

# Explaining Graph Neural Network Predictions for Drug Repurposing

Julie Loesch<sup>1</sup>, Yang Yang<sup>1</sup>, Pelinsu Ekmekci<sup>2</sup>, Michel Dumontier<sup>1</sup> and Remzi Celebi<sup>1</sup>

<sup>1</sup>*Department of Advanced Computing Sciences, Maastricht University, Paul-Henri Spaaklaan 1, Maastricht, 6229 EN, Netherlands*

<sup>2</sup>*Department of Bioengineering, İzmir Institute of Technology, Urla, 35430, İzmir, Turkey*

## Abstract

Graph Neural Networks (GNNs) are powerful tools for graph-related tasks, excelling in processing graph-structured data while maintaining permutation invariance. However, their challenge lies in the obscurity of new node representations, hindering interpretability. This paper introduces a framework addressing this limitation by explaining GNN predictions. The proposed method takes any GNN prediction, for which it returns a concise subgraph as explanation. Utilizing Saliency Maps, an attribution gradient-based technique, we enhance interpretability by assigning importance scores to entities within the knowledge graph via backpropagation. Evaluated on the Drug Repurposing Knowledge Graph, Graph Attention Network achieved a Hits@5 score of 0.451 and a Hits@10 score of 0.672. GraphSAGE demonstrated notable results with the highest recall rate of 0.992. Our framework underscores GNN efficacy and interpretability, which is crucial in complex scenarios like drug repurposing. Illustrated through an Alzheimer's disease case study, our approach provides meaningful and comprehensible explanations for GNN predictions. This work contributes to advancing the transparency and utility of GNNs in real-world applications.

## Keywords

Knowledge Graphs (KGs), Drug Repurposing, Graph Neural Networks (GNNs), Explainable AI (XAI), Saliency Maps, Alzheimer's Disease

## 1. Introduction

In recent years, significant attention has been directed towards the development and utilization of extensive Knowledge Graphs (KGs) [1] for computational drug discovery tasks [2, 3, 4]. KGs serve as a unified data model to integrate knowledge from diverse sources, enabling the discovery of complex and hidden relationships between biological entities.


Drug repurposing using KGs is often modeled as link prediction problem, where the objective is to predict the connection between a candidate drug and a target disease. Several methods, including embedding-, rule- and network-based approaches, have been employed to address this problem. While Graph Neural Networks (GNNs) have gained prominence due to their capabilities of complex reasoning over expansive KGs, their application introduces interpretability challenges. Specifically, the inherent opacity of GNNs, stemming from their aggregation


---

SWAT4HCLS'24, Feb 26 – 29, Leiden, the Netherlands

✉ julie.loesch@maastrichtuniversity.nl (J. Loesch); rena.yang@maastrichtuniversity.nl (Y. Yang); michel.dumontier@maastrichtuniversity.nl (M. Dumontier); remzi.celebi@maastrichtuniversity.nl (R. Celebi)

ORCID 0000-0001-7290-5851 (Y. Yang); 0000-0003-4727-9435 (M. Dumontier); 0000-0001-7769-4272 (R. Celebi)

 © 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

 CEUR Workshop Proceedings (CEUR-WS.org)

strategy, hinders the traceability of computing new node representations, posing difficulties in interpreting the final predictions.

To address this interpretability challenge, Chari et al. [5] introduced an explanation ontology, categorizing various literature-derived explanation types. This paper proposes a framework generating explanations falling within the "statistical" and "scientific" explanation categories. Our approach draws inspiration from statistical explanation principles to enhance the interpretability of GNN predictions. Moreover, the explanations given for the case study on Alzheimer's disease are supported by existing literature. Through empirical evidence obtained from our experiments, we aim to provide clear insights into the predictions generated by GNNs.

Our framework takes a prediction from a trained GNN as input, for which it produces an explanation in the form of a concise subgraph. By employing Saliency Maps, an attribution gradient-based method, our method assigns importance scores to nodes via backpropagation. The method generates an explanatory subgraph by summing node scores to determine the rank or importance of specific triples.

