

Automating RDF Dataset Transformation and Enrichment

Mohamed Ahmed Sherif^(✉), Axel-Cyrille Ngonga Ngomo, and Jens Lehmann

Department of Computer Science, University of Leipzig, 04109 Leipzig, Germany
{sherif,ngonga,lehmann}@informatik.uni-leipzig.de

Abstract. With the adoption of RDF across several domains, come growing requirements pertaining to the completeness and quality of RDF datasets. Currently, this problem is most commonly addressed by manually devising means of enriching an input dataset. The few tools that aim at supporting this endeavour usually focus on supporting the manual definition of enrichment pipelines. In this paper, we present a supervised learning approach based on a refinement operator for enriching RDF datasets. We show how we can use exemplary descriptions of enriched resources to generate accurate enrichment pipelines. We evaluate our approach against eight manually defined enrichment pipelines and show that our approach can learn accurate pipelines even when provided with a small number of training examples.

1 Introduction

Over the last years, the Linked Data principles have been used across academia and industry to publish and consume linked data [16]. With this adoption of Linked data come novel challenges pertaining to the integration of these datasets for dedicated applications such as tourism, question answering, enhanced reality and many more. Providing consolidated and integrated datasets for these applications demands the specification of data enrichment pipelines, which describe how data from different sources is to be integrated and altered so as to abide by the precepts of the application developer or data user. Currently, most developers implement customized pipelines by compiling sequences of tools manually and connecting them via customized scripts. While this approach most commonly leads to the expected results, it is time-demanding and resource-intensive. Moreover, the results of this effort can most commonly only be reused for new versions of the input data but cannot be ported easily to other datasets. Over the last years, a few frameworks for RDF data enrichment such as LDIF¹ and DEER² have been developed. The frameworks provide enrichment methods such as entity recognition [22], link discovery [15] and schema enrichment [4]. However, devising appropriate configurations for these tools can prove a difficult endeavour, as the tools require (1) choosing the right sequence of enrichment functions

¹ <http://ldif.wbsg.de/>.

² <http://aksw.org/Projects/DEER.html>.

and (2) configuring these functions adequately. Both the first and second task can be tedious.

In this paper, we address this problem by presenting a supervised machine learning approach for the automatic detection of enrichment pipelines based on a refinement operator and self-configuration algorithms for enrichment functions. Our approach takes pairs of concise bounded descriptions (CBDs) of resources $\{(k_1, k'_1) \dots (k_n, k'_n)\}$ as input, where k'_i is the enriched version of k_i . Based on these pairs, our approach can learn sequences of atomic enrichment functions that aim to generate each k'_i out of the corresponding k_i . The output of our approach is an enrichment pipeline that can be used on whole datasets to generate enriched versions.

Overall, we provide the following core contributions: (1) We define a supervised machine learning algorithm for learning dataset enrichment pipelines based on a refinement operator. (2) We provide self-configuration algorithms for five atomic enrichment steps. (3) We evaluate our approach on eight manually defined enrichment pipelines on real datasets.

2 Preliminaries

Enrichment: Let \mathcal{K} be the set of all RDF knowledge bases. Let $K \in \mathcal{K}$ be a finite RDF knowledge base. K can be regarded as a set of triples $(s, p, o) \in (\mathcal{R} \cup \mathcal{B}) \times \mathcal{P} \times (\mathcal{R} \cup \mathcal{L} \cup \mathcal{B})$, where \mathcal{R} is the set of all resources, \mathcal{B} is the set of all blank nodes, \mathcal{P} the set of all predicates and \mathcal{L} the set of all literals. Given a knowledge base K , the idea behind *knowledge base enrichment* is to find an *enrichment pipeline* $M : \mathcal{K} \rightarrow \mathcal{K}$ that maps K to an enriched knowledge base K' with $K' = M(K)$. We define M as an ordered list of *atomic enrichment functions* $m \in \mathcal{M}$, where \mathcal{M} is the set of all atomic enrichment functions. $2^{\mathcal{M}}$ is used to denote the power set of \mathcal{M} , i.e. the set of all enrichment pipelines. The order of elements in M determines the execution order, e.g. for an $M = (m_1, m_2, m_3)$ this means that m_1 will be executed first, then m_2 , finally m_3 . Formally,

$$M = \begin{cases} \phi & \text{if } K = K', \\ (m_1, \dots, m_n), \text{ where } m_i \in \mathcal{M}, 1 \leq i \leq n & \text{otherwise,} \end{cases} \quad (1)$$

where ϕ is the empty sequence. Moreover, we denote the number of elements of M with $|M|$. Considering that a knowledge base is simply a set of triples, the task of any atomic enrichment function is to (1) determine a set of triples Δ^+ to be added to the source knowledge base and/or (2) determine a set of triples Δ^- to be deleted from the source knowledge base. Any other enrichment process can be defined in terms of Δ^+ and Δ^- , e.g. altering triples can be represented as combination of addition and deletion.

