PAPER

Drug Repositioning using Consilience of Knowledge Graph Completion Methods

Roger Tu,[1,†] Meghamala Sinha,[1,†] Carolina González,[1] Eric Hu,[1]
Shehzaad Dhuliawala,[2] Andrew McCallum[2] and Andrew I. Su[1,*]

[1]Department of Integrative Structural and Computational Biology, The Scripps Research Institute, 10550 North Torrey Pines Road, 92037, CA, USA, [2]College of Information of Computer Sciences, University of Massachusetts, Amherst, 140 Governors Drive, 01003, MA, USA and [†]Denotes equal contribution.

[*]Corresponding author. asu@scripps.edu

Abstract

Motivation: While link prediction methods in knowledge graphs have been increasingly utilized to locate potential associations between compounds and diseases, they suffer from lack of sufficient evidence to explain why a drug and a disease may be indicated. This is especially true for knowledge graph embedding (KGE) based methods where a drug-disease indication is linked only by information gleaned from a vector representation. Complementary pathwalking algorithms can increase the confidence of drug repositioning candidates by traversing a knowledge graph. However, these methods heavily weigh the relatedness of drugs, through their targets, pharmacology or shared diseases. Furthermore, these methods rely on arbitrarily extracted paths as evidence of a compound to disease indication and lack the ability to make predictions on rare diseases.

Results: In this paper, we evaluate seven link prediction methods on a vast biomedical knowledge graph for drug repositioning. We follow the principle of consilience, and combine the reasoning paths and predictions provided by path-based and KGE methods to not only demonstrate a significant ranking performance improvement but also identify putative drug repositioning indications. Finally, we highlight the utility of our approach through a potential repositioning indication.

Availability: The MIND dataset can be found at 10.5281/zenodo.8117748. The python code to reproduce the entirety of this analysis can be found at https://github.com/SuLab/{KnowledgeGraphEmbedding, CBRonMRN}.

Contact: Andrew I. Su at asu@scripps.edu

Supplementary information: Supplementary data are available at *The Journal Title* online.

Key words: drug repositioning, biological networks, computational biology, heterogenous networks, knowledge graph completion, case-based reasoning, knowledge graph embedding

Introduction

The estimated cost to bring a drug to market increased from $802 million to $2.7 billion between 2003 and 2013 due to drug attrition, longer development timelines and changing regulatory requirements, adjusting for inflation (1). Drug repositioning, the process of identifying a new indication for an advanced clinical compound or approved drug, has become increasingly more attractive by leveraging prior work characterizing a drug candidate's safety and efficacy profile; resulting in a concomitant decrease in time to market, risk of failure, and investment costs (2).

An emerging area in computational drug repositioning exploit features in knowledge graphs, networks consisting of a set of entities including but not limited to "drugs", "genes" and "diseases" bound together by relationships like "associated with" or "treats", to identify new connections between a drug and a disease (3). Common drug repositioning method have identified prospective candidates through exploiting drug-drug and disease-disease similarities in a knowledge graph (4; 5; 6). Alternative approaches have utilized graph traversal algorithms like shortest path and random walks, or path ranking algorithms like degree weighted path count to prioritize paths linking a drug and a disease (5; 7; 8; 9; 10; 11). More specific graph traversal methods have applied deep learning methods by formalizing drug repositioning as a reinforcement learning task for link prediction (12; 13).

Link prediction methods like Knowledge Graph Embedding (KGE) and Case based reasoning (CBR) have enjoyed success in identifying relationships in large semantic based knowledge graphs and can be utilized to establish links between a drug and a disease as drug repositioning candidates. KGE algorithms

learn a vector representation of each entity and relation in n-dimensional space for link prediction in a knowledge graph, and represent a class of novel methods for computational drug repositioning (14; 15; 16; 17; 18). As knowledge graphs by nature are inherently incomplete, KGE are a powerful method for drug repositioning because they can extrapolate relationships between compounds and diseases in embedding space without the restriction of traversing a knowledge graph. This is particularly useful for identifying drug repositioning candidates in rare diseases as they are not well characterized.

