

A Principled Approach to Bridging the Gap between Graph Data and their Schemas

Marcelo Arenas^{1,2}, Gonzalo Díaz¹, Achille Fokoue³,
Anastasios Kementsietsidis³, Kavitha Srinivas³

¹Pontificia Universidad Católica de Chile

²University of Oxford

marenas@ing.puc.cl, gdiacz@uc.cl

³IBM T.J. Watson Research Center

{achille, ksrinivs}@us.ibm.com, tasosk@ca.ibm.com

ABSTRACT

Although RDF graph data often come with an associated schema, recent studies have proven that real RDF data rarely conform to their perceived schemas. Since a number of data management decisions, including storage layouts, indexing, and efficient query processing, use schemas to guide the decision making, it is imperative to have an accurate description of the *structuredness* of the data at hand (how well the data conform to the schema).

In this paper, we have approached the study of the structuredness of an RDF graph in a principled way: we propose a framework for specifying structuredness functions, which gauge the degree to which an RDF graph conforms to a schema. In particular, we first define a formal language for specifying structuredness functions with expressions we call rules. This language allows a user to state a rule to which an RDF graph may fully or partially conform. Then we consider the issue of discovering a refinement of a sort (type) by partitioning the dataset into subsets whose structuredness is over a specified threshold. In particular, we prove that the natural decision problem associated to this refinement problem is NP-complete, and we provide a natural translation of this problem into Integer Linear Programming (ILP). Finally, we test this ILP solution with three real world datasets and three different and intuitive rules, which gauge the structuredness in different ways. We show that the rules give meaningful refinements of the datasets, showing that our language can be a powerful tool for understanding the structure of RDF data, and we show that the ILP solution is practical for a large fraction of existing data.

1. INTRODUCTION

If there is one thing that is clear from analyzing real RDF data, it is that the data rarely conform to their assumed schema [5]. One example is the popular type of DBpedia persons (in this paper, we will use the term *sort* as a synonym of *type*), which includes all people with an entry in Wikipedia. According to the sort definition, each person in DBpedia can have 8 properties, namely, a *name*, a *givenName*, a *surName*, a *birthDate*, a *birthPlace*, a *deathDate*, a *deathPlace*, and a *description*. There are 790,703 people listed in

DBpedia, and while we expect that a large portion of them are alive (they do not have a death date or death place) we do expect that we know at least when and where these people were born. The statistics however are very revealing: Only 420,242 people have a birthdate, only 323,368 have a birthplace, and for only 241,156 do we have both. There are $\sim 40,000$ people for whom we do not even know their last name. As for death places and death dates, we only know those for 90,246 and 173,507 people, respectively.

There is actually nothing wrong with the DBpedia person data. The data reflect the simple fact that the information we have about any domain of discourse (in this case people) is inherently incomplete. But while this is the nature of things in practice, sorts in general go against this trend since they favor uniformity, i.e., they require that the data *tightly* conform to the provided sorts. In our example, this means that we expect to have all 8 properties for every DBpedia person. So the question one needs to address is how to bridge the gap between these two worlds, the sorts and the respective data. In our previous work [5], we considered sorts as being the unequivocal ground truth, and we devised methods to make the data fit these sorts. Here, we consider a complementary approach in which we accept the data for what they are and ask ourselves whether *we can refine the schema to better fit our data* (more precisely, we will seek a *sort refinement* of our data).

Many challenges need to be addressed to achieve our goal. First, we need to define formally what it means for a dataset to fit a particular sort. Our own past work has only introduced one such fitness metric, called *coherence*, but that does not clearly cover every possible interpretation of fitness. In this work, we propose a new set of alternative and complementary fitness metrics between a dataset and a sort, and we also introduce a rule language through which users can define their own metrics.

Second, for a given RDF graph D and a fitness metric σ , we study the problem of determining whether there exists a sort refinement \mathcal{T} of D with a fitness value above a given threshold θ that contain at most k implicit sorts, and we show that the problem is NP-complete. In spite of this negative result, we present several techniques enabling us to solve this problem in practice on real datasets as illustrated in our experimental evaluation section. Our first attack on the problem is to reduce the size of the input we have to work with. Given that typical real graph datasets involve millions of instances, even for a single sort, scalability is definitely a concern. We address this challenge by introducing views of our input data that still maintain all the properties of the data in terms of their fitness characteristics, yet they occupy substantially less space. Using said view, given any fitness metric expressed as a rule r in our language, we formulate the previously defined problem as an Integer Linear Programming (ILP) problem instance. Although ILP is

