



Propositionalization and embeddings: two sides of the same coin

Nada Lavrač^{1,2} · Blaž Škrlič³ · Marko Robnik-Šikonja⁴

Received: 15 February 2019 / Revised: 29 May 2020 / Accepted: 8 June 2020 / Published online: 28 June 2020
© The Author(s) 2020

Abstract

Data preprocessing is an important component of machine learning pipelines, which requires ample time and resources. An integral part of preprocessing is data transformation into the format required by a given learning algorithm. This paper outlines some of the modern data processing techniques used in relational learning that enable data fusion from different input data types and formats into a single table data representation, focusing on the propositionalization and embedding data transformation approaches. While both approaches aim at transforming data into tabular data format, they use different terminology and task definitions, are perceived to address different goals, and are used in different contexts. This paper contributes a unifying framework that allows for improved understanding of these two data transformation techniques by presenting their unified definitions, and by explaining the similarities and differences between the two approaches as variants of a unified complex data transformation task. In addition to the unifying framework, the novelty of this paper is a unifying methodology combining propositionalization and embeddings, which benefits from the advantages of both in solving complex data transformation and learning tasks. We present two efficient implementations of the unifying methodology: an instance-based PropDRM approach, and a feature-based PropStar approach to data transformation and learning, together with their empirical evaluation on several relational problems. The results show that the new algorithms can outperform existing relational learners and can solve much larger problems.

Keywords Inductive logic programming · Relational learning · Propositionalization · Embeddings · Knowledge graphs

Editors: Dimitar Kazakov and Filip Železny.

✉ Blaž Škrlič
blaz.skrlic@ijs.si

Extended author information available on the last page of the article

1 Introduction

Data preprocessing for machine learning is a great challenge for a data scientist faced with large quantities of data in different forms and sizes. Most of the modern data processing techniques enable data fusion from different data types and formats into a single table data representation, which is expected by standard machine learning techniques including rule learning, decision tree learning, support vector machines (SVMs), deep neural networks (DNNs), etc. The key element of the success of modern data transformation methods is that similarities of original instances and their relations are encoded as distances in the target vector space.

Two of the most prominent data transformation approaches outlined in this paper are propositionalization and embeddings. While propositionalization (Kramer et al. 2001; Železný and Lavrač 2006) is a well known data transformation technique used in relational learning (RL) and inductive logic programming (ILP) (Muggleton 1992; Lavrač and Džeroski 1994; De Raedt 2008), embeddings (Mikolov et al. 2013; Wu et al. 2018) have only recently been recognized by RL and ILP researchers as a powerful technique for preprocessing relational and complex structured data. In the relational learning context of this paper, both approaches take as input a relational data set (e.g., a given relational database) and transform it into a single data table format, which is then used as an input to a propositional learning algorithm of choice.

The first aim of this paper is to present a unifying survey of propositionalization and embedding data transformation approaches. While both approaches aim at transforming data into a tabular data format, the approaches use different terminology and task definitions, claim to have different goals, and are used in very different contexts. This paper contributes an improved understanding of these data transformation techniques by presenting a unified terminology and definitions, by explaining the similarities and differences of the two definitions as variants of a unified complex data transformation task, by exploring the apparent differences between the two approaches, and by outlining some of their advantages and disadvantages.

In addition to the unifying survey, the main novelty of this paper is a unifying methodology that combines propositionalization and embeddings, which benefits from the advantages of both in solving complex data transformation and learning tasks. The unifying methodology resulted in two new pipelines, PropDRM and PropStar, which implement an instance-based and a feature-based approach to data transformation and learning, respectively. Both approaches are computationally efficient and can successfully solve much larger tasks than the existing relational learning approaches. We made their code publicly available.

The paper starts by motivating the need for transforming heterogeneous relational data into a tabular format in Sect. 2. Section 3 introduces the data transformation approaches in the context of information representation levels proposed by Gärdenfors (2000). Section 4 presents the related work, focusing on selected propositionalization and embeddings methods relevant to the relational learning context of this paper. Section 5 presents a unifying framework for propositionalization and embeddings, allowing for the analysis of characteristic properties of these data transformation approaches. Section 6 proposes a unifying methodology that combines propositionalization and embeddings, which benefits from the advantages of both, and presents two implementations of the proposed unifying framework: an instance-based embedding approach PropDRM based on the existing Deep Relational Machines (DRM) (Srinivasan et al. 2019; Lodhi 2013), followed by a novel feature-based

embedding approach PropStar proposed in this paper, using the StarSpace entity embedding approach (Wu et al. 2018). Experimental evaluation of the proposed implementations is presented in Sect. 7. The paper concludes by a summary and some ideas for future work in Sect. 8.

2 Motivation

Machine learning is the key enabler for computer systems to progressively improve their performance when helping humans to solve difficult problem solving tasks. Nevertheless, current machine learning approaches only come half-way in helping humans, as humans still have to formulate the problem and prepare the data in the form that is best suited to the powerful machine learning algorithms.

Most of the best performing machine learning algorithms, like Support Vector Machines (SVMs) or deep neural networks, assume numeric data and outperform symbolic approaches in terms of predictive performance, efficiency, and scalability. The dominance of numeric algorithms started in 1980s with the advent of backpropagation and neural networks (Rumelhart et al. 1986), continued in late 1990s and early 2000s with SVMs (Cortes and Vapnik 1995), and finally reached the current peak with deep neural networks (Goodfellow et al. 2016). Deep neural networks are currently considered the most powerful learners for solving many of previously unsolvable learning problems in computer vision (face recognition rivals humans' performance), game playing (a program has beaten a human champion in the game of Go), and natural language processing (successful automatic speech recognition and machine translation).

While the most powerful machine learning approaches are numeric, humans perceive and describe real-world problems mostly in symbolic terms, using various data representation format, such as graphs, relations, texts or electronic health records, all involving discrete representations. However, if we are to harness the power of successful numeric deep learning approaches for discrete learning problems, discrete data should be transformed into a form suitable for numeric learning algorithms. The viewpoint of addressing real-world problems as numeric has a rationale even for discrete domains, as many symbolic learners perform generalizations based on object similarity. For example, in graphs, nodes can represent similar entities or have connections with similar other nodes; in text, words can appear with similar contexts or play the same role in sentences; in medicine, patients may have similar symptoms or similar disease histories. Such similarities are used by numerous machine learning algorithms to generalize and learn, including classical bottom-up learning approaches such as hierarchical clustering, as well as symbolic learners adapted to top-down induction of clustering trees (Blockeel et al. 1998). If we want to exploit the power of modern machine learning algorithms, like SVMs and deep neural networks, to process the inherently discrete data, one has to transform discrete data into (numeric) vectors in such a way that similarities between objects are preserved and encoded as distances in the transformed (numeric) space.

Contemporary preprocessing approaches that prepare numeric vector data for machine learning algorithms are called *embeddings*. Nevertheless, as demonstrated in this paper, symbolic data transformations, as ancestors of the contemporary embedding approaches, remain relevant: the role of *propositionalization*, a symbolic approach to relational data transformation into feature vectors, is not only to enable contemporary machine

learning algorithms to induce better predictive models, but to allow descriptive data mining approaches to discover interesting human-comprehensible patterns in symbolic data.

As this paper demonstrates, albeit propositionalization and embeddings represent different types of data transformations, these approaches actually represent the *two sides of the same coin*. The main unifying element they have in common is that they transform the data into a vector format and encode the relations between objects in the original space as distances in the new vector space.

3 Data transformations and information representation levels

As this section will show, we consider data transformations as a particular subprocess of data preprocessing. Data preprocessing aims to handle missing attribute values, control out-of-range values and impossible attribute-value combinations, or handle noisy or unreliable data, to name just some of the types of data irregularities addressed in processing real-life data. Data preprocessing may include data cleaning, instance selection, normalization, feature engineering (feature extraction and/or feature construction), data transformation, feature selection, etc. The result of data preprocessing is the final training set, which is used as input to a machine learning algorithm.

Data preprocessing can be manual, automated, or semi-automated. We focus on automated transformations of data, present in heterogeneous types and formats, into a uniform tabular data representation. We refer to this specific automated data preprocessing task as *data transformation*, and define it as follows.

Definition 1 (*Data transformation*) *Data transformation* is a step in the data preprocessing task that automatically transforms the input data and the background knowledge into a uniform tabular representation, where each row represents a data instance, and each column represents one of the dimensions in a multi-dimensional feature space.

In the above definition, we decided to distinguish between *data* and *background knowledge*. This is an intentional decision, although it could be argued that in some settings, we could refer to both as data. Let us provide an operational distinction between data and background knowledge. *Data* is considered by the learner as the target data from which the learner should learn a model (e.g., a classifier in the case of class labeled data) or a set of descriptive patterns (e.g., a set of association rules in the case of unlabeled data). *Background knowledge* is any additional knowledge used by the learner in model or pattern construction from the target data. Simplest forms of background knowledge define hierarchies of features (attribute values), such as color *green* being more general than *light green* or *dark green*. More complex background knowledge refers to any other declarative prior domain knowledge, such as knowledge encoded in relational databases, knowledge graphs or domain specific taxonomies and ontologies, such as the Gene Ontology, in its 2020-05-02 release including 44,508 GO terms, 7,765,270 annotations, 1,464,358 gene products and 4,593 species.

This data transformation setting is applicable in various data science scenarios involving relational data mining, inductive logic programming, text mining, graph and network mining as well as tasks that require fusion of data of a variety of data types and formats and their transformation into a joint data representation formalism.

3.1 Information representation levels

As currently the most powerful machine learning (ML) algorithms take as input numeric representations, users of ML algorithms tend to transform other forms of human knowledge into the numeric representation space. Interestingly, even if this is countering a standard RL and ILP viewpoint, this is true also for symbolic representations, which are currently used to store most of the human knowledge.

The distinction between the symbolic and numeric representation space mentioned above can be further clarified in terms of the *levels of cognitive representations*, introduced by Gärdenfors (2000), i.e. the neural, spatial and symbolic representation levels. In his theory, Gärdenfors assumes that when modeling cognitive systems in terms of information processing, all three levels are connected: starting from the sensory inputs at the lowest neural representation level, resulting in spatial representations at the middle conceptual spaces level, up to symbolic representations at the level of language.

<i>Neural</i>	This representation level corresponds to the sub-conceptual connectionist level. At this level, information is represented by activation patterns in densely connected networks of primitive units. This enables concepts to be learned from the observed data by modifying the connection weights between the units.
<i>Spatial</i>	This representation level is modeled in terms of Gärdenfors' conceptual spaces. At this level, information is represented by points or regions in a conceptual space built upon some dimensions that represent geometrical, topological or ordinal properties of the observed objects. In spatial representations, the similarity between concepts is represented in terms of the distances between the points or regions in a multidimensional space, where concepts are learned by modeling the similarity between the observed objects.
<i>Symbolic</i>	At this representation level, information is represented by the language of symbols (words), where the meaning is internal to the representation itself (i.e. symbols have meaning only in terms of other symbols, while their semantics is grounded in the spatial level), and concepts are learned by symbolic generalization rules.

From the perspective of this paper, the above levels of cognitive representations introduced by Gärdenfors (2000) provide a theoretical ground to separate the learning approaches as well as the data transformation approaches into three categories based on the levels of their output representation space: neural, spacial and symbolic. However, given the scope of this paper, we do not consider *neural transformations*, and focus only on two data transformation types:

- *symbolic transformations*, in this paper referred to as *propositionalization*, denoting data transformations into a symbolic representation space, and
- *numeric transformations*, in this paper referred to as *embeddings*, denoting data transformations into a spatial representation space.

These two data transformation approaches are briefly introduced below, and further described in the related work (Sect. 4).

3.2 Transformations into symbolic representation space

The past decades of machine learning were characterized by symbolic learning, where the result of a machine learning or data mining algorithm was a predictive model of a set of patterns described in a symbolic representation language, resulting in symbolic human-understandable patterns and models. Symbolic machine learning approaches include rule learning (Michalski et al. 1986; Clark and Niblett 1989), decision tree learning (Quinlan 1986) and learning logical representations by relational learning and inductive logic programming (ILP) algorithms (Muggleton 1992; Lavrač and Džeroski 1994; De Raedt 2008).

To be able to apply a symbolic learner, the data is typically transformed into a single tabular data format, where each row represents a single data instance, and each column represents an attribute or a feature. Such transformation into symbolic vector space (i.e. a symbolic data table format) is well known in the ILP and relational learning community, where it is referred to as *propositionalization*. Propositionalization approaches are presented in Sect. 4.2.

3.3 Transformations into numeric representation space

In the last 20 years we have been witnessing increasing dominance of statistical machine learning and pattern-recognition methods, including neural network learning (Rumelhart and McClelland 1986), Support Vector Machines (SVMs) (Vapnik 1995; Schölkopf and Smola 2001), random forests (Breiman 2001), and boosting (Freund and Schapire 1997). These statistical approaches are quite different from the symbolic approaches mentioned in Sect. 3.2, however there are many approaches that cross these boundaries, including e.g., the CART decision tree learning algorithm (Breiman et al. 1984), the Bump hunting rule learning algorithm (Friedman and Fisher 1999), which are firmly based in statistics. Moreover, ensemble techniques such as boosting (Freund and Schapire 1997), bagging (Breiman 1996) or random forests (Breiman 2001) also combine the predictions of multiple logical models on a sound statistical basis (Schapire et al. 1998; Mease and Wyner 2008; Bennett et al. 2008). All these are also considered to belong to the family of statistical learning approaches.

To be able to apply a statistical learner, the data is typically transformed into a single tabular data format, where each row represents a single data instance, and each column is a numeric attribute or a numeric feature, with some predefined range of numeric values. Such transformation into numeric vector space (i.e. a numeric data table format) is well known in the deep learning community, where it is referred to as *embedding*. Approaches to embedding relational structures are presented in Sect. 4.3.

4 Related work

In this section we first outline various transformation methods in Sect. 4.1, followed by a more detailed description of the data transformation methods relevant for the context of relational learning, i.e. propositionalization and embeddings, in Sects. 4.2 and 4.3, respectively.

4.1 Outline of data transformation methods

While there are many algorithms for transforming data into a spatial representation, it is interesting that recent approaches rely on deep neural networks, thereby harnessing the neural representation level as the means to transform symbolic representations into the spatial representation. Below we list the main types of approaches that perform transformations between representations.

Community detection and graph traversal methods. Many complex data sets can be represented as graphs, where nodes represent data instances and edges represent their relations. Graphs can be homogeneous (consisting of a single type of nodes and relations) or heterogeneous (consisting of different types of nodes and relations). To encode a graph in a tabular form by preserving the information about the relations, various graph encoding techniques were developed, such as propositionalization via random walk graph traversal, representing nodes via their neighborhoods and communities (Plantié and Crampes 2013). These approaches are frequently used for data fusion in mining heterogeneous information networks. Neural network approaches (presented below) are also very competitive as means for encoding graphs.

Matrix factorization methods. When data is not explicitly presented in the form of relations but the relations between objects are implicit, given by a similarity matrix, the objects can be encoded in a numeric form using matrix factorization. As an example take Latent Semantic Analysis used in text mining, which factorizes a word similarity matrix to represent words in a vector form. Another example is factorization of graph adjacency matrices. These types of embeddings were largely superseded by deep neural networks which, instead of observing similarity between different objects, construct a prediction task and forecast similarity. For example, for text, given a word, the word-2vec embedding method (Mikolov et al. 2013) predicts words in its neighborhood.

