Building a Temporal Knowledge Graph for Electronic Health Records

Ricardo M.S. Carvalho*, Andreia Sofia Teixeira and Catia Pesquita

LASIGE, Faculdade de Ciências, Universidade de Lisboa, Portugal

Abstract

The integration of Machine Learning (ML) into healthcare solutions has been accelerated by the proliferation of Electronic Health Records (EHRs). While various studies have leveraged ML to predict clinical outcomes, they often overlook the semantic context of clinical data. However, this context could be easily captured using ontologies and knowledge graphs (KG) as tools in ML development.

EHRs are particularly suitable for KG representation, affording resources to improve predictive clinical AI models and generate new data discoveries. Traditional KGs often neglect the temporal dynamics crucial for maintaining understanding over time, however, recent advancements suggest incorporating temporal dimensions into KGs, termed Temporal Knowledge Graphs (TKGs), as a promising approach for modeling patient data. Despite the potential for TKGs, few resources exist to produce such representations within the healthcare domain.

We propose a method to address this gap by establishing a framework for constructing general edge-centric clinical TKGs. We employ lexical and transformer-based methods for semantic annotation, mapping EHR concepts to biomedical ontologies. These annotations are paired with temporal stamping, enabling tracking of patient progression through time points and intervals for KG facts.

This work provides a comprehensive approach for creating TKGs from EHRs, setting the foundation for general clinical TKG construction whilst allowing the extension to multiple clinical datasets, thereby advancing further ML applications in healthcare.

Keywords

Ontologies, Semantic Annotation, Clinical Temporal Knowledge Graphs, Temporal Annotation

1. Introduction

The surge in Machine Learning (ML) applications within healthcare, particularly for clinical decision-making, has been notable due to the rise of Electronic Health Records (EHRs) systems [1, 2]. Studies have successfully applied ML to predict various clinical outcomes, including mortality, sepsis, and risk of readmission [3, 4, 5, 6], using data mining methods ranging from random forests to recurrent neural networks. However, most of these approaches largely ignore the context of clinical data, relying primarily on direct EHRs data extraction.

Ontologies play a crucial role in connecting data with scientific context by representing domain entities and their relationships in a formalized manner [7]. By linking EHR data to ontologies through semantic annotations, one can provide additional information about the meaning of the clinical data to ML systems. When data is organized according to the schema of one or more ontologies, it forms a Knowledge Graph (KG). A KG is composed of entities and relations structured as triples in graph format [8, 9, 10]. KGs go beyond mere knowledge bases, as they can represent complex entities and relations in a way that reflects real-world domains. This enables the incorporation of domain ontology knowledge and mappings across real data entities [11], enhancing domain representation, interpretability, and explainability.

Health-related data such as EHRs are particularly well-suited for representation through KGs. These KGs can enhance predictive clinical AI models by adding semantic information [2, 12, 3, 13], potentially revealing new data correlations and improving performance. However, one critical oversight in traditional KG approaches is the neglect of time as a dynamic factor. Real-world data, especially in

the clinical field, changes over time and requires temporal consideration to maintain logical coherence and accuracy. Non-KG-based ML techniques have highlighted the significance of including time for more precise predictions [14, 15, 16]. In the clinical field, there is a growing interest in using KGs to represent clinical data temporally [17, 18], particularly for modelling patient data. Still, there is a lack of resources that are able to combine both the semantic richness afforded by biomedical ontologies and temporal aspects.

In this work, we target this critical gap in healthcare data resources by proposing an approach to encapsulate the semantic and temporal context of clinical entities in a KG representation. As a case study, we apply this methodology to the MIMIC-III EHR data set [19]. Our approach integrates several relevant knowledge engineering methods such as ontology selection, semantic annotation and ontology alignment, and in particular, pairs semantic annotation with temporal annotation so that every fact describing a patient has temporal context, enabling the tracking of the patient's progression. Our method focuses on building edge-centric TKGs, given the abundance of downstream approaches built following that paradigm [20, 21, 22, 23, 24], thus aiding in improving data interoperability and future ML applications in healthcare.

The resulting TKG comprises four biomedical ontologies representing about 60,000 hospital stays with more than 27 million temporal facts.

2. Related Work

2.1. KG Construction

The methodology for constructing a KG should be tailored to relevant factors such as the domain, the intended application, any pre-existing data sources, and the construction theme. Building a KG from data sources includes two main steps: entity discovery, where entities acquired from the source data constitute the nodes of the KG, and triple extraction, where these entities are linked through their semantic relationships. These steps are usually preceded by data pre-processing and data cleaning. Many KGs include a semantic backbone composed of ontologies, and as such, the entity discovery step includes linking the entities to their types in the ontologies that compose the KG. Entity linking, also known as semantic annotation, involves associating data objects with ontology entities with well-defined semantics. Both more classical NLP [25, 26, 27] and neural approaches have been employed in entity linking [28, 29, 30, 31, 32].

2.2. Entity linking in EHRs

EHRs can contain both highly structured data such as clinical codes and terminology (e.g., SNOMED-CT), as well as short-form and long form free text variables. While controlled vocabularies allow for a straightforward inclusion in the KG, the annotation of free text poses additional challenges due to the complexity of the biomedical vocabulary, the abundance of polysemy, synonymy, acronyms and abbreviations, which result in ambiguity, and clinical free text variables often contain spelling mistakes. The typical entity linking process takes text as input and enriches it with entity mentions linked to nodes in a KG (or concepts in an ontology). The task is commonly split into entity mention detection and entity disambiguation sub-tasks.

Many clinical semantic annotators rely on a combination of strategies, including text processing, large-scale knowledge bases, semantic similarity measures and ML techniques [33, 34, 35, 36, 37]. Large Language Models (LLMs) have recently played a growing role in this task (e.g., [38, 39, 40]).

