



ExCut: Explainable Embedding-Based Clustering over Knowledge Graphs

Mohamed H. Gad-Elrab^{1,2} ^(✉), Daria Stepanova² , Trung-Kien Tran²,
Heike Adel² , and Gerhard Weikum¹

¹ Max-Planck Institute for Informatics,
Saarland Informatics Campus, Saarbrücken, Germany
{gadelrab,weikum}@mpi-inf.mpg.de

² Bosch Center for Artificial Intelligence, Renningen, Germany
{daria.stepanova,trungkien.tran,heike.adel}@de.bosch.com

Abstract. Clustering entities over knowledge graphs (KGs) is an asset for explorative search and knowledge discovery. KG embeddings have been intensively investigated, mostly for KG completion, and have potential also for entity clustering. However, embeddings are latent and do not convey user-interpretable labels for clusters. This work presents ExCut, a novel approach that combines KG embeddings with rule mining methods, to compute informative clusters of entities along with comprehensible explanations. The explanations are in the form of concise combinations of entity relations. ExCut jointly enhances the quality of entity clusters and their explanations, in an iterative manner that interleaves the learning of embeddings and rules. Experiments on real-world KGs demonstrate the effectiveness of ExCut for discovering high-quality clusters and their explanations.

1 Introduction

Motivation. Knowledge graphs (KGs) are collections of triples of the form $\langle \textit{subject predicate object} \rangle$ used for important tasks such as entity search, question answering and text analytics, by providing rich repositories of typed entities and associated properties. For example, Tedros Adhanom is known as a health expert, director of the World Health Organization (WHO), alumni of the University of London, and many more.

KGs can support analysts in exploring sets of interrelated entities and discovering interesting structures. This can be facilitated by **entity clustering**, using unsupervised methods for grouping entities into informative subsets. Consider, for example, an analyst or journalist who works on a large corpus of topically relevant documents, say on the Coronavirus crisis. Assume that key entities in this collection have been spotted and linked to the KG already. Then the KG can guide the user in understanding what kinds of entities are most relevant. With thousands of input entities, from health experts, geo-locations, political decision-makers all the way to diseases, drugs, and vaccinations, the user is

likely overwhelmed and would appreciate a group-wise organization. This task of computing entity clusters [4,6,16] is the problem we address.

Merely clustering the entity set is insufficient, though. The user also needs to understand the nature of each cluster. In other words, clusters must be explainable, in the form of **user-comprehensible labels**. As entities have types in the KG, an obvious solution is to label each cluster with its prevalent entity type. However, some KGs have only coarse-grained types and labels like “people” or “diseases” cannot distinguish health experts from politicians or virus diseases from bacterial infections. Switching to fine-grained types, such as Wikipedia categories, on the other hand, causes the opposite problem: each entity is associated with tens or hundreds of types, and it is unclear which of these would be a good cluster label. The same holds for an approach where common SPO properties (e.g., *educatedIn UK*) are considered as labels. Moreover, once we switch from a single KG to a set of linked open data (LOD) sources as a joint entity repository, the situation becomes even more difficult.

Problem Statement. Given a large set of entities, each with a substantial set of KG properties in the form of categorical values or relations to other entities, our problem is to jointly tackle: (i) **Clustering**: group the entities into k clusters of semantically similar entities; (ii) **Explanation**: generate a user-comprehensible concise labels for the clusters, based on the entity relations to other entities.

State-of-the-Art and Its Limitations. The problem of clustering relational data is traditionally known as conceptual clustering (see, e.g., [25] for overview). Recently, it has been adapted to KGs in the Semantic Web community [6,16]. Existing approaches aim at clustering graph-structured data itself by, e.g., introducing novel notions of distance and similarity directly on the KG [4,5]. Due to the complexity of the data, finding such universally good similarity notions is challenging [5].

Moreover, existing relational learning approaches are not sufficiently scalable to handle large KGs with millions of facts, e.g., YAGO [26] and Wikidata [30]. Clustering entities represented in latent space, e.g., [12,31], helps to overcome this challenge, yet, the resulting clusters are lacking explanations, clustering process is prone to the embedding quality, and hyperparameters are hard to tune [5]. Explaining clusters over KGs, such as [27,28] focus on the discovery of explanations for given perfect clusters. However, obtaining such high-quality clusters in practice is not straightforward.

Approach. To address the above shortcomings, we present **ExCut**, a new method for computing explainable clusters of large sets of entities. The method uses KG embedding as a signal for finding plausible entity clusters, and combines it with logical rule mining, over the available set of properties, to learn interpretable labels. The labels take the form of concise conjunctions of relations that characterize the majority of entities in a cluster. For example, for the above Coronavirus scenario, we aim at mining such labels as $worksFor(X, Y) \wedge type(Y, health.org) \wedge hasDegreeIn(X, life_sciences)$ for a cluster of health experts, $type(X, disease) \wedge causedBy(X, Y) \wedge type(Y, virus)$ for a cluster of virus diseases, and more. A key point in our approach is that

these labels can in turn inform the entity embeddings, as they add salient information. Therefore, we interleave and iterate the computation of embeddings and rule mining adapting the embeddings using as feedback the information inferred by the learned rules.

Contributions. Our main contributions can be summarized as follows:

- We introduce ExCut, a novel approach for computing explainable clusters, which combines embedding-based clustering with symbolic rule learning to produce human-understandable explanations for the resulting clusters. These explanations can also serve as new types for entities.
- We propose several strategies to iteratively fine-tune the embedding model to maximize the explainability and accuracy of the discovered clusters based on the feedback from the learned explanations.
- We evaluate ExCut on real-world KGs. In many cases, it out-performs state-of-the-art methods w.r.t. both clustering and explanations quality.

We made the implementation of ExCut and the experimental resources publicly available at <https://github.com/mhmgad/ExCut>.

2 Preliminaries

Knowledge Graphs. KGs represent interlinked collections of factual information, encoded as a set of $\langle \textit{subject predicate object} \rangle$ triples, *e.g.*, $\langle \textit{tedros_adhanom directorOf WHO} \rangle$. For simplicity, we write triples as in predicate logics format, *e.g.*, $\textit{directorOf}(\textit{tedros_adhanom}, \textit{WHO})$. A signature of a KG \mathcal{G} is $\Sigma_{\mathcal{G}} = \langle \mathbf{P}, \mathbf{E} \rangle$, where \mathbf{P} is a set of binary predicates and \mathbf{E} is a set of entities, *i.e.*, constants, in \mathcal{G} .

KG Embeddings. KG embeddings aim at representing all entities and relations in a continuous vector space, usually as vectors or matrices called *embeddings*. Embeddings can be used to estimate the likelihood of a triple to be true via a scoring function: $f : \mathbf{E} \times \mathbf{P} \times \mathbf{E} \rightarrow \mathbb{R}$. Concrete scoring functions are defined based on various vector space assumptions: (i) The translation-based assumption, *e.g.*, TransE [1] embeds entities and relations as vectors and assumes $\mathbf{v}_s + \mathbf{v}_r \approx \mathbf{v}_o$ for true triples, where $\mathbf{v}_s, \mathbf{v}_r, \mathbf{v}_o$ are vector embeddings for subject s , relation r and object o , resp. (ii) The linear map assumption, *e.g.*, ComplEx [29] embeds entities as vectors and relations as matrices. It assumes that for true triples, the linear mapping \mathbf{M}_r of the subject embedding \mathbf{v}_s is close to the object embedding \mathbf{v}_o : $\mathbf{v}_s \mathbf{M}_r \approx \mathbf{v}_o$. The likelihood that these assumptions of the embedding methods hold should be higher for triples in the KG than for those outside. The learning process is done through minimizing the error induced from the assumptions given by their respective loss functions.

