InterpretME: A Tool for Interpretations of Machine Learning Models Over Knowledge Graphs

Yashrajsinh Chudasama [a,b,*], Disha Purohit [a,b], Philipp D. Rohde [a,b,c], Julian Gercke [b] and Maria-Esther Vidal [a,b,c]

[a] *TIB Leibniz Information Centre for Science and Technology, Hannover, Germany*
E-mails: yashrajsinh.chudasama@tib.eu, disha.purohit@tib.eu, philipp.rohde@tib.eu, maria.vidal@tib.eu
[b] *Leibniz University, Hannover, Germany*
E-mail: julian.gercke@stud.uni-hannover.de
[c] *L3S Research Center Germany, Hannover, Germany*

Abstract. In recent years, knowledge graphs have been considered pyramids of interconnected data enriched with semantics for complex decision-making. The potential of knowledge graphs and the demand for interpretability of machine learning (ML) models in diverse domains (e.g., healthcare) have gained more attention. The lack of model transparency negatively impacts the understanding and, in consequence, interpretability of the predictions made by a model. Data-driven models should be empowered with the knowledge required to trace down their decisions and the transformations made to the input data to increase model transparency. In this paper, we propose InterpretME, a tool for fine-grained representations in a knowledge graph, of the main characteristics of trained machine learning models. They include data- (e.g., features' definition and SHACL validation) and model-based characteristics (e.g., relevant features and interpretations of prediction probabilities and model decisions). InterpretME allows for defining a model's features over knowledge graphs (KGS) and relational data in various formats, including CSV and JSON; SHACL states domain integrity constraints. InterpretME traces the steps of data collection, curation, integration, and prediction; it documents the collected metadata in the InterpretME KG. InterpretME is publicly available as a tool; it includes a pipeline for enhancing the interpretability of ML models, the InterpretME KG, and an ontology to describe the main characteristics of trained ML models.

Keywords: Interpretability, Knowledge Graphs, Machine Learning Models, SHACL, Ontologies

1. Introduction

Interpretability is the degree to which humans can understand the decisions made by computational frameworks. Specifically, in Artificial Intelligence (AI), the higher the interpretability of predictive models, the easier for humans to understand why a model makes certain decisions. The recent advancements and complexity of machine learning methods have demonstrated their success in forecasting complex problems (e.g., disease diagnosis or progression [1, 2]). Unfortunately, they are often opaque and do not offer interpretations for their predictions, and it is not always possible to comprehend the results. Interpretable predictive models have rapidly become a relevant problem [3]. Nevertheless, although various tools aim to interpret the algorithmic decisions of ML models [4, 5], they are incapable of capturing the knowledge required to translate a model's insights into the application domain. On the contrary, KGs encode data and knowledge; they, together with domain ontologies (e.g., ML Schema [6]),

[*] All the authors have equally contributed to this work.

represent building blocks for increasing the understanding of the behavior and effects of a predictive model.

Our Tool: we propose an analytical tool, named InterpretME, for tracing and explaining the predictive models built over data collected from KGs. InterpretME implements a set of validating integrity constraints that provide a meaningful description of a target entity of a prediction model, and its main properties. InterpretME is designed to work with both knowledge graphs (accessible via Web interfaces, e.g., SPARQL endpoints), and datasets in various formats (e.g., CSV or JSON). InterpretME represents as factual statements all what is collected during the process of training a ML model and its interpretation; the InterpretME KG integrates these statements. To retrieve data for the predictive pipeline, the user must provide either a SPARQL endpoint or the path to a dataset. If the data is retrieved from KGs, InterpretME can validate constraints expressed in SHACL, align the InterpretME KG to the input KGs, and execute federated queries. On the other hand, in input data is collected from files, a user can still explore interpretations with the InterpretME KG. The latest version of InterpretME is customized for supervised machine learning models (e.g., decision trees and random forests) and interpretable tools (e.g., LIME [5]). InterpretME is publicly available as a resource in GitHub[1], Zenodo[2], and in PyPI[3]. We empirically evaluate InterpretME in real-world KGs regarding execution time, interpretation quality, and traceability of target entities. The observed results reveal the key role of Semantic Web technologies in the interpretation of the outcomes of predictive models.

The rest of the paper is structured as follows: Section 2 describes the main concepts and motivates our work. Section 3 defines the InterpretME architecture, while Section 4 reports the results of our experimental studies. Section 5 presents the main characteristics of InterpretME as a resource. Section 6 discusses the state of the art, and our conclusions and future work are outlined in Section 7.

2. Preliminaries and Motivation

2.1. Main Concepts

2.1.1. Predictive Modeling Frameworks.

Predictive Modeling comprises methods to forecast future outcomes based on data encoding what has happened in the past (e.g., historical data) and what is currently happening (e.g., current data). Predictive modeling resorts to a variety of models and algorithms (e.g., random forest, gradient boosted model or decision trees) to solve predictive problems (e.g., classification or outlier detection) [7]. Despite the large spectrum of mature models and well-defined problems, predictive modeling is a complex task that requires user expertise in the domain context and modeling. De Bie et al. [3] conceptualizes predictive modeling in four stages, as depicted in Figure 4.

a) *Data Engineering* comprises the tasks of acquiring, organizing, curating, and preparing data for prediction modeling. **b)** *Data Exploration* includes the understanding of an application domain, and the interpretation of missing values, integrity constraints, and data structures. **c)** *Model Building* involves the selection of the predictive model (e.g., machine or deep learning), target function, hyperparameters, and relevant features. **d)** *Exploitation* covers the transformation of the predictions into the decisions on the application domain.

Predictive modeling complexity and the mandatory requisite of knowledgeable users motivate automated techniques for facilitating the implementation of the stages in Figure 4. De Bie et al. [3] identify three forms of automation: **i)** *Mechanization* is used in well-known tasks that can be performed algorithmically without human intervention. Functions and modules, for example, can be used to perform tasks such as clustering algorithms and standardizing table values. Packages for performing such operations are available as a library that can be easily integrated into any project. **ii)** *Composition* is performed on a series of tasks that are implemented as workflows in high-level programming languages. For example, algorithm selection, hyperparameter optimization, model selection, and so on fall under the stage of *Model Building*, which is covered in Composition. **iii)** *Assistance* implemented to enhance interpretability and user efficiency by presenting visualizations, prediction scores, and explanatory model decisions.

[1] https://github.com/SDM-TIB/InterpretME
[2] https://doi.org/10.5281/zenodo.7684228
[3] https://pypi.org/project/InterpretME/

This includes constantly monitoring what humans are doing during the data science process so that an automated assistant can identify inappropriate choices and make appropriate recommendations.

Current developments in automation have mainly impacted the stages of *Model Building* and *Exploitation*. In particular, Automated Machine Learning (AutoML) systems (e.g., AutoML[4] and AutoWeka [8]) successfully implement mechanization, and can automatize the processes of model and feature selection, and the optimization of the hyperparameters and target functions. Furthermore, interactive tools like RapidMiner[5], the Local Interpretable Model-agnostic Explanations (LIME) framework [5], and SHapely Additive exPlanations (SHAP) [4] provide assistance by reporting explanations of a model's results. Specifically, LIME relies on local surrogate models to explain the prediction of each entity of a test dataset. It computes, for the tested entities, the feature contribution for each target class and the prediction probability. In general, LIME is independent of the original algorithm given to the model and works locally to provide explanations for the prediction relative to each instance. LIME tries to fit the local model using sample data points that are similar to the instance being explained.

2.1.2. *Knowledge Graph Frameworks.*

The Semantic Web [9] aims at humans and machines working cooperatively in data exchange. Technologies capable of encoding semantics have been defined to achieve this goal. One of these technologies is the Resource Description Framework (RDF) [10]; the W3C standard for publishing and exchanging data over the web. An RDF graph $\mathcal{G} = \langle V_{\mathcal{G}}, E_{\mathcal{G}} \rangle$ is a directed graph with labeled edges [11]. The nodes represent subjects and objects of RDF triples, while predicates correspond to the labels of edges between nodes. We will use the terms RDF graph and knowledge graph [12] interchangeably. In the context of machine learning, a knowledge graph [13] can be used to store and reason about the outcome of a machine learning model.

SPARQL [14] is the W3C recommendation language to query RDF graphs. SPARQL 1.1 [15] provides the SERVICE clause and allows for specifying a federated query over various RDF graphs accessible via SPARQL endpoints. However, usually, SPARQL endpoints prohibit using these clauses. Federated query engines offer the possibility to answer federated queries without the need to know from where parts of a query will be answered [16].

The Shapes Constraint Language (SHACL) [17] is the W3C recommendation language to define integrity constraints over RDF data. SHACL integrity constraints are expressed in RDF and modeled as a network of shapes, called *shape schema*. A shape consists of **i)** a definition of the target; usually an RDF class or set of nodes, and **ii)** a set of constraints that are imposed over the instances of the target. Constraints can also link shapes to each other, hence, a shape network. An RDF graph's entity that matches the target definition of a shape, satisfies the shape if it validates all the constraints of the shape; the problem is, in general, intractable [18]. However, algorithms have been proposed and implemented to validate tractable SHACL fragments [19, 20] efficiently.

Definition 2.1 (InterpretME KG)**.** An InterpretME KG is a heterogeneous graph where each node and edge is assigned to a type, i.e., InterpretME KG = (V, E, L, T), where: **i)** V represents the set of nodes, where each node $v \in V_i$ is associated with a type $T_i(v)$ (i.e., Classes in Figure 2); **ii)** E is the set of edges, where each edge $l \in E$ connects two nodes of different types, i.e., $T(v) \neq T(u)$; **iii)** L represents a set of directed edge labels (i.e., Properties in Figure 2); and **iv)** T represents the set of entity types (i.e., Classes in Figure 2), where each entity type T_i is associated with set of entities V_i.