2.3. Temporal Knowledge Graphs

TKGs are an extension of traditional KGs [41] where facts about a domain include a temporal dimension. There are two major paradigms in capturing time in KGs: edge-centric and node-centric. In the edge-centric paradigm, the approaches are based on the notion of time-stamping an edge to capture the
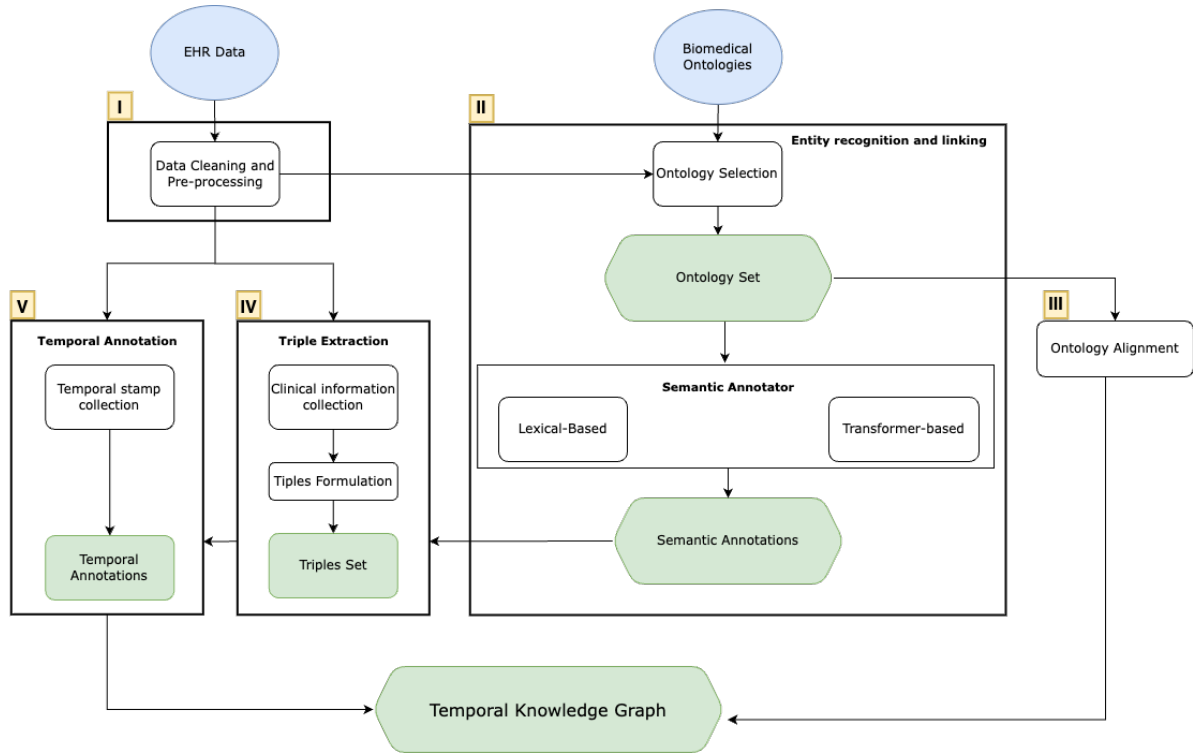
Figure 1: Five step methodology developed in this work. The method first proposes a cleaning of the data, consequently selects and aligns ontologies appropriate for annotation and annotates the entities to generate semantic triples and temporal facts to develop a final TKG.

validity of a fact. Most of the work that employs this paradigm is not specifically concerned with the challenges in modelling the temporal dimension in KGs but rather with how to explore it in the context of mining and ML applications. There are well-established data sets, such as YAGO or DBPedia, that already have published temporal knowledge graph resources for downstream tasks [42, 43, 44].

The node-centric paradigm captures the temporal dimension by modelling facts not as triples but as instances of an event class with properties that capture the validity of the event. An early notable effort in event modelling was the Simple Event Model (SEM) [45]. SEM represents data from the Web as events, in and across various domains, driven by minimal commitment (e.g., no cardinality restrictions, no functional or inverse functional properties, and Simple Knowledge Organization System (SKOS) [46] is used to link to other vocabularies or event models to avoid inheriting their constraints.) SEM has been used as the foundation for several event-based KGs, including EventKG[47] and Open EventKG [48].

3. Methods

Our decision to model the TKG as edge-centric was motivated both by the fact that there is no single central event in a hospital stay (but rather a sequence of several occurrences) and the fact that the majority of temporal KG mining approaches are applicable to edge-centric TKGs [49].

Following [50], we define an edge-centric TKG as $G = (E, R, T, F)$, where E, R, and T are the sets of entities, relations, and timestamps, respectively. $F \in E \times R \times E \times T$ is the set of all possible facts. Respecting this definition, facts can take two forms: *time-point facts*, denoted as $s = \{h, r, t, \tau\}$, where h, r, t, and τ are the head entity, the relation, the tail entity, and the timestamp, respectively; and *time-interval facts*, denoted as $s = \{h, r, t, < \tau_s, \tau_e >\}$, where h, r, t, and τ_s, τ_e are the head entity, the relation, the tail entity, the timestamp for the start of the time interval, and the timestamp for the end of the interval.