Propositionalization methods are used to get tabular data from multi-relational databases as well as from a mixture of tabular data and background knowledge in the form of logic programs or networked data, including ontologies. These transformations were mostly developed within the Inductive Logic Programming and Relational Learning community, and are still actively researched and used. Propositionalisation methods do not perform dimensionality reduction and are most often used with data mining and symbolic machine learning algorithms. We discuss these methods in Sect. 4.2.

Neural networks based methods. In neural networks the information is represented by activation patterns in interconnected networks of primitive units. This enables that concepts are gradually learned from the observed data by modifying the connection weights between the hierarchically organized units. These weights can be extracted from neural networks and used as a spatial representation that transforms relations between entities into distances. Recently, this approach became a prevalent way to build representation for many different types of entities, e.g., texts, graphs, electronic health records, images, relations, recommendations, etc. In Sect. 4.3 we describe the data types and approaches, which are capable of embedding relational structures and are therefore most relevant for the context of this paper. These include knowledge graph embeddings (presented in Sect. 4.3.1), entity embeddings capable of forming (both supervised and unsupervised) representations based on the similarity of entities (presented in Sect. 4.3.2), and Deep Relational Machines methodology that links symbolic representations to deep neural networks (presented in Sect. 4.3.3).

Other embedding methods. Other forms of embeddings were developed by different communities that observed the need to better represent the (symbolic) data. For example, Latent Dirichlet Allocation (LDA) (Blei et al. 2003) used in text analysis learns distributions of words for different topics. These distributions can be used as an effective embedding for words, topics, and documents. Feature extraction methods form a rich representation of instances by projecting them into a high dimensional space (Lewis 1992). Another example of (implicit) transformation into high dimensional space is the kernel convolutional approach proposed by Haussler (1999), which introduces the idea that kernels can be used for discrete structures by iteratively applying convolution and kernels to smaller parts of the data structure. Convolutional kernels exist for sets, graphs, trees, strings, logical interpretations, and relations (Cumby and Roth 2003). This allows methods such as SVM or Gaussian Processes to work with relational data. Most of these embeddings are recently superseded or merged with neural networks.

All the above approaches perform data transformations from different data formats to a single table representation. However, their underlying principles are different: while factorization and neural embeddings perform dimensionality reduction, resulting in lower-dimensional feature vector representations capturing the semantics of the data, propositionalization results in a vector representation using relational features with a higher generalization potential than the features used in the original data representation. Note that there exist also other approaches to data transformation and fusion, including HINMINE (Kralj et al. 2018), metapath2vec (Zhu et al. 2018) and OhmNet (Žitnik and Leskovec 2017), which are out of the main scope of this paper.

4.2 Propositionalization

In propositionalization, relational feature construction is the most common approach to data transformation. LINUS (Lavrač et al. 1991) was one of the pioneering propositionalization approaches using automated relational feature construction. LINUS was restricted to generation of features that do not allow recursion and existential local variables, which means that the target relation cannot be many-to-many and self-referencing. The second limitation was more serious: the queries could not contain joins (conjunctions of literals). The LINUS descendant SINUS (Lavrač and Flach 2001) incorporates more advanced feature construction techniques inspired by IBC (Flach and Lachiche 1999). The LINUS approach had many followers, including relational subgroup discovery system RSD (Železný and Lavrač 2006), which is outlined also in the list of propositionalization approaches below. Alternatives to relational feature construction include the construction of aggregation queries.

In this section we first clearly define the distinction between attributes and features, followed by an outline of selected propositionalization approaches and of the specific Wordification approach used in the algorithms developed in this work.

4.2.1 Features

To be able to apply a symbolic propositional learner, the data should be represented in a single table data format, where each row represents a single data instance, and each column represents an attribute or a feature. For the sake of clarity, let us distinguish between *attributes* and *features* below.

Attributes that describe the data instances can be either numeric variables (with values like 7 or 1.5) or nominal/discrete variables (with values like *red* or *female*). In contrast to attributes, a *feature* describes the presence or absence of some property of an instance. As a result, features are always Boolean-valued (values *true* or *false*). For example, for attribute *gender* with values *female* and *male*, two separate features can be constructed: f_1 : $gender=female$ and f_2 : $gender=male$, and only one of these features is assumed to be *true* for an individual data instance. Note that features are different even from binary-valued attributes: e.g., for a binary attribute a_i with values *true* and *false*, there are two corresponding features: f_3 : $a_i = true$ and f_4 : $a_i = false$. Furthermore, features can test a value of a single attribute, like $a_j > 3$, or they can represent complex logical and numerical relations, integrating properties of multiple attributes, like f_5 : $a_k < 2 \cdot (a_j - a_i)$.

Previous feature types are referred to as *propositional features*. On the other hand, *relational features* relate the values of different attributes to each other. In the simplest case, for example, they test for the equality or inequality of the values of two attributes of the same type, such as *Length* and *Height*. More complex relational features can use the background relations, e.g., f_6 : $adjacent(NodeX, NodeY)$. Even more advanced, relational features can introduce new variables. For example, if relations are used to encode a graph, a relational feature such as f_7 : $color(CurrentNode, blue) \wedge link(CurrentNode, NewNode) \wedge color(NewNode, red)$, can introduce a new variable *NewNode* to subsequently test whether there exists a previously not visited node in the graph that is colored red.

Take a simple toy trains example learning problem illustrated in Appendix A, and two complex relational features describing trains:

$$f_8: \text{hasCar}(T,C) \wedge \text{carLength}(C,\text{short}) \wedge \text{carRoof}(C,\text{peaked})$$

$$f_9: \text{hasCar}(T,C1) \wedge \text{carLength}(C1,\text{short}) \wedge \text{hasCar}(T,C2) \wedge \text{carRoof}(C2,\text{peaked})$$

Feature f_8 is a single complex relational feature, while f_9 contains two distinct relational features. Formally, a feature is defined as a minimal set of literals such that it introduces at most one local (i.e. existential) variable in the feature set composing the relational feature.

The main point of relational features is that they localize variable sharing: this can be made explicit by naming the features:

$$f_{10}: \text{hasShortCar}(T) \leftarrow \text{hasCar}(T,C) \wedge \text{clength}(C,\text{short})$$

$$f_{11}: \text{hasPeakedroofCar}(T) \leftarrow \text{hasCar}(T,C) \wedge \text{carRoof}(C,\text{peaked})$$

The propositionalization approach to relational learning captures exactly this idea: generating complex features, such as f_8 , f_{10} and f_{11} , which will allow multi-relational data representation of properties of target instances (such as trains T) through representations of properties of their components (such as cars C). Selected propositionalization approaches, which use complex feature construction in the automated multi-relational data transformation process are outlined below.

4.2.2 Outline of selected propositionalization algorithms

Below we outline a selection of propositionalization approaches, while an interested reader can find extensive overviews of different feature construction approaches in the work of Kramer et al. (2001) and Krogel et al. (2003).

Relaggs (Krogel and Wrobel 2001) stands for *relational aggregation*. It is a propositionalization approach that takes the input relational database schema as a basis for a declarative bias, using optimization techniques usually used in relational databases (e.g.,

indexes). The approach employs aggregation functions in order to summarize non-target relations with respect to the individuals in the target table.

IBC (Flach and Lachiche 1999) strives to enable the propositional naive Bayes classifier to handle relational data. It does so by a transformation in which a set of first-order conditions is generated and then used as attributes in the naive Bayes classifier. The transformation, however, is done in a dynamic manner, as opposed to standard propositionalization, which is performed as a static step of data preprocessing. This approach is extended by IBC2 (Lachiche and Flach 2003), which allows distributions over sets, tuples, and multisets, thus enabling the naive Bayes classifier to consider also structured individuals.

Tertius (Flach and Lachiche 2001) is a top-down rule discovery system, incorporating first-order clausal logic. The main idea is that no particular prediction target is specified beforehand, hence Tertius can be seen as an ILP system that learns rules in an unsupervised manner. Its relevance for this survey lies in the fact that Tertius encompasses IBC, i.e. relational data is handled through IBC transformation.

RSD (Železný and Lavrač 2006) is a relational subgroup discovery algorithm composed of two main steps: the propositionalization step and the (optional) subgroup discovery step. The output of the propositionalization step can be used also as input to other propositional learners. RSD effectively produces an exhaustive list of first-order features that comply with the user-defined mode constraints, similar to those of Progol (Muggleton 1995) and Aleph (Srinivasan 2007). Furthermore, RSD features satisfy the connectivity requirement, which imposes that no feature can be decomposed into a conjunction of two or more features. Mode declarations define the algorithm's syntactic bias, i.e. the space of possible features.

HiFi (Kuželka and Železný 2008) is a propositionalization approach that constructs first-order features with hierarchical structure. Due to this feature property, the algorithm performs the transformation in polynomial time of the maximum feature length. Furthermore, the resulting features are the shortest in their semantic equivalence class. The algorithm is shown to perform several orders of magnitude faster than RSD for higher feature lengths.

RelF (Kuželka and Železný 2011) is the most relevant of the algorithms in the Tree-Liker software (Kuželka and Železný 2011). It constructs a set of tree-like relational features by combining smaller conjunctive blocks. RelF preserves the monotonicity of feature reducibility and redundancy (instead of the typical monotonicity of frequency), which allows the algorithm to scale far better than other state-of-the-art propositionalization algorithms.

Cardinalization (Ahmed et al. 2015) is specifically designed to enable more than just categorical attributes in propositionalization. Specifically, it can handle a threshold on numeric attribute values and a threshold on the number of objects satisfying the condition on the attribute simultaneously. Cardinalization can be seen as an implicit form of discretization. While in discretization one sets a threshold on a numeric attribute and see how many objects satisfy the threshold later, and the cardinality follows implicitly from the attribute value threshold; on the other hand, in cardinalization, we set a threshold on the cardinality, and let an attribute-value learner decide where the threshold value on the numerical attribute should lie. Hence, Cardinalization allows for context-aware discretization. Quantiles (Ahmed et al. 2015) is a variant of Cardinalization. Instead of choosing an absolute number as cardinality threshold, Quantiles uses a relative number. CARAF (Charnay et al. 2015) approaches the problem of large relational feature search space by aggregating base features into complex compounds, which makes CARAF

similar to Relaggs. Complex aggregates run the risk of overfitting. While Relaggs tackles this problem by restricting itself to relatively simple aggregates, the distinguishing feature of CARAF is that instead it incorporates more complex aggregates into a random forest, which ameliorates the overfitting effect.

Aleph (Srinivasan 2007) is the most popular ILP algorithm and is actually an ILP toolkit with many modes of functionality: learning of theories, feature construction, incremental learning, etc. Aleph uses mode declarations to define the syntactic bias. Input relations are Prolog clauses, defined either extensionally or intensionally. Aleph's feature construction functionality also means it is a propositionalization approach.

Wordification (Perovšek et al. 2013, 2015) is a propositionalization method inspired by text mining that can be viewed as a transformation of a relational database into a corpus of text documents. The distinguishing property of Wordification is its efficiency when used on large relational data sets and the potential for using text mining approaches on the transformed propositional data. While most of the outlined propositionalization algorithms construct complex relational features including variables in the arguments of relational features, Wordification constructs simple, easily interpretable features that are treated as 'words' in the transformed Bag-Of-Words representation. It constructs features of the kind $a_i = v_{ij}$ (formulated as $a_{i_v_{ij}}$). In addition to such simple features, it constructs also conjuncts (of size 2) of such features, e.g., $a_i = v_{ij} \wedge a_k = v_{kl}$, formulated as $a_{i_v_{ij}_a_k_v_{kl}}$. To avoid confusion in case the same attribute name appeared in several tables, the actual form of features is $t_a_{i_v_{ij}}$ including the indicator of the name of table t in which attribute a_i appears. For a simple example of how such features are generated, the reader is referred to Appendix A.

4.2.3 Wordification

Given that in a previous experimental evaluation of propositionalization algorithms (Perovšek et al. 2013, 2015) the Wordification algorithm was shown to be the most effective, we selected Wordification as the propositionalization algorithm of choice in the proposed implementations combining propositionalization and embeddings in Sect. 6, where the Wordification algorithm was adapted to handle large data sets.

In the Wordification implementation, described in detail in Sect. 6.2.1, the original feature representation *TableName_AttributeName_AttributeValue* was—for implementational convenience—replaced by a tuple representation $(t.name, c, v)$, where $t.name$ refers to a table name, c to a given colon (attribute) in the table t , and v to a given value v of attribute c . Such features will be referred to as *features* or as *relational items* in the algorithm description, as appropriate.

Using this feature representation, Wordification of a multi-relational database can be summarized as the following operation:

$$DB_i = \bigsqcup_{t \in \mathcal{T}} \text{WORDIFY}(t(m(i)))$$

where m maps a given table t 's indices to target (initial) table indices (i) and \mathcal{T} is the set of all tables from which a foreign key path exists to the target table. The \bigsqcup operator represents a disjoint union of multisets (sum), yielding a single multiset (duplicates are allowed).

Foreign keys are designated columns that link data between distinct tables. Value of a foreign key in a given table is referred to as the instance id (the row is uniquely determined by this value). Let C represent the set of all columns that are not foreign keys, ids or target

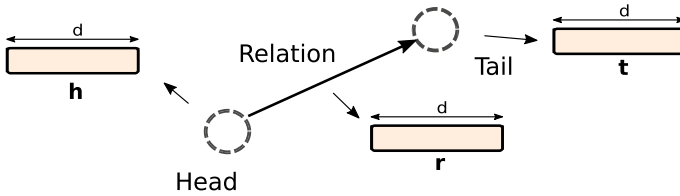


Fig. 1 Schematic representation of knowledge graph embedding. Head-Relation-Tail (h, r, t) triplets are used as inputs. Triplets are embedded in a common d -dimensional vector space

classes. The WORDIFY method returns a multiset (a bag) of relational items (for the i -th instance) constructed as follows:

$$\text{WORDIFY}(t(m(i))) = \bigcup_{v \in t(m(i))[c \in C]} (t.name, c, v)$$

where $t[c]$ represents the values v of table t in column c , and $t.name$ is the name of table t . Thus, Wordification is naïve in the sense that it simply concatenates attribute values across tables by maintaining the column and table name information in constructing features. The original implementation, however, can become spatially intractable (see (Perovšek et al. 2013), proof of complexity) as its spatial complexity is $\mathcal{O}(\text{row} \cdot \text{tables} \cdot 2^{|\text{col}|})$. Details of a more efficient implementation of Wordification are available in Sect. 6.2.1.

4.3 Embedding relational structures

In this section, we discuss methodologies capable of embedding relational structures. We start with an introduction to knowledge graph embeddings, an emerging group of methods that operate on large, real-world, annotated graphs, in Sect. 4.3.1. We proceed by the presentation of entity embeddings, a more general methodology capable of supervised, as well as unsupervised embeddings of many entities, including texts and knowledge graphs in Sect. 4.3.2. Finally, in Sect. 4.3.3, we present Deep Relational Machines, an emerging methodology that links symbolic representations to deep neural networks.

