

Biomedical knowledge graphs driving new target identification and novel drug discovery

Pan Dou^{1#}, Xiaobo Yang^{1#}, Yuling Jiang¹, Xiaowei Xie¹, Xiaoju Geng¹, He Huang¹, Honglin Li^{1*} and Shiliang Li^{1,2*}

¹ Innovation Center for AI and Drug Discovery, School of Pharmacy, East China Normal University, Shanghai 200062, China

² Department of Pain Management, HuaDong Hospital affiliated to Fudan University, Shanghai 200040, China

[#] Authors contributed equally: Pan Dou, Xiaobo Yang

* Correspondence: hlli@hsc.ecnu.edu.cn (Li H); slli@hsc.ecnu.edu.cn (Li S)

Abstract

Biomedical knowledge graphs (KGs) provide a structured and traceable framework for target-oriented, AI-assisted drug discovery. By linking genes, proteins, diseases, phenotypes, compounds, pathways, assays, structures, adverse events, and clinical evidence, biomedical KGs can partially mitigate evidence fragmentation and support hypothesis generation across molecular, cellular, and clinical levels. This review summarizes how biomedical KGs are constructed, reasoned over, evaluated, and translated into practical drug-discovery tasks. We first outline major data resources, entity and relation extraction, knowledge fusion, quality control, and large language model-assisted KG construction, with emphasis on ontology grounding, evidence tracing and hallucination control. We then compare key KG reasoning paradigms, including symbolic rule-based reasoning, representation learning, path- and subgraph-aware reasoning, graph Transformer models, and KG-LLM hybrid reasoning. Their strengths and limitations are discussed in relation to interpretability, scalability, leakage risk, benchmark bias, and prospective validation. We further highlight the need to incorporate three-dimensional structural evidence, including protein structures, binding pockets, variants, conformational states, and ligand compatibility, so that graph-based relevance can be connected with chemical tractability. Finally, we discuss applications of biomedical KGs in target prioritization, druggability assessment, mechanism interpretation, drug-target interaction prediction, molecular generation, lead optimization, ADMET and safety assessment, drug repurposing, combination therapy, biomarker discovery, precision medicine, and drug-resistance modeling. Current limitations include heterogeneous data quality, incomplete causal evidence, static graph representations, privacy constraints, and insufficient experimental or clinical validation. Overall, biomedical KGs should be viewed as evidence-integration and prioritization engines rather than automatic proof-generating systems for drug discovery, guiding rigorous, iterative, and experimentally grounded translational decisions.

Citation: Dou P, Yang X, Jiang Y, Xie X, Geng X, et al. 2026. Biomedical knowledge graphs driving new target identification and novel drug discovery. *Targetome* 2(3): e026 <https://doi.org/10.48130/targetome-0026-0025>

Introduction

Drug discovery remains slow, costly, and failure-prone, with many programs failing because early target hypotheses are not supported by sufficient genetic, mechanistic, structural, safety, or clinical evidence before entering expensive experimental and clinical stages^[1]. Modern drug discovery increasingly depends on genomics, transcriptomics, proteomics, metabolomics, structural biology, phenotypic screening, literature, electronic health records, and clinical trial data^[2–6]. However, these resources differ in identifiers, evidence standards, assay conditions, temporal coverage, and curation bias. As a result, target-disease associations, druggability, safety, and translational evidence are often difficult to connect in a coherent and auditable manner (Fig. 1a). Conventional machine-learning models can rank candidate targets or compounds, but their predictions often lack explicit evidence paths for mechanism tracing, toxicity interpretation, causal assessment, and experimental follow-up.

Biomedical knowledge graphs (KGs) offer a structured way to partially address this evidence-fragmentation problem by representing biomedical entities and their typed relationships as a computable semantic network (Fig. 1b)^[7]. By linking drugs, genes, proteins, diseases, phenotypes, pathways, structures, assays, adverse events, and clinical evidence, KGs provide explicit context for target prioritization, druggability assessment, mechanism interpretation, drug-target interaction prediction, repurposing, and safety evaluation^[8]. Knowledge graph reasoning (KGR) further supports rule-based, embedding-based, path-based, graph-neural, transformer-based,

and LLM-assisted inference over these relationships^[9]. Nevertheless, KG-derived predictions should be interpreted as evidence-supported hypotheses that require benchmark-aware evaluation and experimental or clinical validation, rather than as direct proof of therapeutic efficacy (Fig. 1c).

This review focuses on how biomedical KGs support target identification and translational decision-making in drug discovery. This review first introduces KG concepts, data resources, construction pipelines, quality control, and LLM-assisted extraction. The analysis then compares major KGR paradigms, quantitative benchmarks, leakage risks, validation strategies, and structure-augmented multi-modal reasoning. Subsequent sections discuss applications in target discovery, druggability assessment, mechanism interpretation, drug-target interaction prediction, molecular design, multi-parameter lead optimization, ADMET and safety assessment, drug repurposing, combination therapy, biomarkers, precision medicine, and dynamic KGs for resistance analysis. Finally, current limitations and future directions are summarized.

Concepts and construction of biomedical knowledge graphs

Core concepts

The term 'knowledge graph' was formally proposed by Google in 2012^[10]. Mathematically, a knowledge graph is a labeled directed




Fig. 1 Biomedical knowledge graphs break down data silos and promote drug discovery. (a) Biomedical data silos. Different data sources are isolated; real-world data, including literature, clinical records, genomic data, and drug-related information are heterogeneous and fragmented, hindering cross-semantic and cross-hierarchical integration. (b) Knowledge graph construction pipeline. Human-machine collaboration is used to extract entities and relationships from multi-source data, followed by integration and quality control to form a biomedical knowledge graph. (c) Major applications of biomedical knowledge graphs. Knowledge graph-based reasoning enables drug target discovery, drug repositioning, and mechanistic interpretation.

graph structure used to model knowledge and relationships between entities^[11].

Biomedical knowledge graphs take entities such as genes, proteins, compounds, diseases, phenotypes, and pathways as nodes, and semantic relations, including targeting, treatment, pathogenesis, and participation, as edges. They integrate molecular-level

information (genes, proteins), mesoscale processes (cellular functions, tissue physiology), and macroscale phenomena (disease phenotypes, clinical manifestations), enabling cross-scale integration and providing a computational basis for understanding complex biological mechanisms and cross-scale reasoning.

Data resources

High-quality biomedical knowledge graphs rely on extensive and reliable multi-source data resources, which focus on different biological scales. This paper provides a practical classification based on the core features and typical applications of each resource, rather than a strict semantic division (Supplementary Table S1).

The molecular layer is the basis of the knowledge framework and provides a large number of core entities. In the case of ChEMBL^[12], such databases contain information on small-molecule chemical structures and target binding activity; UniProt^[13] and Ensembl^[14] are important sources of knowledge about gene function and protein function. The interaction layer mainly describes the interaction relationship at a higher level. As for STRING^[15], it provides the interaction network between proteins. The pathway layer records the signal transmission process, and the representative resources are KEGG^[16] and Reactome^[17]. The phenotypic layer links the molecular changes with the phenotypic characteristics and biological outcomes of the disease, and DisGeNET^[18] integrates the disease-related genes and variation information. The clinical layer involves medical terminology, diagnosis and treatment intervention, and drug use data. DrugBank^[19] collects the core clinical information of drugs and records the molecular mechanisms and clinical decisions.

Based on these data, researchers have constructed many biomedical knowledge graphs, and each has its own emphasis on the breadth and depth of data integration. As exemplified by Pharos^[20], it concentrates on druggable genes. There are also some special KGs for specific scenarios such as drug target interaction^[21], COVID-19^[22], and precision medicine^[23]. To compare the construction objectives, data coverage, and typical application scenarios of these KGs, representative resources are summarized in Supplementary Table S2.

KG construction pipeline

The knowledge graph follows the bottom-up construction process and is divided into three stages: knowledge extraction, knowledge fusion, and quality evaluation (Fig. 1b).

Knowledge extraction

Knowledge extraction mainly includes named entity recognition (NER) and relation extraction (RE). The task of NER is to recognize biomedical entities from text. Early methods highly relied on experts to formulate rules and dictionaries^[24], and their generalization ability was limited. Subsequently, traditional machine learning methods such as conditional random fields (CRFs)^[25] and support vector machines (SVMs)^[26] improved the flexibility and generalization ability of the model^[27]. Deep learning further reduces the burden of artificial feature design. Pre-trained language models such as BERT^[28] and BioBERT^[29] significantly improve the accuracy and generalization ability of NER by better capturing context semantics^[30].

RE is to further extract the semantic relationship between entities after identifying entities, but it is often limited by the lack of annotation data. Therefore, such tasks usually adopt a weak supervision method, use unmarked data or an external knowledge base, and reduce the labeling cost through distant supervision^[31].

The traditional pipeline mode is prone to cascade errors. In recent years, researchers have begun to use the end-to-end deep learning method to jointly extract entities and relationships^[32], so as to better model the complex interaction between the two.

Knowledge fusion

Knowledge fusion mainly solves the differences and conflicts in multi-source data integration, and its core tasks include ontology

matching and entity alignment. Different biomedical databases are usually built independently by different institutions, so there are often semantic differences and conceptual gaps. Ontology matching^[33] aims to correspond to the conceptual systems in different data sources. Taking DrugBank's drug level and ChEMBL's compound level as examples, the difficulty lies in the inconsistency between classification criteria and conceptual granularity. Entity alignment^[34] identifies and merges nodes that actually point to the same object for different names or identifiers of the same biological entity in different databases to achieve entity normalization and ambiguity resolution. Early methods mainly relied on string matching and standard thesaurus mapping, and now the alignment method based on embedding^[35,36] has gradually become the mainstream.

Quality assessment and optimization

In order to make the knowledge graph reliably serve the downstream drug discovery, it is necessary to carry out systematic quality assessment and optimization^[37]. The evaluation mainly focuses on accuracy and consistency. The former determines whether the triplet conforms to the facts, and the latter checks whether there is a logical conflict. At present, there are many automatic verification methods, but automatic denoising can only deal with common errors. Due to the high accuracy requirements of biomedical data, expert manual review is still irreplaceable.

The knowledge graph also has the problem of incompleteness, but the completion of the knowledge graph in drug R&D and application is usually the final stage.

LLM-driven construction

The large language model has strong natural language understanding, generation, and context learning capabilities, so it can be used for the automatic construction of a biomedical knowledge graph^[38] and has obvious advantages over traditional machine learning and early pre-training language model methods^[39–41]. In the process of knowledge extraction, the outstanding advantage of LLM lies in its learning ability of zero-shot and few-shot learning ability. Through the design of prompts or instructions, the model can also complete named entity recognition and relation extraction in the absence of annotation data^[42,43], and can directly convert complex biomedical literature into structured triples, potentially reducing intermediate pipeline errors, but only when the extracted triples are subsequently verified^[44]. In knowledge fusion, the fine-tuned LLM is easier to identify the same standard concept behind different expressions, showing great potential in zero-shot scenarios and long-tail concept alignment tasks^[45,46]. In addition, a large amount of prior biomedical knowledge implied in the model parameters^[47] can also provide assistance in KG refinement, quality control, and link prediction^[38].

However, LLM-based KG construction still faces a measurable hallucination risk. Because LLMs do not naturally obey graph-structural, ontological, or evidence constraints, hallucinated entities or relations may be converted into false biomedical triples and then propagated to downstream drug-target, disease-gene, or drug repurposing models. In a systematic evaluation of 12 BioNLP benchmarks, LLaMA 2 13B under zero-shot prompting generated 506 hallucinated responses on ChemProt, approximately 3% of 16,943 instances, and 102 hallucinated responses on HoC (32%); GPT-4 also remained below PubMedBERT on BC5CDR-chemical NER and below BioBERT on ChemProt and DDI2013 RE^[48].

Two practical safeguards are therefore required. First, pure zero-shot extraction should be avoided when possible because one

task-specific example reduced HoC hallucinations from 102 to 0 and inconsistent responses from 69 to 23^[48]. Second, LLM-generated triples should be treated as candidate evidence rather than curated facts, and should be inserted into biomedical KGs only after schema checking, ontology/entity grounding, evidence tracing to sentences or PMIDs, and expert or automated validation.

Knowledge graph reasoning for target and drug discovery

This review classifies KGR methods into five core paradigms according to evidence representation and inference modes: symbolic reasoning, representation learning, path- and graph-structure-aware reasoning, Transformer-based graph reasoning, and KG-LLM integrated reasoning. These paradigms reflect a gradual shift from explicit rule-based inference to latent representation learning, multi-hop structural reasoning, global graph context modeling, and language-guided hybrid reasoning. Structure-augmented multi-modal KG reasoning is discussed as a cross-cutting evidence layer instead of a standalone paradigm, linking graph topology to protein structures, binding pockets, variants, conformational states, and ligand compatibility. Typical quantitative benchmarks and evaluation criteria are summarized, along with prevalent evaluation flaws and translational validation challenges that affect the application of benchmark results to real-world drug discovery.

Symbolic reasoning

Symbolic reasoning derives new knowledge using explicit first-order logic rules (Fig. 2a). For drug repositioning, a rule can be formulated: if drug A inhibits target T and target T drives disease D progression, drug A may treat D. Such rules yield clear causal chains, valuable in interpretable drug R&D.

Representative algorithms include AMIE+^[49] and AnyBURL^[50]. AMIE+ uses top-down search under the partial completeness assumption, while AnyBURL relies on bottom-up path sampling. Though logically rigorous, both suffer from a rule-space explosion in large graphs.