Case based reasoning (CBR) is an established technique from artificial intelligence modeled after human ability to retrieve and apply prior experience to tackling a new but similar challenge, and represents another promising approach to conduct drug repositioning (19; 20; 21; 22; 23). While not specifically applied to drug repositioning, (24) demonstrated the utility and effectiveness of this simple approach for link prediction on knowledge graphs. Applied to drug repositioning, pathways from other similar compounds can be leveraged as evidence to identify a potential disease treatment given a compound of interest. This is beneficial as the analogy to other drugs provides mechanistic insight for the application of a compound to a disease.

In this paper, we apply the concept of consilience, that evidence from independent provenience can converge on a conclusion, to drug repositioning through the use of various knowledge graph completion approaches. We highlight the performance of seven KGE and path-based reasoning methods on a large biomedical knowledge graph using approved drug-disease indications, and conduct a thorough analysis of our combined model prediction results against each algorithm evaluated. Finally, we demonstrate the effectiveness of our approach through manual curation and identifying plausible drug repositioning indications by combining the results of both complementary link prediction methods.

Methods

A Mechanistic Repurposing Network with Indications (MIND) Knowledge Graph

Mechanistic Repositioning Network with Indications (MIND)is a knowledge graph that distinguishes approved drug indications from semantically derived drug-disease relationships. Based on MechRepoNet, a knowledge graph that reflects important drug mechanism relationships identified from a curated biomedical drug mechanism dataset, MIND elevates pre-existing DrugCentral (a curated resource with regulatory approved drug indications) obtained *treat* edges as *indication* edges (25; 26). The *treat* edge represents a weaker link between a drug and disease compared to the *indication* edge as *treat* edges are not substantiated by regulatory approval. In MIND, when both *indication* and *treat* edges exist, *indication* superseded and replaced *treat* edges. In total, MIND consists of 9,652,116 edges, 249,605 nodes, 9 node types and 22 relations. Supplementary Figure 1 (upper) highlights total node to node and (lower) node to relation counts in MIND as a whole.

Link Prediction Algorithms for Drug Repositioning

In this paper, we utilized and evaluated a variety of algorithms for drug repositioning on MIND. These algorithms fall in two classes: knowledge graph embeddings, and path reasoning methods.

Knowledge Graph Embeddings (KGE)

Knowledge graph embedding algorithms model missing graph links by defining a scoring function for each triple. A knowledge graph is made up of a collection of triples (head, relation, tail), where the set of all entities is represented by $h, t \in \mathcal{E}$, and the set of all graph relations is represented by $r \in \mathcal{R}$. Each entity is represented by a vector embedding, and is modeled by a score function $f_r(h, t)$. The goal of knowledge graph embedding algorithms is to score true triplets (h, r, t) higher than corrupted (untrue) triplets (h', r, t) or (h, r, t'). In this paper, we trained knowledge graph embedding models on MIND with the following algorithms: TransE , DistMult , ComplEx and RotatE (17; 27; 28; 29). More information regarding the algorithms and their respective scoring functions are provided in Supplementary Table 1.

Path Reasoning Methods

Path reasoning methods leverage and traverse knowledge graph edges to identify potential drug repositioning candidates. Here we describe two path based reasoning methods we evaluated for drug repositioning: Degree Weighted Path Count (Rephetio) and Case based reasoning.

Degree Weighted Path Count, an algorithm from the Rephetio project, penalizes paths traveling through high-degree nodes when calculating metapath (path based on node type) prevalence; metapaths are incorporated into a logistic regression to calculate the expected probability a compound treats a disease (10). Mayers et. al. expanded on this approach and incorporated rules based path exclusions and hyperparameter optimization schemes to improve path interpretability for drug repositioning (25).