Rules. Let \mathbf{X} be a set of variables. A *rule* r is an expression of the form $\textit{head} \leftarrow \textit{body}$, where *head*, or $\textit{head}(r)$, is an atom over $\mathbf{P} \cup \mathbf{E} \cup \mathbf{X}$ and *body*, or $\textit{body}(r)$, is a conjunction of positive atoms over $\mathbf{P} \cup \mathbf{E} \cup \mathbf{X}$. In this work, we are concerned with *Horn rules*, a subset of first-order logic rules with only

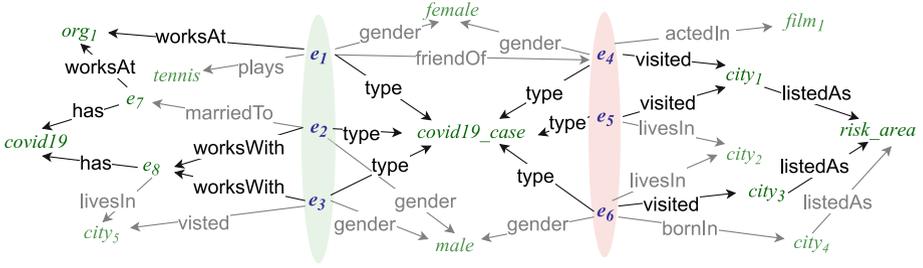


Fig. 1. An example KG with potential COVID-19 cases split into two entity clusters (in green and red). Black edges are relevant for the potential explanations of these clusters.

positive atoms, in which every head variable appears at least once in the body atoms.

Example 1. An example of a rule over the KG in Fig. 1 is $r : has(X, covid19) \leftarrow worksWith(X, Y), has(Y, covid19)$ stating that coworkers of individuals with covid19 infection, potentially, also have covid19.

Execution of rules over KGs is defined in the standard way. More precisely, let \mathcal{G} be a KG, r a rule over $\Sigma_{\mathcal{G}}$, and a an atom, we write $r \models_{\mathcal{G}} a$ if there is a variable assignment that maps all atoms of $body(r)$ into \mathcal{G} and $head(r)$ to the atom a . *Rule-based inference* is the process of applying a given rule r on \mathcal{G} , which results in the extension \mathcal{G}_r of \mathcal{G} defined as: $\mathcal{G}_r = \mathcal{G} \cup \{a \mid r \models_{\mathcal{G}} a\}$.

Example 2. Application of the rule r from Example 1 on the KG \mathcal{G} from Fig. 1 results in $r \models_{\mathcal{G}} has(e_2, covid19)$ and $r \models_{\mathcal{G}} has(e_3, covid19)$. Hence, $\mathcal{G}_r = \mathcal{G} \cup \{has(e_2, covid19), has(e_3, covid19)\}$.

3 Model for Computing Explainable Clusters

Given a KG, a subset of its entities and an integer k , our goal is to find a “good” split of entities into k clusters and compute explanations for the constructed groups that would serve as informative cluster labels. E.g., consider the KG in Fig. 1, the set of target entities $\{e_1, \dots, e_6\}$ and the integer $k = 2$. One of the possible solutions is to put e_{1-3} into the first cluster C_1 and the other three entities into the second one C_2 . Explanations for this split would be that C_1 includes those who got infected via interacting with their coworkers, while the others were infected after visiting a risk area. Obviously, in general there are many other splits and identifying the criteria for the best ones is challenging.

Formally, we define the problem of *computing explainable entity clusters* as follows:

Definition 1 (Computing Explainable Entity Clusters Problem).

Given: (i) a knowledge graph \mathcal{G} over $\Sigma_{\mathcal{G}} = \langle \mathbf{P}, \mathbf{E} \rangle$; (ii) a set $T \subseteq \mathbf{E}$ of target

entities; (iii) a number of desired clusters $k > 1$; (iv) an explanation language L ; and (v) an explanation evaluation function $d : 2^L \times 2^T \times \mathcal{G} \rightarrow [0..1]$

Find: a split $\mathcal{C} = \{C_1, \dots, C_k\}$ of entities in T into k clusters and a set of explanations $\mathcal{R} = \{r_1, \dots, r_k\}$ for them, where $r_i \in L$, s.t. $d(\mathcal{R}, \mathcal{C}, \mathcal{G})$ is maximal.

3.1 Explanation Language

Explanations (i.e., informative labels) for clusters can be characterized as conjunctions of common entity properties in a given cluster; for that Horn rules are sufficient. Thus, our explanation language relies on (cluster) explanation rules defined as follows:

Definition 2 (Cluster Explanation Rules). Let \mathcal{G} be a KG with the signature $\Sigma_{\mathcal{G}} = \langle \mathbf{P}, \mathbf{E} \rangle$, let $C \subseteq \mathbf{E}$ be a subset of entities in \mathcal{G} , i.e., a cluster, and \mathbf{X} a set of variables. A (cluster) explanation rule r for C over \mathcal{G} is of the form

$$r : \text{belongsTo}(X, e_c) \leftarrow p_1(\mathbf{X}_1), \dots, p_m(\mathbf{X}_m), \quad (1)$$

where $e_c \notin \mathbf{E}$ is a fresh unique entity representing the cluster C , $\text{belongsTo} \notin \mathbf{P}$ is a fresh predicate, and $\text{body}(r)$ is a finite set of atoms over \mathbf{P} and $\mathbf{X} \cup \mathbf{E}$.

Example 3. A possible explanation rule for $C_1 = \{e_1, e_2, e_3\}$ in \mathcal{G} from Fig. 1 is

$$r : \text{belongsTo}(X, e_{c_1}) \leftarrow \text{worksWith}(X, Y), \text{has}(Y, \text{covid19})$$

which describes C_1 as a set of people working with infected colleagues.

Out of all possible cluster explanation rules we naturally prefer **succinct** ones. Therefore, we put further restrictions on the explanation language L by limiting the number of rule body atoms (an adjustable parameter in our method).

3.2 Evaluation Function

The function d from Definition 1 compares solutions to the problem of explainable entity clustering w.r.t. their quality, and ideally d should satisfy the following two criteria: (i) **Coverage:** Given two explanation rules for a cluster, the one covering more entities should be preferred and (ii) **Exclusiveness:** Explanation rules for different clusters should be (approximately) mutually exclusive.

The coverage measure from data mining is a natural choice for satisfying (i).

Definition 3 (Explanation Rule Coverage). Let \mathcal{G} be a KG, C a cluster of entities, and r a cluster explanation rule. The coverage of r on C w.r.t. \mathcal{G} is

$$\text{cover}(r, C, \mathcal{G}) = \frac{|\{c \in C \mid r \models_{\mathcal{G}} \text{belongsTo}(c, e_c)\}|}{|C|} \quad (2)$$

Example 4. Consider clusters $C_1 = \{e_1, e_2, e_3\}$, $C_2 = \{e_4, e_5, e_6\}$ shown in Fig. 1. The set of potential cluster explanation rules along with their coverage scores for C_1 and C_2 respectively, is given as follows:

$r_1 : \text{belongsTo}(X, \mathbf{ec}_i) \leftarrow \text{type}(X, \text{covid19_case})$	1	1
$r_2 : \text{belongsTo}(X, \mathbf{ec}_i) \leftarrow \text{gender}(X, \text{male})$	0.67	0.33
$r_3 : \text{belongsTo}(X, \mathbf{ec}_i) \leftarrow \text{worksWith}(X, Y), \text{has}(Y, \text{covid19})$	0.67	0
$r_4 : \text{belongsTo}(X, \mathbf{ec}_i) \leftarrow \text{visited}(X, Y), \text{listedAs}(Y, \text{risk_area})$	0	1

While addressing (i), the coverage measure does not account for the criteria (ii). Indeed, high coverage of a rule for a given cluster does not imply a low value of this measure for other clusters. For instance, in Example 4 r_1 is too general, as it perfectly covers entities from both clusters. This motivates us to favour (approximately) mutually exclusive explanation rules, *i.e.*, explanation rules with high coverage for a given cluster but low coverage for others (similar to [13]). To capture this intuition, we define the *exclusive explanation coverage* of a rule for a cluster given other clusters as follows.

