

Representation Learning for Person or Entity-Centric Knowledge Graphs: An Application in Healthcare

Christos Theodoropoulos*
IBM Research Europe
Dublin, Ireland
KU Leuven
Leuven, Belgium
christos.theodoropoulos@kuleuven.be

Natasha Mulligan
IBM Research Europe
Dublin, Ireland
natasha.mulligan@ie.ibm.com

Thaddeus Stappenbeck
Lerner Research Institute, Cleveland Clinic
Cleveland, Ohio, United States
stappet@ccf.org

Joao Bettencourt-Silva
IBM Research Europe
Dublin, Ireland
jbettencourt@ie.ibm.com

ABSTRACT

Knowledge graphs (KGs) are a popular way to organise information based on ontologies or schemas. Despite advances in KGs, representing knowledge remains a non-trivial task across industries and it is especially challenging in the biomedical and healthcare domains due to complex interdependent relations between entities, heterogeneity, lack of standardization, and sparseness of data. KGs are used to discover diagnoses or prioritize genes relevant to disease, but they often rely on schemas that are not centred around a node or entity of interest, such as a person. Entity-centric KGs are relatively unexplored but hold promise in representing important facets connected to a central node and unlocking downstream tasks beyond graph traversal and reasoning, such as training graph neural networks (GNNs) for a wide range of predictive tasks. This paper presents an end-to-end representation learning framework to extract entity-centric KGs from structured and unstructured data. We introduce a star-shaped ontology to represent the multiple facets of a person and use it to guide KG creation. Compact representations of the graphs are created leveraging GNNs and experiments are conducted using different levels of heterogeneity or explicitness. A readmission prediction task is used to evaluate the results of the proposed framework, showing a stable system, robust to missing data, that outperforms a range of baseline machine learning classifiers. We highlight that this approach has several potential applications across domains and is open-sourced.

CCS CONCEPTS

• **Information systems** → **Information integration; Data structures**; • **Computing methodologies** → **Knowledge representation and reasoning; Machine learning**.

*Work completed while the first author was an intern at IBM Research Europe.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

K-CAP '23, December 05–07, 2023, Pensacola, Florida, USA

© 2023 Association for Computing Machinery.
ACM ISBN 978-1-4503-XXXX-X/18/06... \$15.00
<https://doi.org/XXXXXXXX.XXXXXXX>

KEYWORDS

Entity-Centric Knowledge Graphs, Person-Centric Ontology, Representation Learning, Graph Neural Networks

ACM Reference Format:

Christos Theodoropoulos, Natasha Mulligan, Thaddeus Stappenbeck, and Joao Bettencourt-Silva. 2023. Representation Learning for Person or Entity-Centric Knowledge Graphs: An Application in Healthcare. In *Proceedings of 12th International Conference on Knowledge Capture (K-CAP '23)*. ACM, New York, NY, USA, 8 pages. <https://doi.org/XXXXXXXX.XXXXXXX>

1 INTRODUCTION

Knowledge graphs (KGs) have been widely used to organize information in a structured and flexible way enabling a variety of downstream tasks and applications [28]. KGs consist of nodes (entities) and edges (relations) between them, that represent the information in a particular domain or set of domains. The ability of KGs to support complex reasoning and inference has been explored in a variety of tasks including search [33, 49], recommendation [25, 67], and knowledge discovery [37, 47].

KGs are becoming increasingly used across a wide range of healthcare applications. Knowledge graphs typically rely on information retrieved from biomedical literature and early approaches, such as BioGrid [15], have first depended on manual curation of knowledge bases to map protein-to-protein interactions using genes and proteins as entity types. Semi-automatic and machine-learning approaches have been introduced to assist, for example, in finding associations or relations between important concepts such as diseases and symptoms [53]. Recently, healthcare and clinical applications have otherwise focused on building KGs from electronic health records (EHRs) where nodes represent diseases, drugs, or patients, and edges represent their relations [48]. Most approaches are however limited by the challenging nature of healthcare data, including heterogeneity, sparseness and inconsistent or lacking standardization [5]. Recent calls for future work on KGs have expressed the need to develop new models and algorithms that are able to take these challenges (e.g. missing data) into account [48]. Furthermore, there is a need to be able to accurately represent information from multiple data sources about individual patients. This is not only required by physicians to support routine hospital activities but also for clinical research in, for example, developing novel

predictive models or discovering new important features associated with poor patient outcomes. A holistic representation of individual patients should therefore capture not only their clinical attributes such as diagnoses, procedures, and medication but also other variables that may also be predictors such as demographics, behavioral and social aspects (e.g. smoking habits or unemployment). Incorporating new types of information in models and analyses will allow the creation of better tools for evaluating the effectiveness of therapies and directing them to the most relevant patients.

We present an end-to-end representation learning framework to extract information and organize it into entity-centric KGs from both structured and unstructured data. A star-shaped ontology is designed and used for the purpose of representing the multiple facets of a person and guiding the first stages of a person knowledge graph (PKG) creation. Graphs are then extracted, and compact representations are generated leveraging GNNs. We evaluate our approach using a real-world hospital intensive care unit (ICU) dataset and a hospital readmission prediction task. To the best of our knowledge, the novelty of the proposed approach can be summarised as:

- the first end-to-end framework for PKG extraction in Resource Description Framework (RDF) and PyTorch Geometric applicable format, using structured EHRs as well as unstructured clinical notes.
- the first use of a star-shaped Health & Social Person-centric Ontology (HSPO) [29] to model a comprehensive patient's view, focused on multiple facets (e.g. clinical and social).
- a representation learning approach that embeds personal knowledge graphs (PKGs) using GNNs and tackles the task of ICU readmission prediction using a real-world ICU dataset.
- the implementation is open-sourced¹ and generalizable so that it can be used to undertake other downstream tasks.

The paper is structured in the following way: section 2 discusses related work and section 3 describes the HSPO ontology. Section 4 describes the data preprocessing pipeline for PKG extraction using the ontology and the processed dataset. Section 5 evaluates different graph approaches in tackling downstream tasks anchored by an ICU hospital readmission prediction task. Finally, section 6 discusses the impact and adoption of the study and highlights the applicability of the framework in different downstream tasks and domains.

2 RELATED WORK

Public KGs are perhaps the most pervasive type of KGs today. These are often based on publicly available documents, encyclopedias or domain-specific databases and their schemas describe the features typically found within them. In recent years, especially in the health and biomedical domains, different types of KGs have been proposed from literature or EHRs yet they are not usually centred around the individual. Recent works include the PubMed KG [69], enabling connections among bio-entities, authors, articles, affiliations, and funding, the clinical trials KG [16], representing medical entities found in clinical trials with downstream tasks including drug-repurposing and similarity searches, and the PrimeKG [14], a multimodal KG for precision medicine analyses centred around diseases. Indeed disease-centric KGs have been previously proposed and despite some efforts in overlaying individual patient information [46], these graphs are not centred around the person or patient.