4.3.1 Knowledge graph embeddings

In knowledge graphs (KG), edges correspond to relations between entities (nodes) and the graphs present Subject-Predicate-Object *triplets*. The KG handling algorithms attempt to solve the problems like triplet completion, relation extraction, and entity resolution. The KG embedding algorithms, briefly discussed below, outline some of the key ideas which render these methods highly scalable and useful for large, semantics-rich graphs. For detailed description and a recent, extensive overview of the field, we refer the reader to Wang et al. (2017), from where we next summarize some of the key ideas underlying knowledge graph embedding.

In the below description of KG embedding algorithms, the Subject-Predicate-Object triplet notation is replaced by the (h, r, t) triplet notation, where h is referred to as the *head* of a triplet, t as the *tail*, and r as the *relation* connecting the head and the tail. A schematic representation of triplet embedding is shown in Fig. 1. The embedding

methods briefly outlined below optimize the total plausibility of the input set of triplets, where plausibility of a single triplet is denoted with $f_r(h, t)$.

- The first group of KG embedding algorithms are termed *translational distance models*, as they exploit distance-based scoring functions. They measure the plausibility of a fact as the distance between the two entities, usually after a translation carried out by the relation. One of the representative methods for this type of embedding is transE (Bordes et al. 2013), where the cost function being optimized can be stated as:

$$f_r(h, t) = -\|\mathbf{h} + \mathbf{r} - \mathbf{t}\|^2.$$

For vectors \mathbf{h} , \mathbf{r} , and \mathbf{t} in the obtained embedding, score $f_r(h, t)$ is high if triplet (h, r, t) is present in the data.

- The second group of KG embedding algorithms is not deterministic, as it takes into account the uncertainty of observing a given triplet. A representative method for this type of embeddings is KG2E (He et al. 2015), which models the triplets with multivariate Gaussians. It models individual entities, as well as relations as vectors, drawn from multivariate Gaussians, assuming that \mathbf{h} , \mathbf{r} and \mathbf{t} vectors are normally distributed, with mean vectors $\mu_h, \mu_r, \mu_t \in \mathbb{R}^d$ and covariance matrices $\Sigma_h, \Sigma_r, \Sigma_t \in \mathbb{R}^{d \times d}$, respectively. KG2E uses Kullback-Liebler divergence to directly compare the distributions as follows:

$$\begin{aligned} f_r(h, t) &= \text{KL}(\mathcal{N}(\mu_t - \mu_h), \mathcal{N}(\mu_r)) \\ &= \int \mathcal{N}_x(\mu_t - \mu_h, \Sigma_t + \Sigma_h) \ln \frac{\mathcal{N}_x(\mu_t - \mu_h, \Sigma_t + \Sigma_h)}{\mathcal{N}_x(\mu_r, \Sigma_r)} dx, \end{aligned}$$

where \mathcal{N}_x denotes the probability density function of the normal distribution.

- Semantic matching models exploit similarity-based scoring functions. They measure plausibility of facts by matching latent semantics of entities and relations embodied in their vector space representations. One of the representative algorithms for learning by semantic matching is RESCAL (Nickel et al. 2011). RESCAL optimizes the following expression:

$$f_r(h, t) = \mathbf{r}^T \cdot M_r \cdot \mathbf{t},$$

where \mathbf{h} and \mathbf{t} are representations of entities, and $M_r \in \mathbb{R}^{d \times d}$ is a matrix associated with relations.

- Matching using neural networks. Deep neural networks model triplets via training of neural network architectures. One of the first approaches was Semantic Matching Energy (SME) (Bordes et al. 2014). This method first projects entities and their relations to their corresponding vector embeddings. The relation's representation is next combined with the relation's head and tail entities to obtain $g_1(\mathbf{h}, \mathbf{r})$ and $g_2(\mathbf{t}, \mathbf{r})$ entity-relation representations in the hidden layer. Finally, a dot product is used to score the triplet relation matching

$$f_r(h, t) = g_1(\mathbf{h}, \mathbf{r})^T \cdot g_2(\mathbf{t}, \mathbf{r}).$$

The simplest version of SME defines the g_1 and g_2 as:

$$g_1(\mathbf{h}, \mathbf{r}) = W_1^{(1)} \cdot \mathbf{h} + W_1^{(2)} \cdot \mathbf{r} + b_1$$

$$g_2(\mathbf{t}, \mathbf{r}) = W_2^{(1)} \cdot \mathbf{t} + W_2^{(2)} \cdot \mathbf{r} + b_2.$$

Here, $W_1^{(1)}$, $W_1^{(2)}$, $W_2^{(1)}$ and $W_2^{(2)}$ are $\mathbb{R}^{d \times d}$ dimensional weight matrices and b_1 and b_2 are bias vectors.

Recent advances in embeddings of knowledge graphs show interesting research directions. For example, hyperbolic geometry could be used to better capture latent hierarchies, commonly present in real-world graphs (Nickel and Kiela 2017). Further, KG embedding methods are increasingly tested on large, multi-topic data collections, for example, the Linked Data (LD) which standardize and fuse data from different resources. Knowledge graph embeddings, such as RDF2vec (Ristoski and Paulheim 2016) attempt to exploit vast amounts of information in LD and transform it into a learning-suitable format. As knowledge graphs are not necessarily the only source of available information, algorithms exploit also other information, e.g., textual information available for each triplet (Wang et al. 2014). Recent trends in knowledge graph embeddings also explore how symbolic, logical structures could be used during embedding construction. Approaches such as KALE (Guo et al. 2016) construct embeddings by taking into account logical rules (e.g., Horn clauses) related to the knowledge graph, thus increasing the quality of embeddings. Similar work was proposed by Rocktäschel et al. (2015), where pairs of embeddings were considered during optimization. The same group also showed how relations can be modeled without grounding the head and tail entities for simple implication-like clauses (Demeester et al. 2016). Wang et al. (2015) demonstrated that logical rules can aid in knowledge graph completion on large knowledge bases. They showed that inclusion of rules can reduce the solution space and significantly improve the inference accuracy of embedding models.

4.3.2 Entity embedding with the StarSpace approach

The guiding principle behind all embeddings, described in the previous section, is the persistence of similarity, i.e. that entities which are similar in the knowledge graph must be represented by vectors that are similar in the embedding space. A general approach implementing this principle is to use any similarity function between entities to form a prediction task for a neural network. Below we describe a successful example of this approach, called StarSpace (Wu et al. 2018). As this approach assumes discrete features from a fixed dictionary, it is particularly appealing to relational learning and inductive logic programming.

The idea of StarSpace is to form a prediction task where a neural network is trained to predict the similarity between an entity and its related entity (e.g., its label or some other entity). The resulting neural network can be used for several purposes: directly in classification, to rank instances by their similarity, or weights of the trained network can be used as pretrained embeddings.

In StarSpace, each entity has to be described by a set of discrete features from a fixed-length dictionary and forms a so called Bag-Of-Features. This representation is general enough to cover texts (documents or sentences can be described by bags-of-words or bags-of-n-grams), users (described by bags of documents, movies, or items they like), relations and links in graphs (described by semantic triples), etc. During training, entities of different kinds are embedded *in the same* latent space, suitable for various down-stream learning tasks, e.g., a user can be compared with the recommended items. Note that entities can be embedded along with target classes, resulting in *supervised embedding learning*. This type

of representation learning is the key element of the proposed PropStar algorithm outlined in Sect. 6.1.2 and presented in detail in Sect. 6.2.3.

The StarSpace approach trains a neural network model to predict which pairs of entities are similar and which are dissimilar. Two kinds of training instances are formed, positive $(a, b) \in E^+$, which are task dependent and contain correct relations between entities (e.g., document a with its correct label b), and negative instances $(a, b_1^-), \dots, (a, b_k^-) \in E_a^-$. For each entity a (e.g., a document) appearing in the positive instances, negative instances are formed using k -negative sampling from labels $\{b_i^-\}_{i=1}^k$ as in word2vec (Mikolov et al. 2013). In each batch, the neural network tries to minimize the loss function L , defined as follows:

$$L = \sum_{(a,b) \in E^+} \left(\text{Loss}(\text{sim}(a, b)) + \frac{1}{k} \sum_{\substack{i=1 \\ (a, b_i^-) \in E_a^-}}^k \text{Loss}(\text{sim}(a, b_i^-)) \right)$$

For each batch update in the training of neural network, k negative examples (a parameter) are formed by randomly sampling labels b_i^- from within the set of entities that can appear in b . For example, in the document classification task, document a has its correct label b , while k negative instances have their labels b_i^- sampled from the set of all possible labels. Similarity function sim represents the similarity between the vector representations of the two entities; typically a dot product similarity is used. Within one batch, loss function Loss sums the losses of the positive instance (a, b) and the average of the k negative instances $(a, b_i^-), i \in 1 \dots k$. To assess the loss, margin ranking loss is used, $\text{Loss} = \max(0, m - \text{sim}(a, b'))$, where m is the margin parameter, i.e. the similarity threshold, and b' is a label.

The trained network can be used for several purposes. To classify a new instance a , one iterates over all possible labels b' and chooses $\arg \max_{b'} \text{sim}(a, b')$ as the prediction. For ranking, entities can be sorted by their predicted similarity score. The embedding vectors can also be extracted and used for some other downstream task. Wu et al. (2018) recommend that the similarity function $\text{sim}(\cdot, \cdot)$ is shaped in such a way that it will directly fit the intended application, so that training will be more effective.

A few examples of tasks successfully tackled with the StarSpace feature transformation approach are described below.

- In *multiclass text classification* the positive instances (a, b) are taken from the training set of documents E^+ , represented with bags-of-words and their labels b . For negative instances, entities b_i^- are sampled from the set of possible labels.
- In *recommender systems* users are described with a bag of items they liked (or bought). The positive instances use a single user ID as a and one of the items that user liked as b . Negative instances take b_i^- from the set of possible items. Alternatively, to work for new users, the a part of user representation is composed of all the items that user liked, except one, which is used as b .
- For *link prediction* the concepts in a graph are represented as triples head-relation-tail (h, r, t) , e.g., gene-generates-protein. A positive instance a consists either of h and r , while b consists of t ; alternatively, a consists of h , and b consists of r and t . Negative instances b_i^- are sampled from the set of possible concepts. The trained network can then predicted links, e.g., gene-generates-what.

Fig. 2 Example input data for a deep relational machine that operates on the instance level

Instance	$f_1 \wedge f_2$	$f_3 \wedge f_2$	$f_1 \wedge f_3$	$f_5 \wedge f_2$	$f_4 \wedge f_1 \wedge f_5$	Class
1	[1	1	1	1	0]	+
2	[0	1	0	0	1]	+
3	[0	0	1	0	0]	-
4	[0	1	0	0	1]	-
5	[1	0	0	0	1]	-

- For *sentence embedding* in an unsupervised fashion, a collection of documents, containing sentences, is turned into a training set. For positive instances, a and b are sentences from the same document (or are close together in a document), while for negative instances, sentences b_i^- are coming from different documents. This definition of a task tries to capture the semantic similarity between sentences in a document.

In the PropStar algorithm proposed in this work, we use StarSpace similarly to the first case mentioned above (multiclass text classification). Namely, Wordification returns a bag of features (relational items) for each instance in the target table. The embeddings are learned for each feature separately, and class labels are also embedded in the same space. During classification, representations of relational items associated with a given instance (bag of features) are averaged to obtain the representation of the instance—a similar idea as in the document representation adopted in the highly efficient doc2vec branch of algorithms aimed at document classification (Le and Mikolov 2014). The embedded instances, now located in the same vector space as the embeddings of class labels, are directly used for classification. The label, closest to the representation of a given target instance is selected as the final prediction.

4.3.3 Deep relational machines

Deep neural networks are effective learners in numeric space, capable of constructing intermediate knowledge constructs and thereby improve semantics of baseline input representation. Training deep neural networks on propositionalized relational data were explored by Srinivasan et al. (2019), following the work of Lodhi (2013), where Deep Relational Machines (DRMs) were first introduced. In Lodhi's work, the DRMs used bodies of first order Horn clauses as input to restricted Boltzmann machines, where conjuncts of bonds and other molecular structure information compose individual complex features; when all structural properties are present in a given instance, the target's value is true, and false otherwise. For example, consider the following propositional representation of five instances (rows), where complex features are comprised of conjuncts of atoms f_i , as illustrated in Fig. 2.

Note that the propositionalized data set P is usually a sparse matrix, which can represent additional challenge for neural networks. The DRMs proposed by Lodhi (2013) were used for prediction of protein folding properties, as well as mutagenicity assessment of small molecules. This approach used feature selection with information theoretic measures such as information gain as the sparse matrix resulting from the propositionalization was not

suitable as an input to the neural network. The initial studies regarding DRMs explored how deep neural networks could be used as an extension of relational learning.

Recently, promising results were demonstrated in the domain of molecule classification (Dash et al. 2018) using ILP learner Aleph in its propositionalization mode for feature construction. After obtaining propositional representation of data, the obtained data table was fed into a neural network that associated such representations with the output space (e.g., a molecule's activity). Again, sparsity and size of the propositionalized representation is a problem for deep neural networks. Again, stochastic feature selection of relational features that are used as input to deep relational machines can improve the performance and interpretability (Dash et al. 2019).

The work of Srinivasan et al. (2019) is relevant for the interpretability of deep relational machines, proposing a logical approximation of well-known prediction explanation method LIME (Ribeiro et al. 2016) and showing how it can be efficiently computed.

In summary, DRMs address the following issues at the intersection of deep learning and relational learning:

- DRMs demonstrated that deep learning on propositionalized relational structures is a sensible approach to relational learning.
- Their input is comprised of logical conjuncts, offering the opportunity to obtain human-understandable explanations.
- DRMs were successfully employed for classification and regression.
- Emerging ideas in the area of representation learning have only recently been explored in the ILP context (Dumančić et al. 2018), indicating there are many possible improvements both in terms of execution speed, as well as more informative feature construction on the symbolic side of computation.

We further discuss DRMs in the context of efficiency of their implementation in Sects. 6.1.1 and 6.2.2. Development of DRMs that are efficient with respect to both space and time is an ongoing research effort. Building on the ideas of DRMs, we implemented a variant of this approach, capable of learning directly from large, sparse matrices that are returned from Wordification of a given relational database, rather than using feature selection or the output of Aleph's feature construction approach. Our novel, efficient DRM implementation is presented in Sect. 6.2.2.

5 Unifying framework for propositionalization and embeddings

The connection we made between different information representation levels and different transformation techniques shows that propositionalization and embeddings are two sides of the same coin. If we view embeddings as transformations for texts, graphs, recommendations, electronic health records, and other entities with defined similarity function, we can conclude that all these transformation present a multifaceted approach to feature construction.

To this end, the paper contributes a novel understanding of these data transformation techniques. In Sect. 5.1, we first present a unified terminology and definitions, and explain the apparent differences between the definitions of propositionalization and embeddings as variants of a complex data transformation task. In further sections we explore the apparent differences between the two approaches. In Sects. 5.2, 5.3, and 5.4 we discuss

Table 1 Unifying and differentiating aspects of propositionalization and embeddings in terms of data representation

Representation	Propositionalization	Embeddings
Vector space	Symbolic	Numeric
Features/variables	Symbolic	Numeric
Feature values	Boolean (0 or 1)	Numeric
Sparsity	Sparse	Dense
Space complexity	Space consuming	Mostly efficient
Interpretability	Interpretable features	Non-interpretable

differences in data representation, learning, and use. Finally, in Sect. 5.5 we summarize strengths and limitations of propositionalization and embeddings.

