
A Survey on Knowledge Graphs for Healthcare: Resources, Applications, and Promises

Hejie Cui¹ Jiaying Lu¹ Shiyu Wang^{1*} Ran Xu^{1*} Wenjing Ma^{1*} Shaojun Yu^{1*} Yue Yu^{2*} Xuan Kan^{1*}
Chen Ling¹ Joyce Ho¹ Fei Wang³ Carl Yang¹

Abstract

Healthcare knowledge graphs (HKGs) have emerged as a promising tool for organizing medical knowledge in a structured and interpretable way, which provides a comprehensive view of medical concepts and their relationships. However, challenges such as data heterogeneity and limited coverage remain, emphasizing the need for further research in the field of HKGs. This survey paper serves as the first comprehensive overview of HKGs. We summarize the pipeline and key techniques for HKG construction (i.e., from scratch and through integration), as well as the common utilization approaches (i.e., model-free and model-based). To provide researchers with valuable resources, we organize existing HKGs¹ based on the data types they capture and application domains, supplemented with pertinent statistical information. In the application section, we delve into the transformative impact of HKGs across various healthcare domains, spanning from fine-grained basic science research to high-level clinical decision support. Lastly, we shed light on the opportunities for creating comprehensive and accurate HKGs in the era of large language models, presenting the potential to revolutionize healthcare delivery and enhance the interpretability and reliability of clinical prediction.

1. Introduction

A knowledge graph (KG) is a data structure that captures the relationships between different entities and their at-

^{*}Equal contribution ¹Department of Computer Science, Emory University, Atlanta, GA ²School of Computational Science and Engineering, Georgia Institute of Technology, Atlanta, GA ³Department of Population Health Sciences, Weill Cornell Medicine, New York, NY, USA. Correspondence to: Carl Yang <j.carlyang@emory.edu>.

¹The resource is available at <https://github.com/lujiaying/Awesome-HealthCare-KnowledgeBase>

tributes (Ji et al., 2021; Nicholson & Greene, 2020). KG models and integrates data from various sources, including structured and unstructured data, and has been studied to support a wide range of applications such as search engines (Wang et al., 2019a), recommendation systems (Wang et al., 2019b; Zhou et al., 2020), and question answering (Lin et al., 2019; Yasunaga et al., 2021; Yan et al., 2021; Kan et al., 2021). Particularly for healthcare, KG facilitates an interpretable representation of medical concepts such as drugs and disease, which enables context-aware insights and enhances clinical research, decision-making, and healthcare delivery (Santos et al., 2022; Chandak et al., 2023).

On the data side, Healthcare knowledge graphs (HKGs) are usually built on the landscape from complex medical systems such as electronic health records, medical literature, clinical guidelines, and patient-generated data (Bouayad et al., 2017; Rajkomar et al., 2018). However, these data resources are often heterogeneous and distributed, making it challenging to integrate and analyze them effectively (Mehta & Pandit, 2018). The data heterogeneity can also lead to incomplete or inconsistent data representations within HKGs, limiting their usefulness for downstream healthcare tasks (Dash et al., 2019). Additionally, the current use of domain-specific knowledge graphs may result in limited coverage and granularity of the knowledge captured across different levels, hindering the ability to identify correlations and relationships between medical concepts from multiple domains. These challenges underscore the need for continued research on HKGs to fully realize their potential.

On the model side, HKG can be constructed from scratch or through the integration of existing dataset resources, where many key steps such as entity and relation extraction can be optimized depending on natural language processing tools and algorithms. Recent progress in general domain knowledge extraction has leveraged advances in pre-trained large language models, (i.e. BERT (Devlin et al., 2019), GPT-3 (Brown et al., 2020)). These models revolutionize the field of natural language processing, enabling the efficient and effective integration of heterogeneous medical data from various sources. The use of pre-trained models has also allowed for the development of more accurate and comprehensive

medical ontologies and taxonomies (Zhang et al., 2021a; Yu et al., 2020; Wang et al., 2021a; Zeng et al., 2021), allowing for the construction of unprecedentedly comprehensive and fine-grained KGs for healthcare (Xu et al., 2020).

A comprehensive and fine-grained healthcare knowledge graph holds the potential to revolutionize healthcare across various levels (Gyrard et al., 2018; Santos et al., 2022; Li et al., 2020). At the micro-scientific level, HKGs can help researchers identify new phenotypic and genotypic correlations and understand the underlying mechanisms of disease (Hassani-Pak & Rawlings, 2017), leading to more targeted and effective treatments (Seneviratne et al., 2021; Chandak et al., 2023). At the clinical care level, HKGs can be used to develop clinical decision support systems that provide clinicians with relevant information, improving clinical workflows and patient outcomes (Eberhardt et al., 2012; Castaneda et al., 2015). Therefore, conducting an extensive survey of the existing literature on healthcare knowledge graphs becomes an indispensable roadmap and invaluable resource for constructing a comprehensive HKG that can drive transformative advancements in healthcare.

To the best of our knowledge, this survey paper represents the first comprehensive overview of healthcare knowledge graphs (HKGs). The content overview of the survey is depicted in Figure 1, providing a visual summary of the key aspects discussed. We delve into the construction pipelines of HKGs, including both building from scratch and integration approaches, and highlight the key techniques employed in HKG construction. Additionally, we explore two common utilization methods of HKGs, namely model-free and model-based approaches (Section 2). In Section 3, we compile a comprehensive summary of existing HKG resources across various applications, serving as a valuable reference for researchers interested in utilizing or building upon HKGs. Furthermore, we meticulously investigate the literature on mainstream health applications, offering an in-depth overview of the diverse use cases of HKGs in healthcare (Section 4). Finally, we address the unique challenges associated with HKGs and discuss promising research directions, particularly in leveraging large language models to enhance their potential (Section 5). This survey paper targets a wide range of audience, including researchers, practitioners, clinicians, and other experts in healthcare, medical informatics, data science, and artificial intelligence.

2. Backgrounds

2.1. HKG Definition

A healthcare knowledge graph (HKG) is a domain-specific knowledge graph designed to capture medical concepts such as drugs, diseases, genes, phenotypes, and so on, and their relationships in a structured and semantic way.

2.2. HKG Construction

Healthcare knowledge graphs can be constructed from scratch or through the integration of existing data resources.

Constructing HKGs from Scratch. A multi-step pipeline, as in Figure 2, is used to construct HKGs from scratch.

1. The first step is to identify the scope and objectives. In most cases, researchers develop a schema (Guarino et al., 2009; Blagec et al., 2022) or use existing schemas (Guarino et al., 2009; Ashburner et al., 2000; Schriml et al., 2012; Bard et al., 2005) to serve as the formal and explicit specification of a domain, thus ensuring consistent, coherent, and aligned domain knowledge. Unlike the general domain KG, utilizing schemas is a common practice in HKG construction.
2. Secondly, researchers gather data from various sources, including medical literature, clinical trials, and patient-generated data. It's essential to ensure the quality and consistency of the data and to remove any identifiable information to protect patient privacy.
3. The third step is to extract and transform the data into a structured format. This step involves identifying medical entities and creating relationships between them via specialized biomedical Natural Language Processing (NLP) tools (Song et al., 2021; Xing et al., 2020; Hahn & Oleynik, 2020).
4. Next, researchers map the entities and relationships to the chosen ontologies with the help of thesauruses (Bodenreider, 2004) or terminologies (Donnelly et al., 2006; Hirsch et al., 2016). This step ensures that the knowledge graph is interoperable with other healthcare systems and facilitates data integration.
5. Until now, an initial KG has been built. The next step is to populate the KG to infer missing links between entities. This inference can be done using graph databases (Wang et al., 2020) or link prediction models (Bordes et al., 2013; Lu & Yang, 2022).
6. The final step is to continuously update and validate the KG to ensure accuracy and relevance. This step involves incorporating new data and knowledge, refining the schema, and evaluating the quality of the KG.

Constructing HKGs by Integration. Considering significant efforts have been paid to construct and curate HKGs from scratch, it is promising to integrate these data resources to avoid repetitive work. Healthcare KG integration (also called Healthcare KG fusion) refers to the processing of merging two or more HKGs into a single, more comprehensive graph (Himmelstein et al., 2017; Su et al., 2023; Youn et al., 2022). The integration process is challenging because different HKGs may use different terminologies, schemas, or data formats. To address these challenges, researchers have developed various techniques and algorithms for knowledge graph fusion, including ontology

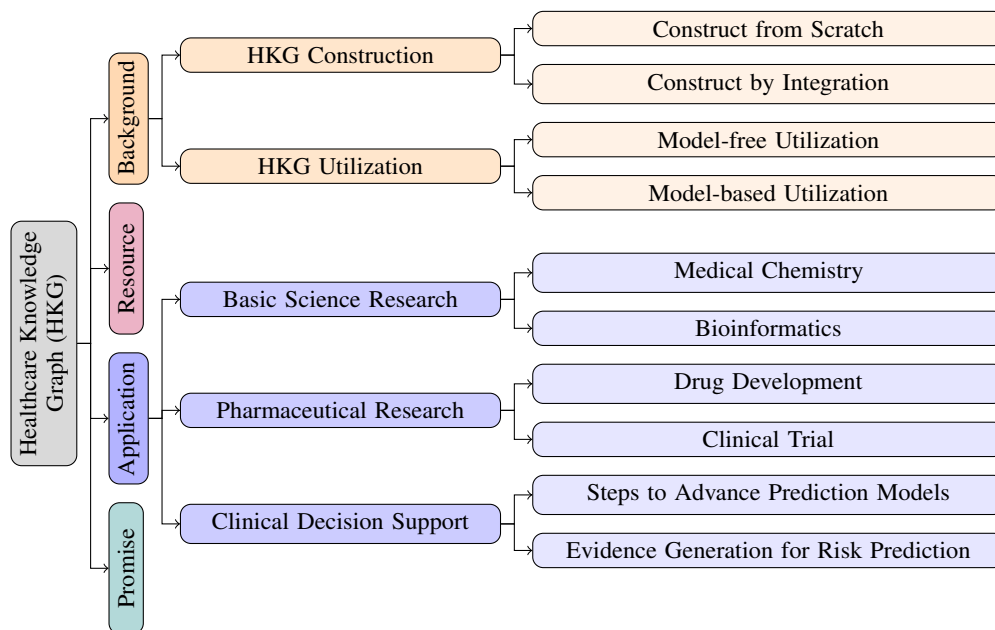


Figure 1. Overview of healthcare knowledge graph in this survey.




Figure 2. The pipeline of constructing HKGs from scratch.

matching (Faria et al., 2014; He et al., 2022), schema alignment (Suchanek et al., 2011; Maaroufi et al., 2014), entity resolution (Bachman et al., 2018; Hu et al., 2021), and conflicts resolution (Ma et al., 2023). These methods aim to identify and reconcile the differences between KGs.