The remainder of the paper is structured as follows: Section 2 delves into previous related works. Section 3 describes the Drug Repurposing Knowledge Graph (DRKG) [6] while Section 4 introduces our proposed XAI framework. Section 5 details the experimental setup and reports results, followed by a discussion in Section 6. Finally, Section 7 concludes this paper. The implementation for analysis is available in our Github repository <https://github.com/jloe2911/ExplainableDrugRepurposing>.

## 2. Related Work

While GNNs have gained traction for their capacity to perform complex reasoning over large KGs, their lack of interpretability poses a challenge for user trust. Jiannan et al. [7] addressed this limitation with their interpretable Deep learning-based Path-reasoning framework for drug repurposing (iDPath), which captures the mechanisms of drug actions (MODA) by identifying the paths from drugs to diseases in the human body. iDPath utilizes a Graph Convolutional Network (GCN) [8] module for global connectivity information of the human molecular network, and a Long-Short Term Memory (LSTM) [9] neural network module to capture detailed drug action mechanisms. Furthermore, iDPath introduces two attention modules, path attention and node attention, to enhance model interpretability. To accurately characterize the MODA, the authors built a multi-layer biological network, integrating a gene regulatory layer, a protein-protein interaction layer, a protein-chemical interaction layer, and a chemical-chemical interaction layer. The authors demonstrated its real-world utility identifying ten potential drugs for treating prostate cancer, of which six have already been proven effective in previous studies.

Similarly, Islam et al. [10] introduced a workflow that leverages a set of explanation rules and extracts explanatory paths instantiating the rules for given (*Compound, Disease*) pairs of interest based on the DRKG [6]. Their approach generates compact ensemble embeddings of the entities and relations using TransE [11], TransH [12], and DistMult [13], followed by training a graph-based deep neural network using the generated ensemble embeddings to predict the probability that a *treat* relation exists for a given (*Compound, Disease*) pair. Achieving an average mean squared error of 0.09 and an average AUC score of 0.96 through 10-fold cross validation,

their method extracts explanation rules from DRKG to generate plausible explanation(s) for the predictions. The rule mining method process involves five major steps: (i) sampling a subset of triples for the *treat* relation, (ii) extracting paths of maximum length three from *Compound* to *Disease* entities for one triple, (iii) computing path scores based on their entity and relation embedding and rank them, (iv) transforming top-ranked paths into rules by replacing entities with variables and update the global rule set, and (v) iterating steps (ii) to (iv) for each sampled triple and outputting the global rule set [14].

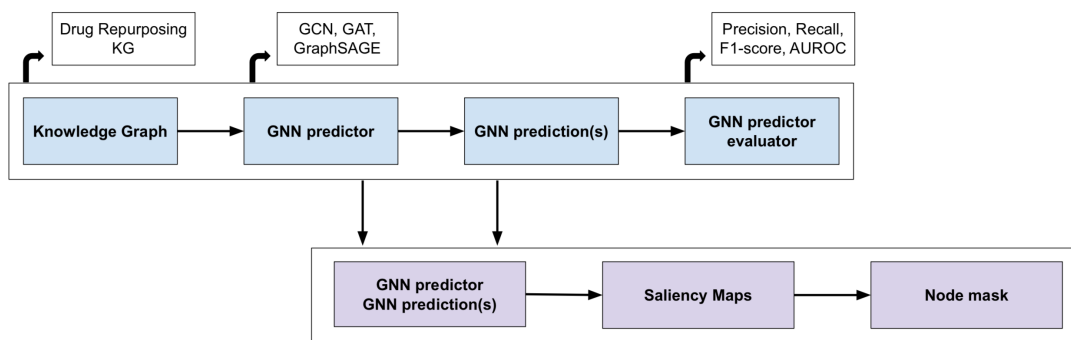
Gottlieb et al. [15] presented the PREDICT method, grounded in the "Guilty by Association" (GBA) approach. PREDICT predicts novel associations between desired drugs and diseases, involving the measurements of similarities among the known drug and disease to drug-disease pairs, given a target query drug and disease. However, lacking reasoning behind the final predictions promoted Ozkan et al. [16] to propose XPREDICT, a KG-based explainable AI approach for drug repurposing. XPREDICT focuses on generating similarity-based explanations, calculating cosine similarity between drug-drug and disease-disease pairs. While effective in generating plausible explanations on a case study on Alzheimer's disease through the OpenPredict API [17], a limitation is the exclusion of relevant entities like genes and pathways in the generated explanation. To address this, we expanded our XAI framework to include gene and pathway nodes.

