-centered drug repurposing. *Nat. Med.*, September 2024.
- [51] Libero Vitiello, Lucia Tibaudo, Elena Pegoraro, Luca Bello, and Marcella Canton. Teaching an old molecule new tricks: Drug repositioning for duchenne muscular dystrophy. 20(23):6053. ISSN 1422-0067. . URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6929176/>.
- [52] Mattia Quattrocchi, Aaron S. Zelikovich, Isabella M. Salamone, Julie A. Fischer, and Elizabeth M. McNally. Mechanisms and clinical applications of glucocorticoid steroids in muscular dystrophy. *Journal of Neuromuscular Diseases*, 8:39, 2021. ISSN 22143602. . URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7902991/>.
- [53] Anand Kumar Gupta and Neetu Sardana. Significance of clustering coefficient over jaccard index. In *2015 Eighth International Conference on Contemporary Computing (IC3)*. IEEE, August 2015.
- [54] Omar F Robledo, Xiu-Xiu Zhan, Alan Hanjalic, and Huijuan Wang. Influence of clustering coefficient on network embedding in link prediction. *Appl. Netw. Sci.*, 7(1), December 2022.
- [55] A. Saranya and R. Subhashini. A systematic review of explainable artificial intelligence models and applications: Recent developments and future trends. *Decision Analytics Journal*, 7:100230, 6 2023. ISSN 2772-6622. .
- [56] Chirag Agarwal, Satyapriya Krishna, Eshika Saxena, Martin Pawelczyk, Nari Johnson, Isha Puri, Marinka Zitnik, and Himabindu Lakkaraju. Openxai: Towards a transparent evaluation of model explanations. *36th Conference on Neural Information Processing Systems (NeurIPS 2022) Track on Datasets and Benchmarks*, 6 2022. URL <https://arxiv.org/abs/2206.11104v3>.
- [57] Thomas Fel, Lucas Hervier, David Vigouroux, Antonin Poche, Justin Plakoo, Remi Cadene, Mathieu Chalvidal, Julien Colin, Thibaut Boissin, Louis Bethune, Agustin Picard, Claire Nicodeme, Laurent Gardes, Gregory Flandin, and Thomas Serre. Xplique: A deep learning explainability toolbox. 2022.
- [58] Hao Yuan, Haiyang Yu, Shurui Gui, and Shuiwang Ji. Explainability in graph neural networks: A taxonomic survey. *IEEE Trans. Pattern Anal. Mach. Intell.*, 45(5):5782–5799, May 2023.
- [59] Claudio Borile, Alan Perotti, and André Panisson. Evaluating link prediction explanations for graph neural networks. August 2023.
- [60] Daniel Daza, Cuong Xuan Chu, Trung-Kien Tran, Daria Stepanova, Michael Cochez, and Paul Groth. Explaining graph neural networks for node similarity on graphs. July 2024.
- [61] Mehwish Alam, Frank van Harmelen, and Maribel Acosta. Towards semantically enriched embeddings for knowledge graph completion. July 2023.
- [62] 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, September 2017.
- [63] Xinyu Fu, Jiani Zhang, Ziqiao Meng, and Irwin King. Maggn: Metapath aggregated graph neural network for heterogeneous graph embedding. *The Web Conference 2020 - Proceedings of the World Wide Web Conference, WWW 2020*, pages 2331–2341, 2 2020. . URL <http://arxiv.org/abs/2002.01680><http://dx.doi.org/10.1145/3366423.3380297>.
- [64] Michael Mayers, Roger Tu, Dylan Steinecke, Tong Shu Li, Nria Queralt-Rosinach, and Andrew I Su. Design and application of a knowledge network for automatic prioritization of drug mechanisms. *Bioinformatics*, 38(10):2880–2891, April 2022.
- [65] Ayush Noori, Michelle M Li, Amelia L M Tan, and Marinka Zitnik. Metapaths: similarity search in heterogeneous knowledge graphs via meta-paths. *Bioinformatics*, 39(5), May 2023.
- [66] Ana Jimenez, Mara Jos Merino, Juan Parras, and Santiago Zazo. Explainable drug repurposing via path based knowledge graph completion. *Sci. Rep.*, 14(1):16587, July 2024.