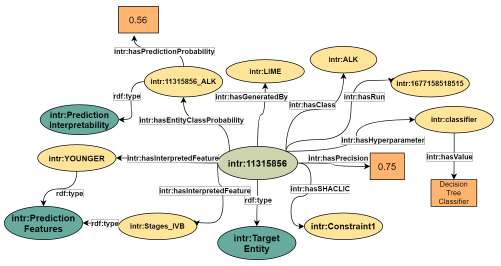

Fig. 1.: **InterpretME KG** shows an instance of a target entity highlighted in light green. This entity contains several key features highlighted in yellow, including hyperparameters, LIME and SHACL.

The target entity is a specific node in V_i of the input KG, which is of particular interest, and the goal might be to predict the properties of that node based on the local context in the knowledge graph. For instance, Figure 3a demonstrates the classification task of a lung cancer patient being positive for biomarker ALK. In this case, the target

[4] https://www.automl.org/
[5] https://rapidminer.com/

ML Schema	
Class	Run Implementation ModelEvaluation HyperParameter
Properties	executes realizes hasInput implements hasOutput specifiedBy hasValue hasHyperParameter

(a) ML schema vocabulary

InterpretME		
SPARQL endpoint	Class	Endpoint
	Properties	hasEndpoint
Cross Validation	Class	CrossValidation
	Properties	hasCrossValidation
Prediction Classes	Class	PredictionClasses
	Properties	hasClasses
Feature Definition	Class	FeatureDefinition
	Properties	hasFeature hasDefinition hasRun
Entity Alignment	Class	TargetEntity
	Properties	hasEntity owl:sameAs hasRun
Important Features	Class	ImportantFeature
	Properties	hasImportantFeature
Model Evaluation	Class	PrecisionRecall
	Properties	hasClasses hasPrecision hasRecall hasF1score hasSupport hasRun
Sampling Strategy	Class	SamplingStrategy
	Properties	hasSampling

(b) InterpretME vocabulary

InterpretME		
SHACL Constraints	Class	SHACLValidation
	Properties	hasSHACLResult hasSHACLSchema hasSHACLShape hasSHACLConstraints hasConstraintID hasSHACLIC hasRun
LIME Interpreted Features	Class	PredictionFeatures FeaturesWeights TestedTargetEntity
	Properties	hasEntity hasFeatureWeight hasFeature hasWeight hasGeneratedBy hasInterpretedFeature hasRun
LIME Prediction Probablities	Class	PredictionInterpretability TestedTargetEntity
	Properties	hasEntity hasClass hasPredictionProbability hasGeneratedBy hasRun hasEntityClassProbability

(c) InterpretME vocabulary

Fig. 2. **Vocabulary** Figure 2a shows the properties of the classes used to describe the machine learning models. Figure 2b and Figure 2c showing the extension of ML schema used to describe the predictive model characteristics and also incorporates the SHACL validation results. The required properties and classes are utilize for presenting the traced knowledge of the predictive models in the InterpretME KG.

entity will be a specific patient with known characteristics. The combination of the relationships and attributes of a target entity in the input KG with the predicted relationships and attributes generated by the trained predictive model provides more comprehensive interpretation of the knowledge represented in the InterpretME KG.

2.2. *Motivating Example*

The motivation for our work originates from the lack of automated assistance despite the great potential of integrating knowledge graphs with predictive modeling frameworks. Tracing and explaining the predictive models built over data collected from the KGs in order to provide assistance is the main goal of our tool InterpretME. Even though the state-of-the-art successfully developed automated machine learning systems, through mechanization and composition, pipelines for predictive modeling are unable to generate human- and machine-readable decisions to assist users and enhance their efficiency. Figure 3a depicts a predictive modeling pipeline, where an automated machine learning system (e.g., AutoML) is utilized for model and feature selection, and hyperparameters' optimization. Moreover, interpretable tools (e.g., LIME) provides interpretable results. Figure 3a illustrates the pipeline **B**; an input dataset **A** is collected from an **RDF knowledge graph** that integrates data about lung cancer patients. The *SPARQL query* is used to extract features from the KG, describing the main characteristics of a lung cancer patient, i.e., patient identifier (a.k.a. EHR is electronic health record of a patient), Gender, Age, cancer stage (a.k.a. CancerStage), smoking habits (a.k.a. SmokingHabit), and lung cancer Biomarkers. Additionally, the dataset comprises information about the cancer types of the relatives of the lung cancer patients; this information is maintained in the features CancerType and FamilialDegree. The predictive task is a binary classification to predict if a patient will be positive for the biomarker ALK or by any other biomarker. AutoML performs model selection and hyperparameter optimization **C**; based on AutoML recommendations, the random forests and decision tree models are selected to implement the classification problem. Further, the interpretable surrogate tools LIME and Decision trees are utilized to provide local interpretations of each patient in the test dataset. Decision trees yield the relevant features which contribute to the model outcomes **D**. **E** depicts an exemplar entity where LIME determines a prediction probability of 0.7 to belong to the target class ALK (i.e., Class 0), otherwise, 0.3 target class Others (i.e., Class 1). LIME also identifies the top 10 relevant features for the target entity and assigns weights. These outcomes allow for understanding the quality of the implemented pipeline. Nevertheless, when they are reported to oncologists, many questions may still arise Figure 3b: **Q1)** Who is this patient interpreted by LIME? **Q2)** How does the Stages feature

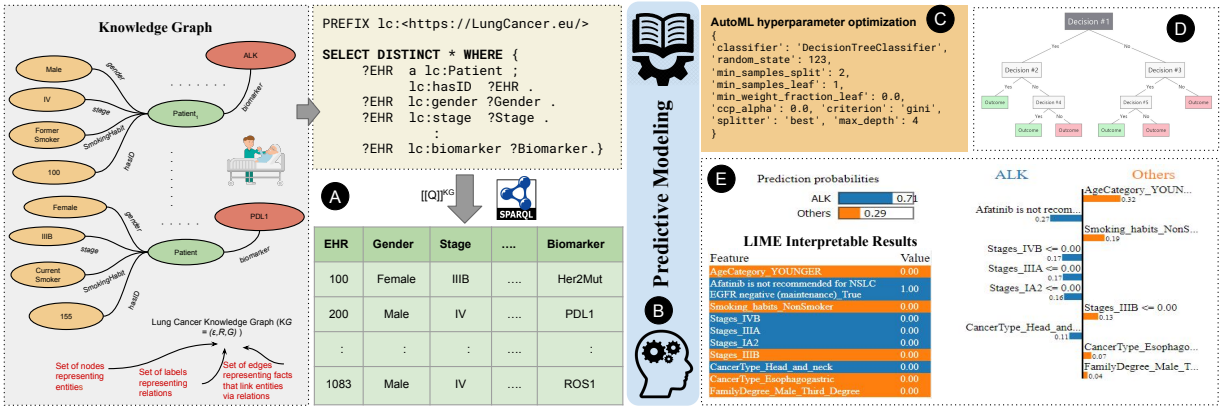

(a) Motivating Example

	Example Questions
Q1.	Who is this patient interpreted by LIME?
Q2.	How does the Stages feature contribute to classification of this patient in the class ALK?
Q3.	Which other features are relevant for this classification?
Q4.	Does this patient satisfy the domain integrity constraints?
Q5.	What are the main characteristics of this patient?

(b) Example Questions

Fig. 3. **Lung Cancer KG Example:** Figure 3a depicts the motivation of a pipeline that integrates knowledge graphs and data-driven frameworks, i.e., AutoML, Decision Trees, LIME. SPARQL queries are utilized to extract the data from KG, the classification task is to predict whether a patient is positive for biomarker and Figure 3b illustrates the questions reported by domain experts regarding the issue in interpreting such data-driven frameworks, shows that such frameworks still lack in terms of interpretability and traceability.

contribute to the classification of this patient in the target class ALK? **Q3)** Which other features are relevant for this classification? **Q4)** Does this patient satisfy the domain integrity constraints? **Q5)** What are the main characteristics of this patient? The oncologists gave us the feedback that InterpretME facilitates their understanding of the reasons for the classification. Notably, the users comprehended better the answers to the queries over the InterpretME KG and the Lung Cancer KG than simply the answers of LIME. Although, the user would have been able to interpret the results produced by LIME, he/she would need to trace back these results to the original data properties to know if, for example, the reported patient violates the domain constraints or not. For instance, 'If a patient is given a drug for the treatment, then he/she is cured'. InterpretME combines local explanations with knowledge graphs and SHACL constraints to ensure traceability. In order to support our justification a user case study which will be discussed in the later sections was also carried out. The assessment outcomes suggest that following the Linked Data principles to document, in a KG, facts- about a target entity and its links to the KGs from where it is collected- contributes to a better understanding of the decisions made by an ML model.

3. InterpretME: An Interpretable Pipeline for Predictive Modeling over Knowledge Graphs

3.1. Tracing Metadata in Machine Learning Pipelines

We aim to collect metadata during the different stages of a predictive model pipeline. Figure 4 shows, in a pictorial view, the characteristics of the data collected towards the improvement of automation *assistance*. In the data engineering stage, InterpretME captures metadata from the input KG (e.g., features' definition, endpoint, and target classes) and records the SHACL constraints, used for data validation. During the model building, suggestions from AutoML systems can be considered. InterpretME traces the optimized hyperparameters and estimated features' relevancy, and records the model performance metric outcomes (e.g., precision) for a particular run.

SHACL validation reports are stored too. InterpretME also exploits decision trees and visualizes the validation report to enhance exploration. Lastly, for exploitation, InterpretME aligns target entities with entities in KGs, where the model's features are defined. These facts allow for **i)** tracing back the main properties of target entities in the input KG; **ii)** understanding the interpretability models' results (e.g., LIME prediction probabilities, interpreted features, and associated weights); and **iii)** checking SHACL validation reports. InterpretME identifies uniquely all the collected metadata. As a result, tracing the collected metadata enhances explanations about an ML model behavior.

