Explaining Drug Repositioning: A Case-Based Reasoning Graph Neural Network Approach

Carolina González-Cavazos,*
The Scripps Research Institute
La Jolla, CA, 92037
agonzalez@scripps.edu

Roger Tu,
The Scripps Research Institute
La Jolla, CA, 92037
rogertu@scripps.edu

Meghamala Sinha,
The Scripps Research Institute
La Jolla, CA, 92037
msinha@scripps.edu

Andrew I. Su
The Scripps Research Institute
La Jolla, CA, 92037
asu@scripps.edu

Abstract

Drug repositioning, the identification of novel uses of existing therapies, has become an attractive strategy to accelerate drug development. Knowledge graphs (KGs) have emerged as a powerful representation of interconnected data within the biomedical domain. While link prediction on biomedical can ascertain new connections between drugs and diseases, most approaches only state *whether* two nodes are related. Yet, they fail to explain *why* two nodes are related. In this project, we introduce an implementation of the semi-parametric Case-Based Reasoning over subgraphs (CBR-SUBG), designed to derive a drug query's underlying mechanisms by gathering graph patterns of similar nodes. We show that our adaptation outperforms existing KG link prediction models on a drug repositioning task. Furthermore, our findings demonstrate that CBR-SUBG strategy can provide interpretable biological paths as evidence supporting putative repositioning candidates, leading to more informed decisions.

1 Introduction

Traditional drug discovery and development is a time-consuming and resource-intensive process [1]. Repositioning existing drugs for new therapeutic uses has emerged as a strategy aimed at reducing drug development costs by decreasing failure rates and time to market[2].

Over recent years, there has been an increasing interest in utilizing biomedical knowledge graphs (KGs) for drug repositioning [3, 4, 5, 6]. Briefly, biomedical KGs consist of nodes representing biological concepts (such as genes, drugs, diseases, and pathways) and edges describing their relationship (such as drugs treating diseases, or diseases associated with genes) [7].

A growing number of drug repositioning link prediction approaches on biomedical KGs have emerged to predict connections of drugs to candidate diseases [8, 9, 10, 11]. However, these methods focus on mapping nodes and relations to vector representations via translational distance models, matrix factorization, or neural network models. Then, a scoring function is defined to measure the correctness of a triple in the embedding space. While these projects have been successful predicting novel indication candidates, the limitation of all these methods is that they neither consider nor explain the biological mechanism of drugs and their relationship with a disease.

*Code available at the following Github repository:github.com/Carolina1396/CBR-SUBG_MIND.git

NeurIPS 2023 AI for Science Workshop.

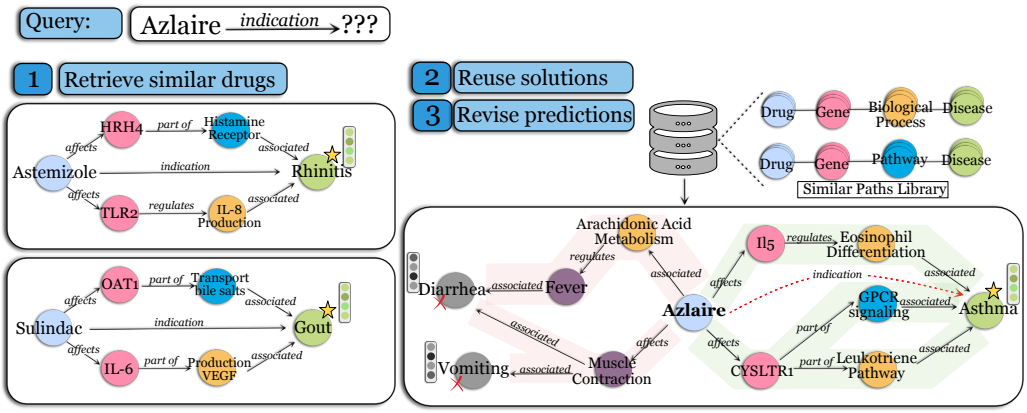

Figure 1: Explaining CBR-SUBG workflow for drug repositioning. Given a query (azlaire, *indication*, ?), (1) CBR-SUBG first **retrieves** k-NN similar drug nodes (astemizole, sulindac) to the input query and gathers their reasoning paths that lead to the corresponding disease they treat. (2) Next, these paths are **reused** to address the query of interest. (3) Finally, all gathered nodes by reasoning chains are **revised**. The disease answer node (star) is the node within the query subgraph that exhibits the highest similarity to the disease answer nodes of the k-NN nodes. Note that the reused paths required to answer the query repeat over similar drugs (highlighted in green), and the reasoning chains not obtained from similar drugs (highlighted in red) lead to incorrect answers.