In this article we cover two problems: (1) how to create self-configurable atomic enrichment functions $m \in \mathcal{M}$ capable of enriching a dataset and (2) how to automatically generate an enrichment pipeline M . As a running example, we use the portion of *DrugBank* shown in Fig. 1. The goal of the enrichment here is

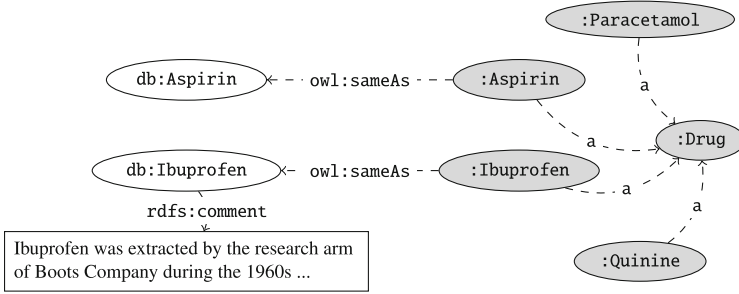


Fig. 1. RDF graph of the running example. Ellipses are RDF resources, literals are rectangular nodes. Gray nodes stand for resources in the input knowledge base while nodes with a white background are part of an external knowledge base.

to gather information about companies related to drugs for a market study. To this end, the `owl:sameAs` links to *DBpedia* (prefix `db`) need to be dereferenced. Their `rdfs:comment` then needs to be processed using an entity spotter that will help retrieve resources such as the **Boots Company**. Then, these resources need to be attached directly to the resources in the source knowledge base, e.g., by using the `:relatedCompany` property. Finally, all subjects need to be conformed under one subject authority (prefix `ex`).

Refinement Operators: Below, we give definitions of refinement operators and their properties. Refinement operators have traditionally been used, e.g. in [11], to traverse search spaces in structured machine learning problems. Their theoretical properties give an indication of how suitable they are within a learning algorithm in terms of accuracy and efficiency.

Definition 1 (Refinement Operator and Properties). *Given a quasi-ordered space (S, \preceq) an upward refinement operator r is a mapping from S to 2^S such that $\forall s \in S : s' \in r(s) \Rightarrow s \preceq s'$. s' is then called a generalization of s . A pipeline $M_2 \in \mathcal{M}$ belongs to the refinement chain of $M_1 \in \mathcal{M}$ iff $\exists i \in \mathbb{N} : M_2 \in r^i(M_1)$, where $r^0(M) = M$ and $r^i(M) = r(r^{i-1}(M))$. A refinement operator r over the quasi-ordered space (S, \preceq) can abide by the following criteria. r is finite iff $r(s)$ is finite for all $s \in S$. r is proper if $\forall s \in S, s' \in r(s) \Rightarrow s \neq s'$. r is complete if for all s and s' , $s' \preceq s$ implies that there is a refinement chain between s and s' . A refinement operator r over the space (S, \preceq) is redundant if two different refinement chains can exist between $s \in S$ and $s' \in S$.*

3 Knowledge Base Enrichment Refinement Operator

Our refinement operator expects the set of atomic enrichment functions \mathcal{M} as input and returns an enrichment pipeline M as output. Each positive example $e \in \mathcal{E}$ is a pair of CBDs (k, k') , with $k \subseteq K$ and $k' \subseteq K'$, the K' stands




Fig. 2. Ibuprofen concise bound description before and after enrichment

for the enriched version of K . Note that we model CBDs as sets of RDF triples. Moreover, we denote the resource with the CBD k as $resource(k)$. For our running example, the set \mathcal{E} could contain the pair shown in Fig. 2a as k and in Fig. 2b as k' .

The set of all first elements of the pairs contained in \mathcal{E} is denoted $source(\mathcal{E})$. The set of all second elements is denoted $target(\mathcal{E})$. To compute the refinement pipeline M , we employ an upward refinement operator (which we dub ρ) over the space $2^{\mathcal{M}}$ of all enrichment pipelines. We write $M \supseteq M'$ when M' is a subsequence of M , i.e., $m'_i \in M' \rightarrow m'_i = m_i$, where m_i resp. m'_i is the i^{th} element of M resp. M' .

Proposition 1 (Induced Quasi-Ordering). \supseteq induces a quasi-ordering over the set $2^{\mathcal{M}}$.

Proof. The reflexivity of \supseteq follows from each M being a subsequence of itself. The transitivity of \supseteq follows from the transitivity of the subsequence relation. Note that \supseteq is also antisymmetric. \square

We define our refinement operator over the space $(2^{\mathcal{M}}, \supseteq)$ as follows:

$$\rho(M) = \bigcup_{\forall m \in \mathcal{M}} M ++ m \quad (++ \text{ is the list append operator}) \quad (2)$$

We define precision $P(M)$ and recall $R(M)$ achieved by an enrichment pipeline on \mathcal{E} as

$$P(M) = \frac{\left| \bigcup_{k \in source(\mathcal{E})} M(k) \cap \bigcup_{k' \in target(\mathcal{E})} k' \right|}{\left| \bigcup_{k \in source(\mathcal{E})} M(k) \right|}, R(M) = \frac{\left| \bigcup_{k \in source(\mathcal{E})} M(k) \cap \bigcup_{k' \in target(\mathcal{E})} k' \right|}{\left| \bigcup_{k' \in target(\mathcal{E})} k' \right|}. \quad (3)$$

The F-measure $F(M)$ is then

$$F(M) = \frac{2P(M)R(M)}{P(M) + R(M)}. \quad (4)$$

Using Fig. 2a from our running example as source and Fig. 2b as target with the CBD of `:Iboprofen` being the only positive example, an empty enrichment pipeline $M = \phi$ would have a precision of 1, a recall of $\frac{3}{4}$ and an F-measure of $\frac{6}{7}$. Having defined our refinement operator, we now show that ρ is finite, proper, complete and not redundant.

Proposition 2. ρ is finite.

Proof. This is a direct consequence of \mathcal{M} being finite. \square

Proposition 3. ρ is proper.

Proof. As the quasi order is defined over subsequences, i.e. the space $(2^{\mathcal{M}}, \supseteq)$, and we have $|M'| = |M| + 1$ for any $M' \in \rho(M)$, ρ is trivially proper. \square

Proposition 4. ρ is complete.