Differentiable rule learning methods combine symbolic structures with parametric learning. NeuralLP^[51] models rule application as differentiable matrix multiplication with attention; DRUM^[52] uses bidirectional RNNs to share parameters and avoid overconfidence in erroneous rules. These methods improve scalability and robustness to noise.

However, strict rule matching limits generalization. Even with soft matching, performance heavily depends on graph completeness, severely restricting utility in sparse biomedical graphs.

Representation learning

Knowledge graph embedding (KGE) models represent triple semantics *via* vector algebra or geometry, divided into translational, semantic matching, and geometric transformation methods (Fig. 2b).

Translational models, represented by TransE^[35], model relations as vector translations but struggle with complex drug-target associations; later models introduce relation-specific projections to address this. Semantic matching models use bilinear transformation or tensor decomposition: DistMult^[53] performs well on symmetric relations but lacks directionality, while ComplEx^[54] uses complex embeddings for asymmetric relations such as gene regulation. Geometric models, including RotatE^[55], model relations as rotations in complex space, supporting diverse relation patterns.

In biomedicine, entity functions depend on local network topology. Traditional KGEs underutilize high-order structure, leading to the adoption of graph neural networks (GNNs) such as R-GCN^[56], which propagates relation-aware information.

Representation learning performs well in practice, but traditional KGEs fail to generalize to unseen nodes, and GNNs suffer from over-smoothing that erases fine-grained biological differences.

Path and structural reasoning

Path and structure reasoning is centered on the multi-hop connectivity pattern between entities and mainly includes explicit path features, reinforcement learning path search, subgraph inductive reasoning, and neural path reasoning. The core criterion distinguishing this class of methods from representation learning methods is whether they actively leverage graph structural information (Fig. 2c). Early methods directly enumerate multi-hop paths or local subgraphs and convert them into discrete features for link prediction^[57], which is intuitive and interpretable, but is prone to a combination explosion in large-scale sparse graphs. Reinforcement learning path reasoning models link prediction as a sequential decision problem^[58]. However, due to the graph sparsity and long-tail features, it often faces the problems of sparse rewards and a too-large search space.

Subgraph reasoning extracts the enclosing subgraph around the target entity pair and completes the prediction with the help of relationship-aware GNN^[59,60], which can be generalized to unseen entities, but the computational cost is high. Neural path reasoning implicitly models multi-hop dependency through differentiable message passing^[61,62], taking into account structural expression, representation learning, and strong interpretability, which is more suitable for biological hypothesis generation. Overall, this direction is moving from explicit enumeration to neuro-implicit modeling.

Transformer-based graph reasoning

Many biological processes have long-distance and cross-level connections, which are not easy to capture in terms of remote gene regulation, protein allosteric effects, and cross-pathway coupling. When the layers of the graph neural network are deepened, the problems of over-smoothing and over-compression often occur, so the modeling ability of remote biological association is limited. In contrast, transformers can directly establish the connection between any node by relying on the self-attention mechanism, so it has a better chance of breaking through the restriction of local message delivery (Fig. 2d).

However, the standard transformer is not naturally suitable for graph data. Therefore, researchers have proposed a variety of improved methods with structural bias. As exemplified by Graphormer, this model explicitly encodes node centrality and shortest path distances and incorporates structural information into the attention mechanism^[63]. GPS and other frameworks combine local graph convolution and global transformer to achieve a balance between local structure modeling and long-range dependency capture^[64]. Recent studies have begun to explore the state space model (SSM), represented by Graph-Mamba, which reduces the computational overhead with linear complexity and retains strong long-range dependency modeling capability^[65].

LLM-KG integrated reasoning

The core goal of the combination of LLM and KG is to connect the ability of structural constraints with the ability of semantic understanding^[38]. The earliest research mainly transformed




Fig. 2 Mainstream knowledge graph reasoning paradigms for target and drug discovery. (a) Symbolic reasoning. Performs inference via random walk sampling and first-order logic rules with high interpretability. (b) Representation learning. Maps entities and relations to low-dimensional vector spaces via KGE or GNNs and scores triples. (c) Path and structural reasoning. Infers links by mining enclosing subgraphs or multi-hop path patterns. (d) Transformer-based graph reasoning. Integrates node features and structural encoding into self-attention to capture long-range dependencies. (e) LLM-KG integrated reasoning. Retrieves structured subgraphs as context to guide LLMs for traceable logical inference. Structure-derived features such as pockets, residues, conformational states, variants, and ligand poses can serve as multimodal evidence layers across these paradigms.

structured knowledge into text^[66], but the expression ability of complex multi-hop structures is limited. Subsequent studies mapped graph topology to semantic space through structure coding^[67], but the intermediate reasoning process was still not transparent enough.

A graph retrieval enhancement model is an important development in this direction. It relies on a knowledge graph to build a structured context and uses factual evidence to constrain language model generation^[68] (Fig. 2e). In addition, the multi-agent collaboration framework can also alleviate the instability of single-model reasoning. It reduces false associations and semantic drift through interactive debate, adaptive graph search, self-reflection, and reverse reasoning^[69]. Overall, such LLM + KG frameworks have gone beyond the task of traditional knowledge graph completion and are more suitable for complex inductive reasoning in drug discovery and precision medicine. These five paradigms mainly model graph topology, semantic relations, and evidence paths. For drug discovery, however, target relevance is not enough: tractability also depends on three-dimensional structure, pocket accessibility, conformational dynamics, variant effects, and ligand compatibility. Structure-augmented KG reasoning should therefore be viewed as a cross-cutting multimodal evidence layer rather than a sixth independent paradigm.

Multimodal structure-augmented KG reasoning

Structure-augmented reasoning is not a separate paradigm, but a multimodal evidence layer shared by the paradigms above: rules, embeddings, subgraph models, graph transformers, and KG-LLM systems. Disease-gene-drug paths suggest relevance; structure indicates whether the target has an accessible, ligandable, variant-sensitive, and stable pocket.

Sources include AlphaFold2^[70], RoseTTAFold^[71], ESMFold^[72], and AlphaFold3^[73] for protein or complex models, plus Fpocket^[74],

SiteMap^[75], and MDPocket^[76] for pocket geometry, druggability, and dynamics. Because cryptic allosteric pockets may be transient^[77], annotations should keep source, conformational state, and confidence.

A structure-aware KG should add structural nodes (structures, residues, variants, pockets, conformational states, ligands, conformers), edges (protein-has-structure, structure-has-pocket, variant-perturbs-pocket, pocket-binds-ligand, ligand-compatible-with-pocket), and attributes (pLDDT/PAE, pocket descriptors, MD occupancy, docking pose, assay/literature evidence). This enables disease-gene-protein-pocket-ligand paths.

Existing work shows feasibility: MKGE adds SMILES text and molecular graphs to KG embeddings^[78], and BioBLP combines protein sequences, drug SMILES, disease text, and missing attributes^[79]. Extending these designs to pockets, residues, and conformational states would connect KG reasoning with structure-based drug discovery.

Without structure, predictions may be biologically plausible but chemically weak. Topology-only models can over-rank targets with buried or ligand-incompatible pockets and miss targets enabled by cryptic pockets^[80]. Graph-only links remain hypotheses; links also supported by pocket geometry, dynamics, ligand compatibility, and experiments have stronger mechanistic credibility.

Benchmarking and comparative interpretation of KGR paradigms

After defining the major paradigms and the need for structure-aware evidence, subsequent comparisons focus on the evaluation of KGR methods. Table 1 provides a comprehensive overview of three link prediction benchmarks: FB15k-237, Hetionet, and ogbl-ppa, as they offer the most widely reported metrics across various paradigms (more benchmarks can be obtained from Open Graph

Table 1. Representative quantitative benchmarks for knowledge graph reasoning paradigms.

Dataset/benchmark	Model	Paradigm	MRR	Hits@1	Hits@10	Hits@100
FB15k-237 ^[82] : general multi-relational KGC benchmark	Neural-LP ^[51]	Symbolic reasoning	0.237	0.173	0.361	NR
	TransE ^[35]	Representation learning	0.312	0.212	0.51	NR
	RotatE ^[55]		0.338	0.241	0.533	NR
	TuckER ^[83]		0.358	0.266	0.544	NR
	CompGCN ^[84]	Path and structure-aware reasoning	0.355	0.264	0.535	NR
	NBFNet ^[61]		0.415	0.321	0.599	NR
	HittER ^[85]	Transformer-based graph reasoning	0.373	0.279	0.558	NR
	KnowFormer ^[86]		0.43	0.343	0.608	NR
	SimKGC ^[87]	Hybrid KG-LLM reasoning	0.338	0.252	0.511	NR
	SAT ^[88]		0.354	0.319	NR	NR
	SRP-KGC ^[89]		0.431	0.353	0.585	NR
	AnyBURL ^[50]	Symbolic reasoning	0.322	0.229	0.553	NR
Hetionet ^[90] : biomedical KG filtered tail prediction	Noisy-or rule aggregation ^[91]		0.472	0.377	0.642	NR
	TransE ^[35]	Representation learning	0.205	0.099	0.444	NR
	CompLex ^[54]		0.25	0.152	0.47	NR
	RotatE ^[55]		0.257	0.185	0.403	NR
	CompGCN ^[84]	Path and structure-aware reasoning	0.292	0.172	0.543	NR
	MINERVA ^[92]		0.37	0.264	0.593	NR
	PoLo ^[93]		0.43	0.337	0.641	NR
	HittERnc ^[85]	Transformer-based graph reasoning	0.453	0.316	0.74	NR
	Common neighbor ^[94]	Symbolic reasoning	NR	NR	NR	0.2765
	Resource allocation ^[95]		NR	NR	NR	0.4933
ogbl-ppa ^[81] : large-scale biological protein association link prediction	Matrix factorization ^[81]	Representation learning	NR	NR	NR	0.3229
	SEAL ^[96]	Path and structure-aware reasoning	NR	NR	NR	0.488
	NGNN + SEAL ^[97]		NR	NR	NR	0.5971
	Neural common neighbor ^[98]		NR	NR	NR	0.6119
	GraphGPT ^[67]	Transformer-based graph reasoning	NR	NR	NR	0.7655

Benchmark^[81]). MRR measures how highly the correct entity or relation is ranked, whereas Hits@K measures whether the correct answer appears among the top K candidates. Filtered ranking is commonly used in KGC to avoid penalizing other known true triples.

Downstream biomedical tasks often require task-specific metrics. Binary classification tasks such as drug-target interaction, disease-target association, drug repurposing, adverse-event prediction, and combination-synergy classification commonly use AUROC, AUPRC, F1-score, precision, recall, specificity, balanced accuracy, and Matthews correlation coefficient (MCC). Ranking-oriented screening tasks often report recall@K, precision@K, enrichment factor, BEDROC, MRR, or Hits@K. Binding-affinity and dose-response prediction tasks may use RMSE, MAE, Pearson or Spearman correlation, and concordance index. Therefore, Table 1 should be interpreted as a representative KGC/link-prediction comparison, not as a complete evaluation framework for all downstream drug-discovery tasks.

The metrics support several paradigm-level observations. Symbolic reasoning remains competitive when rule-like graph regularities are informative: Resource Allocation reaches a Hits@100 of 0.4933 on ogbl-ppa, and Noisy-or rule aggregation achieves an MRR of 0.472 and Hits@10 of 0.642 on Hetionet. Representation learning provides scalable baselines, but richer structural models are often stronger; for example, TuckER reaches an MRR of 0.358 on FB15k-237, whereas NBFNet improves this to 0.415. Path and structure-aware methods are particularly useful when multi-hop topology is informative, as shown by PoLo (pruned) on Hetionet (MRR = 0.430, Hits@10 = 0.641). Transformer-based graph reasoning further improves long-range or relation-aware modeling, with KnowFormer reaching an MRR of 0.430 on FB15k-237 and HittERnc reaching an MRR of 0.453 and Hits@10 of 0.740 on Hetionet. Hybrid KG-LLM models show promising general KGC results, represented by SRP-KGC on FB15k-237 (MRR = 0.431, Hits@10 = 0.585), but their biomedical utility still requires leakage control, task-specific validation and prospective evidence.

Quantitative analysis also has certain limitations, as high MRR, Hits@K, AUROC, or F1 may still reflect dataset-specific shortcuts

rather than prospective biological utility. The five paradigms are therefore further compared from a practical drug-discovery perspective, focusing on use cases, strengths, biomedical limitations, and validation strategies (Table 2).

Benchmarking and validation of biomedical knowledge graph reasoning models

The evaluation of biomedical KGs is affected by negative-sample design, data splitting, and benchmark bias. Uniform random negative sampling often creates biologically implausible or trivially false entity pairs, which lowers task difficulty and may inflate performance. More realistic strategies, including type-constrained, relation-aware, degree-aware, and hard-negative sampling, better reflect real discovery settings but also make prediction more challenging. Because unobserved biomedical relations are usually unlabeled rather than truly negative, negative sets may contain undiscovered true associations. Evaluation protocols should therefore respect entity and relation ontologies, avoid impossible combinations, and report the positive-to-negative ratio, resampling strategy, and whether different models use the same negative set.

Data splitting can also cause leakage. Random train-test splits may share inverse or duplicated edges, ontology parent-child relations, similar drugs or proteins, overlapping disease hierarchies, literature co-occurrence, or similar entity neighborhoods across splits. These shortcuts allow models to infer test links without learning generalizable biological mechanisms. Temporal splits, target- or disease-disjoint splits, scaffold-based drug splits, relation holdout, and cold-start entity splits are more suitable for testing real-world generalization. For target identification and drug discovery, time-cutoff-based prospective evaluation is especially important because it asks whether models trained on historical knowledge can predict later evidence.