¹<https://github.com/IBM/hspo-ontology>

The idea behind entity-centric KGs and particularly person-centric graphs is relatively unexplored. One of the first efforts to define personal knowledge graphs (PKGs) was that of Balog and Kenter [3] where the graph has a particular “spiderweb” or star-shaped layout and every node has to have a direct or indirect connection to the user (i.e. person). The main advantage of having such a star-shaped representation lies in the fact the graph itself becomes personalized, enabling downstream applications across health and wellbeing, or personal assistants [3]. Similarly, a review paper [56] discussed the idea of a Person Health KG as a way to represent aggregated multi-modal data including all the relevant health-related personal data of a patient in the form of a structured graph but several challenges remained and implementations are lacking especially in representation learning. Subsequent works have proposed research PKGs [13] to represent information about the research activities of a researcher, and personal attribute knowledge bases [40] where a pre-trained language model with a noise-robust loss function aims to predict personal attributes from conversations without the need for labeled utterances. A knowledge model for capturing dietary preferences and personal context [60] has been proposed for personalized dietary recommendations where an ontology was developed for capturing lifestyle behaviors related to food consumption. A platform that leverages the Linked Open Data [7] stack (RDF, URIs, and SPARQL) to build RDF representations of Personal Health Libraries (PHLs) for patients is introduced in [1] and is aimed at empowering care providers in making more informed decisions. A method that leverages personal information to build a KG to improve suicidal ideation detection on social media is introduced by [12]. In the latter, the extracted KG includes several user nodes as the social interaction between the users is modeled. UniSKGRep [61] is a unified representation learning framework of KGs and social networks. Personal information of athletes and scientists is extracted to create two different universal KGs. The framework is used for node classification [57] and link prediction [37]. A graph-based approach that leverages the interconnected structure of personal web information, and incorporates techniques to update the representations as new data are added is proposed by [58]. The approach captures personal activity-based information and supports the task of email recipient prediction [54].

Despite the above efforts and idiosyncrasies, person-centric KGs have not been extensively used for predictive or classification tasks, especially those involving graph embeddings and GNNs. Similarly, ontologies that support the creation of entity-centric, star-shaped PKGs are not well established and there is no published research on representation learning of person-centric KGs using GNNs. To the best of our knowledge, this paper is the first to propose a framework for learning effective representations of an entity or person (i.e., graph classification setting [11, 68]) using PKGs and GNNs that can be applied to different downstream predictive tasks.

3 HSPO ONTOLOGY

The HSPO ontology has been designed to describe a holistic view of an individual spanning across multiple facets. HSPO defines a schema for a star-shaped PKG [3] with a person as a central node of the graph and corresponding characteristics (e.g. computable phenotype) linked to the central node. This view is unique as it is designed to continue to be expanded with additional domains and

facets of interest to the individual including but not limited to the clinical, demographics, and social factors in a first version but also being able to expand future versions to include behavioral, biological, or gene information. Representing a holistic view of individuals with new information generated beyond the traditional healthcare setting is expected to unlock new insights that can transform the way in which patients are treated or services delivered.

Previous ontologies have been built to harmonize disease definitions globally (MONDO [64]), to provide descriptions of experimental variables such as compounds, anatomy, diseases, and traits (EFO [42]), or to describe evidence of interventions in populations ([43]). Other ontologies have focused on specific contexts or diseases, such as SOHO, describing a set of social determinants affecting individual’s health [36] or an ontology that describes behavior change interventions [45]. Further to this, not all ontologies provide mappings for their entities into standard biomedical terminologies and vocabularies. The HSPO aims to address these challenges by being the first to create a person-centric view linking multiple facets together, leveraging existing ontological efforts, and providing mappings to known terms of biomedical vocabularies when appropriate.

HSPO has been built incrementally with the main objective of providing an accurate and clear representation of a person and their characteristics across multiple domains of interest. Domains of interest were identified and prioritized together with domain experts and a well-established methodology [50] was followed to guide development. HSPO is built not only to ensure that questions may be asked from the generated KGs but also that derived graphs may be used to train neural networks for downstream tasks.




Figure 1: Main classes defined in HSPO.

4 PERSON-CENTRIC KG EXTRACTION

Generating Person-Centric KGs requires a data preprocessing pipeline to prepare the dataset before graphs can be extracted. This section describes the use case rationale and data, the data preprocessing pipeline, and the steps taken to extract KGs.

4.1 Use Case Rationale and Dataset

In this study, we use EHRs on ICU admissions provided by MIMIC-III [31, 32], a well-established publicly available and de-identified dataset that contains EHRs from a large hospital’s ICU. MIMIC-III, therefore, includes data and a structure similar to most hospitals, containing both tabular data and clinical notes in the form of free text. The data covers the demographic (e.g. marital status, ethnicity, etc.), clinical (e.g. diagnoses, medication, etc.), and some aspects of the social view of the patient embedded in text notes. Detailed results of lab tests and metrics of monitoring medical devices are also provided. MIMIC-III is not only an appropriate dataset because of its structure and the types of data that it contains but also because of its population characteristics, including the reasons for hospital admission. More than 35% of the patients in this dataset were admitted with cardiovascular conditions [31] which are broadly relevant

across healthcare systems globally. Indeed, following their first hospital discharge, nearly 1 in 4 heart failure patients are known to be readmitted within 30 days, and approximately half are readmitted within 6 months [34]. These potentially avoidable subsequent hospitalizations are on the increase and reducing 30-day readmissions has now been a longstanding target for governments worldwide to improve the quality of care (e.g. outcomes) and reduce costs [38].

Therefore, both the experiments carried out and knowledge graphs generated in this paper describe a use case on patients admitted with cardiovascular conditions and a downstream prediction task to identify potential 30-day readmissions. This task can be further generalised to other conditions as readmissions are a widely used metric of success in healthcare, and other outcome metrics (e.g. specific events, mortality) are also possible using the approach proposed in this paper and accompanying open-source code repository.

4.2 Preprocessing Pipeline




Figure 2: Preprocessing: The EHRs are preprocessed and a JSON file is extracted, that consists of admission records describing the clinical, demographic, and social patient’s view.

4.2.1 Data Selection and Completion. The goal of the data preprocessing is to prepare the dataset for the PKG extraction. Due to the nature of MIMIC-III, each PKG represents the state of a patient during a single hospital admission. In the first step of the data preprocessing pipeline (Fig. 2), we select and aggregate the relevant data. In order to construct an efficient and applicable representation for a range of downstream tasks, the data selection step is necessary to include a subset of the EHR data, that is concise. We exclude the detailed lab test results (e.g. full blood count test values), as we assume that the diagnoses, medication, and procedures data are sufficiently expressive to represent the clinical view of the patient. The inclusion of fine-grain information poses additional challenges in the encoding of the PKG and the representation learning process.