As a solution, we explore Case-Based Reasoning (CBR) — an artificial intelligence approach that learns from past experiences to solve new problems [12]. At a high level, the CBR system comprises the following main steps: **retrieve** cases that are similar to the given problem, **reuse** solutions of similar cases for the problem in hand, and **revise** if the solutions are appropriate for solving the given problem [13].

In this paper we introduce a novel semi-parametric adaptation of CBR strategy over subgraphs (CBR-SUBG) focused on drug repositioning. Previously, Das et al. [14] demonstrated the utility and effectiveness of this semi-parametric approach for natural language queries [15]. Here, we leverage reasoning paths from similar drugs to identify potential repositioning candidates for a given drug of interest. Our framework utilizes a Graph Neural Network (GNN) trained on a queried drug's CBR extracted subgraph patterns. As the reasoning patterns of similar drug cases (solutions) reoccur within the subgraph of a drug query of interest, a disease answer node can be found analogously located within the subgraph.

Our CBR-SUBG implementation, not only outperforms sophisticated GNN-based models but also enables interpretable biological mechanisms identification to support a given drug's repositioning. Moreover, we present a case study on azlaire and Parkinson's disease, demonstrating the potential of CBR-SUBG for drug repositioning and its effectiveness in uncovering intermediate biological nodes that explain novel drug-disease associations.

2 Related work

Case-based reasoning models A simple non-parametric CBR approach for KG completion was proposed by Das et al. [16, 17]. Simple CBR finds reasoning chains for a query of interest by retrieving similar cases. However, it only handles symbolic matching to find similarities between a query and similar cases. More recently, Das et al. proposed a CBR-SUBG model that uses soft-matching to compare representations of answer queries and similar answer cases [14, 15]. Our CBR-SUBG implementation for drug repositioning task takes advantage of derived reasoning patterns from similar answer cases to explain putative drug reasoning candidates.

Graph Neural Networks In knowledge graph completion frameworks, different GNN architectures, including relational-GCN (R-GCN) [18], and composition-based multi-relational GCN (CompGCN) [19], serve as encoders of the entire immediate neighborhood of each node, along with a scoring

function such as DistMult [20], and ConvE [21]. While CBR-SUBG is not dependent on any specific Message Passing Neural Networks (MPNN) scheme, in contrast to other frameworks, a model like CBR-SUBG generates tailored subgraphs, enabling the collection of relevant chains for a query.

3 Methods

Notation A KG is a directed labeled multi-graph $G = (\mathcal{E}, \mathcal{R})$, defined by a set of nodes $(\mathcal{E})$ and a set of edges $(\mathcal{R})$. A KG is a collection of facts represented as triplets of the form (e_1, r, e_2) where $e_1, e_2 \in \mathcal{E}$ and $r \in \mathcal{R}$. A path (p) in a KG is defined as a sequence of alternating nodes and edges connecting a starting node and an ending node. Formally, a path is represented as $p = (e_1, r_1, e_2, ...r_n, e_{n+1})$.

Task Description We formulate drug repositioning as a series of query answering on a biomedical KG. As described, in CBR a new problem is solved based on similar previously observed cases. In our setting, we define a query problem (q) as $(e_1, r_1, ?)$ (e.g. Azlaire, *indication*, ?), where e_1 denotes a particular drug entity and r_1 the missing relation "*indication*" that links to a disease node. Here we define a case (c) as an observed indication (e_1, r, e_2) in the KG along with a set of biological reasoning paths (p) that connect a drug e_1 to the corresponding disease e_2 it treats.