Benchmark bias is another major concern. Public databases are enriched for well-studied targets, common diseases, approved drugs, and high-degree entities. Models may therefore reproduce historical research attention rather than discover new mechanisms,

Table 2. Practical comparison of KGR paradigms for target and drug discovery.

Reasoning paradigm	Core methodology	Key applications in drug discovery	Critical comparison in biomedical scenarios
Symbolic reasoning	Logical rule learning, differentiable rule reasoning, random walk path sampling	Target identification, mechanism of action analysis, high-confidence association prediction	These methods are most useful when transparent evidence chains and expert-auditable hypotheses are required. They can be competitive on biomedical link-prediction benchmarks, but their coverage may be limited in sparse or noisy KGs, and rule search can become expensive in large heterogeneous graphs.
Representation learning	Knowledge graph embedding (KGE), relation-aware message passing, attention-based feature aggregation	Drug-target interaction prediction, disease-target association identification, molecular property prediction	These models scale well to large biomedical KGs and provide strong graph-fitting baselines for candidate ranking. Their main limitations are limited mechanistic interpretability, sensitivity to negative sampling and random splits, possible hub-node or literature-density bias, and over-smoothing in deep GNNs.
Path and structure-aware reasoning	Explicit path enumeration, reinforcement learning for path search, subgraph inductive reasoning, neural implicit path modeling	Multi-hop mechanistic inference, cold-start target/molecule prediction, drug repurposing path mining	These models are attractive when the task requires interpretable multi-hop evidence or generalization to sparsely annotated diseases and drugs. Their reliability depends strongly on KG completeness, path quality, disease-split design and candidate-universe definition.
Transformer-based graph reasoning	Structural encoding + self-attention, spectral feature embedding, local convolution + global transformer, linear-complexity state space models	Long-range target regulation prediction, cross-pathway association inference, large-scale KG reasoning	These models are conceptually suitable when distant graph context, cross-pathway information or multi-source molecular features are expected to influence prediction. However, biomedical benchmark evidence is often task-specific rather than evaluated under a shared KG benchmark protocol, and performance may be inflated when similar drugs, homologous proteins or related diseases appear across train and test splits.
Hybrid KG-LLM reasoning	Retrieval-augmented generation, subgraph prompting, multi-agent collaborative reasoning, chain-of-thought graph reasoning	Complex mechanism analysis, low-resource drug discovery, interpretable clinical translational reasoning	These methods are promising for natural-language evidence synthesis and explanation-rich biomedical reasoning. However, their evaluation should include evidence of faithfulness, path validity, calibration, prompt robustness, and hallucination control rather than predictive accuracy alone.

and may exploit shortcuts from ontology structure, gene-family similarity, chemical analogs, literature frequency, or curation patterns. This partly explains why many biomedical KG models perform well on internal benchmarks but fail in prospective validation.

Accordingly, metric and model selection should match dataset scale, relation heterogeneity, entity-type constraints, graph topology, and the biomedical task. Resources such as Open Graph Benchmark, OpenBioLink, KG-Hub, and PyKEEN support reproducible comparison. However, MRR and Hits@K only measure ranking quality, not translational value. Prospective validation should train models on data available before a defined time point and test predictions against later literature, experiments, or clinical progress. Low-throughput assays, perturbation screens, and expert review remain necessary to confirm biological relevance, feasibility, and safety.

Target discovery

Disease-target association discovery and prioritization

In the early stage of drug discovery, researchers need to screen candidate targets from a large number of genes, proteins, pathways, and disease-related molecules. Traditional target discovery mostly relies on single experimental results or manually collated disease-gene relationships, with limited sources of evidence. Knowledge graphs can integrate literature, omics, genetics, protein-protein interactions, pathways, phenotypes, drugs, and clinical annotations. PrimeKG, DisGeNET, and TarKG have demonstrated this multi-level integration.

Early network methods mainly relied on network proximity and disease modules. Candidate targets are usually evaluated according to their distance from known disease-associated genes or proteins^[99,100] (Fig. 3a). This method can find candidate targets in disease-related network regions, but can only provide evidence of association. With the development of knowledge graph representation learning and graph neural networks, disease-target discovery has gradually turned to multi-relational reasoning and graph completion. KDGene^[101], MiGCN^[102], KG2ML^[103], Progeni^[104], TarKG^[105], and other methods can capture the complex relationships among diseases, genes, proteins, drugs, pathways, phenotypes, and literature evidence.

Target ranking needs to pay attention to the level of evidence. Literature co-occurrence, network proximity, and knowledge graph link prediction can propose candidate targets, but they are vulnerable to publication bias, hub node effect, and incomplete database. GWAS signals, colocalization analysis, Mendelian randomization, and Open Targets Genetics^[106] can provide stronger genetic support. CRISPR screening, RNAi perturbation, Perturb-seq, loss-of-function experiments, and dependency mapping analysis^[107] can further determine whether the perturbed target changes the disease-related phenotype.

At present, target ranking increasingly focuses on multidimensional evidence, including disease relevance, genetic support, functional dependence, tractability, safety, tissue specificity, and clinical relevance (Fig. 3b). Open Targets Genetics^[106], resources and models such as dependency map analysis^[107], GraphEvo^[108], KGDRP^[109], and PubMed Knowledge Graph 2.0^[110] illustrate that joint evaluation of multiple evidences can improve the reference value of target ranking. However, diseases or proteins with more research may have higher scores due to more literature, more annotations, and more

connections; less studied targets may also be underestimated due to insufficient evidence. Therefore, target ranking also needs to combine structural tractability, ligandability, tissue distribution, safety, and experimental validation results.

Target druggability assessment

After the candidate targets are sorted, the druggability and tractability need to be further evaluated (Fig. 3c). Disease relevance alone is not sufficient to support target development. A therapeutic target should not only participate in the disease process but also be effectively regulated by small molecules, antibodies, peptides, RNA drugs, or targeted protein degradation technology. Target tractability emphasizes safe and effective regulation, while druggability pays more attention to the development of drug-like small molecules or biological agents^[111,112].

Structural biology is the core of drug evaluation. In the context of small molecule discovery, major considerations include binding pockets, molecular volume, cavity depth, hydrophobicity, electrostatic complementarity, conformational flexibility, and allosteric sites^[113]. For antibody or protein drugs, greater emphasis is placed on extracellular accessibility, membrane topology, epitope exposure, and tissue-specific expression. Structural resources such as PDB^[114], AlphaFold2/3, ESMFold, and RoseTTAFold can provide important information, but the predicted structures need to be interpreted carefully when the confidence is low, the key conformation is missing, or there is a ligand-induced conformational change.

Knowledge graphs can transform structural information and tractability into nodes and edges. A protein can be linked to a protein structure, binding pocket, ligand-binding events, mutation annotation, known inhibitors, tissue expression, essentiality score, and safety risks. Pocket descriptors such as fpocket score, SiteMap volume, DoGSiteScorer score^[115], hydrophobicity, flexibility, and allosteric potential can also be added to the knowledge graph. This structure-aware representation helps to distinguish 'disease-related targets' from 'therapeutically tractable targets'. DrugnomeAI^[116] and DrugTar^[117] also show that gene, network, protein language model, and Gene Ontology information can jointly improve the druggability prediction performance.

Safety is also part of the tractability assessment. If the target is important for normal tissue function or widely expressed in key organs, regulating it may bring target-related toxicity. Therefore, tissue expression, human genetic constraints, essential gene status, pathway redundancy, and disease background specificity need to be considered when sorting targets.

Mechanism interpretation and pathway reasoning

After the candidate targets are sorted, it is also necessary to explain how they participate in the disease mechanism (Fig. 3d). Knowledge graphs can put candidate targets into disease-related biological processes, signaling pathways, cellular phenotypes, and molecular cascades, and form traceable relationship paths and subgraphs^[118]. With the increase of multi-omics data, mechanism analysis needs cross-level integration. Similarity network fusion has been used for disease typing and identification of key regulatory factors^[119]. Breast cancer research also shows that integrating molecular, pathway, and clinical information is helpful in discovering regulatory factors and disease-driving pathways^[120]. GNNRAI further combines multi-omics data with biomedical knowledge to improve the interpretability of disease prediction^[121].

Path-based knowledge graph reasoning can connect candidate targets with disease biology. Metapath, random walk, shortest path,

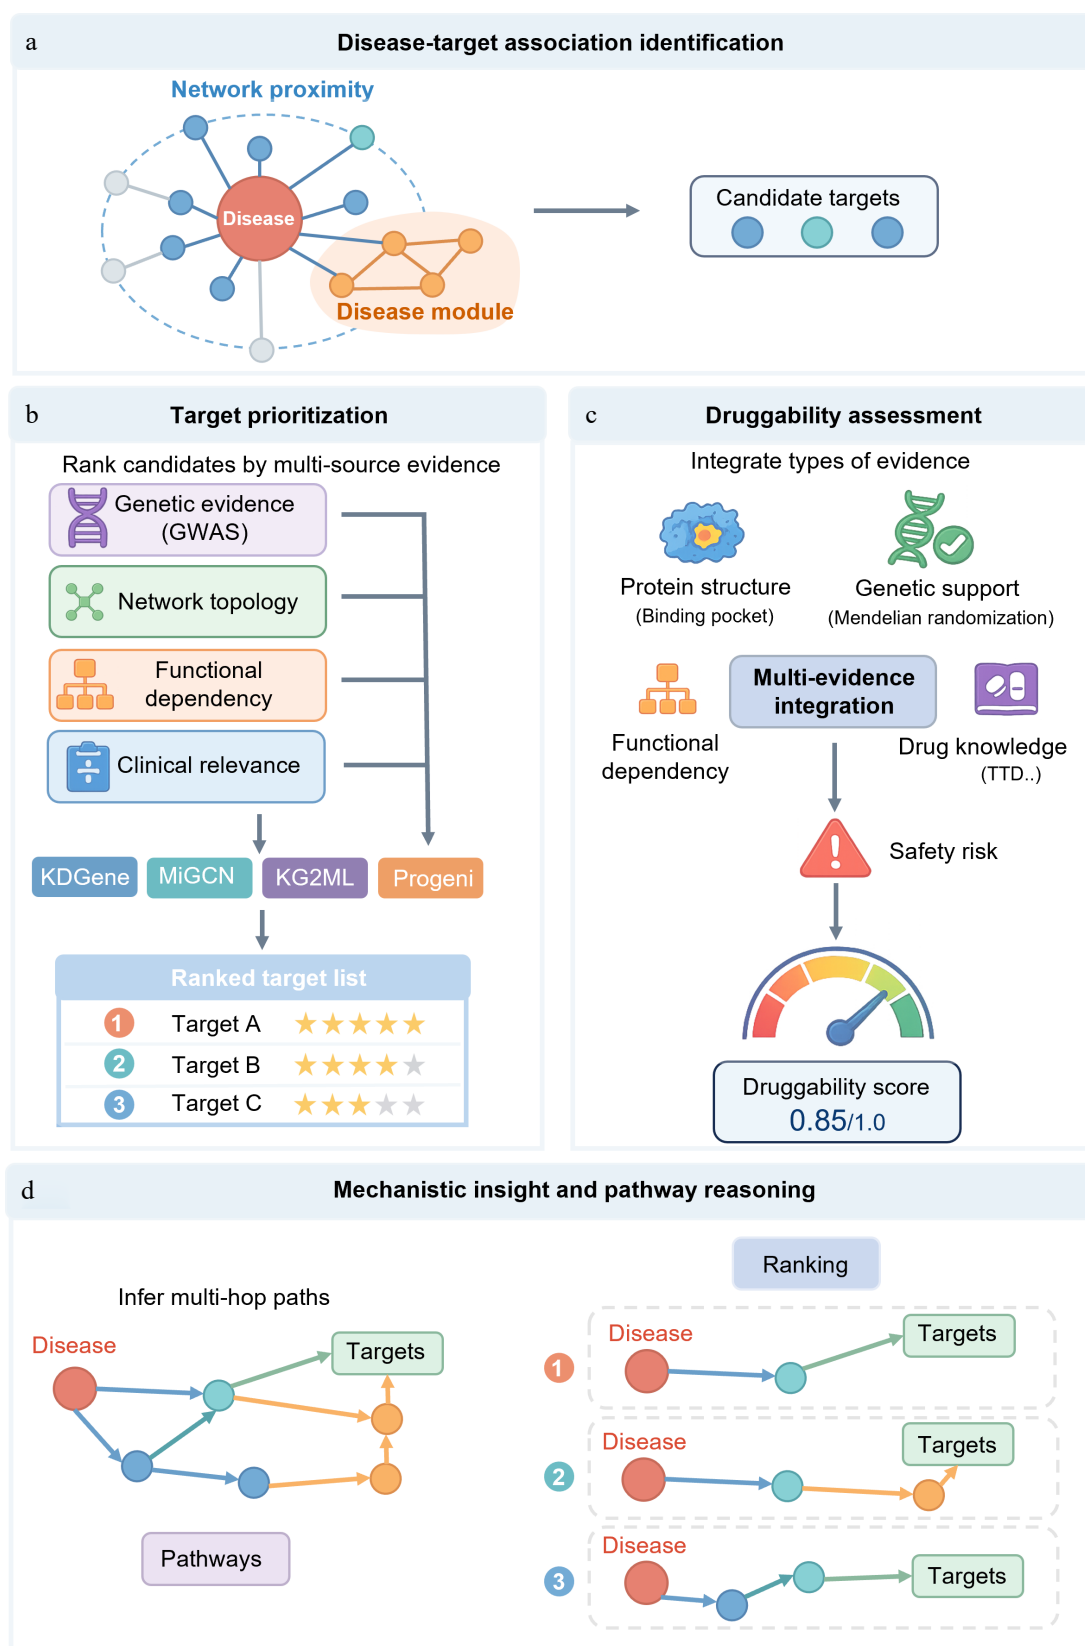


Fig. 3 A knowledge graph-based framework for disease-related target discovery, prioritization, druggability assessment, and mechanism analysis. (a) Identification of disease-related targets. Candidate targets associated with disease-related molecules are identified using network proximity, disease module expansion, and multi-relational knowledge graph representation learning. (b) Target prioritization. Candidate targets are ranked by integrating genetic evidence, network topology, functional dependencies, and clinical relevance. (c) Druggability assessment. The development potential of targets is evaluated by combining structural suitability, genetic support, functional networks, existing drug knowledge, and safety risks. (d) Mechanism elucidation and pathway inference. Infer and rank multi-hop pathways connecting diseases to candidate targets within the context of disease-related networks and pathways to reveal potential molecular mechanisms.

subgraph interpretation, and neural path reasoning can identify the relationship chain among diseases, targets, pathways, phenotypes, drugs, and molecular functions. BioPathNet improves link prediction ability through learning path representation and shows key reasoning paths^[122].

