

Clinical Data Goes MEDS? Let's OWL make sense of it

Alberto Marfoglio^{1,2}, Jong Ho Jhee² and Adrien Coulet^{2,*}

¹Dept. of Computer Science and Engineering - DISI, University of Bologna, Bologna, Italy

²Inria, Inserm, Université Paris Cité, HeKA, UMR 1346, Paris, France

Abstract

The application of machine learning on healthcare data is often hindered by the lack of standardized and semantically explicit representation, leading to limited interoperability and reproducibility across datasets and experiments. The Medical Event Data Standard (MEDS) addresses these issues by introducing a minimal, event-centric data model designed for reproducible machine-learning workflows from health data. However, MEDS is defined as a data-format specification and does not natively provide integration with the Semantic Web ecosystem. In this article, we introduce MEDS-OWL, a lightweight OWL ontology that provides formal concepts and relations to represent MEDS datasets as RDF graphs. Additionally, we implemented `meds2rdf`, a Python conversion library that transforms MEDS events into RDF graphs, ensuring conformance with the ontology. We evaluate the proposed approach on two datasets: a synthetic clinical cohort describing care pathways for ruptured intracranial aneurysms, and a real-world subset of MIMIC-IV. To assess semantic consistency, we performed a SHACL validation against the resulting knowledge graphs. The first release of MEDS-OWL comprises 13 classes, 10 object properties, 20 data properties, and 24 OWL axioms. Combined with `meds2rdf`, it enables data transformation into FAIR-aligned datasets, provenance-aware publishing, and interoperability of event-based clinical data. By bridging MEDS with the Semantic Web, this work contributes a reusable semantic layer for event-based clinical data and establishes a robust foundation for subsequent graph-based analytics.

Keywords

MEDS, Semantic Web, Knowledge graph, Healthcare interoperability, OWL ontology, FAIR data

1. Introduction

Machine learning (ML) in the healthcare domain is hindered by the lack of a standardized representation for structured Electronic Health Record (EHR) data, resulting in dataset-specific preprocessing practices, limited reproducibility, and poor model portability across datasets. The Medical Event Data Standard (MEDS) [1, 2] addresses this gap by introducing a unified, event-centric schema for longitudinal health data that is explicitly designed to support reproducible and efficient ML workflows. MEDS encodes each event as a minimal tuple consisting of a subject identifier, a code, and an optional temporal and value attributes. This minimal and pragmatically oriented design makes MEDS particularly suitable for benchmarking, multi-center studies, and the development of foundational models. For these reasons, MEDS is being adopted in various ML projects as a common schema for EHR data and has been utilized in a transformer-based model for predicting future events in patient healthcare trajectories [3, 4].

In its native form, MEDS is a data-format specification and does not provide ontology-level semantics. The Semantic Web paradigm offers a complementary approach to address this limitation [5, 6, 7]. First, ontologies expressed in Resource Description Framework (RDF) and Web Ontology Language (OWL) provide an interoperable, machine-interpretable lingua franca for representing entities and their relationships. Second, they support semantic querying and formal reasoning. Third, they support the construction of knowledge graphs (KGs) that facilitate the explicit linkage of dataset elements to external terminologies such as LOINC, SNOMED CT, and ATC [8]. Projecting MEDS datasets into RDF would thus enable their semantic integration and validation, while also creating opportunities to explore graph-based ML on event-oriented datasets. Indeed, patient representations can be learned from graph representations with several advantages. KG embeddings map graph nodes and edges

SWAT4HCLS 2026

*Corresponding author.

✉ alberto.marfoglio2@unibo.it (A. Marfoglio); jong-ho.jhee@inria.fr (J. H. Jhee); adrien.coulet@inria.fr (A. Coulet)

🆔 0009-0000-5857-2376 (A. Marfoglio); 0000-0001-8887-8149 (J. H. Jhee); 0000-0002-1466-062X (A. Coulet)



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

into a continuous vector space [9], enabling the representation of complex relational structures in a form amenable to computational analysis. Such approaches have been successfully applied to tasks including link prediction, node classification, and clustering [10]. In the clinical domain, relational models such as Relational Graph Convolutional Networks (RGCNs) have been shown to outperform tabular baselines when the underlying schema aligns with task-specific requirements; for instance, compact and patient-centered representations for outcome prediction [11]. These observations motivate the present work.

While several data models and ontologies, such as FHIR [12], OMOP CDM [13], and Phenopackets [14], have been proposed to harmonize clinical data, RDF transformations of these models have seen limited adoption [13, 15], in part due to loosely specified semantics or incomplete coverage of the clinical domain. Patient-centric RDF-based representations such as SPHN [16] and CARE-SM [17] illustrate how KGs can facilitate interoperability, provenance modeling, and temporal reasoning. However, these models are often tightly coupled to specific use cases or semantic frameworks.

In this study, we propose to bridge MEDS with the Semantic Web formalism. We report on the development of two complementary components: MEDS-OWL, an OWL ontology that captures the core concepts and relations of MEDS [18], and `meds2rdf`, a Python library that transforms MEDS-compliant datasets into RDF graphs in accordance with the ontology [19]. Together, these components enable interoperable, reproducible, and FAIR-aligned representations of event-based clinical data.

The remainder of this article is organized as follows. Section 2 describes the methodology adopted for the development of the MEDS-OWL ontology and the accompanying conversion library. Section 3 presents the resulting ontology and the `meds2rdf` library in action. Finally, Section 4 discusses the results and outlines directions for future work.

2. Methodology

The MEDS-OWL ontology was developed following the principled and iterative methodology described by Noy and McGuinness in “Ontology Development 101” [20], using the Protégé ontology editor [21]. The scope of the ontology encompasses the representation of EHR and claims-style medical events as defined by the MEDS conceptual model, together with the dataset- and experiment-level metadata required to support reproducible ML workflows. As part of this process, we formulated a set of competency questions to guide the modeling choices and to ensure that the ontology adequately supports the intended use cases. An example of such a competency question is: “*Given a medical event, which subject does it pertain to, when did it occur, and which canonical code and value modalities are associated with it?*”. These competency questions informed the identification and formalization of the core MEDS classes, namely `Subject`, `Event`, `Code`, `DatasetMetadata`, `SubjectSplit`, `SubjectLabel`, and `ValueModality`. A detailed description of these Seven core classes, together with illustrative examples, are provided in Table 1.

To clarify how these ontology classes relate to the original MEDS data model, Table 2 summarizes the correspondence between the five primary MEDS schema components and the MEDS-OWL classes derived from them. In MEDS, each schema is defined as an Apache Arrow schema, specifying the required columns and data types for MEDS data files.

