

KG-MDL: Mining Graph Patterns in Knowledge Graphs with the MDL Principle

Francesco Bariatti ^{a,c,*}, Peggy Cellier ^b and Sébastien Ferré ^a

^a *Univ Rennes, CNRS, IRISA*

E-mails: francesco.bariatti@irisa.fr, sebastien.ferre@irisa.fr

^b *Univ Rennes, INSA, CNRS, IRISA*

E-mail: peggy.cellier@irisa.fr

^c *LIACS, Leiden University, Leiden, The Netherlands*

Abstract. Nowadays, increasingly more data are available as knowledge graphs (KGs). While this data model supports advanced reasoning and querying, they remain difficult to mine due to their size and complexity. Graph mining approaches can be used to extract patterns from KGs. However this presents two main issues. First, graph mining approaches tend to extract too many patterns for a human analyst to interpret (pattern explosion). Second, real-life KGs tend to differ from the graphs usually treated in graph mining: they are multigraphs, their vertex degrees tend to follow a power-law, and the way in which they model knowledge can produce spurious patterns. Recently, a graph mining approach named GRAPHMDL+ has been proposed to tackle the problem of pattern explosion, using the *Minimum Description Length* (MDL) principle. However, GRAPHMDL+, like other graph mining approaches, is not suited for KGs without adaptations.

In this paper we propose KG-MDL, a graph pattern mining approach based on the MDL principle that, given a KG, generates a *human-sized* and *descriptive* set of graph patterns, and so in a parameter-less and anytime way. We report on experiments on medium-sized KGs showing that our approach generates sets of patterns that are both small enough to be interpreted by humans and descriptive of the KG. We show that the extracted patterns highlight relevant characteristics of the data: both of the schema used to create the data, and of the concrete facts it contains. We also discuss the issues related to mining graph patterns on knowledge graphs, as opposed to other types of graph data.

Keywords: Knowledge Graph, Pattern Mining, Graph Mining, Minimum Description Length

1. Introduction

With the growing amount and diversity of data available as knowledge graphs (KGs), there is an increasing need for analyzing them and making sense of their contents. A first layer of analysis is provided by schemas (RDFS), ontologies (OWL), and constraints (SHACL and ShEx). These are logical statements expressing absolute truths about the data and usually need to be produced by human experts. In this paper we are interested in knowledge intermediate between raw facts and logical statements, namely *frequent patterns*, i.e. recurring structures in data. Such frequent structures can be indicators of regularities in the data, which may in turn be expressed as schemas or constraints, allowing the user to improve the modeling of the KG; or can help optimising query evaluation by highlighting different archetypes of entities [1]. Also, since patterns are extracted from raw facts, they can highlight the presence of errors in the data or in the modeling when they do not conform with the expected schema.

*Corresponding author.

1 *Data mining* is the process of discovering interesting patterns in data [2]. First applied to transactional data, where
2 transactions and patterns are sets of items, it was later extended to sequential data (e.g., logs, texts) [3], and to graph
3 data (e.g., molecules, text corpora) [4]. *Graph pattern mining* was found much more difficult than for transactions
4 and sequences, due to the intrinsic complexity of graph algorithms.

5 Existing graph mining algorithms cannot easily be applied to KGs because they were designed for very different
6 kinds of graphs: large collections of small undirected simple graphs (e.g. molecules). KGs, on the contrary, generally
7 have a single connected component, and they are directed multi-graphs. Another major difference is that in KGs,
8 node degrees tend to follow a power law [5]. All those differences make graph pattern mining in KGs a very
9 challenging problem that has not been sufficiently addressed so far. Many existing works applying data mining to
10 KGs have addressed the problem by converting graph data into transactions in order to lower the complexity [6–8].
11 A few approaches have been proposed to mine rules from KGs [9, 10]. These are relevant because the rules use a
12 graph pattern as the body (and generally a single triple pattern as the head). However, the shape of graph patterns is
13 strongly limited, generally to chain patterns.

14 Pattern mining—and in particular graph mining—generally suffers from the large number of extracted patterns.
15 Several kinds of approaches have been proposed to tackle this problem (e.g., constraints or condensed representa-
16 tions) [4]. Among them, the methods based on the *Minimum Description Length* (MDL) principle showed that it is
17 possible to drastically reduce the number of patterns by selecting a small set of descriptive patterns [11]. The MDL
18 principle [12, 13] comes from the domain of information theory, and states that the *model* that describes the data
19 the best is the one that compresses the data the best, i.e. compression is used as a metric to evaluate models, with
20 sets of patterns used as models in pattern mining. Several methods based on the MDL principle have been proposed
21 to tackle the number of patterns issue for graph mining [14–18]. However, among them only GRAPHMDL+ [18]
22 allows to generate patterns without restriction on their shape. Even so, none of them is directly applicable to KGs.
23 To our knowledge, only one MDL-based method, called KGist [19], has been proposed to target KGs specifically,
24 extracting sets of rules describing the data. However, the shapes of the rules are restricted to rooted trees.

25 The contribution of this paper is threefold. First, we propose KG-MDL, a MDL-based method to extract a *human-*
26 *sized* set of *descriptive* patterns from KGs and labeled multi-graphs in general. It is based on the GRAPHMDL+
27 approach, but its MDL-related definitions have been revisited to adapt to KGs. The approach is parameter-free
28 and anytime, and does not put any restriction on the shape of the extracted patterns. Second, we discuss important
29 distinctions between KGs and the graphs that are usually encountered in graph mining. Even though both are called
30 “graphs”, KGs have many peculiarities that must be overcome in order to bring the two domains together. Finally,
31 we evaluate our approach experimentally on medium-sized KGs, showing that it is effective at extracting a small
32 amount of descriptive patterns. We also show that the extracted patterns provide useful insights about both the
33 schema used to create the KG, and the raw facts contained in it. Because of that, they can be a useful asset for
34 exploring unknown KGs, refining ontologies, and introducing constraints.

35 The rest of this paper is structured as follows: Section 2 presents related work on graph mining and KG mining.
36 Section 3 presents preliminary notions and the notations that we use in the rest of the paper. In Section 4 we present
37 our contributions, and in Section 5 we evaluate it on real-life datasets both from a quantitative and qualitative point
38 of view. Section 6 concludes.

41 2. Related Work

42
43 *Graph mining.* The first graph mining approaches that have been proposed [4] were *complete approaches*, i.e.
44 with the aim to extract *all* possible patterns that follow a certain acceptance rule, usually a minimum frequency
45 of appearance (e.g. gSpan [20], Gaston [21]). A common drawback of this kind of approaches is that they extract
46 a large amount of patterns, which makes human analysis difficult. Approaches have been proposed to reduce the
47 amount of extracted patterns using constraints or condensed representations, with mitigated results [4].

48 More recently, graph mining approaches have been proposed that use the *Minimum Description Length* (MDL)
49 principle [12, 13] to more drastically reduce the amount of extracted patterns, and to obtain graph patterns that are
50 descriptive of the data. In this family of methods, Subdue [14] can extract patterns with any shape, but its encoding
51 entails a loss of information about the graph structure. It is possible to know which patterns appear in the data, but

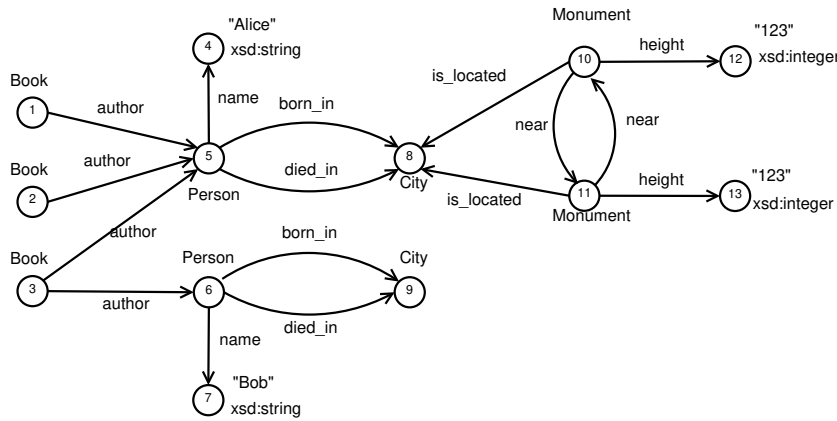


Fig. 1. An example of a directed multi-graph, used as example data throughout the paper.

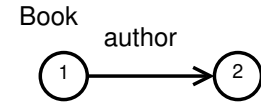


Fig. 2. An example pattern which has four occurrences in the graph of Fig 1.

not *how they connect to each other*, which is a major disadvantage for analysis. VoG and MeGS [16, 17] are methods targeted at summarizing unlabeled graphs that have no such loss of information, but only consider a pre-defined vocabulary of pattern shapes. While this vocabulary can be expanded, doing so requires to have prior knowledge about the type of structures that can be observed in the data. The approach proposed in [15] allows to extract *motifs* from unlabeled graphs, which are similar in the idea to patterns. This approach interprets MDL encodings as probability distributions, which are needed to select motifs. However it only applies to unlabeled graphs, and only when an *a priori hypothesis* about their shape exists. Finally, GRAPHMDL+ [18] can mine patterns with arbitrary shapes from both directed and undirected graphs. The approach introduces the notion of “ports” to avoid loss of information about the graph structure. However, all these approaches can not be applied directly to KG mining. This is in part because the type of graphs that they support does not encompass labeled multi-graphs (which knowledge graphs are), and in part because —as we detail in Section 4.1— KGs have peculiarities that differentiates them from the graphs found in a classic graph mining context, namely their highly skewed degree distribution and the presence of connections between semantically unrelated elements.