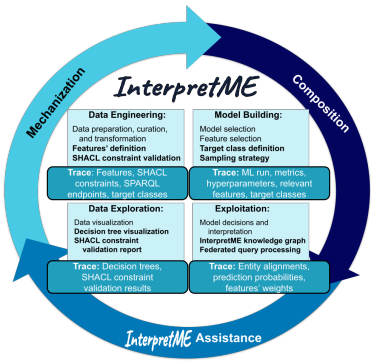

Fig. 4.: **Tracing Metadata** depicts traced metadata (e.g., model characteristics) of the trained predictive models at each layer of the pipeline, i.e., *Mechanization*, *Composition*, *Assistance*.

3.2. The InterpretME Architecture

Figure 5 depicts the InterpretME architecture. InterpretME is a tool for fine-grained representations of the main characteristics of a trained predictive modeling framework. The architecture of InterpretME deals with training the predictive models and collecting information generated as output of predictive models (i.e., model accuracy, list of important features, prediction probabilities, and classified classes for each instance), and provides assistance (Figure 4) to the user by tracing metadata to generate instances of the InterpretME KG and executes federated queries on top of the InterpretME KG and the input KG to trace back, and answers questions in Figure 3b. InterpretME takes a JSON file as input (i.e., endpoints of KGs, features' definition, target definition, SHACL constraints, sampling strategy, and class definition)[6]; a SPARQL query is generated based on the feature definition given by the user and the query is used to retrieve the application domain data from the input KGs. Users perform the step of providing a SPARQL endpoint or path to a dataset and its corresponding feature definition in the input JSON file. InterpretME performs additional steps automatically, such as extracting data based on the query, running predictive models, saving prediction results, and so on. The structure of the query enables entities in the input KGs can be aligned to the identifiers of the instances in the ML models' datasets. These alignments facilitate the SHACL validation and federated query processing over the InterpretME KG and the input KGs.

1. InterpretME evaluates the SHACL constraints over the nodes of the input KGs and outputs a validation report per constraint and target entity. These results indicate whether an entity validates or invalidates the constraints defined by the user. This validation reports state the validity of data used by the predictive models[7]. Thus, SHACL constraints are a crucial part of the InterpretME KG, to identify if a particular entity is violating the integrity constraints, where *True* represents the particular entity is valid, inversely represented by *False*.

2. Data preparation includes transforming data collected from the input KGs into a form that can be used in the training of a predictive model. Many machine learning models cannot handle categorical values directly, InterpretME generates one hot-encodings using the Python library sklearn to transform the model's categorical features into binary features usable for predictive models[8]; it makes training data more expressive, and can be easily re-scaled. Also, the target class can be defined in InterpretME. The Lung Cancer KG has many categorical features, e.g., Age, Gender, Biomarker, CancerType, and FamilyDegree.

Sampling strategies (e.g., under-sampling or over-sampling) are also received to reduce data imbalance so that the machine learning algorithms can perform better. Many machine learning algorithms, like decision trees, random forests, and neural networks resort to class distribution in the training dataset to compute the probability of instances in each class when the model will be used to make predictions. Model building can be done based on automated systems' selections (e.g., AutoML/OpenML) or based on a user's preferences. Automated tools (e.g., Optuna [21]) can be used to generate optimized hyperparameters for predictive models. Here, the automated model **3.** can also perform stratified shuffle split cross-validation with random forest, and identify the relevant features; they are used to train a decision tree classifier to predict and visualize the outcomes[9].

[6] https://github.com/SDM-TIB/InterpretME/blob/v1.2.0/example/example_kg_french_royalty.json
[7] https://github.com/SDM-TIB/InterpretME/blob/v1.2.0/InterpretME/pipeline.py#L153-L185
[8] https://github.com/SDM-TIB/InterpretME/blob/v1.2.0/InterpretME/preprocessing_data.py#L57-L84
[9] https://github.com/SDM-TIB/InterpretME/blob/v1.2.0/InterpretME/classification.py

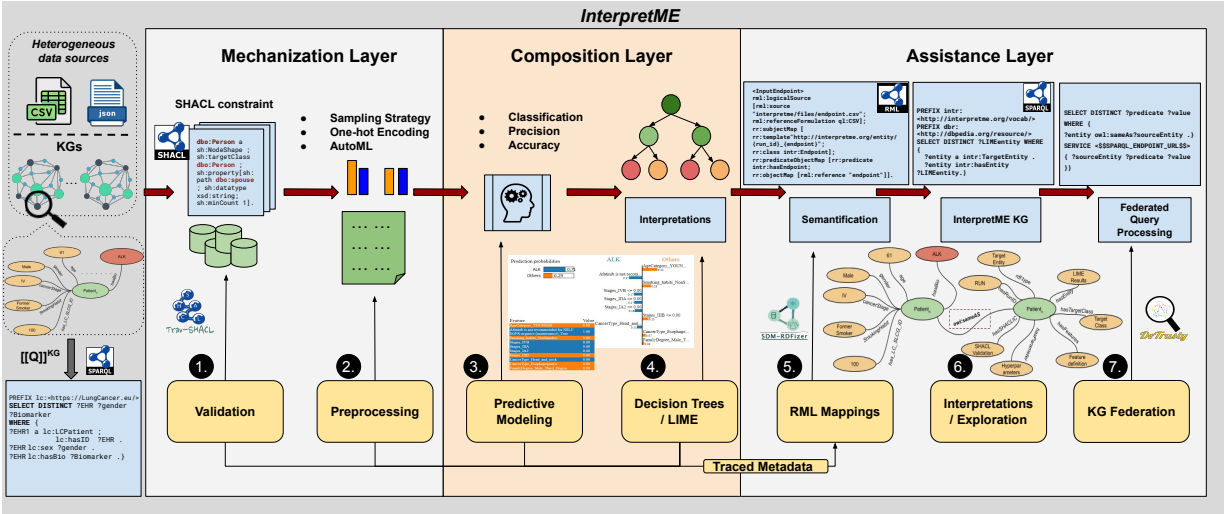

Fig. 5. **The InterpretME Architecture.** The input of InterpretME is either Knowledge graphs or datasets in formats, i.e., CSV or JSON. In *Mechanization* layer, SHACL validation, data preprocessing, and AutoML are utilized. The *Composition* layer performs the predictive task, and the interpretable tools like Decision trees and LIME are implemented for understanding the predictions. In *Assistance* layer, InterpretME generates an InterpretME KG with the traced metadata of the trained predictive model, to provide users with more enhanced and reliable interpretations. Generating and Exploiting the InterpretME KG involves RML mappings and Federated Query Processing.

4. Interpretation is utilized to enhance the understanding of the trained predictive model. Also, a visualization of the decision tree can be obtained with associated constraints. This module provides more insights of a validated/in-validated feature, and interpretable surrogate tools like LIME [5] (as in Figure 3a) yields the most relevant features and local explanations list. They reflect the contribution of each feature to the prediction.

5. Knowledge graph creation is a module of InterpretME, where mapping rules are defined with the RDF Mapping Language (RML) [22] using the data collected from the predictive models. RML triple map; it is composed of *rml:logicalSource*, *rr:subjectMap*, *rr:predicateObjectMap*. The mapping rule component **(1)** states the logical source using the term *rml:logicalSource*. **(2)** The *rr:subjectMap* element defines the class of subject. **(3)** The term *rr:predicateObjectMap* establishes how the predicate *intr:hasDefinition* will be populated with the collected data.

6. SPARQL queries can be executed to explore the InterpretME results of a particular target entity. Users can identify patterns and correlations between the given features and also perform statistical analysis of a particular target entity. The InterpretME KG provides clarification and eases the interpretation of the model's prediction of a particular entity aligned with the SHACL validation results.

7. Federated query processing is utilized to query both the input KGs and the InterpretME KG. At this point, the user is involved in executing federated SPARQL queries to gain more insights from the predictive model's decision. As a result, the advantage of comparing the predictive model's characteristics in the InterpretME KG to the characteristics of a specific entity in the input KGs is obtained. Mostly interpretable tools results are related to the target entity in the dataset, which is hard to interpret and understand the characteristics of the target entity. Thus, a user can trace back via query federation to interpret the results for a specific target entity with input characteristics aligned with SHACL validation results. Exemplar queries are shown in InterpretME GitHub repository[10].

3.3. Running example

To exemplify different components of the InterpretME architecture, we have used the following running example (Figure 6). The prerequisites to run an example of the **French Royalty KG** with InterpretME is available[11].