Techniques for HKG Constructions. Traditionally, each step of HKG construction involves one specially designated model. For instance, Hidden Markov Models and Recurrent Neural Networks are widely used for healthcare named entity recognition, relation extraction, and other sequence tagging tasks, while Translational Models and Graph Neural Networks are used for HKG completion and conflicts resolution tasks. Recently, large language models (LLMs) have shown great utility to serve as a uniform tool for constructing KGs (Ye et al., 2022). Several key steps of constructing KGs, such as named entity recognition (Liang et al., 2020; Chen et al., 2023a; Huang et al., 2022; Liu et al., 2022a), relation extraction (Zhuang et al., 2022; Lu et al., 2022; Yang et al., 2022), entity linking (De Cao et al.; Mrini et al., 2022; Cho et al., 2022), and KG completion (Geng et al., 2022; Saxena et al., 2022; Xie et al., 2022; Shen et al., 2022), have been successfully tackled by these large foundation models. Early explorations of construction HKG with large foundation models show that healthcare entity normalization (Zhang et al., 2023; Agrawal et al., 2022),

healthcare entity recognition (Fries et al., 2021; Hu et al., 2023), healthcare entity linking (Zhu et al., 2022a), and healthcare knowledge fusion (Lu et al., 2023) can also be performed, without extensive training on expensive healthcare annotated corpus. On the other hand, researchers start to construct KGs under the open-world assumption (Shi & Weninger, 2018; Das et al., 2020; Niu et al., 2021; Li et al., 2022; Lu & Yang, 2022), thus getting rid of the dependency on pre-defined schemas and exhaustive entity&relation normalization. Although open-world KGs greatly increase the coverage, ensuring the quality of extracted knowledge is still an open research challenge, especially for explainable and trustworthy HKGs.

2.3. HKG Utilization

Model-free Utilization. Various query languages can be used for KGs, such as SPARQL, Cypher, and GraphQL (Wang et al., 2020). These query languages allow users to query healthcare KGs using a standardized syntax, thus enabling users to retrieve, manipulate, and analyze data in a structured and consistent way. More complex applications can be further supported by graph queries. For instance, automatic healthcare question answering can be tackled by Natural Language Question-to-Query (NLQ2Query) approach (Kim et al., 2022), where natural language questions are first translated into executable graph queries and then answered by the query responses. HKGs can also be utilized as an up-to-date and trustworthy augmentation to large language models (LLMs) for many applications. Some pioneering studies (Liu et al., 2022b; Guu et al., 2020; Xu et al., 2023b; Shi et al., 2023) show that retrieved knowl-

edge triples can improve the reliability of LLMs in various knowledge-intensive tasks, by addressing the nonsensical or unfaithful generation. Moreover, KGs can be a useful tool for fact-checking (Techechmedjiev et al., 2019; Vedula & Parthasarathy, 2021; Mayank et al., 2022) as they provide a structured representation of information that can be used to quickly and efficiently verify the accuracy of claims. Researchers have explored the utility of HKGs in identifying ingredient substitutions of food (Shirai et al., 2021), COVID-19 fact-checking (Mengoni & Yang, 2022), etc.

Model-based Utilization. Utilizing HKGs in complex reasoning tasks often involves utilizing machine learning models. HKG embeddings (Yu et al., 2021; Su et al., 2022) have shown great potential to tackle these tasks. In particular, HKG embedding models are a class of machine learning models that aim to learn low-dimensional vector representations, or embeddings, of the entities and relations in a knowledge graph. After obtaining HKG embeddings, they can be plugged into any kind of deep neural network and further fine-tuned toward downstream objectives. On the other hand, symbolic logic models represent another prominent approach for KG reasoning due to their interpretability. More specifically, symbolic reasoning models first mine logical rules from existing knowledge by inductive logic programming (Muggleton, 1992), association rule mining (Galárraga et al., 2013), or Markov logic networks (Kok & Domingos, 2005). These mined rules are used to infer new facts, make logical deductions and answer complex queries. Recently, researchers start to explore combining logical rules into KG embedding to further improve the generalization and performance of HKG reasoning (Alshahrani et al., 2017; Zhu et al., 2022b).

3. Resources

In this section, we aim to provide a detailed resource overview of the current state of healthcare knowledge graphs (HKGs), as a reference for researchers and healthcare professionals interested in developing and applying HKGs. A vast range of HKG resources from various domains is organized in Table 1, with attribute information including HKG name, node types, edge types, statistics, and their applications.

4. Applications

4.1. Basic Science Research

Several previous biological terms can also be considered as knowledge graphs such as ontology (e.g., gene ontology, cell ontology, disease ontology), network (e.g., gene regulatory network), etc. We use the original biological terms as they are more popular according to historical reasons.

4.1.1. MEDICINAL CHEMISTRY

Topics related to medicinal chemistry involve drug-drug interactions (DDIs) and drug-target interactions (DTIs), which will be discussed in this section.

Drug-drug interactions (DDIs) refer to changes in the actions, or side effects, of drugs when they are taken at the same time or successively (Giacomini et al., 2007). In general, DDIs are a significant contributor to life-threatening adverse events (Su et al., 2022; Pang et al., 2022; Yu et al., 2023), and their identification is one of the key tasks in public health and drug development. The existence of diverse datasets on drug-drug interactions (DDIs) and biomedical KGs has enabled the development of machine learning models that can accurately predict DDIs (Zhong et al., 2023). Yu et al. (2021) develop SumGNN, a model that includes a subgraph extraction module to efficiently extract relevant subgraphs from a KG, a self-attention-based summarization scheme to generate reasoning paths within the subgraph, and a multichannel module for integrating knowledge and data, resulting in significantly improved predictions of multi-typed DDIs. Su et al. (2022) propose DDKG, an attention-based KG representation learning framework that involves an encoder-decoder layer to learn the initial embeddings of drug nodes from their attributes in the KG. Karim et al. (2019) compare various techniques for generating KG embeddings with different settings and conclude that a combined convolutional neural network and LSTM yield the highest accuracy when predicting drug-drug interactions (DDIs). Dai et al. (2021) propose a new KG embedding framework by introducing adversarial autoencoders based on Wasserstein distances and Gumbel-Softmax relaxation for DDI tasks. Lin et al. (2020) develop KGNN that resolves the DDI prediction by capturing drug and its potential neighborhoods by mining their associated relations in KG.

Drug-target interactions (DTIs) is just as important as DDIs (Chen et al., 2016). Machine learning models can leverage knowledge graphs constructed from various types of interactions, such as drug-drug, drug-disease, protein-disease, and protein-protein interactions, to aid in the prediction of DTIs. For instance, Li et al. (2023) utilize the KG transfer probability matrix to redefine the drug-drug and target-target similarity matrix, thus constructing the final graph adjacent matrix to learn node representations by VGAE and augmenting them by utilizing dual Wasserstein Generative Adversarial Network with gradient penalty. Zhang et al. (2021c) propose a new hybrid method for DTI prediction by first constructing DTI-related KGs and then employing graph representation learning model to obtain feature vectors of the KG. Wang et al. (2022b) construct a knowledge graph of 29,607 positive drug-target pairs by DistMult embedding strategy, and propose a Conv-Conv module to extract features of drug-target pairs. Ye et al.

Table 1. Resource of Existing Healthcare Knowledge Graphs (HKGs).

Name	Node Types	Edge Types	Statistic	Application
HetioNet (Himmelstein & Baranzini, 2015)	11 (e.g., drug, disease)	24 (e.g., drug-disease)	#N: 47.0 K, #E: 2.3 M	Medicinal Chemistry
DrKG (Ioannidis et al., 2020)	13 (e.g., disease, gene)	107 (e.g., disease-gene)	#N: 97 K, #E: 5.8 M	Medicinal Chemistry
PrimeKG (Chandak et al., 2023)	10 (e.g., phenotypes)	30 (e.g., disease-phenotype)	#N: 129.4 K, #E: 8.1 M	Medicinal Chemistry
Gene Ontology ³ (Ashburner et al., 2000)	3 (e.g., biological process)	4 (e.g., partOf)	#N: 43 K, #E: 7544.6K	Bioinformatics
KEGG ³ (Kanehisa & Goto, 2000)	16 (e.g., pathway)	4 (e.g., partOf)	#N: 48.5 M, #E: unknown	Bioinformatics
STRING ⁴ (Szklarczyk et al., 2023)	1 (e.g., protein)	4 (e.g., interactions)	#N: 67.6 M, #E: 20 B	Bioinformatics
Cell Ontology ⁵ (Diehl et al., 2016)	1 (i.e., cell type)	2 (e.g., subclassOf)	#N: 2.7 K, #E: 15.9 K	Bioinformatics
GEFA (Ranjan et al., 2022)	510 (e.g., kinases)	2 (e.g., drug-drug)	#N: 0.5 K, #E: 30.1 K	Drug Development
Reaction (Li & Chen, 2022)	2 (e.g., reactant & normal)	19 (e.g., reaction paths)	#N: 2192.7 K, #E: 932.2 K	Drug Development
ASICS (Jeong et al., 2022)	2 (e.g., reactant & product)	1 (e.g., reactions)	#N: 1674.9 K, #E: 923.8 K	Drug Development
Hetionet (Jeong et al., 2022)	11 (e.g., biological process)	24 (e.g., disease-associates-gene)	#N: 47.0 K, #E: 2250.2 K	Drug Development
LBD-COVID (Zhang et al., 2021b)	1 (i.e., concept)	1 (i.e., SemMedDB relation)	#N: 131.4 K, #E: 1016.1 K	Drug Development
GP-KG (Gao et al., 2022)	7 (e.g., drug)	9 (e.g., disease-gene)	#N: 61.1 K, #E: 1246.7 K	Drug Development
DRKF (Zhang & Che, 2021)	4 (e.g., drug)	43 (e.g., drug-disease)	#N: 12.5 K, #E: 165.9 K	Drug Development
DDKG (Ghorbanali et al., 2023)	2 (i.e., drug & disease)	1 (e.g., drug-disease)	#N: 551, #E: 2.7 K	Drug Development
Disease Ontology ⁶ (Schrimal et al., 2012)	1 (i.e., disease)	2 (e.g., subclassOf)	#N: 11.2 K, #E: 8.8 K	Clinical Decision Support
DrugBank (Wishart et al., 2018)	4 (e.g., drug, pathway)	4 (e.g., drug-target)	#N: 7.4 K, #E: 366.0 K	Clinical Decision Support
KnowLife (Ernst et al., 2014)	6 (e.g., genes)	14 (e.g., gene-diseases)	#N: 2.9 M, #E: 11.4 M	Clinical Decision Support
PharmKG (Zheng et al., 2021)	3 (e.g. diseases)	3 (e.g. chemical-diseases)	#N: 7601, #E: 500958	Clinical Decision Support
ROBOKOP ⁷ (Bizon et al., 2019)	54 (e.g., genes, drugs)	1064 (e.g. biolink, CHEBI)	#N: 8.6M, #E: 130.4 M	Clinical Decision Support
iBKH ⁸ (Su et al., 2023)	11 (e.g., anatomy, disease)	18 (e.g., anatomy-gene)	#N: 2.4 M, #E: 48.2 M	Clinical Decision Support

(2021b) learn a low-dimensional representation for various entities in the KG, and then integrate the multimodal information via neural factorization machine.

4.1.2. BIOINFORMATICS RESEARCH

Knowledge graphs have been in use in Bioinformatics research for quite some time. In a Bioinformatics setting, a knowledge graph is a type of resource that represents biomedical knowledge in a structured and interconnected way. It is a graph-based representation where nodes are biomedical entities (such as mutations, genes, proteins, metabolites, diseases, and biological pathways) and edges are their relationships (such as associations, interactions, regulations)(Nicholson & Greene, 2020).

Conventional Bioinformatics Resource: While not explicitly referred to as such, we consider many conventional biomedical resources, such as Gene Ontology (Ashburner et al. (2000)), STRING (Szklarczyk et al. (2023)), KEGG (Kanehisa & Goto (2000)) can be classified as knowledge graphs. These resources have already been extensively utilized and have been shown to make significant advancements in current biomedical research. Gene Ontology (GO) is a knowledge graph that contains information about the functions of genes and their products. It is widely used in gene annotation and functional analysis. STRING is another example of a biomedical knowledge graph. It is a database that contains information about protein-protein interactions (PPIs). STRING integrates information from multiple sources, including experimental data, literature, and databases, to provide a comprehensive view of PPIs. KEGG (Kyoto Encyclopedia of Genes and Genomes) is a knowledge graph that contains information about biological pathways and networks. KEGG integrates information from multiple sources, including genomics, metabolomics,

and systems biology, to provide a comprehensive view of cellular processes.