Knowledge graph mining. Some data mining approaches have been proposed to target knowledge graphs specifically. Some approaches such as SWApriori [6], SWARM [7] and the one proposed in [8] convert the KG to a set of transactions to lower the complexity of the task, and then apply classic itemset-mining algorithms to it. However, this conversion limits the shape of the patterns that can be extracted. Most other KG mining approaches, such as AMIE [9], AnyBURL [10], RuDiK [22], and KGist [19] (which also uses the MDL principle) mine association rules whose bodies are graph patterns. However, in these approaches the shape of the patterns is also limited, often to rooted trees, or requires a minimum connectivity on the vertices in the pattern. Another common drawback of KG mining methods is that they are often developed for (and evaluated on) a specific task (e.g. link prediction, error correction), and do not perform other tasks, in contrast with the open-minded goal of data exploration that is central to data mining [23].

3. Preliminaries

In this section we define some notions that are used extensively in the rest of the paper.

3.1. Graphs

In this paper, we define *graphs* using the following structure:

Definition 1. A graph $G = (V, V_{\mathcal{L}}, E_{\mathcal{L}})$ over a set of label symbols \mathcal{L} is a data structure composed of a set of vertices V , a set of vertex labels $V_{\mathcal{L}} \subseteq V \times \mathcal{L}$, and a set of edge labels —or labeled edges— $E_{\mathcal{L}} \subseteq V \times V \times \mathcal{L}$.

This definition encompasses all graphs that are usually found in the literature: it allows for both directed and undirected graphs (depending on whether $(u, v, l) \in E_{\mathcal{L}}$ is equivalent to $(v, u, l) \in E_{\mathcal{L}}$ or not); it allows vertices to have any number of labels; and finally it allows for multi-graphs, i.e. graphs where multiple edges with different labels can exist between two vertices. Note that in Section 4.1 we discuss how to represent KGs (defined as set of triples) with this definition, and important considerations to take into account in doing so when targeting a graph mining context. Fig. 1 shows an example graph, with 13 vertices (named 1 to 13), 17 vertex labels, and 16 edge labels, from a set of 16 label symbols. This graph will be used as example data throughout this paper. Note that many existing graph mining approaches are limited in the kind of graphs that they can accept. For example, most accept only simple graphs, i.e. graphs that contain no self loops (edges from a vertex to itself) and a maximum of one edge between any pair of vertices. Many also require vertices to have exactly one label.

Definition 2 (Pattern occurrence). *Let G^P (the “pattern”) and G^D (the “data”) be graphs. An occurrence—or embedding—of G^P in G^D is an injective function $\varepsilon \in V^P \rightarrow V^D$ such that: (1) $(v, l) \in E_{\mathcal{L}}^P \implies (\varepsilon(v), l) \in E_{\mathcal{L}}^D$, (2) $(v_1, v_2, l) \in E_{\mathcal{L}}^P \implies (\varepsilon(v_1), \varepsilon(v_2), l) \in E_{\mathcal{L}}^D$.*

Informally, a pattern occurrence is a function that maps all vertices of the “pattern” to some of the vertices of the “data”. It conveys the information that the structure of the pattern can be found at some specific point in the data. Note that this definition represents the occurrences under *isomorphism*, meaning that each pattern vertex must be mapped to a *different* data vertex. This is the definition commonly used in graph mining, whereas the SPARQL query language uses *homomorphism* (i.e. two different variables in a SPARQL graph pattern can be bound to the same RDF node)¹. Fig. 2 shows an example pattern that has four occurrences in the graph of Fig. 1: three occurrences mapping vertex 1 of the pattern to respectively vertices 1, 2, or 3 of the data and vertex 2 of the pattern to vertex 5 of the data; and one occurrence mapping vertex 1 of the pattern to vertex 3 of the data and vertex 2 to vertex 6.

3.2. The MDL Principle

The *Minimum Description Length* (MDL) principle [12, 13] comes from the domain of information theory. It allows to identify the model—from a family of models—that best describes some data. It has been successfully used in pattern mining to select a small and descriptive set of patterns on some data, by considering sets of patterns as possible models for the data [11].

Definition 3 (MDL principle). *Given a family of models \mathcal{M} and some data D , the best model $M \in \mathcal{M}$ to describe the data is the one minimizing the description length $L(D, M) = L(M) + L(D|M)$, where $L(M)$ is the description length of the model, and $L(D|M)$ is the description length of the data encoded with the model.*

The MDL principle is generic for any kind of data and models. Each MDL-based approach must devise its own definitions for the different description lengths. However, a few primitives are commonly used [24]. We list here the ones used in this paper (note that in the whole paper, \log indicates the base 2 logarithm):

- An element x chosen with uniform probability from a set \mathcal{X} has a description length of $\log(|\mathcal{X}|)$ bits.
- k elements chosen with uniform probability from a set \mathcal{X} have a description length of $\log\binom{|\mathcal{X}|}{k}$ bits.
- An element $x \in \mathcal{X}$ appearing $\#x$ times in data has a description length of $L_{\mathcal{X}}(x) = -\log\left(\frac{\#x}{\sum_{x_i \in \mathcal{X}} \#x_i}\right)$ bits. This is often called a *prefix code* [12, 24].
- An integer $n \in \mathbb{N}$ with no known upper bound has a description length of $L_{\mathbb{N}}(n)$ bits, where $L_{\mathbb{N}}$ is a *universal integer encoding* [25]².
- A sequence S of elements drawn from a set \mathcal{X} has a description length of $L_{preq}(\mathcal{X}, S)$, where L_{preq} is a *prequential plugin code* [12] (see Appendix B for details)³.

¹It is possible to use isomorphism in SPARQL graph patterns by adding a FILTER clause expressing that all variables must be pairwise distinct.

²In our implementation we use the *Elias delta encoding*, shifted by 1 so that it can encode 0.

³In our implementation we initialize the usage of elements in the prequential encoding to $\epsilon = 0.5$, according to literature.

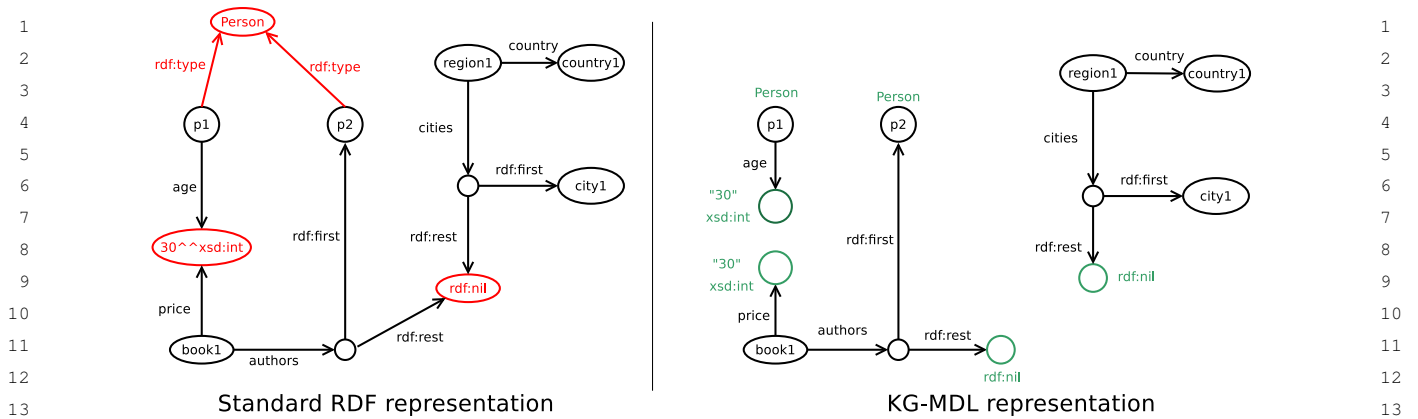


Fig. 3. Left: a KG represented using the usual RDF graph representation. Right: the same KG represented following KG-MDL’s mapping. The differences between the two representations are emphasized with colours.

The act of associating a description length to an element is called *encoding*. Note that in MDL-based approaches what matter are the *theoretical* description lengths as stated above, and the actual encodings do not need to be materialised [12]. In particular, it is common and acceptable for description lengths to be non-integer amounts. However, the MDL principle still requires encodings to be (at least theoretically) reversible: meaning that no information should be lost. This is the case for all primitives listed above.

4. KG-MDL

In this section, we formally define our approach, KG-MDL, which is an adaptation of GRAPHMDL+ [18] to KGs. First in Section 4.1 we discuss the need for a pre-processing stage before a KG can be used in a pattern mining task, showing that the standard graph representation of KGs presents some issues in a graph mining context, and we propose an alternative (and reversible) representation. In Section 4.2 we present the MDL encoding and related description lengths that are used in KG-MDL in order to evaluate sets of patterns. Then, in Section 4.3 we present the search algorithm used to search for the best set of patterns according to the MDL criterion. Finally, in Section 4.4 we discuss some issues that arise when applying our approach to real-life KGs, due to the differences between those KGs and the graphs usually found in the graph mining domain, and we present our solution to avoid them.