[10] https://github.com/SDM-TIB/InterpretME/tree/v1.2.0/example/queries/template
[11] https://github.com/SDM-TIB/InterpretME/v1.2.0/example/InterpretME_french_royalty.ipynb

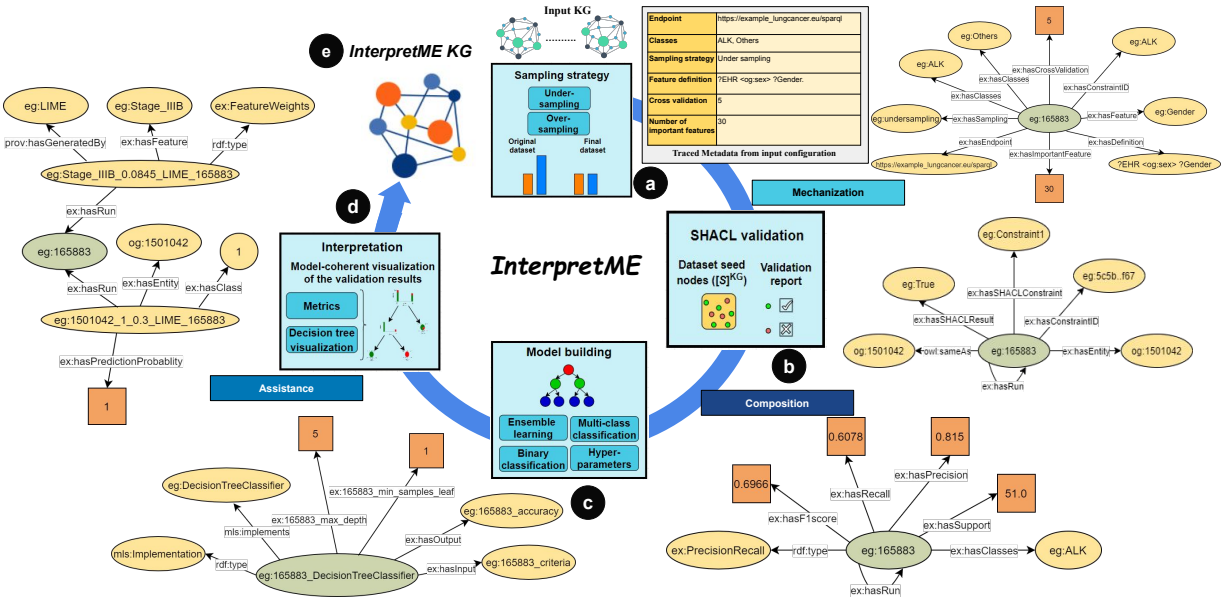

Fig. 6. **Running example** demonstrates the InterpretME internal architecture over the Lung Cancer KG. At each stage, all that InterpretME traces about the features' definition, SHACL validation, and the predictive model characteristics are shown as an RDF graph. InterpretME integrates all these traced metadata to generate the InterpretME KG.

a (Lung Cancer KG). This KG integrates the features' and class target definitions about lung cancer patients; their constraints are defined in terms of SHACL. Features' definitions are classified into independent and dependent variables. Independent variables are features that the user selects for analyzing the model's prediction. On the other hand, a dependent variable is a feature under investigation that is expected to change in response to changes in the independent variable. They are used later in the predictive modeling pipeline and are defined as follows:

```
{
  "Endpoint": "https://example_lungcancer/sparql",
  "Index_var": "EHR",
  "Independent_variable": {
    "Gender": "?EHR <https://LungCancer.eu/vocab/sex> ?Gender."
  },
  "Dependent_variable": {
    "Biomarker": "?EHR <https://LungCancer.eu/vocab/hasBio> ?Biomarker."
  }
}
```

where **EHR** is the electronic health record of a patient, the first part states the feature name (i.e., *Gender*), and the later part (i.e., *<https://LungCancer.eu/vocab/sex>*) describes the feature in the KG. Figure 6 shows the traced input configuration, i.e., features' definitions, sampling strategies, class definitions, number of important features, SPARQL endpoint that are represented in the form of *RDF triples*, later used for the **InterpetME KG**. In the current running example Figure 6, the target class (Biomarker) is imbalanced (i.e., *patient with target class ALK is 872 instances while a patient with Others as target class is 432*). Therefore, the under-sampling technique[12] is selected for this use case to tackle the problem of imbalance.

b Considering our running example in Figure 6, the SHACL shapes represent integrity constraints. A SHACL schema defines integrity constraints, used by InterpretME for validation results and for uncovering their impact on a model's decisions. Here, the *"constraint": "Afatinib is not recommended for NSCLC EGFR negative (hasDrug)"* is applied over the extracted data from input KG. The above constraint states a medical protocol that Afatinib is a

[12] https://github.com/SDM-TIB/InterpretME/blob/v1.2.0/InterpretME/sampling_strategy.py

drug that is not recommended for a patient having EGFR negative biomarker. Figure 6 depicts the traced metadata about SHACL validation, i.e., *EHR: 1501042 with EGFR positive satisfying the constraint*, an entity alignment is performed to trace original entity of input KG with SHACL validation results and predictive modeling pipeline.

c The preprocessed and sampled data with optimized hyperparameters is then fed to the automated tools for model building (e.g., ensemble learning) to perform predictive task. Here, the automated model can also perform stratified shuffle split cross-validation with models like *Random forest, AdaBoost classifier, Gradient boost classifier*, and identify the relevant features; they are used to train a decision tree classifier to predict and visualize the outcomes[13]. InterpretME stores the metadata about the evaluation of the model for a particular run.

d The trained predictive model in **c** can be visualized by means of *Decision trees*; they can be visualized with SHACL constraints to observe which subtree validates or invalidates the protocol[14]. To understand the predictive model's outcomes, InterpretME resorts to interpretable tools, e.g., LIME, to provide a local explanation of each entity. LIME provides the prediction probability of the target class *ALK (i.e., Class 0)* as 0.65 and target class *Others (i.e., Class 1)* as 0.35. The relevant features list is often used to illustrate which features may cause a change in the prediction of the trained model. LIME results are also represented in the InterpretME KG.

e Traced metadata can be semantified using RML mappings. RML mapping rules specify the role of transforming collected metadata into RDF triples for the InterpretME KG[15]. The unified schema of InterpretME represents the meaning of data in the InterpretME KG[16]. The following RML syntax is used to define one of the mapping file:

```
<InputEndpoint>
  rml:logicalSource [ rml:source "interpretme/files/endpoint.csv";
                      rml:referenceFormulation ql:CSV; ];
  rr:subjectMap [
    rr:template "http://interpretme.org/entity/{run_id}_{endpoint}";
    rr:class intr:Endpoint ];
  rr:predicateObjectMap [
    rr:predicate intr:hasEndpoint;
    rr:objectMap [ rml:reference "endpoint" ] ].
```

SDM-RDFizer [23] integrates entities collected during the training of a ML model, into the InterpretME KG as factual statements. A user can query factual statements of the InterpretME KG to compute statistics about the results of a data-driven pipeline. In the appendix 7 of the paper, we have included some queries to be executed on the SPARQL endpoint of InterpretME whose answers provide insights into the behavior of the trained ML models. RML provides a declarative definition of the classes and properties of the InterpretME ontology and ML schema to the collected entities. As a result, the InterpretME KG comprises machine and human statements that document the trained ML behavior. The generated RDF data is uploaded to an instance of *Virtuoso*. Entities in the InterpretME KG and the input KGs are aligned to ensure traceability. InterpretME can accept input from multiple KGs and therefore the process of entity alignment as shown in the figure Figure 7 plays a vital role. The federation of KGs requires the identification of a specific entity, as well as the source of the entity. As a result, the URI of an entity in the input KGs differs from the URI in the InterpretME KG. The InterpretME federated query engine allows for tracing back the ML models' results, i.e., main properties of target entities, LIME prediction probabilities, and SHACL validation reports. InterpretME provides an advantage of analyzing target entities not only based on ML models but also according to how all the properties that characterize the entity. The short tutorial is available to demonstrate the use of InterpretME in the example[17] of the GitHub repository.

4. Empirical Evaluation

We evaluate InterpretME with the goal of answering the following research questions: **RQ1)** What is the impact of integrating predictive modeling frameworks with KGs to enhance interpretability? **RQ2)** Can InterpretME trace

[13] https://github.com/SDM-TIB/InterpretME/blob/v1.2.0/InterpretME/classification.py
[14] https://github.com/SDM-TIB/InterpretME/tree/v1.2.0/images
[15] https://github.com/SDM-TIB/InterpretME/tree/v1.2.0/InterpretME/mappings
[16] https://github.com/tibonto/InterpretME
[17] https://github.com/SDM-TIB/InterpretME/blob/v1.2.0/example/InterpretME_french_royalty.ipynb

decisions made by predictive models? **RQ3)** To which extent does the InterpretME KG satisfy standard quality criteria? **RQ4)** How much is the observed overhead at each stage of InterpretME? The experimental settings utilized to evaluate InterpretME are as follows:

Benchmarks. Table 1 presents our KGs in numbers. **The French Royalty KG [24]** is fully curated. For each person in the KG, we added the class `dbo:Person`, the number of children, the number of predecessors, and triple-related counts like the number of distinct predicates. Furthermore, we added the ground truth for the relationship `hasSpouse`. Forecasting whether a member of the French Royalty has a spouse is our predictive task. We created ten SHACL constraints of which some are inspired by the logical rules presented in [24], e.g., 'if I have a child, this child has another parent which is not me'. The other constraints follow a similar pattern, e.g., if a person has two different parents, the parents are likely to be married. We use a private KG from the biomedical domain – **The Lung Cancer KG [25]** – which integrates lung cancer patients mentioned as in Section 2. This KG includes information describing the main characteristics of a lung cancer patient.The prediction task is a binary classification to predict the biomarker of the patient, which can be 'ALK' or 'others'. The SHACL constraints used are the medical protocols that recommend when treatments should be prescribed according to a patient's biomarkers; we defined four different SHACL constraints stating that 'EGFR negative' patients should not take *Afatinib* or *Gefitinib*.

Machine Learning Models. InterpretME resorts to ensemble learning techniques to solve the previously described prediction tasks. Random forests create multiple decision trees by bootstrapping datasets created from the KGs, and randomly selecting a subset of variables at each step of the decision tree. A random forest generates the list of relevant features iteratively with a stratified shuffle split cross-validator (5 fold-validation with training 70% and testing 30%). A decision tree uses this list to train the predictive model. The generated classification decision tree is easy to understand due to its visualization of the decision process. Accuracy without the random-forest stratified shuffle split is 69% in the Lung Cancer KG predictive task and 95% in the French Royalty KG, but with our approach, accuracy is increased to 77.50% and 96.17%, respectively.

