

“© 2025 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.”

# The Role of Explainable AI in Knowledge Graph-Based Drug Repurposing: Bridging Trust and Discovery

Ethical AI and regulatory challenges

Marwa Mustafa  
Engineering and Information  
Technology  
University of Technology Sydney  
Sydney, Australia  
[marwa.r.mustafa@student.uts.edu.au](mailto:marwa.r.mustafa@student.uts.edu.au)  
[0009-0009-5346-6254](tel:0009-0009-5346-6254)

Firas Al-Doghman  
Engineering and Information  
Technology  
University of Technology Sydney  
Sydney, Australia  
[Firas.Al-Doghman@uts.edu.au](mailto:Firas.Al-Doghman@uts.edu.au)  
[0000-0002-1020-8097](tel:0000-0002-1020-8097)

Farookh Hussain  
Engineering and Information  
Technology  
University of Technology Sydney  
Sydney, Australia  
[Farookh.Hussain@uts.edu.au](mailto:Farookh.Hussain@uts.edu.au)  
[0000-0003-1513-8072](tel:0000-0003-1513-8072)

Mohamed Awadallah  
Engineering and Information  
Technology  
University of Technology Sydney  
Sydney, Australia  
[Mohamed.Awadallah@uts.edu.au](mailto:Mohamed.Awadallah@uts.edu.au)  
[0000-0002-2317-4527](tel:0000-0002-2317-4527)

Mohammed Ali  
School of Engineering  
DSTI School of Engineering  
Paris, France  
[mohammed.ali@edu.dsti.institute](mailto:mohammed.ali@edu.dsti.institute)  
[0000-0002-2317-4527](tel:0000-0002-2317-4527)

**Abstract**— Drug repurposing presents a cost-effective and accelerated alternative to traditional drug development. With Knowledge Graphs (KGs) structuring biomedical relationships and AI models performing predictive link analysis, Explainable AI (XAI) emerges as a critical enabler for interpretability, transparency, and trust. This paper explores the central role of XAI in enhancing the reliability, ethical transparency, and regulatory compliance of AI-powered, KG-based drug repurposing. One of XAI’s most important contributions is its ability to explain why a drug is predicted to treat a particular disease - by revealing key features, paths, or biological relationships that support the prediction. This allows researchers and clinicians to verify the biological plausibility of AI-driven suggestions, which is crucial for preclinical validation and clinical decision-making. The paper also discusses key XAI techniques - including feature attribution, counterfactual analysis, and path reasoning - and the challenges and ethical implications tied to biomedical use cases.

**Keywords**—*Explainable AI; Knowledge Graphs; Drug Repurposing; Graph Neural Networks, Link Prediction; Biomedical AI;*

## I. INTRODUCTION

The process of bringing a new drug to market is highly challenging, requiring substantial time and investment. On average, it takes  $8.3 \pm 2.8$  years and costs around \$374.1 million, though costs can climb to \$1.336 billion [1] when accounting for failures, as over 90% of drug development efforts do not succeed [2] [3] [4]. The drug approval process itself is rigorous and complex, requiring clinical trials in three phases [5]. Phase I focuses on toxicology to assess safety, with approval rates ranging from 3.4% to 32.6%. Phase II evaluates the minimum effective dose and lists potential side effects, yet nearly 60% of drugs fail to progress past this phase due to insufficient efficacy or excessive adverse effects. Phase III then measures the drug’s effectiveness to determine if it should be approved for production and sale. Altogether, moving from Phase I to Phase III typically takes 10 to 15 years [5], delaying drug availability and incurring significant costs.

Due to these high development costs, pharmaceutical companies often prioritize research on diseases with profitable markets, which restricts innovation for rare diseases with fewer patients. In response to these challenges, drug repurposing (also known as drug repositioning) has emerged as an alternative approach. Drug repurposing involves finding new applications for existing drugs or compounds, including those that may not have passed the final stages of clinical trials but could be successful in other therapeutic contexts [3]. By focusing on drugs that have already undergone some evaluation, repurposing can accelerate the discovery of treatments, reduce development costs, and lower the risk of failure, offering a more efficient and cost-effective pathway to meet urgent therapeutic needs [2].

The healthcare industry has faced significant challenges in recent years, highlighted by the urgent need for effective treatments during the COVID-19 pandemic [5]. As the virus spread rapidly, the search for existing drugs that could be repurposed to treat COVID-19 became a priority. Drug repurposing, the practice of finding new therapeutic uses for existing drugs, offers a valuable strategy to accelerate drug discovery, allowing researchers to bypass many of the costly and time-intensive steps required in conventional drug development. Drug repurposing now increasingly leverages computational methods. Knowledge Graphs (KGs), which structure biomedical data into interconnected nodes and edges, and Artificial Intelligence (AI) techniques, including Machine Learning (ML), are revolutionizing the field [1].