5.1 Unifying definitions

Below we present a unified view on the definitions of propositionalization and embedding tasks, as instances of a general data transformation task defined in Sect. 1 via Definition 1.

Definition 2 (*Propositionalization*)

Given: Input data of a given data type and format, and heterogeneous background knowledge of various data types and formats.

Find: A tabular representation of the data enriched with the background knowledge, where each row represents a single data instance, and each column represents a feature in a d -dimensional symbolic¹ vector space F^d .

Definition 3 (*Embedding*)

Given: Input data of a given data type and format, and heterogeneous background knowledge of various data types and formats.

Find: A tabular representation of the data enriched with the background knowledge, where each row represents a single data instance, and each column represents one of the dimensions in the d -dimensional numeric vector space \mathbb{R}^d .

5.2 Unifying propositionalization and embeddings in terms of data representation

Both data transformation techniques result in a vector space representation. The unifying dimensions of propositionalization and embeddings in terms of data representation, which are summarized in Table 1, are explained below.

In propositionalization, the transformation results in a binary matrix of sparse binary vectors, where rows corresponds to training instances and columns correspond

¹ In the case of binary valued features, each value in each column is $\in \{0, 1\}$.

Table 2 Unifying and differentiating aspects of propositionalization and embeddings in terms of learning context

Learning	Propositionalization	Embeddings
Meaning capturing	Via symbols	Via distances
Search strategy	Heuristic search	Greedy
Search goal	Global optimum	Local optimum
Typical algorithms	Symbolic, linear regression, SVM	Deep neural networks
Parameters	Few	Many
Hardware	CPU	CPU/GPU

to symbolic features constructed by a particular propositionalization algorithm. These features are human interpretable, as they are either simple logical features (such as attribute values), conjunctions of such features, relations among simple features (such as e.g., a test for the equality or inequality of values of two attributes of the same type), or relations among entities (such as links among nodes in a graph). Given that the number of constructed features is usually large, such transformation results in a sparse binary matrix with few non-zero elements.

Embeddings output is usually a dense matrix of a user-defined dimensionality, composed of vectors of numeric values, one for each object of interest. For neural network based embeddings, vectors usually represent the activation of neural network nodes of one or more levels of a deep neural network. Given a relatively low dimensionality of these vectors (from 100 to 1000) this dense representation is efficient in terms of space. However, the features/dimensions are non-interpretable, therefore a separate explanation mechanisms and visualizations are required.

5.3 Unifying propositionalization and embeddings in terms of learning

For both data transformation techniques, the resulting vector space representation is used as an input to a learning algorithm of the user's choice. The unifying dimensions of propositionalization and embeddings in terms of most frequently used learners (summarized in Table 2) are explained below.

After propositionalization, any learner capable of processing symbolic features can be used. Typical learners include rule learning, decision tree learning, random forests for a supervised setting, or association rules and symbolic clustering algorithms applied in a non-supervised learning setting. Learners usually use heuristic search to find a global optimum in terms of heuristics to be optimized (exceptions being, e.g., association rule learners using exhaustive search with constraints). Typical algorithms are decision tree learners, rule learners, linear regression and SVMs. Learners require some parameter tuning to achieve optimal results, but parameters are relatively few. Learning is typically performed on CPUs.

The embedded vectors are best suited for distance-based learners, such as neural networks, and to a lesser degree for kernel methods or logistic regression. Deep neural networks use greedy search to find locally optimal solutions, and are usually trained on GPUs, but can be used for prediction on both CPUs or GPUs. As a weakness, deep learning algorithms require substantial (hyper)parameter tuning.

Table 3 Unifying and differentiating aspects of propositionalization and embeddings in terms of use

Use	Propositionalization	Embeddings
Problems/context	Relational	Tabular, texts, graphs
Data type fusion	Enabled	Enabled
Explanation	Directly interpretable	Special approaches

5.4 Unifying propositionalization and embeddings in terms of use

The unifying dimensions of propositionalization and embeddings in terms of their use (summarized in Table 3) are explained below.

Propositionalization (Kramer et al. 2001) is one of the established methodologies used in relational learning (Džeroski and Lavrač 2001; De Raedt 2008) and ILP (Muggleton 1992; Lavrač and Džeroski 1994; De Raedt 2008) (see the propositionalization methods outlined in Sect. 4.2). The propositionalization approach was applied also in the semantic data mining where ontologies are used as a background knowledge in relational learning (Podpečan et al. 2011; Lavrač et al. 2009; Vavpetič and Lavrač 2011).

The embedding technologies are mostly used in the context of deep learning for various data formats, including tabular data, texts, images, and graphs (including knowledge graphs). In addition to knowledge graph embedding approaches (see Sect. 4.3.1), we outline some other approaches to graph embeddings below.

The first studies of graph embeddings were influenced by embedding construction from textual data. For example, the well known skip-gram model, initially used as part of word2vec (Mikolov et al. 2013) was successfully applied to learn node representations. DeepWalk (Perozzi et al. 2014) was one of the first learners that treats short random walks in graphs as sentences (or short phrases) to learn latent node embeddings. DeepWalk was revisited as node2vec (Grover and Leskovec 2016) to take into account different types of random walks, parameterized by breadth, as well as depth-first search. LINE (Tang et al. 2015b) performs similarly well for the tasks of classification and link prediction by attempting to optimize both local, as well as global network structure.

As for fusing heterogeneous data types, a propositionalization approach was proposed as a mechanism for heterogeneous data fusion (Grčar et al. 2013). As for data type fusion using embedding-based methods, PTE (Tang et al. 2015a) exploits heterogeneous networks of texts for supervised embedding construction. NetMF (Qiu et al. 2018) is a generalization of Deepwalk, node2vec, LINE and PTE, re-formulating them as a matrix factorization problem. Furthermore, struc2vec (Ribeiro et al. 2017) builds on two main ideas: representations of two nodes must be close if the two nodes are structurally similar, and the latent node representation should not depend on any node or edge attribute, including the node labels. Examples of approaches to heterogeneous graph embeddings include HIN-MINE (Kralj et al. 2018), metapath2vec (Zhu et al. 2018) and OhmNet (Žitnik and Leskovec 2017), an extension of node2vec to a heterogeneous biological setting. Heterogeneous data embeddings (Chang et al. 2015) of images, videos and text were also formulated as a task of heterogeneous graph embedding.

Concerning the interpretability of results, propositionalization approaches are mostly used with symbolic learners whose results can be interpretable, given the interpretability of features used in the transformed data description. For embedding-based methods, given the non-interpretable numeric features/dimensions, specific mechanisms need to be implemented to ensure results explanation (Robnik-Šikonja and Kononenko 2008; Štrumbelj and

Kononenko 2014). A recent well-known approach, which can be used in a post-processing phase of an arbitrary prediction model, is named SHAP (Lundberg and Lee 2017). In this approach, Shapley values offer insights into instance-level predictions by assigning fair credit to individual features for participation in prediction-explaining interactions. Explanation methods such as SHAP are commonly used to understand and debug black-box models. We refer the reader to Lundberg and Lee (2017) for a detailed overview of the method.

5.5 Summary of strengths and limitations of propositionalization and embeddings

Let us summarize the unified presentation of propositionalization and embeddings by presenting the strengths and weaknesses of the two approaches. The main strength of propositionalization is the interpretability of the constructed features, while the main strength of embeddings is high performance of classifiers learned from embeddings due to their compact representation in a vector space.

In terms of their strengths, both approaches to data transformation are: (a) automated, (b) fast, (c) semantic similarity of instances is preserved in the transformed instance space (as a remark, due to a more compact representation, embeddings preserve semantic similarity of features even better than propositionalization), (d) transformed data can be used as input to standard propositional learners, as well as to contemporary approaches.

In addition to these characteristics, embeddings have other favorable properties: (a) embedded vectors representations allow for transfer learning, e.g., for cross-lingual applications in text mining or image classification from different types of images, (b) cover a very wide range of data types (text, relations, graphs, images, time series), and (c) have a very wide community of developers and users, including industry.

In terms of their limitations when used in a multi-relational setting, both approaches to data transformation: (a) are limited to 1-many relationships (cannot handle many-to-many relationships between the connected data tables), (b) cannot handle recursion, and (c) cannot be used for predicate invention.

In addition to these characteristics, limitations of propositionalization include: (a) only boolean values are used in the transformed vector space, (b) generated sparse vectors can be memory inefficient, (c) limited range of data types are handled (relations, graphs), and (d) a small community of developers and users (mainly from ILP).

Embeddings also have several limitations: (a) loss of explainability of features and consequently of the models trained on the embedded representations, (b) many user-defined hyper-parameters, (c) high memory consumption due to many weights in neural networks, and (d) requirement for specialized hardware (GPU) for efficient training of embeddings, which may be out of reach for many researchers.

6 Proposed unification methodology and its two implementations

The unifying aspects analyzed in Sect. 5 can be used as a basis for a unifying methodology that combines propositionalization and embeddings, and benefits from the advantages of both. The propositionalization successfully captures relational information through complex relational feature construction, but results in a sparse symbolic feature vector representation. This weakness can be successfully overcome by embedding the constructed feature

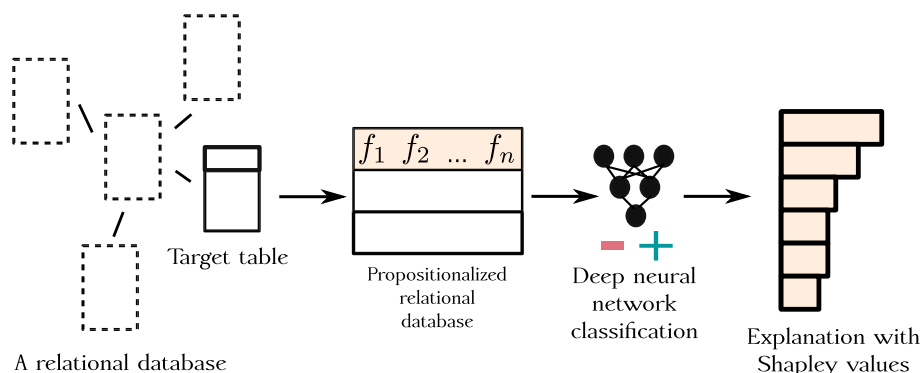


Fig. 3 Overview of the PropDRM instance-based embedding methodology, based on DRMs. Note that features in the propositionalized relational database represent either single features f_i or conjuncts of features, e.g., $f_i \wedge f_j$, given that Wordifications constructs both feature forms. For simplicity, the propositionalized database shows only two instances

vectors into a lower dimensional numeric vector space, resulting in a condensed numeric feature vector representation appropriate for use by modern deep learning algorithms.

To this end, we describe two novel data transformation algorithms, combining propositionalization and embedding based transformations into a joint data transformation framework. We first briefly outline the two approaches in Sect. 6.1, followed by their detailed descriptions in Sect. 6.2.

6.1 Outline of proposed data transformation and learning methods

We first overview the proposed unifying data transformation approaches. The first, named PropDRM, is an instance-based data transformation approach. The second one is a feature-based data transformation pipeline, called PropStar. The approaches are outlined in the next two subsections.

6.1.1 PropDRM: an instance-based approach

The first unifying approach for embedding of multi-relational databases is based on Deep Relational Machines (Dash et al. 2018) (DRMs), presented in Sect. 4.3.3. Rather than using the output of Aleph’s feature construction approach, as was the case in the DRM implementation of Dash et al. (2018), we implemented a variant of this approach, capable of learning directly from large, sparse matrices that are returned by the Wordification (Perovšek et al. 2015) approach to propositionalization of relational databases. In this work, following the paradigm of propositionalization by Wordification, each instance is described by a bag (a multiset that allows for multiple appearances of its elements) of features of the form *TableName_AttributeName_Value*. Wordification treats these simple easily interpretable features as ‘words’ in the transformed Bag-Of-Words representation. In this work, they represent individual ‘relational items’ and we use the notation (table.name, column.name, value).

Relational representations are thus obtained for individual instances, resulting in embeddings of instances (e.g., molecules, persons, companies etc). Batches of instances are then

fed to a neural network, which performs the desired down-stream task, such as classification or regression. Schematically, the approach is illustrated in Fig. 3.²

Note that although propositionalization and subsequent learning are conceptually two distinct steps, they are not necessarily separated when implemented in practice: as neural networks operate with small batches of input data, if propositionalization is capable of similar batch functionality, relational features can be generated in a lazy manner when needed by the neural network. The technical details of the proposed PropDRM implementation are presented in Sect. 6.2.2.

When compared to our PropStar algorithm presented in Sects. 6.1.2 and 6.2.3 below, the key difference of the outlined DRM-based implementation of the unifying methodology is the type of embeddings: PropDRM embeds instances (i.e. whole bags of constructed features), whereas PropStar embeds features along with the class values in the same vector space.

6.1.2 PropStar: a feature-based approach

In this section, we outline the proposed PropStar algorithm for classification via feature embedding. Its details and implementation are presented in Sect. 6.2.3. Unlike the PropDRM algorithm, where each embedding vector represents a single data *instance*, the idea of PropStar is to use embedding vectors to represent the *features* of the data set. Here, individual relational features, obtained as the result of propositionalization by Wordification, are used by a supervised embeddings learner to obtain representations, co-located with instance labels. This approach is conceptually different in the sense that representations are not learned for individual instances (as is the case of DRMs); instead, they are learned for every single relational feature that is the output of the selected propositionalization algorithm (i.e. Wordification).

The fact that PropStar produces vector representations of features means that the labels (label=true and label=false) are also represented by vectors in the same dense space as the other vectors. This leads to an intuitive direct classification of new examples. We can observe the set of vectors representing the relational items present in the itemset representing the new example. To classify a new instance, the embeddings of the set of its features (i.e. true values) are averaged and the result is compared to the embedding of class labels. The nearest class label is chosen as the predicted value.

Figure 4 illustrates how new instances are classified by direct comparison of the representations of their features in the latent dense semantics-preserving space that also contains the information on labels. The classification is based on the proximity to a given label (in the latent space). If the center of feature vectors of a given instance is closer to the vector representing the feature label=true, then the example is classified as positive.

In contrast to the instance-based embeddings discussed in Sect. 6.1.1, which relies on batches, the whole data set is needed to obtain representations for individual features. To avoid high spatial complexity, this class of algorithms would ideally operate on sparse inputs. An example of feature-based embeddings are items that are to be recommended to users, where the representation of a given item is obtained by jointly optimizing the item's co-occurrence with other items, as well as other user's properties. In

² As its last step, the methodology includes the explanation of results using the SHAP approach. However, as Sect. 6 focuses on our research contributions, this well known approach and its results are presented in Appendix D.