Following this strategy, we select the demographic information: gender, age, marital status, religion, ethnicity, and the essential clinical data: diagnoses, medication, and procedures. We create a different record for each admission with the corresponding information using the JSON format. The diagnoses and procedures in MIMIC-III are recorded using the international classification of diseases, ninth revision (ICD-9) [51] coding schema with a hierarchical structure (with a format $xXXX.YYY$) where the first three or four digits of the code represent the general category, and the subsequent three digits after the dot separator represent more specific subcategories. We group the diagnoses and procedures using the corresponding general family category ($xXXX$) to reduce the number of different textual descriptions while defining substantially each diagnosis and procedure. For example, the diagnoses *acute myocardial infarction of anterolateral wall* (ICD-9 code: 410.0) and *acute myocardial infarction of inferolateral wall* (ICD-9 code: 410.2) are grouped under the general family diagnosis *acute myocardial infarction* (ICD-9 code: 410). This grouping is important, otherwise,

the encoding of the graph and training of a GNN to solve a downstream task would be very challenging, as some diagnoses and procedures are rare and underrepresented in the limited dataset of the study. In detail, more than 3,900 diagnoses and 1,000 procedures have a frequency of less than 10 in the dataset, while after grouping the numbers drop to less than 300 and 280 respectively.

4.2.2 Data Sampling. As the diagnoses and procedures are given using ICD-9 coding, we add the textual descriptions of the ICD-9 codes to the dataset. We sample the data records that are appropriate for the downstream task: 30-day ICU readmission prediction. The day span is a hyperparameter. The ICU admission records of patients that passed away during their admission to the hospital, or in a day span of fewer than 30 days after their discharge, are excluded.

4.2.3 Clinical Notes Integration. The clinical notes of the dataset contain information related to the clinical, demographic, and social aspects of the patients. We use the MetaMap [2] annotator to annotate the clinical notes with UMLS entities (UMLS Metathesaurus [8, 44]) and sample the codes that are related to certain social aspects such as employment status, household composition, and housing conditions. In this study, we focus on these three social aspects because social problems are known to be non-clinical predictors of poor outcomes but also they are often poorly recorded yet previous works have identified these specific three social problems available in MIMIC-III [41]. The extracted UMLS annotations of the clinical notes are integrated into the processed dataset.

4.2.4 Summary of Processed Dataset. After the data preprocessing steps, the total number of admission records is 51,296. From these, in 47,863 (93.3%) of the cases, the patients are not readmitted to the hospital, while 3,433 (6.7%) of the patients returned to the hospital within 30 days. In the downstream task, we focus on patients diagnosed with heart failure or cardiac dysrhythmia. Hence, the final dataset consists of 1,428 (9.2%) readmission cases and 14,113 (90.8%) non-readmission cases. Hence, the dataset is highly imbalanced as the readmission cases are underrepresented. In addition, EHR data is often incomplete [4, 30, 62, 66] and MIMIC-III, as it consists of real-world EHRs, is no exception. Information may be missing for multiple reasons, such as the unwillingness of the patient to share information, technical and coding errors, among several others [5]. More precisely, we observe that for some fields, such as religion, marital status, and medication there is a significant percentage of missing information. Tab. 1 shows the number of missing records per field for all admissions. We highlight that the social information (employment, housing conditions, and household composition), extracted from the unstructured data is scarce. This is an indication

Table 1: Missing information in the processed dataset.

Information	Records with missing information
Gender	0
Religion	17,794 (34.69%)
Marital Status	9,627 (18.77%)
Race/Ethnicity	4,733 (9.23%)
Diseases/Diagnoses	10 (0.02%)
Medication	8,032 (15.66%)
Procedures	6,024 (11.74%)
Employment	25,530 (49.77%)
Housing conditions	49,739 (96.96%)
Household composition	42,958 (83.75%)

that the clinical notes focus predominantly on the clinical view of the patients, without paying attention to aspects that can be connected to the social determinants of health, as reported [6, 41].

MIMIC-III is protected under HIPPA regulations. Thus, the detailed data distribution per field (e.g. diseases, etc.) cannot be shared publicly. The implementation for the extraction of the distributions is available in the official repository accompanying this paper and can be used when dataset access is officially granted² to the user.

4.3 Extraction

The HSPO ontology provides the knowledge schema used to create the PKGs. A PKG is extracted for every admission record of the processed data. The ontology represents the different classes (e.g. Religion, Disease, etc.), instances/individuals of the classes (e.g. Acute Kidney Failure), relations between the classes (e.g. hasDisease, hasSocialContext, etc.), and data properties (e.g. the instances of the class Age has a data property *age_in_years* that is an integer number). We use the rdflib python library for the implementation. The extracted knowledge graphs follow the RDF [19] format (Fig. 3), which is the primary foundation for the Semantic Web. Thus, SPARQL Query Language [18, 52] can be used to query, search and perform operations on the graphs.

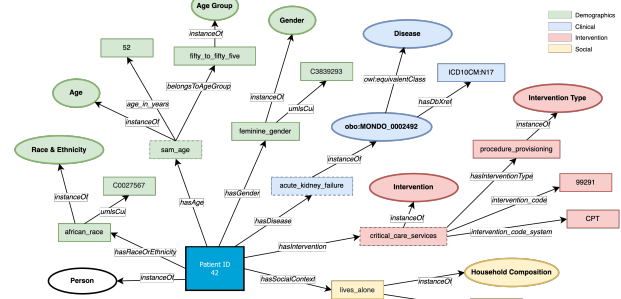


Figure 3: Person-Centric KG: The graph represents the demographics, clinical, intervention, and social patient’s aspects.

5 EVALUATION

This section reflects on the reliability and effectiveness of the person-centric graphs in downstream tasks. We evaluate the patient representation learning using person-centric graphs and a GNN on a specific ICU readmission prediction task. We provide insights into the applicability, benefits, and challenges of the proposed solution.

5.1 Data Transformation

The extracted graphs in RDF format cannot be used to train GNNs. Hence, a transformation step is implemented to transform graphs into the format (.dt files) of PyTorch Geometric [21], as we adopt this framework to build the models. The transformed graphs consist of the representations (initialized embeddings) of the nodes and the adjacency matrices for each relation type. They follow a triplet-based format, where the triplets have the following format: $[entity\ 1, relation\ type, entity\ 2]$ (e.g. [patient, hasDisease, acute renal failure]).