AI and ML algorithms enable the analysis of complex and voluminous biomedical datasets, facilitating the identification of potential drug-disease relationships with higher precision and efficiency [3]. These technologies are instrumental in automating link prediction, feature extraction, and scalability within KG frameworks, thereby significantly enhancing the overall effectiveness of drug repurposing efforts. These graphs support drug repurposing by revealing connections that might otherwise remain hidden. However, current KG frameworks often lack real-time data integration and

interpretability, making it difficult for healthcare providers to trust and act on AI-generated predictions [1].

## II. THE IMPERATIVE OF EXPLAINABILITY IN AI-DRIVEN KNOWLEDGE GRAPH PIPELINES FOR DRUG REPURPOSING

Explainable AI (XAI) has become a critical requirement in biomedical applications where interpretability, trust, and accountability are essential. Drug repurposing - also known as drug repositioning - is the strategy of identifying new therapeutic uses for existing drugs outside the scope of their original medical indication. It provides a lower-cost, time-efficient alternative to de novo drug development, especially when safety profiles are already well understood [6] [7] [8]. In this context, where novel therapeutic hypotheses are generated from large-scale biomedical data, black-box AI models can hinder clinical adoption.

Knowledge Graphs (KGs) are structured representations of biomedical entities and their relationships, helping to unify diverse sources such as drugs, diseases, genes, and clinical outcomes into a machine-readable format [2]. While their foundational structure is beyond this paper's scope, KGs serve as essential backbones that allow AI models to reason over biomedical data - enabling tasks like link prediction and relationship inference relevant to drug repurposing. These graphs allow AI models to identify latent relationships, such as new drug-disease associations, using tasks like link prediction or node classification.

However, the increasing use of complex AI models such as Graph Neural Networks (GNNs) [9] in these graphs introduces a critical challenge: their predictions are often difficult to interpret. This is especially problematic in drug repurposing, where incorrect predictions can lead to wasted resources or safety concerns.

Explainable AI (XAI) addresses this issue by clarifying why an AI model suggests a particular drug for a disease. It reveals important nodes, paths, or rules in the graph that justify the prediction - ensuring that the decision is not only accurate but also biologically and clinically reasonable.

This paper explores how these three components - KGs, AI models, and XAI - interact to build an interpretable and ethical pipeline for drug repurposing. We show how XAI transforms opaque AI predictions into transparent and auditable recommendations, aligning with both scientific rigor and regulatory expectations. XAI also helps mitigate risks such as unintended biases or unjustifiable predictions by ensuring every AI-driven recommendation is grounded in traceable, domain-relevant logic. This is particularly vital in biomedical contexts, where ethical oversight and clinical accountability are non-negotiable.

Moreover, ethical and regulatory frameworks increasingly demand transparent AI decision-making, particularly in healthcare. Regulations from bodies like the FDA and EMA call for interpretable models, especially when AI is used to recommend treatments. XAI directly supports compliance by illuminating model behaviour and justifying predictions. This paper places XAI at the centre of the KG-based drug repurposing pipeline and outlines how various XAI methods enhance both predictive insight and ethical governance [10] [11] [12].

Fig. 1 presents a high-level overview of the proposed XAI-enhanced pipeline for knowledge graph-based drug repurposing. The process begins with integrating diverse

biomedical data sources to construct a comprehensive knowledge graph, followed by the application of AI models - such as graph neural networks - to perform link prediction tasks. To ensure transparency and interpretability, explainable AI techniques like GraphSHAP [13] and rule extraction methods are then employed. This layered approach not only facilitates the discovery of novel drug-disease associations

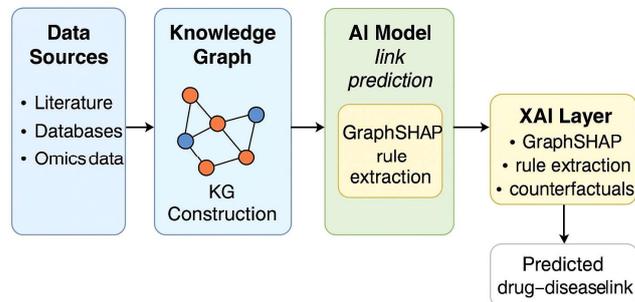


Fig. 1. Overview of XAI-Enhanced KG-Based Drug Repurposing Pipeline

but also provides interpretable justifications that support regulatory and clinical trust.