The OWL classes and properties we developed closely follow the MEDS data model, while reusing as much as possible established external ontologies to maximize interoperability. Specifically, we reused the Data Catalogue Vocabulary (DCAT), DCTERMS, and PROV ontology (PROV-O) to represent datasets, ETL (Extraction, Transformation, Loading) processes, and provenance metadata, respectively, thereby ensuring alignment with widely adopted ontologies. Finally, for MEDS fields specified as optional (nullable), such as time, the ontology captures this constraint through upper-bound cardinality restrictions (`owl:maxCardinality 1`) without corresponding minimum bound or existential constraints, thus permitting zero or one value, while forbidding multiple values.

We developed a Python library named `meds2rdf` [19] to implement the mapping rules necessary to transform a MEDS dataset into an RDF graph that follows the terms of the MEDS-OWL ontology.

Table 1
Core MEDS-OWL classes and their interpretation.

Class	Description	Example
Subject	Primary entity under observation, typically corresponding to a patient, admission, or longitudinal unit to which events are attributed.	Patient #10234
Event	Atomic clinical observation representing a single MEDS row, characterized by a subject, a coded concept, and optional temporal and value attributes.	Administration of nimodipine at admission
Code	An entity used to type events, which includes canonical identifiers, descriptions, and hierarchical relations, with optional linkage to external terminologies.	ATC:C08CA06 (Nimodipine)
DatasetMetadata	Dataset-level description and provenance information, capturing identity, versioning, creation time, and ETL-related metadata.	X dataset, in MEDS format, v1.0
SubjectSplit	A grouping of subjects in several disjoint sets aimed for an ML experiment.	Split for experiment 5 into a train, eval, and test sets
SubjectLabel	A supervised learning instance defined by a subject identifier, an as-of prediction time (the latest time up to which data may be used), and exactly one observed ground truth label value y .	Subject #10234, prediction time = discharge date, categorical label = rehabilitation
ValueModality	Generic, extensible entity that links an Event to additional value modalities (for example, images, waveform recordings, audio clips, or other non-tabular payloads).	Chest X-ray image (DICOM series URI)

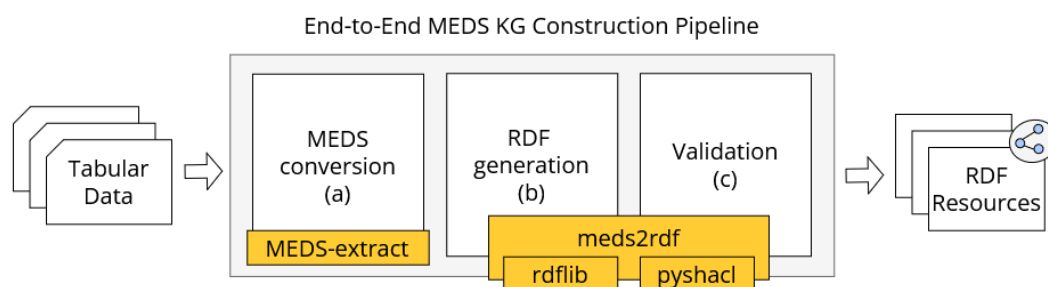


Figure 1: Overview of the end-to-end MEDS KG construction pipeline, illustrating (a) normalization and extraction of tabular data into MEDS event tuples (optional, if data are already in MEDS), (b) transformation into RDF KGs compliant with the MEDS-OWL ontology, and (c) SHACL-based validation and serialization into standard RDF formats.

Figure 1 presents the end-to-end conversion workflow implemented leveraging `meds2rdf`. The pipeline consists of three main stages. (a) First, source EHR data are ingested by the `MEDS-extract` component [2], which applies a declarative conversion template to map dataset-specific features into MEDS event tuples. Note that this first step is optional, and only necessary when healthcare data are not already available in MEDS. (b) In the subsequent RDF generation stage, MEDS-formatted data are transformed into a KG conforming to the MEDS-OWL ontology using the `meds2rdf` library. The mapping module of `meds2rdf` implements a deterministic, table-driven transformation from MEDS tuples to RDF. It exploits manually curated mapping tables that relate elements of the MEDS schema to corresponding classes and properties in MEDS-OWL. These tables constitute a central contribution of this work, as they provide an implementation-independent and ontology-aligned specification that formalizes the MEDS-to-RDF mapping. Moreover, in this phase, `meds2rdf` focuses on syntactic correctness and local mapping guarantees, such as required-field checks and type normalization. The output of this stage is an in-memory RDF graph constructed using the `rdflib`. Graph data structure. (c) Next, the

Table 2

Mapping from MEDS schema components to MEDS-OWL classes. See Table 1 for class definitions and examples.

MEDS schema	Brief description	Mapped MEDS-OWL class(es)
DataSchema	Core medical event table describing sequences of subject observations, where each row corresponds to a single clinical event.	<code>meds:Event</code> as the primary representation of each MEDS row. <code>meds:ValueModality</code> is used to model additional or non-tabular payloads associated with an event (e.g., images, waveforms, signals). See Table 3 for row-to-triple mappings.
DatasetMetadataSchema	Dataset-level metadata captures dataset identity, versioning, creation time, licensing, and ETL provenance information.	<code>meds:DatasetMetadata</code> , aligned with <code>dcterms</code> and <code>dcat</code> , and linked to ETL activities via <code>prov:wasGeneratedBy</code> . See Table 5.
CodeMetadataSchema	Metadata describing the controlled vocabularies and codes used to type medical events, including canonical identifiers and hierarchical relations.	<code>meds:Code</code> , with <code>meds:codeString</code> as a canonical identifier and optional <code>meds:parentCode</code> relations to represent code hierarchies. See Table 4.
SubjectSplitSchema	Specification of how subjects are assigned to mutually exclusive subsets (e.g., training, tuning, held-out) within a given machine learning experiment.	<code>meds:SubjectSplit</code> , with subjects linked via <code>meds:assignedSplit</code> . Multiple splits may coexist across experiments, while each subject belongs to exactly one subset per split. See Table 6.
LabelSchema	Definition of supervised learning labels associated with a subject at a specific prediction time, including the observed ground truth value.	<code>meds:SubjectLabel</code> , capturing the prediction cutoff time (<code>meds:predictionTime</code>) and exactly one label value using datatype-specific properties. See Table 6.
(subject identifier)	Subject identifiers appear across multiple MEDS schemas and provide the logical join key between events, splits, and labels.	<code>meds:Subject</code> , representing the canonical subject entity via <code>meds:subjectId</code> , and serving as the hub linking <code>Event</code> , <code>SubjectLabel</code> , and <code>SubjectSplit</code> .

generated graph is validated against a SHACL schema using the `pySHACL` validator, which is integrated within `meds2rdf`. The schema declaratively encodes the semantic constraints of MEDS-OWL, and the validation process ensures the enforcement of both global and cross-entity invariants. These include mutually exclusive property patterns, cardinality constraints, and referential consistency across entities. This approach allows semantic validation to be performed before graph materialization. Finally, RDF graphs that successfully pass SHACL validation may be serialized as an optional final step in standard RDF formats, including Turtle, RDF/XML, and N-Triples.