Training GNNs using person-centric graphs as input is a relatively unexplored field and defining *a priori* the most useful graph structure is not a trivial task. We experiment with 4 different graph versions (Fig. 4) to find the most suitable structure for the downstream task given the available data. The strategy to define the

²<https://physionet.org/content/mimiciii/1.4/>

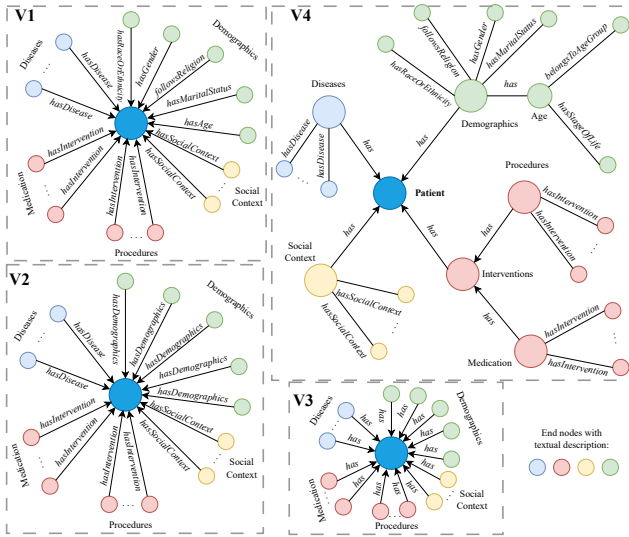


Figure 4: The 4 directed graph structures of the study are presented. The colors of the nodes represent the different types of information and are consistent across the versions.

different graph structures is described as follows: starting from a detailed version, we progressively simplify the graph by reducing the heterogeneity with relation grouping. We highlight that finding the optimal level of heterogeneity is an open and challenging research question as it depends on the available data, the downstream task, and the model architecture. More precisely, the first version is aligned with the schema provided by HSPO and includes 8 relation types (*hasDisease*, *hasIntervention*, *hasSocialContext*, *hasRaceOrEthnicity*, *followsReligion*, *hasGender*, *hasMaritalStatus*, and *hasAge*). The detailed demographic relations are grouped under the *hasDemographics* relation in the second version. The third version is the most simplified containing only the *has* relation type.

Lastly, we present the fourth version with the inclusion of group nodes to explore the effectiveness of graph expressivity on learnt representation. The design of the fourth version is based on the assumption that the summarization of the corresponding detailed information (e.g. the disease information is summarized in the representation of the *Disease* node) and the learning of a grouped representation during training can boost the performance of the models. To explore this assumption, we introduce a node type *group node* and add 7 group nodes (*Diseases*, *Social Context*, *Demographics*, *Age*, *Interventions*, *Procedures*, and *Medication*). For the inclusion of the group nodes in the graph, we introduce the general relation type *has*. Fig. 4 presents the four directed graph structures. We also experiment with the undirected corresponding versions.

The nodes of each graph are initialized using the bag-of-words (BOW) approach. Hence, the initial representations are sparse. We create the vocabulary using the textual descriptions of the end nodes (diagnoses names, medication, etc.) and the description of the group nodes (patient, social context, diseases, interventions, medication, procedures, demographics, and age). The final vocabulary consists of 3,723 nodes. Alternatively to the introduction of sparsity using the BOW approach, a language model, such as BioBERT [39], PubMedBERT [24], and CharacterBERT [20, 63], can be used for the node initialization. We provide this capability in our open-source framework.

5.2 GNNs and Baseline Models

We experiment with two different GNN architectures and each of them has two variations. The difference between the variations lies in the final layer of the model which is convolutional or linear (Fig. 5). The first model (PKGSage) is based on Sage Graph Convolution Network [26] and the second (PKGA) utilizes the Graph Attention Network (GAT) architecture [10, 65]. Given a graph structure G with N nodes, a sequence of transformations T is applied, and the final prediction p of the model is extracted as follows:

$$X_{k,i} = \sigma(T_i(X_{k,i-1}, G)), \text{ with } k \in [1, \dots, N] \text{ and } i \in [1, 2, 3], \quad (1)$$

$$p = \sigma(X_{n,3}), \quad (2)$$

where T_i is the transformation (Sage or GAT Convolution, or linear) of the i^{th} layer, $X_{k,i}$ is the k node representation after the T_i transformation, $X_{n,3}$ is the final output of the last layer for the *patient* node, and σ is the activation function. ReLU is used as the activation function for the first two layers and Sigmoid for the last layer. In principle, the graphs are multi-relational and this can lead to rapid growth in the number of model parameters. To address this issue, we apply basis decomposition [59] for weight regularization.

We incorporate a set of baseline classifiers to compare the graph-based approach with traditional machine-learning algorithms. Particularly, k -nearest neighbors (KNN) (k is set to 5 based on the average performance in the validation set), linear and non-linear Support Vector Machines (L-SVM and RBF-SVM respectively) [17], Decision Tree (DT) [9], AdaBoost [22, 27], Gaussian Naive Bayes (NB), and Gaussian Process (GP) [55] with radial basis function (RBF) kernel are included in the study. We apply one-hot encoding to transform the textual descriptions into numerical features for the diagnoses, medication, and procedures. The remaining features (gender, religion, marital status, age group, race, employment, housing conditions, and household composition) are categorical, so we encode each category using mapping to an integer number. For a fair comparison, feature engineering or feature selection techniques are not implemented, since the graph-based models include all the available information extracted from the EHRs.




Figure 5: The input is a graph with N number of nodes and edge index per relation type. The dimension of the initial representation of the nodes is 3,723. The PKGSage models leverage the Sage Convolution while the PKGA models include the GAT Convolution module. The final output of the models is the readmission probability of the patient.

5.3 Experimental Setup

We create 10 different balanced dataset splits for experimentation to overcome the imbalance problem. In detail, the 14,113 non-readmission cases are randomly divided into 10 folds. Combined with the 1,428 readmission cases, these folds constitute the final 10 balanced splits. For each balanced split, 5-fold cross-validation is applied. We highlight that we use the same splits across the different experimental settings to have a fair comparison. The Adam optimizer [35] is used and the learning rate is set to 0.001. The models are trained for 100 epochs and the best model weights are stored based on the performance in the validation set (15% of the training set). The number of bases, for the basis decomposition method [59], is 3, and the batch size is set to 32.

5.4 Results

We present the average results, across the different runs and folds of 5-fold cross-validation, of the models using different graph versions (Tab. 2) and compare the performance of the best models with various machine learning classifiers (Tab. 3). Starting with the intra-model comparison, the PKGSage model performs best using the first directed graph version with 62.16% accuracy and the third undirected graph version with 68.06% F1-score. The PKGA model achieves 61.69% accuracy and 67.49% F1-score using the third undirected graph version. Overall, the end-to-end convolution strategy is advantageous as the second variation of the models (PKGSage₂, PKGA₂) performs better than the first variation with the linear final layer in most cases. Nonetheless, the best performance is achieved by the PKGSage₁ model. The inter-model comparison reveals that both models (PKGSage, PKGA) achieve similar results.

The graph's structure is essential since the performance varies according to the graph version. More precisely, the results are comparable with less notable differences using the undirected/directed first version of the graph, and the undirected second and third versions. However, the performance degradation is significant when the directed second and third versions are utilized. In both versions, we observe that all links point to the central *Patient* node (Fig. 4), and the representations of the remaining nodes are not updated during training, imposing a challenge on the trainability of the models. However, the same structure is present in the directed first version where no performance drop is noticed. Another possible reason is the level of heterogeneity of the graph. The first version has 8 relation types, while the second and third versions contain 4 and 1 relation types respectively. Given these, we conclude that the direction of the links and graph heterogeneity are crucial for

Table 2: Results: models with different graph versions.

