

Building Relatedness Explanations from Knowledge Graphs

Giuseppe Pirro^{a,*}

^aICAR-CNR, Italian National Research Council, Rende (CS), Italy

E-mail: pirro@icar.cnr.it

Abstract. Knowledge graphs (KGs) are a key ingredient for searching, browsing and knowledge discovery activities. Motivated by the need to harness knowledge available in a variety of KGs, we face the problem of building relatedness explanations. Relatedness explanations are graphs that can explain how a pair of entities is related in a KG and can be used in a variety of tasks; from exploratory search to query answering. We formalize the notion of relatedness explanation and present two algorithms. The first, E4D (Explanations from Data) assembles explanations starting from paths in the data. The second algorithm E4S (Explanations from Schema), starting from the KG schema, allows to find paths driven by a specific relatedness perspective. We describe different criteria to build explanations based on information-theory, diversity and their combination. As a concrete application, we introduce relatedness-based KG querying, which revisits the query-by-example paradigm from the perspective of relatedness explanations. We implemented all machineries in the `RECAP` tool, which is based on RDF and SPARQL. We discuss an evaluation of our explanation building algorithms and a comparison of `RECAP` with related systems on real-world data.

Keywords: Knowledge Graphs, Explanations, Patterns, Relatedness-based querying

1. Introduction

Knowledge graphs (KGs) maintaining structured data about entities are becoming a common support for browsing, searching and knowledge discovery activities [1]. Search engines like Google, Yahoo! and Bing complement the classical search results with facts about entities in their KGs. An even large number and variety of KGs, based on the Resource Description Framework (RDF) data format, stem from the Linked Open Data project [2]. DBpedia, Freebase, DBLP and Yago are just a few examples that witness the popularity and spread of open KGs. KGs are a type of heterogeneous information network [3] where nodes represent entities and edges different types of relationships. Extracting knowledge from KGs has applications in many domains including information retrieval [4], radicalization detection [5], twitter analysis [6], recommendation [7], clustering [8], entity resolution [9], and generic exploratory search [10].

One common-need for many classes of knowledge discovery tasks is that explaining the relatedness between entities. The availability of (visual) explanations helps in understanding why entities are related while at the same time allowing to discover/browse other entities of interest. The availability of explanations is useful in several areas including: terrorist networks, to uncover the connections between two suspected terrorists [11]; co-author networks, to discover interlinks between researchers [12]; bioinformatics, to discover relationships between biomedical terms [13]; generic exploratory search [10].

Fig. 1 (e) shows an excerpt of explanation for the pair of entities (Metropolis, F. Lang) obtained from the DBpedia KG. The explanation includes relationships among entities of different types (e.g., Film, Actors, Cinematographer). We can see that R. Klein-Rogge starred Metropolis and other films (e.g., *Destiny*, *Spione*) all directed by F. Lang. We can also see that T. von Harbou has been married with both R. Klein-Rogge and F.Lang and that she has written some films with the latter. We can also notice that W. Ruttmann has been the cinematographer for both Metropolis and *Die Nibelungen*.

*Part of this work was done while the author was working in the WeST group at the University of Koblenz-Landau.

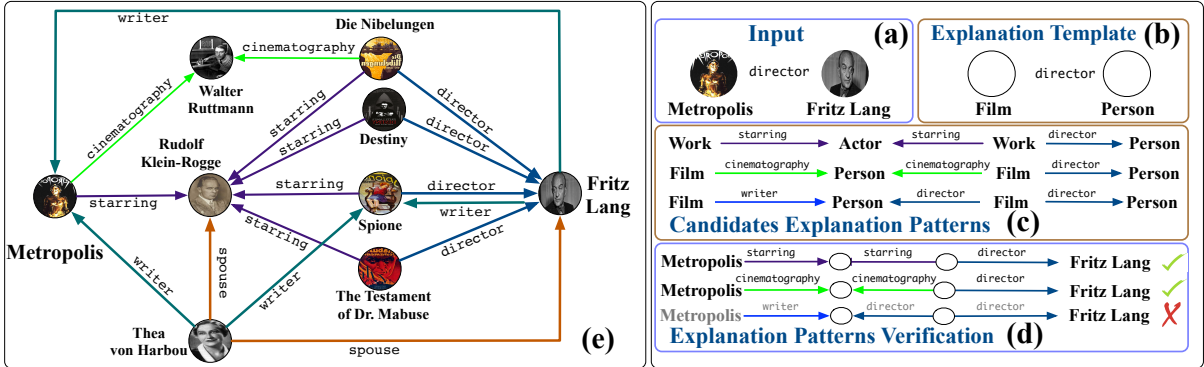


Fig. 1. An excerpt of explanation from the DBpedia KG (left). Overview of the E4S algorithm (right).

The computation of relatedness explanations is usually tackled by discovering paths that link entities in the data (e.g., [14–19]). Our first explanation-building algorithm called Explanations From the Data (E4D presented in Section 3) [10] follows the same philosophy. To guarantee better generality and scalability E4D is SPARQL-based; it reduces the problem of finding paths to that of evaluating a set of (automatically generated) SPARQL queries and can leverage any SPARQL-enabled (remote) KG. However, working directly at the data level may have some shortcomings. In fact, E4D and related research (e.g., [16, 19, 20]) do not allow to build focused explanations, that is, explanations concerning a particular relatedness perspective. It would be interesting, for instance, to build explanations by providing one or more target predicates that drive the process toward data that concern these and closely related aspects only.

Therefore, we introduce a second algorithm called Explanations from the Schema (E4S presented in Section 4) that starts from the KG schema. The idea is to start from explanation templates built from the entity types of the input pair along with some target predicates. Fig. 1 (b) shows the explanation template for the entity pair (Metropolis, F. Lang) when considering director as a target predicate. The explanation template is then used to generate candidate explanation patterns of bounded-length, that is, schema-compliant paths having as endpoints the *types* (e.g., Film, Person) instead of concrete entity instances (e.g., Metropolis, F. Lang). As the search space for candidate patterns can be very large, E4S leverages a predicate relatedness measure to isolate the portion of the schema that is most related to the target predicates. Fig. 1 (c) shows some candidate explanation patterns of length 3. These are all plausible explanation patterns for the pair in input. We can see for instance, that a Work (Metropolis is a Film,

which is a subclassOf Work) has some Actor starring it. The same Actor starred other Works whose director is a Person. Nevertheless, not all explanation patterns will have a counterpart in the data; they need to be verified to extract knowledge useful to build an explanation. Fig. 1 (d) shows that only the first two patterns are actually verified in the data. This can be seen by looking at Fig. 1 (e) where we notice that the first pattern leads from Metropolis to F. Lang via R. Klein-Rogge and Spione or Destiny or The Testament of Dr. Mabuse. As for the second pattern, it leads from Metropolis to F. Lang via W. Rotmann and Die Nibelungen.

The separation between candidate explanation pattern generation and verification allows to choose only those patterns that are of most interest and proceed to their verification according to some ranking (e.g., average relatedness between predicates in the pattern and the target predicate). The richness of node and edge types, the usage of meta information and the possibility to focus on the most interesting explanations first, offers a significant advantage wrt related research. We are not aware of any previous work that starts from the KG schema to build relatedness explanations.

Contributions. We contribute: (i) a SPARQL-based algorithm called E4D, which starting from a pair of entities can retrieve paths interlinking them directly looking at the data; (ii) a KG-schema-based algorithm E4S, which leverages the schema to build candidate explanation patterns and an algorithm for their verification in the data; (iii) different ways to rank paths (found by E4D and/or E4S) based on informativeness, diversity and combinations of them; (iv) a (visual) tool called RECAP; (v) an approach to query KGs by relatedness; (vi) an experimental evaluation.

A preliminary version of this work, which does not include the E4S algorithm, appeared in ISWC2015 [10].

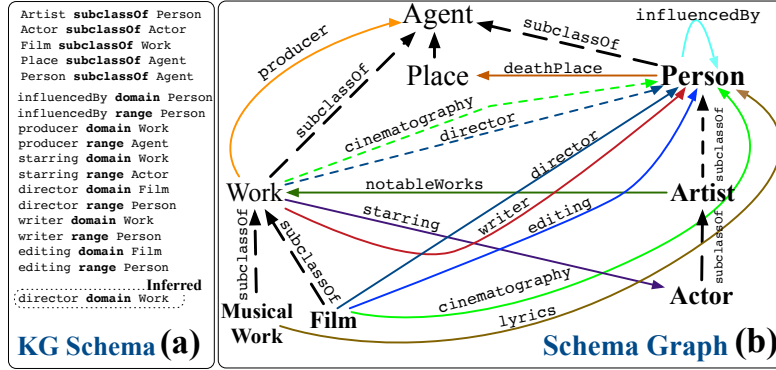


Fig. 2. A KG schema (a) and its corresponding schema graph (b).

Outline. The remained of the paper is organized as follows. We introduce the problem along with the necessary background in Section 2. Section 3 introduces the E4D algorithm. Section 4 introduces our second algorithm E4S, which builds explanations starting from the KG schema. Section 5 describes different strategies to rank/combine paths in order to build explanations. We describe the RECAP tool implementing our algorithms in Section 6. Section 7 discusses the experimental evaluation. Related Work is treated in Section 8. We conclude in Section 9.

2. Problem Description and Background

We now formalize the notions of KG and knowledge base. Although there are several KGs today available (e.g., Yahoo!, Google) we will focus on those encoded in RDF¹ widely and openly accessible on the Web for querying via the standard SPARQL language [21].

Let \mathcal{U} (URIs) and \mathcal{L} (literals) be countably disjoint infinite sets. An RDF *triple* is a tuple of the form $\mathcal{U} \times \mathcal{U} \times (\mathcal{U} \cup \mathcal{L})$ whose elements are referred to as *subject*, *predicate* and *object*, respectively. As we are interested in discovering explanations in terms of nodes and edges carrying semantic meaning for the user, we do not consider blank nodes. We also assume an infinite set \mathcal{V} of variables disjoint from the sets \mathcal{U} and \mathcal{L} . In the SPARQL language, variables in \mathcal{V} are prefixed with a question mark (e.g., $?x$ is a variable).

Definition 1. (Triple and Graph Pattern). A triple pattern τ is a triple $(x_s, y_p, z_o) \in (\mathcal{U} \cup \mathcal{V})(\mathcal{U} \cup \mathcal{V})(\mathcal{U} \cup \mathcal{L} \cup \mathcal{V})$. We denote by $var(\tau)$ the variables in τ . A graph pattern of length k is a set of triple patterns τ_1, \dots, τ_k where $(var(\tau_i) \cap var(\tau_{i+1}) \neq \emptyset), i \in [1, k - 1]$.

A Knowledge Graph (KG) is a directed node and edge labeled multi-graph $G = (V, \mathcal{E}, T)$ where V is a set of uniquely identified vertices representing entities (e.g., D. Lynch), \mathcal{E} a set of predicates (e.g., director) and T a set of triples of the form (s, p, o) representing directed labeled edges, where $s, o \in V$ and $p \in \mathcal{E}$.

To structure knowledge, KGs resort to an underlying schema S , which can be seen as a set of triples defined by using some ontology vocabulary (e.g., RDFS, OWL). In this paper, we focus on RDFS and specifically on entity *types* (and their hierarchy) and *domain* and *range* of predicates. We denote by $type(e)$ the set of types of an entity e . Our approach leverages RDFS inference rules to construct the RDFS-closure of the KG schema [22, 23]. From the closure of the schema we build the corresponding schema graph closure.

Definition 2. (Schema Graph Closure). Given a KG schema closure S , its schema graph closure is defined as $G_S = (V_s, \mathcal{E}_s, T_s)$, where each $v_i \in V_s$ is a node denoting an entity type, each $p_i \in \mathcal{E}_s$ denotes a predicate and each $(v_s, p_i, v_t) \in T_s$ is a triple (asserted or inferred via RDFS reasoning) where v_i (resp., v_t) is the domain (resp., range) of the predicate p_i .

Fig. 2 (a) shows an excerpt of the DBpedia schema while Fig. 2 (b) of its schema closure; here, dashed lines represent inferred triples. For instance, the triple (Work, director, Person) is inferred from the triples (Film, subclassOf, Work) and (Film, director, Person). Given a predicate p , $Dom(p)$ (resp., $Range(p)$) denotes the set of nodes (i.e., entity types) in G_s having an outgoing (resp., incoming) edge labeled with p . We refer to the schema graph closure simply as the schema graph.

Definition 3. (Knowledge Base). A knowledge base is a tuple of the form $K = \langle G, G_s, A \rangle$ where G is a knowledge graph, G_s is the schema graph and A is a query endpoint used to access data in G .

¹A list is available at <http://lod-cloud.net>

2.1. Problem Description

Motivation. The high-level objective of this paper is to tackle the problem of explaining knowledge in KGs. As discussed in the Introduction, the availability of explanations can support several applications scenarios. The need for relatedness explanations also emerged in the context of the SENSE4US FP7 project², which aimed at creating a toolkit to support information gathering, analysis and policy modeling. Here, relatedness explanations were useful to investigate and show to the user topic connectivity³, which enabled to find out previously unknown information that is of relevance, understand how it is of relevance, and navigate it.