4.1. Knowledge Graphs Representation for Pattern Mining

The RDF standard defines a way to convert a KG expressed as a collection of triples to a graph representation. It consists in creating a vertex for each element appearing as subject or object, and adding an edge for each triple — labeled with the triple’s predicate— connecting the vertex corresponding to the subject to the vertex corresponding to the object. An example of such a representation is shown on the left of Fig. 3.

While we could use this representation directly (as the graph definition given in Section 3.1 covers it), we show here that this standard representation is not suited when the resulting graph is to be used in a graph mining context. A first problem that arises with this representation is about literals: in Fig. 3 we can see that in the standard representation a single vertex is used for all instances of the integer literal 30. As a consequence, vertices `book1` and `p1` are neighbours in the graph, even though the value “30” represents a price for the former and an age for the latter. This is not a problem when answering queries (hence why the standard uses this representation), but becomes a problem if the graph is used to generate patterns: we would see spurious patterns such as “a book that costs as many dollars as a person is old in years”, which relate entities just because they happen to have the same numerical value for semantically different quantities. An equivalent problem happens with the `rdf:nil` entity, which is used to mark the end of *all* lists. If not treated carefully, patterns can be generated relating semantically unrelated lists, just because of them being lists. Finally, `rdf:type` edges can also generate spurious patterns by making all enti-

	G^P	$ \Pi_P $	Π_P
	Pattern structure	Port count	Port IDs
$P1$		1	{1}
$P2$		3	{1,2,3}
$P3$		2	{1,2}
$P4$		1	{1}
$P5$		1	{1}

Fig. 4. A KG-MDL's code table.

ties of the same type neighbours in the graph. In graph mining, attributes such as the type of an entity are instead represented as vertex labels.

In conclusion, before mining patterns, KG-MDL converts the data KG to what is shown on the right of Fig 3: the object of triples with predicate `rdf:type` is added as a label to the vertex corresponding to the triple's subject; for each triple having a literal object, a new *distinct* vertex is created for the literal, labeled with the literal's type and value (and language for language-tagged strings); for each triple having `rdf:nil` as subject, a new distinct vertex is created, labeled `rdf:nil`. For other triples, an edge is added between the vertices corresponding to the subject and object. Note that this conversion is reversible: it is always possible to convert the graph back to the original KG⁴. In particular, the patterns extracted from the graph can be converted to SPARQL queries that —when executed on the original KG— will give mappings corresponding to the patterns' occurrences. For example, the pattern of Fig. 2 can be converted to `SELECT ?x1 ?x2 WHERE { ?x1 a <Book> ; <author> ?x2 .}`.

4.2. MDL Encoding for Knowledge Graphs

In this section we present the MDL encoding used by our approach for computing the MDL description length (see Section 3.2) needed to compare possible sets of patterns. This encoding is based on the one used by GRAPHMDL+ [18], from which it adopts the same high-level structure. However, since GRAPHMDL+ considers both the data and the patterns to be simple graphs —i.e. graphs that have no edges from a vertex to itself and no more than one edge between each pair of vertices— its encoding does not support the full expressiveness of knowledge graphs. Instead, new description lengths must be defined.

The encoding used by our approach is composed of two parts: the *code table* (corresponding to MDL's "model") and the *rewritten graph* (corresponding to MDL's "data encoded with the model").

4.2.1. The Model: Code Table

KG-MDL's model is called a *code table* (CT): it is a collection of patterns, with some associated information that will be used when encoding the data with those patterns. Fig. 4 shows an example code table that can be used to encode the graph of Fig. 1. Each row of the code table is a tuple $P = (G^P, \Pi_P)$ representing the characteristics of a pattern: its *graph structure* $G^P = (V^P, V_L^P, E_L^P)$ and its set of *ports* $\Pi_P \subseteq V^P$. The notion of ports, first introduced in [26], is key for encoding graph data with graph patterns. The ports of a pattern are the vertices that the pattern's

⁴On the condition of storing which labels were originally types and which were literal datatypes and values, which is an easy task.

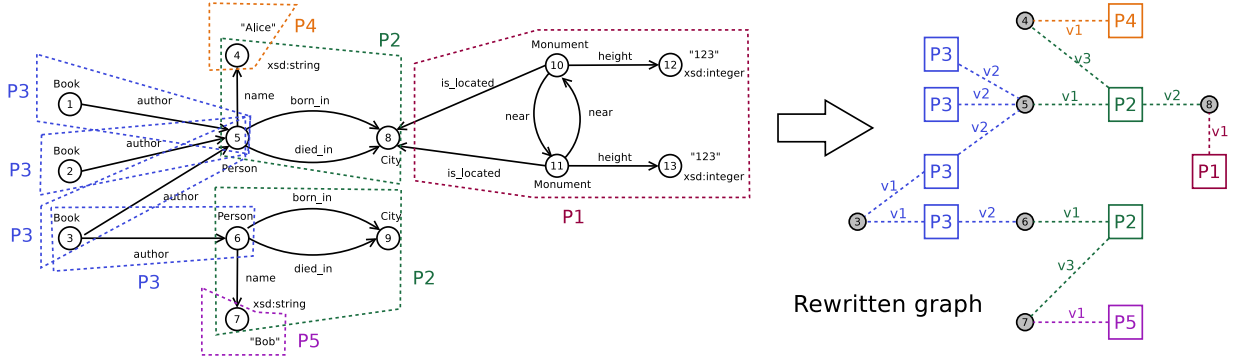


Fig. 5. How the patterns of the code table of Fig. 4 can be used to encode the data graph of Fig. 1, and the resulting rewritten graph.

occurrences share with other pattern occurrences when the CT is used to encode the data (see below). For example, the first row of Fig. 4 represents the pattern “two monuments of height 123 that are near to each other and located in the same place”. This pattern has only one port, its vertex 1, meaning that when the CT is used to encode the data, occurrences of this pattern can only be connected to other pattern occurrences by overlapping on their vertex 1.

4.2.2. The Encoded Data: Rewritten Graph

In KG-MDL, the “data encoded with the model” is represented by what is called a *rewritten graph*. The rewritten graph —first introduced in [26]— is a data structure allowing to represent a data graph as a composition of pattern occurrences, without any loss of information about the structure of the data. Fig. 5 shows a rewritten graph corresponding to the encoding of the data graph of Fig. 1 with the patterns in the code table of Fig. 4.

Definition 4 (from [26], with adapted notation). A rewritten graph $G^R = ((V_{emb}^R, V_{port}^R), V_{\mathcal{L}}^R, E_{\mathcal{L}}^R)$ is a special graph where: V_{emb}^R is the set of embedding vertices, V_{port}^R is the set of port vertices, $V_{\mathcal{L}}^R = V_{emb}^R \rightarrow CT$ a function associating CT patterns to embedding vertices, and $E_{\mathcal{L}}^R \subset V_{emb}^R \times V_{port}^R \times \bigcup_{P \in CT} \Pi_P$ a set of edges between embedding vertices and port vertices, labeled with pattern ports. Such edges are subject to the condition $(v_e, v_p, \pi) \in E_{\mathcal{L}}^R \implies \pi \in \Pi_P$ with $P = V_{\mathcal{L}}^R(v_e)$, i.e. edges can only be labeled with port identifiers for the pattern associated to their embedding vertex.

The rewritten graph is composed of two disjoint sets of vertices: each *embedding vertex* $v_e \in V_{emb}^R$ (squares in Fig. 5) represents an occurrence of a pattern of the CT, which describes part of the data graph. Each embedding vertex has one *label* telling to which pattern $P \in CT$ it corresponds. *Port vertices* $v_p \in V_{port}^R$ (circles in Fig. 5) represent those data vertices which are “covered” by at least two pattern occurrences associated to embedding vertices. Finally, *edges* $(v_e, v_p, \pi) \in E_{\mathcal{L}}^R$ can connect embedding vertices to port vertices, telling that a specific embedding vertex uses a specific port vertex. These edges have a label $\pi \in \Pi_P$ —where Π_P are the possible ports of the pattern associated to v_e — telling which of the pattern’s vertices is mapped to the data vertex corresponding to v_p . For example, in the top-left of Fig. 5 we see that part of the data (around vertices 1 and 5) is described as an occurrence of pattern P_3 , which shares its second vertex (vertex 5 of the data) with other pattern occurrences. In the rewritten graph, this is represented with an embedding vertex labeled P_3 connected with an edge to a port vertex that represents vertex 5 of the data, and the edge is labeled v_2 to indicate that it corresponds to vertex 2 of the pattern. Other pattern occurrences that also use the same data vertex are represented as embedding vertices also connected to the same port vertex. Data vertices that are not shared between pattern occurrences do not appear in the rewritten graph.

The rewritten graph has two uses: firstly it is needed for applying the MDL principle, as the “data encoded with the model” is part of the MDL description length. But it can also be given to the user as a companion to the selected set of patterns: in this way the user can not only analyze the selected patterns, but also how they connect and interact with each other. For example, in Fig. 5 there are four occurrences of pattern P_3 (“a book has an author who is a person”), and for three of them their vertex 2 is mapped to the same port vertex, surfacing the knowledge that most of the books share the same author in the data graph.