Integrating SHACL validation within `meds2rdf` enables the ontology and its constraints to evolve independently of the mapping logic. This design choice is crucial given that MEDS-OWL is currently at its first public version and is expected to undergo further revisions in alignment with the evolving MEDS specification. By leveraging the declarative nature of SHACL, future updates to the ontology or constraint definitions can be accommodated by modifying the SHACL schema alone, without requiring changes to the conversion code.

3. Results and examples of usage

MEDS-OWL in its first release comprehends 13 classes, 10 object properties, 20 data properties, and 24 OWL axioms, and it is accompanied by example RDF instances and a SHACL validation suite. MEDS-OWL is illustrated in Figure 2 and its documentation is available on Zenodo and GitHub [18].

Mappings between MEDS core data schema elements and MEDS-OWL are reported in Tables 3-6. Table 3 reports those to the `meds:Event` class. In MEDS-OWL, events are required to be associated

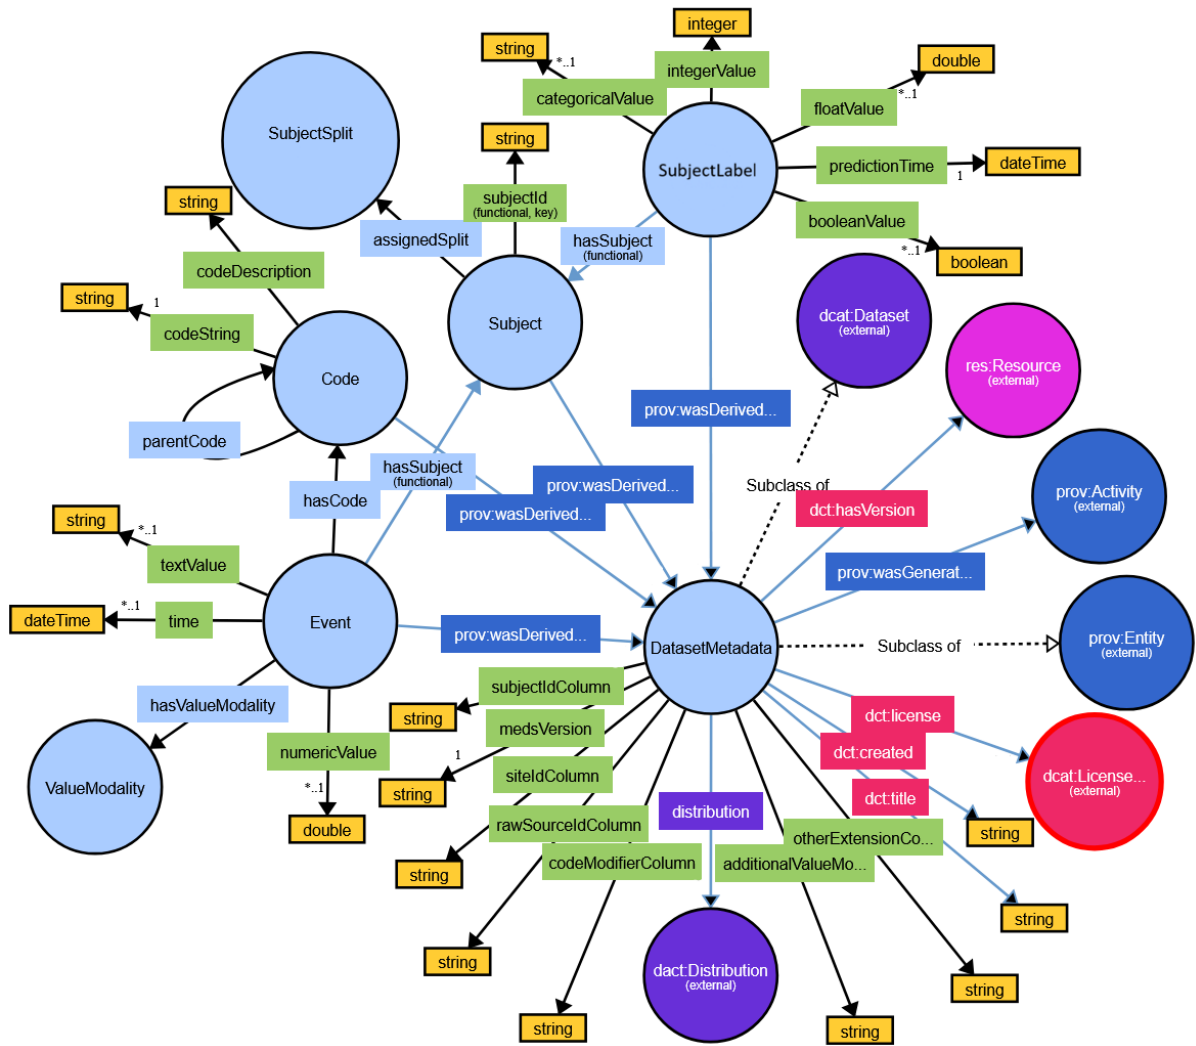


Figure 2: Overview of the MEDS-OWL ontology. Azure nodes and edges denote locally defined classes and object properties, respectively. Black edges with green labels are datatype properties. Nodes and edges colored differently indicate imported entities from either the Data Catalogue Vocabulary (dcat prefix), Dublin Core Terms (dct prefix) or PROV-O (prefix prov). Dotted edges represent subclassOf predicate. The figure is drawn with WebVOWL [22].

with exactly one subject and exactly one code, while temporal and value-related attributes are optional and for this reason constrained by cardinality restrictions. Table 4 reports mappings to the `meds:Code` class. Code metadata is modeled explicitly as individuals of type `meds:Code`, which enables reusability, hierarchical organization, and linkage to external terminologies. Dataset-level provenance and conversion metadata are represented by the `meds:DatasetMetadata` class, which is aligned with both `dcat:Dataset` and `prov:Entity`. Cardinality constraints are enforced at the OWL level. The corresponding mappings are shown in Table 5. Finally, Table 6 reports the mapping for dataset splits and supervised learning labels. Dataset partitions are modeled as predefined individuals of class `meds:SubjectSplit`, while label samples are represented as instances of class `meds:SubjectLabel` with exactly one label value. The **meds2rdf** library that implements the mapping rules is available on Zenodo and GitHub [19].

Neurovasc example of usage. To evaluate `meds2rdf`, we first employed a synthetic EHR dataset, here named *neurovasc*, that describes care pathways of 10,000 patients hospitalized for a ruptured intracranial aneurysm [11]. *neurovasc* is an open dataset generated based on real-world EHR from Nantes University Hospital. Its patients are associated with 30 clinical features, including 8 time-stamped events representing the interval from hospital admission to their first occurrence, and 22 non-temporal

Table 3

Mapping of the elements of the MEDS DataSchema tables to `meds:Event` and `meds:Subject` classes, with descriptions. For each tuple (or row) of DataSchema, a list of associated RDF triples is generated. Card. stands for the cardinality of the relationship.