Proof. Let M resp. M' be an enrichment pipeline of length n resp. n' with $M' \supseteq M$. Moreover, let m'_i be the i^{th} element of M' . Per definition, $M \uplus m'_{n+1} \in \rho(M)$. Hence, by applying ρ $n' - n$ times, we can generate M' from M . We can thus conclude that ρ is complete. \square

Proposition 5. ρ is not redundant.

Proof. ρ being redundant would mean that there are two refinement chains that lead to a single refinement pipeline M . As our operator is equivalent to the list append operation, it would be equivalent to stating that two different append sequences can lead to the same sequence. This is obviously not the case as each element of the list M is unique, leading to exactly one sequence that can generate M . \square

4 Learning Algorithm

The learning algorithm is inspired by refinement-based approaches from inductive logic programming. In these algorithms, a search tree is iteratively built up using heuristic search via a fitness function. We formally define a node N in a search tree to be a triple (M, f, s) , where M is the *enrichment pipeline*, $f \in [0, 1]$ is the F-measure of M (see Eq. 4), and $s \in \{\text{normal}, \text{dead}\}$ is the status of the node. Given a search tree, the heuristic selects the fittest node in it, where fitness is based on both F-measure and complexity as defined below.

4.1 Approach

For the automatic generation of enrichment pipeline specifications, we created a learning algorithm based on the previously defined refinement operator. Once provided with training examples, the approach is fully automatic. The pseudo-code of our algorithm is presented in Algorithm 4.1.

Our learning algorithm has two inputs: a set of positive examples \mathcal{E} and a set of atomic enrichment operators \mathcal{M} . \mathcal{E} contains pairs of (k, k') where each k contains a CBD of one resource from an arbitrary source knowledge base K and k' contains the CBD of the same resource after applying some manual enrichment. Given \mathcal{E} , the goal of our algorithm is to learn an enrichment pipeline M that maximizes $F(M)$ (see Eq. 4).

As shown in Algorithm 4.1, our approach starts by generating an empty refinement tree τ which contains only an empty root node. Using \mathcal{E} , the algorithm then accumulates all the original CBDs in k ($\text{SOURCE}(\mathcal{E})$). Using the same procedure, k' is accumulated from \mathcal{E} as the knowledge base containing the enriched version of k ($\text{TARGET}(\mathcal{E})$). Until a termination criterion holds (see Sect. 4.3), the algorithm keeps expanding the most promising node (see Sect. 4.2). Finally, the algorithm ends by returning the best pipeline found in τ : ($\text{GetPipeline}(\text{GetMaxQualityNode}(\tau))$).

Having a *most promising node* t at hand, the algorithm first applies our refinement operator (see Eq. 2) against the most promising enrichment pipeline M_{old} included in t to generate a set of atomic enrichment functions $\mathcal{M} \leftarrow \rho(M_{old})$. Consequently, using both k_{old} (as the knowledge base generated by applying M_{old} against k) and k' , the algorithm applies the self configuration process of the current atomic enrichment function $m \leftarrow \text{SELFCONFIG}(m, k_{old}, k)$ to generate a set of parameters P (a detailed description for this process is found in Sect. 5). Afterwards, the algorithm runs m against k_{old} to generate the new enriched knowledge base $k_{new} \leftarrow m(k_{old}, P)$. A dead node $N \leftarrow \text{CREATENODE}(M, 0, \text{dead})$ is created in two cases: (1) m is inapplicable to k_{old} (i.e., $P == \text{null}$) or (2) m does no enrichment at all (i.e., k_{new} is isomorphic³ to k_{old}). Otherwise, the algorithm computes the F-measure f of the generated dataset k_{new} . M along with f are then used to generate a new search tree node $N \leftarrow \text{CREATENODE}(M, f, \text{normal})$. Finally, N is added as a child of t ($\text{ADDCHILD}(t, N)$).

4.2 Most Promising Node Selection

Here we describe the process of selecting the most promising node $t \in \tau$ as in $\text{GETMOSTPROMISINGNODE}()$ subroutine in Algorithm 4.1. First, we define *node complexity* as linear combination of the node’s children count and level. Formally,

Definition 2. (Node Complexity). *The complexity of a node $N = (M, f, s)$ in a refinement tree τ is a function $c : N \times \tau \rightarrow [0, 1]$, where $c(N, \tau) = \alpha \frac{|N_d|}{|\tau|} + \beta \frac{N_l}{\tau_d}$, $|N_d|$ is number of all N ’s descendant nodes, $|\tau|$ is the total number of nodes in τ , N_l is N ’s level, τ_d is τ ’s depth, α is the children penalty weight, β is the level penalty weight and $\alpha + \beta = 1$. Seeking for simplicity, we will use the $c(N)$ instead of $c(N, \tau)$ in the rest of this paper.*

We can then define the fitness $f(N)$ of a *normal* node N as the difference between its enrichment pipeline F-measure (Eq. 4) and weighted complexity. $f(N)$ is zero for *dead* nodes. Formally,

Definition 3. (Node Fitness). *Let $N = (M, f, s)$ be a node in a refinement tree τ , N ’s fitness is the function*

$$f(N) = \begin{cases} 0 & \text{if } s = \text{dead}, \\ F(M) - \omega \cdot c(N) & \text{if } s = \text{normal}. \end{cases} \tag{5}$$

³ <http://www.w3.org/TR/rdf11-concepts/>.

where M is the enrichment pipeline contained in the node N , ω is the complexity weight and $0 \leq \omega \leq 1$.

Note, that we use the *complexity* of pipelines as second criterion, which makes the algorithm (1) more flexible in searching less explored areas of the search space, and (2) leads to simpler specification being preferred over more complex ones (Occam's razor [3]). The parameter ω can be used to control the trade-off between a greedy search ($\omega = 0$) and search strategies closer to breadth first search ($\omega > 0$). The fitness function can be defined independently of the core learning algorithm.