### 3. Dataset

The input to our framework is the DRKG [6], initially developed for the COVID-19 drug repurposing project. DRKG is a comprehensive KG that interconnects various biological entities, including compounds, diseases, genes, and side effects. It collects approximately 97 thousand entities and over 5 million edges. Curated from six different data sources, namely STRING, DrugBank, Hetionet, Global Network of Biomedical Relationships (GNBR), IntAct and Drug Gene Interaction database (DGIdb), DRKG comprises 107 different edge-types. Given the goal of drug repurposing to discover new uses for existing drugs, we approach the challenge as a link prediction task, aiming to predict missing edges between Compounds and Diseases: *Compound*  $\xrightarrow{\text{treats}}$  *Diseases*. Figure 1 illustrates the entire pipeline, encompassing two major steps.

### 4. Methodology

This section outlines the proposed framework for explaining predictions generated by GNNs. The initial phase of the pipeline focuses on training a GNN-based link prediction model. The fundamental concept behind a GNN layer is to compress a set of vectors into a single one. This involves a two-step process: (i) message computation and (ii) aggregation. During the first step, each node calculates a message, represented as a node embedding. Mathematically, for each node  $u \in \{N(v) \cup v\}$ , the message at layer  $l$  is computed as  $m_u^{(l)} = \text{MSG}^{(l)}(h_u^{(l-1)})$ . Moving to the second step, each node aggregates the messages from its neighbours. Formally,  $h_v^{(l)}$  is determined by  $\text{AGG}^{(l)}(\{m_u^{(l)}, u \in N(v)\}, m_v^{(l)})$ . Consequently, the  $l$ -th GNN layer takes node embeddings from the node itself,  $h_v^{(l-1)}$ , and those from neighbouring nodes,  $h_{u \in N(v)}^{(l-1)}$ , producing



**Figure 1:** Pipeline of generating GNN prediction(s) and its node mask.

the node embedding  $h_v^{(l)}$ . In this study, various GNN variants were employed, including Graph Convolutional Network (GCN) [8], GraphSAGE [18], and Graph Attention Network (GAT) [19]. The key distinction lies in how these variants aggregate information across the different layers.

In the second phase of the pipeline, the process involves utilizing a trained GNN-based link prediction model and a prediction  $\hat{y}_i$  to generate an explanation  $G_S \subseteq G_C$  in the form of a compact subgraph derived from the input graph. To achieve this, we employed Saliency Maps [20], an attribution gradient-based method proficient in assigning importance scores to entities within the KG via the process of backpropagation.

By highlighting or masking specific nodes to reveal their contributions to the final prediction, we can easily extract a concise subgraph comprising the most important nodes that significantly influenced the prediction(s). Specifically, we applied Saliency Maps for each prediction  $\hat{y}_i, \forall i \in N$  to obtain importance scores for all nodes influential in determining the outcome of that particular prediction  $\hat{y}_i$ . Following the computation of gradients of the output concerning the input (i.e., node embedding), the attribution value using Saliency Maps along the  $i^{th}$  dimension for an input  $x \in \mathbb{R}^n$  is defined as the absolute value of the gradient:  $\text{Saliency}_i(x) = \left| \frac{\partial F(x)}{\partial x_i} \right|$  where  $F(x)$  denotes the output of a GNN model for input  $x$ .

Utilizing the scores assigned by Saliency Maps to each node in the network, we can construct an explanatory subgraph employing Algorithm 1. The algorithm generates explanatory subgraphs based on node attribution scores computed by Saliency Maps (SM). It starts by ranking Gene entities based on SM scores. For each Gene, various relations, including participation in pathways, association with diseases and binding with compounds are extracted. These relations are ranked based on the cumulative scores assigned to their respective nodes. In the end, top  $k$  triples are included in the final explanatory subgraph.