The text-enhanced knowledge graph model can also connect the graph relationships with semantic evidence of the literature, thereby improving the interpretability^[123]. However, a disease-target pathway can only explain that the mechanism may be reasonable and cannot directly prove causality. Incomplete databases, uneven literature coverage, duplicate pathway annotations, and hub node interference may affect pathway interpretation.

Perturbation evidence can test mechanistic hypotheses. CRISPR screening, RNAi experiments, single-cell perturbation, drug perturbation, transcriptomics, and pathway activity analysis can determine whether downstream pathways or disease-related phenotypes change after regulating candidate targets. For COVID-19 research, virus-host interaction networks are adopted to reconstruct infection-related molecular networks and identify potential intervention targets^[124]. MMiKG and MicrobiomeKG also support mechanism-related pathway discovery by integrating microbial, host, and phenotypic information^[125,126].

DTI prediction and molecular design

Drug-target interaction prediction

Early DTI predictions mainly relied on chemical similarity, molecular fingerprints, target annotation, and handcrafted descriptors. When the target has a known ligand, this kind of method has a certain effect; however, in data-sparse or cold-start scenarios, the effect is easily limited. With the development of deep learning, DTI prediction has gradually turned to representation learning. DeepDTA^[127] and DeepAffinity^[128] use protein sequences and SMILES sequences to predict binding affinity. GraphDTA represents compounds as molecular graphs, thus capturing richer chemical topology information^[129] (Fig. 4a).

It is still not enough to rely only on chemical structures and protein sequences. The prediction results may also be affected by target family, pathway background, disease relevance, drug category, off-target spectrum, and experimental assay evidence. A knowledge graph can put drugs, proteins, diseases, pathways, phenotypes, and experimental data into the same relationship space (Fig. 4b). DTINet^[130] and TriModel^[131] illustrate that heterogeneous network integration helps to improve DTI prediction. It is recommended that the systems-level model and the multi-relational graph completion method can also improve the prediction effect in the weak supervision and cold-start scenarios through the information propagation between related entities.

In recent years, research has put more emphasis on the joint modeling of chemistry, protein, semantics, and graph context. DTKGIN improves DTI recognition ability in sparse settings^[132]. KGE-UNIT shows that cross-task knowledge sharing can improve the effect of relationship prediction^[133]. CBKG-DTI extracts complex network features through hierarchical embedding and knowledge distillation^[134]. LKE-DTA further integrates protein sequences and biomedical knowledge by combining the sequence semantic representation of large language models with the embedding of knowledge graphs^[135] (Fig. 4c). It should be noted that the predicted drug-target edge may only reflect the similarity, common annotation, literature co-occurrence or database curation rules, which still need to be verified by molecular docking, molecular dynamics,

biophysical experiments, biochemical experiments, and cellular target engagement assays.

Structure-aware knowledge graph for small molecule discovery

Three-dimensional structure is a key factor in the discovery of small molecules. The physical binding depends on the geometric complementarity between the compound and the binding site as well as the physicochemical complementarity. It is difficult to judge whether the ligand can enter the pocket, form effective interactions, avoid steric clashes, or stabilize the functional conformation based on sequence similarity or graph proximity alone.

The structure-aware knowledge graph can connect chemical nodes, target nodes, protein structure nodes, binding pocket nodes, experimental nodes, mutation nodes, and disease nodes. As exemplified by target proteins, they can be linked to PDB experimental structures, predicted structures generated by AlphaFold2/3, ESMFold or RoseTTAFold, binding pockets, known ligands, drug-resistance mutations, allosteric sites, and bioactivity assays. Pocket volume, depth, hydrophobicity, electrostatic properties, flexibility, ligandability, and allosteric potential can also be added to the graph as features. Tools such as fpocket, SiteMap, and DoGSiteScorer^[115] can provide relevant descriptors. This structural representation can distinguish 'biologically relevant targets' from 'chemically tractable targets'. Because the predicted structure is affected by flexible regions, alternative conformations, ligand-induced fit, membrane environment, and disordered regions, the structure-aware model should also incorporate structural confidence and combine docking, molecular dynamics, or experimental validation.

KG-guided molecular generation

After the ranking of candidate targets and potential interactions, drug discovery enters the phase of molecular generation and lead optimization. Knowledge graphs can connect target biology, pathway safety, off-target risk, structure-activity relationship, ADMET properties, toxicity signals, and disease background, provide constraints for molecular design, and make molecular generation pay more attention to therapeutic relevance than chemical novelty.

Previous studies have used external knowledge to guide molecular generation and reinforcement learning to focus the model on target-specific chemical space^[136]. KGDiff introduces protein-ligand binding information into the diffusion denoising process. DiffPhore guides 3D conformation generation through pharmacophore matching and supports virtual screening and lead discovery^[137] (Fig. 4d). KGG added chemical knowledge at the orbital level to graph self-supervised learning, showing potential in low-data scenarios^[138]. However, the generated molecules may be chemically unreasonable, difficult to synthesize, unstable, difficult to purify, or approximate known compounds. High predicted binding affinity also does not necessarily translate into selectivity, cellular activity, permeability, metabolic stability, or safety.

Lead compound optimization and multi-parameter optimization

Lead compound optimization is not to improve the target binding ability alone, but to find a balance between activity, selectivity, solubility, permeability, metabolic stability, hERG risk, CYP-mediated metabolism or CYP inhibition, toxicity, and synthetic accessibility. Multi-parameter optimization (MPO) has become an important strategy to improve the quality of candidate compounds and reduce the risk of late failure^[139].




Fig. 4 Knowledge graph-based prediction of small-molecule-target interactions, molecular design, and translational applications. (a) Prediction of small-molecule-protein interactions. This has evolved from predictions based on chemical similarity and ligand-protein pairs to deep learning-based affinity prediction using molecular graph and protein sequence features. (b) Knowledge-guided relationship inference. Integrating heterogeneous entities such as drugs, proteins, diseases, pathways, genes, and side effects into a knowledge graph enables models to shift from local pairwise judgments to relationship inference within a biomedical context. (c) Graph reasoning and multimodal learning. With the introduction of graph reasoning, multi-task knowledge sharing, and multimodal representation, drug-target identification has evolved from simple relationship completion to joint modeling within complex biological contexts. (d) Molecular generation. Knowledge graphs are not only used for candidate screening but also provide conditional guidance for molecular generation and lead optimization, shifting design from rule-driven to information-driven approaches. (e) Drug repurposing and combination therapy. Knowledge graphs support drug reuse, synergistic effect prediction, and drug interaction modeling, extending candidate discovery to system-level inference based on multi-source evidence and contextual constraints. (f) Safety assessment and precision medicine. Knowledge graphs can integrate ADMET, toxicity, adverse reactions, and patient multimodal data for safety risk assessment and precision medicine analysis.

Knowledge graphs can connect compounds, targets, off-target proteins, ADMET endpoints, metabolic enzymes, transporters, toxic phenotypes, and disease pathways, thus providing constraints for multi-parameter optimization. In the case of transporter protein knowledge graphs, the integration of AutoInt_KG and MolGPT_KG enables drug-transporter relationship prediction and transporter information-guided molecular generation^[140]. Graph-based molecular Pareto optimization uses graph representation and the non-dominated sorting genetic algorithm for Pareto frontier optimization^[141]. ChemMORT uses deep learning and multi-objective particle swarm optimization to optimize multiple ADMET endpoints, indicating that multi-objective ADMET optimization has become an important direction of computer-aided lead optimization^[142]. Therefore, the value of knowledge graphs is not only to improve the accuracy of single-property prediction but also to help screen more balanced candidate molecules.

Translational applications in the drug discovery process

Target discovery and validation

In translational drug discovery, target discovery should not rest on the ranking of candidate genes or proteins, but more importantly, on whether candidate targets can enter the validation cohort. Knowledge graphs can integrate disease association, genetic evidence, functional perturbation, pathway location, tissue expression, drug resistance, safety risk, and biomarker support around each target. In the case of target screening, targets with GWAS colocalization, disease-associated expression, perturbation evidence, and ligandable structures are generally more amenable to downstream research than those supported merely by literature co-occurrence or network proximity.

Previous studies have shown the possibility of the knowledge graph target hypothesis entering experimental validation. Rosalind predicts the disease-gene relationship based on heterogeneous knowledge graphs and tensor decomposition and validates candidate targets of rheumatoid arthritis in the experiment of patient-derived fibroblast-like synoviocytes, resulting in five potential genes^[143]. Progeni constructed a probabilistic knowledge graph, and its top-ranked targets in melanoma and colorectal cancer were supported by knockdown experiments and pharmacological validation^[104]. In idiopathic pulmonary fibrosis, Standigm ASK™ proposed AMFR, MDFIC, and NR5A2, and verified their functional relevance through siRNA knockdown experiments^[144]. BenevolentAI identified baricitinib for COVID-19, and subsequent ACTT-2 and COV-BARRIER studies provided clinical evidence^[145–147]. These cases show that the knowledge graph is more suitable for establishing a validation cohort, and further confirmation is still needed in the follow-up by CRISPR, RNAi, organoids, animal models, patient-derived samples, or pharmacological perturbation.

Hit discovery and DTI prediction

After the target ranking and preliminary validation, the next step is to find compounds or therapeutic methods that can regulate the target. DTI prediction based on a knowledge graph can sort candidate drug-target pairs, help less-studied targets find potential ligands, and suggest possible off-target interactions. DeepDTA^[127], DeepAffinity^[128], GraphDTA^[129], DTINet^[130], TriModel^[131], DTKGIN^[132], KGE-UNIT^[133], CBKG-DTI^[134], and LKE-DTA^[135], which reflects that DTI prediction has gradually shifted from sequence and

structure modeling to relational reasoning combined with biological background. The priority of candidate compounds should not only depend on predicted binding but also on the location of disease pathways, off-target risks, and experimental clues.

Lead optimization and molecular design

Lead compound optimization links computational molecular design and drug development. Real optimization can not only improve the target binding ability, but also consider the potency, selectivity, solubility, permeability, metabolic stability, clearance, hERG risk, CYP inhibition, toxicity, tissue distribution, and synthetic feasibility. Molecules that only pursue affinity may fail due to poor ADMET properties, off-target toxicity, insufficient selectivity, or difficulties in synthesis.

Knowledge graphs can connect compounds, chemical series, targets, off-target proteins, experiments, metabolic enzymes, transporters, toxicity endpoints, adverse events, pathways, and clinical risks and support multi-target scores and Pareto-frontier ranking. For molecular optimization, compounds with enhanced potency accompanied by hERG liability, CYP inhibition or hepatotoxicity demand deliberate prioritization, whereas those with moderate potency and superior druggable properties are more suitable for further optimization. The knowledge graph can also point out that the limiting factors are insufficient target affinity, off-target activity, metabolic instability, poor permeability, structural alerts, or pathway toxicity, thus providing directions for medicinal chemistry optimization.

Safety, ADMET, and risk assessment

The safety assessment should enter the target and compound prioritization as early as possible. Risks at the target level include target-related toxicity, critical organ expression, essential gene status, human genetic constraints, developmental functions, pathway redundancy, and disease background specificity. Compound-level risks include hERG risk, CYP inhibition, drug-induced liver injury, cardiotoxicity, genotoxicity, reactive substructure, and transporter-related risks (Fig. 4f).

Knowledge graphs can integrate FAERS^[148], SIDER^[149], CTD^[150], DrugBank, ToxCast^[151], Open Targets Safety^[152], electronic health records, toxicological genomics data, and pathway annotations. Bean et al. constructed a knowledge graph containing drugs, targets, indications, and adverse reactions, and verified the high-confidence prediction with electronic health records^[153]. Subsequent studies have improved ADR risk identification ability through representation learning^[154] and extended it to multi-drug adverse reaction prediction^[155]. MVDSA integrates multi-source knowledge, relational semantics, and multi-perspective features for joint association prediction^[156]; graph representation and transformer model have also been used for multi-label ADR prediction^[157].

ADMET evaluation focuses on absorption, distribution, metabolism, excretion, and toxicity. Machine learning frameworks are gradually replacing traditional empirical rules^[158]. However, the pharmacovigilance database may have reporting bias, confounding indications, the impact of multiple drug combinations, and a lack of denominator information. Therefore, the associative signal is more suitable for interpretation as a risk association rather than a causal safety conclusion.