Multi-Omics Applications: In recent years, the field of multi-omics analysis has become increasingly important for understanding biological systems, such as genomics, transcriptomics, proteomics, metabolomics, and epigenomics.

HKGs have been used to identify disease-associated mutations, genes, proteins, and metabolites by integrating multi-omics data with existing biological knowledge. This approach has led to the discovery of novel biomarkers and therapeutic targets for various diseases and interpreting the functional effects of genetic elements. Quan et al, built a comprehensive multi-relational HKG, called AIMedGraph, providing interpretation of impact of genetic variants on disease or treatment (Quan et al.). GenomicsKG (Jha et al., 2019) is an HKG to analyze and visualize multi-omics data. GenomicsKG can be used to improve drug development based on clinical genomics correlations and personalized drug customization in the extended version based on interactive relationships.

Single-Cell: Cells are fundamental and essential units of living organisms. With high-throughput sequencing technologies advancing to measure genomic profiles in a single-cell resolution, cell functions (inside cells) and cell-cell interactions (between cells) are revealed (Linnarsson & Teichmann, 2016). The gene regulatory mechanism, visualized by gene regulatory networks (GRNs), plays a crucial role in cell functions, impacting gene expression, cell differentiation, and disease progression. GRNs depict interactions between genes and their regulators and can be largely expanded by mining the whole-genome scale measurements provided by single-cell sequencing data. For example, GRNdb provides detailed regulon and TF-target pairs information from different human and mouse tissues under different conditions

by analyzing existing sequencing data (Fang et al., 2021). GenomicKB integrates existing datasets along with genome annotations and formulates the data into a KG to emphasize the relationships among genomic entries (Feng et al., 2023). On the other hand, cell-cell interactions also help understand cell cycles, cell fate decisions, tissue development, etc. Among several approaches through that cells can interact, cell-cell communication or cell signaling is of the most interest (Eltzschig et al., 2006), cells send signaling molecules called ligands and receive them through receptors located on cell surfaces. Recent advancements in spatial sequencing technologies incorporate colocalization information to better model and score ligand-receptor interactions (LRIs) as shown in recent studies (Liu et al., 2022c; Dimitrov et al., 2022). Although many LRIs databases (Efremova et al., 2020; Shao et al., 2021) have been constructed and applied to infer cell-cell communication, only recently, SpaTalk (Shao et al., 2022) integrates CellTalkDB, KEGG pathways, Reactome and TFs from AnimalTFDB to construct a ligand-receptor-target KG to help improve the inference of cell-cell communication. The success of SpaTalk indicates that correctly utilizing the KG could provide helpful information in cell-cell communication tasks. With the continuous collection of cell type information in Cell Ontology (Diehl et al., 2016), adding cell types as nodes into the biomedical KG could potentially provide biomedical researchers with cell-type-specificity and higher-resolution information as Bioteque does (Fernández-Torras et al., 2022). We hope that a comprehensive KG constructed by mining single-cell sequencing data could help revolutionize the understanding of biological mechanisms in different cells.

4.2. Pharmaceutical Research Development

4.2.1. DRUG DEVELOPMENT

Drug development identifies novel chemical compounds that can effectively treat or alleviate human diseases. Despite the growing trend of computer-assisted drug discovery (Wang et al., 2022a; Zeng et al., 2022; Pan et al., 2022a; Du et al., 2021; Wang et al., 2022c; Zhang & Zhao, 2021), there remains a key question regarding how to effectively integrate data and extract valuable insights from the vast chemical dataset. To approach this question, KGs have been employed for drug discovery due to various advantages (Zeng et al., 2022): (1) In contrast to traditional methods that capture only one type of relationship, KGs are capable of providing heterogeneous information that includes diverse entities (e.g., scaffolds, proteins and genes); (2) KGs can handle multiple types of relationships between various types of entities, such as drug-target pairs; and (3) KGs can provide unstructured semantic relationships between entities. In such graphs, entities are represented as nodes while their relationships are represented as edges, by which complex relations in biochemical systems can be easily handled.

In general, the field of drug development encompasses two main areas: *drug design* and *drug repurposing*. Drug design creates novel and diverse drug molecules with desirable pharmaceutical properties (Jing et al., 2018; Fu et al., 2021), whereas drug repurposing identifies new uses for existing approved drugs that were originally developed for a different indication (Pan et al., 2022b; Huang et al., 2020).

Drug Design. Knowledge Graphs are widely employed in drug design, particularly in the generation of novel molecules that hold promise as potential drug candidates for various diseases (Ranjan et al., 2022; Li & Chen, 2022). Ranjan et al. (2022) utilize Gated Graph Neural Network (GGNN) to generate novel molecules that target the coronavirus (i.e., SARS-CoV-2) (Hasöksüz et al., 2020) and integrate KGs into their approach to reduce the search space. Specifically, KGs were leveraged to discard non-binding molecules before inputting them into the Early Fusion model, thus optimizing the efficiency of the drug design process. In addition to employing deep learning for direct structure design, KGs are also utilized in the analysis of chemical synthesis. Quantitative estimation of molecular synthetic accessibility plays a critical role in prioritizing the molecules generated from generative models. For instance, Li & Chen (2022) utilize reaction KGs to construct classification models for compound synthetic accessibility. By leveraging KGs that capture information about reactions, including reaction types, substrates, and reaction conditions, they are able to train machine learning models that could predict the synthetic accessibility of compounds. Jeong et al. (2022) introduce an intelligent system that integrates generative exploration and exploitation of reaction knowledge base to support synthetic path design.

Drug Repurposing. Compared to drug design, KGs are more commonly utilized to expedite the drug re-purposing process (Zhu et al., 2020; MacLean, 2021; Himmelstein et al., 2017; Zhang et al., 2021b; Gao et al., 2022; Xu et al., 2019; Zhang & Che, 2021; Fang et al., 2023; Ghorbanali et al., 2023). Many applications on drug re-purposing that utilize KGs are primarily focused on link prediction tasks (MacLean, 2021). To re-purpose promising drug candidates for new indications, many methods employ predictive models that focus on predicting drug-treats-disease relationships within pharmacological knowledge graphs KGs. Himmelstein et al. (2017) use a degree-normalized pathway model on the hetionet KG which includes genes, diseases, tissues, pathophysiologies, and multimodal edges, to identify potentially repurposable drugs for epilepsy. Xu et al. (2019) develop a multi-path random walk model on a network that incorporates gene-phenotype associations, protein-protein interactions, and phenotypic similarities for training and prediction purposes. Zhang et al. (2021b) introduce an integrative and literature-based discovery model for identifying potential drug candidates from COVID-19-

focused research literature, including PubMed and other relevant sources. Gao et al. (2022) construct a knowledge graph (KG) by integrating multiple genotypic and phenotypic databases. They then learn low-dimensional representations of the KG and utilize these representations to infer new drug-disease interactions, providing insights into potential drug repurposing opportunities. Zhang & Che (2021) introduce a model for drug re-purposing in Parkinson’s disease that leverages a local medical knowledge base incorporating accurate knowledge along with medical literature containing novel information. Ghorbanali et al. (2023) present the DrugRep-KG method, which utilizes a KG embedding approach for representing drugs and diseases in the process of drug repurposing.

4.2.2. CLINICAL TRIAL

The major goal of clinical trials is to assess the safety and effectiveness of drug molecules on human bodies. A novel drug molecule needs to pass three phases of clinical trials before it is approved by Food and Drug Administration (FDA) and enters the drug market. The whole process is prohibitively time-consuming and expensive, costing 7-11 years and two billion dollars on average (Martin et al., 2017).

Clinical Trial Optimization targets on identifying eligible patients for clinical trials based on their medical history and health conditions (Rivera et al., 2020; He et al., 2020). Recently, with massive electronic health records (EHR) data and trial eligibility criteria (EC), data-driven methods have been studied to automatically assign appropriate patients for clinical trials (Yuan et al., 2019; Tseo et al., 2020; Liu et al., 2021b). However, it is often hard to fully capture and represent the complex knowledge present in unstructured ECs and EHR data, as ECs may only provide general disease concepts. In contrast, patient EHR data contain more specific medical codes to represent patient conditions. To better capture the interactions among different medical concepts from EHR records and ECs, Gao et al. (2020a) enhance patient records with hierarchical taxonomies to align medical concepts of varying granularity between EHR codes and ECs. Besides, Fu et al. (2022) leverage additional knowledge-embedding modules along with drug pharmacokinetic and historical trial data to improve the patient trial optimization process, and Wang et al. (2023) leverage the knowledge graphs to learn static trial embedding and further designed meta-learning module to generalize well over the imbalanced clinical trial distribution.

4.3. Clinical Decision Support

Nowadays, abundant Electronic Health Record (EHR) data enables better computational models for accurate diagnoses and treatments. EHR contains essential patient information such as disease diagnoses, prescribed medications, and test

results. Due to this valuable information, EHRs are extensively utilized to identify patterns in patient health and assist healthcare providers in making informed clinical decisions.

To facilitate automatic clinical predictions, various deep-learning-based approaches have been adopted including recurrent neural networks (RNN) (Choi et al., 2016; Ma et al., 2017; Yin et al., 2019; Fu et al., 2019; Raket et al., 2020; Gao et al., 2020b; Chen et al., 2019; Guo et al., 2021), graph neural networks (GNN) (Choi et al., 2020; Wang et al., 2021b; Zhu & Razavian, 2021; Xu et al., 2022; Nelson et al., 2022; Mao et al., 2022; Cai et al., 2022; Nikolentzos et al., 2023; Xu et al., 2023a), and transformers (Ma et al., 2020; Prakash et al., 2021; Antikainen et al., 2023; Labach et al., 2023; Chen et al., 2023b; Shickel et al., 2023). However, the sparsity of EHR data typically allows for only a small fraction of medical codes to be learned effectively, thereby restricting the ability of deep learning approaches. To overcome this drawback, KGs have been applied to incorporate prior medical knowledge for these deep learning models, which augment the representation of medical codes to better support the downstream tasks.

4.3.1. STEPS TO ADVANCE PREDICTION MODELS

ICD Coding aims to extract diagnosis and procedure codes from clinical notes which often consisted of raw text (Mullenbach et al., 2018; Zhang et al., 2020; Vu et al., 2021; Dong et al., 2022). It is often challenging, as the size of the candidate target codes can be large and their distribution is often long-tailed (Kim & Ganapathi, 2021). To overcome this, Xie et al. (2019) and Cao et al. (2020) propose to leverage knowledge graphs as *distant supervision* (Min et al., 2013; Zhang et al., 2022) and inject the label information via structured *knowledge graph propagation* by leveraging graph convolution networks to learn the correlations among medical codes. Besides, Lu et al. (2020) propose to leverage knowledge graphs as well as the co-occurrence graph among clinical nodes simultaneously with a knowledge aggregation module to boost the performance of ICD coding further. Overall, injecting additional knowledge with graph neural networks offers a way to mitigate the imbalanced label distribution and thus better.

Entity and Relation Extraction from Health Records.