This work is licensed under the Creative Commons Attribution-NonCommercial-NoDerivs 3.0 Unported License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc-nd/3.0/>. Obtain permission prior to any use beyond those covered by the license. Contact copyright holder by emailing info@vldb.org. Articles from this volume were invited to present their results at the 40th International Conference on Very Large Data Bases, September 1st - 5th 2014, Hangzhou, China. *Proceedings of the VLDB Endowment*, Vol. 7, No. 8. Copyright 2014 VLDB Endowment 2150-8097/14/04.

also known to be NP-hard in the worst case, in practice, highly optimized commercial solvers (e.g. IBM ILOG CPLEX) exist to efficiently solve our formulation of the sort refinement problem (see experimental evaluation for more details). In particular, we study two complementary formulations of our problem: In the first alternative, we allow the user to specify a desired fitting value θ' , and we compute a smallest set of implicit sorts, expressed as a partition $\{D_1, D_2, \dots, D_n\}$ of the input dataset D , such that the fitness of each D_i is larger than or equal to θ' . In the second alternative, we allow the user to specify the desired number k of implicit sorts, and we compute a set of k implicit sorts such that the minimum fitness across all implicit sorts is maximal amongst all possible decompositions of the sort that involve k implicit sorts. Both our alternatives are motivated by practical scenarios. In the former alternative, we allow a user to define a desirable fitness and we try to compute a sort refinement with the indicated fitness. However, in other settings, the user might want to specify a maximum number of sorts to which the data should be decomposed and let the system figure out the best possible sort and data decomposition.

A clear indication of the practical value of this work can be found in the experimental section, where we use different rules over real datasets and not only provide useful insights about the data themselves, but also automatically discover sort refinements that, in hindsight, seem natural, intuitive and easy to understand. We explore the correlations between alternative rules (and sort refinements) over the same data and show that the use of *multiple* such rules is important to fully understand the nature of data. Finally, we study the scalability of the ILP-based solution on a sample of explicit sorts extracted from the knowledge base YAGO, showing that it is practical for all but a small minority of sorts.

Our key contributions in this paper are fivefold: (1) we propose a framework for specifying structuredness functions to measure the degree to which an RDF graph conforms to a schema; (2) we study the problem of discovering a refinement of a sort by partitioning the dataset into subsets whose structuredness is greater than a given threshold and show that the decision problem associated with this sort refinement problem is NP-complete; (3) we provide a natural translation of an instance of the sort refinement problem into an ILP problem instance; (4) we successfully test our ILP approach on two real world datasets and three different structuredness functions; (5) we study the scalability of our solution on a sample of sorts extracted from the knowledge base YAGO, showing that the solution is practical in a large majority of cases.

The remainder of the paper is organized as follows. Section 2 presents a brief introduction to RDF and sample structuredness functions, while the syntax and the formal semantics of the language for specifying structuredness functions is defined in Section 3. Section 4 introduces the key concepts of signatures and sort refinements. Section 5 presents the main complexity result of the sort refinement problem, while Section 6 describes the formulation of the problem as an ILP problem. In Section 7, we present our experimental evaluation on three real world datasets. Finally, we review related work in Section 8, and conclude in Section 9.

An extended version of this paper (with proofs of the theoretical results) is available at <http://arxiv.org/abs/1308.5703>.

2. PRELIMINARIES

2.1 A schema-oriented graph representation

We assume two countably infinite disjoint sets U and L of URIs, Literals, respectively. An RDF triple is a tuple $(s, p, o) \in U \times U \times (U \cup L)$, and an RDF graph is a finite set of RDF triples. Given an

RDF graph D , we define the sets of subjects and properties mentioned in D , respectively denoted by $S(D)$ and $P(D)$, as:

$$\begin{aligned} S(D) &= \{s \in U \mid \exists p \exists o \text{ s.t. } (s, p, o) \in D\}, \\ P(D) &= \{p \in U \mid \exists s \exists o \text{ s.t. } (s, p, o) \in D\}. \end{aligned}$$

Given an RDF graph D and $s, p \in U$, we say that s has property p in D if there exists $o \in U$ such that $(s, p, o) \in D$.