- [67] Qianwen Wang, Kexin Huang, P Chandak, Nils Gehlenborg, and M Zitnik. Interactive visual explanations for deep drug repurposing. <https://icml.cc/virtual/2021/workshop/8358>, 2021. Accessed: 2024-10-8.
- [68] Daniel S Himmelstein, Michael Zietz, Vincent Rubinetti, Kyle Kloster, Benjamin J Heil, Faisal Alquaddoomi, Dongbo Hu, David N Nicholson, Yun Hao, Blair D Sullivan, Michael W Nagle, and Casey S Greene. Hetnet connectivity search provides rapid insights into how biomedical entities are related. *Gigascience*, 12, December 2022.
- [69] Qianwen Wang, Kexin Huang, Payal Chandak, Marinka Zitnik, and Nils Gehlenborg. Extending the nested model for user-centric xai: A design study on gnn-based drug repurposing. *IEEE Transactions on Visualization and Computer Graphics*, 29(1):1266–1276, 2023. .
- [70] Mona Alshahrani, Mohammad Asif Khan, Omar Maddouri, Akira R Kinjo, Nria Queralt-Rosinach, and Robert Hoehndorf. Neuro-symbolic representation learning on biological knowledge graphs. *Bioinformatics*, 33(17):2723–2730, 04 2017. ISSN 1367-4803. . URL <https://doi.org/10.1093/bioinformatics/btx275>.
- [71] Deepak Unni, Vasundra Tour, Philip Krauss, Katrin Crameri, and Sabine sterle. SPHN strategy to unravel the semantic drift between versions of standard terminologies. December 2023.
- [72] Harshad Hegde, Jennifer Vendetti, Damien Goutte-Gattat, J Harry Caufield, John B Graybeal, Nomi L Harris, Naouel Karam, Christian Kindermann, Nicolas Matentzoglou, James A Overton, Mark A Musen, and Christopher J Mungall. A change language for ontologies and knowledge graphs. 2024.
- [73] Mark D. Wilkinson, Michel Dumontier, IJsbrand Jan Aalbersberg, Gabrielle Appleton, Myles Axton, Arie Baak, Niklas Blomberg, Jan Willem Boiten, Luiz Bonino da Silva Santos, Philip E. Bourne, Jildau Bouwman, Anthony J. Brookes, Tim Clark, Merc Crosas, Ingrid Dillo, Olivier Dumon, Scott Edmunds, Chris T. Evelo, Richard Finkers, Alejandra Gonzalez-Beltran, Alasdair J.G. Gray, Paul Groth, Carole Goble, Jeffrey S. Grethe, Jaap Heringa, Peter A.C. t Hoen, Rob Hoof, Tobias Kuhn, Ruben Kok, Joost Kok, Scott J. Lusher, Maryann E. Martone, Albert Mons, Abel L. Packer, Bengt Persson, Philippe Rocca-Serra, Marco Roos, Rene van Schaik, Susanna Assunta

- 1 Sansone, Erik Schultes, Thierry Sengstag, Ted Slater, George Strawn, Morris A. Swertz, Mark Thompson, Johan Van Der Lei, Erik Van
2 Mulligen, Jan Velterop, Andra Waagmeester, Peter Wittenburg, Katherine Wolstencroft, Jun Zhao, and Barend Mons. The FAIR Guiding
3 Principles for scientific data management and stewardship. *Scientific Data* 2016 3:1, 3(1):1–9, mar 2016. ISSN 2052-4463. . URL
4 <https://www.nature.com/articles/sdata201618>.
- [74] Luiz Olavo Bonino da Silva Santos, Kees Burger, Rajaram Kaliyaperumal, and Mark D. Wilkinson. Fair data point: A fair-oriented
5 approach for metadata publication. *Data Intelligence*, 5:163–183, 12 2023. ISSN 2641435X. . URL [https://dx.doi.org/10.1162/dint_a_](https://dx.doi.org/10.1162/dint_a_00160)
6 [00160](https://dx.doi.org/10.1162/dint_a_00160).