Health records contain rich unstructured or semi-structured data, making it difficult for clinicians to analyze relevant information. Entity and relation extraction helps convert this unstructured text into structured data that can be more easily processed, understood, and utilized. Specifically, Varma et al. (2021) transfer structural knowledge from the knowledge base to the medical domain, which improved the disambiguation accuracy of rare entities. Fries et al. (2021) leverage clinical ontologies to provide *weak supervision* sources to create additional training data for clinical entity disambiguation. Yuan et al. (2023) inject additional

knowledge from the knowledge graphs for entity linking, and proposed Post-pruning and thresholding techniques to reduce the effect of unlinkable entity mentions. Besides, Fei et al. (2021); Roy & Pan (2021) propose additional post-training steps to align the language models with biomedical knowledge. Hong et al. (2021) construct embeddings for a wide range of codified concepts from EHRs to identify relevant features related to a disease of interest, and Lin et al. (2022) design a co-training scheme to jointly learn from text and knowledge graphs for extracting disease-disease relations. Fusing knowledge graphs with deep language models can flexibly accommodate missing data types and brings additional performance gains, especially for those rare entities and relations.

Clinical Report Summarization. Numerous studies have focused on transforming raw patient visit data into concise yet informative medical reports to enhance the automatic diagnosis process. Although standard models for medical report generation have achieved promising performances, there is usually no guarantee of the clinical informativeness of the generated text. To improve the faithfulness of the summarized text, Biswal et al. (2020) exploit the anchor words of relevant disease phenotypes from the external knowledge base to ensure the clinical accuracy of the generated report. Liu et al. (2021a) use an additional memory-augmented module to distill the fine-grained knowledge preserved in the knowledge graph to acquire accurate report generation. Besides, another specific issue of the clinical report is the *missing data*: some attributes are inevitably missing (Cismondi et al., 2013) when scored by domain experts. To combat the missing data issue, Xi et al. (2021) design a knowledge-aware encoder-decoder structure that injects structural information from knowledge graphs during the encoding stage and infers patients' links to clinical outcomes during the decoder stage. In summary, incorporating external knowledge graphs into clinical report summarization enhances the content's factual accuracy and alleviates the impact of missing data by grounding the generated text in verified knowledge.

4.3.2. EVIDENCE GENERATION FOR RISK PREDICTION

Disease Prediction aims to predict the potential diseases of a given patient with his past clinical records. To assist the diagnosis with additional knowledge, GRAM (Choi et al., 2017) and KAME (Ma et al., 2018) utilize a medical ontology (Dubberke et al., 2006) where the leaf nodes are the medical codes found in EHR data, and their ancestors are more general categories. By incorporating information from medical ontologies into deep learning models via neural attention, these approaches learn better embeddings for different medical concepts to alleviate the data scarcity bottleneck. Yin et al. (2019); Zhang et al. (2019) further consider the domain-specific knowledge graph KnowLife (Ernst et al.,

2015) to enrich the embeddings of medical entities with their neighbors on the knowledge graph. These approaches mainly directly update the embeddings of different concepts to improve feature learning, but ignore the high-level order information from the knowledge graph. To tackle this drawback, Ye et al. (2021a) explicitly exploit *paths* in KG from the observed symptoms to the target disease to model the personalized information for diverse patients with a relational-guided attention mechanism. Xu et al. (2021) design a self-supervised learning approach to pre-train a graph attention network for learning the embedding of medical concepts and completing the knowledge graph simultaneously. These approaches better harness the structure information, and often lead to better performance than the pure embedding-based knowledge integration techniques.

Treatment Recommendation aims to recommend personalized medications to patients based on their individual health conditions, which can help physicians to select the most effective medications for their patients, and improve treatment outcomes (Zhang et al., 2017; Bhoi et al., 2021; Shang et al., 2019b). To effectively exploit external knowledge, Shang et al. (2019a) use drug ontologies to design additional pre-training loss and directly improve the representation of drugs, and several studies (Wu et al., 2022; Tan et al., 2022; Yang et al., 2021) attempt to extract the additional drug interaction graphs to model the negative side effects of specific drug pairs and reduce the possibility of recommending negative drug-drug interaction combinations.

5. Promise and Outlook

Several promising research directions in computer science are poised to generate large and accurate healthcare knowledge graphs (HKGs) in the near future. One such direction involves the development of advanced entity and relation extraction techniques that can effectively capture and represent complex biomedical knowledge. Additionally, large language models (*a.k.a.* foundation models) like GPT-3 have demonstrated promise in capturing the semantics and context of biomedical language (Agrawal et al., 2022; Singhal et al., 2022; Nath et al., 2022), allowing researchers to better understand and interpret complex biomedical data (Moor et al., 2023). Another direction includes the use of graph-based learning paradigms capable of effectively integrating heterogeneous data sources and learning from the complex relationships within HKGs. These approaches can facilitate the creation of comprehensive and fine-grained HKGs that capture a diverse range of biomedical knowledge, ultimately promoting the interpretability of biomedical research and clinical decision-making.

The potential impact of comprehensive and fine-grained HKGs on biomedical research and clinical practice is significant. By integrating vast amounts of biomedical knowledge

from multiple domains, HKGs can facilitate the discovery of new disease mechanisms and identification of novel drug targets. They also hold the potential to enable personalized medicine by identifying patient subgroups with shared disease mechanisms and guiding the selection of targeted interventions based on individual patient characteristics. Additionally, HKGs can enhance clinical decision-making by providing access to up-to-date and pertinent biomedical knowledge, thereby improving efficiency and accuracy. Overall, the creation and utilization of HKGs present a promising avenue for promoting transparency and interpretability for clinical practice.

6. Conclusion

In conclusion, healthcare knowledge graphs (HKGs) offer a promising approach to capturing and organizing medical knowledge in a structured and interpretable way, providing a comprehensive and fine-grained view of medical concepts and relationships. Despite challenges like data heterogeneity and limited coverage, recent technical advancements have enabled the creation of comprehensive and precise HKGs. This survey provides a comprehensive overview of the current state of HKGs, covering their construction, utilization models, and applications in healthcare. We also discuss potential future developments, emphasizing the importance of HKGs in facilitating efficient and effective healthcare delivery. With the emergence of large language models, the potential for creating even more comprehensive and precise HKGs is unprecedented. In conclusion, healthcare knowledge graphs (HKGs) hold great potential in improving the interpretability of healthcare, enabling transparent and informed decision-making, and promoting evidence-based practices for better patient outcomes.

References

- Agrawal, M., Heggelmann, S., Lang, H., Kim, Y., and Son-tag, D. Large language models are few-shot clinical information extractors. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pp. 1998–2022, 2022.
- Alshahrani, M., Khan, M. A., Maddouri, O., Kinjo, A. R., Queralt-Rosinach, N., and Hoehndorf, R. Neuro-symbolic representation learning on biological knowledge graphs. *Bioinformatics*, 33(17):2723–2730, 2017.
- Antikainen, E., Linnosmaa, J., Umer, A., Oksala, N., Eskola, M., van Gils, M., Hernesniemi, J., and Gabbouj, M. Transformers for cardiac patient mortality risk prediction from heterogeneous electronic health records. *Scientific Reports*, 13(1):3517, 2023.
- Ashburner, M., Ball, C. A., Blake, J. A., Botstein, D., Butler, H., Cherry, J. M., Davis, A. P., Dolinski, K., Dwight, S. S., Eppig, J. T., et al. Gene ontology: tool for the unification of biology. *Nature genetics*, 25(1):25–29, 2000.
- Bachman, J. A., Gyori, B. M., and Sorger, P. K. Famplex: a resource for entity recognition and relationship resolution of human protein families and complexes in biomedical text mining. *BMC bioinformatics*, 19:1–14, 2018.
- Bard, J., Rhee, S. Y., and Ashburner, M. An ontology for cell types. *Genome biology*, 6(2):1–5, 2005.
- Bhoi, S., Lee, M. L., Hsu, W., Fang, H. S. A., and Tan, N. C. Personalizing medication recommendation with a graph-based approach. *ACM Transactions on Information Systems (TOIS)*, 40(3):1–23, 2021.
- Biswal, S., Xiao, C., Glass, L. M., Westover, B., and Sun, J. Clara: clinical report auto-completion. In *Proceedings of The Web Conference 2020*, pp. 541–550, 2020.
- Bizon, C., Cox, S., Balhoff, J., Kebede, Y., Wang, P., Morton, K., Fecho, K., and Tropsha, A. Robokop kg and kgb: integrated knowledge graphs from federated sources. *Journal of chemical information and modeling*, 59(12):4968–4973, 2019.
- Blagec, K., Barbosa-Silva, A., Ott, S., and Samwald, M. A curated, ontology-based, large-scale knowledge graph of artificial intelligence tasks and benchmarks. *Scientific Data*, 9(1):322, 2022.
- Bodenreider, O. The unified medical language system (umls): integrating biomedical terminology. *Nucleic acids research*, 32(suppl_1):D267–D270, 2004.
- Bordes, A., Usunier, N., Garcia-Duran, A., Weston, J., and Yakhnenko, O. Translating embeddings for modeling multi-relational data. *Advances in neural information processing systems*, 26, 2013.
- Bouayad, L., Ialynytchev, A., and Padmanabhan, B. Patient health record systems scope and functionalities: literature review and future directions. *Journal of medical Internet research*, 19(11):e388, 2017.
- Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- Cai, D., Sun, C., Song, M., Zhang, B., Hong, S., and Li, H. Hypergraph contrastive learning for electronic health records. In *Proceedings of the 2022 SIAM International Conference on Data Mining (SDM)*, pp. 127–135. SIAM, 2022.