A natural way of storing RDF data in a relational table, known as the horizontal database [11], consists in defining only one relational table in which each row represents a subject and there is a column for every property. With this in mind, given an RDF graph D , we define an $|S(D)| \times |P(D)|$ matrix $M(D)$ (or just M if D is clear from the context) as follows: for every $s \in S(D)$ and $p \in P(D)$,

$$M(D)_{sp} = \begin{cases} 1 & \text{if } s \text{ has property } p \text{ in } D \\ 0 & \text{otherwise.} \end{cases}$$

It is important to notice that $M(D)$ corresponds to a view of the RDF graph D in which we have discarded a large amount of information, and only retained information about the structure of the properties in D . Thus, in what follows we refer to $M(D)$ as the *property-structure view of D* .

In an RDF graph, to indicate that a subject s is of a specific sort t (like person or country), the following triple must be present: (s, type, t) , where the constant $\text{type} = \text{http://www.w3.org/1999/02/22-rdf-syntax-ns\#type}$ (note that $\text{type} \in U$).

Given a URI t , we define the following RDF subgraph $D_t \subseteq D$: $D_t = \{(s, p, o) \in D \mid (s, \text{type}, t) \in D\}$. This subgraph consists of all triples whose subject s is explicitly declared to be of sort t in D . With this subgraph D_t , we can mention its set of subjects, $S(D_t)$, which is also the set of subjects of sort t in D , and its set of properties $P(D_t)$, which is the set of properties set by some subject of sort t . We will use the term *sort* to refer to the constant t , the RDF subgraph D_t , and sometimes the set $S(D_t)$.

2.2 Sample structuredness functions

As there are many alternative ways to define the fitness, or *structuredness*, of a dataset with respect to a schema, it is convenient to define structuredness initially in the most general way. As such, we define a *structuredness function* σ to be any function which assigns to every RDF graph D a rational number $\sigma(D)$, such that $0 \leq \sigma(D) \leq 1$. Within the context of our framework a structuredness function will only produce rational numbers. In what follows, we offer concrete examples of structuredness functions which gauge the structuredness of RDF graphs in very different ways.

2.2.1 The coverage function

Duan et. al. defined the COVERAGE function [5] σ_{Cov} to test the fitness of graph data to their respective schemas. The metric was used to illustrate that though graph benchmark data are very relational-like and have high fitness (values of $\sigma_{\text{Cov}}(D)$ close to 1) with respect to their sort, real graph data are fairly unstructured and have low fitness ($\sigma_{\text{Cov}}(D)$ less than 0.5). Using the property-structure view $M(D)$, the coverage metric of [5] can be defined as: $\sigma_{\text{Cov}}(D) = (\sum_{sp} M(D)_{sp}) / (|S(D)| |P(D)|)$. Intuitively, the metric favors *conformity*, i.e., if one subject has a property p , then the other subjects of the same sort are expected to also have this property. Therefore, the metric is not forgiving when it comes to *missing* properties. To illustrate, consider an RDF graph D_1 consisting of N triples: (s_i, p, o) for $i = 1, \dots, N$ (i.e. all N subjects have the same property p). The matrix $M(D_1)$ for D_1 is shown in Figure 1a. For this dataset, $\sigma_{\text{Cov}}(D_1) = 1$. If we insert a new triple (s_1, q, o) for some property $q \neq p$, resulting dataset $D_2 = D_1 \cup \{(s_1, q, o)\}$ whose matrix is shown in Figure 1b. Then, the structuredness of

$$\begin{array}{ccc}
\begin{array}{c} p \\ s_1 \begin{pmatrix} 1 \\ 1 \\ \vdots \\ 1 \end{pmatrix} \\ s_2 \\ \vdots \\ s_N \end{array} & \begin{array}{cc} p & q \\ s_1 \begin{pmatrix} 1 & 1 \\ 1 & 0 \\ \vdots & \vdots \\ 1 & 0 \end{pmatrix} \\ s_2 \\ \vdots \\ s_N \end{array} & \begin{array}{cccc} p_1 & p_2 & \cdots & p_N \\ s_1 \begin{pmatrix} 1 & 0 & \cdots & 0 \\ 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 1 \end{pmatrix} \\ s_2 \\ \vdots \\ s_N \end{array} \\
\text{(a) } M(D_1) & \text{(b) } M(D_2) & \text{(c) } M(D_3)
\end{array}$$

Figure 1: Sample matrixes for datasets D_1 , D_2 and D_3

$\sigma_{\text{Cov}}(D_2) \approx 0.5$ (for a large value of N). The addition of the single triple generates a new dataset D_2 in which most of the existing subjects are *missing* property q , an indication of unstructuredness.