- [75] Rafael Ferreira da Silva, Loïc Pottier, Tainã Coleman, Ewa Deelman, and Henri Casanova. Workflowhub: Community framework for
7 enabling scientific workflow research and development – technical report. *Proceedings of 15th Workshop on Workflows in Support of*
8 *Large-Scale Science, WORKS 2020 - Held in conjunction with SC 2020: The International Conference for High Performance Computing,*
9 *Networking, Storage and Analysis*, pages 49–56, 9 2020. . URL [http://arxiv.org/abs/2009.00250http://dx.doi.org/10.1109/WORKS5191](http://arxiv.org/abs/2009.00250http://dx.doi.org/10.1109/WORKS51914.2020.00012)
10 [4.2020.00012](http://arxiv.org/abs/2009.00250http://dx.doi.org/10.1109/WORKS51914.2020.00012).
- [76] H. Hanke and D. Knees. A phase-field damage model based on evolving microstructure. *Asymptotic Analysis*, 101:149–180, 2017. 11
- [77] E. Lefever. A hybrid approach to domain-independent taxonomy learning. *Applied Ontology*, 11(3):255–278, 2016. 12
- [78] P.S. Meltzer, A. Kallioniemi, and J.M. Trent. Chromosome alterations in human solid tumors. In B. Vogelstein and K.W. Kinzler, editors,
13 *The Genetic Basis of Human Cancer*, pages 93–113. McGraw-Hill, New York, 2002. 14
- [79] P.R. Murray, K.S. Rosenthal, G.S. Kobayashi, and M.A. Pfaller. *Medical Microbiology*. Mosby, St. Louis, 4th edition, 2002. 15
- [80] E. Wilson. *Active vibration analysis of thin-walled beams*. PhD thesis, University of Virginia, 1991. 16
- [81] S. Fortunato. Community detection in graphs. *Phys. Rep.-Rev. Sec. Phys. Lett.*, 486:75–174, 2010. 17
- [82] M. E. J. Newman and M. Girvan. Finding and evaluating community structure in networks. *Phys. Rev. E.*, 69:026113, 2004. 18
- [83] C. Vehlow, T. Reinhardt, and D. Weiskopf. Visualizing fuzzy overlapping communities in networks. *IEEE Trans. Vis. Comput. Graph.*,
19 19:2486–2495, 2013. 20
- [84] U. Raghavan, R. Albert, and S. Kumara. Near linear time algorithm to detect community structures in large-scale networks. *Phys. Rev E.*,
21 76:036106, 2007. 22
- [85] L. Šubelj and M. Bajec. Robust network community detection using balanced propagation. *Eur. Phys. J. B.*, 81:353–362, 2011. 23
- [86] H. Lou, S. Li, and Y. Zhao. Detecting community structure using label propagation with weighted coherent neighborhood propinquity.
24 *Physica A.*, 392:3095–3105, 2013. 25
- [87] A. Clauset, M. E. J. Newman, and C. Moore. Finding community structure in very large networks. *Phys. Rev. E.*, 70:066111, 2004. 26
- [88] V. D. Blondel, J. L. Guillaume, R. Lambiotte, and E. Lefebvre. Fast unfolding of communities in large networks. *J. Stat. Mech.-Theory*
27 *Exp.*, 2008:P10008, 2008. 28
- [89] S. Sobolevsky and R. Campari. General optimization technique for high-quality community detection in complex networks. *Phys. Rev.*
29 *E.*, 90:012811, 2014. 30
- [90] S. Fortunato and M. Barthelemy. Resolution limit in community detection. *Proc. Natl. Acad. Sci. U. S. A.*, 104:36–41, 2007. 31
- [91] L. Šubelj and M. Bajec. Unfolding communities in large complex networks: Combining defensive and offensive label propagation for
32 core extraction. *Phys. Rev. E.*, 83:036103, 2011. 33
- [92] X. Wang and J. Li. Detecting communities by the core-vertex and intimate degree in complex networks. *Physica A.*, 392:2555–2563,
34 2013. 35
- [93] J. Li, X. Wang, and J. Eustace. Detecting overlapping communities by seed community in weighted complex networks. *Physica A.*, 392:
36 6125–6134, 2013. 37
- [94] D. R. Fabio, D. Fabio, and P. Carlo. Profiling core-periphery network structure by random walkers. *Sci. Rep.*, 3:1467, 2013. 38
- [95] Q. Chen, T. T. Wu, and M. Fang. Detecting local community structure in complex networks based on local degree central nodes. *Physica*
39 *A.*, 392:529–537, 2013. 40
- [96] S. Zhang, R. Wang, and X. Zhang. Identification of overlapping community structure in complex networks using fuzzy c-means clustering.