Input. The input of our problem is a pair (w_s, w_t) of entities defined in some knowledge base $K=\langle G, G_s, A \rangle$.

Output. Given $K=\langle G, G_s, A \rangle$ and a pair of entities (w_s, w_t) , the output that we seek is a graph $G_e \subseteq G$; we call such a graph the *relatedness explanation* between w_s and w_t .

Challenges. The problem that we tackle in this paper, the algorithmic solutions and their implementation pose several challenges, among which:

- (1) how to meaningfully capture the notion of explanation between entities? How to control the size of an explanation? Which kind of information in G is useful for building an explanation?
- (2) what role can the KG schema have?
- (3) how to make available the machinery discussed in the paper in a variety of KGs?

2.2. Basic Definitions

Before further delving into the discussion of the solutions to the above challenges, we introduce some definitions.

Definition 4. (Explanation). Given a knowledge base $K=\langle G, G_s, A \rangle$ and a pair of entities (w_s, w_t) where $w_s, w_t \in G$, an explanation is a tuple of the form $E=(w_s, w_t, G_e)$ where $w_s, w_t \in G_e$ and $G_e \subseteq G$.

The above definition is very general; it only states that two entities are connected via nodes and edges in a graph G_e , which is a subgraph of the knowledge graph G and has an arbitrary structure. The challenging as-

pect that we face in this paper concerns *how to uncover the structure of G_e* .

To tackle this challenge, we shall characterize the desired properties of G_e . Consider the explanation shown in Fig. 3 (a); G_e contains two types of nodes: nodes such as n_1, n_3, n_4 that belong to some path between w_s and w_t and other nodes such as n_2 that do not.




Fig. 3. An explanation (a) and a minimal explanation (b).

Although the edge (n_2, p_1, n_3) can contribute to better characterize n_3 , such edge is in a sense non-necessary as it does not directly contribute to explain how w_s and w_t are related. Hence, we introduce the notion of necessary edge.

Definition 5. (Necessary Edge). An edge $(n_i, p_k, n_j) \in G$ is necessary for an explanation $E=(w_s, w_t, G_e)$ if it belongs to a simple path (no node repetitions) between w_s and w_t .

The necessary edge property enables to refine the notion of explanation into that of minimal explanation.

Definition 6. (Minimal Explanation). Given a knowledge base $K=\langle G, G_s, A \rangle$ and a pair of entities (w_s, w_t) where $w_s, w_t \in G$, a minimal explanation is an explanation $E=(w_s, w_t, G_e)$ where G_e is obtained as the merge of all simple paths between w_s and w_t .

Fig. 3 (b) shows a minimal explanation. Minimal explanations enable to focus only on nodes and edges that are in some path between w_s and w_t ; hence, minimal explanations preserve connectivity information only. Representing explanations as graphs enables users to have (visual) insights of why/how two entities are related and discover/browse other entities. An investigation of the usefulness of graph visualization supports is out of the scope of this paper; the reader can refer to Herman et al. [24] for a comprehensive discussion about the topic. We are now ready our first algorithm, which discovers paths useful to build different kinds of relatedness explanations by directly looking at the KG data.

²<http://www.sense4us.eu>

³A module of the SENSE4US toolkit extracts topics from policy documents.

3. Explaining Relatedness from the Data (E4D)

Relatedness explanations are graphs that provide a (concise) representation of the relatedness between a pair of entities in a KG in terms of predicates (carrying a semantic meaning) and other entities. The challenging question is how to retrieve explanations. Consider the minimal explanation shown in Fig. 3 (b); it could be retrieved by matching the *pattern graph* G_p shown in Fig. 4 (nodes and edge labels are query variables) against G . Hence, if the structure of G_p were available one could easily find G_e ; however, such structure, that is, the right way of joining query variables representing nodes and edges in G_p is unknown before knowing G_e . Nevertheless, since minimal explanations are built by considering (simple) paths in the data between w_s and w_t , the retrieval of such paths is the first step toward building explanations via the E4D (Explanations from the Data) algorithm.




Fig. 4. The pattern graph for the minimal explanation in Fig. 3 (b).

Generally speaking, paths between entities can have an arbitrary length; in practice it has been shown that for KGs like Facebook the average distance between entities is bound by a value [25]. The choice of considering paths of length k in our approach is reasonable on the light of the fact that we focus on providing explanations of *manageable size* that can be *visualized and interpreted* by the user. Fig. 5 outlines our first explanation building approach. We now describe step (1); step (2) and (3) will be described in Section 5.

3.1. Finding Paths from the Data

A path is a sequence of edges (RDF triples) bound by a length value k . The underlying assumption of the E4D algorithm is to access a knowledge base $K=\langle G, G_s, A \rangle$ only via the (SPARQL) query endpoint A . This allows to work on top of any KG for which a query endpoint is available.

Definition 7. (k -connectivity Pattern). Given a knowledge base $K=\langle G, G_s, A \rangle$ a pair of entities (w_s, w_t) where $w_s, w_t \in G$ and an integer k , a k -connectivity pattern is a tuple of the form $\Xi=\langle w_s, w_t, \mathcal{T}, k \rangle$ where \mathcal{T} is a graph pattern of length k .

Building Explanations from Data (E4D)

Input: A pair (w_s, w_t) , an integer k , the address of the query endpoint A

Output: A graph G_e

- (1) *Find paths*: we are going to describe in Section 3.1 an approach based on SPARQL queries against A to retrieve paths between w_s and w_t of length k .
- (2) *Rank paths*: we describe in Section 5 different mechanisms to rank paths.
- (3) *Select and merge top- m paths*: we discuss in Section 5.4 different ways of selecting ranked paths to build an explanation G_e .

Fig. 5. Building explanations from the data.

An example of graph pattern \mathcal{T} is shown in Fig. 6. Here, both nodes (but w_s and w_t) and edges represent query variables. Note that in the figure, edge directions are not reported; each edge has to be considered both as incoming and outgoing, which corresponds to join triple patterns in all possible ways. Indeed, by joining each of the k triple patterns in Fig. 6 in all different ways it is possible to obtain a set of 2^k graph patterns. These graph patterns can be used to generate SPARQL queries.




Fig. 6. Structure of a query to find paths of length k .

Example 8. (Example of k -connectivity Pattern). The 2-connectivity pattern between Metropolis (:Mt) and F. Lang (:FL) allows to generate the following set of 4 queries:

```
SELECT DISTINCT * WHERE{:Mt ?p1 ?n1. ?n1 ?p2 :FL}
SELECT DISTINCT * WHERE{:Mt ?p1 ?n1. :FL ?p2 ?n1}
SELECT DISTINCT * WHERE{?n1 ?p1 :Mt. :FL ?n1 ?p2}
SELECT DISTINCT * WHERE{?n1 ?p1 :Mt. ?n1 ?p2 :FL}
```

As we are going to discuss in the following, these queries are executed in parallel (by using multi-threading) by E4D thus significantly speeding-up up the retrieval of paths. Note that SPARQL 1.1 supports *property paths* (PPs) [21], that is, a way to discover routes between nodes in an RDF graph. However, since variables *cannot* be used as part of the path specification itself, PPs are not suitable for our purpose; we need information about all the constituent elements of a path (both nodes and edges) to build explanations.

Definition 9. (Path). Given $K=\langle G, G_s, A \rangle$ and a pair of entities (w_s, w_t) , a path ξ is a set of edges (i.e., RDF triples): $\xi(w_s, w_t)=w_s \xrightarrow{p_1} n_1 \xrightarrow{p_2} n_2 \xrightarrow{p_3} \dots \xrightarrow{p_k} w_t$, $n_i \in G \forall i \in [1, q]$, $p_j \in G \forall j \in [1, k]$ and $- \in \{\leftarrow, \rightarrow\}$.

In the above definition, the symbol $-$ models the fact that in a path we may have edges pointing in different directions. We denote by $\text{Labels}(\pi)$ the set of labels (RDF predicates) in ξ . Practically speaking, paths are obtained by evaluating queries deriving from a k -connectivity pattern; hence, n_i and p_j are solutions (variable bindings) obtained from some query evaluated on A . While in a k -connectivity pattern all edges and nodes (but w_s and w_t) are variables, in a path there are no variables. Fig. 7 shows an example of path ξ where $\text{Labels}(\xi)=\{\text{writer, spouse}\}$.



Fig. 7. An example of path between Metropolis and F. Lang.

We want to stress the fact that the set of queries to retrieve paths are automatically generated and evaluated on the (remote) query endpoint A . A sketch of the multi-thread algorithm to find paths is shown in Fig. 8.

Function $\text{FINDPATHS}(w_s, w_t, k, A, m, \text{max}T)$

Input: A pair (w_s, w_t) , an integer k , the address of the query endpoint A , monitor class m , max num. of threads $\text{max}T$.

- 1: $Q=\text{GENQUERIES}(w_s, w_t, k)$ /* queries */
- 2: **for** $q \in Q$ **do**
- 3: $\text{PATHTHREAD } t=\text{CREATETHREAD}(q, A, m)$
- 4: $\text{SUBMITJOB}(t, \text{max}T)$

Function $\text{PATHTHREAD}(q, A, m)$

Input: Query q , address A , monitor m .
/* Methods in m are thread-safe */

- 1: $\text{Res}=\text{EXECQUERY}(q, A)$ /* evaluates the query */
- 2: $m.\text{ADDERESULTS}(\text{Res})$
- 3: **if** $(m.\text{GETQUERYCOUNT}()==0$ **AND**
 $m.\text{ISLASTTHREAD}())$ **then**
- 4: $m.\text{RETURNRESULTS}()$

Fig. 8. Finding paths: an overview.

The algorithm leverages a monitor class (not reported here); an instance such class is denoted with m in the pseudocode. This class controls: (i) the concurrent access to the shared data structures by the threads (e.g., to add paths retrieved by queries); (ii) the termi-

nation of the algorithm that will happen when no more queries have to be executed and no thread is active (see lines 3-4 PATHTHREAD). Moreover, the method SUBMITJOB (not reported here) takes care of submitting threads for the executions to an executor service, which can be configured to create a $\text{max}T$ number of threads. This enables to control the degree of concurrency. We now move to the description of the E4S algorithm, which adopts a different philosophy; instead of starting from the evaluation of queries directly in the KG data to obtain paths between w_s and w_t , it leverages the schema to control which kinds of paths are interesting to look for.

4. Explaining Relatedness from the Schema (E4S)

The choice to build explanations starting from the data via E4D allows to obtain a potential large number of explanations covering different (and varied) relatedness perspectives. Nevertheless, one may be interested in the generation of explanations focused on specific relatedness perspectives. Returning to our example, when explaining the relatedness between the film Metropolis and F. Lang one could focus on building explanations in the domain of movies. This interest can be expressed by considering predicates like director or writer. As anticipated in the Introduction, E4S (Explaining Relatedness from the Schema) leverages the notion of *explanation template*.

Definition 10. (Explanation Template). Given a knowledge graph $G=(V, \mathcal{E}, T)$, a source entity $w_s \in G$, a target entity $w_t \in G$, a schema graph $G_s=(V_s, \mathcal{E}_s, T_s)$ and a target predicate $p^* \in \mathcal{E}_s$, an explanation template is a triple $(\text{type}(w_s), p^*, \text{type}(w_t))$ where $\text{type}(\cdot)$ is one of the types of the source (resp., target) entities.

Explanation templates are the building blocks of explanation patterns.

Definition 11. (Explanation Pattern). Given a schema graph $G_s=(V_s, \mathcal{E}_s, T_s)$, a target predicate $p^* \in \mathcal{E}_s$ and an integer d , an explanation pattern is a sequence $\mathbf{T}(p^*) = \tau_1 \cdot \tau_2 \cdot \dots \cdot \tau_d$, such that $\tau_i=(t_x \otimes p_i, t_y)$ for $i \in \{1, \dots, d\}$, $\otimes \in \{\rightarrow, \leftarrow\}$, $t_x \in \text{Dom}(p^*)$, $t_y \in \text{Range}(p^*)$, $p_i \in \mathcal{E}_s$, and \cdot denotes the composition (i.e., join) operator.

The operator $\otimes \in \{\rightarrow, \leftarrow\}$ determines the direction of an edge in the explanation pattern. By looking at Fig. 2 (b), an example of explanation patterns is $\mathbf{T}(\text{director})=\tau_1 \cdot \tau_2 \cdot \tau_3$ where $\tau_1=(\text{Work}, \rightarrow \text{starring}, \text{Actor})$,

$\tau_2=(\text{Actor}, \leftarrow\text{starring}, \text{Work}), \tau_3=(\text{Work} \leftarrow\text{director}, \text{Person})$. As it can be noted, these patterns only contain *predicates that are semantically related to the target predicate*. Fig. 9 gives an overview of the E4S algorithm.