Case based reasoning (CBR), a problem solving approach analogous to how a doctor prescribes a treatment by relying upon their prior experiences, applies similar solutions from compounds most like a given drug to identify a putative disease treatment. CBR models first retrieve similar entities to the query entity, h_q, that have the specified query relation, r_q given a query (h_q, r_q). Next, the set of relations connecting the similar entities to their answer via r_q are collected and the paths obtained are applied to the query entity h_q to identify the answer (24). Probabilistic case based reasoning (pCBR) extends the original case based reasoning method by using k-nearest neighbors to improve similar entity recognition and by utilizing probabilistic models to estimate the likelihood a retrieved path is correct given the query relation (30).

In this paper, we applied KGE algorithms (TransE, DistMult, ComplEx and RotatE) and path reasoning methods (Mayers *et al.* implementation of Rephetio, CBR and pCBR) to identify putative drug repositioning candidates and their corresponding path reasoning evidence.

Drug Repositioning with Consilience

Consilience, the principle that evidence from unrelated sources converge on strong conclusions, was applied and evaluated by testing the ability of seven knowledge graph completion methods to correctly prioritize approved drug-disease indications over unknown drug-disease links.

Hyperparameter Optimization

Hyperparameters were tuned over MIND train and validation splits and evaluated on the test split utilizing a compound-to-disease indication train/test/validation split of 80/10/10%, respectively; Algorithms were evaluated using hits at k, the

Fig. 1. Consilience inspired Drug Repositioning Schematic. (1) Each knowledge graph completion model is trained on a biomedical knowledge graph and (2) evaluated on approved drug indications (test set). (3) Given each model's putative predictions, (4) the common answer ranks are collected, averaged and (5) resorted to identify drug repositioning candidates.

Table 1. Drug to disease indication validation performance results for seven algorithms. Intersection (7) represents MRR and hits at k calculated with seven algorithms and their respective policies.

Algorithm	MRR	Hits@1	Hits@3	Hits@10
Intersection (7)	**0.9540**	**0.9080**	**1.0000**	**1.0000**
CBR	0.0424	0.0093	0.0480	0.0980
probCBR	0.1814	0.1170	0.1863	0.3064
Rephetio	0.1229	0.1403	0.2055	0.3439
TransE	0.1765	0.0685	0.1996	0.4168
RotatE	0.1682	0.0920	0.1898	0.3229
DistMult	0.0742	0.0176	0.0744	0.1761
ComplEx	0.0904	0.0274	0.0841	0.2329

success rate of identifying the correct answer among the top k predictions, and mean reciprocal rank (MRR), the average of the reciprocal of the positional rank for all correct answers. KGE and CBR hyperparameters were optimized using Optuna software (31). Rephetio was optimized independently through its own hyperparameter optimization pipeline. Hyperparameter optimization selected parameter values can be found at Supplementary Tables 2 and 3, respectively.

Consilience of Seven Algorithm Predictions

The top 100 rank predictions for each drug-disease indication in the test set for each computational method were aggregated into algorithm combinations of sizes two to seven. The disease ranks from each drug prediction by each algorithm in an algorithm combination are then aggregated following an *intersection* operation. Under an *intersection* policy, only diseases reported by all algorithms are retained.

Following aggregation, the predicted disease ranks or reciprocal ranks for each drug in an algorithm combination are averaged and the set of 100 diseases are re-ranked from best to worst. Statistical testing was conducted on the rank distributions of approved vs predicted indications grouped by algorithm combination lengths using the Kruskal-Wallace H-test; a non-parametric method of comparing medians of groups. The Mann-Whitney U-test was applied to observe differences between the distributions of approved and unknown drug to disease indications.

Evaluating Plausibility of Predicted Indications

In order to evaluate the plausibility of drug repositioning candidates, literature curation was conducted for each drug and its top disease candidate (first re-ranked position) from the test set. Every putative indication was categorized into three categories: positive, negative or neutral effect. An indication prediction was labeled "positive" if literature review suggested a drug treated or improved the disease condition; these predicted edges have the potential to become drug indications and warrant further study or are already used clinically off-label. "Negative" labels are predicted indications that might induce, negatively impact or exacerbate the effect of the drug

on the disease. Finally, "neutral" labeled edges have no literature support for the given drug and disease prediction, or have been found to neither treat nor modify disease outcomes. The curation results can be found in the supplementary file.