Hence, the proposed methodology explains GNN predictions by encompassing two primary phases. The initial phase involves training a GNN-based link prediction model. Variants like GCN, GraphSAGE, and GAT are employed to compute node embeddings. The second phase leverages the trained GNN-based link prediction model to generate explanations for GNN predictions employing Saliency Maps. Saliency Maps highlight the most influential nodes contributing to the final predictions. The algorithm defined in the framework uses these scores

---

**Algorithm 1** Algorithm used to generate explanations.

---

```
1: function GENERATEEXPLANATORYSUBGRAPH( $SM\_scores, k$ )    ▷ where  $SM\_scores$  - scores
   derived from Saliency Maps,  $k$  - number of triples included in the explanation
2:   Let  $g_1, g_2, \dots, g_n$  be ranked gene entities based on  $SM\_scores$ 
3:   RankedTriples = []
4:   for  $g_i = 1$  to  $n$  do
5:     PathwayRel = ExtractRelations( $g_i, "ParticipatesIn", "Pathway"$ )
6:     DiseaseRel = ExtractRelations("Disease", "Associates",  $g_i$ )
7:     CompoundRel = ExtractRelations("Compound", "BindsTo",  $g_i$ )
8:     RankedPathwayRel = RankRelations(PathwayRel)
9:     RankedDiseaseRel = RankRelations(DiseaseRel)
10:    RankedCompoundRel = RankRelations(CompoundRel)
11:    RankedTriples.append(RankedPathwayRel[: $k$ ], RankedDiseaseRel[: $k$ ], RankedCom-
   poundRel[: $k$ ])
12:   end for
13:   ExplanatorySubgraph = BuildExplanatorySubgraph(RankedTriples)
   return ExplanatorySubgraph
14: end function
```

---

to rank paths involving genes, pathways, diseases, and compounds, which are then used to construct an explanatory subgraph.

## 5. Evaluation and Results

### 5.1. Graph Neural Networks

To evaluate the proposed framework, we employed GCN [8], GraphSAGE [18], and GAT [19] on the DRKG [6]. The assessment of GNNs was based on Precision, Recall, and Hits@ $k$  metrics.

In order to train the GNNs, we implemented a train-test split of 80% for training samples and 20% for testing samples. The node embeddings were set with a dimension size of 100 to represent the features and relationships within the graph structures. For optimization, the Adam optimizer was selected to efficiently update model weights and enhance convergence during the training process. To iteratively refine the model's parameters, we conducted 500 epochs, enabling the network to learn intricate patterns and relationships within the data.

Precision measures the ratio of true positives to predicted positives, offering insights into the cost of false positives. Conversely, recall calculates the ratio of true positives to actual positives, indicating the cost of false negatives. Hits@ $k$  measures how frequently the correct head or tail is among the top  $k$  predictions. Table 1 highlights that GAT outperformed, achieving a Hits@5 score of 0.451 and a Hits@10 score of 0.672. The trade-off between precision and recall is evident, with a high emphasis on recall in drug repurposing scenarios to minimize falsely predicted *treat*-links between *drugs* and *diseases*. Notably, GraphSAGE achieved the highest recall rate of 0.992.

**Table 1**

Results for predicting link (Compound, treats, Disease).

|             | GNN variant | Precision    | Recall       | Hits@5       | Hits@10      |
|-------------|-------------|--------------|--------------|--------------|--------------|
| <b>DRKG</b> | GraphSAGE   | 0.287        | <b>0.992</b> | 0.298        | 0.385        |
|             | GCN         | 0.361        | 0.871        | 0.280        | 0.409        |
|             | GAT         | <b>0.834</b> | 0.610        | <b>0.451</b> | <b>0.672</b> |

## 5.2. Use Case: Alzheimer’s disease

We applied our explainability framework to interpret predictions generated by GNNs in a case study focused on Alzheimer’s disease. The analysis involved a trained GNN, namely GraphSAGE having the highest recall rate of 0.992, and a disease-filtered KG, where only Alzheimer’s disease-related information was retained. In particular, we excluded all triples (Compound, treats, Disease) where the disease differs from Alzheimer’s. Notably, our focus was solely on explaining correct predictions, disregarding classification errors made by the trained GNN.