MEDS element	Description	RDF triple	Card.
event_row e	The row representing an event e of the subject s.	?e a meds:Event.	1
subject_id id	The identifier of the subject s (typically the patient).	?s a meds:Subject. ?s meds:subjectId ?id^^xsd:string. ?e meds:hasSubject ?s.	1
time t	The time of the measurement, nullable for static measurements.	?e meds:time ?t^^xsd:dateTime.	0..1
code c	The categorical descriptor of the measurement (e.g., lab test, diagnosis code).	?c a meds:Code. ?e meds:hasCode ?c.	1
code_string cs	The string representation of the code.	?e meds:codeString ?cs^^xsd:string.	1
numeric_value nv	A numeric value associated with the measurement (e.g., lab test result), nullable.	?e meds:numericValue ?nv^^xsd:double.	0..1
text_value tv	A text value associated with the measurement (e.g., clinical note, text-based test result), nullable.	?e meds:textValue ?tv^^xsd:string.	0..1

Table 4

Mapping of the elements of the MEDS CodeMetadataSchema table to `meds:Code` class. For each tuple (or row) of CodeMetadataSchema, a list of associated RDF triples is generated.

MEDS element	Description	RDF triple	Card.
code_row c	The row representing a code c in the dataset.	?c a meds:Code.	1
code_string cs	The categorical descriptor of a possible measurement; links to the DataSchema code column.	?c meds:codeString ?cs^^xsd:string.	1
description d	A human-readable description of what the code represents.	?c meds:codeDescription ?d^^xsd:string.	0..1
parent_codes p	One or more higher-level or parent codes in an ontological hierarchy, possibly linking to other codes in MEDS or external vocabularies (e.g., OMOP CDM).	?p a meds:Code. ?c meds:parentCode ?p.	0..*

demographic or historical features. Patient outcomes are categorized in one of three values: Back Home, Rehabilitation, or Death, depending on the patient's status at the end of the hospital stay. The dataset preserves realistic correlations and event transitions, making it suitable for testing ontology-driven data representation of clinical pathways and basic analytics.

The conversion of the *neurovasc* dataset illustrates the practical utility of MEDS-OWL. As shown in Figure 3, the workflow depicts the output of each conversion stage, beginning with a simplified representation of the original *neurovasc* data. To provide visual clarity, only three patient features are shown: age, administration of nimodipine (ATC code: C08CA06), and a tracheostomy procedure (ICD-10-PCS code: 0BH17EZ). (a) Following an initial preprocessing step, (b) the data are passed to

Table 5

Mapping metadata of a transformed MEDS dataset (`DatasetMetadataSchema`) to the `meds:DatasetMetdata` class. For each dataset generated, a list of associated “meta” RDF triples is generated.

MEDS element	Description	RDF triple	Card.
dataset ds	The dataset converted to RDF.	?ds a meds:DatasetMetadata.	1
dataset_name n	The name of the dataset ds.	?ds dct:title ?n^^xsd:string.	1
dataset_version dsv	The version of the dataset.	?ds dct:hasVersion ?dsv.	0..1
meds_version mv	The version of the MEDS schema.	?ds meds:medsVersion ?mv^^xsd:string.	1
created_at	ISO 8601 creation date and time.	?ds dct:created ?time^^xsd:dateTime.	1
license	License under which the dataset is released.	?li a dct:LicenseDocument. ?ds dct:license ?li.	0..1
location_uri	URI(s) where the dataset is hosted.	?ds dcat:distribution ?dist. ?dist dcat:downloadURL ?dw.	0..*
description_uri	URI(s) referencing a detailed description of the dataset.	?dist dcat:accessURL ?au.	0..*
etl_name etln	Name of the ETL process etl that generated the dataset.	?etl a prov:Activity. ?etl rdfs:label ?etln. ?ds prov:wasGeneratedBy ?etl.	0..1
etl_version etlv	Version of the ETL process.	?etl dct:hasVersion ?etlv.	0..1

MEDS-extract and converted into the MEDS tabular format, where each feature is represented as a distinct event optionally linked to external code systems. (c) In the final stage, the MEDS root folder containing the ETL output is provided as input to the `meds2rdf` library, which systematically transforms MEDS events into RDF graphs by instantiating the corresponding MEDS-OWL ontology classes. The resulting graphs preserve source information and enable semantic validation through SHACL shapes. The exemplified use of our tools to transform *neurovasc* to MEDS and then to RDF is publicly available at <https://github.com/albertomarfoglia/meds-to-owl-examples>.

Figure 4 further illustrates how the events shown in Figure 3 are linked to provenance metadata: each event references the originating `meds:DatasetMetadata` instance via `prov:wasDerivedFrom`, while the dataset metadata captures creation time, MEDS schema version, dataset title, and is associated with the ETL activity (`prov:Activity`) that generated it, in this case MEDS-extract. By integrating PROV-O and Dublin Core concepts with MEDS-OWL, this approach provides a semantically rich and traceable representation of clinical data and the associated processing workflow.

To provide a more objective characterization of the transformation process and the resulting KG, we report quantitative statistics at each stage of the *neurovasc* pipeline shown in Figure 3. The original dataset contains 10,000 rows, each describing 30 patient-level features. After preprocessing and conversion with MEDS-extract, the dataset is transformed into 245,015 MEDS events, reflecting the expansion of patient-level features into event-based representations. These events are distributed across the standard ML splits as follows: 196,064 events in the train set, 24,474 in the validation set, and 24,477 in the test set.

The subsequent conversion of MEDS events into RDF using the `meds2rdf` library results in a KG comprising 1,330,351 RDF triples. The graph contains 255,135 distinct subjects, 18 predicates, and 46,722 distinct objects, corresponding to 255,161 distinct IRIs and 36,594 literals, with no blank nodes generated. In terms of instantiated ontology elements, the graph includes 245,015 individuals of class `meds:Event`, corresponding one-to-one with the MEDS events, as well as 10,000 `meds:Subject` instances and 115 `meds:Code` instances reused across events. Dataset-level metadata is represented by a single `meds:DatasetMetadata` individual, while three predefined `meds:SubjectSplit` individuals capture the train, validation, and test partitions. No label instances are generated for this dataset, as outcome labels are not represented in the current transformation.

Finally, we analyzed the structural footprint of events in the resulting RDF graph. Each `meds:Event`

Table 6

Mapping of the elements of MEDS SubjectSplitSchema and LabelSchema tables to `meds:SubjectSplit` and `meds:SubjectLabel` classes, respectively. For each tuple (or row) of LabelSchema, a list of associated RDF triples is generated.