Consequently, the most promising node is the node with the maximum fitness through the whole refinement tree τ . Formally, the most promising node t is defined as $t = \arg \max_{\forall N \in \tau} f(N)$, where N is not a *dead* node. Note that if several nodes achieve a maximum fitness, the algorithm chooses the shortest node as it aims to generate the simplest enrichment pipeline possible.

Algorithm 4.1. ENRICHMENTPIPELINELEARNER(\mathcal{E}^+ , \mathcal{M})

comment: initialize τ

$\tau \leftarrow \text{CREATEROOTNODE}()$

$k \leftarrow \text{SOURCE}(\mathcal{E})$

$k' \leftarrow \text{TARGET}(\mathcal{E})$

repeat

comment: Expand most promising node of τ

$t \leftarrow \text{GETMOSTPROMISINGNODE}(\tau)$

$M_{old} \leftarrow \text{GETPIPELINE}(t)$

$\mathcal{M} \leftarrow \rho(M_{old})$

comment: Create a child of t for each $m \in \mathcal{M}$

for each $m \in \mathcal{M}$

 do {

$k_{old} \leftarrow M_{old}(k)$

$P \leftarrow \text{SELFCONFIG}(m, k_{old}, k')$

$k_{new} \leftarrow m(k_{old}, P)$

if $P == \text{null}$ **or** $k_{new} == k_{old}$

then { $N \leftarrow \text{CREATENODE}(M, 0, \text{dead})$

else { $f \leftarrow F(m)$

$N \leftarrow \text{CREATENODE}(M, f, \text{normal})$

$\text{ADDCHILD}(t, N)$

 }

until $\text{TERMINATIONCRITERIONHOLDS}(\tau)$

return ($\text{GETPIPELINE}(\text{GETMAXQUALITYNODE}(\tau))$)

4.3 Termination Criteria

The subroutine $\text{TERMINATIONCRITERIONHOLDS}()$ in Algorithm 4.1 can check several termination criteria depending on configuration: (1) optimal enrichment

pipeline found (i.e., a fixpoint is reached), (2) maximum number of iterations reached, (3) maximum number of refinement tree nodes reached, or a combination of the aforementioned criteria. Note that the termination criteria can be defined independently of the core learning algorithm.

5 Self-Configuration

To learn an appropriate specification from the input positive examples, we need to develop self-configuration approaches for each of our framework’s atomic enrichment functions. The input for each of these self-configuration procedures is the same set of positive examples \mathcal{E} provided to our pipeline learning algorithm (Algorithm 4.1). The goal of the self-configuration process of an enrichment function is to generate a set of parameters $P = \{(mp_1, v_1), \dots, (mp_m, v_m)\}$ able to reflect \mathcal{E} as well as possible. In cases when insufficient data is contained in \mathcal{E} to carry out the self-configuration process, an empty list of parameters is returned to indicate inapplicability of the enrichment function.

5.1 Dereferencing Enrichment Functions

The idea behind the self-configuration process of the enrichment by *dereferencing* is to find the set of predicates D_p from the enriched CBDs that are missing from source CBDs. Formally, for each CBD pair (k, k') construct a set $D_p \subseteq \mathcal{P}$ as follows: $D_p = \{p' : (s', p', o') \in k'\} \setminus \{p : (s, p, o) \in k\}$. The dereferencing enrichment function will *dereference* the object of each triple of k_i given that this object is an external URI, i.e. all o in k_i with $(s, p, o) \in k_i$, $o \in \mathcal{R}$ and o is not in the local namespace of the dataset will be dereferenced. Dereferencing an object returns a set of triples. Those are filtered using the previously constructed property set D_p , i.e. when dereferencing o the enrichment function only retains triples with subject o and a predicate contained in D_p . The resulting set of triples is added to the input dataset.

We illustrate the process using our running example: In the first step, we compute the set $D_p = \{:\text{relatedCompany}, \text{rdfs:comment}\}$ which consists of the properties occurring in the target but not in the source CBD. In the second step, we collect the set of resources to dereference, which only consists of the element `db:Ibuprofen`. In the third step, we perform the actual dereferencing operation and retain triples for which the subject is `db:Ibuprofen` and the predicate is either `:relatedCompany` or `rdfs:comment`. In our example, no triples with predicate `:relatedCompany` exist, but we will find the desired triple (`db:Ibuprofen`, `rdfs:comment`, "`Ibuprofen ...`"), which is then added to the input dataset.

5.2 Linking Enrichment Function

The aim of *link discovery* is as follows: Given two sets $R_s \subseteq \mathcal{R}$ of source resources and $R_t \subseteq \mathcal{R}$ of target resources, we aim to discover links $L \subseteq R_s \times R_t$ such that for any $(s, t) \in L$ we have $\delta(s, t) \leq \theta$ where δ is a similarity function and θ a

threshold value. The goal of the linking enrichment function is to learn so called *link specifications* including a similarity function δ and a threshold θ . To this aim, we rely on an unsupervised hierarchical search approach, which optimizes a target function akin to F-measure. The search space of all link specifications is split into a grid and the approach computes the objective function for all points in the grid. Thereafter, the region surrounding the point which achieves the highest score is selected as new search space. This approach is applied iteratively until a stopping condition (e.g., a maximal number of iterations) is reached. More details can be found at [18].