41 *Physica A.*, 374:483–490, 2007. 42
- [97] T. Nepusz, A. Petróczy, L. Négyessy, and F. Bazsó. Fuzzy communities and the concept of bridgeness in complex networks. *Phys. Rev.*
43 *E.*, 77:016107, 2008. 44
- [98] B. Fabricio and Z. Liang. Fuzzy community structure detection by particle competition and cooperation. *Soft Comput.*, 17:659–673, 2013. 45
- [99] P. Sun, L. Gao, and S. Han. Identification of overlapping and non-overlapping community structure by fuzzy clustering in complex
46 networks. *Inf. Sci.*, 181:1060–1071, 2011. 47
- [100] W. Wang, D. Liu, X. Liu, and L. Pan. Fuzzy overlapping community detection based on local random walk and multidimensional scaling.
48 *Physica A.*, 392:6578–6586, 2013. 49
- [101] I. Psorakis, S. Roberts, M. Ebden, and B. Sheldon. Overlapping community detection using bayesian non-negative matrix factorization.
50 *Phys. Rev. E.*, 83:066114, 2011. 51
- [102] Y. Zhang and D. Yeung. Overlapping community detection via bounded nonnegative matrix tri-factorization. In *In Proc. ACM SIGKDD*
51 *Conf.*, pages 606–614, 2012.
- [103] J. Liu. Fuzzy modularity and fuzzy community structure in networks. *Eur. Phys. J. B.*, 77:547–557, 2010.
- [104] T. C. Havens, J. C. Bezdek, Ramamohanarao K. Leckie, C., and M. Palaniswami. A soft modularity function for detecting fuzzy commu-
nities in social networks. *IEEE Trans. Fuzzy Syst.*, 21:1170–1175, 2013.
- [105] M. E. J. Newman. Network data. <http://www-personal.umich.edu/~mejn/netdata/>, 2013.
- [106] L. Šubelj and M. Bajec. Ubiquitousness of link-density and link-pattern communities in real-world networks. *Eur. Phys. J. B.*, 85:1–11,
2012.

- 1 [107] A. Lancichinetti, S. Fortunato, and F. Radicchi. Benchmark graphs for testing community detection algorithms. *Phys. Rev. E.*, 78:046110, 1
2 2008. 2
- 3 [108] W. Liu, M. Pellegrini, and X. Wang. Detecting communities based on network topology. *Sci. Rep.*, 4:5739, 2014. 3
- 4 [109] L. Danon, A. Diaz-Guilera, J. Duch, and A. Arenas. Comparing community structure identification. *J. Stat. Mech.-Theory Exp.*, page 4
5 P09008, 2005. 5
- 6 [110] S. Gregory. Fuzzy overlapping communities in networks. *J. Stat. Mech.-Theory Exp.*, page P02017, 2011. 6
- 7 [111] A. Lancichinetti and S. Fortunato. Benchmarks for testing community detection algorithms on directed and weighted graphs with over- 7
lapping communities. *Phys. Rev. E.*, 80:016118, 2009. 7
- 8 [112] E. Hullermeier and M. Rifqi. A fuzzy variant of the rand index for comparing clustering structures. In *in Proc. IFSA/EUSFLAT Conf.*, 8
9 pages 1294–1298, 2009. 9
- 10 [113] . 10
- 11 [114] Rare diseases. URL https://ec.europa.eu/info/research-and-innovation/research-area/health-research-and-innovation/rare-diseases_en. 11
- 12 [115] U. Schara, J. Mortier, and W. Mortier. Long-Term Steroid Therapy in Duchenne Muscular Dystrophy-Positive Results versus Side Effects. 12
Journal of clinical neuromuscular disease, 2(4):179–183, 2001. ISSN 1537-1611. . URL <https://pubmed.ncbi.nlm.nih.gov/19078632/>. 12
- 13 [116] Salvatore Crisafulli, Janet Sultana, Andrea Fontana, Francesco Salvo, Sonia Messina, and Gianluca Trifirò. Global epidemiology of 13
14 duchenne muscular dystrophy: an updated systematic review and meta-analysis. 15(1):141. ISSN 1750-1172. . URL <https://doi.org/10.1186/s13023-020-01430-8>. 14