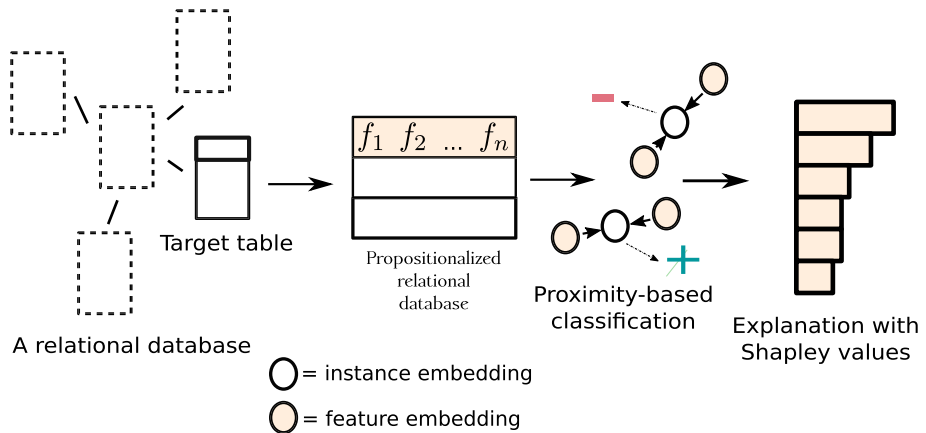


Fig. 4 Overview of the proposed feature-based embedding methodology PropStar. Note that embedded features represent embeddings of single features f_i or of conjuncts of features, e.g., $f_i \wedge f_j$, given that Wordifications constructs both feature forms. For simplicity, the propositionalized database shows two instances Blank and shaded circles correspond to embedded representations of instances and features, respectively

a relational setting considered in this work, we follow the paradigm of propositionalization by Wordification, where each instance is described by a bag of features of the form (table.name, column.name, value). Consequently, in the PropStar approach the embeddings represent bags of such features and their conjuncts (of size 2). There are as many embeddings as there are unique *features* in the propositionalized representation of a given relational database. As such embeddings by themselves do not contain any information which relates them to the desired output space, target values get embedded alongside other features in a supervised manner.

6.2 Detailed description of proposed data transformation and learning methods

This section presents the implementations of the proposed methods, preceded by the description of the updates to the Wordification algorithm (Perovšek et al. 2015 for multi-propositionalization algorithm presented in Sect. 6.2.1. In Sect. 6.2.2 we discuss how Deep Relational Machines (described briefly in Sect. 4.3.3), which use neural networks for learning from relational databases, were adapted to operate on sparse matrices generated by an improved Wordification algorithm. In Sect. 6.2.3 we describe a novel algorithm, called PropStar, which embeds relational features, extracted as part of propositionalization.

6.2.1 Improving the efficiency of Wordification

In this work we significantly extend the ideas proposed in Wordification (Perovšek et al. 2013, 2015) with the aim to maintain the classification performance, yet improve its *scalability*. Both proposed algorithms build on the idea of Wordification, yet its use in our algorithms is differentiated by the following design decisions:

1. Inputs do not need to be hosted in relational databases. PropStar operates on .sql files directly. The algorithm supports SQL conventions, as commonly used in the ILP com-

- munity.³ This modification renders the method completely local, enabling offline execution without additional overhead. Such setting also offers easier parallelism across computing clusters.
- Algorithm is implemented in Python 3 with minimum dependencies for computationally more intense parts, such as the Scikit-learn (Pedregosa et al. 2011), Pandas, and Numpy libraries (Van Der Walt et al. 2011). All database operations are implemented as array queries, filters or similar, unlocking the potential to run PropDRM and PropStar also on GPUs.
 - As shown by Perovšek et al. (2015), Wordification’s caveat is extensive sampling of (all) tables. We relax this constraint to close (up to second order) foreign key neighborhood, notably speeding up the relational item sampling part, but with some loss in terms of relational item diversity. For larger databases, minimum relational item frequency can be specified, constraining potentially noisy parts of the feature space.

One of the original Wordification’s most apparent problems is its spatial complexity. In this work we address this issue as follows:

- Relational items are hashed for minimal spatial overhead during sampling.
- During construction of the final representation, a sparse matrix is filled based on relational item occurrence.
- The matrix is serialized directly into list-like structures, suitable for StarSpace algorithm and thus we maintain minimal spatial overhead.
- Only the final representation is stored as a low-dimensional (e.g., 32) dense matrix.

6.2.2 Detailed description of the proposed PropDRM implementation

The novelty of the proposed implementation of DRM instance-based embedding, inspired by the work of Dash et al. (2018), concerns its capability to effectively handle the sparseness of the data with deep neural networks. The main novelty of the proposed implementation is that it is indeed capable of operating on larger, sparse matrices directly. Such capability is necessary for DRMs to be compatible with propositionalization, which yields large sparse matrices as the main output. Below we discuss the neural network architecture and its adaptations.

Let P represent a sparse item matrix, as returned by Wordification (discussed in Sects. 4.2.3 and 6.2.1). Note that Wordification is unsupervised, and thus does not include any information on instance labels. The neural network we use (termed ω) represents the mapping $\omega : P \rightarrow \mathcal{C}$, where \mathcal{C} is the set of classes. In this work, we experimented with dense feed-forward neural networks, regularized using dropout (Srivastava et al. 2014), and ELU activation function (Clevert et al. 2016) (of intermediary weights). The output weights are activated using sigmoid activation (σ) in order to obtain binary predictions.

$$ELU(x) = \begin{cases} c(e^x - 1), & \text{for } x < 0 \\ x & \text{for } x \geq 0 \end{cases},$$

where c is the user-specified constant. For a given input matrix P , an example of a single hidden-layer neural network is defined as follows.

³ <https://relational.fit.cvut.cz/>.

$$\omega = \sigma(W_o^T(\text{ELU}(\text{Drop}(W_1^T P + b_1))) + b_o).$$

Here, the σ is a sigmoid activation, defined as $\sigma(x) = \frac{1}{1+e^{-x}}$. The W_1 is the weight matrix, P the sparse input space, and b_l the bias vector of a given layer $l \in \{0, 1\}$. The described neural network returned satisfactory results, hence, we did not perform neuroevolution or similar large-scale search for potentially better performing architectures. Throughout this work, we use the binary cross-entropy loss, referred to as *Loss*. For a given probabilistic classifier, which returns a probability p_{ij} of an instance i belonging to a class j , the loss function is defined as follows:

$$\text{Loss}(i) = \sum_{j \in \mathcal{C}} y_{ij} \cdot \log p_{ij}.$$

Here y_{ij} is a binary value (0 or 1) indicating whether class j is the correct class label assigned to instance i , and \mathcal{C} is a set of all the target classes. In the case of DRMs, where the instances of a relational database (one of the tables) are classified, each of the $|\mathcal{C}|$ output neurons predicts a single probability p_{ij} for a given target class $j \in \mathcal{C}$. If the neural networks are trained in small batches, the results of the Loss function are averaged to obtain the overall loss of a given batch of instances.

Neural networks are adapted for dense inputs such as images and texts, and are not necessarily suitable for large sparse matrices, as considered in this work (i.e. P). The proposed variant of DRMs is adapted as follows. Once the batch size bs (a free parameter) is determined, propositionalized representation P is traversed (in chunks of bs instances). Note that each instance is effectively a d -dimensional vector. As the neural network operates with dense batches, each batch is converted to a dense matrix of $bs \cdot d$ elements that is used during matrix multiplication within the neural network. The spatial complexity is thus at most $\mathcal{O}(bs \cdot d)$. We observed that even by considering batch size of one, the DRMs are stable and efficient.

6.2.3 Detailed description of the PropStar algorithm

We next present the novel feature-based embedding algorithm that can operate directly on the propositionalized relational databases. The proposed PropStar algorithm merges symbolic and non-symbolic representations as part of a single procedure for obtaining real-valued representations of features in arbitrary relational databases. The pseudocode of the PropStar algorithm is given in Algorithm 1.

Algorithm 1: The pseudocode of PropStar algorithm.

Data: A Relational database R , foreign key map m
Parameters : Entity embedding parameter set \mathcal{E} , representation dimension d , target table T , target attribute t

```

1 itemContainer ← empty bag of items ;                                ▷ Begin Wordification.
2 foreach instance  $i \in T$  do
3   relationalItems ← {} ;
4   candidateKeys ← getForeignKeys( $R, i$ ) ;                            ▷ Links to other tables.
5   candidateTables ← getCandidateTables(candidateKeys,  $R$ ) ;          ▷ Linked tables.
6   foreach table  $\in$  candidateTables do
7     while not final number of items do
8       bagOfItems ← WORDIFY(table( $m(\text{instance})$ )) ;
9       add bagOfItems to relationalItems ;                            ▷ Store sampled items.
10    end
11  end
12  itemContainer[ $i$ ].add(relationalItems) ;                            ▷ Store relational items.
13 end
14 relevantFeatures ← frequencySelection(itemContainer.values,  $d$ ) ;
15 symRep ← [] ;                                                       ▷ Sparse vector representations of instances.
16 foreach instance  $i \in$  targetTable do
17   instanceItems ← itemContainer[ $i$ ] ;
18   propRep ← RelationalFeatures(relevantFeatures, instanceItems) ;
19   symRep.append(propRep) ;
20 end
21 featureEmbeddings ← StarSpace(symRep,  $T[t], \mathcal{E}$ ) ;                 ▷ Input is a sparse matrix.
22 return featureEmbeddings ;

```

The algorithm consists of two main steps. First, a relational database is transformed into sets of features describing individual instances. The WORDIFY method constructs features of the form (table.name, column.name, value) and uses them to describe each individual instance (see Sect. 6.2.1 for a detailed formulation of this step).

Second, sets of relational items (features) are used as input to the StarSpace entity embedding algorithm (described in Sect. 4.3.2), producing embeddings for each distinct relational item. The StarSpace embeddings are computed using efficient C++ implementation. We wrote a wrapper which seemingly integrates the first part of PropStar (sampling and propositionalization) with the embedding construction. The problem is formulated as a multiclass classification, where the positive pair generator comes directly from a training set of labeled data specifying $(a, b) \in E^+$ pairs where a are relational item ‘documents’ and b are labels (singleton features). Negative entities b_i^- are sampled from the set of possible labels. Inputs can be described as (multi) sets comprised of both relational items f_i , their conjuncts, as well as class labels c_i . For example,

$$\{f_1, f_2, f_6, f_6 \wedge f_2, c_1\}$$

represents a simple input consisting of three relational items, a conjunct and the target label c_1 . Note that we apply StarSpace in such manner that the representations are learned for *individual relational items*. A representation matrix of dimension $\mathbb{R}^{|W| \times d}$ is produced as the final output ($|W|$ represents the number of unique relational items considered). Intuitively, the embedding construction can be understood as determining relational item locations in a latent space based on their co-occurrence with other items present in all training instances. The wrapper can be called via ‘fit’ and ‘predict’ methods, commonly used in contemporary data science and machine learning. In this work, we consider the inner product similarity

between a pair of vectors e_1, e_2 for the construction of embeddings.⁴ i.e. The complexity of obtaining $\text{sim}(e_1, e_2) = e_1^T \cdot e_2$. As the class labels are embedded in *the same* space as individual relational items, classification of novel bags of relational items is possible by direct comparison, as common tasks operating on word embeddings. We discuss this classification below.

Let M represent a novel instance to be classified. Note that M (without additional index) is considered a multiset of relational items. For prediction purposes, StarSpace averages the representations of relational items present in a given input instance (a bag). The representation is normalized (as during training) and compared to label embeddings in the common space. Representation of a relational bag e_M is computed (with considered hyperparameters) as:

$$e_M = \frac{\bigoplus_{f_i \in M} e_{f_i}}{\sqrt{|M_{\text{unique}}|}}$$

which is a d -dimensional, real-valued vector. Note that \bigoplus in this expression denotes element-wise summation. The M_{unique} represents the set of all (unique) relational features currently considered. Note that original bags of features can be redundant (multisets), yet representations are learned for unique features. Next, the similarity of this vector is compared to the label embeddings in the same space. The label that is the most similar to e_M is the top-ranked prediction, the second most similar label is the second-ranked prediction, etc. In this work we consider only the top-ranked prediction, resulting in the following label assignment:

$$\text{label}(e_M) = \arg \max_{c \in \mathcal{C}} [\text{sim}(e_M, e_c)].$$

The complexity of obtaining a single prediction is hence $\mathcal{O}(|\mathcal{C}|)$, not taking the complexity of scalar product for function sim into account. The PropStar algorithm first samples the relational items with respect to the target table (lines 2-11 in Algorithm 1). Binary indicator function (relationalFeatures) is applied to obtain the propositionalized representation of the target table (line 12). Here, zeros represent absence of a given relational items, and ones their presence.⁵ Finally, StarSpace is used to embed the table into a low-dimensional, real valued embedding (line 19).

The spatial complexity of PropStar is linear with respect to the number of non-zero elements in the propositionalized version of a relational database. The exact spatial complexity can be formulated as follows. Let row represent the average number of rows per table. Let n_t represent the number of tables and col the average number of columns per table. We improve the original spatial complexity of $\mathcal{O}(\text{rows} \cdot n_t \cdot 2^{\text{col}})$ by introducing a constraint, which determines the maximum number of relational items that can be considered. The exponential term in the initial complexity thus reduces to col times some constant, yielding the complexity of $\mathcal{O}(\text{rows} \cdot \text{col} \cdot n_t)$. This formulation yields a scalable propositionalization.

⁴ Note that e_1, e_2 represent vector representations of relational items (i.e. features) in the output of propositionalization.

⁵ Note that in the actual implementation CSR format of sparse matrices is used to reduce the spatial overhead of storing zeros.

7 Experimental evaluation

In this section we describe the implementation details of the proposed methods, the relational data sets used in the experiments, and the experimental evaluation of the proposed methods.

7.1 Implementation and hyperparameters

We discuss how the proposed methods were implemented, along with the hyperparameters explored. Both new methods (PropDRM and PropStar) are implemented in Python, with the following exceptions. In PropDRM, the DRMs are implemented in PyTorch. For PropStar we used the efficient StarSpace implementation written in C++, for which we build a wrapper offering basic embedding training and prediction functionality.

We used 10-fold stratified cross validation, which was conducted for individual hyperparameter settings. The best setting is reported, other are discussed in ablation studies. Experiments were performed on an of-the-shelf workstation with no GPUs (even though PropDRM and PropStar can exploit them). We intentionally omit the GPU-based training to explore the minimum hardware, required to perform competitively on the selected data sets—ILP baselines, such as Aleph and RSD are Prolog-based, and are to our knowledge not able to use multiple GPU threads simultaneously. The machine on which experiments were conducted had 128GB of RAM and 12 CPUs (Intel i8 series).

In PropDRM, we varied the dropout rate, learning rate, number of epochs, and the hidden layer size. In PropStar, we varied the number of negative samples, embedding dimension, learning rate, and the number of epochs.

The source code of our implementation is publicly available⁶.

7.2 Relational data sets

Five relational database sources⁷ (Motl and Schulte 2015) were used in the experiments. Their characteristics are summarized in Table 4.

Trains (Michie et al. 1994) data set is used in the East-West trains challenge problem, which is well-known in ILP. The East-West trains challenge is to predict whether a train is eastbound or westbound, based on the properties of eastbound and westbound cars. Trains contain variable number of cars, each having one of various shapes and carrying various loads.

Carcinogenesis (Srinivasan et al. 1997) task is to predict carcinogenicity of a diverse set of chemical compounds. The data set was obtained by testing different chemicals on rodents, where each trial would take several years and hundreds of animals. The data set consists of 329 compounds, of which 182 are carcinogens.