Definition 4 (Exclusive Explanation Rule Coverage). Let \mathcal{G} be a KG, let \mathcal{C} be a set of all clusters of interest, $C \in \mathcal{C}$ a cluster, and r an explanation rule. The exclusive explanation rule coverage of r for C w.r.t. \mathcal{C} and \mathcal{G} is defined as

$$\text{exc}(r, C, \mathcal{C}, \mathcal{G}) = \begin{cases} 0, & \text{if } \min_{C' \in \mathcal{C} \setminus C} \{ \text{cover}(r, C, \mathcal{G}) - \text{cover}(r, C', \mathcal{G}) \} \leq 0 \\ \text{cover}(r, C, \mathcal{G}) - \frac{\sum_{C' \in \mathcal{C} \setminus C} \text{cover}(r, C', \mathcal{G})}{|\mathcal{C} \setminus C|}, & \text{otherwise.} \end{cases} \quad (3)$$

Example 5. Consider $\mathcal{C} = \{C_1, C_2\}$, $\mathcal{R} = \{r_1, r_2, r_3, r_4\}$ from Example 4 and the KG \mathcal{G} from Fig. 1. We have $\text{exc}(r_1, C_1, \mathcal{C}, \mathcal{G}) = \text{exc}(r_1, C_2, \mathcal{C}, \mathcal{G}) = 0$, which disqualifies r_1 as an explanation for either of the clusters. For r_2 , we have $\text{exc}(r_2, C_1, \mathcal{C}, \mathcal{G}) = 0.34$ making it less suitable for the cluster C_1 than r_3 with $\text{exc}(r_3, C_1, \mathcal{C}, \mathcal{G}) = 0.67$. Finally, r_4 perfectly explains C_2 , since $\text{exc}(r_4, C_2, \mathcal{C}, \mathcal{G}) = 1$.

Similarly, we can measure the *quality* of a collection of clusters with their explanations by averaging their per-cluster exclusive explanation rule coverage.

Definition 5 (Quality of Explainable Clusters). Let \mathcal{G} be a KG, $\mathcal{C} = \{C_1, \dots, C_k\}$ a set of entity clusters, and $\mathcal{R} = \{r_1, \dots, r_k\}$ a set of cluster explanation rules, where each r_i is an explanation for C_i , $1 \leq i \leq k$. The explainable clustering quality q of \mathcal{R} for \mathcal{C} w.r.t. \mathcal{G} is defined as follows:

$$q(\mathcal{R}, \mathcal{C}, \mathcal{G}) = \frac{1}{|\mathcal{C}|} \sum_{i=1}^{|\mathcal{C}|} \text{exc}(r_i, C_i, \mathcal{C}, \mathcal{G}) \quad (4)$$

Realizing the function d in Definition 1 by the above measure allows us to conveniently compare the solutions of the explainable clusters discovery problem.

Example 6. Consider \mathcal{G} from Fig. 1, the set of target entities $T = \{\mathbf{e}_1, \dots, \mathbf{e}_6\}$, $k = 2$, language L of cluster explanation rules with at most 2 body atoms, and the evaluation function d given as q from Definition 5. The best solution to the respective problem of computing explainable entity clusters is $\mathcal{C} = \{C_1, C_2\}$, $\mathcal{R} = \{r_3, r_4\}$, where C_1, C_2, r_3, r_4 are from Example 4. We have that $q(\mathcal{R}, \mathcal{C}, \mathcal{G}) = 0.83$.

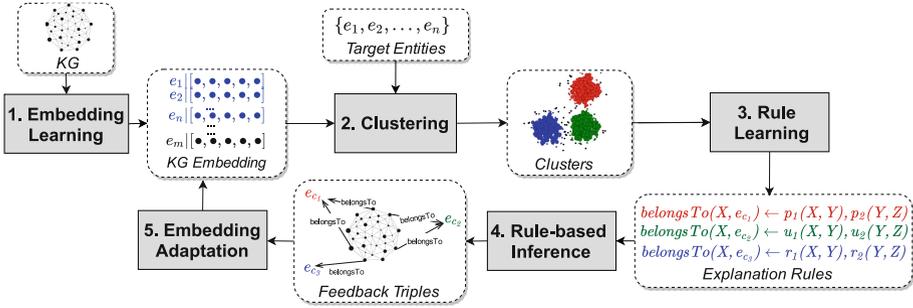


Fig. 2. ExCut pipeline overview.

4 Method

We now present our method ExCut, which iteratively utilizes KG *Embedding-based Clustering* and *Rule Learning* to compute explainable clusters. More specifically, as shown in Fig. 2, ExCut starts with (1) *Embedding Learning* for a given KG. Then, it performs (2) *Clustering* of the entities in the target set over the learned embeddings. Afterwards, (3) *Rule Learning* is utilized to induce explanation rules for the constructed clusters, which are ranked based on the exclusive coverage measure. Using the learned explanation rules, we perform (4) *Rule-based Inference* to deduce new entity-cluster assignment triples reflecting the learned structural similarities among the target entities. Then, ExCut uses the rules and the inferred assignment triples in constructing feedback to guide the clustering in the subsequent iterations. We achieve that by fine-tuning the embeddings of the target entities in Step (5) *Embedding Adaptation*.

In what follows we present the detailed description of ExCut’s components.

4.1 Embedding Learning and Clustering

Embedding Learning. ExCut starts with learning vector representations of entities and relations. We adopt KG embeddings in this first step, as they are well-known for their ability to capture semantic similarities among entities, and thus could be suited for defining a robust similarity function for *clustering relational data* [5]. Embeddings are also effective for dealing with data incompleteness, *e.g.*, predicting the potentially missing fact *worksWith*(e_1, e_7) in Fig. 1. Moreover, embeddings facilitate the inclusion of unstructured external sources during training, *e.g.*, textual entity descriptions [33].

Conceptually, any embedding method can be used in our approach. We experimented with TransE [1] and ComplEx [29] as prominent representatives of translation-based and linear map embeddings. To account for the context surrounding the target entities, we train embeddings using the whole KG.

Clustering. The *Clustering* step takes as input the trained embedding vectors of the target entities and the number k of clusters to be constructed. We perform

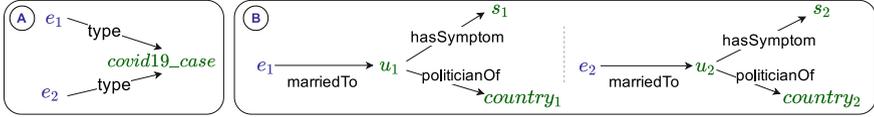


Fig. 3. KG fragments.

clustering relying on the embeddings as features to compute pairwise distances among the target entities using standard distance functions, *e.g.*, *cosine distance*. Various classical clustering approaches or more complex embedding-driven clustering techniques [31] could be exploited here too. In this paper, we rely on the traditional *Kmeans* method [17] as a proof of concept.

For KGs with types, the majority of embedding models [1, 29] would map entities of a certain type to similar vectors [31]. For example, e_1 and e_2 in Fig. 3A are likely to be close to each other in the embedding space, and thus have a high chance of being clustered together. An ideal embedding model for explainable clustering should follow the same intuition even if types in the KG are missing. In other words, it should be capable of assigning similar vectors to entities that belong to structurally similar subgraphs of certain pre-specified complexity. For instance, in Fig. 3B, both e_1 and e_2 belong to subgraphs reflecting that these entities are married to politicians with some *covid19* symptom, and hence should be mapped to similar vectors.

Despite certain attempts to consider specific graph patterns (*e.g.*, [15]), to the best of our knowledge none of the existing embedding models is general enough to capture patterns of arbitrary complexity. We propose to tackle this limitation (see Sect. 4.3) by passing to the embedding model feedback created using cluster explanation rules learned in the Step 3 of ExCut.

4.2 Explanation Mining

KG-Based Explanations. KG embeddings and the respective clusters constructed in Steps 1 and 2 of our method are not interpretable. However, since KG embeddings are expected to preserve semantic similarities among entities, the clusters in the embedding space should intuitively have some meaning. Motivated by this, in ExCut, we aim at decoding these similarities by learning rules over the KG extended by the facts that reflect the cluster assignments computed in the *Clustering* step.