## III. KNOWLEDGE GRAPHS IN BIOMEDICAL RESEARCH

Recent real-world deployments of knowledge graphs have underscored their critical value during health emergencies such as the COVID-19 pandemic. Drug repurposing played a pivotal role in the pandemic response by allowing rapid exploration of existing drug candidates without starting from scratch. One notable effort, the Drug-CoV KG [5], integrated interaction data from DrugBank and PubChem and used AI models with multi-relational embeddings to prioritize drug candidates targeting SARS-CoV-2. This KG-driven approach accelerated hypothesis generation for potential treatments and exemplified the power of explainable KG frameworks to adapt swiftly in times of crisis.

Knowledge Graphs (KGs) are gaining popularity across various fields, including the life sciences, as structured representations of domain-specific knowledge designed to organize information in a machine-readable format [2]. They often employ semantic structures like Tbox (schema-level) and Abox (instance-level) [1, 14], and represent data using RDF triples (e.g., Drug A - targets - Protein B). A KG is formally defined as a directed, edge-labelled graph  $G = (V, E)$ , where  $V$  is a set of nodes (e.g., drugs, diseases, genes) and  $E$  represents edges or relationships among them [15, 16]. For example, the relationship:

'Drug A → targets → Protein X → linked to → Disease Y'

It illustrates how KGs connect biomedical concepts to support drug discovery.

Biomedical Knowledge Graphs such as Hetionet [17], DRKG [18], BioKG [16], and PrimeKG [8] have proven effective in drug repurposing. These KGs integrate extensive data on drugs, diseases, genes, and interactions, enabling powerful AI-driven predictions by capturing complex biomedical relationships.

To visually support this discussion, two representative biomedical knowledge graphs are illustrated in Fig. 2 [17] and Fig. 3 [18]. These examples demonstrate the typical structure and entity types found in contemporary KGs used for drug repurposing. The graphs highlight critical relationships among

biomedical entities such as genes, diseases, compounds, pathways, and side effects. By visualizing such relational structures, these images underscore how knowledge graphs encode semantic links that can be leveraged for explainable AI models and link prediction tasks in biomedical reasoning.

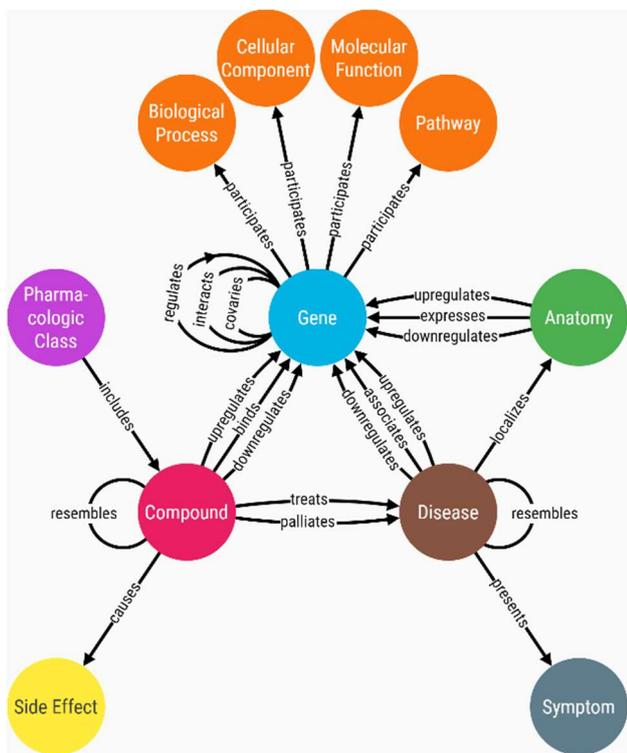


Fig. 2. Hetionet Knowledge Graph.

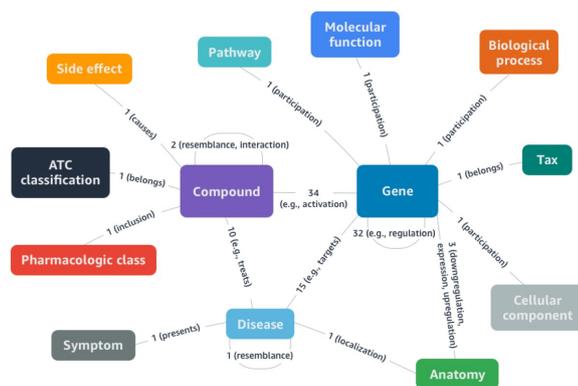


Fig. 3. DRKG Knowledge Graph.