Explanations from Schema (E4S)

Input: An explanation template ($\text{type}(w_s), p^*, \text{type}(w_t)$), an integer k , and integer d , the address of the query endpoint A

Output: A graph G_e

- (1) *Find top-k related predicates:* we are going to describe in Section 4.1 a data-driven approach to compute relatedness between predicates, which will be used to find predicates most related to p^* .
- (2) *Find Candidate Explanation Patterns:* we describe in Section 4.2 an algorithm that generates candidate explanation patterns of length d leveraging the KG schema, the types of source and target entities and predicates found in (1).
- (3) *Verify Candidate Explanation Patterns:* we discuss in Section 4.3 a SPARQL-based algorithm to verify candidate patterns that are used to build explanations.
- (4) *Select and merge paths:* paths found from verified patterns can be ranked and merged as described in Section 5.

Fig. 9. Building Explanations from the Schema.

4.1. Predicate Relatedness

The first ingredient of the E4S algorithm is a predicate relatedness measure, which given a target predicate p^* allows to focus on the part of the schema most related to p^* during the candidate explanation generation, and thus of the data during candidate verification. Given a knowledge graph $G=(V, \mathcal{E}, T)$ and a pair of predicates $(p_i, p_j) \in \mathcal{E}$, the relatedness measure takes into account the co-occurrence of p_i and p_j in the set of triples T and weights it by the predicate popularity [26]. We now introduce its building blocks.

$$PF(p_i, p_j) = \log(1 + C_{i,j}) \quad (1)$$

where $C_{i,j}$ counts the number of s and o in triples of the form (s, p_i, o) and (s, p_j, o) or (o, p_i, s) and (o, p_j, s) .

IPF is defined as follows:

$$IPF(p_j, E) = \log \frac{|E|}{|\{p_i : C_{i,j} > 0\}|} \quad (2)$$

Having co-occurrences for each pair (p_i, p_j) weighted by the IPF allows to build a co-occurrence matrix:

$$w_{i,j}(p_i, p_j, E) = PF(p_i, p_j) \times IDF(p_j, E) \quad (3)$$

The relatedness between p_i and p_j is :

$$Rel(p_i, p_j) = \text{Cosine}(W_i, W_j) \quad (4)$$

where W_i (resp., W_j) is the row of p_i (resp., p_j).

To give some hint about the results obtained by the predicate relatedness measure, Table 1 shows the top-5 most related predicates when giving as input some predicates (in bold) defined in the DBpedia ontology. We can see, for instance, that **director** is very related to **writer** and **producer**; that **spouse** is very related to **child** and **relative** and so forth. The predicate relatedness measure seems to reasonably output semantically related predicates given a target predicate.

Table 1
Examples of related predicates in DBpedia.

director	writer	spouse	influenced	birthPlace
writer	producer	child	influencedBy	deathPlace
producer	director	relative	author	residence
musicComposer	musicComposer	parent	notableWork	country
executiveProducer	composer	relation	spouse	restingPlace
starring	author	starring	child	hometown

4.2. Finding Candidate Explanation Patterns

We now describe how the E4S algorithm generates candidate explanation patterns. This process is handled via Algorithm 1, which takes as input a schema graph, a target predicate p^* , an integer d to bound the length of the pattern, an integer k to pick the top- k most related predicates to p^* , and uses a priority queue to store candidate explanation patterns ranked by their relatedness wrt the target predicate p^* .

The algorithm for each of the most related predicates (line 4) in parallel: performs a traversal of the schema graph by checking that predicate p_i and predicate p^* both have as source node t_r (lines 7) (resp., t_s (line 11)); this allows to build length-1 explanation patterns that are added to the results if they allow to reach the same node t_s (line 9) (resp., t_r (line 13)).

Input: A schema graph $G_S=(V_s, \mathcal{E}_s, T_s)$, a target predicate p^* , maximum depth d , number of related predicates k

Output: Candidate Explanation Patterns \mathbb{P}

```

1  $\mathbb{P} = \emptyset$ ; /* priority queue based on average relatedness to
    $p^*$  */
2  $R = \text{getRelatedPredicates}(p^*, k)$ 
3 Parallel execution:
4 for each predicate  $p_i \in R$  do
5   Let  $\Delta_1 = \emptyset$ 
6   for  $(t_r, p_i, t_s) \in T_s$  do
7     if  $t_r \in \text{sourceNodes}(p^*)$  then
8        $\Delta_1 = \Delta_1 \cup \{(t_r, \rightarrow, p_i, t_s)\}$ 
9       if  $t_s \in \text{targetNodes}(p^*)$  then
10         $\mathbb{P} = \mathbb{P} \cup \{(t_r, \rightarrow, p_i, t_s)\}$ 
11      end
12    end
13    if  $t_s \in \text{sourceNodes}(p^*)$  then
14       $\Delta_1 = \Delta_1 \cup \{(t_r, \leftarrow, p_i, t_s)\}$ 
15      if  $t_r \in \text{targetNodes}(p^*)$  then
16         $\mathbb{P} = \mathbb{P} \cup \{(t_r, \leftarrow, p_i, t_s)\}$ 
17      end
18    end
19  end
20  for  $j = 2, \dots, d$  do
21    Let  $\Delta_j = \emptyset$ ;
22    for each explanation pattern  $\mathbf{T}_i \in \Delta_{j-1}$  do
23      for each entity type  $n_s \in \text{outNodes}(\mathbf{T}_i)$  do
24        for each triple  $(n_s, p, n_t) \in T_s$  s.t.
25           $p \in R$  do
26             $\Delta_j = \Delta_j \cup \{\tau \cdot (n_s, \rightarrow, p, n_t)\}$ 
27          end
28          for each triple  $(n_t, p, n_s) \in T_s$  s.t.
29             $p \in R$  do
30               $\Delta_j = \Delta_j \cup \{\tau \cdot (n_s, \leftarrow, p, n_t)\}$ 
31            end
32          end
33        for each explanation pattern  $\mathbf{T}_j \in \Delta_j$  do
34          if  $(\text{checkTypes}(\tau_j, p^*) = \text{true})$  then
35             $\mathbb{P} = \mathbb{P} \cup \{\mathbf{T}_j\}$ 
36          end
37        end
38  end
39 return  $\mathbb{P}$ 

```

Algorithm 1: findCandidateExplPatterns

Length-1 explanation patterns are expanded (line 11-25). The algorithm takes a q -length (partial) pattern (with $q < d$) and expands it by only considering the top- k most related predicates to p^* (lines 19 and 21).

The expansion is done both in forward and reverse direction as the schema graph is treated as undirected. Candidate explanation patterns of at most length d are added to the results (line 25) after checking that the pattern starts from one of the types in $\text{Dom}(p^*)$ and ends in one of the types in $\text{Range}(p^*)$ (line 24).

Theorem 12. Algorithm 1 runs in time $O(d \times |E_s|)$.

Proof sketch. Algorithm 1 runs (in parallel) a d -step traversal of the graph and at each step at most $|E_s|$ edges (note that the algorithm in practice considers a subset R of all predicates) can be visited. `checkTypes` can be done in constant time by hashing domain(s) and range(s) of predicates. \square

4.3. Verifying Explanation Patterns

Candidate explanation patterns found via Algorithm 1 represent schema-compliant ways in which `type(w_s)` and `type(w_t)` could be connected in the underlying data. To understand which explanation patterns actually contribute to build a relatedness explanation there is the need to verify them. A pattern is verified if starting from w_s it is possible to reach w_t by traversing the sequence of edges in it. In other words, to verify a pattern we need to instantiate its endpoints with w_s and w_t and then check if the resulting pattern has a solution in the data. As an example, by instantiating the first explanation pattern in Fig. 1 (c), we obtain the explanation pattern in Fig. 1 (d). This pattern is actually verified as can be seen in Fig. 1. On the other hand, the last pattern in Fig. 1 (e) is not verified. We now outline our approach for candidate patterns verification, which builds upon the notion of path pattern.

Definition 13. (Path Pattern). Given a source entity w_s a target entity w_t and a candidate explanation pattern $\mathbf{T}(p^*) = \tau_1 \dots \tau_d$ with $\tau_i = (t_x, \otimes p_i, t_y)$ for $i \in \{1, \dots, d\}$, the pattern Π corresponding to $\mathbf{T}(p^*)$ is defined as follows: $\Pi(w_s, w_t, \mathbf{T}(p^*)) = (w_s, \xrightarrow{p_1}, ?v_1) \cdot (?v_1, \xrightarrow{\text{type}}, t_1) \cdot (?v_1, \xrightarrow{p_2}, ?v_2) \cdot (?v_2, \xrightarrow{\text{type}}, t_2) \dots (?v_d, \xrightarrow{p_d}, w_t) \cdot (?v_d, \xrightarrow{\text{type}}, t_d)$, where $\forall i \in \{1, 2, \dots, d\}$: $?v_i$ is a query variable, $t_i \in V_s$ an entity type and \cdot denote the composition (join) operator. Moreover, $\forall \otimes p_i \in \tau_i$ we have that $\xrightarrow{p_i}$ in $\mathbf{T}(p^*)$ becomes $\xrightarrow{p_i}$ in $\Pi(w_s, w_t, \mathbf{T}(p^*))$ if $\otimes = \rightarrow$ (resp., \leftarrow becomes $\xleftarrow{p_i}$ if $\otimes = \leftarrow$).

A path pattern basically transforms a candidate explanation pattern into a SPARQL graph pattern [27] where: (i) individual triple patterns (viz. τ_i) are joined

thanks to shared variables; (ii) an additional triple pattern checks that, for each τ_i , the bindings (viz. entities in the KG) of the variable $?v_i$ are of the right `type`; (iii) the endpoints are replaced with the source (w_s) and target entity (w_t). As an example, the candidate explanation pattern $\mathbf{T}(\text{director})=\tau_1 \cdot \tau_2 \cdot \tau_3$ where $\tau_1=(\text{Work}, \rightarrow\text{starring}, \text{Actor})$, $\tau_2=(\text{Actor}, \leftarrow\text{starring}, \text{Work})$, $\tau_3=(\text{Work} \leftarrow\text{director}, \text{Person})$ gives the path pattern $\Pi(\mathbf{T}(\text{Metropolis}, \text{F. Lang}, \text{director}))=\text{Metropolis starring } ?v_1. ?v_1 \text{ type Actor. } ?v_2 \text{ starring } ?v_1. ?v_2 \text{ type Work. } ?v_2 \text{ director F. Lang.}$

4.3.1. Building a SPARQL Query for the Verification

At this point, from a path pattern we can build a SPARQL query whose evaluation allows to check whether the original candidate explanation pattern is verified and, if so, to get a set of paths that conform to it. The prototypical SPARQL query has the form:

```
SELECT DISTINCT * WHERE {
   $\Pi(w_s, w_t, \mathbf{T}(p^*))$ 
}
```

As an example, the SPARQL query obtained from the path pattern $\Pi(\text{Metropolis}, \text{F. Lang}, \mathbf{T}(\text{director}))$ is:

```
SELECT DISTINCT * WHERE {
  Metropolis starring ?v1. ?v1 type Actor.
  ?v2 starring ?v1. ?v2 type Work.
  ?v2 starring F. Lang.
}
```

E4S has the advantage of requiring the evaluation of SPARQL queries more specific than those required by E4D (see Definition 7) as predicates are fixed and the bindings of variables are forced to be of specific types. Moreover, note that since candidate explanation patterns work at the level of schema, they do not need to be generated from scratch every time. Indeed, they can be reused when the types of the input pair and the target predicate is the same.

5. Ranking Paths to build Explanations

We have described two different algorithms to find paths connecting two entities w_s and w_t in a KG. In both cases we may have that the number of such paths can be prohibitively large, which can hinder the process of making sense of an explanation. As an example, considering the merge of all paths, as done in minimal explanations (see Definition 6) may be an obstacle toward providing concise explanations to the users. To cope with this issue, we are going to introduce differ-

ent criteria to rank paths. Having a rank among paths gives flexibility in choosing the set of paths that will form an explanation. Cheng et al. provide an overview on the topic [28].

5.1. Path Informativeness

The first approach to rank paths leverages the notion of informativeness. Given a graph G , the informativeness of a path connecting a pair of entities $(w_s, w_t) \in G$ is estimated by investigating the informativeness of its constituent RDF predicates. This reasoning is analogous to that of associating relative importance to words in a document [29] contained in a set of documents. We leverage the notion of *Predicate Frequency Inverse Triple Frequency* (`pf itf`) [26].

Definition 14. (Predicate Frequency). Given $G = (V, \mathcal{E}, T)$, an entity $w \in V$ and a predicate $p \in \mathcal{E}$ appearing in some triple involving w , the Predicate Frequency (`pf`) quantifies the informativeness of p wrt w . We distinguish between incoming $\text{pf}_i^w(p)$ and outgoing $\text{pf}_o^w(p)$:

$$\text{pf}_i^w(p, G) = \frac{|T_i|_p}{|T_i|} \quad \text{pf}_o^w(p, G) = \frac{|T_o|_p}{|T_o|} \quad (5) \quad (6)$$

where $|T_i|_p$ (resp., $|T_o|_p$) is the number of triples in G where the predicate p is incoming (resp., outgoing) in w and $|T_i|$ (resp., $|T_o|$) is the total number of incoming (resp., outgoing) triples in which w is involved.