Code table description length

$$L(M) = L(CT) = \sum_{P \in CT} \overbrace{L(G^P)}^{\text{Pattern structure}} + \overbrace{\log(|V^P| + 1)}^{\text{port count } |\Pi_P|} + \overbrace{\log\left(\frac{|V^P|}{|\Pi_P|}\right)}^{\text{ports } \Pi_P}$$

$$L(G^P) = \underbrace{L_{\mathbb{N}}(|V^P|)}_{\text{vertex count}} + \underbrace{\log(|\mathcal{L}|)}_{\text{label symbols count}} + \sum_{\{l \in \mathcal{L} \mid l \text{ is in } P\}} \left[\underbrace{L_{\mathcal{L}}(l)}_{\text{symbol identity}} + \underbrace{L_{\mathbb{N}}(\#occ)}_{\text{occurrences count}} + \underbrace{\log\left(\frac{|V^P|^n}{\#occ}\right)}_{\text{occurrences}} \right]$$

Rewritten graph description length

$$L(D|M) = L(G^R) = \underbrace{L_{\mathbb{N}}(|V_{emb}^R|)}_{V_{emb}^R} + \underbrace{\log(|V^D| + 1)}_{V_{port}^R} + \underbrace{L_{preq}(CT, (-, P) \in V_{\mathcal{L}}^R)}_{V_{\mathcal{L}}^R} +$$

$$\sum_{P \in CT} \left[\underbrace{\sum_{v_e \in \{E_{\mathcal{L}}^R|P\}} \log(|\Pi_P| + 1)}_{\text{edge sources}} + \underbrace{L_{preq}(\Pi_P, (-, -, \pi) \in \{E_{\mathcal{L}}^R|P\})}_{\text{edge labels}} \right] + \underbrace{L_{preq}(V_{port}^R, (-, v_p, -) \in E_{\mathcal{L}}^R)}_{\text{edge destinations}}$$

Fig. 6. MDL description length used by KG-MDL to evaluate sets of patterns.

4.2.3. Description Lengths

In this section we define the MDL description length that KG-MDL uses to compare different sets of patterns. The MDL description length is composed of two parts (see Section 3.2): the model length $L(M)$ and the encoded data length $L(D|M)$. In KG-MDL the model is a code table, and the encoded data is a rewritten graph. All formulas are shown in Fig. 6.

In this section we use the following shorthand notations: 1) $\{E_{\mathcal{L}}^R|P\} = \{(v_e, v_p, \pi) \in E_{\mathcal{L}}^R \mid V_{\mathcal{L}}^R(v_e) = P\}$ is the set of rewritten graph edges that start from an embedding vertex of pattern P , 2) for a set of tuples $X = \{(a_1, b_1), \dots\}$ we write $(-, b) \in X$ to mean the *sequence* composed of all values b_i appearing in X , in an arbitrary but determined order. This notation can be generalised to tuples of any size and to any of their components.

Code table description length. The description length of a code table corresponds to the sum of the description length of each of its rows. For each row the information that needs to be encoded is: the **pattern's structure** G^P , and the **pattern's ports** Π_P .

In order to encode the **pattern's structure**, we first need to encode the number of vertices $|V^P|$ of the pattern. Since a pattern can have any number of vertices, we use a universal integer encoding $L_{\mathbb{N}}$ for this information. Then, we describe all the labels that appear in the pattern. For that, we first encode how many different label symbols can be found in the pattern: since a pattern can contain between one⁵ and as many different symbols as there are in the data (i.e. \mathcal{L}), this information can be encoded with $\log(|\mathcal{L}|)$ bits. Then, for each label symbol present in the pattern, we encode: its identity, how many occurrences it has in the pattern, and which ones they are. In order to indicate which label symbol $l \in \mathcal{L}$ we are describing, we use a prefix code $L_{\mathcal{L}}(l)$ based on the number of occurrences that each

⁵A pattern with zero labels would not make sense, as it would correspond to a structure with no edges and no vertex labels.

label has in the data graph. Then, we encode how many occurrences ($\#occ$) this symbol has in the pattern with a universal integer encoding. Finally, we indicate what those occurrences are: in a pattern with $|V^P|$ vertices, a vertex label has $|V^P|$ possible occurrences, and an edge label has $|V^P|^2$. Thus, we can encode them using $\log\left(\binom{|V^P|^n}{\#occ}\right)$ bits, with n either 1 or 2 depending on the nature of the label symbol. Note that this encoding does not put any limit on the shape of a pattern, and can encode any graph structure.

Once the pattern's structure has been encoded, encoding the **pattern's ports** is simple, as the ports are a subset of the pattern's vertices. First, the number of ports $|\Pi_P|$ (between 0 and all vertices) is encoded with $\log(|V^P| + 1)$ bits, then the choice of which of the vertices are ports is encoded with $\log\left(\binom{|V^P|}{|\Pi_P|}\right)$ bits.

Rewritten graph description length. The description length of a rewritten graph corresponds to the sum of the description length of each of its elements: **the vertices**, **the vertex labels**, and **the edges**.

In order to encode **the vertices** of the rewritten graph it suffices to encode how many embedding vertices and how many port vertices there are. The number of embedding vertices can be anything and is thus encoded with a universal integer encoding, whereas the number of ports is bounded by the number of vertices in the data (at maximum patterns can superpose on all of them) and can therefore be encoded with $\log(|V^D| + 1)$ bits. For the example of Fig. 5, this corresponds to $L_{\mathbb{N}}(9) + \log(14)$.

In order to encode the **vertex labels**, since we encoded the set of embedding vertices and each one of them has one label, we can use a prequential code with elements drawn from the set of patterns in the CT. For the example of Fig. 5, this corresponds to $L_{preq}(\{P_1, P_2, P_3, P_4, P_5\}, \langle P_3, P_3, P_3, P_3, P_4, P_2, P_2, P_5, P_1 \rangle)$

We encode **the edges** $(v_e, v_p, \pi) \in E_{\mathcal{L}}^R$ one element at a time. First, for each pattern we encode the number of edges starting from each of its embedding vertices v_e . Since an embedding vertex can use at maximum all ports Π_P of its pattern (and at minimum 0), its number of edges can be encoded with $\log(|\Pi_P| + 1)$ bits. Then, for each pattern we encode the label π of all edges of its embedding vertices (whose number and sources have been encoded in the previous step), which can be done using a prequential code with elements drawn from the set of ports Π_P of the pattern. For example, for pattern P_3 of Fig. 5, this corresponds to $4 \log(3) + L_{preq}(\{v_1, v_2\}, \langle v_2, v_2, v_2, v_1, v_1, v_2 \rangle)$. Finally, the destinations v_p of all edges in the rewritten graph can be encoded using a prequential code with elements drawn from the set of port vertices V_{port}^R . For the example of Fig. 5, this corresponds to $L_{preq}(\{3, 4, 5, 6, 7, 8\}, \langle 5, 5, 5, 3, 3, 6, 4, 4, 5, 8, 6, 7, 7, 8 \rangle)$.

In total, the the description length of the rewritten graph of Fig. 5 is given by:

$$\begin{aligned} & L_{\mathbb{N}}(9) + \log(14) + L_{preq}(\{P_1, P_2, P_3, P_4, P_5\}, \langle P_3, P_3, P_3, P_3, P_4, P_2, P_2, P_5, P_1 \rangle) \\ & + \underbrace{4 \log(3) + L_{preq}(\{v_1, v_2\}, \langle v_2, v_2, v_2, v_1, v_1, v_2 \rangle)}_{P_3} + \underbrace{\log(2) + L_{preq}(\{v_1\}, \langle v_1 \rangle)}_{P_4} + \underbrace{\dots}_{P_2, P_5, P_1} \\ & + L_{preq}(\{3, 4, 5, 6, 7, 8\}, \langle 5, 5, 5, 3, 3, 6, 4, 4, 5, 8, 6, 7, 7, 8 \rangle) \end{aligned}$$

4.3. Search Algorithm

The algorithm that KG-MDL uses in order to search for the best possible code table is shown in Alg. 1. It is the same search algorithm used by GRAPHMDL+ [18]. The algorithm starts (line 1) with an initial code table CT_0 composed of singleton patterns only, i.e. patterns having either a single vertex with a single label, or a single edge between unlabeled vertices. Then it proceeds iteratively: it computes the rewritten graph for the current code table (line 5), and looks at which patterns appear sharing ports in the rewritten graph. Each pair of embedding vertices that are neighbours in the rewritten graph generates a *candidate pattern* corresponding to the merge of the patterns of the two embedding vertices, along their shared ports (line 6). Such candidates are ordered following a heuristic (line 7), and then one by one the algorithm tries to add them to the current code table (lines 8 and 9): if one is found where the description length of the new code table is smaller than the previous one (which in MDL terms means it is better), a new iteration is started with this new code table (lines 10 and 11). If no candidate gives a better code table the algorithm stops, returning the best code table found (line 14).

This algorithm is parameterless and anytime: it requires only the data graph as input, and when it is stopped it can return the best code table found so far. Patterns are generated bottom-up as merges of smaller patterns. The only

Algorithm 1 Search algorithm used by KG-MDL.