Method overview As depicted in **Figure 1**, given a drug query input (q), CBR-SUBG first retrieves k-nearest neighbor drug cases (k-NN$_q$). For each retrieved case c, a collection of reasoning chains p leading to its corresponding disease answers were collected. Next, paths retrieved were reused on the input query to form a subgraph where a GNN encodes the underlying subgraph structure into a node representation [22]. Lastly, CBR-SUBG reasons across the drug query and retrieved k-NN$_q$ drug's subgraphs to identify the correct query answer.

3.1 Model Description

In this section we describe CBR-SUBG's composition and its implementation for drug repositioning task following three main steps of CBR framework; retrieving similar cases, reusing solutions, and revising results.

1- Retrieve Similar Drug Cases Given the input drug query q, CBR-SUBG first retrieves k-NN$_q$ similar drug cases c using a pre-computed similarity matrix $S \in \mathbb{R}^{\mathcal{E} \times \mathcal{E}}$ that stores the similarity score between all pairs of nodes. To model this, each drug entity is parameterized in an *m*-hot vector $e \in \mathbb{R}^{\mathcal{R}}$, with dimensions equal to the number of relation types on the KG. An entry in the vector is set to 1 if a drug entity has at least one edge with that relation type, otherwise is set to 0. The similarity score between two drug queries is given by the cosine similarity between their normalized vector representations. Naturally, two drug nodes that capture the same relation types should have a high similarity score. For each retrieved drug case, CBR-SUBG gathers the paths in the graph that connect the drug entity to the corresponding disease it treats. Note that since the number of collected paths between two nodes can grow exponentially, we only considered 1000 randomly sampled paths of length up to three around each drug.

2- Reuse Drug Solutions Next, the gathered reasoning paths of retrieved similar drug cases are reused for the query of interest. Beginning with the drug query node, the sequence of relation types (path types) is applied. An example is depicted in Figure 1 where astemizole is retrieved as one of the k-NN$_q$ similar drugs to our query, azlaire. The recovered path: (astemizole, *affects*, HRH4, *part of*, histamine receptor, *associated*, rhinitis), is reused by gathering the sequence of relation types such as: *affects* (drug, gene) $\wedge$ *affects* (gene, pathway) $\wedge$ *associated* (pathway, disease). After iterating over all similar retrieved cases, the collected chains form a subgraph for the drug query.

3- Revise Disease Predictions Once the subgraphs for the drug query and the corresponding k-NN$_q$ similar drugs are defined, CBR-SUBG reasons across them. For this, drug representations are encoded with a GNN to incorporate its respective local subgraph structure. Finally, the most similar disease node in the query subgraph to the answers in the k-NN$_q$ subgraph drugs is retrieved as the answer.

Encoding node representations Considering that biomedical KGs are heterogeneous graphs with labeled edges, we employed the multi-relational R-GCN model [18]. We followed the general MPNN scheme that iteratively updates the representation of each node by aggregating the representation of its neighbors [23]. In particular the lth layer of GNN is given by,

a_e^l = AGGREGATE$^l(h_s^{l-1}$: $s \in \mathcal{N}(e)$, $h_e^{l-1})$ (1)

h_e^l = COMBINE$^l(h_e^{l-1}, a_e^l)$ (2)

where, a_e^l is the aggregated message form the neighbors, h_e^l denotes the representation of node *e* in the lth layer, and $\mathcal{N}(e)$ denotes the set of immediate neighbors of node *e*. For multi-relational R-GCN model, these steps are defined as:

$a_e^l = \sum_{r=1}^{\mathcal{R}} \sum_{e \in \mathcal{N}_r(e)} W_r^l h_e^{l-1}$(3)

h_e^l = ReLU($W_{self}^l h_e^{l-1} + a_e^l$) (4)

where, $\mathcal{R}$ is the total number of relation types captured in the KG, $\mathcal{N}_r(e)$ denotes the immediate outgoing neighbors of node *e* under relation *r*; W_r^l is the transformation matrix used to propagate the message in the lth layer for relation *r*.