Definition 15. (Inverse Triple Frequency). Given $G=(V, \mathcal{E}, T)$ and a predicate $p \in \mathcal{E}$, the inverse triple frequency of p `itf`(p) is defined as:

$$\text{itf}(p, G) = \log \frac{|T|}{|T|_p} \quad (7)$$

where $|T|$ is the total number of triples in G and $|T|_p$ the total number of triples containing p .

Definition 16. (Predicate Frequency Inverse Triple Frequency). The Predicate Frequency Inverse Triple Frequency is defined as $\text{pf itf}(p, G)=\text{pf} \times \text{itf}$. One can consider the *incoming* $\text{pf itf}_i(p, G)=\text{pf}_i \times \text{itf}$ and *outgoing* $\text{pf itf}_o(p, G)=\text{pf}_o \times \text{itf}$ case.

An example of `pf itf` computation is shown in Fig. 10. Having a way to assess the relative informativeness of a predicate wrt to an entity in a path, we can now define the informativeness of a path.

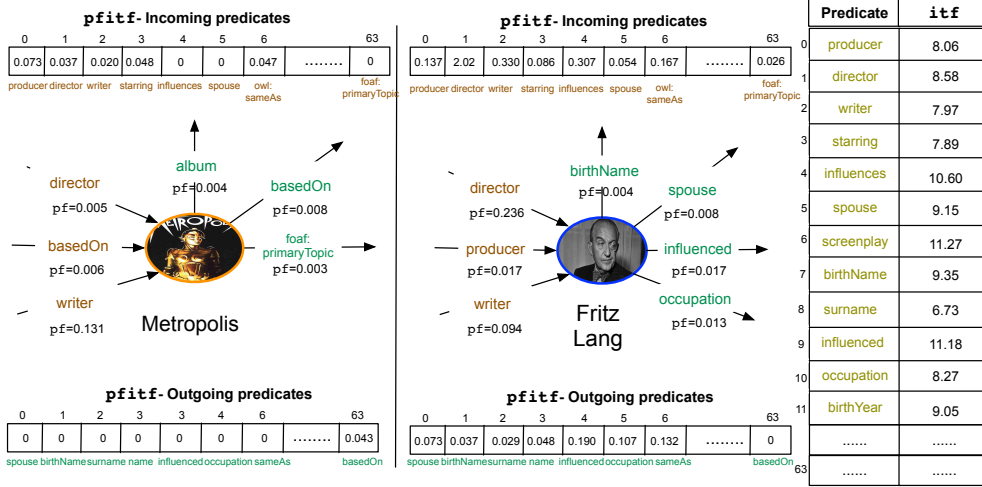


Fig. 10. An example of pfitf computation.

Definition 17. (Path Informativeness). Let ξ be a path of length k . For $k=1$, the informativeness of ξ is:

$$I(\xi, G) = [\text{pfitf}_o^{w_s}(p) + \text{pfitf}_i^{w_t}(p)]/2 \quad (8)$$

It considers the predicate p as outgoing from w_s and incoming to w_t . The second case for $k=1$ is $\xi(w_s, w_t) = w_s \xleftarrow{p} w_t$ where p is an incoming edge for w_s and outgoing from w_t and can be treated in a similar way. For paths of length $k > 1$, the informativeness is computed as follows:

$$I(\xi, G) = \frac{I(\xi(w_s, w_1)) + \dots + I(\xi(w_k, w_t))}{k} \quad (9)$$

5.2. Explanation Pattern Informativeness

We have discussed a method to compute path informativeness based on pfitf values obtained by considering nodes (entities), edges (predicates) and their directions in a path. We now introduce a way of assessing informativeness using path patterns (see Definition 13). Intuitively, a path pattern *summarize*, via variables, all paths that conform to a specific sequence of predicates (according to their directions.)

As an example, the path pattern (Metropolis, starring, $?v_1$) · ($?v_1$, type, Actor) · ($?v_2$, starring, $?v_1$) · ($?v_2$, type, Work) · ($?v_2$, director, F. Lang) captures different paths where actors (e.g., R. Klein-Rogge) starred other films directed by F. Lang besides Metropolis. According to DBpedia there are 10 such paths that involve actors like A. Abel and B. Helm. In the case of E4D, a set of paths having the same set of edges along with their di-

rections and a different set of nodes can be abstracted into a path pattern by replacing nodes with variable. In the case of E4S, (verified) path pattern are simply obtained from candidate patterns (see Definition 13). We now define path pattern informativeness.

Definition 18. (Path pattern informativeness). Let P be the set of path patterns and G a KG. The informativeness of a path pattern $\Pi \in P$ is defined as:

$$I(\Pi, G) = \log \frac{|P|}{|(\Pi, G)|} \quad (10)$$

where $|(\Pi, G)|$ is the number of paths sharing Π in a graph G . In other words, $|(\Pi, G)|$ is the number of different bindings in the solutions to the SPARQL query used to verify Π . The usage of path patterns enables to abstract information to be included into an explanation by only focusing on specific patterns.

5.3. Path Diversity

We have discussed two approaches for ranking paths based on informativeness. We now introduce a different approach, which takes into account the *variety* of predicates in a set of paths. As an example, among the set of paths of length 3 between Metropolis and F. Lang often include predicates related to the fact that there are actors that (besides Metropolis) starred in other films directed by F.Lang; this will potentially rule out other predicates that appear in some paths if the informativeness of such paths is lower.




Fig. 11. Path ranking examples: (a) *most* informative paths; (b) *least* informative path patterns; (c) most *diverse* path.

Definition 19. (Path Diversity). Given a source entity w_s , a target entity w_t and two paths $\xi_1(w_s, w_t)$ and $\xi_2(w_s, w_t)$ we define path diversity as:

$$\delta(\xi_1, \xi_2) = \frac{|\text{Labels}(\xi_1) \cap \text{Labels}(\xi_2)|}{|\text{Labels}(\xi_1) \cup \text{Labels}(\xi_2)|} \quad (11)$$

The above equation computes the Jaccard index between paths in terms of their constituent RDF predicates. The above definition can be used to compute the pairwise diversity of a set of paths \mathcal{P} . This in general requires $O(|\mathcal{P}|^2)$ computations. To avoid pairwise computations, one can use min-wise hashing [30].

5.3.1. Path Ranking Strategies: an example

Consider the entity pair (Metropolis, F. Lang) and the graph G shown in Fig. 1 (d), which includes a set of paths some of which resulting from the verification of the candidate explanations patterns in Fig. 1 (d). In some cases, presenting such a set of paths altogether may hinder the interpretation of the explanation by the user. Hence, the ranking of paths results very useful.

Fig. 11 shows examples of paths obtained according to the three strategies for path ranking introduced before. In particular, Fig. 11 (a) shows the two most informative paths according to the `pfidf` (see Section 5.1). For $k = 1$ the path involves the predicate `director` while for $k = 2$, the informativeness takes into account the fact that Metropolis has been written by T. von Harbour who was married to F. Lang.

Fig. 11 (b) shows two path patterns. The first pattern involves the predicates `producer` with intermediate nodes representing People and Films. Information about the 6 movies produced by E. Pommer and directed by F. Lang is conveniently abstracted by the pattern plus the type of entity abstracted, which enables to reduce the size of the explanation shown to the user. The second path pattern allows to abstract the fact that 11 movies have been co-written by T. von Harbour and F. Lang.

Fig. 11 (c) shows the most diverse path, which includes the predicates `basedOn` and `author` that were discarded when considering path informativeness. As it can be observed, the three strategies for path ranking

give great flexibility in selecting the set of paths to form an explanation.

5.4. Selecting and Merging Paths

Table 2 summarizes different strategies for building explanations. The six strategies reported in the table (but E^\cup), given a value m , select a subset of paths (pattern) according to one of the three strategies described in Section 5. Moreover, two strategies combine path (pattern) informativeness and diversity. The last line in Table 2 refers to all paths without merging them. The merge of paths consists in considering all entities (resp. nodes) sharing the same identifier as the same node (resp. edge) in the explanation graph G_e .

Table 2

Path selection strategies.

Symbol	Meaning
E^\cup	Merge all of paths without any pruning
E_m^ξ	Merge the <i>top-m</i> most informative paths
E_m^{ξ}	Merge paths belonging to the <i>top-m</i> most informative path <i>patterns</i>
E^δ	Merge pairs of paths whose value of diversity falls in $[max, (max - r)]$ where <i>max</i> is the max diversity value over all pairs of paths and <i>r</i> is a % value.
$E^{\xi, \delta}$	Merge the results of E_m^ξ and E^δ
$E^{\bar{\xi}, \delta}$	Merge the results of $E_m^{\bar{\xi}}$ and E^δ
\mathcal{P}	Set of all paths (without merge)

6. The RECAP tool and its application for Querying Knowledge Graphs by Relatedness

We now outline the implementation of our explanation framework (Section 6.1) and then describe a concrete application of explanations for querying KGs (Section 6.2).

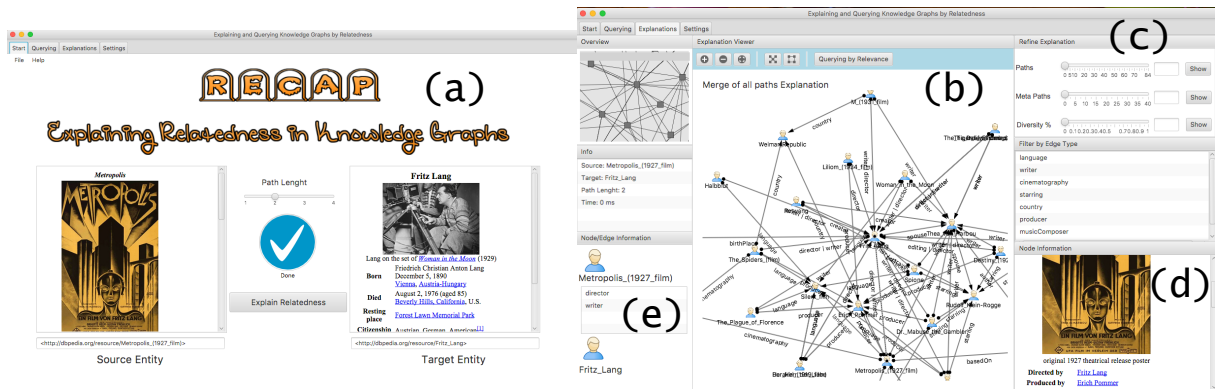


Fig. 12. The RECAP tool.

6.1. Implementation

We have implemented our explanation framework in Java in a tool called RECAP, which leverages the Jena⁴ framework to handle RDF data and SPARQL to access KGs via their endpoint. As previously described, our explanation algorithms make usage of queries with the SELECT query form for path finding and candidate explanation pattern verification; SPARQL queries with COUNT are used for path ranking. The tool along with all the datasets are available upon request.

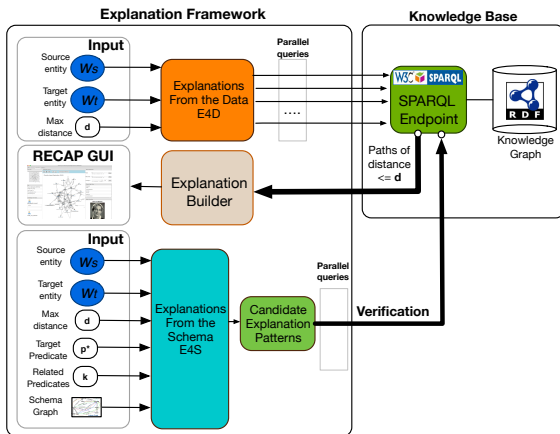


Fig. 13. The RECAP architecture.

Fig. 13 gives an overview of the architecture of the system. It consists of three main components. The *E4D* component implements the E4D algorithm (Section 3) and is responsible for both generating and executing (in parallel) queries to retrieve paths. The *E4S*

component implements the E4S algorithm (Section 4). In this case, candidate explanation patterns are generated after building the schema graph from a given KG schema in input. The verification is done (in parallel) via SPARQL and allows to retrieve a set of paths. The *Explanation Builder* takes a set of paths either from E4D or E4S and ranks them according to the input method chosen by the user. Paths are then merged to build an explanation, which is passed to the RECAP GUI. Fig. 12 (a) shows the main GUI of the tool. The user can specify the pair of input entities and get a short description.

Fig. 12 (b) shows the part of the GUI that deals with explanations. It is possible to select different kinds of explanations (Fig. 12 (c)), and obtain a description of each entity (Fig. 12 (d)) and edge (Fig. 12 (e)) in the explanation graph. RECAP goes beyond related approaches (e.g., REX [18], Explan [20]) that provide visual information about connectivity as it allows to build different types of explanations (e.g., graphs, sets of paths), thus controlling the amount of information visualized. RECAP has the advantage of not requiring any data preprocessing; information is obtained by querying (remote) SPARQL endpoints.