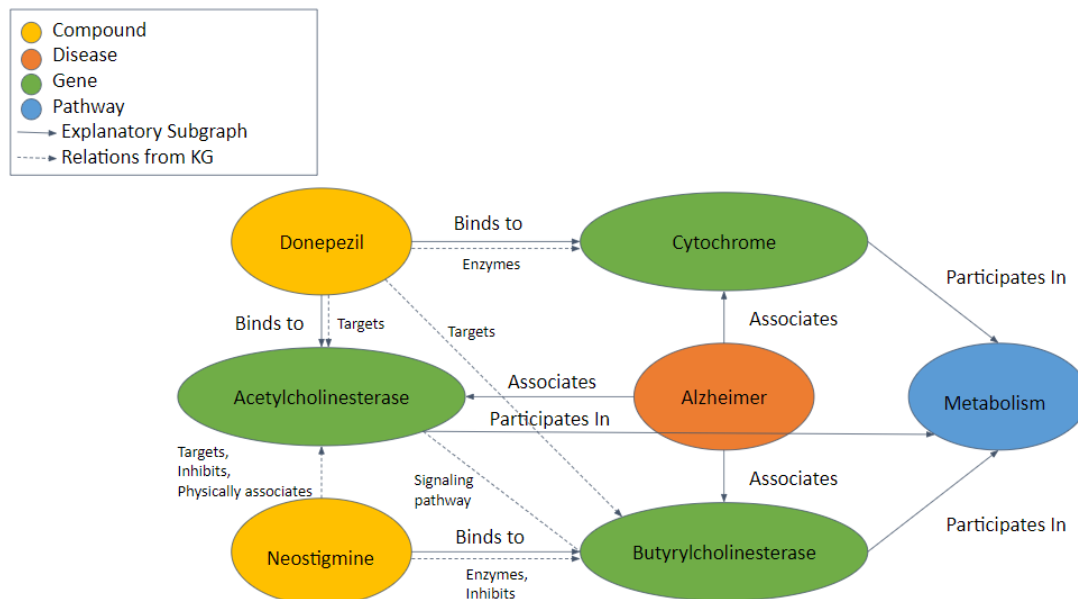
To assess the efficacy of our approach, we generated explanatory subgraphs for specific instances, such as (Donepezil, treats, Alzheimer) and (Memantine, treats, Alzheimer). Figure 2 illustrates the explanatory subgraph for Donepezil. The literature reveals that the primary goal of Alzheimer’s drugs, including Donepezil, is to maintain elevated acetylcholine (ACh) levels, thereby compensating for the loss of functioning cholinergic brain cells [21, 22]. Figure 2 emphasized the crucial role of Donepezil binding to acetylcholinesterase (AChE) and butyrylcholinesterase (BChE). Both BChE and AChE are involved in ACh metabolism and thus are important for the cholinergic function in the brain [23].

Similarly, we explored the explanatory subgraph for (Memantine, treats, Alzheimer) as depicted in Figure 3. Memantine, commonly prescribed for moderate to severe Alzheimer’s disease, is believed to help prevent excess levels of the substance glutamate from damaging the brain<sup>1</sup>. Thus, according to our explainability framework and existing literature, important genes associated with Memantine’s efficacy include glutamate receptor and acetylcholinesterase. This comprehensive interpretation provides valuable insights into the complex mechanisms underlying the treatment of Alzheimer’s disease with these medications.

## 6. Discussion

For experts equipped with biological knowledge, evaluating AI-based model explanations, particularly those generated by GNNs, involves several key considerations. First, assessing the accuracy and consistency of predictions and explanations is crucial, comparing their performance to existing methods or baselines. Second, scrutinizing the biological validity and plausibility of the returned subgraphs becomes paramount, examining how well they align with known or novel associations between nodes and target predictions. The comprehensibility and simplicity of the explanations play a pivotal role, offering insights into how effectively they aid users in grasping the reasoning and logic behind predictions, along with the underlying