Before applying XAI techniques, AI models are commonly used on Knowledge Graphs for tasks such as link prediction, node classification, and path ranking. These tasks power drug repurposing by uncovering hidden relationships, classifying biomedical entities, and reasoning across multi-step biological pathways. Techniques such as Graph Neural Networks (GCN, R-GCN) [19, 20], embedding-based methods like TransE, ConvE [16, 21], and reinforcement learning agents such as MINERVA and RLPPath [22, 23] are employed to learn meaningful patterns. Additionally, large language models (LLMs) like TxGemma [24] have been used

to extract biomedical triples and support graph construction. While powerful, these models introduce complexity that necessitates transparency – hence the increasing role of XAI to justify and audit their outputs.

In this review, we primarily focus on the application of XAI in the context of link prediction tasks, as these are the most common and impactful use cases for uncovering novel drug-disease associations within biomedical Knowledge Graphs.

### 3. Explainable AI (XAI) in Knowledge Graph-Based Drug Repurposing

Explainable Artificial Intelligence (XAI) aims to make AI systems more transparent by offering human-understandable insights into their decisions [1] [2] [10] [15] [25]. This is especially crucial in biomedical fields, where trust and accountability are essential [11].

A distinction must be drawn between interpretability and explainability in AI. Interpretability is about how easily a human can grasp the internal logic of a model, while explainability concerns how well external methods can clarify or justify the outputs of an opaque system [26].

Deep learning models are often referred to as "black boxes" due to their complex internal mechanisms, which are not readily understood despite clear inputs and outputs [27]. In biomedical contexts, such opacity creates significant challenges in ensuring safety, trust, and regulatory compliance [11].

Unlike "black-box" models, XAI helps explain why a drug is suggested for repurposing, identifies potential biases, and enhances model reliability [15]. It differs from interpretability, which relates to understanding the model's logic, by focusing on external explanations of model behaviour [26].

XAI is especially valuable in drug repurposing, where its ability to generate interpretable justifications can build confidence among researchers and clinicians. It helps reveal why a model recommends a drug for a new indication, mitigates risks from hidden biases, and improves the reliability of AI-driven predictions [15].

#### XAI types:

- Ante-hoc (intrinsic) These are models that are interpretable by design and do not require additional methods to explain their outputs. Common examples include methods like Decision Trees and Logistic Regression [11, 12].
- Post-hoc These methods are applied after training more complex, non-interpretable ("black box") models-such as deep neural networks-to interpret or justify their predictions, such as SHAP, LIME, and Counterfactuals [1, 27].

In this review, we focus primarily on post-hoc methods due to their widespread applicability in black-box models such as Graph Neural Networks (GNNs) [28] and embedding-based knowledge graph completion systems. Post-hoc methods allow for greater flexibility and can be applied to a wide range of biomedical AI models, making them particularly useful in the complex domain of drug repurposing.

#### IV. XAI ALGORITHMS FOR DRUG REPURPOSING USING KNOWLEDGE GRAPHS

Below is an expanded list of XAI techniques tailored to Knowledge Graph-based drug repurposing:

- **GNNExplainer**: Identifies a minimal subgraph and set of node features that are most important for a given prediction by a Graph Neural Network (GNN). This helps highlight key biological relationships such as drug-gene-disease links [29].
- **GraphLIME**: A localized linear model that perturbs the neighbourhood of a node and explains predictions based on changes in the local graph structure. Useful for identifying which molecular descriptors contribute most to drug-disease predictions [30].
- **GraphSHAP**: Applies Shapley value theory to graph structures by assigning importance scores to graph elements (nodes, edges, features). It helps quantify the influence of various biomedical entities on model output [13].
- **eXpath**: Traces and compresses semantic paths in a KG into interpretable rules, such as Drug A  $\rightarrow$  targets  $\rightarrow$  Gene X  $\rightarrow$  associated\_with  $\rightarrow$  Disease Y. These rules explain how a prediction was reached [31].
- **Power-Link**: Constructs multi-hop paths and scores them using GNN-based models. It helps surface semantically meaningful explanations by identifying biologically relevant paths supporting predictions [32].
- **PGM-Explainer**: Builds a probabilistic graphical model (e.g., Bayesian network) to mimic GNN decisions. It reveals probabilistic dependencies between drugs, genes, and diseases that led to a prediction [33].
- **GraphSVX**: Uses random sampling and surrogate models to approximate graph model behavior. Shapley-based scores reveal which nodes and features most impact predictions in noisy biomedical KGs [34].
- **XGNN**: Generates hypothetical graphs and observes model responses to perturbations. This helps determine which edges or nodes are critical for sustaining a prediction [35].
- **CF-GNNExplainer**: Produces counterfactual explanations by identifying the minimal change needed in the graph to alter a prediction. This is useful for model auditing and robustness testing [36].
- **AMIE+**: Uses rule mining to extract logical relationships in biomedical KGs, such as: Drug A targets Gene B, which is associated with Disease C. These rules are directly interpretable and domain-relevant [37].
- **RDF2Rules**: Applies frequent pattern mining to RDF triples in a KG to discover rules. For example: Drug A has side effect X  $\rightarrow$  Symptom X is part of Disease B  $\rightarrow$  Drug A may treat Disease B [38].
- **SeXAI**: Translates model outputs into First-Order Logic rules using ontologies. This aligns AI predictions with formal biomedical knowledge, improving trust and interpretability [10].