5.3 NLP Enrichment Function

The basic idea here is to enable the extraction of all possible named entity types. If this leads to the retrieval of too many entities, the unwanted predicates and resources can be discarded in a subsequent step. The self-configuration of the NLP enrichment function is parameter-free and relies on FOX [17]. The application of the NLP self configuration to our running example generates all possible entities included in the literal object of the `rdfs:comment` predicate. The result is a set of related named entities all of them related to our `ex:Iboprufen` object by the default predicate `fox:relatedTo` as shown Fig. 3a. In the following 2 sections we will see how our enrichment functions can refine some of the generated triples and delete others.

5.4 Conformation Enrichment Functions

The *conformation*-based enrichment currently allows for both *subject-authority-based conformation* and *predicate-based conformation*. The self-configuration process of *subject-authority-based conformation* starts by finding the most frequent subject authority rk in $source(\mathcal{E})$. Also, it finds the most frequent subject authority rk' in the target dataset $target(\mathcal{E})$. Then this self-configuration process generates the two parameters: (`sourceSubjectAuthority`, rk) and (`targetSubjectAuthority`, rk'). After that, the self-configuration process replaces each subject authority rk in $source(\mathcal{E})$ by rk' .

Back to our running example, the authority self-conformation process generates the two parameters (`sourceSubjectAuthority`, ":") and (`targetSubjectAuthority`, "ex:"). Replacing each ":" by "ex:" generates, in our example, the new conformed URI "ex:Iboprufen".

We define two predicates $p_1, p_2 \in \mathcal{P}$ to be *interchangeable* (denoted $p_1 \rightleftharpoons p_2$) if both of them have the same subject and object. Formally, $\forall p_1, p_2 \in \mathcal{P} : p_1 \rightleftharpoons p_2 \iff \exists s, o \mid (s, p_1, o) \wedge (s, p_2, o)$.

The idea of the self-configuration process of the *predicate conformation* is to change each predicate in the source dataset to its *interchangeable* predicate in the target dataset. Formally, find all pairs $(p_1, p_2) \mid \exists s, p_1, o \in k \wedge \exists s, p_2, o \in k' \wedge (s, p_1, o) \in k \wedge (s, p_2, o) \in k'$. Then, for each pair (p_1, p_2) create two self-configuration parameters (`sourceProperty`, p_1) and (`targetProperty`, p_2). The predicate conformation will replace each occurrence of p_1 by p_2 .




Fig. 3. Ibuprofen CBD after NLP and predicate conformation enrichment

In our example, let us suppose that we ran the NLP-based enrichment first then we got a set of related named entities all of them related to our `ex:Ibuprofen` object by the default predicate `fox:relatedTo` as shown in Fig. 3a. Subsequently, applying the predicate conformation self-configuration will generate `(sourceProperty, fox:relatedTo)` and `(targetProperty, ex:relatedCompany)` parameters. Consequently, the predicate conformation module will replace `fox:relatedTo` by `ex:relatedCompany` to generate Fig. 3b.

5.5 Filter Enrichment Function

The idea behind the self-configuration of *filter*-based enrichment is to preserve only valuable triples in the source CBDs k and discard any unnecessary triples so as to achieve a better match to k' . To this end, the self-configuration process starts by finding the intersection between source and target examples $I = \bigcup_{(k,k') \in \mathcal{E}} k \cap k'$. After that, it generates an enrichment function based on a SPARQL query which is only preserving predicates in I . Formally, the self-configuration results in the parameter set $P = \bigcup_{p \in K \cap K' \cap \mathcal{P}} p$.

Back to our running example, let us continue from the situation in the previous section (Fig. 3b). Performing the self-configuration of filters will generate $P = \{\text{fox:relatedTo}\}$. Actually applying the filter enrichment function will remove all unrelated triples containing the predicate `fox:relatedTo`. Figure 4 shows a graph representation for the whole learned pipeline for our running example.




Fig. 4. Graph representation of the learned pipeline of our running example, where d_1 is the positive example source presented in Fig. 2a and d_6 is the positive example target presented in Fig. 2b.

6 Evaluation

The aim of our evaluation was to quantify how well our approach can automate the enrichment process. We thus assumed being given manually created training examples and having to reconstruct a possible enrichment pipeline to generate target CBDs from the source CBDs. In the following, we present our experimental setup including the pipelines and datasets used. Thereafter, we give an overview of our results, which we subsequently discuss in the final part of this section.

6.1 Experimental Setup

We used three publicly available datasets for our experiments:

1. From the biomedical domain, we chose *DrugBank*⁴ as our first dataset. We chose this dataset because it is linked with many other datasets⁵, from which we can extract enrichment data using our atomic enrichment functions. For our experiments we deployed a manual enrichment pipeline M_{manual} , in which we enrich the drug data found in *DrugBank* using abstracts dereferenced from *DBpedia*, then we conform both *DrugBank* and *DBpedia* source authority URIs to one unified URI. For *DrugBank* we manually deployed two experimental pipelines:
 - $M_{DrugBank}^1 = (m_1, m_2)$, where m_1 is a dereferencing function that dereferences any `dbpedia-owl:abstract` from *DBpedia* and m_2 is an authority conformation function that conforms the *DBpedia* subject authority⁶ to the target subject authority of *DrugBank*⁷.
 - $M_{DrugBank}^2 = M_{DrugBank}^1 \dashv\dashv m_3$, where m_3 is an authority conformation function that conforms *DrugBank*'s authority to the *Example* authority⁸.
2. From the music domain, we chose the *Jamendo*⁹ dataset. We selected this dataset as it contains a substantial amount of embedded information hidden in literal properties such as `mo:biography`. The goal of our enrichment process is to add a geospatial dimension to *Jamendo*, e.g., the location of a recording or place of birth of a musician. To this end, we deployed a manual enrichment pipeline, in which we enrich *Jamendo*'s music data by adding additional geospatial data found by applying the NLP enrichment function

⁴ *DrugBank* is the Linked Data version of the DrugBank database, which is a repository of almost 5000 FDA-approved small molecule and biotech drugs, for RDF dump see http://wifo5-03.informatik.uni-mannheim.de/drugbank/drugbank_dump.nt.bz2.