Rule Learning Procedure. After augmenting \mathcal{G} with *belongsTo*(e, e_{c_i}) facts for all entities e clustered in C_i , we learn Horn rules of the form (1) from Definition 2. There are powerful rule-learning tools such as AMIE+ [8], AnyBurl [18], RLVLR [20, 21] and RuDiK [22]. Nevertheless, we decided to develop our own rule learner so that we could have full control over our specific scoring functions and their integration into the learner’s search strategy. Following [8], we model rules as sequences of atoms, where the first atom is the head of the rule (*i.e.*,

$belongsTo(X, \mathbf{e}_{c_i})$ with C_i being the cluster to be explained), and other atoms form the rule's body.

For each cluster C_i , we maintain an independent queue of intermediate rules, initialized with a single rule atom $belongsTo(X, \mathbf{e}_{c_i})$, and then exploit an iterative breadth-first search strategy. At every iteration, we expand the existing rules in the queue using the following *refinement operators*: (i) *add a positive dangling atom*: add a binary positive atom with one fresh variable and another variable appearing in the rule, *i.e.*, *shared variable*, *e.g.*, adding $worksAt(X, Y)$, where Y is a fresh variable not appearing in the current rule; (ii) *add a positive instantiated atom*: add a positive atom with one argument being a constant and the other one a shared variable, *e.g.*, adding $locatedIn(X, usa)$, where usa is a constant, and X appears elsewhere in the rule constructed so far.

These operators produce a set of new rule candidates, which are then filtered relying on the given explanation language L . Suitable rules with a minimum coverage of 0.5, *i.e.*, rules covering the majority of the respective cluster, are added to the output set. We refine the rules until the maximum length specified in the language bias is reached. Finally, we rank the constructed rules based on the *exclusive explanation coverage* (Definition 4), and select the top m rules for each cluster.

Example 7. Assume that for \mathcal{G} in Fig. 1, and $T = \{\mathbf{e}_1, \dots, \mathbf{e}_6\}$, the embedding-based clustering resulted in the following clusters $C_1 = \{\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_4\}$ and $C_2 = \{\mathbf{e}_5, \mathbf{e}_6, \mathbf{e}_3\}$, where \mathbf{e}_4 and \mathbf{e}_3 are incorrectly placed in wrong clusters. The top cluster explanation rules for C_2 ranked based on *exc* measure from Definition 4 are:

$r_1 : belongsTo(X, \mathbf{e}_{c_2}) \leftarrow visited(X, Y)$	0.67
$r_2 : belongsTo(X, \mathbf{e}_{c_2}) \leftarrow gender(X, \mathbf{male})$	0.33
$r_3 : belongsTo(X, \mathbf{e}_{c_2}) \leftarrow visited(X, Y), listedAs(Y, \mathbf{risk_area}).$	0.33

Inferring Entity-Clusters Assignments. In the *Rule-based Inference* (Step 4 in Fig. 2), we apply the top- m rules obtained in the *Rule Learning* step on the KG to predict the assignments between the target entities and the discovered clusters over $belongsTo$ relation using standard deductive reasoning techniques. The computed assignment triples are ranked and filtered based on the *exc* score of the respective rules that inferred them.

Example 8. Application of the rules from Example 7 on \mathcal{G} w.r.t. the target entities \mathbf{e}_{1-6} results in the cluster assignment triples: $\{belongsTo(\mathbf{e}_3, \mathbf{e}_{c_2}), belongsTo(\mathbf{e}_4, \mathbf{e}_{c_2}), belongsTo(\mathbf{e}_2, \mathbf{e}_{c_2})\}$. Note that based on r_1 , \mathbf{e}_4 is assigned to C_2 instead of C_1 .

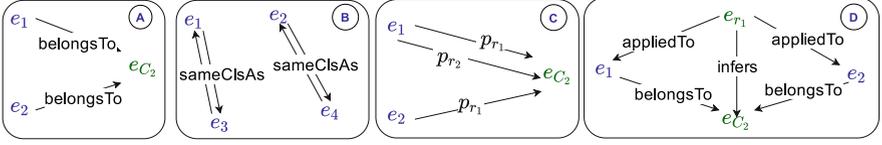


Fig. 4. Inferred clusters assignment triples modeling options.

4.3 Embedding Adaptation

Learned explanation rules capture explicit structural similarities among the target entities. We propose to utilize them to create feedback to guide the embedding-based clustering towards better explainable clusters. This feedback is passed to the embedding model in the form of additional training triples reflecting the assignments inferred by the learned rules. Our intuition is that such added triples should potentially help other similarities of analogous nature to be discovered by the embeddings, compensating for the embedding-based clustering limitation discussed in Sect. 4.1.

Specifically, the embedding adaptation (Step 5 in Fig. 2) is summarized as follows: (a) From the *Rule Learning* and *Rule-based Inference* steps, described above, we obtain a set of *cluster assignment triples* of the form $belongsTo(e, e_C)$ together with rules inferring them, where e is an entity in the input KG \mathcal{G} and e_C is a new entity uniquely representing the cluster C . (b) We then model the cluster assignments from (a) and rules that produce them using one of our four strategies described below and store the results in \mathcal{G}^{inf} . (c) A subset $\mathcal{G}^{context}$ of \mathcal{G} consisting of triples that surround the target entities is then constructed. (d) Finally, we fine-tune the embedding model by training it further on the data compiled from \mathcal{G}^{inf} and $\mathcal{G}^{context}$.

Modeling Rule-Based Feedback. Determining the adequate structure and amount of training triples required for fine-tuning the embedding model is challenging. On the one hand, the training data should be rich enough to reflect the learned structure, but on the other hand, it should not corrupt the current embedding. We now present our proposed four strategies for representing the inferred cluster-assignments along with the corresponding rules as a set of triples \mathcal{G}^{inf} suitable for adapting the embedding. The strategies are listed in the ascending order of their complexity.

- **Direct:** As a straightforward strategy, we directly use the inferred entity-cluster assignment triples in \mathcal{G}^{inf} as shown in Fig. 4A, e.g., $belongsTo(e_1, e_{C_2})$.
- **Same-Cluster-as:** In the second strategy, we model the inferred assignments as edges only. As shown in Fig. 4B, we compile \mathcal{G}^{inf} using triples of $sameClsAs$ relations between every pair of entities belonging to the same cluster as the learned rules suggest, e.g., $sameClsAs(e_1, e_2)$. Modeling the cluster assignments using fresh relations allows us to stress the updates related to the target entities, as no extra entities are added to the KG in this strategy.

- **Rules as Edges:** Third, we propose to model the inferred assignments together with the rules which led to their prediction. More precisely, for every rule r which deduced $belongsTo(\mathbf{e}, \mathbf{e}_{c_i})$, we introduce a fresh predicate p_r and add a triple $p_r(\mathbf{e}, \mathbf{e}_{c_i})$ to the training set \mathcal{G}^{inf} , as illustrated in Fig. 4C. This allows us to encode all conflicting entity-cluster assignments (*i.e.*, assignments, in which an entity belongs to two different clusters) and supply the embedding model with richer evidence about the rules that predicted these assignments.
- **Rules as Entities:** Rules used in the deduction process can also be modeled as entities. In the fourth strategy, we exploit this possibility by introducing additional predicates *infers* and *appliedTo*, and for every rule r a fresh entity \mathbf{e}_r . Here, each $belongsTo(\mathbf{e}, \mathbf{e}_{c_i})$ fact deduced by the rule r is modeled in \mathcal{G}^{inf} with two triples $infers(\mathbf{e}_r, \mathbf{e}_{c_i})$ and $appliedTo(\mathbf{e}_r, \mathbf{e})$ as shown in Fig. 4D.