These algorithms collectively enhance understanding, transparency, and accountability in AI-driven drug repurposing by clarifying why and how decisions are made.

#### V. HYBRID AND ADVANCED XAI APPROACHES IN KNOWLEDGE GRAPH-BASED DRUG REPURPOSING

In biomedical AI applications like drug repurposing using Knowledge Graphs (KGs), combining multiple XAI methods can significantly enhance interpretability and trust. Instead of relying on a single technique, hybrid approaches that integrate path-based, rule-based, and feature attribution methods offer more robust and transparent explanations of AI decisions [39].

A common and powerful hybrid strategy merges knowledge graph embeddings with multi-hop reasoning and logical rule extraction. This combined method not only supports accurate predictions but also ensures those predictions are traceable and understandable to biomedical researchers and clinicians [39].

For example, the KGML-xDTD model integrates reinforcement learning, knowledge graphs, and adversarial training to identify interpretable paths between drugs and diseases. It addresses black-box limitations by providing path-based justifications for its link predictions in biomedical scenarios [40].

Rule-based reasoning is also applied in explainable COVID-19 drug repurposing. In this context, logical rules are extracted from KGs to explain how drugs are connected to diseases via intermediary entities such as genes and proteins - thereby generating biologically meaningful explanations [39].

The KGExplainer framework enhances interpretability by extracting symbolic rules from knowledge graph embeddings (KGEs). Using a five-step process and surrogate models like MDI, K-Lasso, and HSIC-Lasso, it generates instance-based and rule-based explanations without retraining the original model. This method scales well to large biomedical KGs and maintains fidelity to the source model [27].

Another notable method is Example (Explanation Graph), which identifies training examples in latent KGE space that influenced a prediction. It reconstructs semantic subgraphs to show how similar cases contributed to the model's decision, helping illuminate the reasoning behind predictions in link prediction tasks [41].

TxGNN is a clinician-focused, graph-based foundation model designed for zero-shot drug repurposing, especially in rare or poorly understood diseases. It employs metric learning with GNNs and offers multi-hop explanations directly from medical knowledge graphs to increase transparency and usability in clinical settings [42].

Similarly, rd-explainer targets rare disease drug repurposing by combining domain-specific KGs with GNNExplainer. It supports its predictions with semantic subgraph-based explanations, ensuring alignment with biomedical knowledge and improved interpretability [43].

Table I provides a comparative summary of selected models that integrate Explainable AI (XAI) techniques within Knowledge Graph (KG)-based drug repurposing systems. Each model leverages a distinct explainability mechanism - ranging from local feature attribution (e.g., GNNExplainer, GraphSHAP) to rule-based logic extraction (e.g., AMIE+, RDF2Rules). The table highlights the type of explanation each

method offers (such as path-based, rule-based, or counterfactual), along with its specific biomedical application. By outlining the core principles and use cases, this table helps illustrate the diversity and growing maturity of XAI approaches in this emerging research domain.

TABLE I. TABLE TYPE STYLES

Model	Technique Used	Explanation Output	Application Scenario
KGML-xDTD	Reinforcement Learning + Paths	Path-based explanation for drug-disease links	COVID-19 repurposing
TxGNN	GNN + Metric Learning	Multi-hop explanations for zero-shot tasks	Rare disease predictions
KGExplainer	Rule extraction from embeddings	Instance & rule-based explanations	Interpretability in large KGEs
Example	Latent influence tracing	Semantic subgraphs based on similar examples	Transparent link prediction
rd-Explainer	GNNExplainer + biomedical KG	Semantic subgraph highlighting key paths	Rare disease drug validation

## VI. CHALLENGES IN XAI AND KG-BASED DRUG REPURPOSING

The integration of Knowledge Graphs (KGs) and Explainable AI (XAI) for drug repurposing introduces multifaceted challenges across technical, computational, and practical dimensions.

### 1. Data Integration and Standardization:

Integrating data from structured (e.g., DrugBank) and unstructured sources (e.g., scientific literature) with inconsistent schemas and terminologies remains a significant hurdle. Effective use of NLP and entity resolution tools is needed to harmonize these sources [44].