MEDS element	Description	RDF triple	Card.
split	The name of the subset a subject is assigned to, following a dataset split (e.g., train, validation, test).	?set a meds:SubjectSplit. ?s a meds:Subject. ?s meds:assignedSplit ?set.	0..*
label_row	A prediction label (sample) for a subject in a supervised learning task. This is the ground truth label extracted from the dataset according to the task specification.	?l a meds:SubjectLabel. ?l meds:hasSubject ?s.	0..*
prediction_time	The timestamp representing the endpoint of the data that can be used to predict the label for the given sample. This is not necessarily the time of the event itself nor the deployment prediction time. It defines the "as-of" cutoff for input data.	?l meds:predictionTime ?pt^^xsd:dateTime.	1
label_value	The value for the label. It can be one of the following types depending on the task: binary tasks, ordinal or continuous numeric tasks, and categorical tasks.	?l meds:booleanValue ?bv^^xsd:boolean. ?l meds:integerValue ?iv^^xsd:integer. ?l meds:floatValue ?fv^^xsd:double. ?l meds:categoricalValue ?cv^^xsd:string.	0..1

node is described, on average, by 5.31 RDF triples (standard deviation 0.46), with a minimum of 5 and a maximum of 6 triples per event. This low variance reflects the regular structure imposed by MEDS-OWL cardinality constraints, where subject and code associations are mandatory, while temporal and value-related properties are optional. These statistics demonstrate that the transformation produces a compact RDF representation of events while preserving semantic richness and provenance information.

MIMIC-IV example of usage. To further demonstrate the scalability and applicability of `meds2rdf` to real-world EHRs, we applied the library to a subset of the publicly available MIMIC-IV demo dataset [23] already formatted in MEDS [24]. MIMIC-IV is a large, openly accessible critical care dataset containing de-identified EHRs for ICU patients from Beth Israel Deaconess Medical Center. It is widely used in clinical informatics research for benchmarking predictive models, phenotyping, and analyzing temporal data. For this experiment, we selected a subset of 100 patients to illustrate the transformation process while keeping the workflow tractable for visualization and reproducibility.

The conversion generated 916,166 MEDS events distributed across the ML splits as follows: 803,992 in the train set, 69,620 in the tuning set, and 42,554 in the held-out set. The resulting KG comprises 6,408,529 RDF triples, including 925,706 distinct subjects, 19 predicates, 126,680 distinct objects, 925,733 distinct IRIs, and 117,132 literals, with no blank nodes generated. Ontology instantiation includes 916,166 individuals of class `meds:Event`, 100 `meds:Subject` individuals, 9,435 `meds:Code` individuals, three `meds:SubjectSplit` individuals, and one `meds:DatasetMetadata` individual. No label or `ValueModality` instances were included in this subset.

Each `meds:Event` in the MIMIC-IV subset is represented by an average of 6.97 RDF triples (median 7.0, minimum 5, maximum 8, standard deviation 0.42), indicating a slightly richer structure per event compared to *neurovasc*, due to the higher number of numeric and text values associated with MIMIC events. These results confirm that `meds2rdf` can handle real-world MEDS datasets efficiently while preserving the semantic and provenance structure defined by MEDS-OWL.

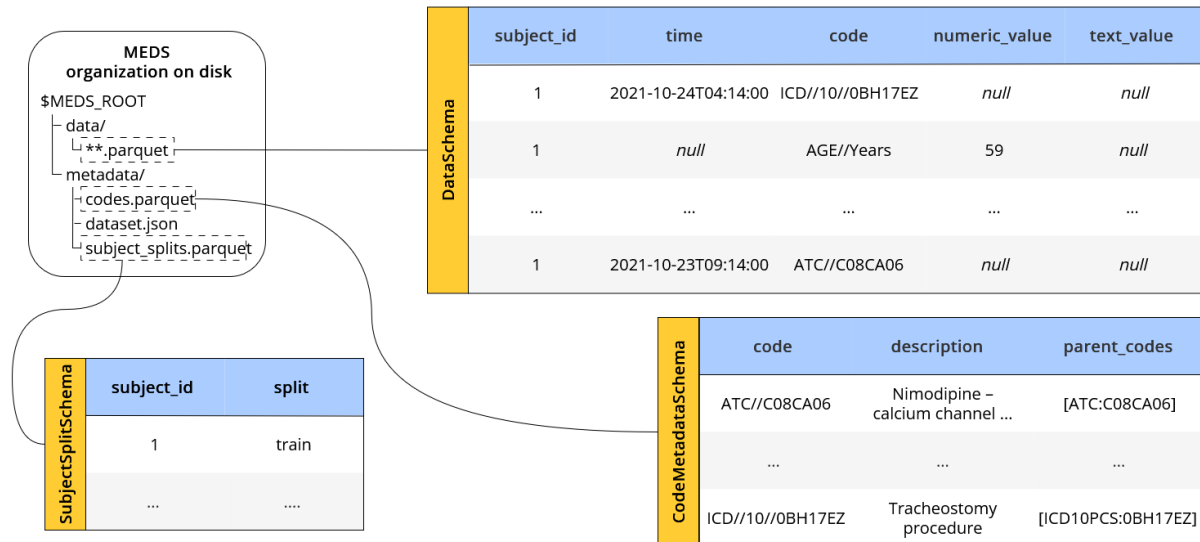
a. Neurovasc (source)

Produced after python script preprocessing

Neurovasc	subject_id	nimodipine	...	age	iot
	1	2021-10-23T09:14:00	...	59	2021-10-24T04:14:00

b. Neurovasc on MEDS (intermediate)

Produced passing neurovasc as input to MEDS-extract



c. Neurovasc on RDF (output)