6.2. Querying Knowledge Graphs by Relatedness

As a concrete application of explanations, we now describe their usage for querying KGs. An explanation $E = \langle w_s, w_t, G_e \rangle$ can be seen as a graph including paths that link a source and target entity. The idea is to turn an explanation graph G_e for an entity pair (w_s, w_t) into an explanation pattern possibly shared by other entities. Fig. 15 summarized our approach for querying KGs using relatedness explanations.

⁴<http://jena.apache.org>

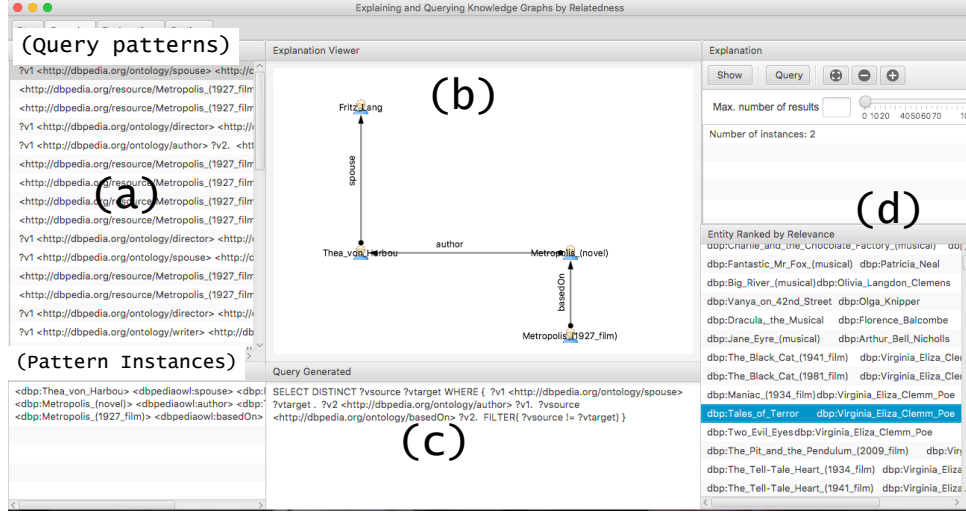


Fig. 14. Querying Knowledge Graphs Using Relatedness Explanations. Path patterns (a), explanation (b), SPARQL query (c), and suggested entities (d) ranked by popularity (PageRank [31] in this case).

Knowledge Graph Querying using Explanation

Input: A pair (w_s, w_t) of entities, an integer k , the address of a query endpoint A

Output: A set of ranked (pairs of) entities

- (1) Find an explanation $E=(w_s, w_t, G_e)$ between w_s and w_t by using E4D or E4S.
- (2) Build the entity query pattern Q_e associated to G_e .
- (3) Query the KG with Q_e (via A) and get a set of (pairs of) entities.
- (4) Rank the answers to Q_e .

Fig. 15. An overview of the query answering algorithm.

We now describe step (2) and (3) that allow to obtain the entity query pattern Q_e from an explanation graph G_e and, from it, a set of entity pairs. Step (4) is described in Section 6.2.

Definition 20. (Explanation Pattern). Let $E=(w_s, w_t, G_e)$ be a relatedness explanation with $G_e=(V^e, \mathcal{E}^e, T^e)$. Let $f_n: V^e \rightarrow \mathcal{V}$ a function such that $f_n(v_i)=?v_i$ where \mathcal{V} is a set of variables. An explanation pattern is a tuple $\bar{E}=(?w_s, ?w_t, G_e^v)$ where $?w_s, ?w_t \in \mathcal{V}$ and $G_e^v = (V^v, \mathcal{E}, T^v)$ with $V^v = \{f_n(v_i), \forall v_i \in V^e\}$ and $T^v = \{(f_n(s_i), p_i, f_n(o_i)), \forall (s_i, p_i, o_i) \in T^e\}$

G_e^v is the query graph obtained by replacing all nodes in G_e with query variables. Basically, an explanation pattern generalizes the structure of an explanation by keeping edge labels only. Explanation patterns are used to generate entity query patterns.

Definition 21. (Entity Query Pattern). Let $\bar{E}=(?w_s, ?w_t, G_e^v)$ be an explanation pattern with $G_e^v = (V^v, \mathcal{E}, T^v)$ and $t_i \in T^v, i \in [1, k]$. An entity query pattern Q_e is a SPARQL query of the form:

```
SELECT DISTINCT ?w_s ?w_t WHERE { t_1 . t_2 . ... t_k }
```

In the above definition, t_1, \dots, t_k are triple patterns and $.$ denotes the join operator; moreover, the variables $?w_s$ and $?w_t$ used in lieu of the entities in input. Note that query patterns are automatically derived from explanations; our approach neither requires familiarity with SPARQL nor with the underlying data/schema. The evaluation of a query pattern returns a set of pairs of entities sharing the same relatedness perspective of the input pair.

6.2.1. Ranking of Results

As the number of results of Q_e can be large, we describe an approach for ranking results. The problem of ranking results of SPARQL queries has been already studied (e.g., [32, 33]) and is not the main purpose of the present paper. Inspired by the Google KG, we consider a simple result ranking mechanism based on the popularity of entities; specifically, we leverage the PageRank [31] algorithm. Given a pair of entities (w_1, w_2) returned when evaluating Q_e , we estimate their popularity as $(PR(w_1) + PR(w_2)) / 2$, where $PR(w_i)$ is the PageRank value of the entity i .

Fig. 14 shows the GUI that handles querying answering using relatedness explanations. In particular, Fig. 14 (a) shows path patterns (see Definition 13) while Fig. 14 (b) shows the instantiation of the pat-

tern selected. We can see that T. von Harbou was married with F. Lang and she wrote the novel Metropolis on which the film Metropolis was based. Fig. 14 (c) shows the automatically generated SPARQL query (see Definition 21). The underlying idea of our approach is to look for other pairs of entities that share a similar relatedness perspective. In other words, are there other pairs beside (Metropolis, F. Lang) for which there exist a spouse that wrote a novel that inspired a film? Fig. 14 (d) shows entity pairs that share the same relatedness perspective. As an example, the pair (Tales of Terror, V. E. C. Poe) shares the perspective since the movie Tales of Terror was based on the novel Tales of Terror written by E. A. Poe with whom V. E. Clemm Poe was married. Ditto for the pair (The Black Cat, V. E. Clemm Poe).

7. Evaluation

We describe an experimental evaluation of our explanation framework having the following main goals:

- G1:** Investigate the performance and scalability of E4D, our knowledge-graph-agnostic approach for building explanations, which allows to: (i) use any SPARQL-enabled KG as a source of background knowledge; (ii) avoid to set up complex processing infrastructures (in terms of data storage and computing power); (iii) work with fresh data without having to replicate and keep them up to date locally.
- G2:** Compare E4D, which builds explanations directly by looking at paths in the data with E4S, which starts from candidate explanation patterns found in the KG schema. This will allow to understand the trade-off between working with all paths (found by E4D) and performing a later refinement and selecting explanations patterns that are of interest beforehand and then verify only those in the data (via E4S).
- G3:** Investigate the usability of the RECAP tool as compare to related tools. This allows to have an account for the practical usability of our tool.

We describe the datasets and experimental setting in Section 7.1. Then, in Section 7.2 we tackle **G1**. We tackle **G2** in Section 7.3. Finally, **G3** is tackled in Section 7.4.

7.1. Datasets and Experimental Setting

There is no standard benchmark for evaluating relatedness explanation algorithms and tools. Neverthe-

less, some query-sets have been defined to evaluate the discovering of semantic associations, that is, paths between an entity pair. We considered a dataset of 74 entity pairs⁵, by extending a dataset of 24 pairs defined by Cheng et al. [20]. The dataset covers a fair large number of knowledge domains (i.e., entertainment, sport, movies, companies) and thus can give insights about the performance of our approach in different scenarios. Table 3 summarized the algorithmic parameters considered.

Table 3
Algorithm parameters and their values

Parameter	Value	Meaning
k	1-3	Max. path length
m	5	Top- m most informative paths
p	5	Top- q related predicates used in E4S
r	0.5	Diversity range

Note that $k \leq i$, $i \in [1,3]$ means that in the same run we consider paths of length up to i . As we will discuss in Section 7.4, paths of length $k > 3$ provide little insight on the relatedness of the input pair as there is a very large number of such paths that involve intermediate nodes and edges that have very little in common.

All the experiments were performed on a MacBook Pro with a 2.8 GHz i7 CPU and 16GBs RAM. Results are the average of 3 runs.

Table 4
Knowledge Graphs used in the evaluation.

KG	#triples	Query Endpoint (A)
DBpedia	~438M	http://dbpedia.org/sparql
Wikidata	~5B	http://query.wikidata.org/sparql

7.1.1. Knowledge Graphs Used

We evaluated our explanation building algorithms by considering two different knowledge graphs (summarized in Table 4). Table 4 also reports the address of the remote (SPARQL) endpoint used to access the KG. The first is KG is DBpedia, which includes more than 400 millions of triples⁶ while the second one is Wikidata, which includes more than 5 billions of triples⁶.

⁵Reported in reported in Appendix B

⁶This number has been obtained by submitting a SPARQL count query to the endpoint.

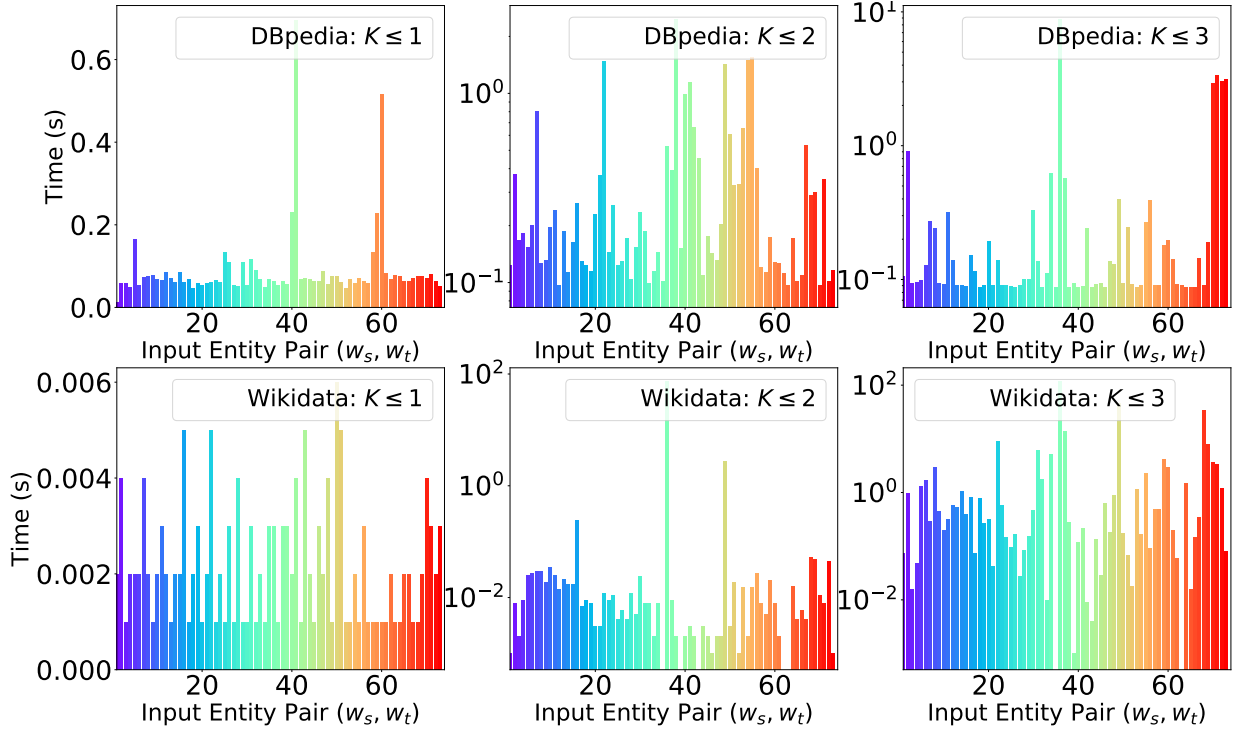


Fig. 16. E4D: Time to find paths of increasing length in DBpedia and Wikidata.

7.2. E4D: Finding Explanations From the Data

We start our analysis of the performance of E4D by looking at the running times of finding paths of increasing length.

7.2.1. E4D: Running time

Fig. 16 shows the running times of E4D while Fig. 17 plots the number of paths found for different values of k . Recall that for each k our algorithm retrieves paths of length up to k . In the figures, each input pair (x -axis) is identified by a different color. For $k \leq 1$ the running time is below one second on both DBpedia and Wikidata. However, we can note in Fig. 17 that only $\sim 1/3$ of the input pairs have at least a path of length 1 (i.e., a direct link). For $k \leq 2$, on DBpedia, we can see that the number of paths stays in the order of tens for most pairs with a few pairs having a relatively large number of paths (e.g., pair 38, that is, (France, Paris)).