Metrics. *Execution Time*: The time elapsed at each component of the pipeline is reported. The reported times correspond to absolute wall-clock system time in seconds as reported by Python's `time.time()` method. The average execution time of five runs is reported. The experiments are run on a machine equipped with an Intel® Core(TM)i7- 10850H at 2.71 GHz and 16 GiB RAM.

4.1. Degree Distribution of the InterpretME KG

In graph theory and complex networks, an entity's degree (or the number of neighbors ignoring edges' direction) is defined as the number of relations with other entities. Increasing the degree of an entity expresses a richer, more accurate view of the particular entity. Better context means better interpretability, not just volume. For instance, in the predictive task about the patient positive for biomarker ALK. The predictive model states the decision of a particular target entity, and it is obscure how the decisions are been made, with just knowing the node of a patient and the features. While InterpretME tries to increase more contextual edges of a patient node via annotating the other contextual information and behavior about the patient (e.g., Validation) in the pipeline. In a nutshell, increasing the degree of a node with more context makes not only human-understandable easier but also machine-readable. InterpretME traces the entities of the target classes (e.g., ALK) defined as independent variables via SPARQL queries over the input KGs. In the InterpretME KG, these entities are described in terms of metadata collected by InterpretME components (e.g., hasInterpretedFeature, hasFeatureWeight, hasEntityClassProbability, hasPredictionProbability, hasSHACLResult, hasSHACLConstraint). For instance, Figure 7 shows entity alignment. Given a predictive task over the Lung Cancer KG, we know the main characteristics of a target entity `lc:11315856` (i.e., LCPatient, Age, smoking habit, biomarker) in the input KG as well the characteristics from the predictive model (i.e., prediction, accuracy, relevant features, feature contribution, hyperparameter's, model, target class), and also the knowledge about the entity satisfying the domain constraint. The entity in the input KG (`lc:11315856`) is aligned with the target entity via `owl:sameAs`. InterpretME utilizes Leibniz's Inference Law [26], $E[X := R] = E1$, where $E1$ is the result of simultaneously replacing in E all the instances of $x_i \in X$ by the corresponding expression $R_i \in R$, InterpretME entails all the characteristics of the entity `lc:11315856`. Here, the outdegree of entity `lc:11315856` is 6 in the input KG, while in the InterpretME KG is 11 with the traced metadata, while in the federation of the

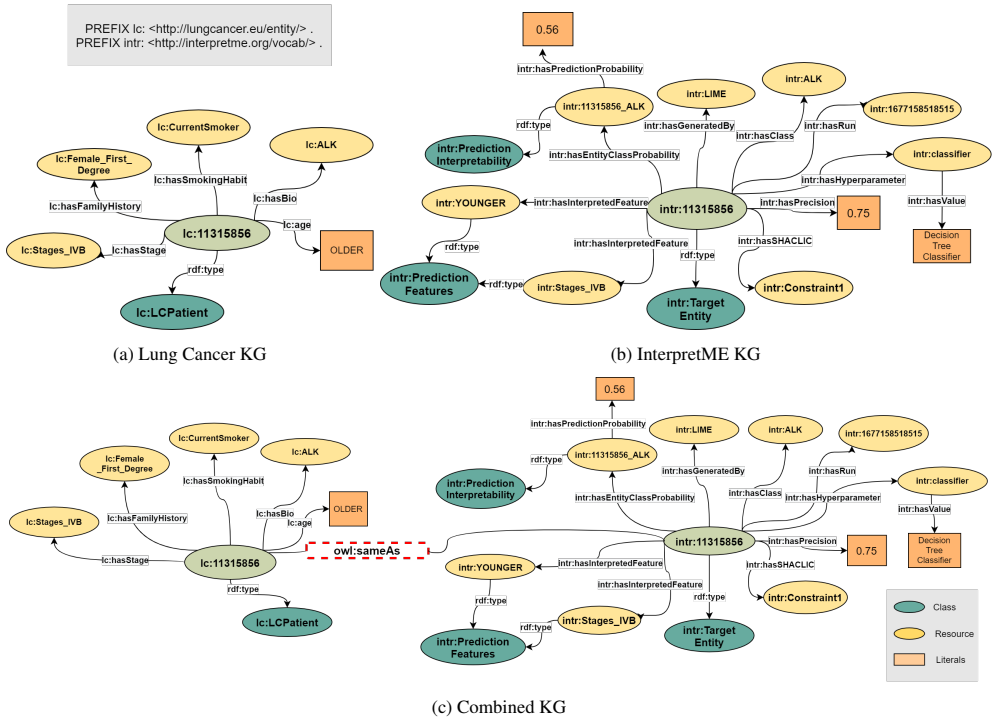

(a) Lung Cancer KG (b) InterpretME KG

(c) Combined KG

Fig. 7. **Entity Alignment**. Figure 7a shows a target entity `lc:11315856` from the Input KG with the features, i.e., age, gender, smoking habit, and biomarker. Figure 7b illustrates the InterpretME KG of a target entity (`intr:11315856`) with the traced metadata from the predictive model. Figure 7c demonstrates the entity from the Lung Cancer KG is aligned with the target entity from the InterpretME KG via `owl:sameAs`. The out-degree distribution of `lc:11315856` is 6 in the input KG, 11 in the InterpretME KG, and goes to 17 in the Combined KG.

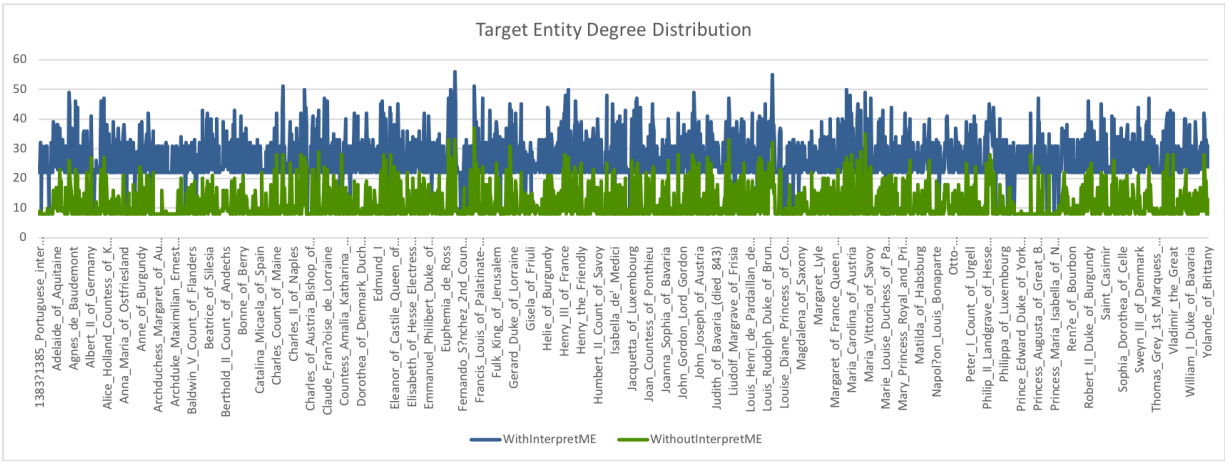

Fig. 8. Degree distributions of the entity targets in the French Royalty KG. The distribution in blue depicts the degrees of the target entities semantically enhanced with InterpretME, while green shows their original degree distribution.

input KG and the InterpretME KG, the outdegree distribution increases to 17. Thus, more contextual edges are represented in the Combined KG and provide a user with more detailed interpretations of the model's decision.

Figure 8 depicts the degree distribution results of French Royalty KG; the average number of neighbors is *With-InterpretME*: 27.19 (with a standard deviation of 6.13) and 11.39 (with a standard deviation of 5, 06) *WithoutInterpretME*. InterpretME increases the number of RDF triples that describe a target entity; a user can query these triples

to explore the predictive model decisions. The InterpretME KG properties are retrieved. These values of degree distribution quantify the information gained for each target entity in terms of all that InterpretME traces. The execution of queries 1 and 4 [18] over the InterpretME KG retrieves the values of these properties for interpretME:Louis_XIV. As a result of tracing the machine learning pipeline, the number of properties of interpretME:Louis_XIV (i.e., out-degree) goes from 32 in the French Royalty KG to 55, and thus, the entity's degree is increased. Based on these outcomes, we can answer the **RQ1)**, and InterpretME enhances the interpretability of a target entity dbr:Louis_XIV and provide the user more contextual insights into the prediction task.

Table 1
Statistics about Input KG and InterpretME KG

Knowledge Graph		#triples	#entities	#predicates	#objects	#triples /#entities
French Royalty	Input KG	31,599	3,439	133	4,390	9.18
	InterpretME KG	245,981	30,914	152	35,622	7.96
Lung Cancer	Input KG	77,466,343	21,420,170	357	15,053,863	2.30
	InterpretME KG	62,393	8,877	152	11,546	7.03

4.2. Traceability

We evaluate InterpretME in terms of the traceability of a target entity. Table 2 reports on the average number of answers to the type of questions presented in our motivating example (Section 2). InterpretME efficiently traces the target entity and provides the user with additional information about the prediction probability of the entity. Also, it helps users to uncover relevant features of an entity that contribute to the prediction, with assigned weight distribution for the top-10 features. The federated query engine, DeTrusty [27], evaluates SPARQL queries to retrieve data from the original KG, the InterpretME KG, or both. Instances in the InterpretME KG are linked to the entity in the original KG via `owl:sameAs`. The KGs are accessible via SPARQL endpoints. In Table 2, the questions presented in Section 2 are expressed as SPARQL queries over the two KGs. For the French Royalty KG, InterpretME models LIME interpretations in terms of 10 RDF triples (on average), and the contribution of a given feature is documented with 2 RDF triples on average for the binary classification task. Moreover, target entities satisfy the 10 domain integrity constraints, and the average number of main characteristics of the target entities in the French Royalty KG is 12.24. The average number of answers in the Lung Cancer KG is higher because a lung cancer patient is described in more detail than a person in the French Royalty KG. Based on these outcomes, we can positively answer **RQ2)** because InterpretME can translate back to the original KGs, the decisions made by the predictive models.