- 15 [117] Karen T. Nozoe, Ricardo T. Akamine, Diego R. Mazzotti, Daniel N. Polesel, Luís F. Grossklauss, Sergio Tufik, Monica L. Andersen, and 15
16 Gustavo A. Moreira. Phenotypic contrasts of duchenne muscular dystrophy in women: Two case reports. 9(3):129–133. ISSN 1984-0659. 16
17 . URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5241604/>. 17
- 18 [118] Tomoko Lee, Yasuhiro Takeshima, Noriko Kusunoki, Hiroyuki Awano, Mariko Yagi, Masafumi Matsuo, and Kazumoto Iijima. Differ- 18
ences in carrier frequency between mothers of duchenne and becker muscular dystrophy patients. 59(1):46–50. ISSN 1434-5161. . URL 19
<https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3970902/>. 20
- 21 [119] Sylvie Tuffery-Giraud, Christophe Bérout, France Leturcq, Rabah Ben Yaou, Dalil Hamroun, Laurence Michel-Calemard, Marie-Pierre 21
22 Moizard, Rafaëlle Bernard, Mireille Cossée, Pierre Boisseau, Martine Blayau, Isabelle Creveaux, Anne Guiochon-Mantel, Bérengère 22
23 de Martinville, Christophe Philippe, Nicole Monnier, Eric Bieth, Philippe Khau Van Kien, François-Olivier Desmet, Véronique Hum- 23
bertclaude, Jean-Claude Kaplan, Jamel Chelly, and Mireille Claustres. Genotype–phenotype analysis in 2,405 patients with a dys- 24
trophinopathy using the UMD–DMD database: a model of nationwide knowledgebase. 30(6):934–945. ISSN 1098-1004. . URL 24
<https://onlinelibrary.wiley.com/doi/abs/10.1002/humu.20976>. 25
_eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/humu.20976>. 25
- 26 [120] M. Liu, N. Chino, and T. Ishihara. Muscle damage progression in duchenne muscular dystrophy evaluated by a new quantitative computed 26
tomography method. 74(5):507–514. ISSN 0003-9993. . 27
- 28 [121] M. Sekiguchi. The role of dystrophin in the central nervous system: a mini review. 24(2):93–97. ISSN 1128-2460. 28
- 29 [122] Dongsheng Duan, Nathalie Goemans, Shin'ichi Takeda, Eugenio Mercuri, and Annemieke Aartsma-Rus. Duchenne muscular dystrophy. 29
7(1):1–19. ISSN 2056-676X. . URL <https://www.nature.com/articles/s41572-021-00248-3>. Number: 1 Publisher: Nature Publishing 30
Group. 30
- 31 [123] David J Birnkrant, Katharine Bushby, Carla M Bann, Benjamin A Alman, Susan D Apkon, Angela Blackwell, Laura E Case, Linda 31
32 Cripe, Stasia Hadjiyannakis, Aaron K Olson, Daniel W Sheehan, Julie Bolen, David R Weber, and Leanne M Ward. Diagnosis and 32
33 management of duchenne muscular dystrophy, part 2: respiratory, cardiac, bone health, and orthopaedic management. 17(4):347–361, . 33
34 ISSN 1474-4422. . URL <https://www.sciencedirect.com/science/article/pii/S1474442218300255>. 34
- 35 [124] David J Birnkrant, Katharine Bushby, Carla M Bann, Susan D Apkon, Angela Blackwell, David Brumbaugh, Laura E Case, Paula R 35
36 Clemens, Stasia Hadjiyannakis, Shree Pandya, Natalie Street, Jean Tomezsko, Kathryn R Wagner, Leanne M Ward, and David R Weber. 36
37 Diagnosis and management of duchenne muscular dystrophy, part 1: diagnosis, and neuromuscular, rehabilitation, endocrine, and gas- 37
38 trointestinal and nutritional management. 17(3):251–267, . ISSN 1474-4422. . URL <https://www.sciencedirect.com/science/article/pii/S1474442218300243>. 38
- 39 [125] Erik Landfeldt, Rachel Thompson, Thomas Sejersen, Hugh J. McMillan, Janbernd Kirschner, and Hanns Lochmüller. Life expectancy at 39