Results

In this section, we applied several knowledge graph completion algorithms on the MIND dataset to identify drug repositioning candidates using consilience. First, we trained and evaluated the performance of each knowledge graph completion algorithm independently on DrugCentral indications in MIND and compared against our consilience approach. Next, we studied the *intersection* policy's ability to rank true indications and unknown edge predictions. Following, we investigated the plausibility of our putative top drug repositioning candidates indications by gathering supporting evidence through manual curation in literature. Finally, we explored the mechanism of action of a potential drug repositioning indication.

Knowledge Graph Indication Prediction Performance Comparison

Each algorithm's prediction efficacy on approved indications was evaluated by training and testing on MIND. We applied and evaluated the consilience principle on approved indications following the aforementioned intersection operation. Table 1 highlights each algorithm's prediction performance. Among the path traversal methods, we observed that probCBR performed the best with a MRR of 0.1814. Meanwhile, the highest performing embedding-based approach, TransE, performed slightly worse than probCBR with a MRR of 0.1765. We observed the intersection policy performed the best overall for this task, outshining individual algorithm predictions with an MRR of 0.9540.

Consilience Policy Effect on Predicted Putative Indications

After comparing our consilience approach against individual algorithms' ability to identify approved indications, we investigated the effects of consilience policy on indication ranking performance and on potential drug repositioning candidates counts. Figure 2 illustrates the MRR prediction performance

Fig. 2. Approved indication prediction MRR and indication counts with respect to algorithm length used for consilience. Solid line denotes average MRR for algorithm combinations of lengths two to seven. Total potential prediction counts are demarcated by dashed lines and its y-axis is log scaled.

on approved indications and candidate prediction counts with respect to the number of algorithms considered for consilience.

Following an intersection policy (solid line), a concomitant increase in MRR is observed with increased algorithm combination length. This approach increases MRR performance from 0.2812 to 0.9540 when consilience length increases from two to seven, respectively. Regarding putative indication counts (dashed line), the intersection policy's prospective indications drop precipitously as algorithm combination length increases. Despite the significant improvement in MRR performance under an intersection policy, only 87 indications remain when algorithm combination length is seven.

Next we visually and statistically explore the consilience inspired approach on ranking approved indications and prospective indications across varying algorithm combinations to study how consilience policy affects ranking distributions. Figure 3 demonstrates the approved indication (Indication) and prospective indication (Not Indication) rank distributions for intersection policy across varying algorithm combination lengths.

Subject to an intersection policy, Figure 3 highlights strong ranking performance among both "Indication" and "Not Indication", with approved indications perceptibly ranking better than putative drug repositioning candidates. The median indication rank was 1 with 87 counts, whereas the median predicted candidate (Not Indication) rank was 2 with 68 counts at a combination length of seven. Applying the Kruskal-Wallis H-test showed at least one distribution median was statistically different from the others; subsequent statistical testing using the Mann-Whitney U-test demonstrated each pair of distributions for each algorithm combination were statistically significant ($p<0.001$). Supplementary Tables 4 and 5 show the Kruskal-Wallis H-test and Mann-Whitney U-test statistics and p-value, respectively.

As a whole, these results suggest that intersection policy significantly improves the ranking performance on approved indications but decreases the number of potential drug repositioning candidates.