Mutagenesis (Debnath et al. 1991) task addresses the problem of predicting mutagenicity of aromatic and heteroaromatic nitro compounds. Predicting mutagenicity is an important task as it is very relevant to the prediction of carcinogenesis. The compounds from the data are known to be more structurally heterogeneous than in any other ILP

⁶ <https://github.com/SkBlaz/PropStar>.

⁷ Freely accessible at <https://relational.fit.cvut.cz/>.

Table 4 Properties of the experimental data tables

Trains	#rows	#attributes
Cars	63	10
trains	20	2
Carcinogenesis	#rows	#attributes
atom	9,064	5
canc	329	2
sbond_1	13,562	4
sbond_2	926	4
sbond_3	12	4
sbond_7	4,134	4
Mutagenesis 42	#rows	#attributes
atoms	1,001	5
bonds	1,066	5
drugs	42	7
ring_atom	1,785	3
ring_strucs	279	3
rings	259	2
Mutagenesis 188	#rows	#attributes
atoms	4,893	5
bonds	5,243	5
drugs	188	7
ring_atom	9,330	3
ring_strucs	1,433	3
rings	1,317	2
IMDB	#rows	#attributes
actors	7,118	4
directors	130	3
directors_genres	1,123	4
movies	166	4
movies_directors	180	3
movies_genres	408	3
roles	7,738	4
MovieLens	#rows	#attributes
actors	99,129	3
directors	2,201	3
movies	3,832	5
movies2actors	152,532	3
movies2directors	4,141	3
u2base	946,828	3
users	6,039	4

data set of chemical structures. The database contains 230 compounds of which 138 have positive levels of mutagenicity and are labeled as ‘active’. Others have class value ‘inactive’ and are considered to be negative examples. We took the data sets from the original paper (Debnath et al. 1991), where the data was split into two subsets: a 188 compound data set and a smaller data set with 42 compounds.

IMDB database is publicly available in the SQL format.⁸ This database contains tables of movies, actors, movie genres, directors, and director genres. The data set used in our experiments encompasses only movies whose titles and years of production appear in the IMDB’s top-250 and bottom-100 chart (Snapshot taken on July 2, 2012). The snapshot contains 166 movies, along with all of their actors, genres and directors. We designate movies present in the IMDB top-250 chart as positive examples, and those in the bottom-100 as negatives.

MovieLens data set from the UC Irvine machine learning repository.⁹ The data set is similar to IMDB above, however is much larger. Overall, the database consists of more than 1.2 million instances. The task is to predict gender of the movie database’s users.

7.3 Results

We present the results of the empirical evaluation of the proposed methodologies on the presented set of standard benchmark ILP data sets. The accuracies of individual learners are given in Table 5, and the AUC scores are reported in Table 6. The results for Aleph, RSD, RelF and Wordification were taken from previous work of Perovšek et al. (2015).

It can be observed that the proposed unifying approaches perform competitively on most data sets. We can observe a distinct difference in performance on the Mutagenesis data sets, where both PropDRM as well as PropStar do not outperform the baselines on the smaller data set (Mut42), yet notably outperform the (best) baselines on the larger one (Mut188). Further, minor improvement over the baseline algorithms is also achieved on Carcinogenesis data set.

In terms of spatial complexity, the proposed methodology greatly outperforms the alternatives under a given set of constraints. Only PropDRM and PropStar scale to very large relational databases without specialized hardware. Detailed studies regarding the sensitivity of PropDRM and PropStar to their parameters are discussed in Appendices B and C, respectively.

We consider the presented results as very favorable for the two proposed approaches. In particular, PropStar is better than current state-of-the-art methods on 3 out of 6 data sets, and is therefore a method to take into consideration when attempting to solve any new relational problem.

7.3.1 Study of propositionalization projections

The considered propositionalization is entirely *unsupervised*. Only once the symbolic representations of instances are obtained, PropDRM and PropStar learn the associations to

⁸ <http://www.webstepbook.com/supplements/databases/imdb.sql>.

⁹ <https://relational.fit.cvut.cz/dataset/MovieLens>.

Table 5 Classification accuracy on different relational data sets

Propositionalization	Learner	Carc.	IMDB	Mut188	Mut42	Trains	MovieLens
	MajorityVote	0.55	0.73	0.67	0.69	0.50	0.72
Aleph (Perovšek et al. 2015)	J48	0.55	0.73	0.60	0.69	0.55	–
Aleph (Perovšek et al. 2015)	SVM	0.55	0.73	0.60	0.69	0.70	–
RSD (Perovšek et al. 2015)	J48	0.60	0.75	0.68	0.98	0.60	–
RSD (Perovšek et al. 2015)	SVM	0.56	0.73	0.71	0.69	0.80	–
RelF (Perovšek et al. 2015)	J48	0.60	0.70	0.75	0.76	0.65	–
RelF (Perovšek et al. 2015)	SVM	0.56	0.73	0.69	0.76	0.80	–
Wordification (Perovšek et al. 2015)	J48	0.62	0.82	0.67	0.98	0.50	–
Wordification (Perovšek et al. 2015)	SVM	0.61	0.73	0.82	0.79	0.50	–
Aleph (replicated)	J48	0.55	–	0.80	0.76	0.70	–
Aleph (replicated)	SVM	0.55	–	0.80	0.79	0.60	–
RSD (replicated)	J48	0.56	0.84	0.88	0.92	0.60	–
RSD (replicated)	SVM	0.60	0.82	0.89	0.84	0.80	–
Wordification (replicated)	J48	0.47	0.85	0.91	0.88	0.90	0.60
Wordification (replicated)	SVM	0.39	0.80	0.83	0.33	0.50	0.72
Treeliker	J48	0.58	–	0.77	0.81	0.50	–
Treeliker	SVM	0.60	–	0.90	0.80	0.70	–
PropDRM		0.63	0.73	0.91	0.86	0.70	0.72
PropStar		0.66	0.74	0.92	0.90	0.80	0.74

The best score for each dataset is in bold

For the proposed methods, we report average performance over 5 runs. The runs, marked with—were unable to finish in 12 h

individual classes. A good representation, however, already contains relevant information on the instance space. In Fig. 5, we projected the propositionalized Mutagenesis 188 and Trains instance space to two dimensions to qualitatively explore whether instances group or any meaningful patterns emerge. Understanding whether the symbolic space exhibits distinct structure on its own could offer insights into why the proposed methods perform well. For projecting the 10,000 dimensional space to two dimensions we used UMAP, a recently introduced non-linear dimensionality reduction method based on insights from manifold theory (McInnes et al. 2018).

We can observe an apparent distinction in the clustering of the UMAP projections of the two propositionalized data sets. The Mutagenesis 188 data set consists of two distinct clusters that, when colored according to the class labels, approximately correspond to the two classes (Fig. 5a). On the other hand, the clustering is not apparent in the case of the Trains data set (Fig. 5b), where the instances do not group distinctly. The purpose of the considered visualizations is twofold. First, we show how the symbolic space can exhibit clustering properties, related to properties of instances such as class labels. Next, we show that projections do not necessarily exhibit such properties, indicating potentially harder classification problems. We believe that UMAP and similar tools offer insights into representation structure.

Table 6 AUC scores on individual data sets

Propositionalization	Learner	Carc.	IMDB	Mut188	Mut42	Trains	Movies
Aleph (from Perovšek et al. 2015)	J48	0.50	0.50	0.68	0.50	0.55	—
Aleph (from Perovšek et al. 2015)	SVM	0.50	0.50	0.68	0.50	0.70	—
RSD (from Perovšek et al. 2015)	J48	0.59	0.59	0.54	0.96	0.60	—
RSD (from Perovšek et al. 2015)	SVM	0.52	0.50	0.58	0.50	0.80	—
ReIF (from Perovšek et al. 2015)	J48	0.59	0.66	0.68	0.68	0.75	—
ReIF (from Perovšek et al. 2015)	SVM	0.52	0.50	0.54	0.62	0.75	—
Wordification (from Perovšek et al. 2015)	J48	0.61	0.75	0.55	0.96	0.95	—
Wordification (from Perovšek et al. 2015)	SVM	0.58	0.50	0.78	0.65	0.50	—
Alpeh (replicated)	J48	0.50	—	0.71	0.72	0.70	—
Aleph (replicated)	SVM	0.50	—	0.75	0.73	0.60	—
RSD (replicated)	J48	0.55	0.71	0.87	0.92	0.60	—
RSD (replicated)	SVM	0.58	0.65	0.90	0.73	0.80	—
Wordification (replicated)	J48	0.48	0.72	0.90	0.86	0.90	0.52
Wordification (replicated)	SVM	0.42	0.62	0.81	0.50	0.50	0.50
Treeliker	J48	0.58	—	0.75	0.71	0.50	—
Treeliker	SVM	0.58	—	0.88	0.68	0.70	—
PropDRM		0.63	0.68	0.90	0.87	0.80	0.54
PropStar		0.63	0.66	0.87	0.87	0.95	0.56

The best score for each dataset is in bold

We report average performance over 5 runs. The runs, marked with—were unable to finish in 12 h

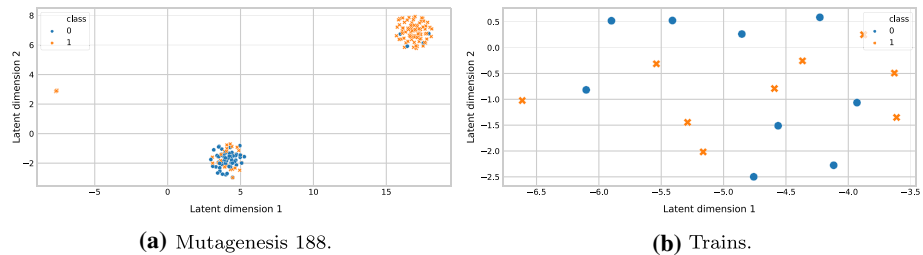


Fig. 5 Two UMAP projections of selected propositionalized data sets

7.3.2 Statistical comparison of PropDRM and PropStar

In previous sections, we demonstrated that both PropDRM and PropStar perform well on the considered data sets, indicating that both approaches are successfully unifying propositionalization and embeddings. We further study the differences in performances of the two approaches. For this purpose, we employ the hierarchical Bayesian t-test, a Bayesian test capable of comparing a pair of classifiers across multiple data sets (Benavoli et al. 2017; Corani et al. 2017). For this comparison, we selected the overall best performing hyperparameter sets for each method, and conducted ten repetitions of stratified ten-fold cross validation (for each data set). The results are visualized as probability distributions across the space of both classifiers and the ‘rope’ region (region of practical equivalence) within which the two classifiers perform the same. The size of this region is a free parameter of the hierarchical

t-test, and was set to 0.05 in this work. Other parameters of the test were left as defaults. The exact methodology for the interested reader is explained by Benavoli et al. (2017).

In terms of AUC, the probabilities returned by the Bayesian test were as follows: $p(\text{PropStar}) = 0.07$ and $p(\text{PropDRM}) = 0.54$, and in terms of classification accuracy, $p(\text{PropStar}) = 0.96$ and $p(\text{PropDRM}) = 0.04$. The results of statistical analysis indicate that with respect to AUC performance, the two approaches perform similarly, even though the probability that PropDRM will outperform PropStar is higher. With respect to the classification accuracy, PropStar outperforms PropDRM in majority of comparisons. Thus, considering the 95% or higher as the probability denoting significance boundary, we can determine that PropStar is (significantly) more suitable choice if accuracy is being optimized for. As Bayesian comparisons are computationally expensive, we compared the two methods using default hyperparameter settings. The PropStar's default configuration is not necessarily optimal when AUC is considered.

8 Conclusions and further work

This paper first provides a critical survey of propositionalization and embedding techniques, especially relevant for relational learning and inductive learning programming. While both data approaches, propositionalization and embeddings, aim at transforming data into the tabular data format, the research papers describing the approaches use different terminology and task definitions, claim to have different goals, and are used in very different contexts. In this paper, we define the main categories of data transformation techniques based on the representation they use and approaches employed. Propositionalization approaches produce tabular data from multirelational databases as well as from a mixture of tabular data and background knowledge in the form of logic programs or networked data, including ontologies. Knowledge stored in graphs can be assessed with community detection and graph traversal methods. Relations described with similarity matrices are encoded in a numeric form using matrix factorization. Currently, the most promising approach to data transformations are neural networks based methods which can be applied to text, graphs, and other entities for which we can define a suitable similarity function.

One of the main strategic problems machine learning has to solve is better integration of knowledge and models across different domains and representations. While the research area of embeddings can unify different representations in a numeric space, symbolic learning may be an essential ingredient for integration of different knowledge areas. We see our PropStar approach that combines advantages of propositionalization and neural embeddings in the same data fusion pipeline as a step in that direction.

The first minor contribution of the paper is that our exposition is based on three cognitive representation levels introduced by Gärdenfors (2000), i.e. neural, spatial, and symbolic. As most of human knowledge is stored in the symbolic form, while the most powerful machine learning algorithms take as input spatial representations, this explains a plethora of techniques that transform other forms of human knowledge into the spatial representation space. The next contribution is the unifying framework in which we describe propositionalization and embedding techniques in terms of their joint properties and their differences. We show how the propositionalization techniques can be merged with deep neural network based embedding to produce a joint embedding, such that spatial representation can be used with any deep learning algorithm and the predictions can be comprehensively explained. The main contributions of our work are thus the two implementations that merge propositionalization and embeddings in the same unifying methodology. The first is an efficient reimplementa-

of existing Deep Relational Machines, while the second one is the novel Deep Propositionalization algorithm. We also contribute an experimental evaluation of the two algorithms and show favorable results in terms of predictive performance, as well as time and space requirements. The source code of both algorithms, DeepProp and PropDRM, is publicly available.¹⁰

In further work, it is worth investigating the integration of symbolic and deep learning, considering them as multitask learning approaches which try to integrate many different learning tasks and use embeddings as input representations. The issue is that different embedding methods have so far only been used in isolation. We already address this challenge in the current work of the authors, where we combine complementary embedding methods from different classes: in particular, to use network traversal methods to produce initial embeddings that are then refined using a deep neural network (Škrlić et al. 2019).

Acknowledgements We acknowledge the financial support of the Slovenian Research Agency through core research programmes P2-0103 and P6-0411 and project *Semantic Data Mining for Linked Open Data* (financed under the ERC Complementary Scheme, N2-0078). The authors have received funding also from the European Union's Horizon 2020 research and innovation programme under grant agreement No 825153 (EMBEDDIA). The work of the second author was funded by the Slovenian Research Agency through a young researcher grant. We wish to thank Jan Kralj for his insightful comments on the formulation of the proposed framework and for mathematical proofreading. Further, we are grateful to Vid Podpečan and Nika Eržen for their help with the implementation of the new version of the PyRDM library. Finally, we would like to thank the anonymous reviewers for careful reading, many insightful observations, and the encouragements to expand the initial work.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Appendix A: Wordification example

The Wordification approach is illustrated on a modified and substantially simplified version of the well-known East-West Trains domain (Michie et al. 1994). Our input database consists of just two tables shown in Fig. 6, where we have only one east-bound and one west-bound train, each with just two cars with certain properties¹¹.