Require: A data graph D

1: $CT_0 \leftarrow$ initial singleton-only code table from labels in D

2: **return** ITERATION(D, CT_0)

3:

4: **function** ITERATION(D, CT)

5: $G^R \leftarrow$ rewritten graph from D and CT

6: $\mathcal{C} \leftarrow$ candidates generated from pairs of embedding vertices in G^R

7: $\mathcal{C}_{ranked} \leftarrow \mathcal{C}$ sorted by heuristic

▷ See [18]

8: **for all** $C \in \mathcal{C}_{ranked}$ **do**

9: $CT' \leftarrow CT \cup C$

10: **if** $L(D, CT') < L(D, CT)$ **then**

11: **return** ITERATION(D, CT')

12: **end if**

13: **end for**

14: **return** CT

15: **end function**

difference between the algorithm used by KG-MDL and the one used by GRAPHMDL+ resides in the formulas used to compute the MDL description length $L(D, CT)$: KG-MDL uses the formulas introduced in Section 4.2.3, which can handle multigraphs (and thus KGs).

4.4. When Theory Meets Practice

When applying KG-MDL on real-life KGs we observed a high memory usage and a high runtime that were not observed when applying GRAPHMDL+ to classical graph mining datasets. Based on our analysis, those problems arose because the basic implementation of the search algorithm would compute for each pattern in the CT all of its occurrences in the data, and go through them all in order to compute the rewritten graph. In real-life KGs some pattern emerged that have an exponentially high number of possible occurrences. An example of such patterns is “N elements located in the same country”: if the data graph contains hundreds of countries with hundreds of elements each, the number of possible occurrences of this pattern is exponentially high. This type of structure seems very typical of KGs, which often have predicates used to “list” elements, which results in some vertices with a particularly high number of edges around them, following a power-law [5]. In contrast, in the datasets usually used in graph mining there is often an upper bound on the degree of a vertex: e.g. when mining molecules, the maximum number of edges is the maximum number of bonds for an atom, which is a small value. This is yet another difference between KGs and classic graph mining datasets, which must be catered for when applying data mining techniques to KGs.

Our solution to limit this problem has two parts. First, we avoid computing (and storing) *all* occurrences of each pattern, as only few of them are required to create the rewritten graph. Indeed, while the rewritten graph allows pattern to overlap on vertices, it forbids edge overlaps [26], therefore a large number of possible occurrences can be discarded. Second, we add the possibility to specify a maximum time to spend on each code table row when creating the rewritten graph. We call this parameter the *row cover timeout*. This allows to enforce that a fairer share of time is allocated to all patterns, avoiding that some complex patterns slow down the whole computation. This is important in the context of an anytime approach such as KG-MDL, since time is an important resource and as such it is crucial to avoid spending too much time on a small part of the search space. We evaluate this solution experimentally in Section 5.

5. Experimental Evaluation

In this section we evaluate KG-MDL on five real-life datasets. We first evaluate the approach quantitatively, and show that KG-MDL is able to extract a small set of patterns which covers the data efficiently, and minimize

Table 1
Characteristics of the datasets used in the experiments.

Dataset	#triples	KG-MDL graph modeling			
		$ V $	$ V_{\mathcal{L}} $	$ E_{\mathcal{L}} $	$ \mathcal{L} $
SemanticBible-NTNames	4k	2k	5k	3k	1k
Lemon-dbpedia-en	30k	14k	19k	17k	2k
Mondial-europe	50k	22k	28k	34k	3k
Taaable	138k	82k	140k	115k	17k
Mondial	168k	74k	95k	114k	8k

the MDL description length. Then, we present some of the extracted patterns, showing that they convey relevant information about both the schema of the data, and the concrete facts contained in the data.

5.1. Settings and Datasets

In order to evaluate KG-MDL we developed a prototype in Java, which is available as a git repository⁶. Experiments were run on a single core (our prototype is not multi-threaded) of an Intel Xeon E5-2630 v3 CPU, and had a RAM limit of 32GB.

We conducted the experiments on five datasets. The **Lemon-dbpedia-en** dataset⁷ is a lexicalisation of the most common DBpedia classes and properties using the “Lemon” ontology. It contains information such as “*The term island is a common noun whose sense is the `dbo: Island` class*”. In this dataset —because of the way it is used— we considered the `partOfSpeech` predicate equivalent to the `rdf:type` predicate. The **Mondial** dataset⁸ contains geographical information, such as “*Mt Everest is a mountain in the Himalayas*”. The **Mondial-europe** dataset is the subset covering Europe only. Since these two datasets contain a large number of numerical values, we removed the concrete values and encoded them with a vertex labeled by the type of the value only (e.g. `xsd: Integer`). The **SemanticBible-NTNames** dataset⁹ is a semantic web representation of named entities in the Bible with associated information, such as “*Adam is a man, spouse of Eve, parent of Abel, Cain and Seth*”. The **Taaable** dataset¹⁰ corresponds to an RDF representation of the Wikitaable semantic wiki [27], containing recipe and food information such as “*Bread is a baked good made using 2 teaspoons of salt*”.

The characteristics of the datasets are presented in Table 1. For each dataset we report the number of triples in its original RDF representation and the number of vertices, vertex labels, edges, and the total number of label symbols in their KG-MDL representation (presented in Section 4.1). For example the Lemon dataset is composed of 30 thousand triples, which in KG-MDL correspond to 14 thousand vertices having a total of 19 thousand vertex labels, connected by 17 thousand edges; in total there are 2 thousand unique label symbols in the dataset.

The results of our experiments are available online on the GraphMDL Visualizer website¹¹, which interactively displays for each dataset global statistics as well as all the extracted patterns.

5.2. Quantitative Evaluation

Table 2 presents the results obtained when running KG-MDL on the different datasets. For each dataset three experiments were run with a row cover timeout of respectively 200ms, 500ms, and 1s. Each experiment was given 8 hours (the Lemon experiments actually terminated after 4h having finished their exploration). We only report in the table for each dataset the experiment that gave the best (i.e. the lowest) description length at the end of the 8

⁶<https://gitlab.inria.fr/fbariatt/graphmdl>

⁷https://lemon-model.net/lexica/dbpedia_en/

⁸<https://www.dbis.informatik.uni-goettingen.de/Mondial/>

⁹<http://www.semanticbible.com/ntn/ntn-overview.html>

¹⁰We obtained the data from the authors of [27].

¹¹<http://graphmdl-viz.irisa.fr/>

Table 2

Results of running KG-MDL on the different datasets for 8h with a row cover timeout of 200ms, 500ms and 1s. Only the timeout yielding the lowest description length is reported for each dataset.

Dataset	Row cover timeout	#patterns	Data labels described by patterns with		Compression ratio $L\% = \frac{L(CT_{kgmdl}, D)}{L(CT_0, D)}$	Time to reach 95% of L% gain
			≥ 2 labels	≥ 10 labels		
NTNames	200ms	234	86.6%	44.8%	50.8%	6 min
Lemon	200ms	168	94.9%	87.5%	20.3%	4 min
Mondial-europe	500ms	118	92.0%	79.1%	34.1%	1 h 45 min
Taaable	200ms	93	83.8%	76.1%	41.0%	2 h 25 min
Mondial	500ms	78	90.2%	55.4%	34.6%	2 h 32 min

hours (we evaluate the impact of the row cover timeout further along this section). The first column of the table reports the dataset, and the second one reports the timeout for that experiment.

Number of generated patterns. The third column of Table 2 (#patterns) reports the number of patterns in the best code table found by KG-MDL. We observe that on all datasets, even on the biggest ones, the number of generated patterns is in the lower hundreds. We argue that it is plausible for a human analyst to comprehend such an amount of patterns, and observe them in order to get an idea of the structure of the data. Note that classic graph mining algorithms would tend to extract tens of thousands or hundreds of thousands patterns for datasets of this size. Thus, we can conclude that KG-MDL is effective at its goal of generating a *small and human-sized* set of patterns.

Labels described by patterns. The fourth column of Table 2 reports how many of each dataset’s labels are described in the rewritten graph by instances of the patterns generated by KG-MDL, the other labels being described by the trivial one-label patterns of the CT_0 . We observe that on all datasets more than 80% of the labels are described by a pattern, which shows that KG-MDL patterns are descriptive enough to cover most of the data. When looking at the data in detail, we notice that most of the non-covered labels correspond to concrete literal values (e.g. “123” or “Alice”). This is consistent with the idea of finding patterns to describe the *structure* of the data, without caring about the concrete values of the different properties.

For the fifth column of Table 2 we take a stricter limit of only considering the data labels covered by patterns with at least ten labels in their structure. We observe that despite limiting to the occurrences of these large patterns, the amount of data labels described remains high. This shows that the data contains large regularities that KG-MDL is able to capture, which would probably interest the analyst.

Description length. The sixth column of Table 2 reports the compression ratio $L\% = \frac{L(CT_{kgmdl}, D)}{L(CT_0, D)}$, corresponding to the ratio between the description length given by the code table found at the end of the experiment and the one given by the initial code table CT_0 . This is the MDL measure that KG-MDL tries to minimize. We observe that the approach is efficient in doing so, obtaining a compression ratio that ranges from 20% to 50% depending on the dataset. From an information theory point of view, that means that KG-MDL manages to find a set of patterns that makes it possible to describe the data with a fifth to half as much information than without those patterns.