Produced passing MEDS root folder as input to med2rdf library

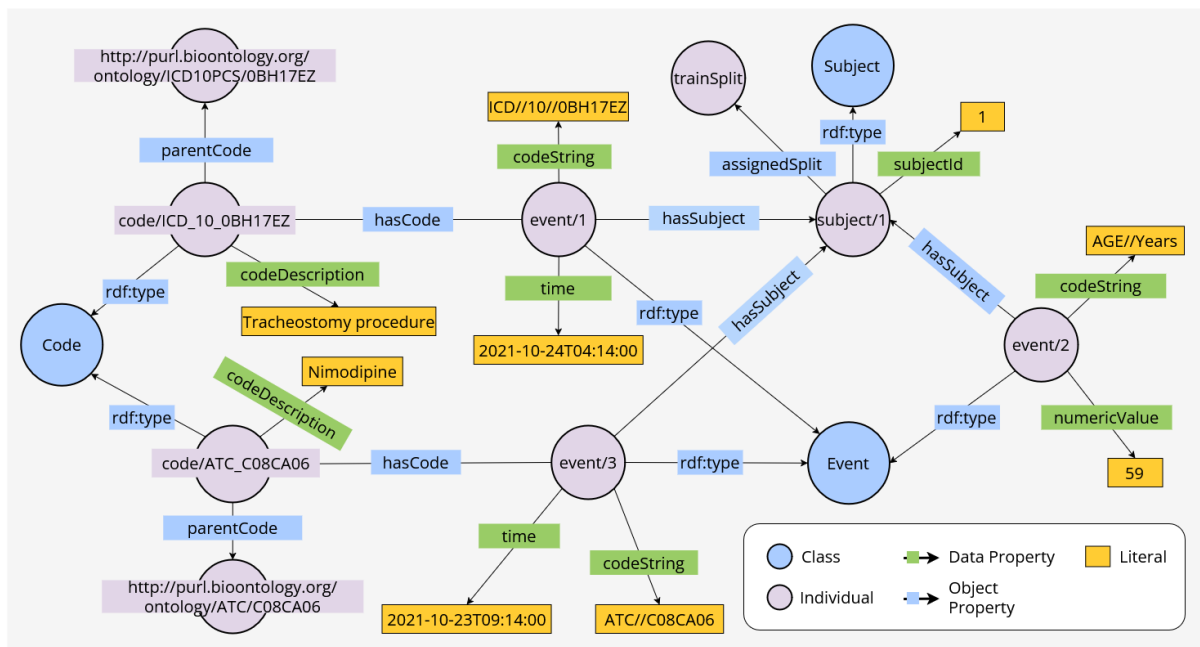


Figure 3: MEDS-OWL conversion workflow for the *neurovasc* dataset. (a) Initial preprocessing of three patient features: age, nimodipine (drug administration), and iot (tracheostomy procedure). (b) Conversion to MEDS tabular format via MEDS-extract, with each feature as a distinct event. (c) Transformation of MEDS output into RDF graphs using `meds2rdf`.

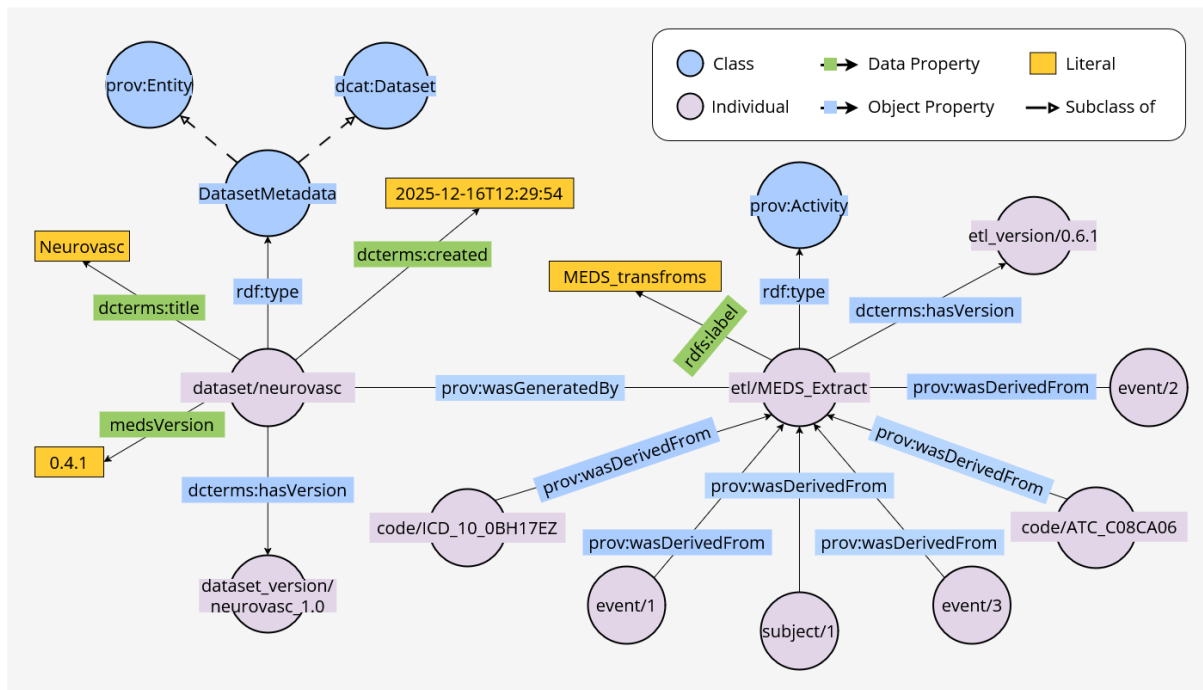


Figure 4: Provenance linkage of MEDS events, codes, and subjects using MEDS-OWL, PROV-O, and Dublin Core. Each event references its source `meds:DatasetMetadata` (here, *neurovasc*) via `prov:wasDerivedFrom` and is associated with the generating ETL activity (`prov:Activity`), here MEDS-extract. This approach provides a semantically rich, traceable representation of clinical data workflows.

4. Discussion

Our evaluation on two complementary datasets, the synthetic *neurovasc* cohort and a real-world subset of MIMIC-IV, demonstrates that MEDS-OWL can faithfully represent MEDS event data and associated provenance in RDF while remaining compact and operationally pragmatic. The resulting RDF graphs are SHACL-validated, supporting semantic and structural consistency. Moreover, exporting MEDS datasets as KGs enables the direct application of graph representation learning techniques, thereby opening new opportunities for downstream applications such as graph-based ML. At the same time, the KG representation allows us to leverage the strengths of both paradigms: the event-centric, ML-oriented design of MEDS and the semantic expressivity of KGs. By explicitly modeling entities, relations, and provenance, MEDS-OWL provides semantic enrichment of the underlying datasets, making complex, multi-relational dependencies between clinical events accessible to downstream prediction models. Importantly, because MEDS-OWL preserves the temporal structure and ordering of events, it naturally supports the construction of Temporal Knowledge Graphs (TKGs), enabling longitudinal modeling of patient trajectories and opening new directions for temporally aware clinical prediction tasks.

Compared to existing approaches for representing clinical data in RDF, such as FHIR-RDF, SPHN, and CARE-SM, MEDS-OWL addresses several gaps. Unlike FHIR-RDF, which suffers from loosely specified semantics, or SPHN and CARE-SM, which are tightly coupled to specific initiatives or aims, MEDS-OWL follows the minimality of MEDS, its ML-oriented design, and a well-defined OWL semantics reusing ontologies. This facilitates interoperability across heterogeneous datasets through explicit mapping points (e.g., `meds:parentCode` or `skos:exactMatch`).