The TRAIN table is the main table and the trains are the instances, with a class label denoting the direction of the train (east of west). As Fig. 7 shows, a multiset (a bag) of features is generated for each train t1 and t5 with the class label appended to the resulting feature vector (bag of features). Both single features and conjunctive features are shown in this example.

¹⁰ <https://github.com/SkBlaz/PropStar>.

¹¹ Note that in the experiments we use the standard version of the East-West Trains domain.

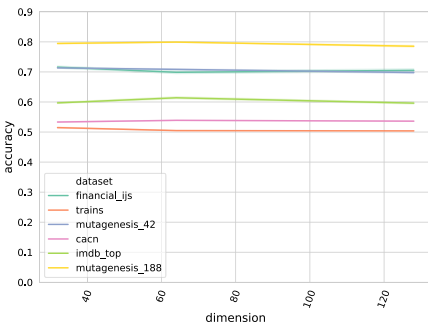
TRAIN		CAR			
trainID	eastbound	carID	shape	roof	train
t1	east	c11	rectangle	none	t1
...	...	c12	rectangle	peaked	t1
...
t5	west	c51	rectangle	none	t5
...	...	c52	hexagon	flat	t5
...

Fig. 6 Example input for Wordification in the East-West Trains domain

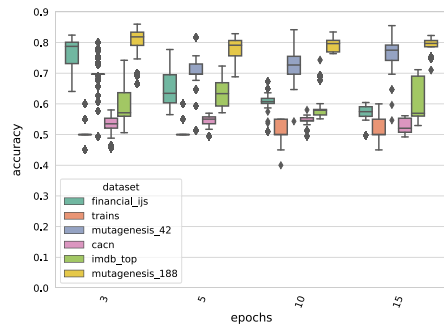
```

t1: {car_shape_rectangle, car_roof_none, car_shape_rectangle__car_roof_none,
    car_shape_rectangle, car_roof_peaked, car_shape_rectangle__car_roof_peaked}, east
...
t5: {car_shape_rectangle, car_roof_none, car_shape_rectangle__car_roof_none,
    car_shape_hexagon, car_roof_flat, car_shape_hexagon__car_roof_flat}, west
...
    
```

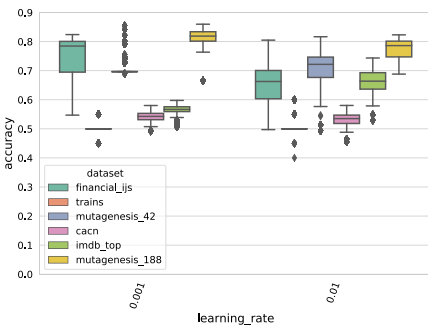
Fig. 7 The database from Fig. 6 in the bag-of-features representation (as in the original Wordification implementation, conjunctions of features are denoted by a long underscore instead of \wedge)



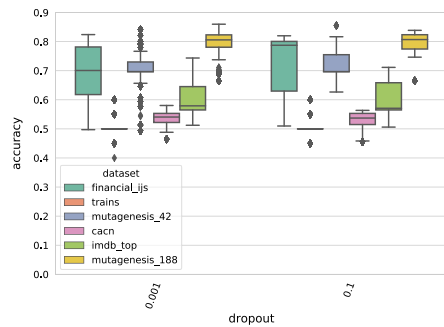
(a) Dependence on hidden layer size.



(b) Dependence on the number of epochs.



(c) Dependence on the learning rate.



(d) Dependence on Dropout.

Fig. 8 Sensitivity of PropDRM to hyperparameter settings

Appendix B: Ablation study—PropDRM

We discuss the impact of individual hyperparameters on the performance of PropDRM. We first visualize the performance of PropDRM w.r.t. individual hyperparameters in Fig. 8.

We can observe that the relevance of individual hyperparameters varies from data set to data set. The learning rate, when too small, decreases the performance. In terms of embedding dimension, even smaller dimensions are sufficient for the considered data sets. This result potentially implies that the considered data sets are relatively small and contain only a small set of relevant features (when propositionalized). Thus, if the neural network detects the correct features as relevant, not many parameters are needed for a successful classification. An alternative explanation would imply that PropDRM learns hierarchical representations efficiently, albeit not optimized with their hierarchical nature in mind, which was previously demonstrated to capture hierarchical relations well (Nickel and Kiela 2017).

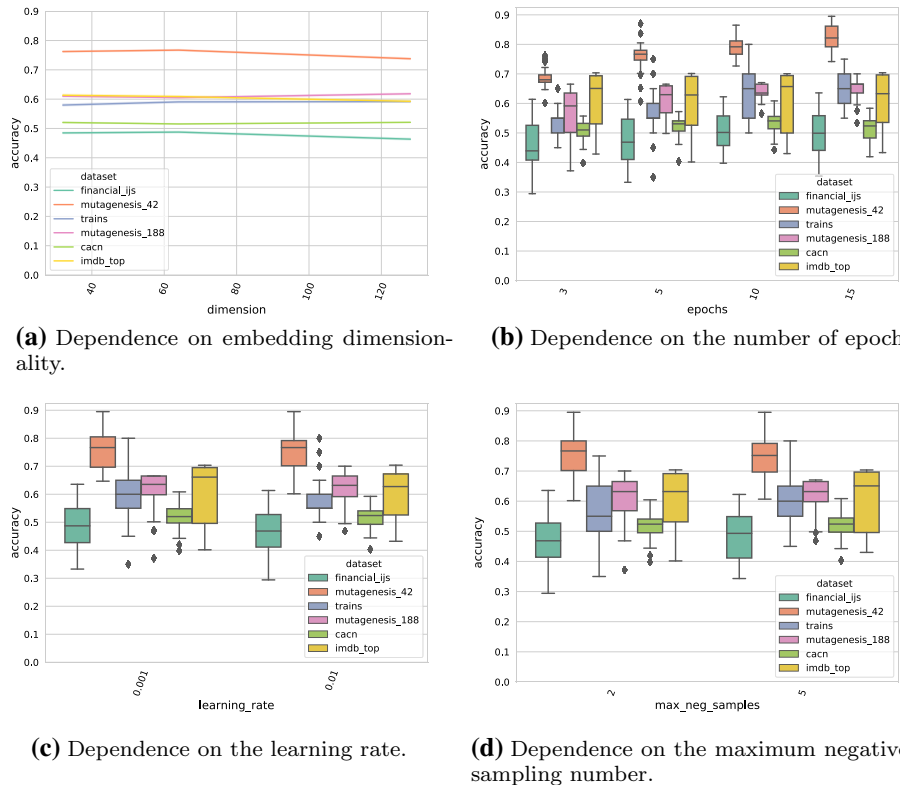


Fig. 9 Sensitivity of PropStar to various hyperparameter settings

Appendix C: Ablation study—PropStar

We first explore the behavior with respect to various hyperparameter settings and visualize them in Fig. 9. We can observe that the amount of negative samples (Subfigure 9d)) impacts the PropStar’s performance the most on the mutagenesis 42 data set, overall reducing the performance, even though a handful of models (outliers marked as dots) perform well. This indicates the importance of negative sample selection. As StarSpace does not use any sophisticated technique for sampling negative examples, the variability in performance could be notable due to this parameter.

It can be observed that a relatively small relational item embedding dimensionality is needed for successful performance. The dependence on other parameters varies from data set to data set. For example, the learning rate does not impact the larger Mutagenesis data set (Mut188) as much as it does the Trains data set. As the proposed methodology is not well adapted to such small data sets (e.g., tens of instances), large variability in performance could be linked to potential overfitting. Further, sufficient number of epochs are needed for PropStar to converge on individual data sets.

Appendix D: Interpretability of embedding-based methods using SHAP

The approximation power of deep neural network which are commonly used with embeddings comes at a cost of lesser interpretability. Compared to symbolic relational (or propositional) learners, one cannot understand the deep relational models’ deductive process by inspecting the model. However, *post hoc* explanation methods for prediction models can be used to better understand which parts of the feature space are relevant for the neural network’s individual predictions. In this work, we leverage the state-of-the-art explanation tool SHAP (Lundberg and Lee 2017), based on the coalitional game theory. This tool captures the importance of interactions between features with Shapley values.

When considered in a feature importance scenario, the contribution of the i -th instance τ_i , is approximated by SHAP with the following expression:

$$\tau_i = \underbrace{\sum_{S \subseteq F \setminus \{i\}} \frac{|S|!(|F| - |S| - 1)!}{|F|!}}_{\text{All possible subsets}} \underbrace{[f(x_{S \cup \{i\}}) - f(x_S)]}_{\text{Difference in predictive performance}}$$

where S is a subset of all features F , f is the used predictive model, and x_S is an instance containing only features from the subset S . Shapley values offer insights into instance-level predictions by assigning fair credit to individual features for participation in interactions. They are commonly used to understand and debug black-box models.

In this work, we use the SHAP kernel approximator, the recently introduced, model-agnostic method for explaining model outputs. The used SHAP kernel explainer is considered an additive feature attribution method. Such methods are characterized as having an explanation model g that is a linear function of binary variables:

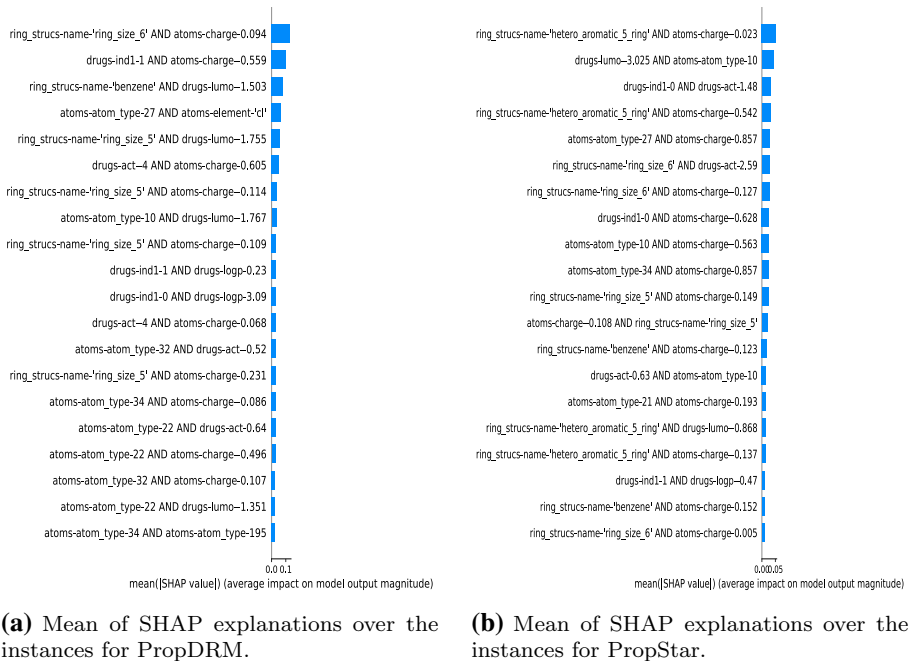


Fig. 10 SHAP Kernel explanations of the two developed approaches

$$g(z') = \phi_0 + \sum_{i=1}^{|F|} \phi_i \cdot z'_i$$

where $z' \in \{0, 1\}^{|F|}$, $|F|$ is the number of input features and $\phi_i \in \mathbb{R}$. This class of models assign an effect ϕ_i to each feature, and summing the effects of all such feature attributions approximates the output $f(x)$ of the original model. Detailed theoretical analysis of how this idea can be extended to approximation of outputs via a kernel is given in Lundberg and Lee (2017).

As an example demonstrating the explainability of the two paradigms, we visualize the Shapley values as explanations of learned classifiers for Mutagenesis 188 problem in Fig. 10. Explanations indicate parts of the feature space that have the largest impact on the model's output. Even though the considered SHAP kernel explainer is known to be a computationally expensive variant of SHAP (it is also the most general one), explanations were obtained in the order of minutes, indicating the potential of this methodology for explanations of predictors in larger relational databases.

References

Ahmed, C. F., Lachiche, N., Charnay, C., Jelali, S. E., & Braud, A. (2015). Flexible propositionalization of continuous attributes in relational data mining. *Expert Systems with Applications*, 42(21), 7698–7709.

- Benavoli, A., Corani, G., Demšar, J., & Zaffalon, M. (2017). Time for a change: A tutorial for comparing multiple classifiers through Bayesian analysis. *Journal of Machine Learning Research*, 18(1), 2653–2688.
- Bennett, K. P., Buja, A., Freund, W. S. Y., Schapire, R. E., Friedman, J., Hastie, T., et al. (2008). Responses to [52]. *Journal of Machine Learning Research*, 9, 157–194.
- Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent Dirichlet allocation. *Journal of Machine Learning Research*, 3(Jan), 993–1022.
- Blockeel, H., Raedt, L. D., & Ramon, J. (1998). Top-down induction of clustering trees. In *Proceedings of the 15th international conference on machine learning*, pp. 55–63. Morgan Kaufmann.
- Bordes, A., Usunier, N., Garcia-Duran, A., Weston, J., & Yakhnenko, O. (2013). Translating embeddings for modeling multi-relational data. *Advances in Neural Information Processing Systems*, pp. 2787–2795.
- Bordes, A., Glorot, X., Weston, J., & Bengio, Y. (2014). A semantic matching energy function for learning with multi-relational data. *Machine Learning*, 94(2), 233–259.
- Breiman, L. (1996). Bagging predictors. *Machine Learning*, 24(2), 123–140.
- Breiman, L. (2001). Random forests. *Machine Learning*, 45(1), 5–32.
- Breiman, L., Friedman, J. H., Olshen, R., & Stone, C. (1984). *Classification and regression trees*. Pacific Grove, CA: Wadsworth & Brooks.
- Chang, S., Han, W., Tang, J., Qi, G. J., Aggarwal, C. C., & Huang, T. S. (2015). Heterogeneous network embedding via deep architectures. In *Proceedings of the 21th ACM SIGKDD international conference on knowledge discovery and data mining*, pp. 119–128. ACM.
- Charnay, C., Lachiche, N., & Braud, A. (2015). CARAF: Complex aggregates within random forests. In *Inductive logic programming—25th international conference, ILP 2015*, Kyoto, Japan, August 20–22, 2015, Revised Selected Papers, pp. 15–29. Springer.
- Clark, P., & Niblett, T. (1989). The CN2 induction algorithm. *Machine Learning*, 3(4), 261–283.
- Clevert, D. A., Unterthiner, T., & Hochreiter, S. (2016). Fast and accurate deep network learning by exponential linear units (ELUs). In *International conference on representation learning, ICLR*. [arXiv:1511.07289](https://arxiv.org/abs/1511.07289).
- Corani, G., Benavoli, A., Demšar, J., Mangili, F., & Zaffalon, M. (2017). Statistical comparison of classifiers through Bayesian hierarchical modelling. *Machine Learning*, 106(11), 1817–1837.
- Cortes, C., & Vapnik, V. (1995). Support-vector networks. *Machine Learning*, 20(3), 273–297.
- Cumby, C. M., & Roth, D. (2003). On kernel methods for relational learning. In *Proceedings of the 20th international conference on machine learning (ICML-03)*, pp. 107–114.
- Dash, T., Srinivasan, A., Vig, L., Orhobor, O. I., & King, R. D. (2018). Large-scale assessment of deep relational machines. In *Proceedings of the international conference on inductive logic programming*, pp. 22–37. Springer, Berlin.
- Dash, T., Srinivasan, A., Joshi, R. S., & Baskar, A. (2019). Discrete stochastic search and its application to feature-selection for deep relational machines. In I. V. Tetko, V. Kůrková, P. Karpov, & F. Theis (Eds.), *Artificial neural networks and machine learning: ICANN 2019—deep Learning* (pp. 29–45). Berlin: Springer.
- De Raedt, L. (2008). *Logical and relational learning*. Berlin: Springer.
- Debnath, A. K., Lopez de Compadre, R. L., Debnath, G., Shusterman, A. J., & Hansch, C. (1991). Structure-activity relationship of mutagenic aromatic and heteroaromatic nitro compounds. Correlation with molecular orbital energies and hydrophobicity. *Journal of Medicinal Chemistry*, 34(2), 786–797.
- Demeester, T., Rocktäschel, T., & Riedel, S. (2016). Lifted rule injection for relation embeddings. In *Proceedings of the 2016 conference on empirical methods in natural language processing*, pp. 1389–1399.
- Dumančić, S., Guns, T., Meert, W., & Blockeel, H. (2018). Auto-encoding logic programs. In *Proceedings of the international conference on machine learning*, Stockholm, Sweden.
- Džeroski, S., & Lavrač, N. (Eds.). (2001). *Relational data mining*. Berlin: Springer.
- Flach, P., & Lachiche, N. (1999). IBC: A first-order Bayesian classifier. In *International conference on inductive logic programming*, pp. 92–103. Berlin: Springer.
- Flach, P., & Lachiche, N. (2001). Confirmation-guided discovery of first-order rules with Tertius. *Machine Learning*, 42(1/2), 61–95.
- Freund, Y., & Schapire, R. E. (1997). A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of Computer and System Sciences*, 55(1), 119–139.
- Friedman, J. H., & Fisher, N. I. (1999). Bump hunting in high-dimensional data. *Statistics and Computing*, 9(2), 123–143.
- Gärdenfors, P. (2000). *Conceptual spaces: The geometry of thought*. Cambridge, MA: MIT Press.
- Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. Cambridge: MIT Press.
- Grčar, M., Trdin, N., & Lavrač, N. (2013). A methodology for mining document-enriched heterogeneous information networks. *The Computer Journal*, 56(3), 321–335.