- Cao, P., Chen, Y., Liu, K., Zhao, J., Liu, S., and Chong, W. Hypercore: Hyperbolic and co-graph representation for automatic icd coding. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 3105–3114, 2020.
- Castaneda, C., Nalley, K., Mannion, C., Bhattacharyya, P., Blake, P., Pecora, A., Goy, A., and Suh, K. S. Clinical decision support systems for improving diagnostic accuracy and achieving precision medicine. *Journal of clinical bioinformatics*, 5(1):1–16, 2015.
- Chandak, P., Huang, K., and Zitnik, M. Building a knowledge graph to enable precision medicine. *Nature Scientific Data*, 2023.
- Chen, R., Stewart, W. F., Sun, J., Ng, K., and Yan, X. Recurrent neural networks for early detection of heart failure from longitudinal electronic health record data: implications for temporal modeling with respect to time before diagnosis, data density, data quantity, and data type. *Circulation: Cardiovascular Quality and Outcomes*, 12(10):e005114, 2019.
- Chen, X., Yan, C. C., Zhang, X., Zhang, X., Dai, F., Yin, J., and Zhang, Y. Drug–target interaction prediction: databases, web servers and computational models. *Briefings in bioinformatics*, 17(4):696–712, 2016.
- Chen, X., Li, L., Fei, Q., Zhang, N., Tan, C., Jiang, Y., Huang, F., and Chen, H. One model for all domains: Collaborative domain-prefix tuning for cross-domain ner. *arXiv preprint arXiv:2301.10410*, 2023a.
- Chen, Z., Siltala-Li, L., Lassila, M., Malo, P., Vilkkumaa, E., Saaresranta, T., and Virkki, A. V. Predicting visit cost of obstructive sleep apnea using electronic healthcare records with transformer. *IEEE Journal of Translational Engineering in Health and Medicine*, 2023b.
- Cho, Y. M., Zhang, L., and Callison-Burch, C. Unsupervised entity linking with guided summarization and multiple-choice selection. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pp. 9394–9401, 2022.
- Choi, E., Bahadori, M. T., Sun, J., Kulas, J., Schuetz, A., and Stewart, W. Retain: An interpretable predictive model for healthcare using reverse time attention mechanism. *Advances in neural information processing systems*, 29, 2016.
- Choi, E., Bahadori, M. T., Song, L., Stewart, W. F., and Sun, J. Gram: graph-based attention model for healthcare representation learning. In *Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and data mining*, pp. 787–795, 2017.
- Choi, E., Xu, Z., Li, Y., Dusenberry, M., Flores, G., Xue, E., and Dai, A. Learning the graphical structure of electronic health records with graph convolutional transformer. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pp. 606–613, 2020.
- Cismondi, F., Fialho, A. S., Vieira, S. M., Reti, S. R., Sousa, J. M., and Finkelstein, S. N. Missing data in medical databases: Impute, delete or classify? *Artificial intelligence in medicine*, 58(1):63–72, 2013.
- Dai, Y., Guo, C., Guo, W., and Eickhoff, C. Drug–drug interaction prediction with wasserstein adversarial autoencoder-based knowledge graph embeddings. *Briefings in bioinformatics*, 22(4):bbaa256, 2021.
- Das, R., Godbole, A., Monath, N., Zaheer, M., and McCallum, A. Probabilistic case-based reasoning for open-world knowledge graph completion. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pp. 4752–4765, 2020.
- Dash, S., Shakyawar, S. K., Sharma, M., and Kaushik, S. Big data in healthcare: management, analysis and future prospects. *Journal of Big Data*, 6(1):1–25, 2019.
- De Cao, N., Wu, L., Papat, K., Artetxe, M., Goyal, N., Plekhanov, M., Zettlemoyer, L., Cancedda, N., Riedel, S., and Petroni, F. Multilingual autoregressive entity linking. *Transactions of the Association for Computational Linguistics*, 10:274–290.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 4171–4186, 2019.
- Diehl, A. D., Meehan, T. F., Bradford, Y. M., Brush, M. H., Dahdul, W. M., Dougall, D. S., He, Y., Osumi-Sutherland, D., Ruttenberg, A., Sarntivijai, S., et al. The cell ontology 2016: enhanced content, modularization, and ontology interoperability. *Journal of biomedical semantics*, 7:1–10, 2016.
- Dimitrov, D., Türe, D., Garrido-Rodriguez, M., Burmedi, P. L., Nagai, J. S., Boys, C., Ramirez Flores, R. O., Kim, H., Szalai, B., Costa, I. G., et al. Comparison of methods and resources for cell-cell communication inference from single-cell rna-seq data. *Nature Communications*, 13(1):3224, 2022.
- Dong, H., Falis, M., Whiteley, W., Alex, B., Matterson, J., Ji, S., Chen, J., and Wu, H. Automated clinical coding: what, why, and where we are? *NPJ digital medicine*, 5(1):159, 2022.

- Donnelly, K. et al. Snomed-ct: The advanced terminology and coding system for ehealth. *Studies in health technology and informatics*, 121:279, 2006.
- Du, Y., Wang, S., Guo, X., Cao, H., Hu, S., Jiang, J., Varala, A., Angirekula, A., and Zhao, L. Graphgt: Machine learning datasets for graph generation and transformation. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*, 2021.
- Dubberke, E. R., Reske, K. A., McDonald, L. C., and Fraser, V. J. Icd-9 codes and surveillance for clostridium difficile-associated disease. *Emerging infectious diseases*, 12(10): 1576, 2006.
- Eberhardt, J., Bilchik, A., and Stojadinovic, A. Clinical decision support systems: potential with pitfalls. *Journal of Surgical Oncology*, 105(5):502–510, 2012.
- Efremova, M., Vento-Tormo, M., Teichmann, S. A., and Vento-Tormo, R. Cellphonedb: inferring cell–cell communication from combined expression of multi-subunit ligand–receptor complexes. *Nature protocols*, 15(4): 1484–1506, 2020.
- Eltzschig, H. K., Weissmüller, T., Mager, A., and Eckle, T. Nucleotide metabolism and cell-cell interactions. *Cell-Cell Interactions: Methods and Protocols*, pp. 73–87, 2006.
- Ernst, P., Meng, C., Siu, A., and Weikum, G. Knowlife: a knowledge graph for health and life sciences. In *2014 IEEE 30th International Conference on Data Engineering*, pp. 1254–1257. IEEE, 2014.
- Ernst, P., Siu, A., and Weikum, G. Knowlife: a versatile approach for constructing a large knowledge graph for biomedical sciences. *BMC bioinformatics*, 16:1–13, 2015.
- Fang, L., Li, Y., Ma, L., Xu, Q., Tan, F., and Chen, G. Grndb: decoding the gene regulatory networks in diverse human and mouse conditions. *Nucleic acids research*, 49(D1): D97–D103, 2021.
- Fang, Y., Zhang, Q., Zhang, N., Chen, Z., Zhuang, X., Shao, X., Fan, X., and Chen, H. Knowledge graph-enhanced molecular contrastive learning with functional prompt. *Nature Machine Intelligence*, pp. 1–12, 2023.
- Faria, D., Jiménez-Ruiz, E., Pesquita, C., Santos, E., and Couto, F. M. Towards annotating potential incoherences in biportal mappings. In *The Semantic Web–ISWC 2014: 13th International Semantic Web Conference, Riva del Garda, Italy, October 19-23, 2014. Proceedings, Part II 13*, pp. 17–32. Springer, 2014.
- Fei, H., Ren, Y., Zhang, Y., Ji, D., and Liang, X. Enriching contextualized language model from knowledge graph for biomedical information extraction. *Briefings in bioinformatics*, 22(3):bbaa110, 2021.
- Feng, F., Tang, F., Gao, Y., Zhu, D., Li, T., Yang, S., Yao, Y., Huang, Y., and Liu, J. Genomickb: a knowledge graph for the human genome. *Nucleic Acids Research*, 51(D1): D950–D956, 2023.
- Fernández-Torras, A., Duran-Frigola, M., Bertoni, M., Locatelli, M., and Aloy, P. Integrating and formatting biomedical data as pre-calculated knowledge graph embeddings in the bioteque. *Nature Communications*, 13(1): 5304, 2022.
- Fries, J. A., Steinberg, E., Khattar, S., Fleming, S. L., Posada, J., Callahan, A., and Shah, N. H. Ontology-driven weak supervision for clinical entity classification in electronic health records. *Nature communications*, 12(1):2017, 2021.
- Fu, T., Hoang, T. N., Xiao, C., and Sun, J. Ddl: Deep dictionary learning for predictive phenotyping. In *28th International Joint Conference on Artificial Intelligence, IJCAI 2019*, pp. 5857–5863, 2019.
- Fu, T., Gao, W., Xiao, C., Yasonik, J., Coley, C. W., and Sun, J. Differentiable scaffolding tree for molecular optimization. *arXiv preprint arXiv:2109.10469*, 2021.
- Fu, T., Huang, K., Xiao, C., Glass, L. M., and Sun, J. Hint: Hierarchical interaction network for clinical-trial-outcome predictions. *Patterns*, 3(4):100445, 2022.
- Galárraga, L. A., Teflioudi, C., Hose, K., and Suchanek, F. Amie: association rule mining under incomplete evidence in ontological knowledge bases. In *Proceedings of the 22nd international conference on World Wide Web*, pp. 413–422, 2013.
- Gao, J., Xiao, C., Glass, L. M., and Sun, J. Compose: cross-modal pseudo-siamese network for patient trial matching. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 803–812, 2020a.
- Gao, J., Xiao, C., Wang, Y., Tang, W., Glass, L. M., and Sun, J. Stagenet: Stage-aware neural networks for health risk prediction. In *Proceedings of The Web Conference 2020*, pp. 530–540, 2020b.
- Gao, Z., Ding, P., and Xu, R. Kg-predict: a knowledge graph computational framework for drug repurposing. *Journal of biomedical informatics*, 132:104133, 2022.

- Geng, Y., Chen, J., Zhang, W., Xu, Y., Chen, Z., Z. Pan, J., Huang, Y., Xiong, F., and Chen, H. Disentangled ontology embedding for zero-shot learning. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 443–453, 2022.
- Ghorbanali, Z., Zare-Mirakabad, F., Akbari, M., Salehi, N., and Masoudi-Nejad, A. Drugrep-kg: Toward learning a unified latent space for drug repurposing using knowledge graphs. *Journal of Chemical Information and Modeling*, 2023.
- Giacomini, K. M., Krauss, R. M., Roden, D. M., Eichelbaum, M., Hayden, M. R., and Nakamura, Y. When good drugs go bad. *Nature*, 446(7139):975–977, 2007.
- Guarino, N., Oberle, D., and Staab, S. What is an ontology? *Handbook on ontologies*, pp. 1–17, 2009.
- Guo, A., Beheshti, R., Khan, Y. M., Langabeer, J. R., and Foraker, R. E. Predicting cardiovascular health trajectories in time-series electronic health records with lstm models. *BMC Medical Informatics and Decision Making*, 21:1–10, 2021.
- Guu, K., Lee, K., Tung, Z., Pasupat, P., and Chang, M. Retrieval augmented language model pre-training. In *International conference on machine learning*, pp. 3929–3938. PMLR, 2020.
- Gyrard, A., Gaur, M., Shekarpour, S., Thirunarayan, K., and Sheth, A. Personalized health knowledge graph. In *CEUR workshop proceedings*, volume 2317, 2018.
- Hahn, U. and Oleynik, M. Medical information extraction in the age of deep learning. *Yearbook of medical informatics*, 29(01):208–220, 2020.
- Hasöksüz, M., Kilic, S., and Sarac, F. Coronaviruses and sars-cov-2. *Turkish journal of medical sciences*, 50(9): 549–556, 2020.
- Hassani-Pak, K. and Rawlings, C. Knowledge discovery in biological databases for revealing candidate genes linked to complex phenotypes. *Journal of integrative bioinformatics*, 14(1), 2017.
- He, Y., Chen, J., Dong, H., Jiménez-Ruiz, E., Hadian, A., and Horrocks, I. Machine learning-friendly biomedical datasets for equivalence and subsumption ontology matching. In *The Semantic Web—ISWC 2022: 21st International Semantic Web Conference, Virtual Event, October 23–27, 2022, Proceedings*, pp. 575–591. Springer, 2022.
- He, Z., Tang, X., Yang, X., Guo, Y., George, T. J., Charness, N., Quan Hem, K. B., Hogan, W., and Bian, J. Clinical trial generalizability assessment in the big data era: a review. *Clinical and translational science*, 13(4):675–684, 2020.
- Himmelstein, D. S. and Baranzini, S. E. Heterogeneous network edge prediction: a data integration approach to prioritize disease-associated genes. *PLoS computational biology*, 11(7):e1004259, 2015.
- Himmelstein, D. S., Lizee, A., Hessler, C., Brueggeman, L., Chen, S. L., Hadley, D., Green, A., Khankhanian, P., and Baranzini, S. E. Systematic integration of biomedical knowledge prioritizes drugs for repurposing. *Elife*, 6: e26726, 2017.
- Hirsch, J., Nicola, G., McGinty, G., Liu, R., Barr, R., Chittle, M., and Manchikanti, L. Icd-10: history and context. *American Journal of Neuroradiology*, 37(4):596–599, 2016.
- Hong, C., Rush, E., Liu, M., Zhou, D., Sun, J., Sonabend, A., Castro, V. M., Schubert, P., Panickan, V. A., Cai, T., et al. Clinical knowledge extraction via sparse embedding regression (keser) with multi-center large scale electronic health record data. *NPJ digital medicine*, 4(1):151, 2021.
- Hu, Y., Nie, T., Shen, D., Kou, Y., and Yu, G. An integrated pipeline model for biomedical entity alignment. *Frontiers of Computer Science*, 15:1–15, 2021.
- Hu, Y., Ameer, I., Zuo, X., Peng, X., Zhou, Y., Li, Z., Li, Y., Li, J., Jiang, X., and Xu, H. Zero-shot clinical entity recognition using chatgpt. *arXiv preprint arXiv:2303.16416*, 2023.
- Huang, K., Fu, T., Glass, L. M., Zitnik, M., Xiao, C., and Sun, J. Deeppurpose: a deep learning library for drug–target interaction prediction. *Bioinformatics*, 36(22-23): 5545–5547, 2020.
- Huang, Y., He, K., Wang, Y., Zhang, X., Gong, T., Mao, R., and Li, C. Copner: Contrastive learning with prompt guiding for few-shot named entity recognition. In *Proceedings of the 29th International conference on computational linguistics*, pp. 2515–2527, 2022.
- Ioannidis, V. N., Song, X., Manchanda, S., Li, M., Pan, X., Zheng, D., Ning, X., Zeng, X., and Karypis, G. Drkg - drug repurposing knowledge graph for covid-19. <https://github.com/gnn4dr/DRKG/>, 2020.
- Jeong, J., Lee, N., Shin, Y., and Shin, D. Intelligent generation of optimal synthetic pathways based on knowledge graph inference and retrosynthetic predictions using reaction big data. *Journal of the Taiwan Institute of Chemical Engineers*, 130:103982, 2022.