G. ¹	D. ²	Model	Accuracy	F1-Score	G. ¹	D. ²	Model	Accuracy	F1-Score	
V1	Undirected	PKGSage ₁	61.72 ± 0.98	67.86 ± 1.31	V3	Undirected	PKGSage ₁	61.69 ± 0.91	68.06 ± 1.14	
		PKGSage ₂	61.27 ± 0.92	67.44 ± 1.34			PKGSage ₂	61.48 ± 0.93	67.5 ± 1.21	
		PKGA ₁	59.73 ± 1.08	66.32 ± 1.12			PKGA ₁	61.69 ± 0.71	66.32 ± 1.56	
	Directed	PKGA ₂	61.39 ± 0.89	66.58 ± 1.56		Directed	PKGA ₂	61.46 ± 0.82	67.49 ± 1.76	
		PKGSage ₁	62.16 ± 0.89	67.93 ± 1.35			PKGSage ₁	50.2 ± 0.51	64.21 ± 1.74	
		PKGSage ₂	61.72 ± 0.73	67.2 ± 1.68			PKGSage ₂	49.93 ± 0.57	61.43 ± 1.83	
V2	Undirected	PKGA ₁	60.5 ± 1.26	66.52 ± 1.51	V4	Undirected	PKGA ₁	50.2 ± 0.63	61.54 ± 1.91	
		PKGA ₂	61.38 ± 0.87	67 ± 1.65			Directed	PKGA ₂	49.98 ± 0.61	61.44 ± 1.87
		PKGSage ₁	60.43 ± 0.63	67.85 ± 1.22				PKGSage ₁	49.5 ± 0.55	59.77 ± 2.54
	PKGSage ₂	60.93 ± 0.81	67.06 ± 1.45	PKGSage ₂		54.5 ± 1.4		59.51 ± 2.36		
	Directed	PKGA ₁	58.95 ± 0.78	66.21 ± 1.36		Directed	PKGA ₁	52.08 ± 1.68	53.1 ± 1.88	
		PKGA ₂	60.24 ± 0.99	67.37 ± 1.31			PKGA ₂	57.9 ± 0.91	58.67 ± 0.9	
PKGSage ₁		51.26 ± 0.81	65.1 ± 1.58	PKGSage ₁	49.45 ± 0.58		58.41 ± 2.75			
Directed	PKGSage ₂	51.61 ± 0.76	61.75 ± 1.26	Directed	PKGSage ₂	54.79 ± 1.79	59.87 ± 2.67			
	PKGA ₁	51 ± 0.81	64.86 ± 1.27		PKGA ₁	51.01 ± 1.29	62.92 ± 1.54			
	PKGA ₂	51.3 ± 0.91	63.84 ± 1.43		PKGA ₂	57.32 ± 1.54	60.16 ± 2.28			

¹ G.: Graph, ² D.: Direction

Table 3: Results: Comparison with baseline models

Metric	DT	AdaBoost	NB	GP
Accuracy	55.34 ± 0.55	60.03 ± 0.65	53.92 ± 1.44	56.82 ± 0.57
F1-Score	57.5 ± 0.65	59.01 ± 1.44	39.93 ± 3.56	52.53 ± 1.62
Metric	KNN	L-SVM	RBF-SVM	PKGSage ₁
Accuracy	57.31 ± 0.54	61.58 ± 0.62	62.11 ± 0.65	62.16 ± 0.89
F1-Score	50.9 ± 1.73	61.45 ± 1.49	64.44 ± 1.49	68.06 ± 1.14

the final performance of the models. Leveraging the fourth graph version is not advantageous since the models achieve worse performance. This indicates that introducing the group nodes and additional expressibility can be an obstacle to the trainability and performance of the models. We observe saturation and stability problems [23] during training when the fourth version is used. The best-performing models of the study also significantly outperform the baseline models. We notice an improvement of 3.62% in F1-Score. Based on the accuracy metric, only the SVM classifiers (linear and non-linear) achieve comparable results. The results illustrate the potential of the PKG-based models in downstream predictive tasks and particularly in ICU readmission prediction.

5.5 Ablation Study

Following the observation that the data is incomplete, we conduct an ablation study to probe the robustness of our approach in handling missing information. The following hypotheses are drawn:

- The patient's clinical view (medication, diseases, procedures) is very important in predicting ICU readmission. (H1)
- The exclusion of additional information results in lower performance. (H2)

We apply the ablation study using the best-performing model PKGSage₁ with the directed first graph version (PKGSage₁DV₁) and the undirected third graph version (PKGSage₁UnV₃). In the first step, we exclude one facet of the data (e.g. medication, procedures) and evaluate performance degradation. To address the H2 hypothesis and to further test model robustness, in the next step, we exclude two facets in a complete manner (as depicted in Tab. 4). The exclusion of the disease information results in the most significant performance decline in the accuracy for both model versions and in the F1-score for PKGSage₁UnV₃. Excluding the medication information leads to 1.13% drop in F1-score for PKGSage₁DV₁. The results of the ablation study only partially support the H1 hypothesis. The robustness of the models in handling missing information is profound as the performance deterioration is limited. In the worst

Table 4: Ablation Study

Excluded Information	PKGSage ₁ undirected V3		PKGSage ₁ directed V1	
	Accuracy	F1-Score	Accuracy	F1-Score
-	61.69	68.06	62.16	67.93
Social aspect	60.68 (↓ 1.01)	67.63 (↓ 0.43)	60.5 (↓ 1.66)	67.64 (↓ 0.29)
Medication	59.7 (↓ 1.99)	66.98 (↓ 1.08)	60.03 (↓ 2.13)	66.8 (↓ 1.13)
Procedures	60.42 (↓ 1.27)	67.43 (↓ 0.63)	59.97 (↓ 2.19)	67.56 (↓ 0.37)
Diseases	59.69 (↓ 2)	66.87 (↓ 1.19)	59.59 (↓ 2.57)	67.48 (↓ 0.45)
Demographics	60.43 (↓ 1.26)	67.43 (↓ 0.63)	60.16 (↓ 2)	67.54 (↓ 0.39)
Social aspect and Medication	60.13 (↓ 1.56)	66.84 (↓ 1.22)	60.01 (↓ 2.15)	66.67 (↓ 1.26)
Social aspect and Procedures	60.13 (↓ 1.56)	67.99 (↓ 0.07)	59.87 (↓ 2.29)	67.34 (↓ 0.59)
Social aspect and Diseases	59.87 (↓ 1.82)	66.62 (↓ 1.44)	59.6 (↓ 2.56)	67.18 (↓ 0.75)
Social aspect and Demographics	60.64 (↓ 1.05)	66.84 (↓ 1.22)	60.14 (↓ 2.02)	66.88 (↓ 1.05)
Medication and Procedures	58.76 (↓ 2.93)	66.59 (↓ 1.47)	59.47 (↓ 2.69)	66.86 (↓ 1.07)
Medication and Diseases	59.49 (↓ 2.2)	65.69 (↓ 2.37)	59.72 (↓ 2.44)	66.06 (↓ 1.87)
Medication and Demographics	59.84 (↓ 1.85)	66.8 (↓ 1.26)	59.9 (↓ 2.26)	66.56 (↓ 1.37)
Procedures and Diseases	59.32 (↓ 2.37)	66.62 (↓ 1.44)	58.53 (↓ 3.63)	66.38 (↓ 1.55)
Procedures and Demographics	60.63 (↓ 1.06)	67.69 (↓ 0.37)	60.01 (↓ 2.15)	67.67 (↓ 0.26)
Diseases and Demographic	59.53 (↓ 2.16)	66.85 (↓ 1.21)	59.56 (↓ 2.6)	66.52 (↓ 1.73)