Embedding Fine-Tuning. At every iteration i of ExCut, we start with the embedding vectors obtained in the previous iteration $i - 1$ and train the embedding further with a set of adaptation triples \mathcal{G}^{adapt} . The set \mathcal{G}^{adapt} is composed of the union of all \mathcal{G}_j^{inf} for $j = 1 \dots i$ and a set of context triples $\mathcal{G}^{context}$. For $\mathcal{G}^{context}$, we only consider those directly involving the target entities as a subject or object. E.g., among the facts in the surrounding context of \mathbf{e}_1 , we have $worksAt(\mathbf{e}_1, \mathbf{org}_1)$ and $plays(\mathbf{e}_1, \mathbf{tennis})$.

Our empirical studies (see the technical report¹) showed that including assignment triples from previous iterations $j < i$ leads to better results; thus, we include them in \mathcal{G}^{adapt} , but distinguish entity and relation names from different iterations. Additionally, considering the context subgraph helps in regulating the change caused by the cluster assignment triples by preserving some of the characteristics of the original embeddings.

5 Experiments

We evaluate the effectiveness of ExCut for computing explainable clusters. More specifically, we report the experimental results covering the following aspects: (i) the quality of the clusters produced by ExCut compared to existing clustering approaches; (ii) the quality of the computed cluster explanations; (iii) the usefulness and understandability of the explanations for humans based on a user study; (iv) the benefits of interleaving embedding and rule learning for enhancing the quality of the clusters and their explanations; and (v) the impact of using different embedding paradigms and our strategies for modeling the feedback from the rules.

5.1 Experiment Setup

ExCut Configurations. We implemented ExCut in Python and configured its components as follows: (i) *Embedding-based Clustering:* We extended the

¹ Code, data and the technical report are available at <https://github.com/mhmgad/ExCut>.

Table 1. Datasets statistics.

	UWCSE	WebKB	Terror.	IMDB	Mutag.	Hepatitis	LUBM	YAGO
Target Entities	209	106	1293	268	230	500	2850	3900
Target Clusters	2	4	6	2	2	2	2	3
KG Entities	991	5906	1392	578	6196	6511	242558	4295825
Relations	12	7	4	4	14	19	22	38
Facts	2216	72464	17117	1231	30805	77585	2169451	12430700

implementation of *TransE* and *ComplEx* provided by Ampligraph [3] to allow embedding fine-tuning. We set the size of the embeddings to 100, and trained a base model with the whole KG for 100 epochs, using stochastic gradient descent with a learning rate of 0.0005. For fine-tuning, we trained the model for 25 epochs with a learning rate of 0.005. *Kmeans* is used for clustering. (ii) *Rule Learning*: We implemented the algorithm described in Sect. 4.2. For experiments, we fix the language bias of the explanations to paths of length two, e.g., $belongsTo(x, e_{c_i}) \leftarrow p(x, y), q(y, z)$, where z is either a free variable or bind to a constant. (iii) *Modeling Rule-based Feedback*: We experiment with the four strategies from Sect. 4.3: direct (*belongToCl*), same cluster as edges (*sameCLAs*), rules as edges (*entExplCl*), and rules as entities (*followExpl*).

Datasets. We performed experiments on six datasets (Tab.1) with a pre-specified set of target entities, which are widely used for relational clustering [4]. Additionally, we considered the following large-scale KGs: (i) *LUBM-Courses*: a subset of entities from LUBM syntactic KG [9] describing the university domain, where target entities are distributed over *graduate* and *undergraduate courses*; and (ii) *YAGO-Artwork* KG with a set of target entities randomly selected from *YAGO* [26]. The entities are uniformly distributed over three types, *book*, *song*, and *movie*. To avoid trivial explanations, type triples for target entities were removed from the KG. Table 1 reports the dataset statistics.

Baselines. We compare ExCut to the following clustering methods: (i) *ReCeNT* [4], a state-of-the-art relational clustering approach, that clusters entities based on a similarity score computed from entity neighborhood trees; (ii) *Deep Embedding Clustering (DEC)* [32], an embedding-based clustering method that performs dimensionality reduction jointly with clustering and (iii) Standard *Kmeans* applied directly over embeddings: *TransE (Kmeans-T)* and *ComplEx (Kmeans-C)*. This baseline is equivalent to a single iteration of our system ExCut. Extended experiments with clustering algorithms that automatically detect the number of clusters can be found in the technical report.

Clustering Quality Metrics. We measure the clustering quality w.r.t. the ground truth with three standard metrics: *Accuracy (ACC)*, *Adjusted Rand Index (ARI)*, and *Normalized Mutual Information (NMI)* (the higher, the better).

Explanation Quality Metrics. The quality of the generated explanations is measured using the coverage metrics defined in Sect. 3.2, namely, *per cluster coverage (Cov)* and *exclusive coverage (Exc)*. In addition, we adapted the “novelty” metric *Weighted Relative Accuracy (WRA)* [14], which represents a trade-off

Table 2. Clustering results of ExCut compared to the baselines.

Methods	UWCSE			IMDB			Hepatitis			Mutagenesis			WebKB			Terrorist			
	ACC	ARI	NMI	ACC	ARI	NMI	ACC	ARI	NMI	ACC	ARI	NMI	ACC	ARI	NMI	ACC	ARI	NMI	
Baselines	ReCeNT	0.90	0.60	0.54	0.61	0.02	0.01	0.51	-0.01	0.01	0.77	0.30	0.24	0.52	0.00	-0.25	0.37	0.10	0.13
	DEC	0.67	0.17	0.12	0.54	0.00	0.01	0.55	0.01	0.01	0.51	0.00	0.00	0.31	0.03	0.05	0.37	0.16	0.26
	Kmeans-T	0.91	0.66	0.51	0.58	0.03	0.08	0.51	0.00	0.00	0.52	0.00	0.00	0.33	0.01	0.06	0.53	0.33	0.44
	Kmeans-C	0.54	0.00	0.01	0.53	0.00	0.00	0.52	0.00	0.00	0.73	0.21	0.18	0.49	0.21	0.34	0.51	0.23	0.28
ExCut-T	belongToCl	0.99	0.96	0.92	1.00	1.00	1.00	0.83	0.43	0.35	0.68	0.12	0.13	0.43	0.13	0.17	0.52	0.27	0.31
	sameClAs	1.00	1.00	1.00	1.00	1.00	1.00	0.56	0.01	0.01	0.65	0.08	0.08	0.36	0.06	0.08	0.35	0.03	0.06
	entExplCl	1.00	1.00	1.00	1.00	1.00	1.00	0.82	0.41	0.33	0.64	0.07	0.08	0.43	0.13	0.20	0.45	0.17	0.23
	followExpl	1.00	1.00	1.00	1.00	1.00	1.00	0.82	0.41	0.33	0.64	0.08	0.08	0.44	0.15	0.22	0.45	0.16	0.22
Excut-C	belongToCl	0.96	0.85	0.77	1.00	1.00	1.00	0.63	0.07	0.05	0.73	0.21	0.18	0.51	0.23	0.37	0.54	0.26	0.29
	sameClAs	0.98	0.91	0.86	1.00	1.00	1.00	0.58	0.02	0.02	0.73	0.21	0.18	0.38	0.08	0.17	0.34	0.03	0.08
	entExplCl	0.97	0.88	0.81	0.65	0.08	0.19	0.69	0.15	0.11	0.73	0.21	0.19	0.52	0.24	0.36	0.53	0.25	0.29
	followExpl	0.99	0.97	0.94	1.00	1.00	1.00	0.66	0.10	0.08	0.73	0.20	0.18	0.51	0.22	0.34	0.52	0.24	0.29

between the coverage and the accuracy of the discovered explanations. We compute the average of the respective quality of the top explanations for all clusters. To assess the quality of the solution to the explainable clustering problem from Definition 1 found by ExCut, we compare the computed quality value to the quality of the explanations computed over the ground truth.

All experiments were performed on a Linux machine with 80 cores and 500 GB RAM. The average results over 5 runs are reported.