Drug repositioning and combination therapy

Drug repositioning aims to find new indications for approved or clinical-stage compounds (Fig. 4e). Existing drugs usually

accumulate safety, pharmacokinetics, manufacturing, and clinical information, so repositioning may shorten the development cycle and reduce uncertainty^[159,160]. PREDICT predicts indications based on drug-disease similarity^[161]. During COVID-19, network proximity, network diffusion, and AI-based prioritization were used to screen drug candidates^[162]. Rephetio integrates multiple biomedical entities and associations for systematic screening of drug-disease combinations^[90]. OREGANO provides a standardized repositioning knowledge base^[163]. DREAMwalk alleviates the problem of uneven gene distribution through semantically guided random walk^[164]. TxGNN framed drug repositioning as a zero-shot prediction task and demonstrated the potential of ranking indications and contraindications^[165]. Coronavirus-related research also used knowledge graphs to integrate literature, databases, and deep learning results to screen drug candidates^[166,167].

Combination therapy needs to distinguish synergy prediction from DDI prediction. Synergy prediction focuses on whether multiple drugs can enhance therapeutic benefit; DDI prediction focuses on whether co-administration increases toxicity, changes pharmacokinetics, or produces unexpected pharmacodynamic effects. SynergyGraph combines cell-line background, knowledge graph representation, and hypergraph modeling^[168]. HyperSynergyX integrates hypergraph modeling and knowledge graph-enhanced retrieval for mechanism interpretation^[169]. MDDI-SCL^[170], HGTSynergy^[171], and Pisces^[172] illustrate that synergy prediction and DDI prediction are moving towards joint modeling of shared knowledge. In DDI prediction, SumGNN^[173], LaGAT^[174], KG-CLDDI^[175], Heterogeneous drug knowledge graph^[176], predicting rich DDI^[177], KA-DDI^[178], and LLM-DDI^[179], which reflects the fusion trend of graph structure, text information, and representation learning. Subsequent models also need to incorporate dose, administration time, tissue exposure, pharmacokinetics, pharmacodynamics, and toxicity overlap.

Biomarkers and precision medicine

Biomarker discovery and patient stratification aim to link molecular characteristics with diagnosis, prognosis, treatment response, and clinical decision-making. Knowledge graphs can integrate mutations, gene expression, proteins, pathways, drug responses, phenotypes, electronic health records, and clinical outcomes, and transform dispersed molecular signals into clinical hypotheses that are easier to interpret.

The Monarch Initiative integrates gene disease and phenotype information to provide the basis for cross-species biomarker discovery^[180]. The Clinical Knowledge Atlas can connect proteomic changes with literature and pathway knowledge to form a more grounded biomarker hypothesis^[181]. In non-small cell lung cancer, incorporating prior knowledge of knowledge graphs into machine learning can improve survival prediction and enhance candidate biomarker interpretability^[182]. The patient-specific knowledge graph can integrate medical history, genetic variation, molecular characteristics, treatment exposure, and health outcomes to build an individualized representation^[183]. PreMedKB integrates diseases, genes, variants, and drugs to provide a unified entry for the interpretation of evidence in precision medicine^[184].

Biomarker evidence needs to be evaluated hierarchically. Retrospective association, cross-cohort validation, prospective validation, clinical utility assessment, and companion diagnostic approval represent different levels of evidence. Cohort bias, batch effects, population differences, overfitting, treatment heterogeneity, and endpoint definitions all affect the reliability of the results. Therefore, in the absence of prospective clinical validation, biomarker

discovery supported by knowledge graphs is more suitable to be regarded as a patient stratification hypothesis.

Dynamic KGs and drug resistance

Most biomedical knowledge graphs are static, but drug responses and disease progression are dynamic processes. Under the pressure of treatment, tumors, pathogens, and diseased tissues may undergo genetic, epigenetic, transcriptional, protein, and microenvironment remodeling. This can be illustrated by drug-resistant mutations. EGFR T790M will change the sensitivity of kinase inhibitors in lung cancer. ALK G1202R is associated with ALK inhibitor resistance. BCR-ABL T315I will affect the response of a variety of tyrosine kinase inhibitors^[185–187]. These mutations may change the binding pocket, reduce drug affinity, change the conformational state, activate bypass signals, or remodel downstream pathways.

The static knowledge graph can record the mutation drug-resistance association, but it is difficult to explain when the mutation appears, under what treatment pressure it is selected, and how the pathway activity changes. Dynamic knowledge graphs, temporal knowledge graphs, versioned knowledge graphs, longitudinal omics, single-cell perturbation data, and treatment response trajectories can represent changes in disease status, clonal composition, pathway activation, target dependence, and drug sensitivity. In the future, knowledge graphs need to move from static evidence integration to time-aware decision support, especially for tumors, infectious diseases, immune diseases, and chronic diseases.

Conclusions and future prospects

In general, biomedical knowledge graphs can break through the barriers of multi-source heterogeneous data and build a multi-level computable knowledge network from molecular to clinical. At present, AI technology has promoted its automatic application and has shown obvious potential in target identification, drug target prediction, drug repurposing, and other tasks.

However, several key challenges remain in this area. First, the quality of multi-source data integration remains problematic. Biomedical knowledge is scattered in different fields, so it is difficult to integrate multi-scale and multi-modal data with high fidelity. In the case of GWAS data mapping, population bias may be introduced. Activity data from different platforms can also undermine edge weight confidence due to disparate experimental conditions^[8]. In the future, it is necessary to establish a more stringent confidence evaluation mechanism in the data fusion stage^[188], filter false positive associations by combining uncertainty modeling^[189], and carry out automatic fact verification relying on domain ontology^[190].

The second is that the generalization ability of the model is still insufficient. At present, reasoning modules are facing problems such as cold start and insufficient out-of-distribution generalization^[191–193]. Methods such as transfer learning and zero-shot learning have been widely studied^[194,195]. In the future, it is also necessary to combine the anomaly detection mechanism to improve the reliability of real-world deployment.

In addition, most current biomedical KGs are still static snapshots. This static nature is a fundamental limitation for resistance prediction and adaptive therapy design, because resistance is not a fixed attribute of a drug-target pair but an evolving process shaped by clonal selection, compensatory pathways, tumor microenvironment, drug pressure, and prior treatment history. A static KG may correctly encode that a drug inhibits a target, but it cannot represent when a resistant clone emerges, which bypass pathway becomes dominant,

or how sequential therapy changes future sensitivity. As a result, predictions based only on fixed associations may lag behind real disease trajectories and may overestimate the durability of a candidate target or regimen.

Self-evolving KGs should therefore be defined more concretely than through continuous updating. They need timestamped entities and edges, versioned evidence, edge confidence that changes with new assays or clinical observations, and mechanisms for retiring contradicted relations. For adaptive treatment, this means representing patient-state transitions, therapy-induced perturbations, emerging resistance mechanisms, and feedback from longitudinal omics, imaging, electronic health records, and trial data. Temporal KG analysis during COVID-19 has shown how scientific understanding can change over time^[196], and incremental embedding methods can maintain representations as graphs grow^[197]. However, the key challenge is not only adding new triples but also enabling reasoning models to compare graph states before and after treatment and to recommend interventions that remain valid under evolving biology.

Finally, privacy protection and federated knowledge graphs are the key to breaking the institutional data silos. Privacy and compliance constraints keep a large amount of data inside the organization, while federated KG can support inter-agency collaboration without disclosing the original data. Existing studies have verified its effectiveness^[198–200]. In the future, it is still necessary to optimize the model architecture, improve the performance of decentralized reasoning, and support secure cross-agency sharing.

With the continuous breakthroughs in these directions, biomedical knowledge graphs are expected to further promote the implementation of AI drug discovery and provide more powerful support for drug discovery and development, and precision medicine.

Ethical statements

Not applicable.

Author contributions

The authors confirm their contributions to the work as follows: conceptualization: Li S, Li H; manuscript writing: Dou P, Yang X; manuscript revision: Jiang Y, Xie X, Geng X, Huang H. All authors reviewed the results and approved the final version of the manuscript.

Data availability

Data sharing is not applicable to this review as no datasets were generated or analyzed.

Acknowledgments

This work was supported in part by the National Natural Science Foundation of China (82425104 to Honglin Li), the National Key R&D Program of China (2022YFC3400504). Shiliang Li is also sponsored by the Shanghai Rising-Star Program (23QA1402800).

Conflict of interest

The authors declare that there is no conflict of interest.

Supplementary information accompanies this paper online at: <https://doi.org/10.48130/targetome-0026-0025>.

Dates

Received 27 April 2026; Revised 31 May 2026; Accepted 3 June 2026; Published online 25 June 2026

References

- [1] Catacutan DB, Alexander J, Arnold A, Stokes JM. 2024. Machine learning in preclinical drug discovery. *Nature Chemical Biology* 20:960–973
- [2] Sadybekov AV, Katritch V. 2023. Computational approaches streamlining drug discovery. *Nature* 616:673–685
- [3] Ocana A, Pandiella A, Privat C, Bravo I, Luengo-Oroz M, et al. 2025. Integrating artificial intelligence in drug discovery and early drug development: a transformative approach. *Biomarker Research* 13:45
- [4] Houssein EH, Hosney ME, Emam MM, Younis EMG, Ali AA, et al. 2023. Soft computing techniques for biomedical data analysis: open issues and challenges. *Artificial Intelligence Review* 56:2599–2649
- [5] Sarvepalli S, Vadarevu S. 2025. Role of artificial intelligence in cancer drug discovery and development. *Cancer Letters* 627:217821
- [6] Li Z, Hua S, Song S, Luo Y, Qian W, et al. 2026. Multi-omics integration identifies *MT2A* as a biomarker and a candidate host target linking zinc dysregulation to COVID-19 mortality. *Targetome* 2:e006
- [7] Peng C, Xia F, Naseriparsa M, Osborne F. 2023. Knowledge graphs: opportunities and challenges. *Artificial Intelligence Review* 56:13071–13102
- [8] Ding K, Zhu Z, Tang Y, Feng K, Zhuang X, et al. 2026. Bridging data and discovery: a survey on knowledge graphs in AI for science. *National Science Review* 13:nwag140
- [9] Chen X, Jia S, Xiang Y. 2020. A review: knowledge reasoning over knowledge graph. *Expert Systems with Applications* 141:112948
- [10] Steiner T, Verborgh R, Troncy R, Gabarro J, Van de Walle R. 2012. Adding realtime coverage to the Google knowledge graph. *ISWC-PD'12: Proceedings of the 2012th International Conference on Posters & Demonstrations Track*. Volume 914. Aachen, DEU: CEUR-WS. org. pp. 65–68 doi: 10.5555/2887379.2887396
- [11] Ji S, Pan S, Cambria E, Martinen P, Yu PS. 2022. A survey on knowledge graphs: representation, acquisition, and applications. *IEEE Transactions on Neural Networks and Learning Systems* 33:494–514
- [12] Mendez D, Gaulton A, Bento AP, Chambers J, De Veij M, et al. 2019. ChEMBL: towards direct deposition of bioassay data. *Nucleic Acids Research* 47:D930–D940
- [13] The UniProt Consortium. 2025. UniProt: the Universal Protein Knowledgebase in 2025. *Nucleic Acids Research* 53:D609–D617
- [14] Martin FJ, Amode MR, Aneja A, Austine-Orimoloye O, Azov AG, et al. 2023. Ensembl 2023. *Nucleic Acids Research* 51:D933–D941
- [15] Szklarczyk D, Kirsch R, Koutrouli M, Nastou K, Mehryary F, et al. 2023. The STRING database in 2023: protein-protein association networks and functional enrichment analyses for any sequenced genome of interest. *Nucleic Acids Research* 51:D638–D646
- [16] Kanehisa M, Furumichi M, Sato Y, Kawashima M, Ishiguro-Watanabe M. 2023. KEGG for taxonomy-based analysis of pathways and genomes. *Nucleic Acids Research* 51:D587–D592
- [17] Milacic M, Beavers D, Conley P, Gong C, Gillespie M, et al. 2024. The reactome pathway knowledgebase 2024. *Nucleic Acids Research* 52:D672–D678
- [18] Piñero J, Ramírez-Angueta JM, Saüch-Pitarch J, Ronzano F, Centeno E, et al. 2020. The DisGeNET knowledge platform for disease genomics: 2019 update. *Nucleic Acids Research* 48:D845–D855
- [19] Knox C, Wilson M, Klinger CM, Franklin M, Oler E, et al. 2024. DrugBank 6.0: the DrugBank knowledgebase for 2024. *Nucleic Acids Research* 52:D1265–D1275
- [20] Nguyen DT, Mathias S, Bologna C, Brunak S, Fernandez N, et al. 2017. Pharos: collating protein information to shed light on the druggable genome. *Nucleic Acids Research* 45:D995–D1002
- [21] Chen X, Yan CC, Zhang X, Zhang X, Dai F, et al. 2016. Drug-target interaction prediction: databases, web servers and computational models. *Briefings in Bioinformatics* 17:696–712