On Wikidata the number of paths is almost an order of magnitude larger with some pair (again pair 38) having a very large number of paths. By comparing DBpedia and Wikidata we can observe a similar trend, although translated in terms of absolute numbers since in the Wikidata KG pairs have a larger number of paths interlink them. The running time in this case stays in

the order of a second on both DBpedia and Wikidata with a few pairs requiring a larger time. In particular, the pair (France, Paris) is somehow a corner case; it is reasonable to have a very large number of path between the very general entities France and Paris. However, we believe that relatedness explanations are much more interesting for other kinds of, less obviously related, input pairs as for instance pair 74 (C. Manson, Beach Boys). For $k \leq 3$, we notice the same correlation between running times and number of paths both on DBpedia and Wikidata. In the case of the latter, running times are now in the order of ten seconds (excluding some outliers). This time the number of paths is in the order of hundreds for most pairs with a larger number of entities having thousands of paths; in Wikidata the number of paths is in general much larger. This comes as no surprise due to the different sizes of the KGs, counting $\sim 0.5M$ and $\sim 5B$ of triples, respectively. We want to stress the fact that running times reported refer to queries on a remote endpoint, with all the network delays included and no data filtering.

The aim of relatedness explanations is to provide concise (visual) graphs that can shed light about why two entities are connected. Therefore, in what follows we investigate the size of explanations.

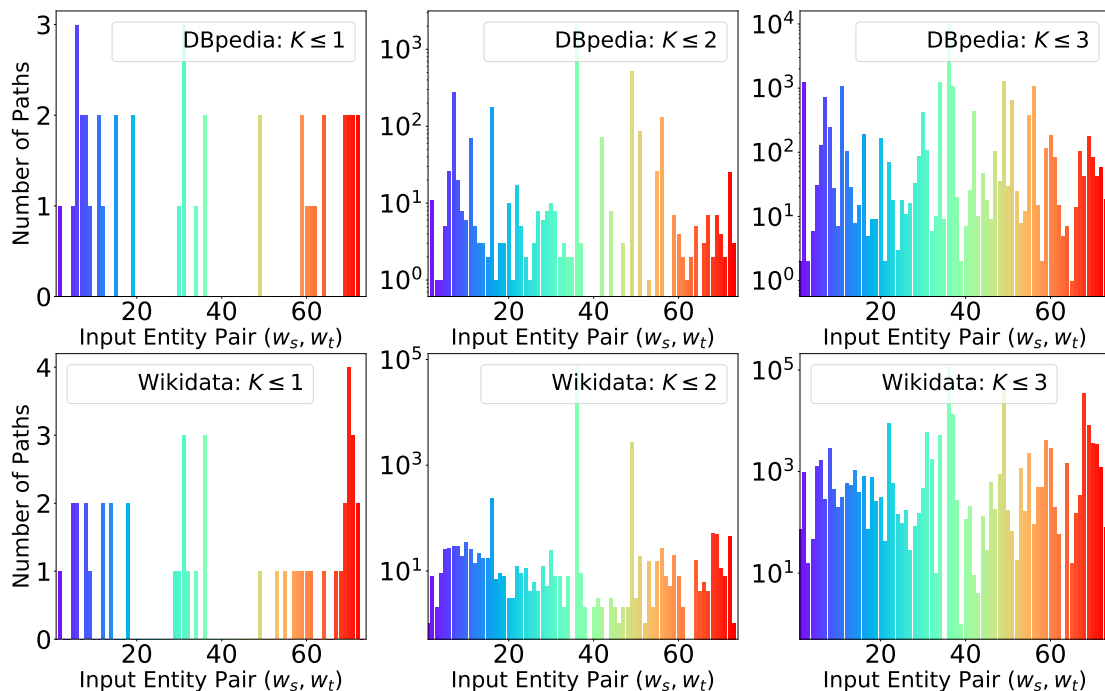


Fig. 17. E4D: Number of paths of increasing length in DBpedia and Wikidata.

7.2.2. E4D: Explanation Size

Fig. 18 shows the explanation size in terms of number of nodes and edges. Explanations considered are built by merging all paths and are referred to as E^U (see Table 2). We can notice that for $k \leq 1$, there are a few available explanations, having between 2 and 4 nodes (and edges). These numbers are consistent in both DBpedia and Wikidata.

On DBpedia, for $k \leq 2$, most explanations have between 10 and 20 nodes with a few exceptions where the number of nodes can be up to 80. On Wikidata, the cluster is centered between 1 and 50 nodes with a few having more than 100 nodes. It is interesting to see that the number of edges is larger in Wikidata than in DBpedia. Nevertheless, for $k \leq 2$ the clusters of nodes and edges mostly have the same shapes.

For $k \leq 3$ explanations have a larger size variety on both DBpedia and Wikidata. We can notice that the number of nodes and edges is in the order of hundreds for most pairs. Note that now the distribution is more scattered, especially on Wikidata where the majority of explanations have between 200 and 600 nodes and between 50 and 150 edges.

We want to stress the fact that already for $k \leq 3$, most explanations become very large. This analysis suggests that considering larger values of k even worsen the situation. Moreover, considering explana-

tions built by the merge of all paths may hinder the understandability of an (visual) explanation already for $k \leq 3$. Therefore, having the possibility to control the amount of information to put into an explanation becomes crucial. Our path ranking strategies specifically deal with this aspect.

7.2.3. Refining Explanations

We now discuss explanations built by considering a subset of paths interlinking an entity pair.

In terms of running time, explanations based on path informativeness E_m^ξ (we discuss the case $m=5$) require to compute `profit` scores; our approach computes these scores in parallel and leverages the merge of all paths as underlying graph (see Definition 14) thus not requiring to perform any remote query. Explanations based on pattern informativeness $E_m^{\tilde{\xi}}$ (we discuss the case $m=5$) are less expensive since they do not analyze the informativeness of all edges in a path. Again, to count the instances satisfying a path we use the merge of all paths as reference graph (see Definition 18). Building these kinds of explanations had a minor impact (i.e., between 0.3 and 1 seconds) on the overall running time. We do not report times since they are almost the same as those already discussed.

What changes is the size of explanations built by considering a subset of paths only. Indeed, E_5^ξ , based

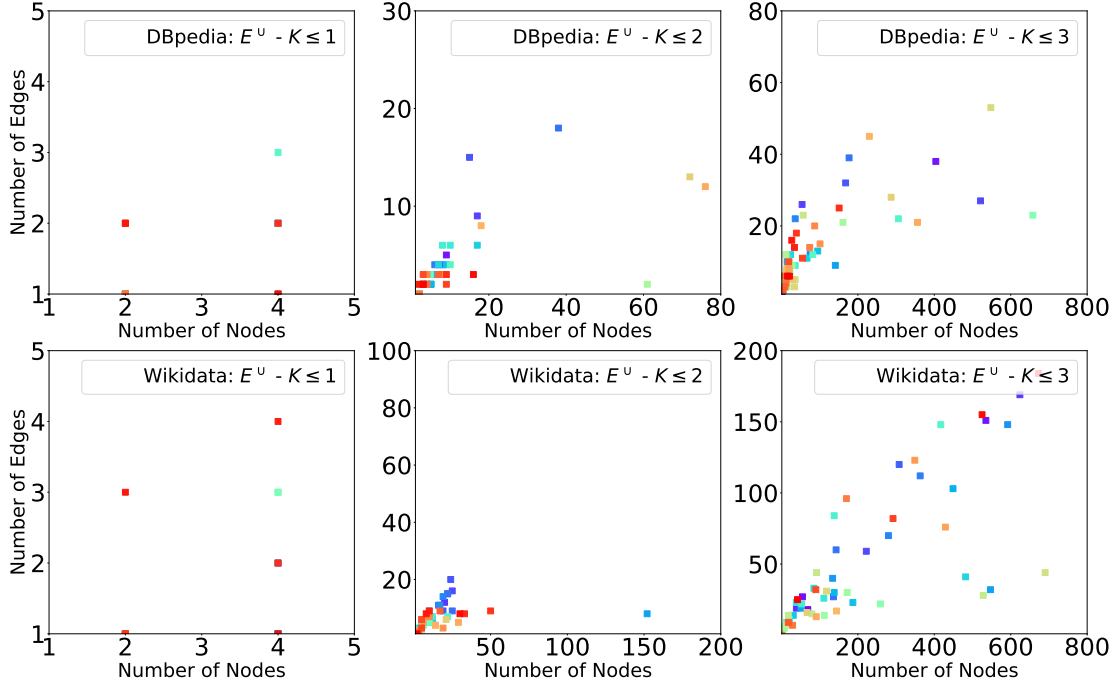


Fig. 18. Size of E^U explanations in DBpedia and Wikidata consisting in the merge of all paths found by E4D.

on path informativeness, are much smaller; the typical size is ~ 8 nodes and ~ 7 edges on DBpedia and ~ 15 nodes and ~ 12 edges on Wikidata, when $k \leq 2$. On the other hand, $E_5^{\bar{\xi}}$, based on path pattern instance, have variable size as it depends on the number of paths for each of the top-5 most informative patterns. In general these are bigger than E_5^{ξ} explanations (~ 15 nodes and ~ 12 edges on DBpedia and ~ 20 nodes and ~ 25 edges on Wikidata). Note that $E_5^{\bar{\xi}}$ explanations enable to focus on specific aspects as they include all the instantiations of each of the top-5 most informative path patterns. We will come back on this aspect in Section 7.3.2 when discussing E4S, which allows to chose patterns before finding paths instead of emerging and selecting such patterns after discovering paths as done by E4D.

When now discuss explanations built by using a diversity criterion, that is, E^δ . E^δ explanations are a bit more expensive (between 1 and 3 seconds) in terms of running time since they require the computation of distances between paths, even if this task leverages a multi-thread approach. In terms of size, these explanations are in the same order of magnitude as $E_5^{\bar{\xi}}$ on DBpedia, while are a bit larger in Wikidata. We hypnotize that this behavior is due to a larger number of predicates present in Wikidata, which allows to consider

a higher number of paths with the same δ value than in DBpedia. The important point of E^δ explanations is that they guarantee to also include rare edges potentially discarded by path or pattern informativeness. We also looked into explanations combining (top-5) path/pattern informativeness and diversity ($r=25\%$). In this case the typical size is ~ 20 nodes and ~ 15 edges in DBpedia and a bit larger in Wikidata. The flexibility of our approach to build explanations allowing to decide the amount of information to be included into an explanation is crucial toward understanding relatedness between entities.

7.3. E4S: Explanations from the Schema

We now discuss the E4S algorithm, which makes usage of the KG schema. We took the latest version of the DBpedia ontology⁷ and considered statements about `rdfs:type`, `rdfs:subClassOf`, `rdfs:subPropertyOf` `rdfs:domain` and `rdfs:range` only. In fact, these are at the core of the notion of schema graph closure that is used by E4S (see Definition 2) to generate candidate explanation patterns. In order to generate candidate explanation patterns, the first of E4S step is to identify seed predicates. While with our RECAP tool the user can

⁷http://downloads.dbpedia.org/2014/dbpedia_2014.owl.bz2




Fig. 19. E4S: Number of queries generated to verify candidate explanation patterns for each input pair.

select one or more seed predicates, to run the experiments we automatized the identification of seed predicates. For each input pair, we took the top-6 most informative predicates (see Definition 15) common to the source of target entity; when there were no predicates in common, we took the top-3 most informative predicates from the source and target entity, respectively. Then, we ran Algorithm 1 by using each of the 6 seed predicates as input and obtained an overall set of candidate explanation patterns; from these, we obtained a set of SPARQL queries used to verify such patterns (see Section 4.3.1).

Fig. 19 shows the number of SPARQL queries generated for each input pair. We observe that the number of queries varies from pair to pair and can reach the order of thousands when $k > 1$. These queries contain predicates that are related to the seed predicate in the “middle” position of triple patterns. In contrast, with E4D we have a fixed number of 2^k queries containing variables in the predicate position of triple patterns.

The peculiarity of E4S is that it allows to reduce the number of candidates queries by focusing on those that are interesting before verifying them in the data; where interestingness can be defined according to e.g., informativeness and/or diversity. In what follows we investigate this aspect.

7.3.1. E4S: Running Time

We now report the running time of E4S in two different settings: (i) when considering all queries shown in Fig. 19; (ii) when only considering the top-5 that have the higher semantic relatedness wrt the seed predicate.

Running times are shown in Fig. 20. We can notice that running times when considering all queries are larger than those obtained by E4D. This comes as no surprise since the number of queries is much larger in this case. Nevertheless, getting reasonably (and visually) understandable explanations from E4D requires to

apply some posteriori filtering criterion like considering top- k paths, patterns, or applying a diversity criterion. Hence, the final cost of building explanations via E4D should also consider the cost of the refinement.