The dual evaluation highlights both strengths and current limitations. While the *neurovasc* dataset illustrates a test on a controlled synthetic cohort, the MIMIC-IV conversion demonstrates that the approach scales to large, heterogeneous, real-world data. The latter confirms that `meds2rdf` can process high event volumes while preserving semantic consistency. It also emphasizes the need for careful performance profiling and analysis to avoid SHACL inconsistencies when handling millions of triples and thousands of distinct clinical codes. Currently, we have not conducted an in-depth analysis of

reasoning performance or triple-store ingestion costs, which remain important directions for future work in large-scale research applications. Moreover, the current MEDS-OWL ontology does not capture certain operational constraints of MEDS datasets. Specifically, MEDS requires subject contiguity, which means that all the features associated with a single subject should appear in a single sort and should be temporally sorted. While temporal ordering becomes largely irrelevant when converting tabular MEDS data to RDF graphs, subject contiguity can be partially represented by associating each subject with a specific ML partition via the `meds:SubjectSplit` class. Importantly, a single subject may participate in multiple independent splits (e.g., for repeated experiments or cross-validation); however, within each split, the subject can belong to only one subset (e.g., training, validation, or test). This ensures that partitioning is mutually exclusive at the level of each experiment while still supporting multiple, independent splits for variability analysis. To address this limitation, future releases of MEDS-OWL will incorporate enhanced axioms, potentially using disjointness rules or qualified cardinality restrictions in description logics, thereby allowing the ontology to more accurately represent MEDS operational constraints.

To substantiate the claimed benefits and broaden the applicability of the ontology, we established an evaluation agenda and identified four steps. (1) Quantifying round-trip fidelity to measure information preservation and loss due to transformations. (2) Aggregating SHACL validation statistics across multiple, heterogeneous datasets to characterize common data-quality issues and to demonstrate how schema validation reduces downstream errors. (3) Profiling performance on large MEDS corpora will facilitate the identification of bottlenecks and enable a rigorous assessment of semantic fidelity and graph generation scalability. (4) Compare predictive analysis on MEDS-OWL data with graph-based ML models, such as GCNs.

Planned extensions of the MEDS-OWL ontology and associated tooling include greater reuse of external ontologies to replace remaining MEDS-specific constructs, thereby reducing bespoke terms and enhancing interoperability. We aim to integrate temporal ontologies (e.g., OWL-Time) and pattern-based representations to more accurately model intervals and episode boundaries, while also adopting unit ontologies to standardize numeric values and support unit conversion and aggregation. Hierarchical relations among codes, currently captured via `meds:parentCode` (see Table 4), will be formalized using `skos:broader` to facilitate cross-vocabulary linking, with optional use of `rdfs:subClassOf` for implementations that model codes as OWL classes. These enhancements aim to increase the semantic richness and usability of MEDS-OWL for KG, enabling systematic cross-vocabulary linking.

5. Conclusion

We have presented MEDS-OWL, an ontology that formalizes the MEDS event model using Semantic Web standards, along with the `meds2rdf` library for converting MEDS datasets into RDF graphs compliant with MEDS-OWL. Our framework preserves MEDS’s ML-oriented representation of medical event sequences while enhancing it with a rich semantic structure and explicit links to biomedical and provenance ontologies. This combination enables symbolic reasoning, interoperability across heterogeneous data sources, and seamless integration with knowledge-based systems.

Preliminary experiments on the synthetic *neurovasc* dataset and a real-world subset of MIMIC-IV demonstrate that MEDS-OWL and `meds2rdf` can accurately transform MEDS events and their associated metadata into semantically validated RDF graphs. Released under an open-source license, these tools promote FAIR-aligned sharing and reuse of medical event data.

By providing a reliable transformation from MEDS to RDF, our framework lays the groundwork for downstream prediction tasks using graph-based ML models. Representing MEDS datasets as KGs enables semantic enrichment of clinical events and their use across multiple analytical tasks. Moreover, because MEDS-OWL preserves temporal ordering, it naturally supports the construction of TKGs, opening new avenues for longitudinal patient modeling and temporally aware prediction tasks that go beyond conventional tabular approaches.

In future work, we plan to systematically compare graph-based ML models with classical ML ap-

proaches to assess whether RDF-based, relational representations provide measurable improvements in predictive performance over traditional tabular methods. Additionally, we aim to evaluate the round-trip fidelity of the framework on other public biomedical datasets formatted in MEDS, thereby assessing information preservation across transformations. Finally, we plan to further align MEDS-OWL with widely adopted biomedical ontologies to strengthen semantic interoperability and facilitate broader reuse across biomedical KGs ecosystems.

Reproducibility Statement

All experiments and dataset conversions presented in the accompanying paper are fully reproducible using the workflows provided in this repository: <https://github.com/albertomarfoglia/meds-to-owl-examples>. The end-to-end pipelines for the *neurovasc* synthetic dataset and the MIMIC-IV Demo MEDS dataset are implemented as executable, notebook-based workflows (main.ipynb). They rely only on open-source tools and publicly available data or scripts. The repository specifies all software dependencies and configuration files. Moreover, randomized steps, such as synthetic data generation or sampling, are executed with fixed seeds to ensure repeatability.

Acknowledgments

This work is supported by CombO (Health Data Hub) project and the Agence Nationale de la Recherche under the France 2030 program, reference ANR-22-PESN-0007 ShareFAIR, and ANR-22-PESN-0008 NEUROVASC.

Declaration on Generative AI

During the preparation of this work, the author(s) used ChatGPT, Grammarly in order to: Grammar and spelling check, Paraphrase and reword. After using this tool/service, the author(s) reviewed and edited the content as needed and take(s) full responsibility for the publication's content.

References

- [1] Bert Arnrich, Edward Choi, Ethan Steinberg, Michael Wornow, Robin van de Wate, Medical Event Data Standard (MEDS): Facilitating Machine Learning for Health, in: ICLR 2024 Workshop on Learning from Time Series For Health, 2024, pp. 03–08. URL: <https://openreview.net/forum?id=IsHy2ebjIG>.
- [2] M. a. c. McDermott, MEDS Extract: ETL Pipelines for Converting EHR Data to MEDS Format, 2024. URL: https://github.com/mmcdermott/MEDS_extract.
- [3] P. Renc, M. K. Grzeszczyk, N. Oufattole, D. Goode, Y. Jia, S. Bieganski, M. B. A. McDermott, J. Was, A. E. Samir, J. W. Cunningham, D. W. Bates, A. Sitek, Foundation model of electronic medical records for adaptive risk estimation, *GigaScience* 14 (2025) giaf107. URL: <https://doi.org/10.1093/gigascience/giaf107>. doi:10.1093/gigascience/giaf107.
- [4] N. Oufattole, T. Bergamaschi, A. Kolo, H. Jeong, H. Gaggin, C. M. Stultz, M. B. A. McDermott, MEDS-Tab: Automated tabularization and baseline methods for MEDS datasets, 2024. URL: <http://arxiv.org/abs/2411.00200>. doi:10.48550/arXiv.2411.00200, arXiv:2411.00200 [cs].
- [5] P. Hitzler, M. Krotzsch, S. Rudolph, Foundations of Semantic Web Technologies, Chapman and Hall/CRC, New York, 2009. doi:10.1201/9781420090512.
- [6] A. Carbonaro, A. Marfoglia, F. Nardini, S. Mellone, CONNECTED: leveraging digital twins and personal knowledge graphs in healthcare digitalization, *Frontiers in Digital Health* 5 (2023). URL: <https://www.frontiersin.org/journals/digital-health/articles/10.3389/fdgth.2023.1322428/full>. doi:10.3389/fdgth.2023.1322428.