Training During training, the disease answer node (d_i) of the query subgraph (G_i) and the disease answer node (d_j) of the k-NN$_q$ subgraphs (G_j) are trained to be more similar to each other in comparison to incorrect disease answer nodes. This similarity is calculated as the cosine similarity between the normalized disease answer representations. Considering that retrieved similar drug case could arrive at a set of answers (D_j), we calculate the mean of the scores between d_i and all answer nodes in G_j as *sim*(d_i, D_j)= $\frac{1}{|D_j|} \sum_{d_j \in D_j}$ ($d_i^\top$ d_j / $\|d_i\|$ $\|d_j\|$). Then, the similarity mean score of all retrieved similar drug cases is aggregated for the given drug query. The loss function used is an extension of the normalized temperature-scaled cross entropy loss used in Chen et al. [24] and introduced by Das et al. [15]:

$$L = -log \frac{\sum_{d_i \in D_i} exp(\sum_{q_j \in kNN_{qi}} sim(d_i, D_j)/\mathcal{T})}{\sum_{x_i \in \mathcal{E}(G_i)} exp(\sum_{q_j \in kNN_{qi}} sim(x_i, D_j)/\mathcal{T})} (5)$$

where, D_i represents the set of all answer nodes in G_i, while D_j denotes the set of all answer nodes of retrieved similar k-NN$_q$ drug cases, $x_i \in G_i$ represents all nodes in the subgraph of drug query and $\mathcal{T}$ denotes the temperature hyperparameter. This way, the loss score the disease answer nodes in D_i higher than all other nodes in G_i with respect to the answer nodes in k-NN$_q$ drug subgraph.

Reasoning Message-passing is run across both the drug query subgraph and the subgraphs of retrieved k-NN$_q$ drugs, resulting in the generation of node representations. A ranking of similarity scores is returned, where the highest score node in the query subgraph with respect to all answer nodes in the retrieved drug subgraphs is returned as a more likely repositioning answer. For this,

$$ans = argmax \sum_{d_i \in D_i} exp(\sum_{q_j \in kNN_{qi}} sim(d_i, D_j)) (6)$$

4 Experimental setup

Dataset Details In this section, we conduct experiments on Mechanistic Repositioning Network with Indications (MIND), a biomedical knowledge graph that integrates two biomedical resources: Mechanistic Repositioning Network (MechRepoNet) and DrugCentral [25]. Briefly, MechRepoNet is a comprehensive biomedical knowledge graph that was constructed by integrating 18 different data sources, consisting of of 9,652,116 edges, 250,035 nodes, 9 node types and 22 relations [4]. DrugCentral, is a publicly available online resource that incorporates information from indications that have received approval from regulatory agencies [26]. Here, we divided MIND into subsets: train (80%, 2087 indications), and test (20%, 390 indications).

Hyperparameter optimisation For CBR-SUBG model, we optimized hyperparameters using Optuna software [27]. Next, our model was trained and tested on datasets mentioned above. We used one GCN layer with 64 dimensions and a dropout of 0.7965, where edges are randomly dropped from the graph while aggregating information from the neighborhood. For training and evaluation we used 5 and 10 nearest neighbors, respectively. We optimized the loss using Adam Optimizer. The learning rate is set to be 0.01 with a temperature value of 0.1053.

Evaluation We compared CBR-SUBG with various state-of the art models using standard ranking metrics: Hits@N metric, which assesses the accuracy of identifying the correct disease answer within the top N predictions, and mean reciprocal rank (MRR), the average of the reciprocal of the positional rank for all correct disease answers.

5 Results

5.1 Performance Comparison

Table 1: Inference performance results on MIND

Model	MRR	Hits@1	Hits@3	Hits@5	Hits@10
RGCN + DistMult	0.1157	0.0180	0.1084	0.2108	0.3494
RGCN + ConvE	0.1197	0.0240	0.1144	0.1867	0.3614
CompGCN + DistMult	0.27735	0.1954	0.30651	0.36398	0.44828
CompGCN + ConvE	0.36112	0.25287	0.41762	0.47126	0.55556
Non-parametric CBR	0.03269	0.0158	0.0317	0.0396	0.0634
CBR-SUBG	**0.3770**	**0.2796**	**0.4329**	**0.4789**	**0.5708**

In this section, we evaluate the performance of CBR-SUBG and other KG completion state-of-the-art models on the MIND dataset to identify drug repositioning candidates. In particular, we compare to other GNN-based link prediction models such as R-GCN [18], and CompGCN [19], as well as the non-parametric CBR model [16] that gathers reasoning chains to find an answer. Evaluation of GNN-based models was conducted using ConvE [21] and DistMult [20] as scoring functions. As demonstrated in **Table 1**, our implementation of CBR-SUBG significantly outperforms all the baselines in all metrics. We further discuss the results in section 6.