Intersection Putative Indications for Annotation

As a rough estimate of prediction performance, we investigate the plausibility of putative predictions for each top predicted drug repositioning candidate with a manual literature search. The effect of each drug on its predicted potential indication

Fig. 3. Indication and non-indication disease rank distribution after reranking. X-axis represents algorithm combinations between CBR, pCBR, Rephetio, TransE, DistMult, ComplEx and RotatE of sizes two through seven. Boxplot distributions highlight the first, second and third quartiles for indication and non-indication disease ranks. The distribution for indication and non-indication sets are statistically significant regardless of algorithm combination length. Violin plots highlight the intersection policy indication and non-indication disease rank distribution; the y-axis is log scaled.

is categorized as a positive, negative or neutral effect. Observing prior defined criteria, we found that of the 87 intersecting potential indications, 25 were ranked in first place. Of the 25, 80% (20) were categorized as positive effects, 20% (5) as neutral effects, and 0 as negative effects. In contrast, 64% (16), 40% (8) and 4% (1) of the predictions from the best validation set performing algorithm, probCBR, were categorized as positive, neutral, or negative effects, respectively. The best performing knowledge graph embedding method, TransE, fared worse with 4%, 96% categorized as positive, and neutral effects, respectively. Through literature curation, we demonstrated our consilience inspired approach outperformed individual algorithm's in predicting prospective drug repositioning indications. Curation summary is illustrated in Table 2. Full curation results can be found in the Supplementary File.

Case Study: Sotalol hydrochloride as a potential treatment for hypertension

Among the most confident predictions made by our consilience inspired approach was the use of sotalol hydrochloride (sotalol) to treat hypertension. This predicted indication was made by all seven algorithms at ranks ranging from 1 to 88 as seen in Supplementary Table 6. Sotalol, an atypical beta blocker, is approved to treat arrhythmias like atrial fibrillation and ventricular tachycardia. While sotalol has not been approved for hypertension, various clinical studies have demonstrated its efficacy in controlling hypertension independently and as an adjunctive therapy with thiazides (32; 33; 34). Utilizing our computational repositioning approach, we propose three mechanisms by which sotalol may moderate hypertension (a condition characterized by increased arterial blood pressure) through potassium channel activity, ADRB1 and FNDC4. Each prospective mechanism retrieved were the top ranked paths by each path reasoning method and is illustrated in Figure 4.

A canonical avenue sotalol may manage hypertension is through inhibition of ADRB1, a β-adrenergic receptors. ADRB1 inhibition mediates adrenaline's effect on the heart,

Table 2. Drug to disease indication validation performance results for seven algorithms. Intersection (7) represents MRR and hits at k calculated with seven algorithms and their respective policies.

	CBR	pCBR	Rephetio	TransE	DistMult	ComplEx	RotatE	Intersection (7)
Positive	44% (11)	64% (16)	44% (11)	4% (1)	60% (15)	68% (17)	64% (16)	80% (20)
Neutral	52% (13)	40% (8)	40% (10)	96% (24)	40% (10)	28% (7)	24% (6)	20% (5)
Negative	4% (1)	4% (1)	16% (4)	-	-	4% (1)	12% (3)	-

Fig. 4. Three prospective mechanisms that sotalol hydrochloride may treat hypertension. Each path shown was the first extracted path suggested by pCBR, Rephetio and CBR for the putative indication. Green, purple, pink and blue circles represent compounds, biological processes, genes & proteins, and diseases, respectively. Orange arrows in between nodes describe node relationships.

resulting in concomitant decreased blood pressure and cardiac output (35). As sotalol and amosulalol both directly block ADRB1, a protein and gene associated with hypertension, and amosulalol has anti-hypertensive effects, it is reasonable that sotalol also has anti-hypertensive effects (36; 37).

Another approach sotalol may modulate hypertension is through potassium channel activity (38). While potassium channel blockers are associated with exacerbating blood pressure as it prevents the outflow of potassium ions, it is plausible that sotalol's inhibition of inward-rectifier potassium channels (repolarization) induces vasodilation and as a result, decreases blood pressure (39; 40; 41).