The methodology we designed includes five components, as shown in Figure 1. The first component is *Data pre-processing and cleaning* where textual terms from selected tables from the EHR database are prepared for entity linking. The second component is *Entity Linking*, which includes selecting the
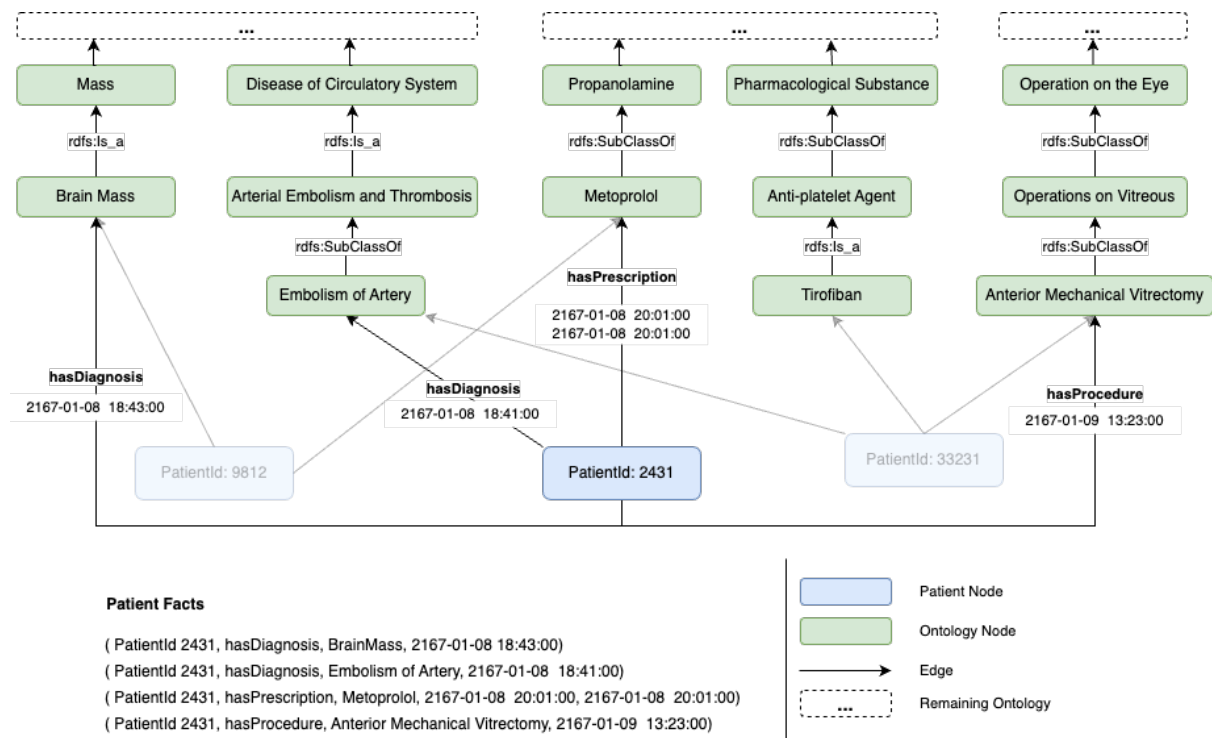
Figure 2: TKG portion with temporal facts for initial diagnosis, procedures and prescriptions. Patient facts are annotated to an ontology class with subsequent ontology relations represented in the TKG.

most appropriate ontologies using an external ontology repository and linking the textual terms to the most appropriate ontology classes. The third component is *Ontology Alignment*, which aligns the selected ontologies to increase their connectivity in the TKG. The fourth component, *Fact Creation*, creates KG facts to represent entries in the EHR, linking a patient stay to the appropriate ontology class that captures that specific entry. Finally, the fifth component is dedicated to the *Temporal annotation* of these facts, where every created fact is annotated with a time-point or time interval.

The resulting KG describes patients under multiple ontologies that cover the appropriate domains, with each fact about a patient including a temporal dimension. A subgraph representing a patient in the final TKG is given in Figure 2.

3.1. Data Pre-processing and Cleaning

Although some information in EHRs is recorded using controlled vocabularies or terminologies, such as the ICD terminology, many entries are filled in manually by clinicians. The goal of this step is to minimize the impact of errors and inconsistencies found in EHR entries in the form of clinical free text or terminology codes.

For clinical free text, we apply the following criteria:

- The empty and inconsistent entries (e.g., very short or with irrelevant symbols) are excluded.
- Incomplete entries are kept because the annotators will try to generate insightful annotations.
- Entries with more than one term are decoupled (e.g., multiple diagnoses split by comma).

For terminological codes, cleaning focuses on adherence to terminological constraints:

- Empty entries are excluded.
- Faulty formatted codes are corrected to follow the terminology guidelines (e.g.: E3924 is corrected to E392.4).

3.2. Entity Recognition and Linking

When EHR entries are in free text, the entities they refer to need to be linked to relevant biomedical ontologies. This step requires both *Ontology Selection* — to identify the relevant ontologies that best describe the text entries — and *Semantic Annotation* — to link each text entry to one or more relevant ontology classes.

3.2.1. Ontology Selection

Our approach employs an ontology selection method that, given a list of input textual terms and an ontology repository, outputs based on coverage, a sorted set of ontologies.

While different ontology selection methods can be used in this step, we employed the BioPortal Recommender platform [51] since it is linked to the largest repository available, which contains more than 1000 biomedical ontologies. The BioPortal Recommender successfully handles misspellings and uses a straightforward dictionary-based approach to create annotations of textual terms to ontology classes, which are then employed to assess how well an ontology describes a set of terms. We utilize the standard text processing capabilities and weight configurations of the Recommender. To operate the BioPortal Recommender, we use input information from a cleaned set, in line with BioPortal's guidelines, and an ontology repository that has been pre-selected to ensure domain relevance. This pre-selection facilitates the choice of an ontology for annotation based on its coverage, selecting the one that offers the best coverage.

3.2.2. Semantic Annotation

We propose two alternatives to the annotation process for the free text variable: a lexical similarity approach designed for full-text search, and a large language model-based approach that uses sentence-transformers for semantic search.

Lexical similarity-based approach

Lexical similarity can be measured through string similarity algorithms such as Jaccard Similarity, Cosine Similarity, Levenshtein Distance, etc. Given the large size of the vocabularies both in the EHR data and for each selected ontology, we developed an annotation approach based on the high-performing search engine ElasticSearch (ES)[52] to find the best matches between an ontology's vocabulary (i.e., its textual component in the form of entity labels) and the EHR textual entries. The search step for each EHR term outputs the six best-matching ontology classes. To select the single best matching class, we measure the Levenshtein Distance between the class label and the input term of every candidate match, and select the best scoring class provided its similarity is above a threshold of 0.6 (empirically determined).