40 birth in duchenne muscular dystrophy: a systematic review and meta-analysis. 35(7):643–653. ISSN 1573-7284. . 40
- 41 [126] Jean-Pierre Jourdan, Ronan Bureau, Christophe Rochais, and Patrick Dallemagne. Drug repositioning: a brief overview. page 41
42 10.1111/jphp.13273. ISSN 0022-3573. . URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7262062/>. 42
- 43 [127] Chao Pang, Annet Solлие, Anna Sijtsma, Dennis Hendriksen, Bart Charbon, Mark De Haan, Tommy De Boer, Fleur Kelpin, Jonathan 43
44 Jetten, Joeri K. Van Der Velde, Nynke Smidt, Rolf Sijmons, Hans Hillege, and Morris A. Swertz. SORTA: a system for ontology-based 44
45 re-coding and technical annotation of biomedical phenotype data. *Database : the journal of biological databases and curation*, 2015, 45
2015. ISSN 1758-0463. . URL <https://pubmed.ncbi.nlm.nih.gov/26385205/>. 45
- 46 [128] Helen Prior, Paul Baldrick, Lolke de Haan, Noel Downes, Keith Jones, Elisabeth Mortimer-Cassen, and Ian Kimber. Reviewing the Utility 46
47 of Two Species in General Toxicology Related to Drug Development. *International Journal of Toxicology*, 37(2):121, mar 2018. ISSN 47
1092874X. . URL <https://pubmed.ncbi.nlm.nih.gov/30811785/>. 48
48 [129] Francesco Sovrano, Fabio Vitali, and Monica Palmirani. Making Things Explainable vs Explaining: Requirements and Challenges under 49
50 the GDPR. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioin-* 50
51 *formatics)*, 13048 LNAI:169–182, oct 2021. . URL <http://arxiv.org/abs/2110.00758>https://doi.org/10.1007/978-3-030-89811-3_12. 51

- [130] Daniel N. Sosa, Alexander Derry, Margaret Guo, Eric Wei, Connor Brinton, and Russ B. Altman. A Literature-Based Knowledge Graph Embedding Method for Identifying Drug Repurposing Opportunities in Rare Diseases. *Pacific Symposium on Biocomputing. Pacific Symposium on Biocomputing*, 25(2020):463, 2020. ISSN 23356936. . URL [/pmc/articles/PMC6937428/](https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6937428/)[/pmc/articles/PMC6937428/?report=abstract](https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6937428/?report=abstract)<https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6937428/>.
- [131] Xiao Wang, Deyu Bo, Chuan Shi, Shaohua Fan, Yanfang Ye, and Philip S. Yu. A Survey on Heterogeneous Graph Embedding: Methods, Techniques, Applications and Sources. *IEEE Transactions on Big Data*, pages 1–1, nov 2020. ISSN 23327790. . URL <https://arxiv.org/abs/2011.14867v1>.
- [132] Xiaoqin Pan, Xuan Lin, Dongsheng Cao, Xiangxiang Zeng, Philip S. Yu, Lifang He, Ruth Nussinov, and Feixiong Cheng. Deep learning for drug repurposing: methods, databases, and applications. feb 2022. . URL <https://arxiv.org/abs/2202.05145v1>.
- [133] Kavitha Agastheeswaramoorthy and Aarti Sevilimedu. Drug REpurposing using AI/ML tools - for Rare Diseases (DREAM-RD): A case study with Fragile X Syndrome (FXS). *bioRxiv*, page 2020.09.25.311142, sep 2020. . URL <https://www.biorxiv.org/content/10.1101/2020.09.25.311142v1><https://www.biorxiv.org/content/10.1101/2020.09.25.311142v1.abstract>.
- [134] Sandra Brasil, Carlota Pascoal, Rita Francisco, Vanessa Dos Reis Ferreira, Paula A. Videira, and Gonalo Valadao. Artificial Intelligence (AI) in Rare Diseases: Is the Future Brighter? *Genes*, 10(12), dec 2019. ISSN 2073-4425. . URL <https://pubmed.ncbi.nlm.nih.gov/31783696/>.