5.2 Propranolol Case Study

Propranolol was first approved by FDA as a beta-blocker to manage heart-related issues, such as hypertension [28]. Since beta blockers were known to relax the cardiac muscle and block nerve impulses to the muscle, it was later approved for essential tremor [29]. Therefore, we investigated if CBR-SUBG could have predicted this repositioning within MIND (**Figure 2**). While hypertension was predicted in the first position, Parkinson's disease was the second prediction. Notably, one of the primary symptoms of Parkinson's disease is tremor, aligning with propranolol's repositioning potential [30].

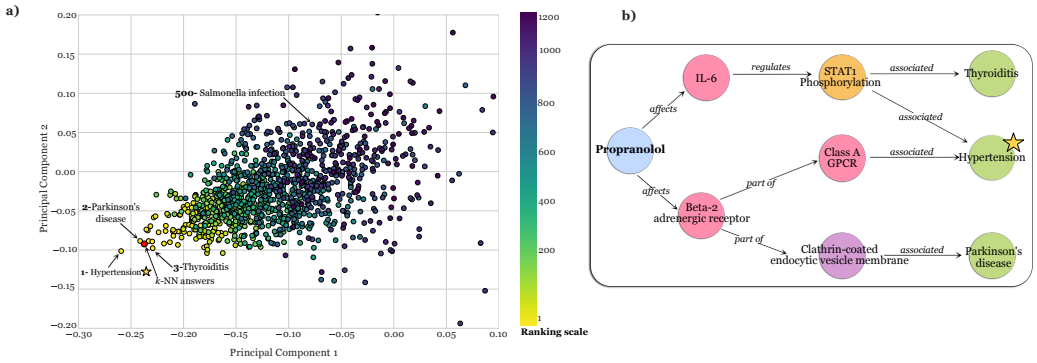

Figure 2: Propranolol case study. a) PCA visualization of propanolol disease predictions. The red dot depicts the mean representation of answer nodes derived from k-NN similar drugs. Color scale illustrates repositioning ranking. Expected true answer (star) and top two predictions are displayed. Note that prediction 500 is shown as scale reference. b) Evidence supporting propranolol prediction for true answer (star) and top two repositioning answers.

Considering the CBR hypothesis, where the representation of answer nodes of drug query q and retrieved k-NN$_q$ drug subgraphs should be similar as the local structure around them share similarities. A PCA visualization of propranolol predictions shows that answer nodes within retrieved k-NN$_q$ drug query subgraphs cluster together alongside the top repositioning answers (**Figure 2a**). As illustrated, as the node representation similarity of disease predictions gets apart w.r.t k-NN$_q$ answers, the ranking of them decreases.

Illustrated in **Figure 2b**, CBR-SUBG model is able to provide interpretable chains for propranolol indication predictions. These reasoning paths enable the identification of shared biological targets among repositioning answers, ultimately enhancing the quality of decision-making with more comprehensive insights. As illustrated, propranolol binds to beta-2 receptor [31], member of GCPR family, playing an important role in regulating heart rate [32]. Furthermore, beta-2 receptor is associated with an increased risk of developing Parkinson [33]. Moreover, we detected that propranolol targets interleukin-6 [34] regulating phosphorylation of transcription factor STAT1. This modulation is documented to take part in the pathogenesis of various cardiac insults [35] and reported to mediate thyroid morphology and function [36].

6 Conclusion

Our work presents a novel adaptation of the CBR paradigm for drug repositioning tasks. While graph computational drug repositioning approaches continue evolving, the interpretable explanation of predictions made continues mostly unsolved. Here, we demonstrated that reasoning chains provided by our implementation of CBR-SUBG approach can provide important domain insights to better assess whether novel drug-disease candidates are plausible.