<sup>1</sup><https://www.ncbi.nlm.nih.gov/books/NBK279356/>



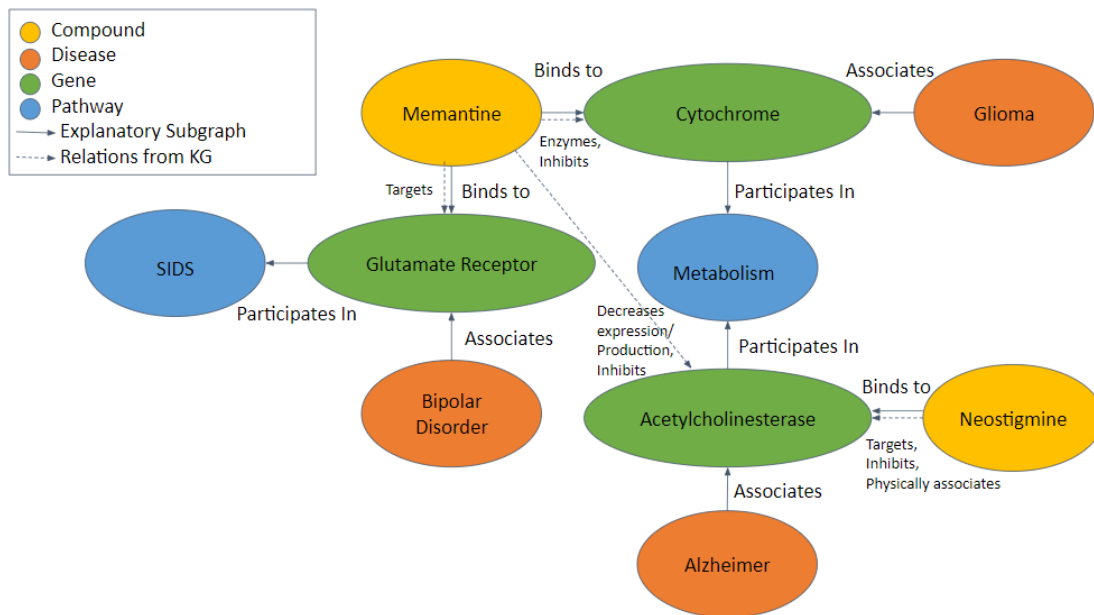
**Figure 2:** Explanatory subgraph for (Donepezil, treats, Alzheimer) including relationships from the original KG.

mechanisms and disease processes. Finally, evaluating the actionability and applicability of the explanations is essential, gauging their capacity to support users in making well-informed decisions, whether related to diagnosis, prognosis, treatment, or prevention of the disease.

While explanations from KGs currently have limitations, relying on simple associations and bindings between drugs, genes, and diseases, there is room for improvement. Existing frameworks often overlook crucial aspects such as causal mechanisms, dose-response relationships, and clinical outcomes. However, insights from biological explanations guide us towards refining our explanation.

Understanding drug interactions at the molecular level, like the binding of Donepezil to enzymes such as CYP2D6, opens avenues for more nuanced explanations. While this impacts pharmacokinetics and pharmacodynamics, it does not directly elucidate the therapeutic action on Alzheimer's disease. Conversely, the binding of Donepezil to ACHE, an enzyme in acetylcholine metabolism, is a crucial mechanism explaining its positive impact on cognitive function in Alzheimer's patients. Similarly, Memantine's binding to CYP2B6, an enzyme in drug metabolism, raises considerations about pharmacokinetics and pharmacodynamics. Memantine's binding to GRIN1, a subunit of the NMDA receptor, presents a critical mechanism. By blocking the NMDA receptor, Memantine protects neurons from excitotoxicity caused by glutamate overstimulation, explaining how it reduces cognitive decline and improves overall function in Alzheimer's patients.

In essence, refining associations with a focus on relevant targets presents an exciting avenue for enhancing explanations related to the therapeutic effects of drugs like Donepezil and Memantine in Alzheimer's treatment. This not only aligns with expert biological insights but also contributes to the wider acceptance and utility of explanations in the domain.



**Figure 3:** Explanatory subgraph for (Memantine, treats, Alzheimer) including relationships from the original KG.

While the proposed framework demonstrates notable strengths in interpretable graph-based predictions, it is essential to recognize certain limitations. Firstly, the framework assumes a static and deterministic nature of input graphs, potentially hindering its applicability to dynamic, temporal, or uncertain graph characteristics. Secondly, the focus on local or subgraph-level explanations might limit the framework's ability to capture global aspects of the graph, including its topology, properties, or embeddings. Furthermore, the framework's single-subgraph approach may restrict its capacity to provide diverse or alternative explanations for the same prediction, highlighting considerations for its application in scenarios requiring comprehensive and varied explainability. Another limitation of this work is that there was no formal evaluation conducted for the explanations.