Description length evolution over time. We expect that a user would probably not want to wait for 8 hours in order to get a result from their experiments. Since KG-MDL is an anytime algorithm, it can be stopped at any point during the search, yielding the best result found so far. In order to evaluate the impact of stopping the algorithm sooner, we show in Fig. 7 the evolution of description length over time. Each point of the plot corresponds to one iteration of the search algorithm, i.e. to a code table that is found which is better than the previous one. We observe that most of the gain is obtained in the first few minutes or hours. For instance, on the Mondial dataset, the compression ratio goes from 100% (the starting model) to around 40% in the first hour, and the last 7 hours are spent going from 40% to 35%. On top of that, as time advances the points tend to be more spaced out, as it takes more time to find a code table that manages to improve the previous one.

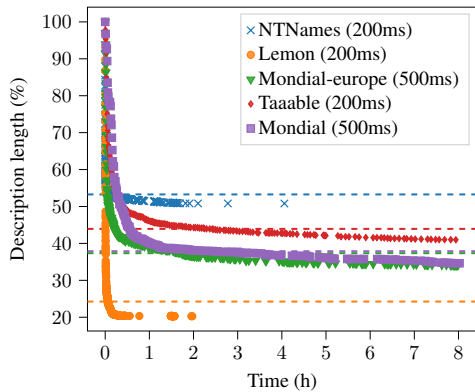


Fig. 7. Evolution of the compression ratio $L\%$ given by the CT found by KG-MDL as a function of time. The dotted line of each dataset represents the compression ratio that is 95% of the way w.r.t. the last one.

Table 3 Compression ratio $L\%$ on the Mondial dataset for different row cover timeouts (8h experiment).

Timeout per CT row during search	Compression ratio $L\%$
100ms	36.0%
200ms	34.6%
500ms	34.6%
1s	37.3%
5s	41.3%
10s	40.7%

We report in the last column of Table 2 the time needed in each experiment to attain a compression ratio that is 95% of the way w.r.t. the full 8 hours experiment. We observe that for all datasets, KG-MDL obtains good results very quickly, which means that the user can usually stop the algorithm early and still obtain good results.

Impact of the row cover timeout. In order to evaluate the impact of the row cover timeout parameter described in Section 4.4 more in depth, we ran some additional experiments on the Mondial dataset, whose results are shown in Table 3. We observe that a timeout too low or too high can negatively impact the final description length (which is also the case for the other datasets of Table 2, as the best description length is not obtained with the same timeout value for all datasets). We believe that too low a timeout can stop the computation before some important occurrences are found for some patterns, which makes the algorithm under-estimate useful patterns. On the other hand, too high a timeout entails the risk of spending so much time on the computation of the occurrences of few complex patterns, that it will limit the amount of different patterns that can be processed in a given time. In conclusion, the cover timeout is a useful parameter, which can make a difference in the performances of the algorithm.

5.3. Qualitative Evaluation

In this section we take a look at the patterns that are selected by KG-MDL and we show that they allow to surface knowledge about both the schema of the KG from which they are extracted, and from the actual raw data, conveying knowledge that is not present in the schema. Note that the patterns are shown here with the graph representation presented in Section 4.1. However, they can be converted seamlessly to SPARQL queries that can be executed on the original data. In Appendix A we show the SPARQL queries corresponding to the patterns in this section.

5.3.1. Schema Knowledge from the Patterns

By looking at the github page of the Lemon dataset¹², we observe that the dataset is created using some predefined structures called *lemon design patterns*. KG-MDL does not have any knowledge of those design patterns, but all of them (around ten different types of design patterns are used in our dataset) are found in the set of extracted patterns, demonstrating the ability of our approach to identify them as important structures in the data. Fig. 8 shows the pattern that appears the most in the KG-MDL rewritten graph on the lemon dataset: it corresponds to the lemon design pattern “ClassNoun”. This structure describes a lexical entry that is a common noun (e.g. “lake”), and whose meaning is a class (e.g. `dbo:Lake`). The structure used in Lemon to encode this entry is not trivial: the pattern contains 7 vertices, 10 vertex labels, and 7 edge labels. The pattern combines three linguistic aspects: lexical (the

¹²<https://github.com/ag-sc/lemon.dbpedia>

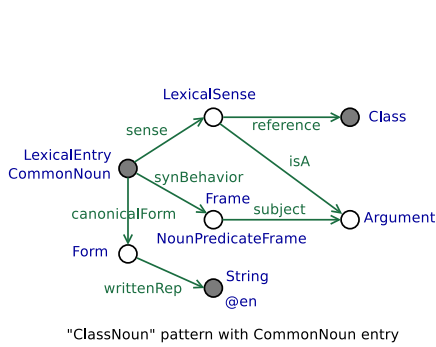


Fig. 8. The pattern that appears the most in the KG-MDL rewritten graph on the Lemon dataset. Vertices filled in grey are the ports of the pattern.

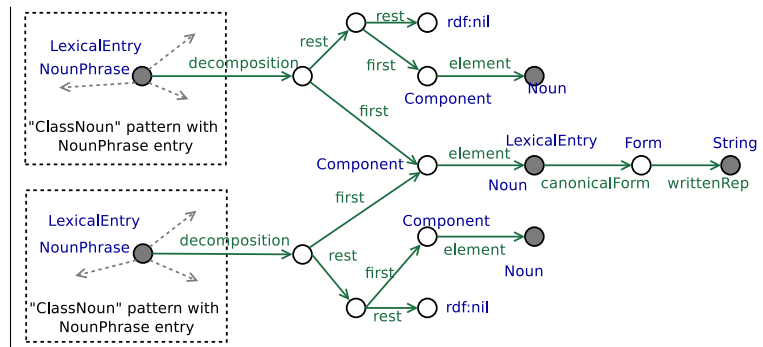


Fig. 9. A pattern extracted by KG-MDL on the Lemon dataset. The squares on the left correspond to the structure of Fig. 8 except for label “CommonNoun” being replaced by label “NounPhrase”.

written representation of the term), syntactic (a noun predicate frame with one argument, the subject), and semantic (the sense is given as a reference to an RDFS class). Syntax and semantics are aligned through the Argument vertex, which links the subject of the frame and the instance of the class. Only three out of the seven pattern vertices—filled in grey in the figure—appear as ports in the rewritten graph: the lexical entry, the class, and the written representation. They correspond to the vertices that differ from one pattern occurrence to another, i.e. they have different additional labels, either vertex labels or edges connecting them to other pattern occurrences. Hence, each pattern occurrence is fully determined by the choice of three data vertices. This corresponds to how the data is created for the lemon dataset, where only those pieces of information need to be provided per instance, the rest of the structure being created automatically. In this case, KG-MDL allowed to highlight the *schema* of the data: the extracted pattern shows the user a characteristic shape from the data that stems from how the data has been designed. This shape could, for instance, be converted to a SHACL constraint in order to ensure that a KG corresponds to the correct design.

Fig. 9 shows another pattern extracted from the Lemon dataset. Contrary to the previous one, this pattern does not correspond directly to a lemon design pattern (what we could call the “schema” of the graph). It represents a pair of lexical entries of type “NounPhrase” (i.e. terms composed by multiple elements) that share the same first element, but whose meaning are distinct classes, e.g. `dbo:BloodVessel` and `dbo:BloodType`, `dbo:RadioProgram` and `dbo:RadioStation`, or `dbo:MountainRange` and `dbo:MountainPass`. This patterns tells us that it is common to have terms with different meanings but whose lexical form starts with the same word. Indeed the pattern has 34 occurrences in the rewritten graph (i.e. it describes 68 lexical entries). Since this pattern shows how to encode this kind of lexical entries, it corresponds to knowledge that could be present in the schema of the data (the design patterns). Therefore, in this case KG-MDL surfaced knowledge that could become an *addition to the schema*, but was not explicitly present in the schema during the design.

Fig. 10 shows one of the patterns extracted by KG-MDL on the NTNAMES dataset. In the OWL description of this dataset `childOf` and `parentOf` are marked as inverse properties, and `siblingOf` is marked as symmetric. KG-MDL generates a pattern that strongly suggests those axioms without having access to the schema itself, just from the data. The algorithm also suggests an axiom that is not present in the schema of the dataset: when two people have the same parent, they are siblings one of the other. In this case KG-MDL shows in the same pattern both knowledge that is present in the schema of the data, and knowledge that could be an addition to the schema.

5.3.2. Factual Knowledge from the Patterns

Patterns are not limited to the schema of the data, but they can also present knowledge about the data that is not present in the schema. While the schema of a dataset describes how facts should be encoded in general, the concrete content of a specific dataset may have some additional structure that stems from the regularity of the real-life facts in the data.

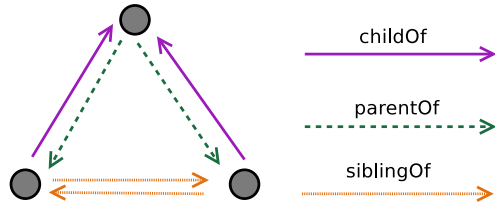


Fig. 10. A pattern extracted by KG-MDL on the NTNames dataset. Vertices filled in grey are the ports of the pattern.