⁵ See <http://datahub.io/dataset/fu-berlin-drugbank> for complete list of linked dataset with *DrugBank*.

⁶ <http://dbpedia.org>.

⁷ <http://wifo5-04.informatik.uni-mannheim.de/drugbank/resource/drugs>.

⁸ <http://example.org>.

⁹ *Jamendo* contains a large collection of music related information about artists and recordings, for RDF dump see <http://moustaki.org/resources/jamendo-rdf.tar.gz>.

against `mo:biography`. For *Jamendo* we deploy manually one experimental pipeline:

- $M_{Jamendo}^1 = \{m_4\}$, where m_4 is an NLP function that find *locations* in `mo:biography`.

3. From the multi-domain knowledge base *DBpedia* [12] we used the class `AdministrativeRegion` for our experiments. As *DBpedia* is a knowledge base with a large ontology, we build a set of five pipelines of increasing complexity:

- $M_{DBpedia}^1 = \{m_5\}$, where m_5 is an authority conformation function that conforms the *DBpedia* subject authority to the *Example* target subject authority.
- $M_{DBpedia}^2 = m_6 ++ M_{DBpedia}^1$, where m_6 is a dereferencing function that dereferences any `dbpedia-owl:ideology`.
- $M_{DBpedia}^3 = M_{DBpedia}^2 ++ m_7$, where m_7 is an NLP function that finds all named entities in `dbpedia-owl:abstract`.
- $M_{DBpedia}^4 = M_{DBpedia}^3 ++ m_8$, where m_8 is a filter function that filters for abstracts.
- $M_{DBpedia}^5 = M_{DBpedia}^3 ++ m_9$, where m_9 is a predicate conformation function that conforms the source predicate `dbpedia-owl:abstract` to the target predicate of `dcterms:abstract`.

Altogether, we manually generated a set of eight pipelines, which we then applied against their respective datasets. The evaluation protocol was as follows: Let M be one of the manually generated pipelines. We applied M to an input knowledge base K and generated an enriched knowledge base $K' = M(K)$. We then selected a set of resources in K and used the CBD pairs of selected resources and their enriched versions as examples E . E was then given as training data to our algorithm, which learned an enrichment pipeline M . We finally compared the triples in K' (which we used as reference dataset) with the triples in $M(S)$ to compute the precision, recall and F-measure achieved by our approach. All generated pipelines are available at the project web site¹⁰.

All experiments were carried out on a 8-core PC running *OpenJDK* 64-Bit Server 1.6.0_27 on *Ubuntu* 12.04.2 LTS. The processors were 8 Hexa-core *AMD Opteron* 6128 clocked at 2.0 GHz. Unless stated otherwise, each experiment was assigned 6 GB of memory. As termination criteria for our experiments, we used (1) a maximum number of iterations of 10 or (2) an optimal enrichment pipeline found.

6.2 Results

We carried out two sets of experiments to evaluate our refinement based learning algorithm. In the first set of experiments, we tested the effect of the complexity weight ω to the search strategy of our algorithm. The results are presented in Table 1. In the second set of experiments, we test the effect of the number of positive examples $|\mathcal{E}|$ on the generated F-measure. Results are presented in Table 2.

¹⁰ https://github.com/GeoKnow/DEER/tree/master/evaluations/pipeline_learner.

Configuration of the Search Strategy. We ran our approach with varying values of ω to determine the value to use throughout our experiments. This parameter is used for configuring the search strategy in the learning algorithm, in particular the bias towards simple pipelines. As shown in Sect. 4.2, this is achieved by multiplying ω with the node complexity and subtracting this as a penalty from the node fitness. To configure ω , we used the first pipeline $M_{DrugBank}^1$. The results suggest that setting ω to 0.75 leads to the best results in this particular experiment. We thus adopted this value for the other studies.

Table 1. Test of the effect of ω on the learning process using the *Drugbank* dataset, where $|\mathcal{E}| = 1$, M is the manually created pipeline, $|M|$ is the complexity of M , M' is the pipeline generated by our algorithm, and I_n is the number of iterations of the algorithm.

ω	$ M $	$ M' $	$ \tau $	I_n	$P(M')$	$R(M')$	$F(M')$
0	3	1	61	10	1.0	0.99	0.99
0.25	3	1	61	10	1.0	0.99	0.99
0.50	3	1	61	10	1.0	0.99	0.99
0.75	3	3	25	4	1.0	1.0	1.0
1.0	3	1	61	10	1.0	0.99	0.99

Effect of Positive Examples. We measured the F-measure achieved by our approach on the datasets at hand. The results shown in Table 2 suggest that when faced with data as regular as that found in the datasets Drugbank, DBpedia and Jamendo, our approach really only needs a single example to be able to reconstruct the enrichment pipeline that was used. This result is particularly interesting, because we do not always generate the manually created reference pipeline described in the previous subsection. In many cases, our approach detects a different way to generate the same results. In most cases (71.4%) the pipeline it learns is actually shorter than the manually created pipeline. However, in some cases (4.7%) our algorithm generated a longer pipeline to emulate the manual configuration. As an example, in case of $M_{Jamendo}^1$ the manual configuration was just one enrichment function, i.e., NLP-based enrichment to find all *locations* in `mo:biography`. Our algorithm learns this single manually configured enrichment as (1) an NLP enrichment function that extracts all named entities types and then (2) a filter enrichment function that filters all non-location triples. Our results also suggest that our approach scales when using a small number of positive example as on average the learning time for one positive example was around 48s.