User Study. To assess the human-understandability and usefulness of the explanation rules, we analyze whether ExCut explanations are the best fitting labels for the computed clusters based on the user opinion. The study was conducted on Amazon MTurk.

More specifically, based on the YAGO KG, we provided the user study participants with: (i) Three clusters of entities, each represented with three entities pseudo-randomly selected from these clusters along with a brief summary for each entity, and a link to its Wikipedia page; (ii) A set of 10 potential explanations composed of the top explanations generated by ExCut and other explanations with high *Cov* but low *Exc*. Explanations were displayed in natural language for the ease of readability. We asked the participants to match each explanation to all relevant clusters.

A *useful* explanation is the one that is *exclusively matched* to the correct cluster by the participants. To detect useful explanations, for every *explanation-cluster* pair, we compute the ratio of responses where the pair is *exclusively matched*. Let $match(r_i, c_m) = 1$ if the user *matched* explanation r_i to the cluster c_m (otherwise 0). Then, r_i is *exclusively matched* to c_m if additionally, $match(r_i, c_j) = 0$ for all $j \neq m$.

5.2 Experiment Results

In seven out of eight datasets, our approach outperforms the baselines with regard to the overall clustering and explanation quality metrics. Additionally, the quality of the computed explanations increases after few iterations.

Table 3. Quality of Clusters Explanations by ExCut compared to the baselines.

Methods	UWCSE			IMDB			Hepatitis			Mutagenesis			WebKB			Terrorist			
	Cov	Exc	WRA	Cov	Exc	WRA	Cov	Exc	WRA	Cov	Exc	WRA	Cov	Exc	WRA	Cov	Exc	WRA	
Baselines	ReCeNT	0.91	0.88	0.14	1.00	0.04	0.01	1.00	0.00	0.00	1.00	0.00	0.00	1.00	1.00	0.00	0.93	0.42	0.06
	DEC	0.73	0.31	0.07	1.00	0.03	0.01	1.00	0.01	0.00	1.00	0.00	0.00	1.00	0.06	0.01	0.60	0.13	0.02
	Kmeans-T	0.83	0.76	0.16	0.74	0.11	0.01	0.81	0.09	0.02	0.75	0.11	0.03	0.75	0.11	0.03	0.49	0.17	0.02
	Kmeans-C	0.59	0.06	0.01	0.73	0.04	0.01	0.61	0.09	0.02	0.87	0.30	0.08	0.98	0.04	0.01	0.64	0.28	0.02
ExCut-T	belongToCl	0.89	0.89	0.19	1.00	1.00	0.11	0.76	0.64	0.13	0.94	0.39	0.09	0.98	0.12	0.01	0.68	0.26	0.03
	sameClAs	0.90	0.90	0.19	1.00	1.00	0.11	0.94	0.45	0.09	0.96	0.50	0.12	0.99	0.04	0.01	0.87	0.49	0.06
	entExplCl	0.90	0.90	0.19	1.00	1.00	0.11	0.75	0.64	0.13	0.99	0.48	0.12	0.99	0.10	0.01	0.94	0.80	0.11
	followExpl	0.90	0.90	0.19	1.00	1.00	0.11	0.75	0.63	0.13	0.98	0.46	0.11	0.99	0.09	0.01	0.95	0.79	0.11
ExCut-C	belongToCl	0.88	0.86	0.18	1.00	1.00	0.11	0.73	0.50	0.12	0.87	0.31	0.08	0.98	0.08	0.01	0.68	0.32	0.02
	sameClAs	0.91	0.89	0.19	1.00	1.00	0.11	0.80	0.45	0.11	0.87	0.30	0.08	0.98	0.10	0.01	0.85	0.61	0.07
	entExplCl	0.88	0.88	0.19	0.73	0.18	0.01	0.85	0.73	0.18	0.87	0.31	0.08	0.97	0.08	0.01	0.68	0.33	0.03
	followExpl	0.90	0.89	0.19	1.00	1.00	0.11	0.81	0.66	0.12	0.87	0.31	0.08	0.97	0.07	0.01	0.67	0.30	0.03
Ground truth	0.92	0.90	0.19	1.00	1.00	0.11	0.92	0.57	0.14	1.00	0.16	0.04	1.00	0.04	0.01	0.64	0.33	0.03	

Clustering Quality. Table 2 presents the quality of the clusters computed by the baselines, in the first 4 rows, followed by ExCut with the four feedback strategies, where *ExCut-T* and *ExCut-C* stand for ExCut with TransE and ComplEx respectively.

For all datasets except for *Mutagenesis*, ExCut achieved, in general, better results w.r.t. the *ACC* value than the state-of-the-art methods. Furthermore, ExCut-T results in significantly better clusters on all datasets apart from *Terrorists* compared to Kmeans-T, *i.e.*, the direct application of *Kmeans* on the TransE embedding model. Since the *Terrorists* dataset contains several attributed predicates (e.g., facts over numerical values), a different language bias for the explanation rules is required.

Our system managed to fully re-discover the ground truth clusters for the two datasets: *UWCSE* and *IMDB*. The accuracy enhancement by ExCut-T compared to the respective baseline (Kmeans-T) exceeds 30% for *IMDB* and *Hepatitis*. Other quality measurements indicate similar increments.

Explanation Quality. Table 3 shows the average quality of the top explanations for the discovered clusters, where the average per cluster coverage (*Cov*) and exclusive coverage (*Exc*) are intrinsic evaluation metrics used as our optimization functions, while the *WRA* measure is the extrinsic one.

The last row presents the quality of the learned explanations for the ground truth clusters; these values are not necessarily 1.0, as perfect explanations under the specified language bias may not exist. We report them as reference points.

ExCut enhances the average *Exc* and *WRA* scores of the clusters' explanations compared to the ones obtained by the baselines. These two measures highlight the exclusiveness of the explanations; making them more representative than *Cov*. Thus, the decrease in the *Cov*, as in *Terrorist*, is acceptable, given that it is in favor of increasing them.

Similar to the clustering results, for *UWCSE* and *IMDB* our method achieved the explanations quality of the ground truth. For other datasets, our method obtained higher explanations quality than the respective baselines. This demon-

Table 4. Quality of the clusters and the explanations found in Large-scale KGs.

Methods	LUBM Courses						Yago Artwork						
	ACC	ARI	NMI	Cov	Exc	WRA	ACC	ARI	NMI	Cov	Exc	WRA	
Baselines	DEC	0.92	0.70	0.66	0.96	0.95	0.19	0.56	0.44	0.57	0.92	0.49	0.11
	Kmeans-T	0.50	0.00	0.00	0.46	0.03	0.01	0.52	0.42	0.58	0.92	0.42	0.11
ExCut-T	belongToCl	1.00	1.00	1.00	1.00	1.00	0.25	0.82	0.63	0.59	0.85	0.70	0.16
	sameCLAs	0.88	0.57	0.53	0.91	0.79	0.19	0.97	0.91	0.90	0.95	0.93	0.21
	entExplCl	1.00	1.00	1.00	1.00	1.00	0.25	0.97	0.92	0.91	0.95	0.93	0.21
	followExpl	1.00	1.00	1.00	1.00	1.00	0.25	0.88	0.73	0.70	0.86	0.78	0.17
Ground truth	-	-	-	1.00	1.00	0.25	-	-	-	0.95	0.93	0.21	-

Table 5. Explanations of clusters *song*, *book*, and *movie* from Yago KG. ($\forall X \in C_i$)