- [22] Ioannidis VN, Zheng D, Karypis G. 2020. Few-shot link prediction via graph neural networks for Covid-19 drug-repurposing. *arXiv* 00:2007.10261
- [23] Chandak P, Huang K, Zitnik M. 2023. Building a knowledge graph to enable precision medicine. *Scientific Data* 10:67
- [24] Aronson AR. 2001. Effective mapping of biomedical text to the UMLS metathesaurus: The MetaMap Program. *Journal of the American Medical Informatics Association* 2001:17–21
- [25] Wang Y, Tang SC. 2022. The race to develop oral SERDs and other novel estrogen receptor inhibitors: recent clinical trial results and impact on treatment options. *Cancer and Metastasis Reviews* 41:975–990
- [26] Pérez A, Weegar R, Casillas A, Gojenola K, Oronoz M, et al. 2017. Semi-supervised medical entity recognition: a study on Spanish and Swedish clinical corpora. *Journal of Biomedical Informatics* 71:16–30
- [27] Bose P, Srinivasan S, Sleeman WC, Palta J, Kapoor R, et al. 2021. A survey on recent named entity recognition and relationship extraction techniques on clinical texts. *Applied Sciences* 11:8319
- [28] Devlin J, Chang MW, Lee K, Toutanova K. 2018. BERT: pre-training of deep bidirectional transformers for language understanding. *arXiv* 00:1810.04805
- [29] Lee J, Yoon W, Kim S, Kim D, Kim S, et al. 2020. BioBERT: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics* 36:1234–1240
- [30] Habibi M, Weber L, Neves M, Wiegandt DL, Leser U. 2017. Deep learning with word embeddings improves biomedical named entity recognition. *Bioinformatics* 33:i37–i48
- [31] Boudjellal N, Zhang H, Khan A, Ahmad A. 2020. Biomedical relation extraction using distant supervision. *Scientific Programming* 2020:8893749
- [32] Luo L, Yang Z, Cao M, Wang L, Zhang Y, et al. 2020. A neural network-based joint learning approach for biomedical entity and relation extraction from biomedical literature. *Journal of Biomedical Informatics* 103:103384
- [33] Shvaiko P, Euzenat J. 2013. Ontology matching: state of the art and future challenges. *IEEE Transactions on Knowledge and Data Engineering* 25:158–176
- [34] Christen P. 2012. *Data matching: concepts and techniques for record linkage, entity resolution, and duplicate detection*. Berlin, Heidelberg: Springer. doi: 10.1007/978-3-642-31164-2
- [35] Bordes A, Usunier N, Garcia-Durán A, Weston J, Yakhnenko O. 2013. Translating embeddings for modeling multi-relational data. *NIPS'13: Proceedings of the 27th International Conference on Neural Information Processing Systems* 2:2787–2795
- [36] Sun Z, Zhang Q, Hu W, Wang C, Chen M, et al. 2020. A benchmarking study of embedding-based entity alignment for knowledge graphs. *Proceedings of the VLDB Endowment* 13:2326–2340
- [37] Paulheim H. 2016. Knowledge graph refinement: a survey of approaches and evaluation methods. *Semantic Web* 8:489–508
- [38] Pan S, Luo L, Wang Y, Chen C, Wang J, et al. 2024. Unifying large language models and knowledge graphs: a roadmap. *IEEE Transactions on Knowledge and Data Engineering* 36:3580–3599
- [39] Ye H, Zhang N, Chen H, Chen H. 2022. Generative knowledge graph construction: a review. *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, Abu Dhabi, United Arab Emirates, 2022*. USA: Association for Computational Linguistics. pp. 1–17 doi: 10.18653/v1/2022.emnlp-main.1
- [40] Zhang J, Qi S, Dong Y, Yuan L, Shen T, et al. 2026. GKG-LLM: a unified framework for generalized knowledge graph construction. *Information Fusion* 128:103956
- [41] Wang Q, Li C, Liu Y, Zhu Q, Song J, et al. 2025. An adaptive framework embedded with LLM for knowledge graph construction. *IEEE Transactions on Multimedia* 27:2912–2923
- [42] Ding Y, Li J, Wang P, Tang Z, Yan B, et al. 2024. Rethinking negative instances for generative named entity recognition. *Findings of the Association for Computational Linguistics: ACL 2024, Bangkok, Thailand*. pp. 3461–3475 doi: 10.18653/v1/2024.findings-acl.206
- [43] Zhang B, Soh H. 2024. Extract, define, canonicalize: an LLM-based framework for knowledge graph construction. *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing, Miami, Florida, USA*. pp. 9820–9836 doi: 10.18653/v1/2024.emnlp-main.548
- [44] Gao S, Yu K, Yang Y, Yu S, Shi C, et al. 2025. Large language model powered knowledge graph construction for mental health exploration. *Nature Communications* 16:7526
- [45] Dobbins NJ. 2025. Generalizable and scalable multistage biomedical concept normalization leveraging large language models. *Research Synthesis Methods* 16:479–490
- [46] Lin Z, Zhang Z, Wu J, Zheng Y, Wu X. 2025. Guiding large language models for biomedical entity linking via restrictive and contrastive decoding. *Findings of the Association for Computational Linguistics: EMNLP 2025, Suzhou, China, 2025*. Stroudsburg, PA, USA: Association for Computational Linguistics. pp. 23745–23759 doi: 10.18653/v1/2025.findings-emnlp.1292
- [47] Singhal K, Azizi S, Tu T, Mahdavi SS, Wei J, et al. 2023. Large language models encode clinical knowledge. *Nature* 620:172–180
- [48] Chen Q, Hu Y, Peng X, Xie Q, Jin Q, et al. 2025. Benchmarking large language models for biomedical natural language processing applications and recommendations. *Nature Communications* 16:3280
- [49] Galárraga L, Teflioudi C, Hose K, Suchanek FM. 2015. Fast rule mining in ontological knowledge bases with AMIE+. *VLDB Journal* 24:707–730
- [50] Meilicke C, Chekol MW, Ruffinelli D, Stuckenschmidt H. 2019. Anytime bottom-up rule learning for knowledge graph completion. *IJCAI'19: Proceedings of the 28th International Joint Conference on Artificial Intelligence*. Macao, China: AAAI Press. pp. 3137–3143 doi: 10.5555/3367471.3367477
- [51] Yang F, Yang Z, Cohen WW. 2017. Differentiable learning of logical rules for knowledge base reasoning. *arXiv* 00:1702.08367
- [52] Sadeghian A, Armandpour M, Ding P, Wang DZ. 2019. DRUM: end-to-end differentiable rule mining on knowledge graphs. *arXiv* 00:1911.00055
- [53] Yang B, Yih WT, He X, Gao J, Deng L. 2014. Embedding entities and relations for learning and inference in knowledge bases. *arXiv* 00:1412.6575
- [54] Trouillon T, Welbl J, Riedel S, Gaussier E, Bouchard G. 2016. Complex embeddings for simple link prediction. *arXiv* 00:1606.06357
- [55] Sun Z, Deng ZH, Nie JY, Tang J. 2019. RotatE: knowledge graph embedding by relational rotation in complex space. *arXiv* 00:1902.10197
- [56] Schlichtkrull M, Kipf TN, Bloem P, van den Berg R, Titov I, et al. 2018. Modeling relational data with graph convolutional networks. *The Semantic Web. Lecture Notes in Computer Science* 10843:593–607
- [57] Lao N, Mitchell T, Cohen WW. 2011. Random walk inference and learning in a large scale knowledge base. *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*. Edinburgh, United Kingdom: Association for Computational Linguistics. pp. 529–539 <https://aclanthology.org/D11-1049/>
- [58] Xiong W, Hoang T, Wang WY. 2017. DeepPath: a reinforcement learning method for knowledge graph reasoning. *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, Copenhagen, Denmark, 2017*. Stroudsburg, PA, USA: Association for Computational Linguistics. pp. 564–573 doi: 10.18653/v1/D17-1060
- [59] Teru KK, Denis E, Hamilton WL. 2019. Inductive relation prediction by subgraph reasoning. *arXiv* 00:1911.06962
- [60] Xu XH, Zhang P, He YQ, Chao CP, Yan CY. 2022. Subgraph neighboring relations infomax for inductive link prediction on knowledge graphs. *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, Vienna, Austria, 2022*. International Joint Conferences on Artificial Intelligence Organization. pp. 2341–2347 doi: 10.24963/ijcai.2022/325
- [61] Zhu ZC, Zhang ZB, Khonneux LP, Tang J. 2021. Neural Bellman-Ford Networks: a general graph neural network framework for link prediction. *arXiv* 00:2106.06935
- [62] Zhu ZC, Yuan X, Galkin M, Khonneux S, Zhang M, et al. 2022. A*Net: a scalable path-based reasoning approach for knowledge graphs. *arXiv* 00:2206.04798
- [63] Ying C, Cai T, Luo S, Zheng S, Ke G, et al. 2021. Do transformers really perform bad for graph representation? *arXiv* 00:2106.05234

- [64] Rampášek L, Galkin M, Dwivedi VP, Luu AT, Wolf G, et al. 2022. Recipe for a general, powerful, scalable graph transformer. *arXiv* 00:2205.12454
- [65] Behrouz A, Hashemi F. 2024. Graph Mamba: towards learning on graphs with state space models. *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Barcelona, Spain, 2024*. New York, NY, USA: Association for Computing Machinery. pp. 119–130 doi: [10.1145/3637528.3672044](https://doi.org/10.1145/3637528.3672044)
- [66] Yao L, Mao C, Luo Y. 2019. KG-BERT: BERT for knowledge graph completion. *arXiv* 00:1909.03193
- [67] Tang J, Yang Y, Wei W, Shi L, Su L, et al. 2024. GraphGPT: graph instruction tuning for large language models. *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval, Washington, DC, USA, 2024*. New York, NY, USA: Association for Computing Machinery. pp. 491–500 doi: [10.1145/3626772.3657775](https://doi.org/10.1145/3626772.3657775)
- [68] Sun J, Xu C, Tang L, Wang S, Lin C, et al. 2023. Think-on-Graph: deep and responsible reasoning of large language model on knowledge graph. *arXiv* 00:2307.07697
- [69] Ma J, Gao ZT, Chai Q, Sun WC, Wang PH, et al. 2025. Debate on Graph: a flexible and reliable reasoning framework for large language models. *Proceedings of the AAAI Conference on Artificial Intelligence* 39:24768–24776
- [70] Jumper J, Evans R, Pritzel A, Green T, Figurnov M, et al. 2021. Highly accurate protein structure prediction with AlphaFold. *Nature* 596:583–589
- [71] Baek M, DiMaio F, Anishchenko I, Dauparas J, Ovchinnikov S, et al. 2021. Accurate prediction of protein structures and interactions using a three-track neural network. *Science* 373:871–876
- [72] Lin Z, Akin H, Rao R, Hie B, Zhu Z, et al. 2023. Evolutionary-scale prediction of atomic-level protein structure with a language model. *Science* 379:1123–1130
- [73] Abramson J, Adler J, Dunger J, Evans R, Green T, et al. 2024. Accurate structure prediction of biomolecular interactions with AlphaFold 3. *Nature* 630:493–500
- [74] Le Guilloux V, Schmidtke P, Tuffery P. 2009. Fpocket: an open source platform for ligand pocket detection. *BMC Bioinformatics* 10:168
- [75] Halgren TA. 2009. Identifying and characterizing binding sites and assessing drugability. *Journal of Chemical Information and Modeling* 49:377–389
- [76] Schmidtke P, Bidon-Chanal A, Luque FJ, Barril X. 2011. MDpocket: open-source cavity detection and characterization on molecular dynamics trajectories. *Bioinformatics* 27:3276–3285
- [77] Bowman GR, Geissler PL. 2012. Equilibrium fluctuations of a single folded protein reveal a multitude of potential cryptic allosteric sites. *Proceedings of the National Academy of Sciences of the United States of America* 109:11681–11686
- [78] Zhang Y, Li Z, Duan B, Qin L, Peng J. 2022. MKGE: knowledge graph embedding with molecular structure information. *Computational Biology and Chemistry* 100:107730
- [79] Daza D, Alivanistos D, Mitra P, Pijnenburg T, Cochez M, et al. 2023. BioBLP: a modular framework for learning on multimodal biomedical knowledge graphs. *Journal of Biomedical Semantics* 14:20
- [80] Kuzmanic A, Bowman GR, Juarez-Jimenez J, Michel J, Gervasio FL. 2020. Investigating cryptic binding sites by molecular dynamics simulations. *Accounts of Chemical Research* 53:654–661
- [81] Hu W, Fey M, Zitnik M, Dong Y, Ren H, et al. 2020. Open graph benchmark: datasets for machine learning on graphs. *arXiv* 00:2005.00687
- [82] Toutanova K, Chen D. 2015. Observed versus latent features for knowledge base and text inference. *Proceedings of the 3rd Workshop on Continuous Vector Space Models and their Compositionality, Beijing, China, 2015*. Beijing, China: Association for Computational Linguistics. pp. 57–66 doi: [10.18653/v1/W15-4007](https://doi.org/10.18653/v1/W15-4007)
- [83] Balazevic I, Allen C, Hospedales T. 2019. TuckER: tensor factorization for knowledge graph completion. *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), Hong Kong, China, 2019*. PA, USA: Association for Computational Linguistics. pp. 5185–5194 doi: [10.18653/v1/D19-1522](https://doi.org/10.18653/v1/D19-1522)
- [84] Vashishth S, Sanyal S, Nitin V, Talukdar PP. 2019. Composition-based multi-relational graph convolutional networks. *arXiv* 00:1911.03082
- [85] Chen S, Liu X, Gao J, Jiao J, Zhang R, et al. 2021. HittER: hierarchical transformers for knowledge graph embeddings. *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, Punta Cana, Dominican Republic, 2021*. Stroudsburg, PA, USA: Association for Computational Linguistics. pp. 10395–10407 doi: [10.18653/v1/2021.emnlp-main.812](https://doi.org/10.18653/v1/2021.emnlp-main.812)
- [86] Liu J, Mao Q, Jiang W, Li J. 2024. KnowFormer: revisiting transformers for knowledge graph reasoning. *arXiv* 00:2409.12865
- [87] Wang L, Zhao W, Wei Z, Liu J. 2022. SimKGC: simple contrastive knowledge graph completion with pre-trained language models. *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), Dublin, Ireland, 2022*. Stroudsburg, PA, USA: Association for Computational Linguistics. pp. 4281–4294 doi: [10.18653/v1/2022.acl-long.295](https://doi.org/10.18653/v1/2022.acl-long.295)
- [88] Liu Y, Cao Y, Lin X, Shang Y, Wang S, et al. 2025. Enhancing large language model for knowledge graph completion via structure-aware alignment-tuning. *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing, Suzhou, China, 2025*. Stroudsburg, PA, USA: Association for Computational Linguistics. pp. 20970–20984 doi: [10.18653/v1/2025.emnlp-main.1061](https://doi.org/10.18653/v1/2025.emnlp-main.1061)
- [89] Hou Y, Zhou S, Liang K, Meng L, Chen X, et al. 2025. Soft Reasoning Paths for Knowledge Graph Completion. *arXiv* 00:2505.03285
- [90] Himmelstein DS, Lizée A, Hessler C, Brueggeman L, Chen SL, et al. 2017. Systematic integration of biomedical knowledge prioritizes drugs for repurposing. *eLife* 6:e26726
- [91] Betz P, Meilicke C, Stuckenschmidt H. 2022. Supervised knowledge aggregation for knowledge graph completion. In *The Semantic Web*, eds Groth P, Vidal ME, Suchanek F, Szekely P, Kapanipathi P, et al. Cham: Springer. pp. 74–92 doi: [10.1007/978-3-031-06981-9_5](https://doi.org/10.1007/978-3-031-06981-9_5)
- [92] Das R, Dhuliawala S, Zaheer M, Vilnis L, Durugkar I, et al. 2017. Go for a walk and arrive at the answer: reasoning over paths in knowledge bases using reinforcement learning. *arXiv* 00:1711.05851
- [93] Liu Y, Hildebrandt M, Joblin M, Ringsquandl M, Raissouni R, et al. Neural Multi-hop Reasoning with Logical Rules on Biomedical Knowledge Graphs. *The Semantic Web*, eds Verborgh R, Hose K, Paulheim H, Champin PA, Maleshkova M, et al. Cham: Springer. pp. 375–391 doi: [10.1007/978-3-030-77385-4_22](https://doi.org/10.1007/978-3-030-77385-4_22)
- [94] Liben-Nowell D, Kleinberg J. 2003. The link prediction problem for social networks. *CIKM '03: Proceedings of the twelfth international conference on Information and knowledge management*. New Orleans, LA, USA: Association for Computing Machinery. pp. 556–559 doi: [10.1145/956863.956972](https://doi.org/10.1145/956863.956972)
- [95] Zhou T, Lü L, Zhang YC. 2009. Predicting missing links via local information. *The European Physical Journal B* 71:623–630
- [96] Zhang M, Chen Y. 2018. Link prediction based on graph neural networks. *arXiv* 00:1802.09691
- [97] Song X, Ma R, Li J, Zhang M, Wipf DP. 2021. Network in graph neural network. *arXiv* 00:2111.11638
- [98] Wang X, Yang H, Zhang M. 2023. Neural common neighbor with completion for link prediction. *arXiv* 00:2302.00890
- [99] Yin T, Chen S, Wu X, Tian W. 2017. GenePANDA—a novel network-based gene prioritizing tool for complex diseases. *Scientific Reports* 7:43258
- [100] Ghiassian SD, Menche J, Barabási AL. 2015. A DIseASE MOdule detection (DIAMOnD) algorithm derived from a systematic analysis of connectivity patterns of disease proteins in the human interactome. *PLoS Computational Biology* 11:e1004120
- [101] Wang X, Yang K, Jia T, Gu F, Wang C, et al. 2024. KDGene: knowledge graph completion for disease gene prediction using interactional tensor decomposition. *Briefings in Bioinformatics* 25:bbae161
- [102] Xie J, Rao J, Xie J, Zhao H, Yang Y. 2024. Predicting disease-gene associations through self-supervised mutual infomax graph convolution network. *Computers in Biology and Medicine* 170:108048
- [103] Kumar P, Metzger VT, Purushotham ST, Kedia P, Bologna CG, et al. 2025. KG2ML: integrating knowledge graphs and positive unlabeled learning for identifying disease-associated genes. *Frontiers in Bioinformatics* 5:1727953