7 Related Work

Linked Data enrichment is an important topic for all applications that rely on a large number of knowledge bases and necessitate a unified view on this data,

Table 2. Test of the effect of increasing number of positive examples in the learning process. For this experiment we set $\omega = 0.75$. M is the manually created pipeline, $|M|$ is the size of M , $T_{M(KB)}$ is the time for applying M to the entire dataset, M' is the pipeline generated by our algorithm, T_l is the learning time, $|\tau|$ is the size of the refinement tree τ , I_n is the number of iterations performed by the algorithm, and all times are in minutes.

M	$ \mathcal{E} $	$ M $	$T_{M(KB)}$	$ M' $	$T_{M'(KB)}$	T_l	$ \tau $	I_n	$P(M')$	$R(M')$	$F(M')$
$M_{DBpedia}^1$	1	1	0.2	1	1.6	1.3	7	1	1.0	1.0	1.0
	2	1	0.2	1	1.8	1.3	7	1	1.0	1.0	1.0
$M_{DBpedia}^2$	1	2	23.3	1	0.1	0.2	7	1	1.0	0.99	0.99
	2	2	15	2	17	0.3	55	9	0.99	1.0	0.99
$M_{DBpedia}^3$	1	3	14.7	3	15.2	6.1	55	9	1.0	0.99	0.99
	2	3	15	2	15.1	0.1	55	9	0.99	0.99	0.99
$M_{DBpedia}^4$	1	4	0.4	2	0.1	0.7	13	2	0.99	0.99	0.99
	2	4	0.6	2	0.3	0.9	13	2	0.99	1.0	0.99
$M_{DBpedia}^5$	1	5	22	2	0.1	0.7	13	2	1.0	1.0	1.0
	2	5	25.5	2	0.2	0.9	13	2	1.0	1.0	1.0
$M_{DrugBank}^1$	1	2	3.5	1	4.1	0.1	61	10	0.99	0.99	0.99
	2	2	3.6	1	3.4	0.1	61	10	0.99	0.99	0.99
$M_{DrugBank}^2$	1	3	25.2	1	0.1	0.1	61	10	1.0	0.99	0.99
	2	3	22.8	1	0.1	0.1	61	10	1.0	0.99	0.99
$M_{Jamendo}^1$	1	1	10.9	2	10.6	0.1	13	2	0.99	0.99	0.99
	2	1	10.4	2	10.4	0.1	7	1	0.99	0.99	0.99

e.g., Question Answering frameworks [13], Linked Education [6] and all forms of semantic mashups [9]. In recent work, several challenges and requirements to Linked Data consumption and integration have been pointed out [14]. For example, the R2R framework [2] addresses those by enabling the publish of mappings across knowledge bases that allow to map classes and defined the transformation of property values. While this framework supports a large number of transformations, it does not allow the automatic discovery of possible transformations. The Linked Data Integration Framework (LDIF) [21], whose goal is to support the integration of RDF data, builds upon R2R mappings and technologies such as SILK [10] and LDSpider¹¹. The concept behind the framework is to enable users to create periodic integration jobs via simple XML configurations. Still these configurations have to be created manually. The same drawback holds for the Semantic Web Pipes¹² [20], which follows the idea of Yahoo Pipes¹³ to enable the integration of data in formats such as RDF and XML. By using Semantic Web Pipes, users can efficiently create semantic mashups by using a number of

¹¹ <http://code.google.com/p/ldspider/>.

¹² <http://pipes.deri.org/>.

¹³ <http://pipes.yahoo.com/pipes/>.

operators (such as `getRDF`, `getXML`, etc.) and connect these manually within a simple interface. KnoFuss [19] addresses data integration from the point of view of link discovery. It begins by detecting URIs that stand for the same real-world entity and either merging them together or linking them via `owl:sameAs`. In addition, it allows to monitor the interaction between instance and dataset matching (which is similar to ontology matching [7]). Fluid Operations' Information Workbench¹⁴ allows to search through, manipulate and integrate datasets for purposes such as business intelligence. [5] describes a framework for semantic enrichment, ranking and integration of web videos, and [1] presents semantic enrichment framework of *Twitter* posts. Finally, [8] tackles the linked data enrichment problem for sensor data via an approach that sees enrichment as a process driven by situations of interest. To the best of our knowledge, the work we presented in this paper is the first generic approach tailored towards learning enrichment pipelines of Linked Data given a set of atomic enrichment functions.

8 Conclusions and Future Work

In this paper, we presented an approach for learning enrichment pipelines based on a refinement operator. To the best of our knowledge, this is the first approach for learning RDF based enrichment pipelines and could open up a new research area. We also presented means to self-configure atomic enrichment pipelines so as to find means to enrich datasets according to examples provided by an end user. We showed that our approach can easily reconstruct manually created enrichment pipelines, especially when given a prototypical example and when faced with regular datasets. Obviously, this does not mean that our approach will always achieve such high F-measures. What our results suggest is primarily that if a human uses an enrichment tool to enrich his/her dataset manually, then our approach can reconstruct the pipeline. This seems to hold even for relatively complex pipelines.

Although we achieved reasonable results in terms of scalability, we plan to further improve time efficiency by parallelising the algorithm on several CPUs as well as load balancing. The framework underlying this study supports directed acyclic graphs as enrichment specifications by allowing to split and merge datasets. In future work, we will thus extend our operator to deal with graphs in addition to sequences. Moreover, we will look at pro-active enrichment strategies as well as active learning.