Transformer-based approach

Lexical matching methods rely on keyword-based linking between input entities and document spaces [53], ignoring contextual and semantic information. However, vocabulary ambiguities, which are very common in clinical data, are challenging and lead to mismatches. An alternative for more informed matches is semantic search.

Semantic search is an alternative to traditional keyword-based searches [54] that incorporates semantic information in the search process [53]. Semantic search employs LLMs to generate numerical vector embeddings that represent textual terms, and uses vector operations to compare them. We used a pre-trained sentence transformer [55] — multi-qa-mpnet-base-dot-v1[1] – to generate dense vector representations for both ontology class labels and EHR terms and computed the similarity for each

---

[1] a MPNet [56] family model designed specifically for semantic search and trained on a large and multi-source question answering set. Fine tuned model available at https://huggingface.co/sentence-transformers/multi-qa-mpnet-base-dot-v1

pair using the dot product. Each EHR entry was annotated to the class ontology with the most similar vector.

3.3. Ontology Alignment

Ontology matching (OM) and alignment are operations to enhance the interoperability and integration of different domain data sources. OM focuses on establishing correspondences between different ontology classes, taking two distinct ontologies, and producing an alignment that is a collection of equivalence semantic relations or mappings of their entities [57, 58].

The necessity for ontology alignment in EHR-related data comes from the diversity of domains and ontologies. The alignment allows a gate between the different ontologies for interoperability and to facilitate downstream tasks that establish paths in a KG.

We employed the AgreementMakerLight (AML) system — one of the best performing OM systems [59] — to generate a fast and reliable alignment between all the biomedical ontologies used. For every alignment, each mapping is transformed into an equivalent statement between the two classes formulated as *'owl: equivalent class'* effectively linking the ontologies.

3.4. Triple Extraction

Triple extraction transforms clinical data structured according to the EHR model into the KG model. All the acquired and annotated entities will be mapped to the KG through relations to the patients and their hospital stays. Each patient is represented by a node (instance) in the KG with associated triples for properties.

For each patient, we first acquire the set of related EHR entries according to their type (e.g., diagnosis, drug prescriptions, etc.). For each entry, if there is a semantic annotation to an ontology class, a simple triple is constructed according to the KG model $t = \{patient, relation, ontology\_class\}$ and using the appropriate relation (see Figure 2).

3.5. Temporal Annotation

The temporal annotation adds temporal information to the patient's triples, capturing the validity period of the relation described by the triple. The process involves stamping the facts with the time stamps contained in the patient's entries on the data set.

On EHR data, the temporal information is represented in *datetime* variables formatted as $Y-m-d, H : M : S$. Each relation is described in the patient's data entry with either one or two time entries representing instantaneous facts or facts with a given duration or validity.

Following the definitions in section 3, to model instantaneous facts, we add a single time stamp to the original triple, resulting in a *time-point fact*, and to model a fact with a specific duration, we add two timestamps marking the beginning and end of the fact's validity, resulting in a *time-interval fact*.

The set of all temporal facts, ontologies and alignments represents the final TKG.

4. Results

We applied the proposed methodology to the generation of an ontology-rich TKG based on the MIMIC-III dataset. We created two versions of the TKG, one describing full hospital stays and the other only ICU stays.

4.1. Data

The medical data used in this project is sourced from the MIMIC-III dataset, a substantial, openly accessible database containing de-identified health data from patients admitted to the critical care units of Beth Israel Deaconess Medical Center between 2001 and 2012. MIMIC-III includes data from 53,423 unique hospital admissions for adult patients (aged 16 and older) and provides a wide array of information

such as demographics, hourly bedside vital sign measurements, laboratory test results, procedures, medications, caregiver notes, imaging reports, and mortality data (including post-discharge) [19].

4.1.1. Data Acquisition, Cleaning and Pre-Processing

The MIMIC-III data is available as a set of CSV files. These files were processed to extract the necessary features to describe each patient. We focus on five specific features to allow a simplistic yet realistic patient representation: Initial Diagnosis, Drug Prescriptions, Laboratory Events, Procedures, and Final Diagnosis.

The diagnosis is collected at admission and at discharge. At admission, the initial diagnosis is provided as a preliminary, free text entry and at discharge, the final diagnosis is coded using The International Classification of Diseases, 9th Revision, Clinical Modification (ICD9CM) glossary. Laboratory Events respect the Logical Observation Identifier Names and Codes (LOINC) terminology and contain information regarding laboratory-based measurements and include both in-hospital and out-of-hospital laboratory measurements from clinics the patient has visited. Drug Prescriptions detail the medication orders attributed to a patient, formatted according to National Drug Code (NDC). Procedures are coded and include the procedures undergone by patients, specifically ICD9CM procedures.

The five types of features extracted from MIMIC-III are thus primarily structured as clinical codes, except the initial diagnosis in short-form free text.

Table 1
Total Entity distribution after the data cleaning process.(*)Entries with multiple diagnoses where split

	MIMIC-III	Clean	Coverage (%)
Initial Diagnosis	58,977	84,746	143*
Drug Prescriptions	4,156,451	4,151,991	99
Laboratory Events	27,854,056	26,851,670	96
Final Diagnosis	651,048	651,048	100
Procedures	240,096	240,096	100

The cleaning process (section 3.1) results showcased in Table 1 point to a reduced data loss. Regarding the initial diagnosis, the cleaning process resulted in more diagnoses, since some entries had more than one diagnosis.

4.1.2. Ontology Selection

The only feature type that required ontology selection is the *Initial Diagnosis*, since all others respect particular controlled vocabularies. To match the clinical domain we pre-select *NCIT*, *SNOMEDCT*, *MeSH*, *MedDRA* and *EFO* as the biomedical ontologies to select from.