Table 2
Average Number of Answers per Target Entity to Questions from Figure 3a

ID	Question	Avg Answers in	
		Lung Cancer KG	French Royalty KG
Q1	Which is the target entity interpreted by LIME?	20.57	10.00
Q2	How does *feature* contribute to the classification of this entity?	2.05	2.00
Q3	Which other features are relevant for this classification?	11.65	4.47
Q4	Does this target entity satisfy the domain integrity constraints?	4.00	10.00
Q5	What are the main characteristics of the target entity?	46.28	12.24

[18] https://github.com/SDM-TIB/InterpretME/blob/v1.2.0/example/queries/french_royalty

Table 3
Evaluation Metrics

Metric	Value	Explanation	Formula (in [28])
Accuracy			Section 2.2.1
Synt. validity of RDF doc	1	RDF files are valid	$m_{synRDF}(g)$
Synt. validity of literals	1	Literals have a label that defines a range of possible values	$m_{synLit}(g)$
Semant. validity of triples	1	Literals are collected from input or can be measured	$m_{semTriple}(g)$
Trustworthiness			Section 2.2.2
KG level	0.25	Automated data curation from structured data sources	$m_{graph}(hg)$
Statement level	0.5	Provenance is traced	$m_{fact}(g)$
Unknown/empty values	0	Empty values are not indicated	$m_{NoVal}(g)$
Relevancy			Section 2.3.1
Ranking of statements	0	Ranking is not useful in this context	$m_{Ranking}(g)$
Ease of Understanding			Section 2.4.1
Description of resources	1	Resources have label and comment	m_{Descr}
Labels in multiple lang	0	Labels are only available in English	$m_{Lang}(g)$
RDF serialization	1	Serialization in Turtle	$m_{uSer}(hg)$
Self-describing URIs	0.5	Partial use of self-describing URIs	$m_{uURI}(g)$
Interlinking			Section 2.5.3
Interlinking via owl:sameAs	1	Use of owl:sameAs for linking the InterpretME KG to the input KG	$m_{Inst}(g)$

Components	French Royalty KG	Lung Cancer KG
SHACL Validation	1.46	67.19
Data Curation	0.30	0.47
Model Training	3.14	4.22
Decision Trees	1.39	1.92
Constraint Visualization	0.01	0.01
LIME	58.09	20.18
Semantification	42.87	15.87

(a) Average execution time (secs.) for KGs per pipeline stage

Interpretations	Avg. Response
LIME	6.23 / 10
InterpretME	8.17 / 10
ML knowledge	3.35 / 5
Interpretability	Familiar (9/17)

(b) User study results (score values are high is better)

Fig. 9. **Evaluation** Figure 9a presents the impact of execution time in seconds at each stage of InterpretME pipeline. Figure 9b reports on a user study conducted with 17 participants ages ranging from 20 to 45, on average the participants were familiar with the concept of Interpretability, and the experience level of machine learning is 3.35 and the participants rated LIME interpretation 6.23, while InterpretME was rated with 8.17.

4.3. The Quality of the InterpretME Knowledge Graph in Numbers

This section answers **RQ3)** and reports the results of the evaluation of the InterpretME KG in terms of the quality metrics proposed by Färber et al. [28]; four data quality categories are considered: **i)** Intrinsic category is independent of the use case context; **ii)** Contextual category depends on the application context of a data consumer; **iii)** Representational category quantities the form how information is available; and **iv)** Accessibility category determines how data can be accessed. Table 3 depicts all the metrics by category from Färber et al. [28] to check the quality of the InterpretME KG. They help the user to better estimate the pros and cons of InterpretME. Färber et al. [28] propose extended quality parameters and categories (i.e, trustworthiness and interlinking). The scores for quality metrics (Table 3) are assigned manually, following the scores' definition by Färber et al. [28]. Since, the InterpretME KG covers a new domain of structured data, the eloquence of several suggested criteria is still relatively low. However, we are optimistic that these values will be increased once real-world applications begin to use InterpretME.

4.4. Execution Time

This section describes Figure 9a and aims at answering **RQ4)**, i.e., the overhead caused by each of the components of InterpretME. The analysis of the execution time reveals the impact of different parameters on the performance of the InterpretME steps. We have calculated the average execution time based on the two use cases, French Royalty and Lung Cancer KGs by taking the average of five runs per use case. The French Royalty KG comprises RDF

triples 31,599, while the Lung Cancer KG comprises 77,466,343 RDF triples (Table 1), i.e., the Lung Cancer KG is 3.8 orders of magnitude larger than the French Royalty KG. The average execution time of the first component, i.e., the SHACL validation, is 1.46 and 67.19 seconds, respectively. The second component – data curation – where data is extracted and preprocessed to be given to the machine learning model, takes on average 0.30 and 0.47 seconds in the two use cases, respectively. The next component is training the model; the average execution time is 3.14 and 4.22 seconds for the French Royalty KG and the Lung Cancer KG, respectively. Further, the decision tree component, where the decision trees for the model are generated, takes an average 1.39 and 1.92 seconds, respectively. For French Royalty, the number of target entities is slightly higher compared to the Lung Cancer KG. Thus, using an interpretable tool (e.g., LIME), generating the interpretable results requires 58.09 and 20.18 seconds in the use cases, respectively. Lastly, the *semantic enrichment* of the traced metadata takes 42.87 and 15.87 seconds for the French Royalty KG and the Lung Cancer KG, respectively. The reported results are consistent with the fact stated by Figuera et al. [20] that the SHACL validation is mainly impacted by the KG size and number of SHACL constraints. The validation of the Lung Cancer constraints takes much longer than for the French Royalty, since the Lung Cancer KG is much larger and the constraints used are more complex, e.g., we are using the concept of linked shapes for the Lung Cancer KG. As a result, the ML pipeline execution time is also impacted, as well as the generation of the InterpretME KG instances.

4.5. User Study

We report the results of a user study[19] about interpretability in InterpretME. We aim to assess how each participant understands the term *interpretability*. Further, we conduct a user experience evaluation of the framework based on usability criteria. The goal is to answer the following questions: **1)** *Is it beneficial to integrate knowledge graphs for interpretability?* The possible answers are: **i)** the impact of knowledge graphs in interpretability; **ii)** without knowledge graph is enough; **iii)** Maybe. **2)** *How clear is the notion of the interpretation from InterpretME on an increasing scale from* $1-10$*?* This question aims to assess the conceptual and practical understanding of the InterpretME KG. **3)** *Do we have a positive or negative attitude towards the tool?* This question reveals the attitude of participants towards InterpretME in terms of usability, reliability, and trustability. During a lecture of data management, we shared the overall 12 questions to **17 participants** with the education ranging from the biomedical and computer science students to professors. The vast majority of our participants had prior knowledge of ML. In the first question, most participants believed that KGs play a significant role in the interpretability of ML models and comply with the fact mentioned in [13]. The second question focuses on the benefit of the interpretation given by the InterpretME KG. Almost all the participants (16/17) struggled in finding the interpretation belonging to which target entity from LIME; the interpretability of LIME output was rated with an average of 6.23. In contrast, the interpretability of InterpretME is placed with an average of 8.17. The third question focuses on assessing the usability and effectiveness of InterpretME. We observed that participants did not understand the interpretations of LIME. They rated InterpetME with a positive attitude, suggesting that the InterpretME KG may provide more insights into a particular prediction. The participants familiar with machine learning and interpretability, better understand LIME and InterpretME interpretations. These findings suggest that the interpretations generated from InterpretME are machine-readable and understandable by humans.

5. InterpretME as a Tool

InterpretME is publicly available as a Python library on PyPI[20]. While the core of the pipeline is entirely new, InterpretME reuses available tools from the community. InterpretME *pipeline()* receives a JSON file as a configuration input from the user to extract all the features' definition, SPARQL endpoints or Datasets, SHACL constraints, target classes, and sampling strategies. Features are defined in the form of independent and dependent variables, provided to the InterpretME pipeline to perform the prediction tasks. For the SHACL validation, InterpretME relies

[19] https://docs.google.com/spreadsheets/d/1UIwwmmOEbNtx-_IN8_5sPg_QzU5WwUXLRgZCowefLtI/edit?resourcekey#gid=836713689
[20] https://pypi.org/project/InterpretME/

on Trav-SHACL [20] since it is capable of validating the shape schema against a SPARQL endpoint and scales better compared to other approaches with the same capability. InterpretME provides pre-processing component for data curation, optimize hyperparameters, and sampling strategy. Model building of InterpretME can provide binary and multi-class classification with ensemble learning techniques. In model building, cross-validation and classification report are obtained for a particular predictive model. Currently, InterpretME uses LIME [5] to create model interpretations. The Mapping Language (RML) [22] defines the process of integrating the traced metadata into the InterpretME KG; it is semantified using SDM-RDFizer [23] (version 4.5.5 [29]), an efficient RML-compliant engine for KG creation. The RDF data is uploaded into an instance of Virtuoso 7.20.3233 representing the InterpretME KG. The following *pipeline()* command executes the whole pipeline; including extracting data and metadata from the input KGs, validating SHACL constraints, preprocessing the data, running predictive models, semantifying the results, and populating the InterpretME KG:

```
from InterpretME import pipeline
results = pipeline(path_config='./example.json', lime_results='./LIME',
    server_url='endpoint of InterpretME KG', username='username to upload to InterpretME KG',
    password='password to upload to InterpretME KG')
```