Finally, sotalol may manage hypertension through upregulating the expression of FNDC4, an anti-inflammatory factor (42). FNDC4 belongs to the fibronectin type III domain-containing protein family, and is highly homologous to irisin, a myokine derived from the proteolytic cleavage of FNDC5 (43; 44). Deletion and overexpression of irisin has been shown to exacerbate and ameliorate cardiac hypertrophy in rats with hypertension, respectively (45). As sotalol and pirinixic acid treatment both upregulate irisin expression, and pirinixic acid has hypertension mediating effects, it is probable that sotalol also exhibits anti-hypertensive properties (46; 47).

Discussion

In this paper, we apply a consilience based approach on seven algorithms and demonstrate an enhanced ability to predict approved and putative indications. By combining path-based and KGE algorithms, our method synergistically bolsters path-based and KGE inspired drug repositioning performance through increased resilience to missing edges and reinforced prospective drug repositioning indications with reasoning chains. Our strategy not only validates KGE derived predictions, but it also allows path-based predictions to prioritize diseases with minimal similar cases in the graph.

Although our method viably identifies potential repositioning candidates, one challenge in our implementation is the restrictive nature of the intersection policy. As the number of algorithms applied for consilience increases, the fewer putative indications remain (as illustrated in Figure 2). This issue can be ostensibly mediated by either increasing the number of total predictions made per indication and/or increasing the number of compounds used for inference. Expanding our study limits for each algorithm's predictions (from 100 to the top 1000 ranked diseases and/or inference compounds from 500 to 1000), barring computational limits, could be considered.

An additional avenue to address the intersection policy's restrictive properties is through a partial-intersection policy; instead of eliminating popular but nonunanimous predictions, the rank is padded with a tunable user specified rank. This approach preserves intersecting prospective indications and penalizes partially-intersected candidates during the re-ranking step. Notably our unreported preliminary study demonstrated a decrease in approved indication performance following a partial-intersection procedure.

Another challenge to our approach is the equal weighting of each algorithm's predictions in its irrespective of prior observed performance on a dataset. Weighting each algorithm's contributions by the predicted indication rank or by incorporating an algorithm's cross validation performance into a logistic regression would potentially improve the prediction efficacy of the algorithm. These modifications would further fuel the addition of diverse algorithms into our implementation.

Finally, our utilization of path-based methods does not mitigate similarity based drug-disease associations, even when the results are filtered by those that also occur in KGE approaches. For example, the pathway identified by our approach in the between sotalol hydrochloride and hypertension, traverses through ADRB1, a commonly shared target by amosulalol and sotalol, in order to treat hypertension. While similarity based paths correctly identify analogous family compounds (sotalol and amosulalol are both beta blockers), adding simple path similarity filters or penalties to the path retrieval mechanism would likely improve both the path retrieval algorithm and our own approach.

Regardless of variation of the consilience strategy applied, our results demonstrate that even with a naive approach, we were able to identify putative drug repositioning candidates that are supported by existing literature. Our strategy synergizes the advantages of both KGE and path-based algorithms by blending their combined predictive power and strengths. Through manual literature curation, we demonstrate that our method is more likely to identify a drug repositioning indication than when utilizing each algorithm alone. Moreover, our method enables human interpretable reasoning chains derived from path-based approaches to support a putative compound to disease indication that would otherwise not be present with KGEs alone.

Acknowledgments

The authors would like to thank Dr. Laura Hughes for visualization feedback, and Dr. Chunlei Wu, Dr. Jean-Christophe Ducom and Dr. Rajarshi Das for computational support. This work was supported by the National Institute of Health (R01 AG066750).