Table 2
Ontologies characteristics.

Ontology	Classes	Properties
NCIT	163,184	97
LOINC	288,940	160
DRON	690,772	131
ICD9CM	22,533	7

After running the BioPortal's Recommender, the ontology with the best coverage is the *NCIT* (Figure 3), reaching about 30% coverage. This apparently low coverage is not unexpected due to the limited dictionary-based approach of the Recommender. The selected set of ontologies to use in the TKG is

then NCIT, LOINC, ICD9CM and the Drug Ontology (DRON)[2]. Their characteristics are described in Table 2

4.2. Semantic Annotation

The annotation process is conducted on the entire set of cleaned and unique diagnoses collected from the MIMIC-III dataset, totaling 10,117 diagnoses. In our implementation, we perform two alternative processes: the Lexical similarity-based approach and the Transformer-based approach. The Initial Diagnosis is matched to the NCIT.
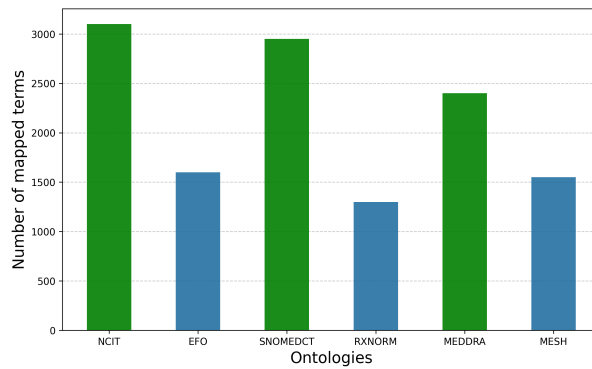

Figure 3: Bioportal Recommender's ontology coverage of the initial diagnosis set.
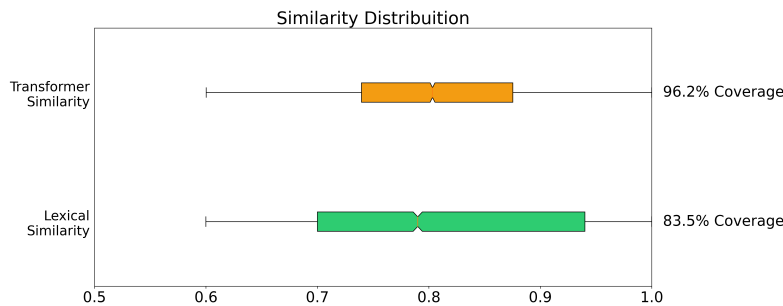

Figure 4: Similarity distribution over the mimic initial diagnosis entries for the transformer-based and lexical-based approaches.

4.2.1. Lexical similarity-based Annotation

The final similarity distribution for the NCIT-Diagnosis matches is shown in Figure 4. Lexical similarity has a median similarity score of 0.78, with the annotations skewed towards the higher lexical similarity values, meaning that the diagnosis are likely matched to a class with an appropriate label. The lexical annotation covers 83.5% of the full diagnosis set.

4.2.2. Transformer-based annotation

The coverage obtained by the Transformer-based annotation is higher than the lexical approach. The similarity distribution for the transformer-based approach has a median similarity score around 0.8 and as in the lexical approach, annotations are skewed towards the higher similarity values. The two approaches measure two different types of similarity. As such one can only argue that; the lexical similarity-based approach is capable of providing annotations that match the ontology class labels and

---

[2]Because NDC is a private repository, we replace it with DRON. Both concern the same information, and DRON is already prepared with information mapping class matches. Despite this change we were able to retain 75% of the full drugs set.

that the transformer approach is able to match entities to conceptually appropriate ontology classes based on their semantics. This recognition paired with the better coverage, motivated our decision to use the transformer annotations to build the TKG.

4.3. Ontology alignment

We generated alignments between all selected ontologies. The alignment results are shown in Table 3 and show a substantially larger alignment between ontologies with more related domains.

Table 3
AML automatic matcher alignments between NCIT, LOINC, ICD9CM, and DRON without any threshold constraints on the alignments.

	NCIT	LOINC	ICD9CM	DRON
NCIT		7,506	1,685	4,760
LOINC			354	3,354
ICD9CM				9

4.4. Temporal Annotation

The initial diagnosis, laboratory events, procedures, and final diagnosis are all instantaneous, and were represented as *time-point facts*. Drug prescriptions, however, have an associated duration and were captured as *time-interval facts*.

Table 4
Temporal Annotation Collected from the mimic data set.

	Hospital Stay		ICU Stay	
	Time Stamped Patient Annotations	**Time interval Patient Annotations**	**Time Stamped Patient Annotations**	**Time interval Patient Annotations**
NCIT	81,520	-	79,644	-
LOINC	21,322,274	-	15,212,374	-
DRON	-	2,773,939	-	1,606,234
ICD9CM D	634,709	-	129,203	-
ICD9CM P	237,947	-	54,605	-
Total	25,046,389	2,773,939	15,475,826	1,606,234

To create the two versions of the TKG we propose — full hospital stay and ICU stay — we extract two sets of temporal facts by filtering the full set according to their time stamps and MIMIC-III's specifications.

We apply the following criteria for the full hospital stay:

- Instantaneous annotations are selected for the full stay set only if their temporal stamp falls within the stay period.
- Prolonged annotations are selected for the full stay set if all or a part of their validity period falls within the stay period.

For the ICU stay, we apply the following criteria:

- Instantaneous annotations are selected to the ICU stay set if their temporal stamp precedes or falls within the period of the patient's ICU stay.
- Prolonged annotations are selected for the ICU stay set if all or a part of their validity period falls within the stay period or if their validity period precedes the ICU stay completely.