- Jha, A., d'Aquin, M., Sahay, R., Zehra, D., Dangwal, S., Khan, Y., Mehmood, Q., Rebholz-Schuhman, D., and Verma, G. Genomicskg: A knowledge graph to visualize poly-omics data. *Journal of Advances in Health*, 01(02):70–84, 2019. doi: <https://doi.org/10.3724/SP.J.2640-8686.2019.0063>.
- Ji, S., Pan, S., Cambria, E., Marttinen, P., and Philip, S. Y. A survey on knowledge graphs: Representation, acquisition, and applications. *IEEE transactions on neural networks and learning systems*, 33(2):494–514, 2021.
- Jing, Y., Bian, Y., Hu, Z., Wang, L., and Xie, X.-Q. S. Deep learning for drug design: an artificial intelligence paradigm for drug discovery in the big data era. *The AAPS journal*, 20:1–10, 2018.
- Kan, X., Cui, H., and Yang, C. Zero-shot scene graph relation prediction through commonsense knowledge integration. In *Machine Learning and Knowledge Discovery in Databases. Research Track: European Conference, ECML PKDD 2021, Bilbao, Spain, September 13–17, 2021, Proceedings, Part II 21*, pp. 466–482. Springer, 2021.
- Kanehisa, M. and Goto, S. Kegg: kyoto encyclopedia of genes and genomes. *Nucleic acids research*, 28(1):27–30, 2000.
- Karim, M. R., Cochez, M., Jares, J. B., Uddin, M., Beyan, O., and Decker, S. Drug-drug interaction prediction based on knowledge graph embeddings and convolutional-lstm network. In *Proceedings of the 10th ACM international conference on bioinformatics, computational biology and health informatics*, pp. 113–123, 2019.
- Kim, B.-H. and Ganapathi, V. Read, attend, and code: Pushing the limits of medical codes prediction from clinical notes by machines. In *Proceedings of the 6th Machine Learning for Healthcare Conference*, volume 149 of *Proceedings of Machine Learning Research*, pp. 196–208, 2021.
- Kim, D., Bae, S., Kim, S., and Choi, E. Uncertainty-aware text-to-program for question answering on structured electronic health records. In Flores, G., Chen, G. H., Pollard, T., Ho, J. C., and Naumann, T. (eds.), *Proceedings of the Conference on Health, Inference, and Learning*, volume 174 of *Proceedings of Machine Learning Research*, pp. 138–151. PMLR, 07–08 Apr 2022.
- Kok, S. and Domingos, P. Learning the structure of markov logic networks. In *Proceedings of the 22nd international conference on Machine learning*, pp. 441–448, 2005.
- Labach, A., Pokhrel, A., Huang, X. S., Zuberi, S., Yi, S. E., Volkovs, M., Poutanen, T., and Krishnan, R. G. Duett: Dual event time transformer for electronic health records. *arXiv preprint arXiv:2304.13017*, 2023.
- Li, B. and Chen, H. Prediction of compound synthesis accessibility based on reaction knowledge graph. *Molecules*, 27(3):1039, 2022.
- Li, G., Sun, W., Xu, J., Hu, L., Zhang, W., and Zhang, P. Ga-ens: A novel drug–target interactions prediction method by incorporating prior knowledge graph into dual wasserstein generative adversarial network with gradient penalty. *Applied Soft Computing*, 139:110151, 2023.
- Li, L., Wang, P., Yan, J., Wang, Y., Li, S., Jiang, J., Sun, Z., Tang, B., Chang, T.-H., Wang, S., et al. Real-world data medical knowledge graph: construction and applications. *Artificial intelligence in medicine*, 103:101817, 2020.
- Li, Q., Wang, D., Song, S. F. K., Zhang, Y., and Yu, G. Oerl: Enhanced representation learning via open knowledge graphs. *IEEE Transactions on Knowledge and Data Engineering*, pp. 1–14, 2022. doi: 10.1109/TKDE.2022.3218850.
- Liang, C., Yu, Y., Jiang, H., Er, S., Wang, R., Zhao, T., and Zhang, C. Bond: Bert-assisted open-domain named entity recognition with distant supervision. In *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining*, pp. 1054–1064, 2020.
- Lin, B. Y., Chen, X., Chen, J., and Ren, X. Kagnet: Knowledge-aware graph networks for commonsense reasoning. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pp. 2829–2839, 2019.
- Lin, X., Quan, Z., Wang, Z.-J., Ma, T., and Zeng, X. Kgnn: Knowledge graph neural network for drug-drug interaction prediction. In *IJCAI*, volume 380, pp. 2739–2745, 2020.
- Lin, Y., Lu, K., Yu, S., Cai, T., and Zitnik, M. Multimodal learning on graphs for disease relation extraction. *arXiv preprint arXiv:2203.08893*, 2022.
- Linnarsson, S. and Teichmann, S. A. Single-cell genomics: coming of age, 2016.
- Liu, A. T., Xiao, W., Zhu, H., Zhang, D., Li, S.-W., and Arnold, A. Qaner: Prompting question answering models for few-shot named entity recognition. *arXiv preprint arXiv:2203.01543*, 2022a.
- Liu, F., You, C., Wu, X., Ge, S., Sun, X., et al. Auto-encoding knowledge graph for unsupervised medical report generation. *Advances in Neural Information Processing Systems*, 34:16266–16279, 2021a.

- Liu, Q., Yogatama, D., and Blunsom, P. Relational memory-augmented language models. *Transactions of the Association for Computational Linguistics*, 10:555–572, 2022b.
- Liu, R., Rizzo, S., Whipple, S., Pal, N., Pineda, A. L., Lu, M., Arneri, B., Lu, Y., Capra, W., Copping, R., et al. Evaluating eligibility criteria of oncology trials using real-world data and ai. *Nature*, 592(7855):629–633, 2021b.
- Liu, Z., Sun, D., and Wang, C. Evaluation of cell-cell interaction methods by integrating single-cell rna sequencing data with spatial information. *Genome Biology*, 23(1):1–38, 2022c.
- Lu, J. and Yang, C. Open-world taxonomy and knowledge graph co-learning. In *4th Conference on Automated Knowledge Base Construction*, AKBC 2022, Nov. 2022.
- Lu, J., Du, L., Liu, M., and Dipnall, J. Multi-label few/zero-shot learning with knowledge aggregated from multiple label graphs. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 2935–2943, 2020.
- Lu, J., Shen, J., Xiong, B., Ma, W., Steffen, S., and Yang, C. Hiprompt: Few-shot biomedical knowledge fusion via hierarchy-oriented prompting. In *Proceedings of The 46th International ACM SIGIR Conference on Research and Development in Information Retrieval - Short Paper*, SIGIR 2023, 2023.
- Lu, Y., Liu, Q., Dai, D., Xiao, X., Lin, H., Han, X., Sun, L., and Wu, H. Unified structure generation for universal information extraction. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 5755–5772, 2022.
- Ma, F., Chitta, R., Zhou, J., You, Q., Sun, T., and Gao, J. Dipole: Diagnosis prediction in healthcare via attention-based bidirectional recurrent neural networks. In *Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and data mining*, pp. 1903–1911, 2017.
- Ma, F., You, Q., Xiao, H., Chitta, R., Zhou, J., and Gao, J. Kame: Knowledge-based attention model for diagnosis prediction in healthcare. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, pp. 743–752, 2018.
- Ma, J., Zhou, C., Chen, Y., Wang, Y., Hu, G., and Qiao, Y. Tetre: A novel temporal conflict resolution method based on temporal knowledge graph embedding. *Information*, 14(3):155, 2023.
- Ma, L., Zhang, C., Wang, Y., Ruan, W., Wang, J., Tang, W., Ma, X., Gao, X., and Gao, J. Concare: Personalized clinical feature embedding via capturing the healthcare context. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pp. 833–840, 2020.
- Maaroufi, M., Choquet, R., Landais, P., and Jaulent, M.-C. Formalizing mappings to optimize automated schema alignment: application to rare diseases. In *e-Health-For Continuity of Care*, pp. 283–287. IOS Press, 2014.
- MacLean, F. Knowledge graphs and their applications in drug discovery. *Expert opinion on drug discovery*, 16(9):1057–1069, 2021.
- Mao, C., Yao, L., and Luo, Y. Medgcn: Medication recommendation and lab test imputation via graph convolutional networks. *Journal of Biomedical Informatics*, 127:104000, 2022.
- Martin, L., Hutchens, M., Hawkins, C., and Radnov, A. How much do clinical trials cost. *Nat Rev Drug Discov*, 16(6):381–382, 2017.
- Mayank, M., Sharma, S., and Sharma, R. Deap-faked: Knowledge graph based approach for fake news detection. In *2022 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, pp. 47–51. IEEE, 2022.
- Mehta, N. and Pandit, A. Concurrence of big data analytics and healthcare: A systematic review. *International journal of medical informatics*, 114:57–65, 2018.
- Mengoni, P. and Yang, J. Empowering covid-19 fact-checking with extended knowledge graphs. In *Computational Science and Its Applications-ICCSA 2022 Workshops: Malaga, Spain, July 4-7, 2022, Proceedings, Part I*, pp. 138–150. Springer, 2022.
- Min, B., Grishman, R., Wan, L., Wang, C., and Gondek, D. Distant supervision for relation extraction with an incomplete knowledge base. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 777–782, 2013.
- Moor, M., Banerjee, O., Abad, Z. S. H., Krumholz, H. M., Leskovec, J., Topol, E. J., and Rajpurkar, P. Foundation models for generalist medical artificial intelligence. *Nature*, 616(7956):259–265, 2023.
- Mrini, K., Nie, S., Gu, J., Wang, S., Sanjabi, M., and Firooz, H. Detection, disambiguation, re-ranking: Autoregressive entity linking as a multi-task problem. In *Findings of the Association for Computational Linguistics: ACL 2022*, pp. 1972–1983, 2022.
- Muggleton, S. *Inductive logic programming*. Number 38. Morgan Kaufmann, 1992.