The InterpretME ontology extends ML Schema [30]. ML Schema is considered a core vocabulary to deal with machine learning algorithms; it can be used to address different algorithms, implementation, model evaluation, and the input and output considered by the algorithms. The ontology represents relationships between machine learning algorithms and their executions. ML Schema also deals with interoperability issues and allows reproducibility. InterpretME reuses 12 concepts in the mappings from ML Schema as shown in the Figure 2a. For instance, the classes `mls:Run`, `mls:Implementation`, `mls:ModelEvaluation`, `mls:HyperParameterSetting`, `etc` and the relations `mls:hasInput`, `mls:hasOutput`, `mls:hasHyperParameter`, `etc` demonstrates the benefit to link machine learning algorithm results into linked data. InterpretME incorporates additional functionality like validating SHACL constraints over the input KG `intr:SHACLValidation`, `intr:hasSHACLConstraint`, `intr:hasSHACLResult`, `etc`, and interpreting the predictions of the ML models `intr:PredictionInterpretability`, `intr:hasPredictionProbability`, `etc` and many more demanded an extension to the ML Schema as shown in Figure 2b and Figure 2c. Following the FAIR principles [31], the InterpretME ontology is publicly available on an instance of VoCoL[21]; it enables the publication, management, and exploration of the ontology.

```
from InterpretME.federated_query_engine import configuration, federated
interpretme_endpoint = 'SPARQL endpoint of InterpretME KG'
input_endpoint = 'SPARQL endpoint of input KG'
config = configuration(interpretme_endpoint, input_endpoint)
query_answer = federated(input_query, config)
```

The data from the original KGs, as well as the InterpretME KG, can be queried using the federated query engine DeTrusty (version 0.11.1 [27]). DeTrusty is based on MULDER [32], i.e., it uses semantic source descriptions during decomposition and planning. Rohde [33] describes the vision of incorporating the SHACL validation result into SPARQL query answers by executing the validation during query processing. While we are not providing an engine fulfilling this vision, a similar outcome can be achieved when querying the original KG together with the InterpretME KG since the SHACL validation result is part of it. InterpretME is a stand-alone framework that works locally on individual systems, with multiple runs that can be combined into a single InterpretME KG to compare interpretations. InterpretME comprises the following resources: **i)** the InterpretME pipeline enhancing interpretability of machine learning models; **ii)** the InterpretME KG describing the predictive model characteristics; and **iii)** the InterpretME ontology specifying the vocabulary for describing the main characteristics of trained predictive models.

InterpretME is utilized on top of the RDF KGs of the following projects: **1)** CLARIFY[22] to define machine learning models to predict biomarkers of a lung cancer patient based on demographic features (e.g., age and gender),

[21] http://ontology.tib.eu/InterpretME/

[22] EU H2020 Funded project https://www.clarify2020.eu/

smoking habits, and relatives with cancer; **2)** ImProVIT[23] to develop models to predict the impact of the immune system and demographic features into the response of Hepatitis B and Influenza vaccines; and **3)** P4-LUCAT[24] to implement predictive models to forecast the relapse after surgery or the disease progression in advanced stages.

6. Related Work

Tools for Supporting Interpretability. Artificial intelligence and machine learning have become global in many domains. The necessity of automated machine learning frameworks with assistance has gained popularity in the research field, in every sector. De Bie et al. [3] discuss the challenges to achieve automated machine learning and highlight that existing tools contribute to mechanization and composition automatization, lacking support in assistance. InterpretME also aims at bridging this gap, and resorts to Semantic Web technologies to enhance users' assistance. Lundberg et al. [4] propose an interpretation framework called SHAP based on coalition game theory (Shapely values). SHAP provides feature contributions for each individual instance, global explanations, and feature importance. Ribeiro et al. [5] present LIME, an approach for local surrogate models, which are used to explain individual predictions of a pipeline. Local surrogate models are trained to approximate the predictions of models locally, instead of training a surrogate model globally. As a result, LIME generates human-friendly explanations for target entities. However, these explanations are not machine-readable and cannot be translated into the domain application. InterpretME overcomes these limitations, and provides fine-grained representations of local interpretations which are linked to the target entities in the domain application KGs.

Semantic Web Technologies in Machine Learning. Semantic Web technologies like ontologies are used to improve the accuracy of predictive models. Ristoski et al. [34] provide a comprehensive survey of the use of Semantic Web Technologies in data mining and knowledge discovery and highlights the potential benefits and challenges of the need for more efficient algorithms and tools in practice. Kulmanov et al. [35] study the role of ontologies in semantic similarity and machine learning, and present ontology embeddings as background knowledge. Further, ontologies can constrain the output of a machine learning model, i.e., making the model consistent with the axioms of the ontology. Min et al. [36] improve the performance of predictive models by using ontological adjustments, i.e., using the hierarchy of an ontology to move samples of rare classes into the next broader concept. On average, Min et al. [36] reduce the area under the receiver operating curve (AROC) by 9.0% in predicting the effectiveness of antidepressants in patients with rare conditions. Haug et al. [37] propose the combination of a large enterprise data warehouse with medical knowledge from a disease-oriented ontology. This combination allows for automating the generation of computable diagnostic models, which aim at supporting researchers in generating and evaluating tools for real-time clinical diagnosis. Similarly, InterpretME resorts to Semantic Web technologies to enhance not only accuracy but interpretability as well.

7. Conclusions and Future Work

InterpretME empowers predictive models with metadata traced along a predictive modeling pipeline. The InterpretME KG reuses concepts defined in the ML Schema and resorts to a state-of-the-art SHACL engine for efficient constraint validation. Empirically, we have observed that InterpretME empowers the description of a predictive model's insights using factual statements and links to the application domain KGs. As a result, InterpretME broadens the portfolio of Semantic Web tools to enhance domain understating by bridging the gap between data meaning and predictive modeling [38]. The current version of InterpretME considers random forests, decision trees, and the LIME interpretable model. Nevertheless, we plan to integrate other models and tools. Furthermore, connecting InterpretME with automated machine learning systems and causal knowledge graphs [39] is part of our future work.

[23] German Funded project https://www.tib.eu/en/research-development/project-overview/project-summary/improvit
[24] EraMed project https://p4-lucat.eu/

References

[1] M.A. Myszczynska, P.N. Ojamies, A.M.B. Lacoste, D. Neil, A. Saffari, R. Mead, G.M. Hautbergue, J.D. Holbrook and L. Ferraiuolo, Applications of machine learning to diagnosis and treatment of neurodegenerative diseases, *Nat Rev Neurol* **16** (2020), 440–456–.

[2] M.F. Pinto, H. Oliveira, S. Batista, L. Cruz, M. Pinto, I. Correia, P. Martins and C. Teixeira, Prediction of disease progression and outcomes in multiple sclerosis with machine learning, *Sci Rep* **10** (2020).

[3] T. De Bie, L. De Raedt, J. Hern_'andez-Orallo, H.H. Hoos, P. Smyth and C.K.I. Williams, Automating Data Science, *Communications of the ACM* **65**(3) (2022), 76–87.

[4] S.M. Lundberg and S.-I. Lee, A Unified Approach to Interpreting Model Predictions, in: *NIPS '17: Proceedings of the 31st International Conference on Neural Information Processing Systems*, Curran Associates Inc., Red Hook, NY, USA, 2017.

[5] M.T. Ribeiro, S. Singh and C. Guestrin, "Why Should I Trust You?": Explaining the Predictions of Any Classifier, in: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, ACM, New York, NY, USA, 2016.

[6] P.I. Nakagawa, L.F. Pires, J.L.R. Moreira, L.O. Bonino da Silva Santos and F. Bukhsh, Semantic Description of Explainable Machine Learning Workflows for Improving Trust, *Applied Sciences* **11**(22) (2021).

[7] R. Prasad, L. Joseph and R.C. Deo, Modeling and Forecasting Renewable Energy Resources for Sustainable Power Generation: Basic Concepts and Predictive Model Results, in: *Translating the Paris Agreement into Action in the Pacific*, Springer, Cham, 2020.

[8] L. Kotthoff, C. Thornton, H.H. Hoos, F. Hutter and K. Leyton-Brown, Auto-WEKA 2.0: Automatic model selection and hyperparameter optimization in WEKA, *J. Mach. Learn. Res.* **18** (2017), 25:1–25:5. https://www.cs.ubc.ca/labs/algorithms/Projects/autoweka/papers/16-599.pdf.

[9] T. Berners-Lee, J. Hendler and O. Lassila, The Semantic Web: A New Form of Web Content That is Meaningful to Computers Will Unleash a Revolution of New Possibilities, 2001. http://www.sciam.com/article.cfm?articleID=00048144-10D2-1C70-84A9809EC588EF21.

[10] O. Lassila and R.R. Swick, Resource Description Framework (RDF) Model and Syntax Specification, 1999. https://www.w3.org/TR/1999/REC-rdf-syntax-19990222/.

[11] M. Arenas, C. Gutierrez and J. Pérez, Foundations of RDF databases, in: *Reasoning Web. Semantic Technologies for Information Systems. Reasoning Web 2009. Lecture Notes in Computer Science*, Vol. 5689, Springer, Berlin, Heidelberg, 2008.

[12] A. Hogan, E. Blomqvist, M. Cochez, C. D'amato, G.D. Melo, C. Gutierrez, S. Kirrane, J.E.L. Gayo, R. Navigli, S. Neumaier and et al., Knowledge Graphs, *ACM Computing Surveys* **54**(4) (2021).

[13] I. Tiddi and S. Schlobach, Knowledge graphs as tools for explainable machine learning: A survey, *Artif. Intell.* (2022).

[14] E. Prud'hommeaux and A. Seaborne, SPARQL Query Language for RDF, 2008. https://www.w3.org/TR/2008/REC-rdf-sparql-query-20080115/.

[15] S. Harris and A. Seaborne, SPARQL 1.1 Query Language, 2013. https://www.w3.org/TR/2013/REC-sparql11-query-20130321/.