The resulting set sizes are shown in Tables 4 and illustrate a substantially large set of temporal facts for the MIMIC-III data set. The difference in size is expected since ICU stays tend to be shorter than the overall hospital stays.

4.5. TKG description

After applying our methodology to the MIMIC-III clinical data set we were able to generate the two TKGs. Table 5 describes some relevant statistics, with totals for the hospital stay TKG (which encapsulates the ICU stay plus additional information recorded outside of the ICU).

Table 5
TKG summary statistics.

			Hospital Stay TKG
Nodes	Classes	1,165,429	
	Hospital stays	58,966	1,224,395
	ICU stays	61,290	
Edges	Mappings	17,668	
	Hospital Stay Time stamped patient annotations	25,046,389	
	Hospital Stay Time interval patient annotations	2,773,939	27,837,996
	ICU Stay Time stamped patient annotations	15,475,826	
	ICU Stay Time interval patient annotations	1,606,234	

The resulting TKG represents patients and their hospital stays (including ICU stays) relating them to their diagnoses, prescribed drugs, conducted procedures and lab analyses which are in turn described by a relevant biomedical ontology. Each of these relations contains a temporal dimension, either a time-point or a time-interval. The TKG contains a total of 1.2 million nodes and more than 27 million edges that describe every patient stay. All code and resource links required to produce the TKGs are freely available [1]. Researchers are required to formally request access to the MIMIC-III dataset via a process documented on the MIMIC website [2].

5. Conclusions

One of the most challenging aspects of building a clinical TKG is the modelling decision regarding temporal data. Our methodology focuses on edge-centric TKGs with both time-point and time-interval facts to tackle the diversity in temporal data formats afforded by EHRs. Also challenging is the handling of text-based EHR entries that are not bound by any terminological constraints. The processing of such data is complex and entity recognition and linking requires the identification of suitable ontologies. Our methodology emphasizes the importance of proper ontology selection and alignment to maximize coverage and interoperability. Semantic and temporal annotation are crucial steps for linking clinical entities to semantically relevant ontology concepts, placing them at the appropriate time points or intervals.

We aim for the MIMIC-III TKG to serve as a resource for the community, supporting the development of TKG mining approaches for predictive tasks. This involves exploring the connections and nuances between the semantic context provided by biomedical ontologies and the temporal aspects extracted from the EHR. In particular, by transforming MIMIC-III data into a KG which includes several biomedical ontologies, we aim to support the application of knowledge-based explanations [60]. We believe this ability is crucial to support the uptake of artificial intelligence approaches in a clinical setting.

---

[1]at https://github.com/liseda-lab/clinical-temporal-kg under a GPL-3.0 license
[2]https://physionet.org/content/mimiciii/1.4/

Acknowledgments

References

[1] P. Jensen, L. Jensen, S. Brunak, Mining electronic health records: Towards better research applications and clinical care, Nature reviews. Genetics 13 (2012) 395–405. doi:10.1038/nrg3208.

[2] B. Goldstein, A. Navar, M. Pencina, J. Ioannidis, Opportunities and challenges in developing risk prediction models with electronic health records data: A systematic review, Journal of the American Medical Informatics Association 24 (2016) ocw042. doi:10.1093/jamia/ocw042.

[3] M. Scherpf, F. Gräßer, H. Malberg, S. Zaunseder, Predicting sepsis with a recurrent neural network using the mimic iii database, Computers in Biology and Medicine 113 (2019) 103395. doi:10.1016/j.compbiomed.2019.103395.

[4] B. Beaulieu-Jones, P. Orzechowski, J. Moore, Mapping patient trajectories using longitudinal extraction and deep learning in the mimic-iii critical care database, Pacific Symposium on Biocomputing. Pacific Symposium on Biocomputing 23 (2018) 123–132.

[5] Y.-W. Lin, Y. Zhou, F. Faghri, M. Shaw, R. Campbell, Analysis and prediction of unplanned intensive care unit readmission using recurrent neural networks with long short-term memory, PLOS ONE 14 (2019) e0218942. doi:10.1371/journal.pone.0218942.

[6] Q. Lu, T. H. Nguyen, D. Dou, Predicting Patient Readmission Risk from Medical Text via Knowledge Graph Enhanced Multiview Graph Convolution, Association for Computing Machinery, New York, NY, USA, 2021, p. 1990–1994. URL: https://doi.org/10.1145/3404835.3463062.

[7] F. Z. Smaili, X. Gao, R. Hoehndorf, OPA2Vec: combining formal and informal content of biomedical ontologies to improve similarity-based prediction, Bioinformatics 35 (2018) 2133–2140. URL: https://doi.org/10.1093/bioinformatics/bty933. doi:10.1093/bioinformatics/bty933. arXiv:https://academic.oup.com/bioinformatics/article-pdf/35/12/2133/48935045/bioinformatics_35_12_2133.pdf.

[8] S. Auer, V. Kovtun, M. Prinz, A. Kasprzik, M. Stocker, M.-E. Vidal, Towards a knowledge graph for science, 2018, pp. 1–6. doi:10.1145/3227609.3227689.

[9] Q. Wang, Z. Mao, B. Wang, L. Guo, Knowledge graph embedding: A survey of approaches and applications, IEEE Transactions on Knowledge and Data Engineering PP (2017) 1–1. doi:10.1109/TKDE.2017.2754499.

[10] L. Ehrlinger, W. Wöß, Towards a definition of knowledge graphs, 2016.

[11] A. Kiryakov, B. Popov, I. Terziev, D. Manov, D. Ognyanoff, Semantic annotation, indexing, and retrieval, Web Semantics: Science, Services and Agents on the World Wide Web 2 (2004) 49–79. doi:10.1016/j.websem.2004.07.005.

[12] L. Huang, A. Shea, H. Qian, A. Masurkar, H. Deng, D. Liu, Patient clustering improves efficiency of federated machine learning to predict mortality and hospital stay time using distributed electronic medical records, Journal of Biomedical Informatics 99 (2019) 103291. doi:10.1016/j.jbi.2019.103291.