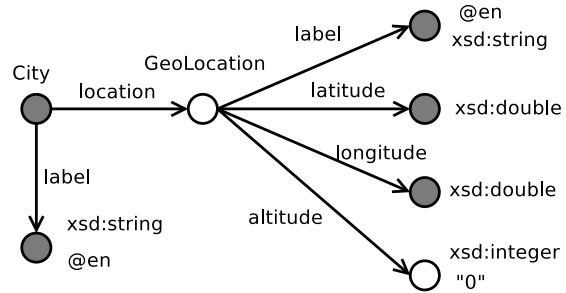


Fig. 11. A pattern extracted by KG-MDL on the NTNames dataset. Vertices filled in grey are the ports of the pattern.

Fig. 11 shows a pattern extracted from the NTNames dataset. By looking at this pattern we can understand how the geographical information of cities is represented in the data. But, more importantly, there is an interesting detail in this pattern: the altitude vertex is not a port and has a value of 0, which suggests that cities tend to have an altitude of 0 (so the information is directly encoded in the pattern). This may seem surprising, but it is actually the case that in this dataset *all* cities have altitude 0, which is probably due to a lack of information. It would not have been possible to extract this knowledge from the schema alone. Since KG-MDL’s patterns are created from the raw facts, they can highlight this type of hidden information that is related to the actual *usage* of properties and classes.

6. Conclusion

In this paper we proposed KG-MDL, a MDL-based method to extract a human-sized set of descriptive patterns from knowledge graphs and labeled multi-graphs. In order to do so, we adapted a MDL-based graph mining approach to knowledge graphs. We showed that the adaptation is not trivial, as KGs have peculiar characteristics that distinguish them from the graphs usually found in the graph mining domain. We evaluated our approach experimentally on five medium-sized datasets. We showed that it is efficient in finding small sets of patterns that minimize the MDL description length in reasonably short time. We analyzed the extracted patterns, showing that they allow to infer knowledge both about the schema used to create the data, and the concrete content of the data. Such patterns may be used to infer the schema from unknown data, create validation constraints (e.g. SHACL shapes), improve the schema by suggesting additions to it, and highlight instances where the raw data is surprising w.r.t. the schema.

As a future work, the approach may be improved to scale better towards large-scale graphs that are common in the semantic web (millions of triples or more). Because the current approach requires to compute a re-encoded representation of the full data for each candidate sets of patterns in order to compute the MDL description lengths, the scalability is affected. In order to scale to larger graphs, a sampling approach might be undertaken. However, this is not a trivial matter, as the choice of the sample may affect the extracted patterns. Therefore the sampled data should be as characteristic as possible w.r.t. the complete data. Another possible extension is to allow for patterns that contain less-specific matches for literal values, such as substring match or range inclusion. This would allow for patterns to be potentially insensitive to small changes in literal values, capturing occurrences of slightly different—but similar—concrete facts.

References

- [1] M. Meimaris, G. Papastefanatos, N. Mamoulis and I. Anagnostopoulos, Extended Characteristic Sets: Graph Indexing for SPARQL Query Optimization, in: *2017 IEEE 33rd International Conference on Data Engineering (ICDE)*, 2017, pp. 497–508. doi:10.1109/ICDE.2017.106.
- [2] U. Fayyad, G. Piatetsky-Shapiro and P. Smyth, From Data Mining to Knowledge Discovery in Databases, *AI Magazine* 17(3) (1996), 37–37. doi:10.1609/aimag.v17i3.1230.

- [3] R. Agrawal and R. Srikant, Mining sequential patterns, in: *Proceedings of the Eleventh International Conference on Data Engineering*, 1995, pp. 3–14. doi:10.1109/ICDE.1995.380415.
- [4] C. Jiang, F. Coenen and M. Zito, A survey of frequent subgraph mining algorithms, *The Knowledge Engineering Review* **28**(1) (2013), 75–105. doi:10.1017/S0269888912000331.
- [5] R. Gil, R. García and J. Delgado, Measuring the semantic web, *AIS SIGSEMIS Bulletin* **1**(2) (2004), 69–72.
- [6] R. Ramezani, M. Saraei, M. Nematbakhsh et al., SWApriori: a new approach to mining Association Rules from Semantic Web Data, *Journal of Computing and Security* (2014).
- [7] M. Barati, Q. Bai and Q. Liu, SWARM: An Approach for Mining Semantic Association Rules from Semantic Web Data, in: *PRICAI 2016: Trends in Artificial Intelligence*, Lecture Notes in Computer Science, Springer International Publishing, 2016, pp. 30–43. doi:10.1007/978-3-319-42911-3_3.
- [8] C. Bobed, P. Maillot, P. Cellier and S. Ferré, Data-driven assessment of structural evolution of RDF graphs, *Semantic Web* **11**(5) (2020), 831–853. doi:10.3233/SW-200368.
- [9] J. Lajus, L. Galárraga and F. Suchanek, Fast and Exact Rule Mining with AMIE 3, in: *The Semantic Web*, Lecture Notes in Computer Science, Springer International Publishing, 2020, pp. 36–52. doi:10.1007/978-3-030-49461-2_3.
- [10] C. Meilicke, M.W. Chekol, D. Ruffinelli and H. Stuckenschmidt, An Introduction to AnyBURL, in: *KI 2019: Advances in Artificial Intelligence*, Lecture Notes in Computer Science, Springer International Publishing, 2019, pp. 244–248. doi:10.1007/978-3-030-30179-8_20.
- [11] E. Galbrun, The minimum description length principle for pattern mining: a survey, *Data Mining and Knowledge Discovery* (2022). doi:10.1007/s10618-022-00846-z.
- [12] P.D. Grünwald, *The minimum description length principle*, MIT press, 2007.
- [13] J. Rissanen, Modeling by shortest data description, *Automatica* **14**(5) (1978), 465–471. doi:10.1016/0005-1098(78)90005-5.
- [14] D.J. Cook and L.B. Holder, Substructure Discovery Using Minimum Description Length and Background Knowledge, *Journal of Artificial Intelligence Research* **1** (1993), 231–255. doi:10.1613/jair.43.
- [15] P. Bloem and S. de Rooij, Large-scale network motif analysis using compression, *Data Mining and Knowledge Discovery* **34**(5) (2020), 1421–1453. doi:10.1007/s10618-020-00691-y.
- [16] D. Koutra, U. Kang, J. Vreeken and C. Faloutsos, Summarizing and understanding large graphs, *Statistical Analysis and Data Mining: The ASA Data Science Journal* **8**(3) (2015), 183–202. doi:10.1002/sam.11267.
- [17] S. Goebel, A. Tonch, C. Böhm and C. Plant, MeGS: Partitioning Meaningful Subgraph Structures Using Minimum Description Length, in: *2016 IEEE 16th International Conference on Data Mining (ICDM)*, 2016, pp. 889–894. doi:10.1109/ICDM.2016.0108.
- [18] F. Bariatti, P. Cellier and S. Ferré, GraphMDL+: Interleaving the Generation and MDL-based Selection of Graph Patterns, in: *SIGAPP Symp. Applied Computing (SAC)*, ACM, 2021.
- [19] C. Belth, X. Zheng, J. Vreeken and D. Koutra, What is Normal, What is Strange, and What is Missing in a Knowledge Graph: Unified Characterization via Inductive Summarization, in: *Proceedings of The Web Conference 2020, WWW '20*, Association for Computing Machinery, 2020, pp. 1115–1126. doi:10.1145/3366423.3380189.
- [20] X. Yan and J. Han, gSpan: Graph-based substructure pattern mining, in: *2002 IEEE International Conference on Data Mining. Proceedings*, Springer Berlin Heidelberg, 2002, pp. 721–724. doi:https://doi.org/10.1109/ICDM.2002.1184038.
- [21] S. Nijssen and J.N. Kok, A Quickstart in Frequent Structure Mining Can Make a Difference, in: *Proceedings of the Tenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '04*, ACM, 2004, pp. 647–652. doi:10.1145/1014052.1014134.
- [22] N. Ahmadi, V.-P. Huynh, V. Meduri, S. Ortona and P. Papotti, Mining Expressive Rules in Knowledge Graphs, *Journal of Data and Information Quality* **12**(2) (2020), 8:1–8:27. doi:10.1145/3371315.
- [23] H. Paulheim, Knowledge graph refinement: A survey of approaches and evaluation methods, *Semantic Web* **8**(3) (2017), 489–508. doi:10.3233/SW-160218.
- [24] T.C.M. Lee, An Introduction to Coding Theory and the Two-Part Minimum Description Length Principle, *International Statistical Review* **69**(2) (2001), 169–183. doi:10.1111/j.1751-5823.2001.tb00455.x.
- [25] P. Elias, Universal codeword sets and representations of the integers, *IEEE Transactions on Information Theory* **21**(2) (1975), 194–203. doi:10.1109/TIT.1975.1055349.
- [26] F. Bariatti, P. Cellier and S. Ferré, GraphMDL: Graph Pattern Selection Based on Minimum Description Length, in: *Advances in Intelligent Data Analysis XVIII*, Lecture Notes in Computer Science, Springer International Publishing, 2020, pp. 54–66. doi:10.1007/978-3-030-44584-3_5.
- [27] A. Cordier et al., Taaable: A Case-Based System for Personalized Cooking, in: *Successful Case-based Reasoning Applications-2*, Studies in Computational Intelligence, 2014, pp. 121–162.