- Mullenbach, J., Wiegrefe, S., Duke, J., Sun, J., and Eisenstein, J. Explainable prediction of medical codes from clinical text. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pp. 1101–1111, 2018.
- Nath, S., Marie, A., Ellershaw, S., Korot, E., and Keane, P. A. New meaning for nlp: the trials and tribulations of natural language processing with gpt-3 in ophthalmology. *British Journal of Ophthalmology*, 106(7):889–892, 2022.
- Nelson, C. A., Bove, R., Butte, A. J., and Baranzini, S. E. Embedding electronic health records onto a knowledge network recognizes prodromal features of multiple sclerosis and predicts diagnosis. *Journal of the American Medical Informatics Association*, 29(3):424–434, 2022.
- Nicholson, D. N. and Greene, C. S. Constructing knowledge graphs and their biomedical applications. *Computational and structural biotechnology journal*, 18:1414–1428, 2020.
- Nikolentzos, G., Vazirgiannis, M., Xypolopoulos, C., Lingman, M., and Brandt, E. G. Synthetic electronic health records generated with variational graph autoencoders. *NPJ Digital Medicine*, 6(1):83, 2023.
- Niu, L., Fu, C., Yang, Q., Li, Z., Chen, Z., Liu, Q., and Zheng, K. Open-world knowledge graph completion with multiple interaction attention. *World Wide Web*, 24: 419–439, 2021.
- Pan, B., Wang, Y., Lin, X., Qin, M., Du, Y., Ghaemi, S., Ding, A., Wang, S., Alkhalifa, S., Minbiole, K., et al. Property-controllable generation of quaternary ammonium compounds. In *2022 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, pp. 3462–3469. IEEE, 2022a.
- Pan, X., Lin, X., Cao, D., Zeng, X., Yu, P. S., He, L., Nussinov, R., and Cheng, F. Deep learning for drug repurposing: Methods, databases, and applications. *Wiley interdisciplinary reviews: Computational molecular science*, 12(4):e1597, 2022b.
- Pang, S., Zhang, Y., Song, T., Zhang, X., Wang, X., and Rodriguez-Patón, A. Amde: a novel attention-mechanism-based multidimensional feature encoder for drug–drug interaction prediction. *Briefings in Bioinformatics*, 23(1):bbab545, 2022.
- Prakash, P., Chilukuri, S., Ranade, N., and Viswanathan, S. Rarebert: Transformer architecture for rare disease patient identification using administrative claims. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pp. 453–460, 2021.
- Quan, X., Cai, W., Xi, C., Wang, C., and Yan, L. AIMed-Graph: A comprehensive multi-relational knowledge graph for precision medicine. 2023:baad006. ISSN 1758-0463.
- Rajkomar, A., Oren, E., Chen, K., Dai, A. M., Hajaj, N., Hardt, M., Liu, P. J., Liu, X., Marcus, J., Sun, M., et al. Scalable and accurate deep learning with electronic health records. *NPJ digital medicine*, 1(1):18, 2018.
- Raket, L. L., Jaskolowski, J., Kinon, B. J., Brasen, J. C., Jönsson, L., Wehnert, A., and Fusar-Poli, P. Dynamic electronic health record detection (detect) of individuals at risk of a first episode of psychosis: a case-control development and validation study. *The Lancet Digital Health*, 2(5):e229–e239, 2020.
- Ranjan, A., Shukla, S., Datta, D., and Misra, R. Generating novel molecule for target protein (sars-cov-2) using drug–target interaction based on graph neural network. *Network Modeling Analysis in Health Informatics and Bioinformatics*, 11:1–11, 2022.
- Rivera, S. C., Liu, X., Chan, A.-W., Denniston, A. K., Calvert, M. J., Ashrafiyan, H., Beam, A. L., Collins, G. S., Darzi, A., Deeks, J. J., et al. Guidelines for clinical trial protocols for interventions involving artificial intelligence: the spirit-ai extension. *The Lancet Digital Health*, 2(10):e549–e560, 2020.
- Roy, A. and Pan, S. Incorporating medical knowledge in bert for clinical relation extraction. In *Proceedings of the 2021 conference on empirical methods in natural language processing*, pp. 5357–5366, 2021.
- Santos, A., Colaço, A. R., Nielsen, A. B., Niu, L., Strauss, M., Geyer, P. E., Coscia, F., Albrechtsen, N. J. W., Mundt, F., Jensen, L. J., et al. A knowledge graph to interpret clinical proteomics data. *Nature Biotechnology*, 40(5): 692–702, 2022.
- Saxena, A., Kochsiek, A., and Gemulla, R. Sequence-to-sequence knowledge graph completion and question answering. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 2814–2828, 2022.
- Schriml, L. M., Arze, C., Nadendla, S., Chang, Y.-W. W., Mazaitis, M., Felix, V., Feng, G., and Kibbe, W. A. Disease ontology: a backbone for disease semantic integration. *Nucleic acids research*, 40(D1):D940–D946, 2012.
- Seneviratne, O., Harris, J., Chen, C.-H., and McGuinness, D. L. Personal health knowledge graph for clinically relevant diet recommendations. *arXiv preprint arXiv:2110.10131*, 2021.

- Shang, J., Ma, T., Xiao, C., and Sun, J. Pre-training of graph augmented transformers for medication recommendation. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*, pp. 5953–5959, 2019a.
- Shang, J., Xiao, C., Ma, T., Li, H., and Sun, J. Gamenet: Graph augmented memory networks for recommending medication combination. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 1126–1133, 2019b.
- Shao, X., Liao, J., Li, C., Lu, X., Cheng, J., and Fan, X. Celltalkdb: a manually curated database of ligand–receptor interactions in humans and mice. *Briefings in bioinformatics*, 22(4):bbaa269, 2021.
- Shao, X., Li, C., Yang, H., Lu, X., Liao, J., Qian, J., Wang, K., Cheng, J., Yang, P., Chen, H., et al. Knowledge-graph-based cell-cell communication inference for spatially resolved transcriptomic data with spatalk. *Nature Communications*, 13(1):4429, 2022.
- Shen, J., Wang, C., Yuan, Y., Han, J., Ji, H., Sen, K., Zhang, M., and Song, D. PALT: Parameter-lite transfer of language models for knowledge graph completion. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pp. 3833–3847, 2022.
- Shi, B. and Weninger, T. Open-world knowledge graph completion. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32, 2018.
- Shi, W., Min, S., Yasunaga, M., Seo, M., James, R., Lewis, M., Zettlemoyer, L., and Yih, W.-t. Replug: Retrieval-augmented black-box language models. *arXiv preprint arXiv:2301.12652*, 2023.
- Shickel, B., Loftus, T. J., Ren, Y., Rashidi, P., Bihorac, A., and Ozrazgat-Baslanti, T. Digital health transformers and opportunities for artificial intelligence-enabled nephrology. *Clinical Journal of the American Society of Nephrology*, 18(4):527–529, 2023.
- Shirai, S. S., Seneviratne, O., Gordon, M. E., Chen, C.-H., and McGuinness, D. L. Identifying ingredient substitutions using a knowledge graph of food. *Frontiers in Artificial Intelligence*, 3:621766, 2021.
- Singhal, K., Azizi, S., Tu, T., Mahdavi, S. S., Wei, J., Chung, H. W., Scales, N., Tanwani, A., Cole-Lewis, H., Pfohl, S., et al. Large language models encode clinical knowledge. *arXiv preprint arXiv:2212.13138*, 2022.
- Song, B., Li, F., Liu, Y., and Zeng, X. Deep learning methods for biomedical named entity recognition: a survey and qualitative comparison. *Briefings in Bioinformatics*, 22(6):bbab282, 2021.
- Su, C., Hou, Y., Zhou, M., Rajendran, S., Maasch, J. R., Abedi, Z., Zhang, H., Bai, Z., Cuturrufo, A., Guo, W., et al. Biomedical discovery through the integrative biomedical knowledge hub (ibkh). *Iscience*, 2023.
- Su, X., Hu, L., You, Z., Hu, P., and Zhao, B. Attention-based knowledge graph representation learning for predicting drug-drug interactions. *Briefings in bioinformatics*, 23(3):bbac140, 2022.
- Suchanek, F. M., Abiteboul, S., and Senellart, P. Paris: probabilistic alignment of relations, instances, and schema. *Proceedings of the VLDB Endowment*, 5(3):157–168, 2011.
- Szklarczyk, D., Kirsch, R., Koutrouli, M., Nastou, K., Mehryary, F., Hachilif, R., Gable, A. L., Fang, T., Doncheva, N. T., Pyysalo, S., et al. The string database in 2023: protein–protein association networks and functional enrichment analyses for any sequenced genome of interest. *Nucleic Acids Research*, 51(D1):D638–D646, 2023.
- Tan, Y., Kong, C., Yu, L., Li, P., Chen, C., Zheng, X., Hertzberg, V. S., and Yang, C. 4sdrug: Symptom-based set-to-set small and safe drug recommendation. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 3970–3980, 2022.
- Tchechmedjiev, A., Fafalios, P., Boland, K., Gasquet, M., Zloch, M., Zapilko, B., Dietze, S., and Todorov, K. Claim-skg: A knowledge graph of fact-checked claims. In *The Semantic Web–ISWC 2019: 18th International Semantic Web Conference, Auckland, New Zealand, October 26–30, 2019, Proceedings, Part II 18*, pp. 309–324. Springer, 2019.
- Tseo, Y., Salkola, M., Mohamed, A., Kumar, A., and Abnoui, F. Information extraction of clinical trial eligibility criteria. *arXiv preprint arXiv:2006.07296*, 2020.
- Varma, M., Orr, L., Wu, S., Leszczynski, M., Ling, X., and Ré, C. Cross-domain data integration for named entity disambiguation in biomedical text. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pp. 4566–4575, 2021.
- Vedula, N. and Parthasarathy, S. Face-keg: Fact checking explained using knowledge graphs. In *Proceedings of the 14th ACM International Conference on Web Search and Data Mining*, pp. 526–534, 2021.
- Vu, T., Nguyen, D. Q., and Nguyen, A. A label attention model for icd coding from clinical text. In *Proceedings of the Twenty-Ninth International Joint Conferences on Artificial Intelligence*, pp. 3335–3341, 2021.