[13] H. Suresh, J. Gong, J. Guttag, Learning tasks for multitask learning: Heterogenous patient populations in the icu (2018). doi:10.1145/3219819.3219930.

[14] R. Zhang, W. Zhang, N. Liu, J. Wang, Susceptible temporal patterns discovery for electronic health records via adversarial attack, in: Database Systems for Advanced Applications: 26th International

Conference, DASFAA 2021, Taipei, Taiwan, April 11–14, 2021, Proceedings, Part III 26, Springer, 2021, pp. 429–444.

[15] F. Xie, H. Yuan, Y. Ning, M. Ong, M. Feng, W. Hsu, B. Chakraborty, N. Liu, Deep learning for temporal data representation in electronic health records: A systematic review of challenges and methodologies, Journal of Biomedical Informatics 126 (2021) 103980. doi:10.1016/j.jbi.2021.103980.

[16] A. Woods, C. Meyer, B. Sauer, B. Cohen, Mining time-stamped electronic health records using referenced sequences, arXiv preprint arXiv:2007.14336 (2020).

[17] J. Schrodt, A. Dudchenko, P. Knaup-Gregori, M. Ganzinger, Graph-representation of patient data: a systematic literature review, Journal of medical systems 44 (2020) 86.

[18] R. M. Carvalho, D. Oliveira, C. Pesquita, Knowledge graph embeddings for icu readmission prediction, BMC Medical Informatics and Decision Making 23 (2023) 12.

[19] A. Johnson, T. Pollard, L. Shen, L.-w. Lehman, M. Feng, M. Ghassemi, B. Moody, P. Szolovits, L. Celi, R. Mark, Mimic-iii, a freely accessible critical care database, Scientific Data 3 (2016) 160035. doi:10.1038/sdata.2016.35.

[20] D. Liben-Nowell, J. Kleinberg, The link prediction problem for social networks, in: Proceedings of the twelfth international conference on Information and knowledge management, 2003, pp. 556–559.

[21] L. Yao, L. Wang, L. Pan, K. Yao, Link prediction based on common-neighbors for dynamic social network, Procedia Computer Science 83 (2016) 82–89.

[22] Z. Han, Y. Ma, Y. Wang, S. Günnemann, V. Tresp, Graph hawkes neural network for forecasting on temporal knowledge graphs, in: Automated Knowledge Base Construction, 2020.

[23] R. Goel, S. M. Kazemi, M. Brubaker, P. Poupart, Diachronic embedding for temporal knowledge graph completion, in: Proceedings of the AAAI conference on artificial intelligence, volume 34, 2020, pp. 3988–3995.

[24] S. S. Dasgupta, S. N. Ray, P. Talukdar, Hyte: Hyperplane-based temporally aware knowledge graph embedding, in: Proceedings of the 2018 conference on empirical methods in natural language processing, 2018, pp. 2001–2011.

[25] J. Jovanovic, E. Bagheri, Semantic annotation in biomedicine: The current landscape, Journal of Biomedical Semantics 8 (2017). doi:10.1186/s13326-017-0153-x.

[26] P. Stenetorp, S. Pyysalo, G. Topić, T. Ohta, S. Ananiadou, J. Tsujii, Brat: a web-based tool for nlp-assisted text annotation, in: Proceedings of the Demonstrations at the 13th Conference of the European Chapter of the Association for Computational Linguistics, 2012, pp. 102–107.

[27] V. Basile, J. Bos, K. Evang, N. Venhuizen, Developing a large semantically annotated corpus, in: LREC 2012, Eighth International Conference on Language Resources and Evaluation, 2012.

[28] W. Shen, Y. Li, Y. Liu, J. Han, J. Wang, X. Yuan, Entity linking meets deep learning: Techniques and solutions, IEEE Transactions on Knowledge and Data Engineering 35 (2021) 2556–2578.

[29] Ö. Sevgili, A. Shelmanov, M. Arkhipov, A. Panchenko, C. Biemann, Neural entity linking: A survey of models based on deep learning, Semantic Web 13 (2022) 527–570.

[30] M. C. Phan, A. Sun, Y. Tay, J. Han, C. Li, Neupl: Attention-based semantic matching and pair-linking for entity disambiguation, in: Proceedings of the 2017 ACM on Conference on Information and Knowledge Management, 2017, pp. 1667–1676.

[31] L. Chen, G. Varoquaux, F. M. Suchanek, A lightweight neural model for biomedical entity linking, in: Proceedings of the AAAI conference on artificial intelligence, volume 35, 2021, pp. 12657–12665.

[32] J. Raiman, O. Raiman, Deeptype: multilingual entity linking by neural type system evolution, in: Proceedings of the AAAI Conference on Artificial Intelligence, volume 32, 2018.

[33] H. Wu, G. Toti, K. I. Morley, Z. M. Ibrahim, A. Folarin, R. Jackson, I. Kartoglu, A. Agrawal, C. Stringer, D. Gale, et al., Semehr: A general-purpose semantic search system to surface semantic data from clinical notes for tailored care, trial recruitment, and clinical research, Journal of the American Medical Informatics Association 25 (2018) 530–537.

[34] E. Soysal, J. Wang, M. Jiang, Y. Wu, S. Pakhomov, H. Liu, H. Xu, Clamp–a toolkit for efficiently building customized clinical natural language processing pipelines, Journal of the American

Medical Informatics Association 25 (2018) 331–336.

[35] R. Gazzotti, C. Faron-Zucker, F. Gandon, V. Lacroix-Hugues, D. Darmon, Injection of automatically selected dbpedia subjects in electronic medical records to boost hospitalization prediction, 2020, pp. 2013–2020. doi:10.1145/3341105.3373932.