Bibliography

References

1. DiMasi,J.A. et al. (2016) Innovation in the pharmaceutical industry: New estimates of R&D costs. J. Health Econ., 47, 20–33.
2. Ashburn,T.T. and Thor,K.B. (2004) Drug repositioning: identifying and developing new uses for existing drugs. Nat. Rev. Drug Discov., 3, 673–683.
3. Pushpakom,S. et al. (2019) Drug repurposing: progress, challenges and recommendations. Nat. Rev. Drug Discov., 18, 41–58.
4. Li,J. and Lu,Z. (2012) A new method for computational drug repositioning using drug pairwise similarity. In, 2012 IEEE International Conference on Bioinformatics and Biomedicine. IEEE, Philadelphia, PA, USA, pp. 1–4.
5. Luo,H. et al. (2016) Drug repositioning based on comprehensive similarity measures and Bi-Random walk algorithm. Bioinformatics, 32, 2664–2671.
6. Luo,H. et al. (2018) Computational drug repositioning using low-rank matrix approximation and randomized algorithms. Bioinformatics, 34, 1904–1912.
7. Luo,H. et al. (2019) Computational Drug Repositioning with Random Walk on a Heterogeneous Network. IEEE-/ACM Trans. Comput. Biol. Bioinform., 16, 1890–1900.
8. Lee,T. and Yoon,Y. (2018) Drug repositioning using drug-disease vectors based on an integrated network. BMC Bioinformatics, 19, 446.
9. Emig,D. et al. (2013) Drug Target Prediction and Repositioning Using an Integrated Network-Based Approach. PLoS ONE, 8, e60618.
10. Himmelstein,D.S. et al. (2017) Systematic integration of biomedical knowledge prioritizes drugs for repurposing. eLife, 6, e26726.
11. Mayers,M. et al. (2019) Time-resolved evaluation of compound repositioning predictions on a text-mined knowledge network. BMC Bioinformatics, 20, 653.
12. Liu,Y. et al. (2021) Neural Multi-hop Reasoning with Logical Rules on Biomedical Knowledge Graphs. In, Verborgh,R. et al. (eds), The Semantic Web, Lecture Notes in Computer Science. Springer International Publishing, Cham, pp. 375–391.
13. Wang,H. et al. (2022) A heterogeneous network-based method with attentive meta-path extraction for predicting drug–target interactions. Brief. Bioinform., 23, bbac184.
14. Bonner,S. et al. (2022) Understanding the performance of knowledge graph embeddings in drug discovery. Artif. Intell. Life Sci., 2, 100036.
15. Celebi,R. et al. (2019) Evaluation of knowledge graph embedding approaches for drug-drug interaction prediction in realistic settings. BMC Bioinformatics, 20, 726.

16. Mohamed,S.K. et al. (2021) Biological applications of knowledge graph embedding models. Brief. Bioinform., 22, 1679–1693.

17. Bordes,A. et al. (2013) Translating embeddings for modeling multi-relational data. In, Proceedings of the 26th International Conference on Neural Information Processing Systems - Volume 2, NIPS'13. Curran Associates Inc., Red Hook, NY, USA, pp. 2787–2795.

18. Toutanova,K. et al. (2016) Compositional Learning of Embeddings for Relation Paths in Knowledge Base and Text. In, Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers). Association for Computational Linguistics, Berlin, Germany, pp. 1434–1444.

19. Schank,R.C. (1983) Dynamic memory: A theory of reminding and learning in computers and people cambridge university press.

20. Kolodner,J.L. (1983) Maintaining organization in a dynamic long-term memory. Cogn. Sci., 7, 243–280.

21. Rissland,E.L. (1983) Examples in Legal Reasoning: Legal Hypotheticals. In, IJCAI., pp. 90–93.

22. Aamodt,A. and Plaza,E. (1994) Case-based reasoning: Foundational issues, methodological variations, and system approaches. *AI Commun.*, **7**, 39–59.

23. Leake,D.B. (1996) CBR in context: the present and future. Case based reasoning experiences-lessons and future experiences. D. Leake Cambridge, MIT Press.

24. Das,R., Godbole,A., Dhuliawala,S., et al. (2020) A Simple Approach to Case-Based Reasoning in Knowledge Bases. arXiv, abs/2006.14198.

25. Mayers,M. et al. (2022) Design and application of a knowledge network for automatic prioritization of drug mechanisms. Bioinforma. Oxf. Engl., 38, 2880–2891.