case, accuracy and F1-score drop by 2.57% and 1.13% respectively. A similar pattern is revealed when we remove two facets of the data as the removal of clinical information is reflected in lower performance. More precisely, the exclusion of the medication and disease information leads to 2.37% and 1.87% drop in F1-score for the PKGSage_{1UnV₃} and PKGSage_{1DV₁} models correspondingly. PKGSage_{1UnV₃} is less accurate by 2.93% when medication and procedures nodes are absent, and PKGSage_{1DV₁} achieves 58.53% accuracy (3.63% reduction) when procedure and disease data are excluded. The performance of the models is robust even when two out of three clinical facets (medication, diseases, and procedures) are unavailable. We highlight that stability and robustness are key properties, especially in healthcare data where missing information is inevitable due to privacy issues, system or human errors.

6 DISCUSSION

The proposed end-to-end framework for person or entity-centric KGs has impacts on the technology and industry and we highlight the benefits and challenges of adopting KG technologies in practice. We focus on the healthcare domain yet present a solution that is, by design, generalisable across domains and it is also not restricted by the final predictive task. The proposed approach was evaluated using complex real-world data, which is imbalanced, heterogeneous, and sparse, and the ontology efforts and approach were reviewed by domain experts from multiple organisations. The proposed star-shaped KG approach and ontology will not only allow data from disparate sources and types to be meaningfully combined, as already demonstrated, but also allow pertinent new research questions to be addressed. This paves the way towards a comprehensive and holistic 360° view of the person to be created from multiple data sources and could be further generalised to patient cohorts or groups of individuals. This approach is also scalable as multiple PKGs can be created, one for each patient, unlike traditional approaches relying on very large general knowledge KGs or others which have reported scalability issues [59]. The topology of traditional KGs used to query and infer knowledge might not be easily applied to learn graph representations using GNNs primarily due to their large size. Our experiments show that reducing the heterogeneity with relation grouping has an effect on F1-score provided that the PKG structure is generally fit for learning patient representations using GNNs. We also observe that the proposed framework is able to generalise even with limited amount of data.

Furthermore, the ontology design process can be used as guidance for the creation of other entity-centric ontologies and the HSPO is continuing to be expanded with new domains and facets. The open-sourced implementation can be reused, with necessary adjustments, to extract entity-centric knowledge graphs in RDF or PyTorch Geometric applicable format. Finally, a wide range of predictive tasks using neural networks, or inferences using the KGs produced can be addressed, even if these will be constrained by the availability and expressiveness of the data and provided annotations. We highlight the applicability of the framework through a readmission prediction task using the PKGs and GNN architectures. Following the proposed paradigm, other classification tasks, such as mortality prediction and clinical trial selection, can be undertaken. The efficient patient representation may also be leveraged for clustering and similarity grouping applications.

7 CONCLUSION AND FUTURE WORK

This paper proposes a new end-to-end representation learning framework to extract entity-centric ontology-driven KGs from structured and unstructured data. We describe how KG technologies can be used in combination with other technologies and techniques (learning entity representation using GNNs, predicting outcomes (prediction tasks) using neural networks and learnt representation) to drive practical industry applications with an example in the healthcare industry. The open-sourced framework and approach are scalable and show stability and robustness when applied to complex real-world healthcare data. We plan to extend and apply the PKG and the proposed framework to new use cases which will further drive adoption. In healthcare, we envisage that this work can unlock new ways of studying complex disease patterns and drive a better understanding of disease across population groups.

REFERENCES

- [1] Nariman Ammar et al. 2021. Using a Personal Health Library-Enabled mHealth Recommender System for Self-Management of Diabetes Among Underserved Populations: Use Case for Knowledge Graphs and Linked Data. *JMIR formative research* 5, 3 (2021), e24738. <https://doi.org/10.2196/24738>
- [2] Alan R Aronson. 2001. Effective mapping of biomedical text to the UMLS Metathesaurus: the MetaMap program. In *Proceedings of the AMLA Symposium*. American Medical Informatics Association, Washington, DC, USA, 17.
- [3] Krisztian Balog and Tom Kenter. 2019. Personal knowledge graphs: A research agenda. In *Proceedings of the 2019 ACM SIGIR International Conference on Theory of Information Retrieval*. Santa Clara, CA, USA, 217–220. <https://doi.org/10.1145/3341981.3344241>
- [4] Brett K Beaulieu-Jones et al. 2018. Characterizing and managing missing structured data in electronic health records: data analysis. *JMIR medical informatics* 6, 1 (2018), e8960. <https://doi.org/10.2196/medinform.8960>
- [5] J Bettencourt-Silva et al. 2012. On creating a patient-centric database from multiple Hospital Information Systems. *Methods of information in medicine* 51, 03 (2012), 210–220. <https://doi.org/10.3414/ME10-01-0069>
- [6] Joao Bettencourt-Silva et al. 2020. Discovering new social determinants of health concepts from unstructured data: framework and evaluation. In *Digital Personalized Health and Medicine*. IOS Press, 173–177. <https://doi.org/10.3233/SHTI200145>
- [7] Christian Bizer et al. 2011. Linked data: The story so far. In *Semantic services, interoperability and web applications: emerging concepts*. IGI global, 205–227. <https://doi.org/10.4018/978-1-60960-593-3.ch008>
- [8] Olivier Bodenreider. 2004. The unified medical language system (UMLS): integrating biomedical terminology. *Nucleic acids research* 32, suppl_1 (2004), D267–D270. <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC2243666/>
- [9] Leo Breiman. 2017. *Classification and regression trees*. Routledge. <https://doi.org/10.1201/9781315139470>
- [10] Shaked Brody et al. 2021. How attentive are graph attention networks? *arXiv preprint arXiv:2105.14491* (2021). <https://doi.org/10.48550/arXiv.2105.14491>
- [11] Hongyun Cai et al. 2018. A comprehensive survey of graph embedding: Problems, techniques, and applications. *IEEE transactions on knowledge and data engineering* 30, 9 (2018), 1616–1637. <https://doi.org/10.1109/TKDE.2018.2807452>
- [12] Lei Cao et al. 2020. Building and using personal knowledge graph to improve suicidal ideation detection on social media. *IEEE Transactions on Multimedia* 24 (2020), 87–102. <https://doi.org/10.1109/TMM.2020.3046867>
- [13] Prantika Chakraborty et al. 2022. Personal Research Knowledge Graphs. In *Companion Proceedings of the Web Conference 2022* (Virtual Event, Lyon, France) (*WWW '22*). Association for Computing Machinery, New York, NY, USA, 763–768. <https://doi.org/10.1145/3487553.3524654>
- [14] Payal Chandak et al. 2023. Building a knowledge graph to enable precision medicine. *Scientific Data* 10, 1 (2023), 67. <https://doi.org/10.1038/s41597-023-01960-3>
- [15] Andrew Chatr-Aryamontri et al. 2012. The BioGRID interaction database: 2013 update. *Nucleic acids research* 41, D1 (2012), D816–D823. <https://doi.org/10.1093/nar/gks1158>
- [16] Ziqi Chen et al. 2022. A knowledge graph of clinical trials (CTKG). *Scientific reports* 12, 1 (2022), 4724. <https://doi.org/10.1038/s41598-022-08454-z>
- [17] Corinna Cortes and Vladimir Vapnik. 1995. Support-vector networks. *Machine learning* 20 (1995), 273–297. <https://doi.org/10.1007/BF00994018>
- [18] Richard Cyganiak. 2005. A relational algebra for SPARQL. *Digital Media Systems Laboratory HP Laboratories Bristol*. HPL-2005-170 35, 9 (2005). <http://shiftleft.com/mirrors/www.hpl.hp.com/techreports/2005/HPL-2005-170.pdf>
- [19] Stefan Decker et al. 2000. The semantic web: The roles of XML and RDF. *IEEE Internet computing* 4, 5 (2000), 63–73. <https://doi.org/10.1109/4236.877487>