- Wang, P., Jiang, H., Xu, J., and Zhang, Q. Knowledge graph construction and applications for web search and beyond. *Data Intelligence*, 1(4):333–349, 2019a.
- Wang, P., Shi, T., and Reddy, C. K. Text-to-sql generation for question answering on electronic medical records. In *Proceedings of The Web Conference 2020*, pp. 350–361, 2020.
- Wang, S., Zhao, R., Chen, X., Zheng, Y., and Liu, B. Enquire one’s parent and child before decision: Fully exploit hierarchical structure for self-supervised taxonomy expansion. In *Proceedings of the Web Conference 2021*, pp. 3291–3304, 2021a.
- Wang, S., Du, Y., Guo, X., Pan, B., Qin, Z., and Zhao, L. Controllable data generation by deep learning: A review. *arXiv preprint arXiv:2207.09542*, 2022a.
- Wang, S., Du, Z., Ding, M., Rodriguez-Paton, A., and Song, T. Kg-dti: a knowledge graph based deep learning method for drug-target interaction predictions and alzheimer’s disease drug repositions. *Applied Intelligence*, 52(1): 846–857, 2022b.
- Wang, S., Guo, X., Lin, X., Pan, B., Du, Y., Wang, Y., Ye, Y., Petersen, A., Leitgeb, A., AlKhalifa, S., et al. Multi-objective deep data generation with correlated property control. *Advances in Neural Information Processing Systems*, 35:28889–28901, 2022c.
- Wang, X., He, X., Cao, Y., Liu, M., and Chua, T.-S. Kgat: Knowledge graph attention network for recommendation. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, pp. 950–958, 2019b.
- Wang, Z., Wen, R., Chen, X., Cao, S., Huang, S.-L., Qian, B., and Zheng, Y. Online disease diagnosis with inductive heterogeneous graph convolutional networks. In *Proceedings of the Web Conference 2021*, pp. 3349–3358, 2021b.
- Wang, Z., Xiao, C., and Sun, J. Spot: Sequential predictive modeling of clinical trial outcome with meta-learning. *arXiv preprint arXiv:2304.05352*, 2023.
- Wishart, D. S., Feunang, Y. D., Guo, A. C., Lo, E. J., Marcu, A., Grant, J. R., Sajed, T., Johnson, D., Li, C., Sayeeda, Z., et al. Drugbank 5.0: a major update to the drugbank database for 2018. *Nucleic acids research*, 46(D1): D1074–D1082, 2018.
- Wu, J., Qian, B., Li, Y., Gao, Z., Ju, M., Yang, Y., Zheng, Y., Gong, T., Li, C., and Zhang, X. Leveraging multiple types of domain knowledge for safe and effective drug recommendation. In *Proceedings of the 31st ACM international conference on information & knowledge management*, pp. 2169–2178, 2022.
- Xi, J., Ye, L., Huang, Q., and Li, X. Tolerating data missing in breast cancer diagnosis from clinical ultrasound reports via knowledge graph inference. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pp. 3756–3764, 2021.
- Xie, X., Xiong, Y., Yu, P. S., and Zhu, Y. Ehr coding with multi-scale feature attention and structured knowledge graph propagation. In *Proceedings of the 28th ACM international conference on information and knowledge management*, pp. 649–658, 2019.
- Xie, X., Zhang, N., Li, Z., Deng, S., Chen, H., Xiong, F., Chen, M., and Chen, H. From discrimination to generation: knowledge graph completion with generative transformer. In *Companion Proceedings of the Web Conference 2022*, pp. 162–165, 2022.
- Xing, R., Luo, J., and Song, T. Biorel: towards large-scale biomedical relation extraction. *BMC bioinformatics*, 21: 1–13, 2020.
- Xu, B., Liu, Y., Yu, S., Wang, L., Dong, J., Lin, H., Yang, Z., Wang, J., and Xia, F. A network embedding model for pathogenic genes prediction by multi-path random walking on heterogeneous network. *BMC Medical Genomics*, 12:1–12, 2019.
- Xu, J., Kim, S., Song, M., Jeong, M., Kim, D., Kang, J., Rousseau, J. F., Li, X., Xu, W., Torvik, V. I., et al. Building a pubmed knowledge graph. *Scientific data*, 7(1):205, 2020.
- Xu, R., Yu, Y., Zhang, C., Ali, M. K., Ho, J. C., and Yang, C. Counterfactual and factual reasoning over hypergraphs for interpretable clinical predictions on ehr. In *Proceedings of the 2nd Machine Learning for Health symposium*, volume 193 of *Proceedings of Machine Learning Research*, pp. 259–278, 2022.
- Xu, R., Ali, M. K., Ho, J. C., and Yang, C. Hypergraph transformers for ehr-based clinical predictions. In *Proceedings of AMIA Informatics Summit*, 2023a.
- Xu, R., Yu, Y., Ho, J. C., and Yang, C. Weakly-supervised scientific document classification via retrieval-augmented multi-stage training. In *the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2023b.
- Xu, X., Xu, X., Sun, Y., Liu, X., Li, X., Xie, G., and Wang, F. Predictive modeling of clinical events with mutual enhancement between longitudinal patient records and medical knowledge graph. In *2021 IEEE International Conference on Data Mining (ICDM)*, pp. 777–786. IEEE, 2021.

- Yan, J., Raman, M., Chan, A., Zhang, T., Rossi, R., Zhao, H., Kim, S., Lipka, N., and Ren, X. Learning contextualized knowledge structures for commonsense reasoning. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pp. 4038–4051, 2021.
- Yang, C., Xiao, C., Ma, F., Glass, L., and Sun, J. Safe-drug: Dual molecular graph encoders for recommending effective and safe drug combinations. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence*, pp. 3735–3741, 2021.
- Yang, X., Chen, A., PourNejatian, N., Shin, H. C., Smith, K. E., Parisien, C., Compas, C., Martin, C., Costa, A. B., Flores, M. G., et al. A large language model for electronic health records. *npj Digital Medicine*, 5(1):194, 2022.
- Yasunaga, M., Ren, H., Bosselut, A., Liang, P., and Leskovec, J. Qa-gnn: Reasoning with language models and knowledge graphs for question answering. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 535–546, 2021.
- Ye, H., Zhang, N., Chen, H., and Chen, H. Generative knowledge graph construction: A review. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pp. 1–17, 2022.
- Ye, M., Cui, S., Wang, Y., Luo, J., Xiao, C., and Ma, F. Medpath: Augmenting health risk prediction via medical knowledge paths. In *Proceedings of the Web Conference 2021*, pp. 1397–1409, 2021a.
- Ye, Q., Hsieh, C.-Y., Yang, Z., Kang, Y., Chen, J., Cao, D., He, S., and Hou, T. A unified drug–target interaction prediction framework based on knowledge graph and recommendation system. *Nature communications*, 12(1): 6775, 2021b.
- Yin, C., Zhao, R., Qian, B., Lv, X., and Zhang, P. Domain knowledge guided deep learning with electronic health records. In *2019 IEEE International Conference on Data Mining (ICDM)*, pp. 738–747. IEEE, 2019.
- Youn, J., Rai, N., and Tagkopoulos, I. Knowledge integration and decision support for accelerated discovery of antibiotic resistance genes. *Nature Communications*, 13(1):2360, 2022.
- Yu, H., Li, K., Dong, W., Song, S., Gao, C., and Shi, J. Attention-based cross domain graph neural network for prediction of drug–drug interactions. *Briefings in Bioinformatics*, pp. bbad155, 2023.
- Yu, Y., Li, Y., Shen, J., Feng, H., Sun, J., and Zhang, C. Steam: Self-supervised taxonomy expansion with mini-paths. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 1026–1035, 2020.
- Yu, Y., Huang, K., Zhang, C., Glass, L. M., Sun, J., and Xiao, C. Sumgnn: multi-typed drug interaction prediction via efficient knowledge graph summarization. *Bioinformatics*, 37(18):2988–2995, 2021.
- Yuan, C., Ryan, P. B., Ta, C., Guo, Y., Li, Z., Hardin, J., Makadia, R., Jin, P., Shang, N., Kang, T., et al. Criteria2query: a natural language interface to clinical databases for cohort definition. *Journal of the American Medical Informatics Association*, 26(4):294–305, 2019.
- Yuan, H., Lu, K., and Yuan, Z. Exploring partial knowledge base inference in biomedical entity linking. *arXiv preprint arXiv:2303.10330*, 2023.
- Zeng, Q., Lin, J., Yu, W., Cleland-Huang, J., and Jiang, M. Enhancing taxonomy completion with concept generation via fusing relational representations. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pp. 2104–2113, 2021.
- Zeng, X., Tu, X., Liu, Y., Fu, X., and Su, Y. Toward better drug discovery with knowledge graph. *Current opinion in structural biology*, 72:114–126, 2022.
- Zhang, J., Song, X., Zeng, Y., Chen, J., Shen, J., Mao, Y., and Li, L. Taxonomy completion via triplet matching network. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pp. 4662–4670, 2021a.
- Zhang, J., Hsieh, C.-Y., Yu, Y., Zhang, C., and Ratner, A. A survey on programmatic weak supervision. *arXiv preprint arXiv:2202.05433*, 2022.
- Zhang, J., Wang, Z., Zhang, S., Bhalerao, M. M., Liu, Y., Zhu, D., and Wang, S. Graphprompt: Biomedical entity normalization using graph-based prompt templates. In *The 37th AAAI Conference on Artificial Intelligence*, 2023.
- Zhang, R., Hristovski, D., Schutte, D., Kastrin, A., Fiszman, M., and Kilicoglu, H. Drug repurposing for covid-19 via knowledge graph completion. *Journal of biomedical informatics*, 115:103696, 2021b.
- Zhang, S., Lin, X., and Zhang, X. Discovering dti and ddi by knowledge graph with mhrw and improved neural network. In *2021 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, pp. 588–593. IEEE, 2021c.
- Zhang, X. and Che, C. Drug repurposing for parkinson’s disease by integrating knowledge graph completion model and knowledge fusion of medical literature. *Future Internet*, 13(1):14, 2021.

- Zhang, X., Qian, B., Li, Y., Yin, C., Wang, X., and Zheng, Q. Knowrisk: an interpretable knowledge-guided model for disease risk prediction. In *2019 IEEE International Conference on Data Mining (ICDM)*, pp. 1492–1497. IEEE, 2019.
- Zhang, Y., Chen, R., Tang, J., Stewart, W. F., and Sun, J. Leap: learning to prescribe effective and safe treatment combinations for multimorbidity. In *proceedings of the 23rd ACM SIGKDD international conference on knowledge Discovery and data Mining*, pp. 1315–1324, 2017.
- Zhang, Z. and Zhao, L. Representation learning on spatial networks. *Advances in Neural Information Processing Systems*, 34:2303–2318, 2021.
- Zhang, Z., Liu, J., and Razavian, N. Bert-xml: Large scale automated icd coding using bert pretraining. In *Proceedings of the 3rd Clinical Natural Language Processing Workshop*, pp. 24–34, 2020.
- Zheng, S., Rao, J., Song, Y., Zhang, J., Xiao, X., Fang, E. F., Yang, Y., and Niu, Z. Pharmkg: a dedicated knowledge graph benchmark for biomedical data mining. *Briefings in bioinformatics*, 22(4):bbaa344, 2021.
- Zhong, Z., Barkova, A., and Mottin, D. Knowledge-augmented graph machine learning for drug discovery: A survey from precision to interpretability. *arXiv preprint arXiv:2302.08261*, 2023.
- Zhou, S., Dai, X., Chen, H., Zhang, W., Ren, K., Tang, R., He, X., and Yu, Y. Interactive recommender system via knowledge graph-enhanced reinforcement learning. In *Proceedings of the 43rd international ACM SIGIR conference on research and development in information retrieval*, pp. 179–188, 2020.
- Zhu, T., Qin, Y., Chen, Q., Hu, B., and Xiang, Y. Enhancing entity representations with prompt learning for biomedical entity linking. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI 2022*, pp. 4036–4042, 2022a.
- Zhu, W. and Razavian, N. Variationally regularized graph-based representation learning for electronic health records. In *Proceedings of the Conference on Health, Inference, and Learning*, pp. 1–13, 2021.
- Zhu, Y., Che, C., Jin, B., Zhang, N., Su, C., and Wang, F. Knowledge-driven drug repurposing using a comprehensive drug knowledge graph. *Health Informatics Journal*, 26(4):2737–2750, 2020.
- Zhu, Z., Galkin, M., Zhang, Z., and Tang, J. Neural-symbolic models for logical queries on knowledge graphs. In *International Conference on Machine Learning*, pp. 27454–27478. PMLR, 2022b.
- Zhuang, Y., Li, Y., Zhang, J., Yu, Y., Mou, Y., Chen, X., Song, L., and Zhang, C. ReSel: N-ary relation extraction from scientific text and tables by learning to retrieve and select. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pp. 730–744, December 2022.