- [7] A. Marfoggia, C. D’Errico, S. Mellone, A. Carbonaro, A knowledge graph-driven framework for deploying AI-powered patient digital twins, *Future Generation Computer Systems* 180 (2026) 108380. URL: <https://www.sciencedirect.com/science/article/pii/S0167739X26000142>. doi:10.1016/j.future.2026.108380.
- [8] A. Hogan, E. Blomqvist, M. Cochez, C. D’amato, G. D. Melo, C. Gutierrez, S. Kirrane, J. E. L. Gayo, R. Navigli, S. Neumaier, A.-C. N. Ngomo, A. Polleres, S. M. Rashid, A. Rula, L. Schmelzeisen, J. Sequeda, S. Staab, A. Zimmermann, *Knowledge Graphs*, *ACM Comput. Surv.* 54 (2021) 71:1–71:37. URL: <https://doi.org/10.1145/3447772>. doi:10.1145/3447772.
- [9] W. L. Hamilton, R. Ying, J. Leskovec, *Representation Learning on Graphs: Methods and Applications*, 2018. URL: <http://arxiv.org/abs/1709.05584>. doi:10.48550/arXiv.1709.05584, arXiv:1709.05584 [cs].
- [10] Q. Wang, Z. Mao, B. Wang, L. Guo, *Knowledge Graph Embedding: A Survey of Approaches and Applications*, *IEEE Transactions on Knowledge and Data Engineering* 29 (2017) 2724–2743. URL: <https://ieeexplore.ieee.org/document/8047276>. doi:10.1109/TKDE.2017.2754499.
- [11] J. H. Jhee, A. Megina, P. Constant Dit Beaufils, M. Karakachoff, R. Redon, A. Gaignard, A. Coulet, *Predicting Clinical Outcomes from Patient Care Pathways Represented with Temporal Knowledge Graphs*, in: *The Semantic Web: 22nd European Semantic Web Conference, ESWC 2025, Portoroz, Slovenia, June 1–5, 2025, Proceedings, Part I*, Springer-Verlag, Berlin, Heidelberg, 2025, pp. 282–300. URL: https://doi.org/10.1007/978-3-031-94575-5_16. doi:10.1007/978-3-031-94575-5_16.
- [12] M. Lehne, S. Luijten, P. Vom Felde Genannt Imbusch, S. Thun, *The Use of FHIR in Digital Health – A Review of the Scientific Literature*, in: *German Medical Data Sciences: Shaping Change – Creative Solutions for Innovative Medicine*, IOS Press, 2019, pp. 52–58. URL: <https://ebooks.iospress.nl/doi/10.3233/SHTI190805>. doi:10.3233/SHTI190805.
- [13] G. Xiao, E. Pfaff, E. Prud’hommeaux, D. Booth, D. K. Sharma, N. Huo, Y. Yu, N. Zong, K. J. Ruddy, C. G. Chute, G. Jiang, *FHIR-Ontop-OMOP: Building clinical knowledge graphs in FHIR RDF with the OMOP Common data Model*, *Journal of Biomedical Informatics* 134 (2022) 104201. URL: <https://www.sciencedirect.com/science/article/pii/S1532046422002064>. doi:10.1016/j.jbi.2022.104201.
- [14] R. Kaliyaperumal, G. Singh, N. Queralt-Rosinach, P.-B. Hoen, M. Roos, *Phenopackets for the Semantic Web*, in: *SWAT4HCLS, 2022*, pp. 29–33.
- [15] A. Marfoggia, F. Nardini, V. A. Arcobelli, S. Moscato, S. Mellone, A. Carbonaro, *Towards real-world clinical data standardization: A modular FHIR-driven transformation pipeline to enhance semantic interoperability in healthcare*, *Computers in Biology and Medicine* 187 (2025) 109745. URL: <https://www.sciencedirect.com/science/article/pii/S0010482525000952>. doi:10.1016/j.combiomed.2025.109745.
- [16] V. Touré, P. Krauss, K. Gnodtke, J. Buchhorn, D. Unni, P. Horki, J. L. Raisaro, K. Kalt, D. Teixeira, K. Cramer, S. Österle, *FAIRification of health-related data using semantic web technologies in the Swiss Personalized Health Network*, *Scientific Data* 10 (2023) 127. URL: <https://www.nature.com/articles/s41597-023-02028-y>. doi:10.1038/s41597-023-02028-y.
- [17] R. Kaliyaperumal, M. D. Wilkinson, P. A. Moreno, N. Benis, R. Cornet, B. dos Santos Vieira, M. Dumontier, C. H. Bernabé, A. Jacobsen, C. M. A. Le Cornec, M. P. Godoy, N. Queralt-Rosinach, L. J. Schultze Kool, M. A. Swertz, P. van Damme, K. J. van der Velde, N. Lalout, S. Zhang, M. Roos, *Semantic modelling of common data elements for rare disease registries, and a prototype workflow for their deployment over registry data*, *Journal of Biomedical Semantics* 13 (2022) 9. URL: <https://doi.org/10.1186/s13326-022-00264-6>. doi:10.1186/s13326-022-00264-6.
- [18] A. Marfoggia, TeamHeKA/meds-ontology, 2025. URL: <https://doi.org/10.5281/zenodo.17953576>. doi:10.5281/zenodo.17953576.
- [19] A. Marfoggia, TeamHeKA/meds2rdf, 2025. URL: <https://doi.org/10.5281/zenodo.17953580>. doi:10.5281/zenodo.17953580.
- [20] N. F. Noy, D. L. McGuinness, *Ontology Development 101: A Guide to Creating Your First Ontology*, 2001.
- [21] M. A. Musen, *The protégé project: a look back and a look forward*, *AI Matters* 1 (2015) 4–12. URL:

<https://doi.org/10.1145/2757001.2757003>. doi:10.1145/2757001.2757003.

- [22] S. Lohmann, V. Link, E. Marbach, S. Negru, Webvowl: Web-based visualization of ontologies, in: International Conference on Knowledge Engineering and Knowledge Management, Springer, 2014, pp. 154–158.
- [23] A. E. W. Johnson, L. Bulgarelli, L. Shen, A. Gayles, A. Shammout, S. Horng, T. J. Pollard, S. Hao, B. Moody, B. Gow, L.-w. H. Lehman, L. A. Celi, R. G. Mark, MIMIC-IV, a freely accessible electronic health record dataset, *Scientific Data* 10 (2023) 1. URL: <https://www.nature.com/articles/s41597-022-01899-x>. doi:10.1038/s41597-022-01899-x.
- [24] M. McDermott, J. Xu, MIMIC-IV MEDS: An ETL pipeline to extract MIMIC-IV data into the MEDS format, 2025. URL: <https://zenodo.org/records/17535580>. doi:10.5281/zenodo.17535580.