Explanations	Kmeans-T			ExCut-T			
	Cov	Exc	WRA	Cov	Exc	WRA	
C_1 <i>created</i> (Y, X), <i>bornIn</i> (Y, Z)	0.94	0.55	0.13	<i>created</i> (Y, X), <i>type</i> (Y, artist)	0.99	0.96	0.21
<i>created</i> (Y, X), <i>type</i> (Y, artist)	0.49	0.45	0.10	<i>created</i> (Y, X), <i>won</i> (Y, grammy)	0.57	0.57	0.12
<i>created</i> (Y, X), <i>type</i> (Y, writer)	0.52	0.44	0.10	<i>created</i> (Y, X), <i>type</i> (Y, person)	0.84	0.48	0.11
C_2 <i>directed</i> (Y, X)	0.92	0.56	0.11	<i>created</i> (Y, X), <i>type</i> (Y, writer)	0.99	0.91	0.19
<i>directed</i> (Y, X), <i>gender</i> (Y, male)	0.89	0.54	0.10	<i>created</i> (Y, X), <i>diedIn</i> (Y, Z)	0.46	0.20	0.04
<i>created</i> (Y, X), <i>type</i> (Y, person)	0.71	0.52	0.06	<i>created</i> (Y, X)	1.00	0.00	0.05
C_3 <i>actedIn</i> (Y, X), <i>type</i> (Y, person)	0.58	0.30	0.07	<i>actedIn</i> (Y, X)	0.81	0.81	0.19
<i>locatedIn</i> (X, Y), <i>hasLang</i> (Y, Z)	0.60	0.29	0.07	<i>actedIn</i> (Y, X), <i>bornIn</i> (Y, Z)	0.79	0.79	0.18
<i>locatedIn</i> (X, Y), <i>currency</i> (Y, Z)	0.60	0.29	0.07	<i>actedIn</i> (Y, X), <i>type</i> (Y, person)	0.78	0.78	0.18

strates the effectiveness of the proposed feedback mechanism in adapting the embedding model to better capture the graph structures in the input KGs.

Results on Large-Scale KGs. Table 4 presents quality measures for clustering and explainability of ExCut running with TransE on *LUBM* and *YAGO*. ExCut succeeds to compute the ground truth clusters on *LUBM*. Despite the noise in *YAGO*, it achieves approximately 40% enhancement of the clustering accuracy. The explanation quality is also improved. ReCent did not scale on *LUBM* and *YAGO* due to memory requirements.

Human-Understandability. For illustration in Table 5, we present the top-3 explanations for each cluster computed by ExCut along with their quality on the *YAGO* KG. In the ground truth, C_1, C_2, C_3 are clusters for entities of the type *Songs*, *Books*, and *Movies* respectively. One can observe that the explanations generated by ExCut-T are more intuitive and of higher quality than those obtained using Kmeans-T. The correlation between the explanation relevance and the used quality metrics can also be observed.

Figure 5 summarizes the results of the 50 responses collected via the user-study. Each bar shows the ratio of responses exclusively matching explanation r_i to each of the provided clusters. The results show that the majority of the participants exclusively matched explanations r_3 and r_{10} to *movies*; r_7 and r_9 to *books*; and r_6 and r_8 to *songs*. The explanations r_3, r_6 , and r_9 have been learned by ExCut. The high relative exclusive matching ratio to the correspond-

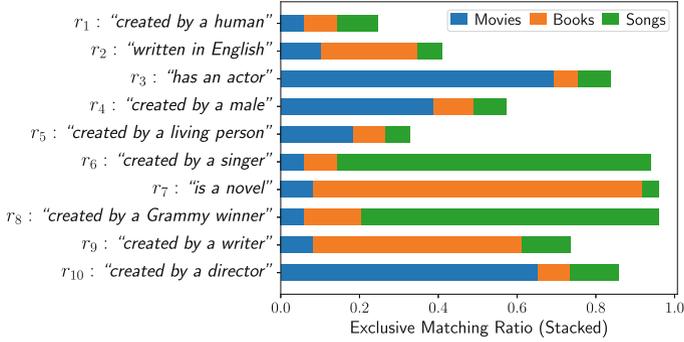


Fig. 5. Ratio of explanation-to-cluster pairs exclusively matched.

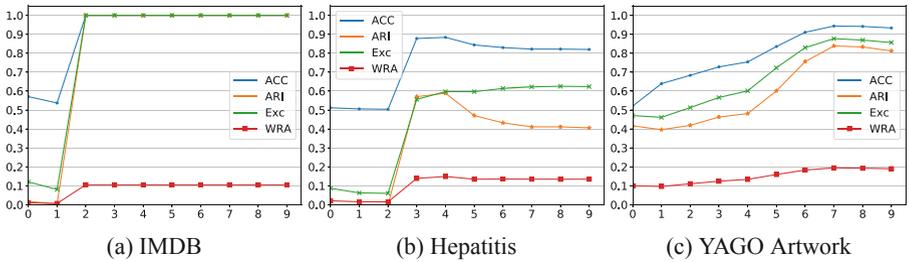


Fig. 6. ExCut-T clustering and explanations quality over the iterations(x-axis).

ing correct cluster for the ExCut explanations demonstrates their usefulness in differentiating between the given clusters.

Results Analysis. In Fig. 6, we present a sample for the quality of the clusters and the aggregated quality of their top explanations over 10 iterations of ExCut-T using the *followExpl* configuration. In general, clustering and explanations qualities consistently improved over iterations, which demonstrates the advantage of the introduced embedding fine-tuning procedure. For *IMDB*, the qualities drop at the beginning, but increase and reach the highest values at the third iteration. This highlights the benefit of accumulating the auxiliary triples for enhancing the feedback signal, thus preventing the embedding tuning from diverging. The charts also show a correlation between the clustering and explanation quality, which proves our hypothesis that the introduced exclusive coverage measure (*Exc*) is useful for computing good clusters.

With respect to the effects of different embeddings and feedback modeling, as shown in Tables 2 and 3, we observe that ExCut with *TransE* is more robust than with *ComplEx* regardless of the feedback modeling method. Furthermore, modeling the feedback using *followExpl* strategy leads to better results on the majority of the datasets, especially for large-scale KGs. This reflects the benefit of passing richer feedback to the embedding, as it allows for better entity positioning in the latent space.

6 Related Work

Clustering relational data has been actively studied (*e.g.*, [4, 6, 7, 16, 25]). The majority of the existing approaches are based on finding interesting features in KGs and defining distance measures between their vectors. Our work is conceptually similar, but we let embedding model identify the features implicitly instead of computing them on the KG directly, which is in spirit of linked data propositionalization [24].

A framework for explaining given high-quality clusters using linked data and inductive logic programming has been proposed in [27, 28]. While [28] aims at explaining existing clusters, we focus on performing clustering and explanation learning iteratively to discover high-quality clusters with explanations. The work [12] targets interpreting embedding models by finding concept spaces in node embeddings and linking them to a simple external type hierarchy. This is different from our method of explaining clusters computed over embeddings by learning rules from a given KG. Similarly, [2] proposes a method for learning conceptual space representations of known concepts by associating a Gaussian distribution over a learned vector space with each concept. In [10, 23] the authors introduce methods for answering logical queries over the embedding space. In contrast, in our work, the concepts are not given but rather need to be discovered.

While the step of explanation learning in our method is an adaptation of [8], the extension of other exact symbolic rule learning methods [18, 22] is likewise possible. In principle, one can also employ neural-based rule learners for our needs, such as [20, 21, 34]; however the integration of our exclusive rule coverage scoring function into such approaches is challenging, and requires further careful investigation.

Several methods recently focused on combining [11, 35] and comparing [5, 19] rule learning and embedding methods. The authors of [11] propose to rank rules learned from KGs by relying both on their embedding-based predictive quality and traditional rule measures, which is conceptually different from our work. In [35] an iterative method for joint learning of linear-map embeddings and OWL axioms (without nominals) has been introduced. The triples inferred by the learned rules are injected into the KG, before the embedding is re-trained from scratch in the subsequent iteration. In contrast, the rule-based feedback generated by ExCut is not limited to only fact predictions, but encodes further structural similarities across entities. Furthermore, we do not re-train the whole model from scratch, but rather adapt the embedding of target entities accounting for the feedback. Finally, unlike [35], the rules that we learn support constants, which allow to capture a larger variety of explanations.

7 Conclusion

We have proposed ExCut, an approach for explainable KG entity clustering, which iteratively utilizes embeddings and rule learning methods to compute

accurate clusters and human-readable explanations for them. Our approach is flexible, as any embedding model can be used. Experiments show the effectiveness of ExCut on real-world KGs.