- Grover, A., & Leskovec, J. (2016). node2vec: Scalable feature learning for networks. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pp. 855–864.
- Guo, S., Wang, Q., Wang, L., Wang, B., & Guo, L. (2016). Jointly embedding knowledge graphs and logical rules. In *Proceedings of the 2016 conference on empirical methods in natural language processing*, pp. 192–202.
- Haussler, D. (1999). Convolution kernels on discrete structures. Tech. rep., Department of Computer Science, University of California.
- He, S., Liu, K., Ji, G., & Zhao, J. (2015). Learning to represent knowledge graphs with Gaussian embedding. In *Proceedings of the 24th ACM international on conference on information and knowledge management*, pp. 623–632. ACM.
- Kralj, J., Robnik-Šikonja, M., & Lavrač, N. (2018). HINMINE: Heterogeneous information network mining with information retrieval heuristics. *Journal of Intelligent Information Systems*, 50(1), 29–61.
- Kramer, S., Lavrač, N., & Flach, P. (2001). Propositionalization approaches to relational data mining. In S. Džeroski & N. Lavrač (Eds.), *Relational data mining* (pp. 262–291). Berlin: Springer.
- Krogl, M. A., & Wrobel, S. (2001). Transformation-based learning using multirelational aggregation. In *Proceedings of international conference on inductive logic programming*, pp. 142–155. Berlin: Springer.
- Krogl, M. A., Rawles, S., Železný, F., Flach, P., Lavrač, N., & Wrobel, S. (2003). Comparative evaluation of approaches to propositionalization. In T. Horvath & A. Yamamoto (Eds.), *Proceedings of the 13th international conference on inductive logic programming (ILP-2003)* (pp. 197–214). Berlin: Springer.
- Kuželka, O., & Železný, F. (2008). HiFi: Tractable propositionalization through hierarchical feature construction. In Železný, F., Lavrač, N. (Eds.) *Late breaking papers, the 18th international conference on inductive logic programming*, pp. 69–74.
- Kuželka, O., & Železný, F. (2011). Block-wise construction of tree-like relational features with monotone reducibility and redundancy. *Machine Learning*, 83(2), 163–192.
- Lachiche, N., & Flach, P. A. (2003). 1BC2: A true first-order Bayesian classifier. *Proceedings of inductive logic programming*, pp. 133–148.
- Lavrač, N., Džeroski, S., & Grobelnik, M. (1991). Learning nonrecursive definitions of relations with LINUS. In *Proceedings of the 5th European working session on learning (EWSL-91)*, pp. 265–281. Springer, Porto, Portugal.
- Lavrač, N., Kralj Novak, P., Mozetič, I., Podpečan, V., Motaln, H., Petek, M., & Gruden, K. (2009). Semantic subgroup discovery: Using ontologies in microarray data analysis. In *Proceedings of the 31st annual international conference of the IEEE EMBS*, pp. 5613–5616.
- Lavrač, N., & Džeroski, S. (1994). *Inductive logic programming: Techniques and applications*. New York: Ellis Horwood.
- Lavrač, N., & Flach, P. (2001). An extended transformation approach to inductive logic programming. *ACM Transactions on Computational Logic*, 2(4), 458–494.
- Le, Q., & Mikolov, T. (2014). Distributed representations of sentences and documents. In *Proceedings of international conference on machine learning*, pp. 1188–1196.
- Lewis, D. D. (1992). An evaluation of phrasal and clustered representations on a text categorization task. In *Proceedings of the 15th annual international ACM SIGIR conference on research and development in information retrieval*, pp. 37–50.
- Lodhi, H. (2013). Deep relational machines. In *Proceedings of the international conference on neural information processing*, pp. 212–219. Berlin: Springer.
- Lundberg, S. M., & Lee, S. I. (2017). A unified approach to interpreting model predictions. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R. (Eds.) *Advances in neural information processing systems*, pp. 4765–4774.
- McInnes, L., Healy, J., Saul, N., & Grossberger, L. (2018). UMAP: Uniform manifold approximation and projection. *The Journal of Open Source Software*, 3(29), 861.
- Mease, D., & Wyner, A. (2008). Evidence contrary to the statistical view of boosting. *Journal of Machine Learning Research*, 9, 131–156.
- Michalski, R. S., Mozetič, I., Hong, J., & Lavrač, N. (1986). The multi-purpose incremental learning system AQ15 and its testing application on three medical domains. In *Proceedings of the 5th national conference on artificial intelligence*, pp. 1041–1045. Philadelphia, PA.
- Michie, D., Muggleton, S., Page, D., & Srinivasan, A. (1994). To the international computing community: A new East-West challenge. Tech. rep., Oxford University Computing laboratory, Oxford, UK.
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., & Dean, J. (2013). Distributed representations of words and phrases and their compositionality. In C. J. C. Burges, L. Bottou, M. Welling, Z.

- Ghahramani, & K. Q. Weinberger (Eds.), *Advances in neural information processing systems 26* (pp. 3111–3119). New York, USA: Curran Associates Inc.
- Motl, J., & Schulte, O. (2015). The CTU Prague relational learning repository. [arXiv:1511.03086](https://arxiv.org/abs/1511.03086).
- Muggleton, S. H. (Ed.). (1992). *Inductive logic programming*. London: Academic Press Ltd.
- Muggleton, S. (1995). Inverse entailment and Prolog. *New Generation Computing*, 13(3–4), 245–286.
- Nickel, M., & Kiela, D. (2017). Poincaré embeddings for learning hierarchical representations. In *Advances in neural information processing systems*, pp. 6338–6347.
- Nickel, M., Tresp, V., & Kriegel, H. P. (2011). A three-way model for collective learning on multi-relational data. *Proceedings of International Conference on Machine Learning*, 11, 809–816.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., et al. (2011). Scikit-learn: Machine learning in python. *Journal of Machine Learning Research*, 12(Oct), 2825–2830.
- Perovšek, M., Vavpetič, A., Cestnik, B., & Lavrač, N. (2013). A wordification approach to relational data mining. In *Proceedings of the international conference on discovery science*, pp. 141–154. Berlin: Springer.
- Perovšek, M., Vavpetič, A., Kranjc, J., Cestnik, B., & Lavrač, N. (2015). Wordification: Propositionalization by unfolding relational data into bags of words. *Expert Systems with Applications*, 42(17–18), 6442–6456.
- Perozzi, B., Al-Rfou, R., & Skiena, S. (2014). Deepwalk: Online learning of social representations. In *Proceedings of the 20th ACM SIGKDD international conference on knowledge discovery and data mining*, pp. 701–710. ACM.
- Plantié, M., & Crampes, M. (2013). Survey on social community detection. In N. Ramzan, R. Zwol, J. S. Lee, K. Clüver, & X. S. Hua (Eds.), *Social media retrieval* (pp. 65–85). London: Springer.
- Podpečan, V., Lavrač, N., Mozetič, I., Kralj Novak, P., Trajkovski, I., Langohr, L., et al. (2011). SegMine workflows for semantic microarray data analysis in Orange4WS. *BMC Bioinformatics*, 12, 416.
- Qiu, J., Dong, Y., Ma, H., Li, J., Wang, K., & Tang, J. (2018). Network embedding as matrix factorization: Unifying DeepWalk, LINE, PTE, and Node2Vec. In *Proceedings of the eleventh ACM international conference on web search and data mining*, WSDM '18, pp. 459–467. ACM.
- Quinlan, J. R. (1986). Induction of decision trees. *Machine Learning*, 1(1), 81–106.
- Ribeiro, L. F., Saverese, P. H., & Figueiredo, D. R. (2017). Struc2vec: Learning node representations from structural identity. In *Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and data mining*, KDD '17, pp. 385–394. New York: ACM.
- Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). 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*, pp. 1135–1144. ACM.
- Ristoski, P., & Paulheim, H. (2016). Rdf2vec: Rdf graph embeddings for data mining. In P. Groth, E. Simperl, A. Gray, M. Sabou, M. Krötzsch, F. Lecue, F. Flöck, & Y. Gil (Eds.), *The semantic web: ISWC 2016* (pp. 498–514). Cham: Springer.
- Robnik-Šikonja, M., & Kononenko, I. (2008). Explaining classifications for individual instances. *IEEE Transactions on Knowledge and Data Engineering*, 20(5), 589–600.
- Rocktäschel, T., Singh, S., & Riedel, S. (2015). Injecting logical background knowledge into embeddings for relation extraction. In *Proceedings of the 2015 conference of the north American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 1119–1129.
- Rumelhart, D. E., & McClelland, J. L. (Eds.) (1986). Parallel distributed processing: Explorations in the microstructure of cognition, vol. 1: Foundations. MIT Press, Cambridge, MA.
- Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning representations by back-propagating errors. *Nature*, 323(6088), 533.
- Schapire, R. E., Freund, Y., Bartlett, P., & Lee, W. S. (1998). Boosting the margin: A new explanation for the effectiveness of voting methods. *The Annals of Statistics*, 26(5), 1651–1686.
- Schölkopf, B., & Smola, A. J. (2001). *Learning with kernels: Support vector machines, regularization, optimization, and beyond*. Cambridge: The MIT Press.
- Škrlić, B., Kralj, J., Konc, J., Robnik-Šikonja, M., & Lavrač, N. (2019). Deep node ranking: An algorithm for structural network embedding and end-to-end classification. [arXiv:1902.03964](https://arxiv.org/abs/1902.03964).
- Srinivasan, A. (2007). Aleph manual. <http://www.cs.ox.ac.uk/activities/machinelearning/Aleph/>.
- Srinivasan, A., King, R. D., Muggleton, S., & Sternberg, M. J. (1997). Carcinogenesis predictions using ILP. In *Proceedings of the international conference on inductive logic programming*, pp. 273–287. Berlin: Springer.
- Srinivasan, A., Vig, L., & Bain, M. (2019). Logical explanations for deep relational machines using relevance information. *Journal of Machine Learning Research*, 20(130), 1–47.

- Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., & Salakhutdinov, R. (2014). Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(1), 1929–1958.
- Štrumbelj, E., & Kononenko, I. (2014). Explaining prediction models and individual predictions with feature contributions. *Knowledge and Information Systems*, 41(3), 647–665.
- Tang, J., Qu, M., & Mei, Q. (2015a). PTE: Predictive text embedding through large-scale heterogeneous text networks. In *Proceedings of the 21th ACM SIGKDD international conference on knowledge discovery and data mining*, pp. 1165–1174. ACM.
- Tang, J., Qu, M., Wang, M., Zhang, M., Yan, J., & Mei, Q. (2015b). LINE: Large-scale information network embedding. In *Proceedings of the 24th international conference on world wide web*, pp. 1067–1077.
- Van Der Walt, S., Colbert, S. C., & Varoquaux, G. (2011). The NumPy array: A structure for efficient numerical computation. *Computing in Science & Engineering*, 13(2), 22.
- Vapnik, V. (1995). *The nature of statistical learning theory*. New York: Springer.
- Vavpetič, A., & Lavrač, N. (2011). Semantic data mining system g-SEGS. In *Proceedings of the workshop on planning to learn and service-oriented knowledge discovery (PlanSoKD-11), ECML PKDD conference*, pp. 17–29.
- Wang, Q., Wang, B., & Guo, L. (2015). Knowledge base completion using embeddings and rules. In *Proceedings of the 24th international joint conference on artificial intelligence*, pp. 1859–1865.
- Wang, Z., Zhang, J., Feng, J., & Chen, Z. (2014). Knowledge graph and text jointly embedding. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pp. 1591–1601.
- Wang, Q., Mao, Z., Wang, B., & Guo, L. (2017). Knowledge graph embedding: A survey of approaches and applications. *IEEE Transactions on Knowledge and Data Engineering*, 29(12), 2724–2743.
- Wu, L. Y., Fisch, A., Chopra, S., Adams, K., Bordes, A., & Weston, J. (2018). Starspace: Embed all the things! In *Proceedings of the 32nd AAAI conference on artificial intelligence*, pp. 5569–5577.
- Železný, F., & Lavrač, N. (2006). Propositionalization-based relational subgroup discovery with RSD. *Machine Learning*, 62, 33–63.
- Zhu, S., Bing, J., Min, X., Lin, C., & Zeng, X. (2018). Prediction of drug–gene interaction by using metapath2vec. *Frontiers in Genetics*, 9.
- Žitnik, M., & Leskovec, J. (2017). Predicting multicellular function through multi-layer tissue networks. *Bioinformatics*, 33(14), i190–i198.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Affiliations

Nada Lavrač^{1,2} · Blaž Škrjč³  · Marko Robnik-Šikonja⁴

Nada Lavrač
nada.lavrac@ijs.si

Marko Robnik-Šikonja
marko.robnik@fri.uni-lj.si

¹ Jožef Stefan Institute, Jamova 39, 1000 Ljubljana, Slovenia

² University of Nova Gorica, Glavni trg 8, 5271 Vipava, Slovenia

³ International Postgraduate School Jožef Stefan, Jožef Stefan Institute, Jamova 39, 1000 Ljubljana, Slovenia

⁴ University of Ljubljana, Faculty of Computer and Information Science, Večna pot 113, 1000 Ljubljana, Slovenia