### 2. Scalability and Computational Requirements:

Managing large biomedical KGs, often containing millions of nodes and edges, demands high-performance computing infrastructure. Likewise, deep learning and LLM-based knowledge extraction further intensify memory and processing needs [45].

### 3. Data Quality and Dynamic Updates:

Biomedical KGs must maintain high data accuracy and consistency while being frequently updated with new discoveries. However, automating these updates without introducing noise or redundancy is challenging and often requires human-in-the-loop systems [46, 47].

### 4. Multimodal and Complex Data Handling:

Incorporating various data types such as genomic sequences, imaging, molecular descriptors, and clinical texts requires sophisticated multimodal alignment strategies to maintain semantic consistency [45].

### 5. Evaluation and Benchmarking:

There is no standardized framework for evaluating the quality, completeness, or explainability of KG-based AI systems. This lack of uniform metrics hinders reproducibility and trust in the outputs [48, 49].

### 6. Security and Privacy Risks:

Providing detailed explanations might inadvertently expose sensitive biomedical data or make systems more vulnerable to adversarial attacks, particularly in healthcare domains [26, 27].

### 7. Audience-Specific Explanations:

Different stakeholders (e.g., clinicians, regulators, researchers) require customized forms of explanation. A one-size-fits-all model is insufficient for practical deployment [1, 2].

### 8. High Computational Cost of XAI

Many XAI methods involve heavy computation, such as surrogate modeling and iterative sampling, which can limit real-time applicability [1] [2].

## VII. DISCUSSION AND FUTURE WORK

Future research should prioritize the seamless integration of Knowledge Graph (KG) construction and Explainable AI (XAI) to advance the effectiveness and trustworthiness of drug repurposing systems. For KG construction, efforts must focus on developing scalable, distributed architectures capable of handling the rapid expansion of biomedical data while ensuring high data quality, completeness, and consistency. Semi-automated update mechanisms - guided by human-in-the-loop validation - will be crucial for maintaining relevance in dynamic domains such as pharmacogenomics and disease progression.

In parallel, XAI approaches must evolve to provide interpretable, secure, and computationally efficient explanations tailored to end-users ranging from clinical researchers to regulatory experts. This includes establishing standardized evaluation frameworks for measuring explanation quality, incorporating privacy-preserving techniques to mitigate data leakage risks, and building adaptive explanation systems that adjust outputs based on user expertise and context.

Furthermore, addressing the high computational cost of current explanation techniques will require novel approximation strategies that maintain fidelity without sacrificing performance. By co-developing these capabilities, future systems will not only uncover novel therapeutic insights more reliably but also provide transparent, evidence-driven justifications necessary for real-world adoption in healthcare and biomedical research.

## REFERENCES

- [1] J. Chen *et al.*, "Knowledge graphs for the life sciences: recent developments, challenges and opportunities. arXiv 5," *arXiv preprint arXiv:2309.17255*, 2023.
- [2] P. Perdomo-Quinteiro and A. Belmonte-Hernández, "Knowledge Graphs for drug repurposing: a review of databases and methods," *Briefings in Bioinformatics*, vol. 25, no. 6, p. bbac461, 2024.