- [135] Andrea Rossi, Donatella Firmani, Antonio Marinata, Paolo Merialdo, and Denilson Barbosa. Knowledge graph embedding for link prediction: A comparative analysis. 15(2):1–49, . ISSN 1556-4681, 1556-472X. . URL <http://arxiv.org/abs/2002.00819>.
- [136] Andrea Rossi, Donatella Firmani, Antonio Marinata, Paolo Merialdo, and Denilson Barbosa. Knowledge graph embedding for link prediction: A comparative analysis. 15(2):1–49, . ISSN 1556-4681, 1556-472X. . URL <http://arxiv.org/abs/2002.00819>.
- [137] DrugBank Online | Database for Drug and Drug Target Info. URL <https://go.drugbank.com/>.
- [138] Ellen E. Ladenheim. Liraglutide and obesity: a review of the data so far. *Drug Design, Development and Therapy*, 9:1867, mar 2015. ISSN 11778881. . URL [/pmc/articles/PMC4386791/](https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4386791/)[/pmc/articles/PMC4386791/?report=abstract](https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4386791/?report=abstract)<https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4386791/>.
- [139] Rex Ying, Dylan Bourgeois, Jiaxuan You, Marinka Zitnik, and Jure Leskovec. GNNExplainer: Generating explanations for graph neural networks. *Adv. Neural Inf. Process. Syst.*, 32:9240–9251, December 2019.
- [140] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space. (arXiv:1301.3781). . URL <http://arxiv.org/abs/1301.3781>. type: article.
- [141] Gnnexplainer - pytorch geometric. URL https://pytorch-geometric.readthedocs.io/en/latest/_modules/torch_geometric/nn/models/gnn_explainer.html.
- [142] Yankai Lin, Zhiyuan Liu, and Maosong Sun. Knowledge Representation Learning with Entities, Attributes and Relations. URL <https://github.com/thunlp/KR-EAR>.
- [143] Visualization function - pytorch geometric. URL https://pytorch-geometric.readthedocs.io/en/latest/_modules/torch_geometric/nn/models/explainer.html#Explainer.visualize_subgraph.
- [144] PyG documentation — pytorch_geometric documentation, . URL <https://pytorch-geometric.readthedocs.io/en/latest/>.
- [145] PyTorch documentation — PyTorch 1.11.0 documentation, . URL <https://pytorch.org/docs/stable/index.html>.
- [146] GNNExplainer tutorial, . URL https://github.com/OpenXAIProject/GNNExplainer-Tutorial/blob/e07571385670462b27ca3274522f46939ea5daa2/gnnexplainer_cor.ipynb. original-date: 2020-08-24T08:11:05Z.
- [147] Raya Khanin and Ernst Wit. How Scale-Free Are Biological Networks. <https://home.liebertpub.com/cmb>, 13(3):810–818, may 2006. ISSN 10665277. . URL <https://www.liebertpub.com/doi/10.1089/cmb.2006.13.810>.
- [148] Farzin Sohraby, Milad Bagheri, and Hassan Aryapour. Performing an In Silico Repurposing of Existing Drugs by Combining Virtual Screening and Molecular Dynamics Simulation. *Methods in molecular biology (Clifton, N.J.)*, 1903:23–43, 2019. ISSN 1940-6029. . URL <https://pubmed.ncbi.nlm.nih.gov/30547434/>.
- [149] Remington Kim. OrphaDRGL: A Novel Graph Deep Learning-based Drug Repositioning Approach for Orphan Diseases.
- [150] Amir Lakizadeh and Sayed Mohammad Hassan Mir-Ashrafi. Drug repurposing improvement using a novel data integration framework based on the drug side effect. *Informatics in Medicine Unlocked*, 23:100523, jan 2021. ISSN 2352-9148. .
- [151] Reka Albert. Scale-free networks in cell biology. *Journal of Cell Science*, 118(21):4947–4957, nov 2005. ISSN 0021-9533. . URL <https://journals.biologists.com/jcs/article/118/21/4947/28519/Scale-free-networks-in-cell-biology>.