There are several directions for future work. Considering more general rules (*e.g.*, with negations) in the *Rule Learning* component of our method or exploiting several embedding models instead of a single one in the *Embedding-based Clustering* step should lead to cleaner clusters. Further questions to study include the analysis of how well our method performs when the number of clusters is very large, and how the feedback from the rules can be used to determine the number of clusters automatically.

References

1. Bordes, A., Usunier, N., Garcia-Duran, A., Weston, J., Yakhnenko, O.: Translating embeddings for modeling multi-relational data. In: NeurIPS, pp. 2787–2795 (2013)
2. Bouraoui, Z., Schockaert, S.: Learning conceptual space representations of interrelated concepts. In: IJCAI, pp. 1760–1766 (2018)
3. Costabello, L., Pai, S., Van, C.L., McGrath, R., McCarthy, N., Tabacof, P.: AmpliGraph: a library for representation learning on knowledge graphs (2019)
4. Dumancic, S., Blockeel, H.: An expressive dissimilarity measure for relational clustering over neighbourhood trees. MLJ (2017)
5. Dumancic, S., García-Durán, A., Niepert, M.: On embeddings as an alternative paradigm for relational learning. CoRR [arXiv:abs/1806.11391v2](https://arxiv.org/abs/1806.11391v2) (2018)
6. Fanizzi, N., d’Amato, C., Esposito, F.: Conceptual clustering and its application to concept Drift and Novelty Detection. In: Bechhofer, S., Hauswirth, M., Hoffmann, J., Koubarakis, M. (eds.) ESWC 2008. LNCS, vol. 5021, pp. 318–332. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-68234-9_25
7. Fonseca, N.A., Costa, V.S., Camacho, R.: Conceptual clustering of multi-relational data. In: ILP, pp. 145–159 (2011)
8. Galárraga, L., Teflioudi, C., Hose, K., Suchanek, F.M.: Fast rule mining in ontological knowledge bases with AMIE+. VLDB J. **24**(6), 707–730 (2015)
9. Guo, Y., Pan, Z., Heflin, J.: LUBM: a benchmark for OWL knowledge base systems. J. Web Semant. **3**(2–3), 158–182 (2005)
10. Hamilton, W.L., Bajaj, P., Zitnik, M., Jurafsky, D., Leskovec, J.: Embedding logical queries on knowledge graphs. In: NeurIPS, pp. 2030–2041 (2018)
11. Ho, V.T., Stepanova, D., Gad-Elrab, M.H., Kharlamov, E., Weikum, G.: Rule learning from knowledge graphs guided by embedding models. In: Vrandečić, D., et al. (eds.) ISWC 2018. LNCS, vol. 11136, pp. 72–90. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-00671-6_5
12. Idahl, M., Khosla, M., Anand, A.: Finding interpretable concept spaces in node embeddings using knowledge bases. In: Cellier, P., Driessens, K. (eds.) ML/KDD, pp. 229–240 (2020)
13. Knobbe, A.J., Ho, E.K.Y.: Pattern teams. In: PKDD, pp. 577–584 (2006)
14. Lavrač, N., Flach, P., Zupan, B.: Rule evaluation measures: a unifying view. In: ILP, pp. 174–185 (1999)
15. Lin, Y., Liu, Z., Luan, H., Sun, M., Rao, S., Liu, S.: Modeling relation paths for representation learning of knowledge bases. In: EMNLP, pp. 705–714 (2015)

16. Lisi, F.A.: A pattern-based approach to conceptual clustering in FOL. In: Schärfe, H., Hitzler, P., Øhrstrøm, P. (eds.) ICCS-ConceptStruct 2006. LNCS (LNAI), vol. 4068, pp. 346–359. Springer, Heidelberg (2006). https://doi.org/10.1007/11787181_25
17. MacQueen, J., et al.: Some methods for classification and analysis of multivariate observations. In: Proceedings of the fifth Berkeley Symposium on Mathematical Statistics and Probability, vol. 1, pp. 281–297 (1967)
18. Meilicke, C., Chekol, M.W., Ruffinelli, D., Stuckenschmidt, H.: Anytime bottom-up rule learning for knowledge graph completion. In: IJCAI, pp. 3137–3143 (2019)
19. Meilicke, C., Fink, M., Wang, Y., Ruffinelli, D., Gemulla, R., Stuckenschmidt, H.: Fine-grained evaluation of rule- and embedding-based systems for knowledge graph completion. In: Vrandečić, D., et al. (eds.) ISWC 2018. LNCS, vol. 11136, pp. 3–20. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-00671-6_1
20. Omran, P.G., Wang, K., Wang, Z.: An embedding-based approach to rule learning in knowledge graphs. *IEEE Trans. Knowl. Data Eng.* 1–1 (2019)
21. Omran, P.G., Wang, K., Wang, Z.: Scalable rule learning via learning representation. In: IJCAI, pp. 2149–2155 (2018)
22. Ortona, S., Meduri, V.V., Papotti, P.: Robust discovery of positive and negative rules in knowledge bases. In: ICDE, pp. 1168–1179. IEEE (2018)
23. Ren, H., Hu, W., Leskovec, J.: Query2box: Reasoning over knowledge graphs in vector space using box embeddings. In: ICLR (2020)
24. Ristoski, P., Paulheim, H.: A comparison of propositionalization strategies for creating features from linked open data. In: 1st Linked Data for Knowledge Discovery Workshop at ECML/PKDDat: Nancy, France (2014)
25. Pérez-Suárez, A., Martínez-Trinidad, J.F., Carrasco-Ochoa, J.A.: A review of conceptual clustering algorithms. *Artif. Intell. Rev.* **52**(2), 1267–1296 (2018). <https://doi.org/10.1007/s10462-018-9627-1>
26. Suchanek, F.M., Kasneci, G., Weikum, G.: Yago: A core of semantic knowledge. In: Proceedings of WWW, pp. 697–706 (2007)
27. Tiddi, I., d’Aquin, M., Motta, E.: Dedalo: Looking for clusters explanations in a labyrinth of linked data. In: Presutti, V., d’Amato, C., Gandon, F., d’Aquin, M., Staab, S., Tordai, A. (eds.) *The Semantic Web: Trends and Challenges (ESWC 2014)*. Lecture Notes in Computer Science, vol. 8465, pp. 333–348. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-07443-6_23
28. Tiddi, I., d’Aquin, M., Motta, E.: Data patterns explained with linked data. In: Bifet, A., et al. (eds.) *Machine Learning and Knowledge Discovery in Databases (ECML PKDD 2015)*. Lecture Notes in Computer Science, vol. 9286, pp. 271–275. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-23461-8_28
29. Trouillon, T., Welbl, J., Riedel, S., Gaussier, É., Bouchard, G.: Complex embeddings for simple link prediction. In: ICML, pp. 2071–2080 (2016)
30. Vrandečić, D., Krötzsch, M.: Wikidata: a free collaborative knowledgebase. *Commun. ACM* **57**(10), 78–85 (2014)
31. Wang, C., Pan, S., Hu, R., Long, G., Jiang, J., Zhang, C.: Attributed graph clustering: a deep attentional embedding approach. In: IJCAI, pp. 3670–3676 (2019)
32. Xie, J., Girshick, R.B., Farhadi, A.: Unsupervised deep embedding for clustering analysis. In: ICML, pp. 478–487 (2016)
33. Xie, R., Liu, Z., Jia, J., Luan, H., Sun, M.: Representation learning of KGs with entity descriptions. In: AAAI, pp. 2659–2665 (2016)

34. Yang, F., Yang, Z., Cohen, W.W.: Differentiable learning of logical rules for knowledge base reasoning. In: NeurIPS, pp. 2319–2328 (2017)
35. Zhang, W., et al.: Iteratively learning embeddings and rules for knowledge graph reasoning. In: WWW, pp. 2366–2377 (2019)