References

1. Abel, F., Gao, Q., Houben, G.-J., Tao, K.: Semantic enrichment of twitter posts for user profile construction on the social web. In: Antoniou, G., Grobelnik, M., Simperl, E., Parsia, B., Plexousakis, D., De Leenheer, P., Pan, J. (eds.) ESWC 2011, Part II. LNCS, vol. 6644, pp. 375–389. Springer, Heidelberg (2011)

¹⁴ <http://www.fluidops.com/information-workbench/>.

2. Bizer, C., Schultz, A.: The R2R framework: Publishing and discovering mappings on the web. In: Proceedings of the COLD (2010)
3. Blumer, A., Ehrenfeucht, A., Haussler, D., Warmuth, M.K.: Occam's razor. *Inf. Process. Lett.* **24**(6), 377–380 (1987)
4. Bühmann, L., Lehmann, J.: Pattern based knowledge base enrichment. In: Alani, H., Kagal, L., Fokoue, A., Groth, P., Biemann, C., Parreira, J.X., Aroyo, L., Noy, N., Welty, C., Janowicz, K. (eds.) ISWC 2013, Part I. LNCS, vol. 8218, pp. 33–48. Springer, Heidelberg (2013)
5. Choudhury, S., Breslin, J.G., Passant, A.: Enrichment and ranking of the youtube tag space and integration with the linked data cloud. Springer, Berlin (2009)
6. Dietze, S., Sanchez-Alonso, S., Ebner, H., Yu, H.Q., Giordano, D., Marenzi, I., Nunes, B.P.: Interlinking educational resources and the web of data: A survey of challenges and approaches. *Progr. Electron. Libr. Inform. Syst.* **47**(1), 60–91 (2013)
7. Euzenat, J., Shvaiko, P.: *Ontology Matching*. Springer, Heidelberg (DE) (2007)
8. Hasan, S., Curry, E., Banduk, M., O'Riain, S.: Toward situation awareness for the semantic sensor web: Complex event processing with dynamic linked data enrichment. *Semantic Sensor Networks*, p. 60 (2011)
9. Hoang, H.H., Cung, T.N.-P., Truong, D.K., Hwang, D., Jung, J.J.: Semantic information integration with linked data mashups approaches. *Int. J. Distrib. Sens. Netw.* **2012**, 12 (2014)
10. Isele, R., Bizer, C.: Learning linkage rules using genetic programming. In: Sixth International Ontology Matching Workshop (2011)
11. Lehmann, J., Hitzler, P.: Concept learning in description logics using refinement operators. *Mach. Learn. J.* **78**(1–2), 203–250 (2010)
12. Lehmann, J., Isele, R., Jakob, M., Jentzsch, A., Kontokostas, D., Mendes, P.N., Hellmann, S., Morsey, M., van Kleef, P., Auer, S., Bizer, C.: DBpedia—a large-scale, multilingual knowledge base extracted from wikipedia. *Semant. Web J.* (2014)
13. Lopez, V., Unger, C., Cimiano, P., Motta, E.: Evaluating question answering over linked data. *Web Semant. Sci. Serv. Agents World Wide Web* **21**, 3–13 (2013)
14. Millard, I., Glaser, H., Salvadores, M., Shadbolt, N.: Consuming multiple linked data sources: Challenges and experiences. In: COLD Workshop (2010)
15. Ngomo, A.-C.N.: On link discovery using a hybrid approach. *J. Data Semant.* **1**(4) 203–217, (December 2012)
16. Ngomo, A.-C.N., Auer, S., Lehmann, J., Zaveri, A.: Introduction to linked data and its lifecycle on the web. In: Koubarakis, M., Stamou, G., Stoilos, G., Horrocks, I., Kolaitis, P., Lausen, G., Weikum, G. (eds.) Reasoning Web. LNCS, vol. 8714, pp. 1–99. Springer, Heidelberg (2014)
17. Ngonga Ngomo, A.-C., Heino, N., Lyko, K., Speck, R., Kaltenböck, M.: SCMS—semantifying content management systems. In: Aroyo, L., Welty, C., Alani, H., Taylor, J., Bernstein, A., Kagal, L., Noy, N., Blomqvist, E. (eds.) ISWC 2011, Part II. LNCS, vol. 7032, pp. 189–204. Springer, Heidelberg (2011)
18. Ngomo, A.-C.N., Lyko, K.: Unsupervised learning of link specifications: deterministic vs. non-deterministic. In: Proceedings of the Ontology Matching Workshop (2013)
19. Nikolov, A., Uren, V., Motta, E., de Roeck, A.: Overcoming schema heterogeneity between linked semantic repositories to improve coreference resolution. In: Gómez-Pérez, A., Yu, Y., Ding, Y. (eds.) ASWC 2009. LNCS, vol. 5926, pp. 332–346. Springer, Heidelberg (2009)

20. Phuoc, D.L., Polleres, A., Hauswirth, M., Tummarello, G., Morbidoni, C.: Rapid prototyping of semantic mash-ups through semantic web pipes. In: WWW, pp. 581–590 (2009)
21. Schultz, A., Matteini, A., Isele, R., Bizer, C., Becker, C.: LDIF—linked data integration framework. In: COLD (2011)
22. Speck, R., Ngonga Ngomo, A.-C.: Ensemble learning for named entity recognition. In: Mika, P., Tudorache, T., Bernstein, A., Welty, C., Knoblock, C., Vrandečić, D., Groth, P., Noy, N., Janowicz, K., Goble, C. (eds.) ISWC 2014, Part I. LNCS, vol. 8796, pp. 519–534. Springer, Heidelberg (2014)