- [104] Liu C, Xiao K, Yu C, Lei Y, Lyu K, et al. 2024. A probabilistic knowledge graph for target identification. *PLoS Computational Biology* 20:e1011945
- [105] Shen X, Yan S, Zeng T, Xia F, Jiang D, et al. 2025. TarKGC: a target identification tool using semantics-enhanced knowledge graph completion with application to CDK2 inhibitor discovery. *Journal of Medicinal Chemistry* 68:1793–1809
- [106] Mountjoy E, Schmidt EM, Carmona M, Schwartzentruber J, Peat G, et al. 2021. An open approach to systematically prioritize causal variants and genes at all published human GWAS trait-associated loci. *Nature Genetics* 53:1527–1533
- [107] Pacini C, Duncan E, Goncalves E, Gilbert J, Bhosle S, et al. 2024. A comprehensive clinically informed map of dependencies in cancer cells and framework for target prioritization. *Cancer Cell* 42:301–316.e9
- [108] Quan Y, Xiong ZK, Zhang KX, Zhang QY, Zhang W, et al. 2023. Evolution-strengthened knowledge graph enables predicting the targetability and druggability of genes. *PNAS Nexus* 2:pgad147
- [109] Ye Q, Zeng Y, Jiang L, Kang Y, Pan P, et al. 2025. A knowledge-guided graph learning approach bridging phenotype- and target-based drug discovery. *Advanced Science* 12:e2412402
- [110] Xu J, Yu C, Xu J, Torvik VI, Kang J, et al. 2025. PubMed knowledge graph 2.0: connecting papers, patents, and clinical trials in biomedical science. *Scientific Data* 12:1018
- [111] Hopkins AL, Groom CR. 2002. The druggable genome. *Nature Reviews Drug Discovery* 1:727–730
- [112] Zhou Y, Zhang Y, Zhao D, Yu X, Shen X, et al. 2024. TTD: *Therapeutic Target Database* describing target druggability information. *Nucleic Acids Research* 52:D1465–D1477
- [113] Wagner JR, Lee CT, Durrant JD, Malmstrom RD, Feher VA, et al. 2016. Emerging computational methods for the rational discovery of allosteric drugs. *Chemical Reviews* 116:6370–6390
- [114] Berman HM, Westbrook J, Feng Z, Gilliland G, Bhat TN, et al. 2000. The protein data bank. *Nucleic Acids Research* 28:235–242
- [115] Volkamer A, Kuhn D, Rippmann F, Rarey M. 2012. DoGSiteScorer: a web server for automatic binding site prediction, analysis and drug-gability assessment. *Bioinformatics* 28:2074–2075
- [116] Raies A, Tulodziecka E, Stainer J, Middleton L, Dhindsa RS, et al. 2022. DrugnomeAI is an ensemble machine-learning framework for predicting druggability of candidate drug targets. *Communications Biology* 5:1291
- [117] Borhani N, Izadi I, Motahharynia A, Sheikholeslami M, Gheisari Y. 2025. DrugTar improves druggability prediction by integrating large language models and gene ontologies. *Bioinformatics* 41:btaf360
- [118] Chen A, Zhu X, Jiang H, Gong M, Cui S, et al. 2025. Combination of silybin and carvedilol synergistically alleviates liver fibrosis by inhibiting Wnt/ β -catenin signaling. *Targetome* 1:e009
- [119] Wang B, Mezlini AM, Demir F, Fiume M, Tu Z, et al. 2014. Similarity network fusion for aggregating data types on a genomic scale. *Nature Methods* 11:333–337
- [120] Chen YX, Chen H, Rong Y, Jiang F, Chen JB, et al. 2020. An integrative multi-omics network-based approach identifies key regulators for breast cancer. *Computational and Structural Biotechnology Journal* 18:2826–2835
- [121] Tripathy RK, Frohock Z, Wang H, Cary GA, Keegan S, et al. 2025. Effective integration of multi-omics with prior knowledge to identify biomarkers via explainable graph neural networks. *npj Systems Biology and Applications* 11:43
- [122] Hu EY, Oleshko S, Firmani S, Cheng H, Zhu Z, et al. 2026. Enhancing link prediction in biomedical knowledge graphs with BioPathNet. *Nature Biomedical Engineering*
- [123] Lan Y, He S, Liu K, Zeng X, Liu S, et al. 2021. Path-based knowledge reasoning with textual semantic information for medical knowledge graph completion. *BMC Medical Informatics and Decision Making* 21:335
- [124] Messina F, Giombini E, Agrati C, Vairo F, Ascoli Bartoli T, et al. 2020. COVID-19: viral-host interactome analyzed by network based-approach model to study pathogenesis of SARS-CoV-2 infection. *Journal of Translational Medicine* 18:233
- [125] Sun H, Song Z, Chen Q, Wang M, Tang F, et al. 2023. MMiKG: a knowledge graph-based platform for path mining of microbiota-mental diseases interactions. *Brief Bioinform* 24:bbad340
- [126] Goetz SL, Glen AK, Glusman G. 2025. MicrobiomeKG: bridging microbiome research and host health through knowledge graphs. *Frontiers in Systems Biology* 5:154432
- [127] Öztürk H, Özgür A, Ozkirimli E. 2018. DeepDTA: deep drug-target binding affinity prediction. *Bioinformatics* 34:i821–i829
- [128] Karimi M, Wu D, Wang Z, Shen Y. 2019. DeepAffinity: interpretable deep learning of compound-protein affinity through unified recurrent and convolutional neural networks. *Bioinformatics* 35:3329–3338
- [129] Nguyen T, Le H, Quinn TP, Nguyen T, Le TD, et al. 2021. GraphDTA: predicting drug-target binding affinity with graph neural networks. *Bioinformatics* 37:1140–1147
- [130] Luo Y, Zhao X, Zhou J, Yang J, Zhang Y, et al. 2017. A network integration approach for drug-target interaction prediction and computational drug repositioning from heterogeneous information. *Nature Communications* 8:573
- [131] Mohamed SK, Novacek V, Nounu A. 2020. Discovering protein drug targets using knowledge graph embeddings. *Bioinformatics* 36:603–610
- [132] Luo Y, Duan G, Zhao Q, Bi X, Wang J. 2024. DTKGIN: predicting drug-target interactions based on knowledge graph and intent graph. *Methods* 226:21–27
- [133] Zhang C, Zang T, Zhao T. 2024. KGE-UNIT: toward the unification of molecular interactions prediction based on knowledge graph and multi-task learning on drug discovery. *Briefings in Bioinformatics* 25:bbae043
- [134] Zhao X, Wang Q, Zhang Y, He C, Yin M, et al. 2025. CBKG-DTI: multi-level knowledge distillation and biomedical knowledge graph for drug-target interaction prediction. *IEEE Journal of Biomedical and Health Informatics* 29:2284–2296
- [135] Mou J, Yan Y, Jiang B, Yang F, Pan Z, et al. 2026. LKE-DTA: predicting drug-target binding affinity with large language model representations and knowledge graph embeddings. *Molecular Diversity* 30:2465–2478
- [136] Ranjan A, Kumar H, Kumari D, Anand A, Misra R. 2023. Molecule generation toward target protein (SARS-CoV-2) using reinforcement learning-based graph neural network via knowledge graph. *Network Modeling and Analysis in Health Informatics and Bioinformatics* 12:13
- [137] Yu JL, Zhou C, Ning XL, Mou J, Meng FB, et al. 2025. Knowledge-guided diffusion model for 3D ligand-pharmacophore mapping. *Nature Communications* 16:2269
- [138] To VT, Van Nguyen PC, Truong GB, Phan TM, Phan TL, et al. 2025. KGG: knowledge-guided graph self-supervised learning to enhance molecular property predictions. *Journal of Chemical Information and Modeling* 65:9443–9458
- [139] Segall MD. 2012. Multi-parameter optimization: identifying high quality compounds with a balance of properties. *Current Pharmaceutical Design* 18:1292–1310
- [140] Chen XH, Ruan Y, Liu YG, Duan XY, Jiang F, et al. 2023. Transporter proteins knowledge graph construction and its application in drug development. *Computational and Structural Biotechnology Journal* 21:2973–2984
- [141] Verhellen J. 2022. Graph-based molecular Pareto optimisation. *Chemical Science* 13:7526–7535
- [142] Yi JC, Yang ZY, Zhao WT, Yang ZJ, Zhang XC, et al. 2024. ChemMORT: an automatic ADMET optimization platform using deep learning and multi-objective particle swarm optimization. *Briefings in Bioinformatics* 25:bbae008
- [143] Paliwal S, de Giorgio A, Neil D, Michel JB, Lacoste AM. 2020. Preclinical validation of therapeutic targets predicted by tensor factorization on heterogeneous graphs. *Scientific Reports* 10:18250
- [144] Han S, Lee JE, Kang S, So M, Jin H, et al. 2024. Standigm ASK™: knowledge graph and artificial intelligence platform applied to target discovery in idiopathic pulmonary fibrosis. *Briefings in Bioinformatics* 25:bbae035
- [145] Richardson P, Griffin I, Tucker C, Smith D, Oechsle O, et al. 2020. Baricitinib as potential treatment for 2019-nCoV acute respiratory disease. *The Lancet* 395:e30–e31