On the other hand, E4S allows to decide a filtering criterion beforehand, thus drastically reducing the number of queries that have to be verified. To further investigate this aspect, we ran additional experiments by verifying only the top-5 queries. In particular, we start the verification from the most semantically related wrt the target predicate and proceed with the query verification until 5 queries were verified (i.e., returned some path). Running times in this case are reported in the bottom part of Fig. 20; these are now in the same order of magnitude as those reported by E4D. This fact is more evident as k grows. As an example, for $k \leq 3$, running times now never exceed ten seconds. Filtering the queries to be verified introduces a significant reduction of the running times of E4S bringing the values to be compatible with those of E4D. Nevertheless, recall that E4D still requires some pruning in order to produce (visually) understandable explanations while those produced by E4S (when applying query filtering) potentially contain less nodes and edges as these are built considering specific predicates. We further investigate this aspect in the next section.

7.3.2. E4S: Explanation Size

Fig. 21 reports that size of explanations found by E4S when considering the verification of the top-5 most semantically related queries to the seed predicate. We note that these explanations contain a few edges while the number of nodes can reach the hundreds (when $k \leq 3$). It is interesting to compare the explanation sizes with those found by E4D (reported in Fig. 18). The size reduction is clear; for instance, when $k \leq 2$ the number of nodes is of at most ~ 15 for E4S while it can be up to ~ 80 for E4D. This difference is




Fig. 20. Running times of E4S when considering all queries (top) and the top-5 most related to the seed predicate (bottom).

more evident when $k \leq 3$ where we have that the number of nodes for E4D can be up to ~ 700 .

We want to mention that explanations found by E4S (when filtering the queries) contain a subset of all nodes and edges that can be found when considering all paths (as done by the merge of all paths via E4D or evaluating all queries generated by E4S). The main point of E4S is exactly to allow the generation of focuses explanations, that is, explanations whose filtering criterion has been decided a priori and that can be directly “parsed” by the user.

7.4. Explanation Building Tools

We now report on the capabilities of the RECAP tool in providing visual relatedness explanations and its comparison with related tools. In particular, we considered RelFinder⁸, and Explass⁹ that are tools publicly available. As both Explass and RelFinder only work on DBpedia, we used DBpedia in the comparison.

⁸<http://www.visualdataweb.org/relfinder/relfinder.php>

⁹<http://ws.nju.edu.cn/explass/>

Setting. Twenty participants were assigned each six random pairs among the 26 entity pairs described in [20]. They were shown how the three systems work and asked to use each system (with no other support) in order to understand the relatedness between entities in a pair. Following the methodology in [20] participants were given a set of six questions (reported in Appendix A); the response to each question was given with an agreement value from 1 (min) to 5 (max). Q6 was not considered in [20]; we included it to understand how users perceive the performance of the systems in terms of running time. Results are reported in Table 5.

According to questions Q1-Q5, users perceived RECAP and Explass as better supports to the explanation task. Users reported that RelFinder does not allow the flexible creation of explanation (e.g., by grouping paths into patterns), which makes it hard to control the amount of information shown. In general, RECAP was judged to be a more comprehensive solution; it provides both a graph-based and pattern-based exploration of results and several ways of controlling the amount of information to be shown. While RECAP and RelFinder quickly provide information immedi-




Fig. 21. E4S: Explanation size on DBpedia for the top-5 most related queries.

Table 5
Questions/responses: means (standard deviation).

Question	RECAP	RelFinder	Expass
Q1: Information overview	4.55(0.65)	3.05(0.77)	3.82(0.75)
Q2: Easiness in finding information	4.45(0.55)	4.05(0.63)	3.85(0.67)
Q3: Easiness in comparing/synthesizing info	4.62(0.62)	3.10(0.82)	4.06(0.61)
Q4: Comprehensive support	4.81(0.73)	3.42(0.77)	4.15(0.79)
Q5: Sufficient support to the task	4.67(0.81)	3.28(0.86)	4.23(0.83)
Q6: Running time	4.82(0.48)	4.12(0.72)	3.18(0.52)

ately after retrieving paths, Expass requires a much longer time. On Q6, Expass was judged to be the less compelling system. RECAP was judged higher than the other two systems in all questions via LSD post-hoc tests ($p < 0.05$). The inter-annotator agreement was of 0.85.

Combining multiple KGs. We asked users to also test the capability of RECAP on the *combination* of DBpedia and Yago. In particular, starting from entities in DBpedia, we looked at `owl:sameAs` links to the corresponding Yago entities. We then merged the explanations obtained from the two KGs. Users ($\sim 75\%$) perceived the combination of multiple KGs as very useful toward more comprehensive explanations. This is especially true when KGs cover the same domain with different levels of detail (DBpedia was judged more comprehensive than Yago). The combination also produces graphs of bigger size. Indeed, the functionality of RECAP allowing to filter information to be included into an explanation was judged very useful (participants thought that the merge of all paths explanations E^U were too big already with $k \leq 2$).

8. Related Work

We review related work along different perspectives. Our analysis ranges from techniques to build explanations and/or similar notions to implemented tools.

8.1. Relatedness Explanations

The problem of finding structures (e.g., paths, subgraphs) connecting entities has been tackled from different perspectives. Faloutsos et al. [34] consider edge weighted graphs without labels and treat the input graph as an electrical network in which each edge represents a resistor with conductance. The sought connectivity subgraph is the one that can deliver as many units of electrical current as possible. Along the same line Ramakrishnan et al. [14] consider multi-relational weighted graphs with weights assigned by looking at the schema (e.g., class and property specificity). To discover the subgraph connecting two entities author resort to the same algorithm as Faloutsos et al. [34].

Our work differs in the fact that we focus on building different kinds of explanations, by also controlling the kind and amount of information to be included, by using two different SPARQL-based algorithms; one that starts from the data (E4D) and the other that starts

from the schema (E4S). Besides, we provide a visual tool (RECAP) with a concrete application of relatedness explanations for query answering.

Heim et al. [35] devised a tool called RelFinder, which given a pair of entities in DBpedia can produce a graph interlinking them. Differently, from our approach that can control the final explanation graph via path ranking and selection strategies, the graph found by RelFinder can quickly become too large to provide useful insights. Moreover, we also provide E4S and described a concrete usage of explanations for query answering. REX [18] finds relationship explanations in the context of the Yahoo! search engine. It leverages (existing) ad-hoc algorithms for explanation enumeration and pruning as well as different measures of interestingness. There are some substantial differences between the work presented in this paper and REX. We make usage of information-theoretic notions to prune the set of paths that will contribute to the explanation. While producing an explanation we take into account *diversity*, which allows to include into an explanation also “rare” edges. We adopt a knowledge-graph agnostic approach and relies on queries to a (remote) endpoint. Moreover, our E4S algorithm allows to start from the KG schema to build personalized explanations. Last but not least, we have described a concrete usage of our approach for querying KGs. This is again achieved by via SPARQL queries to a (remote) endpoint. Voskarides et al. [36] focus on the problem of explaining relatedness making available short text descriptions. In particular, the problem is transformed into that of ranking sentences from documents in a corpus that is related to the knowledge graph. The approach described in this paper is entirely based on structured information and does not require the usage of text corpora. Indeed, our aim is to produce a graph instead of a textual description.

8.2. Discovering Semantic Associations

A number of proposals have tackled the problem of finding semantic associations (paths) interlinking entities. Cheng et al. [20] describe an approach to find a flat list (top-K) of clusters and facet values for refocusing and refining the search in KGs. In particular, the approach can recommend frequent, informative, and small-overlapping patterns by leveraging ontological semantics, query context, and information theory. Our work differs in the fact that we aim at building explanations that are graph based on two different algorithms. Moreover, our path ranking strategies allows to control

the amount of information that will be part of an explanation thus helping the user in precisely understanding why a pair of entities is related. Last but not least, we described a concrete application of explanations for query answering. Tiddi et al. [19] tackle the problem of defining of a cost-function that is able to detect the strongest relationship between two given entities in a supervised way. Authors frame the problem into a genetic programming setting and devise a technique that can learn such cost function by only looking at the topology of the subgraph (set of paths) connecting a pair of entities. In particular, the output of the process is a single path that best (in terms of the cost function) represents the relation between two entities. We start from the perspective of building explanations that are graphs in an unsupervised way. We offer two different algorithms with E4S starting from the KG schema thus giving high flexibility in choosing which kinds of patterns look interesting to build an explanation. Moreover, we present a concrete use case of querying KGs by relatedness along with the RECAP visual tool.

Cheng et al. [37] describe an efficient graph search algorithm for finding associations among a tuple of entities, which prunes the search space by exploiting distances between entities computed based on a distance oracle. After finding a potentially large number of such associations, the approach performs a summarization of notable subgroups. Our approach is SPARQL based and can scale to very large graphs while their work on main memory. Their approach does not use the schema to prune/personalize the search for explanations. Nevertheless, combining our schema-based approach (E4S) with theirs is interesting. Also considering a tuple of entities is an interesting line of future work.

Bianchi et al. [38] focus on the problem of finding the most “interesting” semantic associations by defining a ranking function, which can take into account user preferences. The approach starts from a small sample of semantic associations that are rated by the user and then use it to iteratively learn a personalized ranking function. Our approach is based on two different algorithms (E4D and E4S) that can build relatedness explanations in the form of graphs. Moreover, we have implemented the RECAP tool as we believe that having a way to visually present explanations can help the user in better understanding explanations. Last but not least, as a concrete application of explanations we have implemented a relatedness based querying mechanism. Nevertheless, it would be interesting to inte-

Table 6

Comparison of RECAP with related tools.

System	KG	O	F	Q	L
REX [18]	Yahoo!	Graph	No	No	Yes
RelFinder [35]	DBpedia	Graph	No	No	Yes
Exclass [20]	DBpedia	Paths	Yes (only paths)	No	Yes
RECAP	Any	Graph/Paths	Yes (paths and graphs)	Yes	No

grate our tool with the rating mechanism described by the authors to drive the explanation visualization.

Anyanwu et al. [39], Kochut et al. [40] and Fionda et al. [41] tackle the problem of association discovering from a query language perspective. This work starts from a specific input pair while their from a declarative specification (i.e., a query) the results of which are unknown before evaluation. Fionda et al. [42] focus on the problem of finding graphs at the level of schema starting from the data. Our algorithm E4S does the reverse and focus on building data-level explanations. Finally, other approaches (e.g., [15, 17, 43, 44]) focus on the problem of relation discovery while we focus on building explanations that can be dynamically configured to include the desired amount of information.

8.3. Path Ranking

There is also a solid body of work about path ranking. A series of papers (e.g., [45, 46]) focused on path ranking by leveraging ontological knowledge and information-theoretic techniques with heuristics. Aggarwal et al. [47] focus on ranking all the paths between any two entities by using measures like cohesiveness and specificity. Differently from our approach, these approaches focus on ranking only and not on building explanations. Moreover, we devise a general framework for explanations, two algorithms and a tool. More recently, Cheng et al. [28] perform an empirical comparison of eight techniques to rank semantic associations. The evaluation is based on ground-truth rankings created by human experts. What emerged is that experts generally prefer small semantic associations consisting of entities having similar types. Moreover, semantic associations consisting of uniform relations and those consisting of diverse ones are both preferred. This study suggests that the capability of RECAP to control the size of an explanation (by selecting only a subset of paths) is useful as from the study emerged that users prefer small explanations.

8.4. Comparison of RECAP with related tools

Table 6 summarizes the comparison of our approach with related tools. The comparison is carried out along the following dimensions: KG supported (**KG**), output (**O**), filtering capabilities (**F**), querying capabilities (**Q**), requirement of local data (**L**). RECAP differs from related systems in the following main respects: as for **KG**, RECAP is *KG-independent*; it only requires the availability of a (remote) query endpoint. Moreover, RECAP can combine information from *multiple* KGs (by leveraging `owl:sameAs` links). As for **O** and **F**, RECAP focuses on building *different types* of explanations in the form of graphs or (sets of) paths by leveraging *informativeness* (to estimate the relative importance of edges), *diversity* (to include rare edges) and their combinations. Moreover, RECAP is the only approach that can be used to query KGs (**Q**). As for **L**, *neither* does RECAP assume *local availability* of data *nor* any data *preprocessing*.

9. Concluding Remarks and Future Work

We have introduced a framework to generate different types of relatedness explanations. We have described two different algorithms (E4D and E4S) along with different path ranking strategies. This gives high flexibility in controlling and selecting the preferred amount of information that a (visual) explanation should contain. Explanations are a versatile tool in many contexts, from information retrieval to clustering. This work has also been motivated by the SENSE4US FP7 project, where there was the need to find explanations for groups of topics emerging from textual documents. As a concrete use case, we showed how explanations can be used for querying knowledge graphs in a query-by-example fashion. In the future, we plan to integrate a user feedback mechanism that can help in building more personalized explanations.