Appendix A. SPARQL Queries of Experimental Evaluation Patterns

In this section we show the SPARQL queries that correspond to the patterns shown in Section 5.3 (we omit the pattern of Fig 9 because of its size).

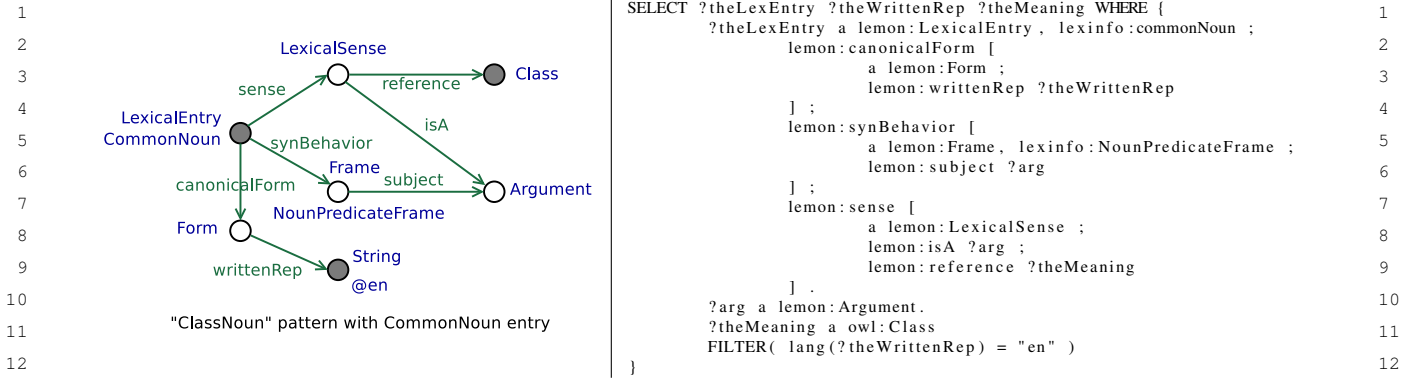


Fig. 12. The SPARQL query corresponding to the pattern of Fig. 8. The chosen query variables are the pattern's ports (in grey in the image).

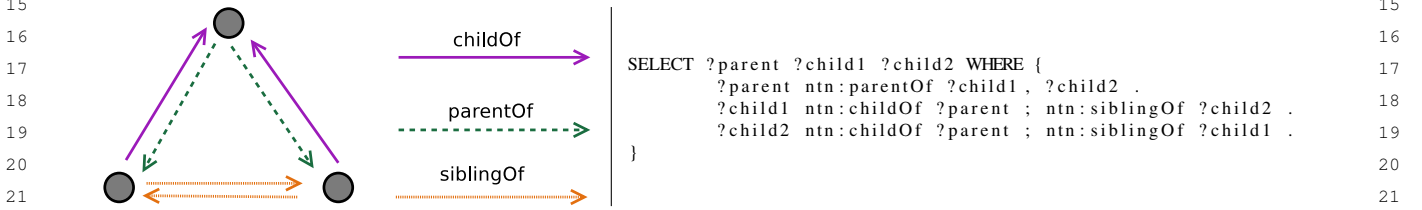


Fig. 13. The SPARQL query corresponding to the pattern of Fig. 10.

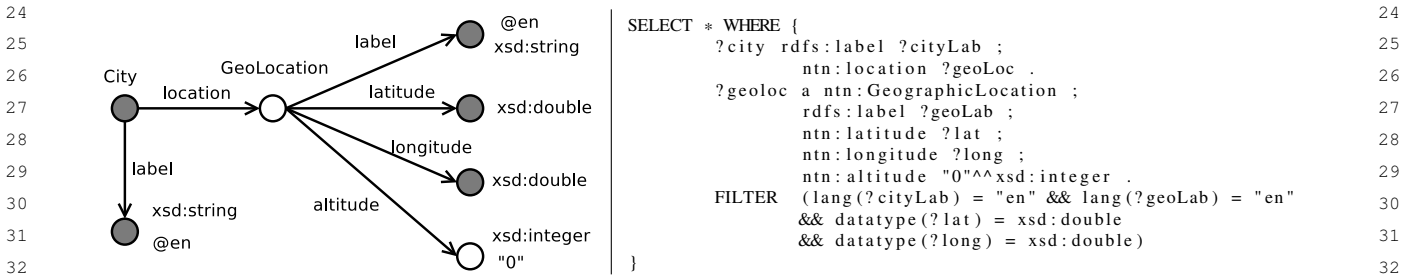


Fig. 14. The SPARQL query corresponding to the pattern of Fig. 11.

Appendix B. Prequential Codes

To compute the MDL description length in Section 4.2 we extensively use *prequential plug-in codes* [12]. A prequential plug-in code allows to encode a sequence of items one item at a time, without needing to know the probability of each item in advance (only the possible items need to be known). Also, prequential codes are asymptotically optimal for large sequences, and —surprisingly— do not depend on the order of items in the sequence (which is fundamental when no canonical order exist on a sequence). In this section we explain how a description length can be computed using prequential codes.

Let \mathcal{X} be a set of items, and $S = \langle s_0, s_1, \dots, s_{|S|-1} \rangle$ a sequence of items $s_i \in \mathcal{X}$. We denote S^i the sub-sequence up to and including position i (e.g. $S^2 = \langle s_0, s_1, s_2 \rangle$), and $usage(x, S)$ the number of occurrences of element x in the sequence S . Each element $s_i \in S$ can be encoded optimally using a prefix code based on the distribution of elements before it [12]. An initial pseudocount ϵ is given to each possible element to avoid elements having an empty code (usually $\epsilon = 0.5$ is used in applications). Therefore the description length for an element s_i based on the previous elements is:

$$L(s_i | S^{i-1}) = -\log \left(\frac{usage(s_i, S^{i-1}) + \epsilon}{\sum_{x \in \mathcal{X}} [usage(x, S^{i-1}) + \epsilon]} \right)$$

Therefore, the description length of the whole sequence is:

$$\begin{aligned}
 L(S) &= \sum_{i=0}^{|S|-1} L(s_i | S^{i-1}) \\
 &= \sum_{i=0}^{|S|-1} \left[-\log \left(\frac{\text{usage}(s_i, S^{i-1}) + \epsilon}{\sum_{x \in \mathcal{X}} [\text{usage}(x, S^{i-1}) + \epsilon]} \right) \right] \\
 (\text{sum of logs is log of product}) &= -\log \left(\prod_{i=0}^{|S|-1} \frac{\text{usage}(s_i, S^{i-1}) + \epsilon}{\sum_{x \in \mathcal{X}} [\text{usage}(x, S^{i-1}) + \epsilon]} \right)
 \end{aligned}$$

Since in a product of fractions the order of each operand in the numerator and denominator does not matter, when computing the description length of a whole sequence using a prequential code, the order of elements in the sequence does not matter.

Note that in the above expression, the numerator for each occurrence of an element $x \in \mathcal{X}$ only depends on the previous occurrences of that same element, because $\text{usage}(x, S)$ yields ϵ for the first occurrence of x , $(1 + \epsilon)$ for the second, and so on, independently of the occurrences of other elements. Therefore the numerator can be made clearer by regrouping occurrences of the same element. The denominator of the above expression only depends on the sum of usages of previous elements (which is the length of S^{i-1} for element i), plus a value of ϵ for each item $x \in \mathcal{X}$. Therefore the above expression can be rewritten as:

$$L(S) = -\log \left(\frac{\prod_{x \in \mathcal{X}} \prod_{i=0}^{\text{usage}(x, S)-1} (i + \epsilon)}{\prod_{i=0}^{|S|-1} (i + |\mathcal{X}|\epsilon)} \right)$$

In the literature, this description length is generally given by employing the gamma function Γ , which is an extension of factorials to real number: $\Gamma(a) = (a-1)\Gamma(a-1)$ (for an integer n , $\Gamma(n) = (n-1)!$). We can introduce the gamma function in the formula above, by writing out all terms of the products. and using the property that $(n)(n-1)(n-2)\dots(k+1) = \frac{\Gamma(n+1)}{\Gamma(k+1)}$. We do not detail the whole transformation here, which gives:

$$\begin{aligned}
 L(S) &= -\log \left(\frac{\prod_{x \in \mathcal{X}} \frac{\Gamma(\text{usage}(x, S) + \epsilon)}{\Gamma(\epsilon)}}{\frac{\Gamma(|S| + |\mathcal{X}|\epsilon)}{\Gamma(|\mathcal{X}|\epsilon)}}} \right) \\
 (\text{log properties}) &= -\sum_{x \in \mathcal{X}} \left[\log \left(\frac{\Gamma(\text{usage}(x, S) + \epsilon)}{\Gamma(\epsilon)} \right) \right] + \log \left(\frac{\Gamma(|S| + |\mathcal{X}|\epsilon)}{\Gamma(|\mathcal{X}|\epsilon)} \right)
 \end{aligned}$$

Which corresponds to the formulas usually given in the literature.

The advantage of this encoding is that the final formula only depends on the usage of each items $x \in \mathcal{X}$, the length of the sequence, and the number of elements of \mathcal{X} . The actual content of the sequence is not needed, which makes the encoding particularly suitable for computing description lengths where the order of elements in the sequence is arbitrary chosen, as the choice does not impact the description length.