- [20] Hicham El Boukkouri et al. 2020. CharacterBERT: Reconciling ELMo and BERT for Word-Level Open-Vocabulary Representations From Characters. In *Proceedings of the 28th International Conference on Computational Linguistics*. International Committee on Computational Linguistics, Barcelona, Spain (Online), 6903–6915. <https://doi.org/10.18653/v1/2020.coling-main.609>
- [21] Matthias Fey and Jan Eric Lenssen. 2019. Fast graph representation learning with PyTorch Geometric. In *ICLR Workshop on Representation Learning on Graphs and Manifolds*.
- [22] Yoav Freund and Robert E Schapire. 1997. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of computer and system sciences* 55, 1 (1997), 119–139. <https://doi.org/10.1006/jcss.1997.1504>
- [23] Xavier Glorot and Yoshua Bengio. 2010. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics (Proceedings of Machine Learning Research, Vol. 9)*, Yee Whye Teh and Mike Titterton (Eds.). PMLR, Chia Laguna Resort, Sardinia, Italy, 249–256. <https://proceedings.mlr.press/v9/glorot10a.html>
- [24] Yu Gu et al. 2021. Domain-specific language model pretraining for biomedical natural language processing. *ACM Transactions on Computing for Healthcare (HEALTH)* 3, 1 (2021), 1–23. <https://doi.org/10.1145/3458754>
- [25] Qingyu Guo et al. 2020. A survey on knowledge graph-based recommender systems. *IEEE Transactions on Knowledge and Data Engineering* 34, 8 (2020), 3549–3568. <https://doi.org/10.1109/TKDE.2020.3028705>
- [26] Will Hamilton et al. 2017. Inductive Representation Learning on Large Graphs. In *Advances in Neural Information Processing Systems*, Vol. 30. Curran Associates, Inc., 1025–1035. <https://proceedings.neurips.cc/paper/2017/file/5dd9db5e033da9c6fb5ba83c7a7eb/ea9-Paper.pdf>
- [27] Trevor Hastie et al. 2009. Multi-class adaboost. *Statistics and its Interface* 2, 3 (2009), 349–360. <https://doi.org/10.4310/SII.2009.v2.n3.a8>
- [28] Aidan Hogan et al. 2021. Knowledge graphs. *ACM Computing Surveys (CSUR)* 54, 4 (2021), 1–37. <https://doi.org/10.1145/3447772>
- [29] HSPO_Team. 2022. Health and Social Person-centric Ontology. <https://github.com/IBM/hspo-ontology>
- [30] Zhen Hu et al. 2017. Strategies for handling missing clinical data for automated surgical site infection detection from the electronic health record. *J. Biomed. Inform.* 68 (2017), 112–120. <https://doi.org/10.1016/j.jbi.2017.03.009>
- [31] Alistair Johnson et al. 2016. MIMIC-III clinical database (version 1.4). *PhysioNet* 10 (2016), C2XW26. <https://doi.org/10.13026/C2XW26>
- [32] Alistair EW Johnson et al. 2016. MIMIC-III, a freely accessible critical care database. *Scientific data* 3, 1 (2016), 1–9. <https://doi.org/10.1038/sdata.2016.35>
- [33] Gjergji Kasneci et al. 2008. Naga: Searching and ranking knowledge. In *2008 IEEE 24th International Conference on Data Engineering*. IEEE, 953–962. <https://doi.org/10.1109/ICDE.2008.4497504>
- [34] Muhammad Shahzeb Khan et al. 2021. Trends in 30-and 90-day readmission rates for heart failure. *Circulation: Heart Failure* 14, 4 (2021), e008335. <https://doi.org/10.1161/CIRCHEARTFAILURE.121.008335>
- [35] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014). <https://doi.org/10.48550/arXiv.1412.6980>
- [36] Navya Martin Kollapally et al. 2022. An Ontology for the Social Determinants of Health Domain. In *2022 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*. IEEE, 2403–2410. <https://doi.org/10.1109/BIBM55620.2022.9995544>
- [37] Ajay Kumar et al. 2020. Link prediction techniques, applications, and performance: A survey. *Physica A: Statistical Mechanics and its Applications* 553 (2020), 124289. <https://doi.org/10.1016/j.physa.2020.124289>
- [38] C Lawson et al. 2021. Trends in 30-day readmissions following hospitalisation for heart failure by sex, socioeconomic status and ethnicity. *EclinicalMedicine* 38 (2021), 101008. <https://doi.org/10.1016/j.eclinm.2021.101008>
- [39] Jinhyuk Lee et al. 2020. BioBERT: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics* 36, 4 (2020), 1234–1240. <https://doi.org/10.1093/bioinformatics/bt2682>
- [40] Yinan Liu et al. 2022. Personal Attribute Prediction from Conversations. In *Companion Proceedings of the Web Conference 2022 (Virtual Event, Lyon, France) (WWW '22)*. Association for Computing Machinery, New York, NY, USA, 223–227. <https://doi.org/10.1145/3487553.3524248>
- [41] Kevin Lybarger et al. 2021. Annotating social determinants of health using active learning, and characterizing determinants using neural event extraction. *J. Biomed. Inform.* 113 (2021), 103631. <https://doi.org/10.1016/j.jbi.2020.103631>
- [42] James Malone et al. 2010. Modeling sample variables with an Experimental Factor Ontology. *Bioinformatics* 26, 8 (2010), 1112–1118. <https://doi.org/10.1093/bioinformatics/btq099>
- [43] Chris Mavergames et al. 2013. Systematic Reviews as an Interface to the Web of (Trial) Data: using PICO as an Ontology for Knowledge Synthesis in Evidence-based Healthcare Research. *SePublica* 994 (2013), 22–6.
- [44] Alexa T McCray. 2003. An upper-level ontology for the biomedical domain. *Comparative and Functional genomics* 4, 1 (2003), 80–84. <https://doi.org/10.1002/cfg.255>