- [3] X. Pan *et al.*, "Deep learning for drug repurposing: Methods, databases, and applications," *Wiley interdisciplinary reviews: Computational molecular science*, vol. 12, no. 4, p. e1597, 2022.
- [4] D. Sun, W. Gao, H. Hu, and S. Zhou, "Why 90% of clinical drug development fails and how to improve it?," *Acta Pharmaceutica Sinica B*, vol. 12, no. 7, pp. 3049-3062, 2022.
- [5] S. Li, K. W. Wong, D. Zhu, and C. C. Fung, "Drug-CoV: a drug-origin knowledge graph discovering drug repurposing targeting COVID-19," *Knowledge and Information Systems*, vol. 65, no. 12, pp. 5289-5308, 2023.
- [6] F. Yang *et al.*, "Machine learning applications in drug repurposing," *Interdisciplinary Sciences: Computational Life Sciences*, vol. 14, no. 1, pp. 15-21, 2022.
- [7] C. Königs, M. Friedrichs, and T. Dietrich, "The heterogeneous pharmacological medical biochemical network PharMeBINet," *Scientific Data*, vol. 9, no. 1, p. 393, 2022.
- [8] P. Chandak, K. Huang, and M. Zitnik, "Building a knowledge graph to enable precision medicine," *Scientific Data*, vol. 10, no. 1, p. 67, 2023.
- [9] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, and P. S. Yu, "A comprehensive survey on graph neural networks," *IEEE transactions on neural networks and learning systems*, vol. 32, no. 1, pp. 4-24, 2020.
- [10] S. Schramm, C. Wehner, and U. Schmid, "Comprehensible artificial intelligence on knowledge graphs: A survey," *Journal of Web Semantics*, vol. 79, p. 100806, 2023.
- [11] M. Mersha, K. Lam, J. Wood, A. AlShami, and J. Kalita, "Explainable artificial intelligence: A survey of needs, techniques, applications, and future direction," *Neurocomputing*, p. 128111, 2024.
- [12] P. Gohel, P. Singh, and M. Mohanty, "Explainable AI: current status and future directions," *arXiv preprint arXiv:2107.07045*, 2021.
- [13] A. Perotti, P. Bajardi, F. Bonchi, and A. Panisson, "Explaining identity-aware graph classifiers through the language of motifs," in *2023 International joint conference on neural networks (IJCNN)*, 2023: IEEE, pp. 1-8.
- [14] V. K. Kommineni, B. König-Ries, and S. Samuel, "From human experts to machines: An LLM supported approach to ontology and knowledge graph construction," *arXiv preprint arXiv:2403.08345*, 2024.
- [15] A. Jiménez, M. J. Merino, J. Parras, and S. Zazo, "Explainable drug repurposing via path based knowledge graph completion," *Scientific Reports*, vol. 14, no. 1, p. 16587, 2024.
- [16] B. Walsh, S. K. Mohamed, and V. Nováček, "Biokg: A knowledge graph for relational learning on biological data," in *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, 2020, pp. 3173-3180.
- [17] D. S. Himmelstein *et al.*, "Systematic integration of biomedical knowledge prioritizes drugs for repurposing," *Elife*, vol. 6, p. e26726, 2017.
- [18] V. N. Ioannidis *et al.*, "Drkg-drug repurposing knowledge graph for covid-19," *arXiv preprint arXiv:2010.09600*, 2020. [Online]. Available: <https://github.com/gnn4dr/DRKG/>.
- [19] R. Yue and A. Dutta, "Repurposing Drugs for Infectious Diseases by Graph Convolutional Network with Sensitivity-Based Graph Reduction," *Interdisciplinary Sciences: Computational Life Sciences*, vol. 17, no. 1, pp. 185-199, 2025.
- [20] M. Schlichtkrull, T. N. Kipf, P. Bloem, R. Van Den Berg, I. Titov, and M. Welling, "Modeling relational data with graph convolutional networks," in *The semantic web: 15th international conference, ESWC 2018, Heraklion, Crete, Greece, June 3-7, 2018, proceedings 15*, 2018: Springer, pp. 593-607.
- [21] Y. Chen, Y. Wu, S. Ma, and I. King, "A literature review of recent graph embedding techniques for biomedical data," in *Neural Information Processing: 27th International Conference, ICONIP 2020, Bangkok, Thailand, November 18-22, 2020, Proceedings, Part V 27*, 2020: Springer, pp. 21-29.
- [22] R. Das *et al.*, "Go for a walk and arrive at the answer: Reasoning over paths in knowledge bases using reinforcement learning," *arXiv preprint arXiv:1711.05851*, 2017.
- [23] L. Chen, J. Cui, X. Tang, Y. Qian, Y. Li, and Y. Zhang, "Rlpath: a knowledge graph link prediction method using reinforcement learning based attentive relation path searching and representation learning," *Applied Intelligence*, pp. 1-12, 2022.
- [24] E. Wang *et al.*, "Txgemma: Efficient and agentic llms for therapeutics," *arXiv preprint arXiv:2504.06196*, 2025.
- [25] S. Ali *et al.*, "Explainable Artificial Intelligence (XAI): What we know and what is left to attain Trustworthy Artificial Intelligence," *Information fusion*, vol. 99, p. 101805, 2023.
- [26] A. B. Arrieta *et al.*, "Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI," *Information fusion*, vol. 58, pp. 82-115, 2020.
- [27] C. Wehner, C. Iliopoulou, and T. R. Besold, "From Latent to Lucid: Transforming Knowledge Graph Embeddings into Interpretable Structures," *arXiv preprint arXiv:2406.01759*, 2024.
- [28] J. Zhou *et al.*, "Graph neural networks: A review of methods and applications," *AI open*, vol. 1, pp. 57-81, 2020.
- [29] Z. Ying, D. Bourgeois, J. You, M. Zitnik, and J. Leskovec, "Gnnexplainer: Generating explanations for graph neural networks," *Advances in neural information processing systems*, vol. 32, 2019.
- [30] Q. Huang, M. Yamada, Y. Tian, D. Singh, and Y. Chang, "Graphlime: Local interpretable model explanations for graph neural networks," *IEEE Transactions on Knowledge and Data Engineering*, vol. 35, no. 7, pp. 6968-6972, 2022.
- [31] Y. Sun, L. Shi, and Y. Tong, "eXpath: Explaining Knowledge Graph Link Prediction with Ontological Closed Path Rules," *arXiv preprint arXiv:2412.04846*, 2024.