26. Ursu,O. et al. (2017) DrugCentral: online drug compendium. Nucleic Acids Res., 45, D932–D939.

27. Yang,B. et al. (2015) Embedding Entities and Relations for Learning and Inference in Knowledge Bases.

28. Trouillon,T. et al. (2016) Complex embeddings for simple link prediction. In, Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48, ICML'16. JMLR.org, New York, NY, USA, pp. 2071–2080.

29. Sun,Z. et al. (2019) ROTATE: KNOWLEDGE GRAPH EMBEDDING BY RELA- TIONAL ROTATION IN COMPLEX SPACE. arXiv, abs/1902.10197, 18.

30. Das,R., Godbole,A., Monath,N., et al. (2020) Probabilistic case-based reasoning for open-world knowledge graph completion. ArXiv Prepr. ArXiv201003548.

31. Akiba,T. et al. (2019) Optuna: A Next-generation Hyperparameter Optimization Framework.

32. Fillastre,J.P. et al. (1979) [Treatment of hypertension with Sotalol (author's transl)]. Sem. Hopitaux Organe Fonde Par Assoc. Enseign. Med. Hopitaux Paris, 55, 1825–1831.

33. Shaw,H.L. (1977) Once daily sotalol in the treatment of hypertension. J. R. Coll. Gen. Pract., 27, 742–745.

34. Jäättelä,A. (1979) The Combination of Sotalol and Hydrochlorothiazide in the Treatment of Hypertension. J. Clin. Pharmacol., 19, 565–570.

35. BERTRIX,L. et al. (1986) Protection against ventricular and atrial fibrillation by sotalol. Cardiovasc. Res., 20, 358–363.

36. Inoue,Y. et al. (1992) Antihypertensive and metabolic effects of long-term treatment with amosulalol in non-insulin dependent diabetics. Curr. Med. Res. Opin., 12, 564–571.

37. Huang,Y. et al. (2017) Downregulation of the $\beta1$ adrenergic receptor in the myocardium results in insensitivity to metoprolol and reduces blood pressure in spontaneously hypertensive rats. Mol. Med. Rep., 15, 703–711.

38. EDVARDSSON,N. et al. (1980) Sotalol-induced delayed ventricular repolarization in man*. Eur. Heart J., 1, 335–343.

39. Thomas,D. et al. (2006) The cardiac hERG/IKr potassium channel as pharmacological target: structure, function, regulation, and clinical applications. Curr. Pharm. Des., 12, 2271–2283.

40. Chen,L. et al. (2016) Cardiac Delayed Rectifier Potassium Channels in Health and Disease. Card. Electrophysiol. Clin., 8, 307–322.

41. Li,C. and Yang,Y. Advancements in the study of inward rectifying potassium channels on vascular cells. Channels, 17, 2237303.

42. Atienzar,F. et al. (2007) Determination of Phospholipidosis Potential Based on Gene Expression Analysis in HepG2 Cells. Toxicol. Sci., 96, 101–114.

43. Bosma,M. et al. (2016) FNDC4 acts as an anti-inflammatory factor on macrophages and improves colitis in mice. Nat. Commun., 7, 11314.

44. Frühbeck,G. et al. (2020) FNDC4, a novel adipokine that reduces lipogenesis and promotes fat browning in human visceral adipocytes. Metabolism, 108, 154261.

45. Li,R.-L. et al. (2018) Irisin alleviates pressure overload-induced cardiac hypertrophy by inducing protective autophagy via mTOR-independent activation of the AMPK-ULK1 pathway. J. Mol. Cell. Cardiol., 121, 242–255.

46. Leibovitz,E. and Schiffrin,E.L. (2007) PPAR Activation: A New Target for the Treatment of Hypertension. J. Cardiovasc. Pharmacol., 50, 120.

47. MA,Y. et al. (2023) Anti-inflammatory Effect of Irisin on LPS-Stimulated Macrophages Through Inhibition of MAPK Pathway. Physiol. Res., 72, 235–249.