From the results, we demonstrated CBR-SUBG outperforms various GNN-based and non-parametric CBR models on drug repositioning tasks. While the non-parametric CBR model has the ability to provide reasoning chains that explain drug-disease predictions, it is limited by the exact symbolic matching done to find an answer. Next, compared to GNN-based models, CBR-SUBG only considers important features that are reused from similar solutions. As described, this is done by leveraging the graph structure patterns from similar drug nodes and reasoning over a subgraph query. In other words, instead of considering all the neighbor nodes around a given query during message-passing scheme, CBR-SUBG only considers relevant features to make an inference.

Acknowledgments The authors would like to thank Dr Rajarshi Das and Dr Chunlei Wu for support and feedback.

References

[1] Fabio Pammolli, Laura Magazzini, and Massimo Riccaboni. The productivity crisis in pharmaceutical r&d. *Nature reviews Drug discovery*, 10(6):428–438, 2011.

[2] Sudeep Pushpakom, Francesco Iorio, Patrick A Eyers, K Jane Escott, Shirley Hopper, Andrew Wells, Andrew Doig, Tim Guilliams, Joanna Latimer, Christine McNamee, et al. Drug repurposing: progress, challenges and recommendations. *Nature reviews Drug discovery*, 18(1):41–58, 2019.

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

[4] Michael Mayers, Roger Tu, Dylan Steinecke, Tong Shu Li, Núria Queralt-Rosinach, and Andrew I Su. Design and application of a knowledge network for automatic prioritization of drug mechanisms. *Bioinformatics*, 38(10):2880–2891, 2022.

[5] Payal Chandak, Kexin Huang, and Marinka Zitnik. Building a knowledge graph to enable precision medicine. *Scientific Data*, 10(1):67, 2023.

[6] Adriana Carolina Gonzalez-Cavazos, Anna Tanska, Michael Mayers, Denise Carvalho-Silva, Brindha Sridharan, Patrick A Rewers, Umasri Sankarlal, Lakshmanan Jagannathan, and Andrew I Su. Drugmechdb: A curated database of drug mechanisms. *Scientific Data*, 10(1):632, 2023.

[7] David N Nicholson and Casey S Greene. Constructing knowledge graphs and their biomedical applications. *Computational and structural biotechnology journal*, 18:1414–1428, 2020.

[8] 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(24):5191–5198, 2019.

[9] Ping Xuan, Yilin Ye, Tiangang Zhang, Lianfeng Zhao, and Chang Sun. Convolutional neural network and bidirectional long short-term memory-based method for predicting drug–disease associations. *Cells*, 8(7):705, 2019.

[10] Mustafa Coşkun and Mehmet Koyutürk. Node similarity-based graph convolution for link prediction in biological networks. *Bioinformatics*, 37(23):4501–4508, 2021.

[11] Haitao Fu, Feng Huang, Xuan Liu, Yang Qiu, and Wen Zhang. Mvgcn: data integration through multi-view graph convolutional network for predicting links in biomedical bipartite networks. *Bioinformatics*, 38(2):426–434, 2022.

[12] Agnar Aamodt and Enric Plaza. Case-based reasoning: Foundational issues, methodological variations, and system approaches. *AI communications*, 7(1):39–59, 1994.

[13] Beatriz López. *Case-based reasoning: a concise introduction*. Springer Nature, 2022.

[14] Rajarshi Das, Manzil Zaheer, Dung Thai, Ameya Godbole, Ethan Perez, Jay-Yoon Lee, Lizhen Tan, Lazaros Polymenakos, and Andrew McCallum. Case-based reasoning for natural language queries over knowledge bases. *arXiv preprint arXiv:2104.08762*, 2021.

[15] Rajarshi Das, Ameya Godbole, Ankita Naik, Elliot Tower, Manzil Zaheer, Hannaneh Hajishirzi, Robin Jia, and Andrew McCallum. Knowledge base question answering by case-based reasoning over subgraphs. In *International conference on machine learning*, pages 4777–4793. PMLR, 2022.

[16] Rajarshi Das, Ameya Godbole, Shehzaad Dhuliawala, Manzil Zaheer, and Andrew McCallum. A simple approach to case-based reasoning in knowledge bases. *arXiv preprint arXiv:2006.14198*, 2020.