- [32] H. Chang, J. Ye, A. Lopez-Avila, J. Du, and J. Li, "Path-based explanation for knowledge graph completion," in *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2024, pp. 231-242.
- [33] M. Vu and M. T. Thai, "Pgm-explainer: Probabilistic graphical model explanations for graph neural networks," *Advances in neural information processing systems*, vol. 33, pp. 12225-12235, 2020.
- [34] A. Duval and F. D. Malliaros, "Graphsvx: Shapley value explanations for graph neural networks," in *Machine Learning and Knowledge Discovery in Databases. Research Track: European Conference, ECML PKDD 2021, Bilbao, Spain, September 13–17, 2021, Proceedings, Part II 21*, 2021: Springer, pp. 302-318.
- [35] H. Yuan, J. Tang, X. Hu, and S. Ji, "Xggn: Towards model-level explanations of graph neural networks," in *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining*, 2020, pp. 430-438.
- [36] A. Lucic, M. A. Ter Hoeve, G. Tolomei, M. De Rijke, and F. Silvestri, "Cf-gnnexplainer: Counterfactual explanations for graph neural networks," in *International Conference on Artificial Intelligence and Statistics*, 2022: PMLR, pp. 4499-4511.
- [37] L. Galárraga, C. Teflioudi, K. Hose, and F. M. Suchanek, "Fast rule mining in ontological knowledge bases with AMIE  $\mathbb{S}^+ \mathbb{S}$ ," *The VLDB Journal*, vol. 24, no. 6, pp. 707-730, 2015.
- [38] Z. Wang and J. Li, "RDF2Rules: Learning rules from RDF knowledge bases by mining frequent predicate cycles," *arXiv preprint arXiv:1512.07734*, 2015.
- [39] M. K. Islam, D. Amaya-Ramirez, B. Maigret, M.-D. Devignes, S. Aridhi, and M. Smail-Tabbone, "Molecular-evaluated and explainable drug repurposing for COVID-19 using ensemble knowledge graph embedding," *Scientific Reports*, vol. 13, no. 1, p. 3643, 2023.
- [40] C. Ma, Z. Zhou, H. Liu, and D. Koslicki, "KGML-xDTD: a knowledge graph-based machine learning framework for drug treatment prediction and mechanism description," *GigaScience*, vol. 12, p. giad057, 2023.
- [41] A. Janik and L. Costabello, "Explaining Link Predictions in Knowledge Graph Embedding Models with Influential Examples," *arXiv preprint arXiv:2212.02651*, 2022.
- [42] K. Huang *et al.*, "A foundation model for clinician-centered drug repurposing," *Nature Medicine*, vol. 30, no. 12, pp. 3601-3613, 2024.
- [43] P. Perdomo-Quinteiro, K. Wolstencroft, M. Roos, and N. Queralt-Rosinach, "Knowledge graphs and explainable ai for drug repurposing on rare diseases," *bioRxiv*, p. 2024.10.17.618804, 2024.
- [44] K. PETROVIČOVÁ, "Knowledge Graphs and Explainable Predictive Models for Drug Repurposing," MASARYK University, 2023.
- [45] G. Tamašauskaitė and P. Groth, "Defining a knowledge graph development process through a systematic review," *ACM Transactions on Software Engineering and Methodology*, vol. 32, no. 1, pp. 1-40, 2023.
- [46] B. Abu-Salih, "Domain-specific knowledge graphs: A survey," *Journal of Network and Computer Applications*, vol. 185, p. 103076, 2021.
- [47] X. Zhu *et al.*, "Multi-modal knowledge graph construction and application: A survey," *IEEE Transactions on Knowledge and Data Engineering*, vol. 36, no. 2, pp. 715-735, 2022.
- [48] I. Tiddi and S. Schlobach, "Knowledge graphs as tools for explainable machine learning: A survey," *Artificial Intelligence*, vol. 302, p. 103627, 2022.
- [49] G. Weikum, X. L. Dong, S. Razniewski, and F. Suchanek, "Machine knowledge: Creation and curation of comprehensive knowledge bases," *Foundations and Trends® in Databases*, vol. 10, no. 2-4, pp. 108-490, 2021.