To propel the capabilities of our framework and address its identified limitations, future work could encompass diverse directions. One avenue involves extending the framework to accommodate various graph types, including directed, weighted, or attributed graphs, enabling a comprehensive evaluation of its performance and interpretability across different graph datasets. Additionally, investigating the integration of domain knowledge, such as biological pathways, gene ontology, or disease phenotypes, into the framework holds potential for enhancing biological interpretability and relevance in the explanations. Furthermore, conducting user studies with domain experts, such as biologists, clinicians, or patients, can provide valuable insights into the usefulness, trustworthiness, and user satisfaction of the framework. Such studies can shed light on potential challenges and limitations, guiding refinements and improvements for future iterations of the framework.



## 7. Conclusion

In this paper, we presented an XAI framework that combines GNNs for link prediction on the DRKG [6] with Saliency Maps to generate explanatory subgraphs, emphasizing the importance of genes in pathways, disease associations, and compound bindings. Among the GNN variants that have been experimented with, GAT demonstrated good performances in drug repurposing, achieving a Hits@5 score of 0.451 and a Hits@10 score of 0.672. In addition, great results were obtained using GraphSAGE, attaining the highest high recall rate of 0.992 and providing effective explanations for a case study on Alzheimer's disease. Future work could involve extending the framework to handle various graph types, exploring alternative methods for generating Saliency Maps, and integrating domain knowledge for enhanced biological interpretability. Additionally, conducting user studies with domain experts can provide valuable insights for refining and improving the framework in future iterations.

## References

- [1] S. Ji, S. Pan, E. Cambria, P. Marttinen, P. S. Yu, A survey on knowledge graphs: Representation, acquisition, and applications, *IEEE Transactions on Neural Networks and Learning Systems* 33 (2022) 494–514. doi:10.1109/TNNLS.2021.3070843.
- [2] F. MacLean, Knowledge graphs and their applications in drug discovery, *Expert Opinion on Drug Discovery* 16 (2021) 1057–1069. URL: <https://doi.org/10.1080/17460441.2021.1910673>. doi:10.1080/17460441.2021.1910673. arXiv:<https://doi.org/10.1080/17460441.2021.1910673>, pMID: 33843398.
- [3] S. Bonner, I. P. Barrett, C. Ye, R. Swiers, O. Engkvist, A. Bender, C. T. Hoyt, W. L. Hamilton, A review of biomedical datasets relating to drug discovery: a knowledge graph perspective, *Briefings in Bioinformatics* 23 (2022). URL: <https://doi.org/10.1093/bib/bbac404>. doi:10.1093/bib/bbac404. arXiv:<https://academic.oup.com/bib/article-pdf/23/6/bbac404/47144248/bbac404.pdf>, bbac404.
- [4] X. Zeng, X. Tu, Y. Liu, X. Fu, Y. Su, Toward better drug discovery with knowledge graph, *Current Opinion in Structural Biology* 72 (2022) 114–126. URL: <https://www.sciencedirect.com/science/article/pii/S0959440X21001354>. doi:<https://doi.org/10.1016/j.sbi.2021.09.003>.
- [5] S. Chari, O. Seneviratne, D. M. Gruen, M. A. Foreman, A. K. Das, D. L. McGuinness, Explanation ontology: A model of explanations for user-centered ai, 2020. arXiv:2010.01479.
- [6] V. N. Ioannidis, X. Song, S. Manchanda, M. Li, X. Pan, D. Zheng, X. Ning, X. Zeng, G. Karypis, Drkg - drug repurposing knowledge graph for covid-19, <https://github.com/gnn4dr/DRKG/>, 2020.
- [7] Y. Jiannan, L. Zhen, K. K. W. William, Y. Shi, X. Zhongzhi, C. Qian, Z. Qingpeng, Deep learning identifies explainable reasoning paths of mechanism of action for drug repurposing from multilayer biological network, 2022. URL: <https://pubmed.ncbi.nlm.nih.gov/36347526/>.
- [8] T. N. Kipf, M. Welling, Semi-supervised classification with graph convolutional networks, *CoRR abs/1609.02907* (2016). URL: <http://arxiv.org/abs/1609.02907>. arXiv:1609.02907.