[16] M. Acosta, M. Vidal, T. Lampo, J. Castillo and E. Ruckhaus, ANAPSID: An Adaptive Query Processing Engine for SPARQL Endpoints, in: *The Semantic Web – ISWC 2011*, 2011.

[17] H. Knublauch and D. Kontokostas, Shapes Constraint Language (SHACL), 2017. https://www.w3.org/TR/2017/REC-shacl-20170720/.

[18] J. Corman, J.L. Reutter and O. Savković, Semantics and Validation of Recursive SHACL, in: *The Semantic Web – ISWC 2018*, Springer, Cham, 2018.

[19] J. Corman, F. Florenzano, J.L. Reutter and O. Savković, Validating SHACL Constraints over a SPARQL Endpoint, in: *The Semantic Web – ISWC 2019*, Springer, Cham, 2019.

[20] M. Figuera, P.D. Rohde and M.-E. Vidal, Trav-SHACL: Efficiently Validating Networks of SHACL Constraints, in: *The Web Conference*, ACM, New York, NY, USA, 2021.

[21] T. Akiba, S. Sano, T. Yanase, T. Ohta and M. Koyama, Optuna: A Next-generation Hyperparameter Optimization Framework, arXiv, 2019.

[22] A. Dimou, M. Vander Sande, P. Colpaert, R. Verborgh, E. Mannens and R. Van de Walle, RML: a generic language for integrated RDF mappings of heterogeneous data, in: *7th Workshop on Linked Data on the Web*, 2014.

[23] E. Iglesias, S. Jozashoori, D. Chaves-Fraga, D. Collarana and M.-E. Vidal, SDM-RDFizer: An RML Interpreter for the Efficient Creation of RDF Knowledge Graphs, in: *CIKM '20: Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, ACM, New York, NY, USA, 2020.

[24] N. Halliwell, F. Gandon and F. Lecue, User Scored Evaluation of Non-Unique Explanations for Relational Graph Convolutional Network Link Prediction on Knowledge Graphs, in: *K-CAP '21: Proceedings of the 11th Knowledge Capture Conference*, ACM, New York, NY, USA, 2021.

[25] M. Vidal, K.M. Endris, S. Jazashoori, A. Sakor and A. Rivas, Transforming Heterogeneous Data into Knowledge for Personalized Treatments - A Use Case, *Datenbank-Spektrum* **19**(2) (2019), 95–106.

[26] D. Gries and F.B. Schneider, *A Logical Approach to Discrete Math*, Texts and Monographs in Computer Science, 1993.

[27] P.D. Rohde, DeTrusty v0.11.1, 2023. doi:10.5281/zenodo.7639170.

[28] M. Färber, F. Bartscherer, C. Menne and A. Rettinger, Linked Data Quality of DBpedia, Freebase, OpenCyc, Wikidata, and YAGO, *Semantic Web Journal* (2017).

[29] E. Iglesias, S. Jozashoori and D. Chaves-Fraga, SDM-RDFizer v4.4.2, 2022.

[30] D. Esteves, A. Ławrynowicz, P. Panov, L. Soldatova, T. Soru and J. Vanschoren, ML Schema Core Specification, 2016. http://www.w3.org/2016/10/mls/.

[31] P. Groth and M. Dumontier, Introduction - FAIR data, systems and analysis, *Data Science* **3**(1) (2020), 1–2.

[32] K.M. Endris, M. Galkin, I. Lytra, M.N. Mami, M.-E. Vidal and S. Auer, Querying Interlinked Data by Bridging RDF Molecule Templates, in: *Transactions on Large-Scale Data- and Knowledge-Centered Systems XXXIX*, Vol. 11310, Springer, Berlin, Heidelberg, 2018.
[33] P.D. Rohde, SHACL Constraint Validation during SPARQL Query Processing, in: *Proceedings of the VLDB 2021 PhD Workshop*, CEUR-WS.org, Aachen, Germany, 2021. http://ceur-ws.org/Vol-2971/paper05.pdf.
[34] P. Ristoski and H. Paulheim, Semantic Web in data mining and knowledge discovery: A comprehensive survey, *J. Web Semant.* (2016).
[35] M. Kulmanov, F.Z. Smaili, X. Gao and R. Hoehndorf, Semantic similarity and machine learning with ontologies, *Briefings in Bioinformatics* **22**(4) (2020).
[36] H. Min, F. Alemi, C.A. Hane and V.S. Nori, Improving the Accuracy of Predictive Models for Outcomes of Antidepressants by Using an Ontological Adjustment Approach, *Applied Sciences* **12**(3) (2022).
[37] P. Haug, J. Holmen, X. Wu, K. Mynam, M. Ebert and J. Ferraro, Ontology-Based Tools to Expedite Predictive Model Construction, *AMIA Summits on Translational Science Proceedings* **2014** (2014). https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4419766/pdf/1859946.pdf.
[38] T. Gebru, J. Morgenstern, B. Vecchione, J.W. Vaughan, H.M. Wallach, H.D. III and K. Crawford, Datasheets for datasets, *Commun. ACM* **64**(12) (2021), 86–92.
[39] U. Jaimini and A.P. Sheth, CausalKG: Causal Knowledge Graph Explainability Using Interventional and Counterfactual Reasoning, *IEEE Internet Comput.* **26**(1) (2022).

Appendix A. SPARQL Queries to represent Statistical Queries over the InterpretME KG

These queries can be executed over the SPARQL endpoint of the InterpretME KG[25] to gain some insights into the behavior of the trained predictive model. Listing 1 presents a SPARQL query that collects the information about the entities that define the class `intr:TargetEntity`. The projected attributes include the entity from the input KG that is classified into a specific target class and the interpretations provided by LIME which include interpreted features and their feature contributions. Based on a particular interpreted feature we can analyze the feature contribution from the highest to lowest value.

```
PREFIX intr: <http://interpretme.org/vocab/>
PREFIX owl: <http://www.w3.org/2002/07/owl#>
PREFIX prov: <http://www.w3.org/ns/prov#>
PREFIX rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#>

SELECT DISTINCT ?entity ?value
    WHERE {
        ?entity rdf:type intr:TargetEntity .
        ?entity owl:sameAs ?sourceEntity .
        ?entity intr:hasEntity ?LIMEentity .
        ?entity intr:hasInterpretedFeature  ?interpretedFeature .
        ?interpretedFeature intr:hasFeatureWeight ?featureWeight .
        ?interpretedFeature prov:hasGeneratedBy ?InterpretableTool .
        ?entity intr:hasEntityClassProbability  ?classProb .
        ?classProb intr:hasPredictionProbability ?probability .
        ?classProb intr:hasClass ?targetClass .
        ?featureWeight intr:hasFeature ?feature .
        ?featureWeight intr:hasWeight ?value .
    FILTER (?feature = <http://interpretme.org/entity/preds_1%20%3C%3D%200.00>)
} ORDER BY DESC(?value)
```

Listing 1 SPARQL Query to Retrieve target entity and the associated feature contributions generated by LIME

Listing 2 depicts a SPARQL query that retrieves the count of the target entities with respect to the target class that entity is classified into and the prediction probability generated by LIME.

Listing 3 performs a SPARQL query that performs the statistical analysis of probability, i.e., MAX, MIN, AVG about the classification of target entities in a specific target class.

[25] https://labs.tib.eu/sdm/InterpretME-wog/sparql

```
PREFIX intr: <http://interpretme.org/vocab/>
PREFIX owl: <http://www.w3.org/2002/07/owl#>
PREFIX prov: <http://www.w3.org/ns/prov#>

SELECT DISTINCT ?targetClass ?probability (count(distinct ?entity) as ?num)
    WHERE {
        ?entity rdf:type intr:TargetEntity .
        ?entity owl:sameAs ?sourceEntity .
        ?entity intr:hasEntity ?LIMEentity .
        ?entity intr:hasInterpretedFeature ?interpretedFeature .
        ?interpretedFeature intr:hasFeatureWeight ?featureWeight .
        ?interpretedFeature prov:hasGeneratedBy ?InterpretableTool .
        ?entity intr:hasEntityClassProbability ?classProb .
        ?classProb intr:hasPredictionProbability ?probability .
        ?classProb intr:hasClass ?targetClass .
        ?featureWeight intr:hasFeature ?feature .
        ?featureWeight intr:hasWeight ?value .
    } GROUP BY ?targetClass ?probability
ORDER BY DESC(?num)
```

Listing 2 SPARQL Query to retrieve target class and the probability ordered in decreasing order

```
PREFIX intr: <http://interpretme.org/vocab/>
PREFIX owl: <http://www.w3.org/2002/07/owl#>
PREFIX prov: <http://www.w3.org/ns/prov#>

SELECT DISTINCT ?targetClass (MAX(?probability) as ?max) (MIN(?probability)
   as ?min) (AVG(?probability) as?avg) (count(distinct ?entity) as ?num)
    WHERE {
        ?entity rdf:type intr:TargetEntity .
        ?entity owl:sameAs ?sourceEntity .
        ?entity intr:hasEntity ?LIMEentity .
        ?entity intr:hasInterpretedFeature ?interpretedFeature .
        ?interpretedFeature intr:hasFeatureWeight ?featureWeight .
        ?interpretedFeature prov:hasGeneratedBy ?InterpretableTool .
        ?entity intr:hasEntityClassProbability ?classProb .
        ?classProb intr:hasPredictionProbability ?probability .
        ?classProb intr:hasClass ?targetClass .
        ?featureWeight intr:hasFeature ?feature .
        ?featureWeight intr:hasWeight ?value .
} GROUP BY ?targetClass
ORDER BY DESC(?num)
```

Listing 3 SPARQL Query to retrieve the MAX and MIN and AVG of the prediction probability generated by LIME for the target entity