[36] Z. Kraljevic, T. Searle, A. Shek, L. Roguski, K. Noor, D. Bean, A. Mascio, L. Zhu, A. A. Folarin, A. Roberts, et al., Multi-domain clinical natural language processing with medcat: the medical concept annotation toolkit, Artificial intelligence in medicine 117 (2021) 102083.

[37] A. Arbabi, D. Adams, S. Fidler, M. Brudno, Identifying Clinical Terms in Free-Text Notes Using Ontology-Guided Machine Learning, 2019, pp. 19–34. doi:10.1007/978-3-030-17083-7_2.

[38] N. Gupta, S. Singh, D. Roth, Entity linking via joint encoding of types, descriptions, and context, in: Proceedings of the 2017 conference on empirical methods in natural language processing, 2017, pp. 2681–2690.

[39] H. Yuan, Z. Yuan, R. Gan, J. Zhang, Y. Xie, S. Yu, Biobart: Pretraining and evaluation of a biomedical generative language model, in: Proceedings of the 21st Workshop on Biomedical Language Processing, 2022, pp. 97–109.

[40] D. Kartchner, J. Deng, S. Lohiya, T. Kopparthi, P. Bathala, D. Domingo-Fernández, C. S. Mitchell, A comprehensive evaluation of biomedical entity linking models, in: Proceedings of the Conference on Empirical Methods in Natural Language Processing. Conference on Empirical Methods in Natural Language Processing, volume 2023, NIH Public Access, 2023, p. 14462.

[41] J. Zhang, S. Liang, Y. Sheng, J. Shao, Temporal knowledge graph representation learning with local and global evolutions, Knowledge-Based Systems 251 (2022) 109234.

[42] Z. Li, X. Jin, W. Li, S. Guan, J. Guo, H. Shen, Y. Wang, X. Cheng, Temporal knowledge graph reasoning based on evolutional representation learning, in: Proceedings of the 44th international ACM SIGIR conference on research and development in information retrieval, 2021, pp. 408–417.

[43] H. Dong, P. Wang, M. Xiao, Z. Ning, P. Wang, Y. Zhou, Temporal inductive path neural network for temporal knowledge graph reasoning, Artificial Intelligence 329 (2024) 104085.

[44] H. Sun, J. Zhong, Y. Ma, Z. Han, K. He, Timetraveler: Reinforcement learning for temporal knowledge graph forecasting, arXiv preprint arXiv:2109.04101 (2021).

[45] W. R. Van Hage, V. Malaisé, R. Segers, L. Hollink, G. Schreiber, Design and use of the simple event model (sem), Journal of Web Semantics 9 (2011) 128–136.

[46] A. Miles, S. Bechhofer, Skos simple knowledge organization system reference, W3C recommendation (2009).

[47] S. Gottschalk, E. Demidova, Eventkg–the hub of event knowledge on the web–and biographical timeline generation, Semantic Web 10 (2019) 1039–1070.

[48] S. Gottschalk, E. Kacupaj, S. Abdollahi, D. Alves, G. Amaral, E. Koutsiana, T. Kuculo, D. Major, C. Mello, G. S. Cheema, et al., Oekg: The open event knowledge graph, in: CLEOPATRA 2021 Cross-lingual Event-centric Open Analytics 2021, April 12 2021, Ljubiljana, Slovenia, volume 2829, Aachen, Germany: RWTH Aachen, 2021.

[49] S. Ji, S. Pan, E. Cambria, P. Marttinen, S. Y. Philip, A survey on knowledge graphs: Representation, acquisition, and applications, IEEE transactions on neural networks and learning systems 33 (2021) 494–514.

[50] B. Cai, Y. Xiang, L. Gao, H. Zhang, Y. Li, J. Li, Temporal knowledge graph completion: A survey, arXiv preprint arXiv:2201.08236 (2022).

[51] M. Martínez-Romero, C. Jonquet, M. O'Connor, J. Graybeal, A. Pazos, M. Musen, Ncbo ontology recommender 2.0: An enhanced approach for biomedical ontology recommendation, Journal of Biomedical Semantics 8 (2017). doi:10.1186/s13326-017-0128-y.

[52] O. Kononenko, O. Baysal, R. Holmes, M. Godfrey, Mining modern repositories with elasticsearch (2014). doi:10.1145/2597073.2597091.

[53] F. Lashkari, E. Bagheri, A. A. Ghorbani, Neural embedding-based indices for semantic search, Information Processing & Management 56 (2019) 733–755. URL: https://www.sciencedirect.com/science/article/pii/S0306457318302413. doi:https://doi.org/10.1016/j.ipm.2018.10.015.

[54] R. Guha, R. McCool, E. Miller, Semantic search, in: Proceedings of the 12th international conference

on World Wide Web, 2003, pp. 700–709.

[55] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz, et al., Huggingface's transformers: State-of-the-art natural language processing, arXiv preprint arXiv:1910.03771 (2019).

[56] K. Song, X. Tan, T. Qin, J. Lu, T.-Y. Liu, Mpnet: Masked and permuted pre-training for language understanding, Advances in neural information processing systems 33 (2020) 16857–16867.

[57] J. Euzenat, P. Shvaiko, J. Euzenat, P. Shvaiko, Matching strategies, Ontology Matching (2013) 149–197.

[58] D. Faria, C. Pesquita, E. Santos, M. Palmonari, I. Cruz, F. Couto, The agreementmakerlight ontology matching system, volume 8185, 2013. doi:10.1007/978-3-642-41030-7_38.

[59] D. Faria, E. Santos, B. S. Balasubramani, M. C. Silva, F. M. Couto, C. Pesquita, Agreementmakerlight, Semantic Web (2023) 1–13.

[60] S. Chari, D. M. Gruen, O. Seneviratne, D. L. McGuinness, Directions for explainable knowledge-enabled systems, in: Knowledge graphs for explainable artificial intelligence: Foundations, applications and challenges, IOS Press, 2020, pp. 245–261.