- [9] S. Hochreiter, J. Schmidhuber, Long short-term memory, *Neural computation* 9 (1997) 1735–1780.
- [10] M. K. Islam, D. Amaya-Ramirez, B. Maigret, M.-D. Devignes, S. Aridhi, M. Smail, Molecular-evaluated and explainable drug repurposing for covid-19 using ensemble knowledge graph embedding, *Scientific Reports* 13 (2023). doi:10.1038/s41598-023-30095-z.
- [11] A. Bordes, N. Usunier, A. Garcia-Duran, J. Weston, O. Yakhnenko, Translating embeddings for modeling multi-relational data, in: *Advances in neural information processing systems*, 2013, pp. 2787–2795.
- [12] Z. Wang, J. Zhang, J. Feng, Z. Chen, Knowledge graph embedding by translating on hyperplanes, *Proceedings of the AAAI Conference on Artificial Intelligence* 28 (2014). URL: <https://ojs.aaai.org/index.php/AAAI/article/view/8870>. doi:10.1609/aaai.v28i1.8870.
- [13] B. Yang, W. tau Yih, X. He, J. Gao, L. Deng, Embedding entities and relations for learning and inference in knowledge bases, 2015. arXiv:1412.6575.
- [14] M. K. Islam, S. Aridhi, M. Smail-Tabbone, Negative sampling and rule mining for explainable link prediction in knowledge graphs, *Knowledge-Based Systems* 250 (2022) 109083. URL: <https://www.sciencedirect.com/science/article/pii/S0950705122005342>. doi:https://doi.org/10.1016/j.knosys.2022.109083.
- [15] A. Gottlieb, G. Y. Stein, E. Ruppin, R. Sharan, Predict: A method for inferring novel drug indications with application to personalized medicine, 2011. URL: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3159979/>.
- [16] E. Ozkan, R. Celebi, A. Yilmaz, V. Emonet, M. Dumontier, Generating knowledge graph based explanations for drug repurposing predictions, 2023, pp. 22–31. URL: <http://ceur-ws.org/Vol-3415/#paper-3>.
- [17] R. Celebi, J. R. Moreira, A. A. Hassan, S. Ayyar, L. Ridder, T. Kuhn, M. Dumontier, Towards fair protocols and workflows: The openpredict case study, 2019. arXiv:1911.09531.
- [18] W. L. Hamilton, R. Ying, J. Leskovec, Inductive representation learning on large graphs, *NIPS'17*, Curran Associates Inc., Red Hook, NY, USA, 2017, p. 1025–1035.
- [19] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Liò, Y. Bengio, Graph attention networks, in: *International Conference on Learning Representations*, 2018. URL: <https://openreview.net/forum?id=rJXMpikCZ>.
- [20] S. Karen, V. Andrea, Z. Andrew, Deep inside convolutional networks: Visualising image classification models and saliency maps, 2013. URL: <https://arxiv.org/abs/1312.6034>. doi:10.48550/ARXIV.1312.6034.
- [21] M. Saeedi, F. Mehranfar, Challenges and approaches of drugs such as memantine, donepezil, rivastigmine, and aducanumab in the treatment, control and management of alzheimer's disease, *Recent patents on biotechnology* 16 (2022) 102–121.
- [22] J. Knowles, Donepezil in alzheimer's disease: An evidence-based review of its impact on clinical and economic outcomes, 2006. URL: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3321665/>.
- [23] S. Sokolow, X. Li, L. Chen, K. D. Taylor, J. I. Rotter, R. A. Rissman, P. S. Aisen, L. G. Apostolova, Deleterious effect of butyrylcholinesterase k-variant in donepezil treatment of mild cognitive impairment, 2017. URL: [https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5534361/#:~:text=Donepezil%20also%20inhibits%20butyrylcholinesterase%20\(BChE, cholinergic%20function%20in%20the%20brain](https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5534361/#:~:text=Donepezil%20also%20inhibits%20butyrylcholinesterase%20(BChE, cholinergic%20function%20in%20the%20brain).