- [45] Susan Michie et al. 2017. The Human Behaviour-Change Project: harnessing the power of artificial intelligence and machine learning for evidence synthesis and interpretation. *Implementation Science* 12, 1 (2017), 1–12. <https://doi.org/10.1186/s13012-017-0641-5>
- [46] Charlotte A Nelson et al. 2019. Integrating biomedical research and electronic health records to create knowledge-based biologically meaningful machine-readable embeddings. *Nature communications* 10, 1 (2019), 3045. <https://doi.org/10.1038/s41467-019-11069-0>
- [47] Markus Nentwig et al. 2017. A survey of current link discovery frameworks. *Semantic Web* 8, 3 (2017), 419–436. <https://doi.org/10.3233/SW-150210>
- [48] David N. Nicholson and Casey S. Greene. 2020. Constructing knowledge graphs and their biomedical applications. *Computational and Structural Biotechnology Journal* 18 (2020), 1414–1428. <https://doi.org/10.1016/j.csbj.2020.05.017>
- [49] Natasha Noy et al. 2019. Industry-scale Knowledge Graphs: Lessons and Challenges: Five diverse technology companies show how it's done. *Queue* 17, 2 (2019), 48–75. <https://doi.org/10.1145/3329781.3332266>
- [50] Natalya F Noy et al. 2001. Ontology development 101: A guide to creating your first ontology. https://corais.org/sites/default/files/ontology_development101_aguide_to_creating_your_first_ontology.pdf
- [51] World Health Organization et al. 1998. *ICD-9-CM: International Classification of Diseases, 9th Revision: Clinical Modification*. PMIC, Practice Management Information Corporation.
- [52] Jorge Pérez et al. 2006. Semantics and Complexity of SPARQL. In *International semantic web conference*. Springer, 30–43. <https://doi.org/10.1145/1567274.1567278>
- [53] Alberto Purpura et al. 2022. Accelerating the Discovery of Semantic Associations from Medical Literature: Mining Relations Between Diseases and Symptoms. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Abu Dhabi, UAE, 77–89. <https://aclanthology.org/2022.emnlp-industry.6>
- [54] Ashequl Qadir et al. 2016. Activity modeling in email. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. 1452–1462. <https://doi.org/10.18653/v1/N16-1171>
- [55] Carl Edward Rasmussen. 2004. *Gaussian Processes in Machine Learning*. Springer Berlin Heidelberg, Berlin, Heidelberg, 63–71. https://doi.org/10.1007/978-3-540-28650-9_4
- [56] Nidhi Rastogi and Mohammed J Zaki. 2020. Personal health knowledge graphs for patients. *arXiv preprint arXiv:2004.00071* (2020). <https://doi.org/10.48550/arXiv.2004.00071>
- [57] Yu Rong et al. [n. d.]. DropEdge: Towards Deep Graph Convolutional Networks on Node Classification. In *International Conference on Learning Representations*.
- [58] Tara Safavi et al. 2020. Toward activity discovery in the personal web. In *Proceedings of the 13th International Conference on Web Search and Data Mining*. 492–500. <https://doi.org/10.1145/3336191.3371828>
- [59] Michael Schlichtkrull et al. 2018. Modeling relational data with graph convolutional networks. In *The Semantic Web: 15th International Conference, ESWC 2018, Heraklion, Crete, Greece, June 3–7, 2018, Proceedings 15*. Springer, 593–607. https://doi.org/10.1007/978-3-319-93417-4_38
- [60] Oshani Wasana Seneviratne et al. 2021. Personal Health Knowledge Graph for Clinically Relevant Diet Recommendations. *ArXiv* (2021). <https://doi.org/10.48550/arXiv.2110.10131>
- [61] Yinghan Shen et al. 2023. UniSKGRep: A unified representation learning framework of social network and knowledge graph. *Neural Networks* 158 (2023), 142–153. <https://doi.org/10.1016/j.neunet.2022.11.010>
- [62] Gregor Stiglic et al. 2019. Challenges associated with missing data in electronic health records: a case study of a risk prediction model for diabetes using data from Slovenian primary care. *Health informatics journal* 25, 3 (2019), 951–959. <https://doi.org/10.1177/1460458217733288>
- [63] Christos Theodoropoulos et al. 2021. Imposing Relation Structure in Language-Model Embeddings Using Contrastive Learning. In *Proceedings of the 25th Conference on Computational Natural Language Learning*. Association for Computational Linguistics, Online, 337–348. <https://doi.org/10.18653/v1/2021.conll-1.27>
- [64] Nicole A Vasilevsky et al. 2022. Mondo: Unifying diseases for the world, by the world. *medRxiv* (2022), 2022–04. <https://doi.org/10.1101/2022.04.13.22273750>
- [65] Petar Veličković et al. 2017. Graph attention networks. *arXiv preprint arXiv:1710.10903* (2017). <https://doi.org/10.48550/arXiv.1710.10903>
- [66] Brian J Wells et al. 2013. Strategies for handling missing data in electronic health record derived data. *Egems* 1, 3 (2013). <https://doi.org/10.13063/2327-9214.1035>
- [67] Shihwen Wu et al. 2022. Graph neural networks in recommender systems: a survey. *Comput. Surveys* 55, 5 (2022), 1–37. <https://doi.org/10.1145/35353101>
- [68] Zonghan Wu et al. 2020. A comprehensive survey on graph neural networks. *IEEE transactions on neural networks and learning systems* 32, 1 (2020), 4–24. <https://doi.org/10.1109/TNNLS.2020.2978386>
- [69] Jian Xu et al. 2020. Building a PubMed knowledge graph. *Scientific data* 7, 1 (2020), 205. <https://doi.org/10.1038/s41597-020-0543-2>