[17] Rajarshi Das, Ameya Godbole, Nicholas Monath, Manzil Zaheer, and Andrew McCallum. Probabilistic case-based reasoning for open-world knowledge graph completion. *arXiv preprint arXiv:2010.03548*, 2020.

[18] Michael Schlichtkrull, Thomas N Kipf, Peter Bloem, Rianne Van Den Berg, Ivan Titov, and Max 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*, pages 593–607. Springer, 2018.

[19] Shikhar Vashishth, Soumya Sanyal, Vikram Nitin, and Partha Talukdar. Composition-based multi-relational graph convolutional networks. *arXiv preprint arXiv:1911.03082*, 2019.

[20] Liwei Cai and William Yang Wang. Kbgan: Adversarial learning for knowledge graph embeddings. *arXiv preprint arXiv:1711.04071*, 2017.

[21] Tim Dettmers, Pasquale Minervini, Pontus Stenetorp, and Sebastian Riedel. Convolutional 2d knowledge graph embeddings. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32, 2018.

[22] Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. The graph neural network model. *IEEE transactions on neural networks*, 20(1):61–80, 2008.

[23] Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. Neural message passing for quantum chemistry. In *International conference on machine learning*, pages 1263–1272. PMLR, 2017.

[24] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020.

[25] Roger Tu, Meghamala Sinha, Adriana Carolina Gonzalez-Cavazos, Eric Hu, Shehzaad Zuzar Dhuliawala, Andrew McCallum, and Andrew I Su. Drug repositioning using consilience of knowledge graph completion methods. *bioRxiv*, pages 2023–05, 2023.

[26] Oleg Ursu, Jayme Holmes, Jeffrey Knockel, Cristian G Bologa, Jeremy J Yang, Stephen L Mathias, Stuart J Nelson, and Tudor I Oprea. Drugcentral: online drug compendium. *Nucleic acids research*, page gkw993, 2016.

[27] Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. Optuna: A next-generation hyperparameter optimization framework. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, pages 2623–2631, 2019.

[28] Mahsa Shahrokhi and Vikas Gupta. Propranolol. 2020.

[29] Pramod Kumar Pal. Guidelines for management of essential tremor. *Annals of Indian Academy of Neurology*, 14(Suppl1):S25, 2011.

[30] Niall J Crosby, Katherine Deane, Carl E Clarke, and Cochrane Movement Disorders Group. Beta-blocker therapy for tremor in parkinson's disease. *Cochrane Database of Systematic Reviews*, 2010(1), 1996.

[31] Amrutur Venkatachar Srinivasan. Propranolol: A 50-year historical perspective. *Annals of Indian academy of neurology*, 22(1):21, 2019.

[32] Stephen SG Ferguson and Ross D Feldman. β-adrenoceptors as molecular targets in the treatment of hypertension. *Canadian Journal of Cardiology*, 30(5):S3–S8, 2014.

[33] Shuchi Mittal, Kjetil Bjørnevik, Doo Soon Im, Adrian Flierl, Xianjun Dong, Joseph J Locascio, Kristine M Abo, Elizabeth Long, Ming Jin, Bing Xu, et al. β2-adrenoreceptor is a regulator of the α-synuclein gene driving risk of parkinson's disease. *Science*, 357(6354):891–898, 2017.

[34] Andrew Steptoe, Amy Ronaldson, Karen Kostich, Antonio I Lazzarino, Livia Urbanova, and Livia A Carvalho. The effect of beta-adrenergic blockade on inflammatory and cardiovascular responses to acute mental stress. *Brain, behavior, and immunity*, 70:369–375, 2018.

[35] Bin Dai, Meng Cui, Mei Zhu, Wen-Ling Su, Ming-Cai Qiu, and Hong Zhang. Stat1/3 and erk1/2 synergistically regulate cardiac fibrosis induced by high glucose. *Cellular Physiology and Biochemistry*, 32(4):960–971, 2013.

[36] Hiroaki J Kimura, Roberto Rocchi, Melissa A Landek-Salgado, Koichi Suzuki, Cindy Y Chen, Miho Kimura, Noel R Rose, and Patrizio Caturegli. Influence of signal transducer and activator of transcription-1 signaling on thyroid morphology and function. *Endocrinology*, 150(7):3409–3416, 2009.