- [146] Marconi VC, Ramanan AV, de Bono S, Kartman CE, Krishnan V, et al. 2021. Efficacy and safety of baricitinib for the treatment of hospitalised adults with COVID-19 (COV-BARRIER): a randomised, double-blind, parallel-group, placebo-controlled phase 3 trial. *Lancet Respiratory Medicine* 9:1407–1418
- [147] Kalil AC, Patterson TF, Mehta AK, Tomashek KM, Wolfe CR, et al. 2021. Baricitinib plus remdesivir for hospitalized adults with Covid-19. *New England Journal of Medicine* 384:795–807
- [148] Potter E, Reyes M, Naples J, Dal Pan G. 2025. FDA adverse event reporting system (FAERS) essentials: a guide to understanding, applying, and interpreting adverse event data reported to FAERS. *Journal of Clinical Pharmacology & Therapeutics* 118:567–582
- [149] Kuhn M, Letunic I, Jensen LJ, Bork P. 2016. The SIDER database of drugs and side effects. *Nucleic Acids Research* 44:D1075–D1079
- [150] Davis AP, Wiegors TC, Johnson RJ, Sciaky D, Wiegors J, et al. 2023. Comparative toxicogenomics database (CTD): update 2023. *Nucleic Acids Research* 51:D1257–D1262
- [151] Dix DJ, Houck KA, Martin MT, Richard AM, Setzer RW, et al. 2007. The ToxCast program for prioritizing toxicity testing of environmental chemicals. *Toxicological Sciences* 95:5–12
- [152] Ochoa D, Hercules A, Carmona M, Suveges D, Gonzalez-Uriarte A, et al. 2021. Open Targets Platform: supporting systematic drug-target identification and prioritisation. *Nucleic Acids Research* 49:D1302–D1310
- [153] Bean DM, Wu H, Iqbal E, Dzahini O, Ibrahim ZM, et al. 2017. Knowledge graph prediction of unknown adverse drug reactions and validation in electronic health records. *Scientific Reports* 7:16416
- [154] Zhang F, Sun B, Diao X, Zhao W, Shu T. 2021. Prediction of adverse drug reactions based on knowledge graph embedding. *BMC Medical Informatics and Decision Making* 21:38
- [155] Yao J, Sun W, Jian Z, Wu Q, Wang X. 2022. Effective knowledge graph embeddings based on multidirectional semantics relations for poly-pharmacy side effects prediction. *Bioinformatics* 38:2315–2322
- [156] Xuan P, Cheng T, Cui H, Gu J, Jin Q, et al. 2025. Multi-knowledge graph and multi-view entity feature learning for predicting drug-related side effects. *Journal of Chemical Information and Modeling* 65:5124–5138
- [157] Yadav M, Ahlawat P, Singh V. 2025. A graph-based transformer neural network for multi-label ADR prediction. *Arabian Journal for Science and Engineering* 50:7323–7341
- [158] Venkataraman M, Chand Rao G, Madavareddi JK, Maddi SR. 2025. Leveraging machine learning models in evaluating ADMET properties for drug discovery and development: review article. *ADMET and DMPK* 13:2772
- [159] Pushpakom S, Iorio F, Eyers PA, Escott KJ, Hopper S, et al. 2019. Drug repurposing: progress, challenges and recommendations. *Nature Reviews Drug Discovery* 18:41–58
- [160] Ashburn TT, Thor KB. 2004. Drug repositioning: identifying and developing new uses for existing drugs. *Nature Reviews Drug Discovery* 3:673–683
- [161] Gottlieb A, Stein GY, Ruppin E, Sharan R. 2011. PREDICT: a method for inferring novel drug indications with application to personalized medicine. *Molecular Systems Biology* 7:MSB201126
- [162] Morselli Gysi D, do Valle Í, Zitnik M, Ameli A, Gan X, et al. 2021. Network medicine framework for identifying drug-repurposing opportunities for COVID-19. *Proceedings of the National Academy of Sciences of the United States of America* 118:e2025581118
- [163] Boudin M, Diallo G, Drancé M, Mougin F. 2023. The OREGANO knowledge graph for computational drug repurposing. *Scientific Data* 10:871
- [164] Bang D, Lim S, Lee S, Kim S. 2023. Biomedical knowledge graph learning for drug repurposing by extending guilt-by-association to multiple layers. *Nature Communications* 14:3570
- [165] Huang K, Chandak P, Wang Q, Havaladar S, Vaid A, et al. 2024. A foundation model for clinician-centered drug repurposing. *Nature Medicine* 30:3601–3613
- [166] Lou P, Fang A, Zhao W, Yao K, Yang Y, et al. 2023. Potential target discovery and drug repurposing for coronaviruses: study involving a knowledge graph-based approach. *Journal of Medical Internet Research* 25:e45225
- [167] Zeng X, Song X, Ma T, Pan X, Zhou Y, et al. 2020. Repurpose open data to discover therapeutics for COVID-19 using deep learning. *Journal of Proteome Research* 19:4624–4636
- [168] Mehrabani M, Lakizadeh A, Siahpirani AF, Salimi M, Zare-Mirakabad F, et al. 2025. SynergyGraph: predicting cell line specific drug combination synergy scores using knowledge graph representation and hypergraph modeling. *Scientific Reports* 16:3131
- [169] Wang Q, Wu B, Xu M, Liu X, Mao Y, et al. 2026. HyperSynergyX: synergistic drug combination prediction via hypergraph modeling and knowledge graph-enhanced retrieval-augmented generation. *IEEE Journal of Biomedical and Health Informatics* 1–15
- [170] Lin S, Chen W, Chen G, Zhou S, Wei DQ, et al. 2022. MDDI-SCL: predicting multi-type drug-drug interactions via supervised contrastive learning. *Journal of Cheminformatics* 14:81
- [171] Wang X, Huang Y, Zhu H, Mao D, Zhu X, et al. 2026. Hgtsynergy: a transfer learning method for predicting anticancer synergistic drug combinations based on a drug-drug interaction heterogeneous graph. *BMC Bioinformatics* 27:35
- [172] Xu H, Lin J, Woicik A, Liu Z, Ma J, et al. 2025. Pisces: a multi-modal data augmentation approach for drug combination synergy prediction. *Cell Genomics* 5:100892
- [173] Yu Y, Huang K, Zhang C, Glass LM, Sun J, et al. 2021. SumGNN: multi-typed drug interaction prediction via efficient knowledge graph summarization. *Bioinformatics* 37:2988–2995
- [174] Hong Y, Luo P, Jin S, Liu X. 2022. LaGAT: link-aware graph attention network for drug-drug interaction prediction. *Bioinformatics* 38:5406–5412
- [175] Zhong J, Zhao H, Zhao Q, Wang J. 2024. A knowledge graph-based method for drug-drug interaction prediction with contrastive learning. *IEEE/ACM Transactions on Computational Biology and Bioinformatics* 21:2485–2495
- [176] Asada M, Miwa M, Sasaki Y. 2023. Integrating heterogeneous knowledge graphs into drug-drug interaction extraction from the literature. *Bioinformatics* 39:btac754
- [177] Wang M, Wang H, Liu X, Ma X, Wang B. 2021. Drug-drug interaction predictions via knowledge graph and text embedding: instrument validation study. *JMIR Medical Informatics* 9:e28277
- [178] Li Y, Liu JM, Li YC, You HR, You ZH, et al. 2026. KA-DDI: a knowledge-adaptive contrastive learning framework for drug-drug interaction prediction. *Expert Systems with Applications* 308:131076
- [179] Li D, Yang Y, Cui Z, Yin H, Hu P, et al. 2026. LLM-DDI: leveraging large language models for drug-drug interaction prediction on biomedical knowledge graph. *IEEE Journal of Biomedical and Health Informatics* 30:773–781
- [180] Putman TE, Schaper K, Matentzoglou N, Rubinetti VP, Alquaddoomi FS, et al. 2024. The Monarch Initiative in 2024: an analytic platform integrating phenotypes, genes and diseases across species. *Nucleic Acids Research* 52:D938–D949
- [181] Santos A, Colaco AR, Nielsen AB, Niu L, Strauss M, et al. 2022. A knowledge graph to interpret clinical proteomics data. *Nature Biotechnology* 40:692–702
- [182] Fang C, Arango Argoty GA, Kagiampakis I, Khalid MH, Jacob E, et al. 2024. Integrating knowledge graphs into machine learning models for survival prediction and biomarker discovery in patients with non-small-cell lung cancer. *Journal of Translational Medicine* 22:726
- [183] Ping P, Watson K, Han J, Bui A. 2017. Individualized knowledge graph: a viable informatics path to precision medicine. *Circulation Research* 120:1078–1080
- [184] Yu Y, Wang Y, Xia Z, Zhang X, Jin K, et al. 2019. PreMedKB: an integrated precision medicine knowledgebase for interpreting relationships between diseases, genes, variants and drugs. *Nucleic Acids Research* 47:D1090–D1101
- [185] Yun CH, Mengwasser KE, Toms AV, Woo MS, Greulich H, et al. 2008. The T790M mutation in EGFR kinase causes drug resistance by increasing the affinity for ATP. *Proceedings of the National Academy of Sciences of the United States of America* 105:2070–2075
- [186] Hatcher JM, Bahcall M, Choi HG, Gao Y, Sim T, et al. 2015. Discovery of inhibitors that overcome the G1202R anaplastic lymphoma kinase resistance mutation. *Journal of Medicinal Chemistry* 58:9296–9308

- [187] Gibbons DL, Priol S, Kantarjian H, Cortes J, Quintás-Cardama A. 2012. The rise and fall of gatekeeper mutations? The *BCR-ABL1* T3151 paradigm. *Cancer* 118:293–299
- [188] Naseem U, Thapa S, Zhang Q, Hu L, Masood A, et al. 2023. Reducing knowledge noise for improved semantic analysis in biomedical natural language processing applications. *Proceedings of the 5th Clinical Natural Language Processing Workshop, Toronto, Canada, 2023*. Stroudsburg, PA, USA: Association for Computational Linguistics. pp. 272–277 doi: [10.18653/v1/2023.clinicalnlp-1.32](https://doi.org/10.18653/v1/2023.clinicalnlp-1.32)
- [189] Zhu Y. Thesis proposal: uncertainty in knowledge graph embeddings. *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 4: Student Research Workshop)*, Albuquerque, USA: Association for Computational Linguistics. pp. 40–47 doi: [10.18653/v1/2025.naacl-srw.4](https://doi.org/10.18653/v1/2025.naacl-srw.4)
- [190] Lin P, Song Q, Wu Y. 2018. Fact checking in knowledge graphs with ontological subgraph patterns. *Data Science and Engineering* 3:341–358
- [191] Wang R, Zhuang C. 2025. Graph neural networks driven acceleration in drug discovery. *Acta Pharmaceutica Sinica B* 15:6163–6177
- [192] Liu C, Jiang D, Cai Y, Li H. 2026. A review of deep learning and large language models for cold start problem in recommender systems. *WIREs Data Mining and Knowledge Discovery* 16:e70068
- [193] Shi Y, Xu W, Hu P. 2025. Out of distribution learning in bioinformatics: advancements and challenges. *Briefings in Bioinformatics* 26:bbaf294
- [194] Pernice WM, Doron M, Quach A, Pratapa A, Kenjeyev S, et al. 2023. Out of distribution generalization via interventional style transfer in single-cell microscopy. *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*. Vancouver, BC, Canada, 2023. US: IEEE. pp. 4326–4335 doi: [10.1109/cvprw59228.2023.00455](https://doi.org/10.1109/cvprw59228.2023.00455)
- [195] Li H, Wang X, Zhang Z, Zhu W. 2025. OOD-GNN: out-of-distribution generalized graph neural network: (extended abstract). *2024 IEEE 40th International Conference on Data Engineering (ICDE)*. Utrecht, Netherlands, 2024. US: IEEE. pp. 5681–5682 doi: [10.1109/ICDE60146.2024.00472](https://doi.org/10.1109/ICDE60146.2024.00472)
- [196] Negro A, Montagna F, Teng MN, Neal T, Thomas S, et al. 2023. Analysis of the evolution of COVID-19 disease understanding through temporal knowledge graphs. *Frontiers in Research Metrics and Analytics* 8:1204801
- [197] Dong L, Zhao D, Zhang X, Li X, Kang X, et al. 2023. Anchors-based incremental embedding for growing knowledge graphs. *IEEE Transactions on Knowledge and Data Engineering* 35:3458–3470
- [198] McMahan HB, Moore E, Ramage D, Hampson S, Agüera y Arcas B. 2016. Communication-efficient learning of deep networks from decentralized data. *arXiv* 00:1602.05629
- [199] Chen MY, Zhang W, Yuan ZG, Jia YT, Chen HJ. 2021. FedE: embedding knowledge graphs in federated setting. *Proceedings of the 10th International Joint Conference on Knowledge Graphs, Virtual Event, Thailand, 2021*. New York, NY, USA: Association for Computing Machinery. pp. 80–88 doi: [10.1145/3502223.3502233](https://doi.org/10.1145/3502223.3502233)
- [200] Yang M, Huang D, Wan W, Jin M. 2025. Federated learning for privacy-preserving medical data sharing in drug development. *Applied and Computational Engineering* 108:7–13



Copyright: © 2026 by the author(s). Published by Maximum Academic Press on behalf of China Pharmaceutical University. This article is an open access article distributed under Creative Commons Attribution License (CC BY 4.0), visit <https://creativecommons.org/licenses/by/4.0/>.