References

- [1] C. Shi, Y. Li, J. Zhang, Y. Sun and S.Y. Philip, A survey of heterogeneous information network analysis, *IEEE Transactions on Knowledge and Data Engineering* **29**(1) (2017), 17–37.
- [2] T. Heath and C. Bizer, *Linked Data: Evolving the Web into a Global Data Space*, 1st edn, Morgan & Claypool, 2011. ISBN 9781608454303. <http://linkeddatabook.com/>.
- [3] Y. Sun and J. Han, Mining heterogeneous information networks: principles and methodologies, *Synthesis Lectures on Data Mining and Knowledge Discovery* **3**(2) (2012), 1–159.
- [4] N. Jayaram, A. Khan, C. Li, X. Yan and R. Elmasri, Querying knowledge graphs by example entity tuples, *IEEE Transactions on Knowledge and Data Engineering* **27**(10) (2015), 2797–2811.
- [5] H. Saif, T. Dickinson, L. Kastler, M. Fernandez and H. Alani, A semantic graph-based approach for radicalisation detection on social media, in: *European Semantic Web Conference*, Springer, 2017, pp. 571–587.
- [6] Y. Wang, M.J. Carman and Y.-F. Li, Using Knowledge Graphs to Explain Entity Co-occurrence in Twitter, in: *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*, ACM, 2017, pp. 2351–2354.
- [7] T.D. Noia, V.C. Ostuni, P. Tomeo and E.D. Sciascio, Sprank: Semantic path-based ranking for top-n recommendations using linked open data, *ACM Transactions on Intelligent Systems and Technology (TIST)* **8**(1) (2016), 9.
- [8] A. Saeedi, E. Peukert and E. Rahm, Using Link Features for Entity Clustering in Knowledge Graphs, in: *Extended Semantic Web Conference.*, 2018.
- [9] I. Hulpus, N. Pranganawarat and C. Hayes, Path-based semantic relatedness on linked data and its use to word and entity disambiguation, in: *International Semantic Web Conference*, Springer, 2015, pp. 442–457.
- [10] G. Pirrò, Explaining and suggesting relatedness in knowledge graphs, in: *International Semantic Web Conference*, Springer, 2015, pp. 622–639.
- [11] A. Sheth, B. Aleman-Meza, I.B. Arpinar, C. Bertram, Y. Warke, C. Ramakrishnan, C. Halaschek, K. Anyanwu, D. Avant, F.S. Arpinar et al., Semantic Association Identification and Knowledge Discovery for National Security Applications, *Journal of Database Management* **16**(1) (2005), 33–53.
- [12] V. Fionda, C. Gutierrez and G. Pirrò, Knowledge Maps of Web Graphs, in: *KR*, 2014.
- [13] Y. Tsuruoka, M. Miwa, K. Hamamoto, J. Tsujii and S. Ananiadou, Discovering and Visualizing Indirect Associations between Biomedical Concepts, *Bioinformatics* **27**(13) (2011).
- [14] C. Ramakrishnan, W.H. Milnor, M. Perry and A.P. Sheth, Discovering informative connection subgraphs in multi-relational graphs, *SIGKDD Newsletter* **7**(2) (2005), 56–63.
- [15] P.N. Mendes, P. Kapanipathi, D. Cameron and A.P. Sheth, Dynamic Associative Relationships on the Linked Open Data Web, in: *Web Science Conference*, 2010.
- [16] P. Heim, S. Hellmann, J. Lehmann, S. Lohmann and T. Stegemann, RelFinder: Revealing Relationships in RDF Knowledge Bases, in: *Semantic Multimedia*, Springer, 2009, pp. 182–187.
- [17] G. Kasneci, S. Elbassuoni and G. Weikum, Ming: Mining Informative Entity Relationship Subgraphs, in: *CIKM*, ACM, 2009, pp. 1653–1656.
- [18] L. Fang, A.D. Sarma, C. Yu and P. Bohannon, REX: Explaining Relationships between Entity Pairs, *VLDB* **5**(3) (2011), 241–252.
- [19] I. Tiddi, M. d’Aquin and E. Motta, Learning to assess linked data relationships using genetic programming, in: *International Semantic Web Conference*, Springer, 2016, pp. 581–597.
- [20] G. Cheng, Y. Zhang and Y. Qu, Expass: Exploring Associations between Entities via Top-K Ontological Patterns and Facets, in: *ISWC*, Springer, 2014, pp. 422–437.
- [21] S. Harris and A. Seaborne, SPARQL 1.1 Query Language W3C Recommendation, 2013.
- [22] S. Munoz, J. Pérez and C. Gutierrez, Simple and efficient minimal RDFS, *Web Semantics: Science, Services and Agents on the World Wide Web* **7**(3) (2009), 220–234.
- [23] E. Franconi, C. Gutierrez, A. Mosca, G. Pirrò and R. Rosati, The logic of extensional RDFS, in: *International Semantic Web Conference*, Springer, 2013, pp. 101–116.
- [24] I. Herman, G. Melançon and M.S. Marshall, Graph Visualization and Navigation in Information Visualization: A Survey, *IEEE Trans. on Visualization and Comp. Graph.* **6**(1) (2000), 24–43.
- [25] J. Ugander, B. Karrer, L. Backstrom and C. Marlow, The Anatomy of the Facebook Social Graph, *arXiv preprint arXiv:1111.4503* (2011).
- [26] G. Pirrò, REWOrD: Semantic Relatedness in the Web of Data., in: *26th Conference on Artificial Intelligence (AAAI)*, 2012.
- [27] J. Pérez, M. Arenas and C. Gutierrez, Semantics and complexity of SPARQL, *ACM Transactions on Database Systems (TODS)* **34**(3) (2009), 16.
- [28] G. Cheng, F. Shao and Y. Qu, An Empirical Evaluation of Techniques for Ranking Semantic Associations, *IEEE Transactions on Knowledge and Data Engineering* **29**(11) (2017), 2388–2401.
- [29] R.A. Baeza-Yates and B.A. Ribeiro-Neto, *Modern Information Retrieval*, ACM Press / Addison-Wesley, 1999. ISBN 0-201-39829-X.
- [30] F. Deng, S. Siersdorfer and S. Zerr, Efficient Jaccard-based Diversity Analysis of Large Document Collections, in: *Proceedings of the 21st ACM International Conference on Information and Knowledge Management*, ACM, 2012, pp. 1402–1411.
- [31] L. Page, S. Brin, R. Motwani and T. Winograd, The PageRank Citation Ranking: Bringing Order to the Web. (1999).
- [32] K. Anyanwu, A. Maduko and A. Sheth, SemRank: ranking complex relationship search results on the semantic web, in: *Proceedings of the 14th international conference on World Wide Web*, ACM, 2005, pp. 117–127.
- [33] S. Magliacane, A. Bozzon and E. Della Valle, Efficient execution of top-k SPARQL queries, in: *The Semantic Web–ISWC 2012*, Springer, 2012, pp. 344–360.
- [34] C. Faloutsos, K.S. McCurley and A. Tomkins, Fast Discovery of Connection Subgraphs, in: *SIGKDD*, ACM, 2004, pp. 118–127.
- [35] P. Heim, S. Lohmann and T. Stegemann, Interactive Relationship Discovery via the Semantic Web, in: *ESWC*, Springer, 2010, pp. 303–317.
- [36] N. Voskarides, E. Meij, M. Tsagkias, M. De Rijke and W. Weerkamp, Learning to explain entity relationships in knowledge graphs, in: *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th*

- International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, Vol. 1, 2015, pp. 564–574.
- [37] G. Cheng, D. Liu and Y. Qu, Efficient algorithms for association finding and frequent association pattern mining, in: *International Semantic Web Conference*, Springer, 2016, pp. 119–134.
- [38] F. Bianchi, M. Palmonari, M. Cremaschi and E. Fersini, Actively learning to rank semantic associations for personalized contextual exploration of knowledge graphs, in: *European Semantic Web Conference*, Springer, 2017, pp. 120–135.
- [39] K. Anyanwu and A. Sheth, ρ -Queries: enabling querying for semantic associations on the semantic web, in: *Proceedings of the 12th international conference on World Wide Web*, ACM, 2003, pp. 690–699.
- [40] K.J. Kochut and M. Janik, SPARQLeR: Extended SPARQL for semantic association discovery, in: *European Semantic Web Conference*, Springer, 2007, pp. 145–159.
- [41] V. Fionda and G. Pirrò, Explaining Graph Navigational Queries, in: *European Semantic Web Conference*, Springer, 2017, pp. 19–34.
- [42] V. Fionda and G. Pirrò, Meta Structures in Knowledge Graphs, in: *International Semantic Web Conference*, Springer, 2017, pp. 296–312.
- [43] N. Nakashole, G. Weikum and F. Suchanek, Discovering and Exploring Relations on the Web, *VLDB* 5(12) (2012), 1982–1985.
- [44] I. Traverso-Ribón, G. Palma, A. Flores and M.-E. Vidal, Considering semantics on the discovery of relations in knowledge graphs, in: *European Knowledge Acquisition Workshop*, Springer, 2016, pp. 666–680.
- [45] B. Aleman-Meza, C. Halaschek, I.B. Arpinar and A.P. Sheth, Context-aware semantic association ranking (2003).
- [46] B. Aleman-Meza, C. Halaschek-Weiner, I.B. Arpinar, C. Ramakrishnan and A.P. Sheth, Ranking Complex Relationships on the Semantic Web, *Internet Computing, IEEE* 9(3) (2005), 37–44.
- [47] N. Aggarwal, S. Bhatia and V. Misra, Connecting the Dots: Explaining Relationships Between Unconnected Entities in a Knowledge Graph, in: *International Semantic Web Conference*, Springer, 2016, pp. 35–39.

Appendix A. Questions used in the user evaluation

Fig. 22 shows the full text of the questions used in the user study and comparison of RECAP with related tools (see Section 7.4. of the paper).

Q	Text
Q1	The system helped me get an overview of all the information
Q2	The system helped me easily find information relevant to this task
Q3	The system helped me easily compare and synthesize all kinds of relevant information
Q4	The system provided me with much support for carrying out this task
Q5	The system provided me with sufficient support for carrying out this task
Q6	The system was fast enough in helping me get useful information

Fig. 22. Questions used in the evaluation.

Appendix B. Dataset

The following tables show the input pairs.

Table 7
Pairs from 1 to 37

Pair	Source Entity (w_s)	Target Entity (w_t)
1	Conchita Martínez	Monica Seles
2	Nike Inc.	Tiger Woods
3	Jack Kirby	Iron Man
4	Apple Inc.	Sequoia Capital
5	Albert Einstein	Giuseppe Peano
6	Ingrid Bergman	Isabella Rossellini
7	John F. Kennedy	Jacqueline Kennedy Onassis
8	Barack Obama	George W. Bush
9	Walt Disney	Roy O. Disney
10	Julia Roberts	Emma Roberts
11	Garry Marshall	Hector Elizondo
12	Andrew Jackson	John Quincy Adams
13	Frank Herbert	Brian Herbert
14	Abraham Lincoln	George Washington
15	Tom Cruise	Katie Holmes
16	Christian Bale	Christopher Nolan
17	Hal Roach	Stan Laurel
18	Danielle Steel	Nora Roberts
19	Richard Gere	Carey Lowell
20	Leonardo DiCaprio	Kate Winslet
21	Capcom	Sega
22	Aldi	Lidl
23	Universal Studios	Paramount Pictures
24	IBM	Hewlett-Packard
25	Last Action Hero	Terminator 2: Judgment Day
26	Forbes	Bloomberg L.P.
27	Charmed	Buffy the Vampire Slayer
28	Nanga Parbat	Broad Peak
29	Manhattan Bridge	Brooklyn Bridge
30	Ford Motor Company	Lincoln MKT
31	Donald Knuth	Stanford University
32	Dennis Ritchie	C (programming language)
33	Ludwig van Beethoven	Symphony No. 5 (Beethoven)
34	Uranium-235	Uranium-238
35	Microsoft Office	C++
36	Steven Spielberg	Minority Report
37	France	Paris

Table 8
Pairs from 38 to 74

Pair	Source Entity (w_s)	Target Entity (w_t)
38	Statoil	Oslo
39	Justin Leonard	Australia
40	York	Mark Prescott
41	Japan	Southern California Open
42	Spain	Kenny Dalglish
43	Northampton	Hampshire County Cricket Club
44	Vienna	Jackie McNamara
45	Mother Teresa	Holy See
46	New York Yankees	John Marzano
47	Trans World Airlines	Long Island
48	Moscow Kremlin	Boris Yeltsin
49	Mexico	Grupo Santander
50	Italy	Venice
51	Jack Charlton	Republic of Ireland
52	NAC Breda	FC Groningen
53	Don Wengert	Baltimore
54	Switzerland	Martina Hingis
55	Alexandra Fusai	Japan
56	Italy	Ferrari
57	East Fife F.C.	Arbroath F.C.
58	Alan Turing	Computer scientist
59	David Beckham	Manchester United F.C
60	Microsoft	Microsoft Excel
61	Steven Spielberg	Catch Me If You Can
62	Boeing C-40 Clipper	Boeing
63	Arnold Palmer	Sportsman of the Year
64	Bjarne Stroustrup	C+
65	Nicole Kidman	Tom Cruise
66	Manchester United F.C	Malcolm Glazer
67	Boeing	Boeing 727
68	David Beckham	A.C. Milan
69	Beijing	2008 Summer Olympics
70	Microsoft	Microsoft Office
71	Metropolis (1927 film)	Fritz Lang
72	David Lynch	Dune (film)
73	Thea von Harbou	Fritz Lang
74	Charles Manson	The Beach Boys