2.2.2 The similarity function

The previous behavior motivates the introduction of a structuredness function that is less sensitive to missing properties. We define the σ_{Sim} structuredness function as the probability that, given two randomly selected subjects s and s' and a random property p such that s has property p in D , s' also has property p in D .

To define the function formally, let $\varphi_1^{\text{Sim}}(s, s', p)$ denote the statement “ $s \neq s'$ and s has property p in D ” and let $\varphi_2^{\text{Sim}}(s', p)$ denote “ s' has property p in D ”. Next, we define a set of total cases $\text{total}(\varphi_1^{\text{Sim}}, D) = \{(s, s', p) \in S(D) \times S(D) \times P(D) \mid \varphi_1^{\text{Sim}}(s, s', p) \text{ holds}\}$, and a set of favorable cases $\text{total}(\varphi_1^{\text{Sim}} \wedge \varphi_2^{\text{Sim}}, D) = \{(s, s', p) \in S(D) \times S(D) \times P(D) \mid \varphi_1^{\text{Sim}}(s, s', p) \wedge \varphi_2^{\text{Sim}}(s', p) \text{ holds}\}$. Finally, define:

$$\sigma_{\text{Sim}}(D) = \frac{|\text{total}(\varphi_1^{\text{Sim}} \wedge \varphi_2^{\text{Sim}}, D)|}{|\text{total}(\varphi_1^{\text{Sim}}, D)|}.$$

Going back to the example in Figure 1, notice that $\sigma_{\text{Sim}}(D_1) = 1$ but also $\sigma_{\text{Sim}}(D_2)$ is still approx. equal to 1 (for large N). Unlike σ_{Cov} , function σ_{Sim} allows certain subjects to have *exotic* properties that either no other subject has, or only a small fraction of other subjects have (while maintaining high values for σ_{Sim}). As another example, consider the RDF graph D_3 in Figure 1c where every subject s_i has only one property p_i , and no two subjects have the same property. This dataset is intuitively very unstructured. Indeed, $\sigma_{\text{Sim}}(D_3) = 0$ while $\sigma_{\text{Cov}}(D_3) \approx 0$ (for a large value of N).

2.2.3 The dependency functions

It is also of interest to understand the correlation between different properties in an RDF graph D . Let $\mathbf{p}_1, \mathbf{p}_2 \in P(D)$ be two fixed properties we are interested in. Define the $\sigma_{\text{Dep}}[\mathbf{p}_1, \mathbf{p}_2]$ function as the probability that, given a random subject $s \in S(D)$ such that s has \mathbf{p}_1 , s also has \mathbf{p}_2 . In the same way as before, we can define a set of total cases and a set of favorable cases, and we define the value of $\sigma_{\text{Dep}}[\mathbf{p}_1, \mathbf{p}_2]$ to be the ratio of the sizes of both sets.

A closely related structuredness function is the symmetric version of $\sigma_{\text{Dep}}[\mathbf{p}_1, \mathbf{p}_2]$, which we call $\sigma_{\text{SymDep}}[\mathbf{p}_1, \mathbf{p}_2]$. It is defined as the probability that, given a random subject $s \in S(D)$ such that s has p_1 or s has p_2 , s has both.

3. A LANGUAGE FOR DEFINING STRUCTUREDNESS MEASURES

We have already shown in Section 2.2 some intuitive structuredness measures that give very different results when applied to the same RDF graphs. As many more natural structuredness functions exist, we do not intend to list all of them in this article, but instead our goal is to introduce a general framework to allow users to define their own custom structuredness measures in a simple way. To this end, we introduce in this section a language for describing

such measures. This language has a simple syntax and a formal semantics, which make it appropriate for a formal study, and it is expressive enough to represent many natural structuredness functions, like the ones presented in Section 2.2. In general, starting from the matrix $M(D)$ of a dataset D , our language can construct statements that involve (i) the contents of the matrix (the cells of the matrix with 0 or 1 values); (ii) the *indices* of the matrix, that correspond to the subjects and properties of the dataset; and (iii) Boolean combinations of these basic building components.

3.1 Syntax of the language

To define the syntax of the language, we need to introduce some terminology. From now on, assume that \mathbf{V} is an infinite set of variables disjoint from \mathbf{U} . We usually use c, c_1, c_2, \dots to denote the variables in \mathbf{V} , as each one of these variables is used as a pointer to a *cell* (or position) in a matrix associated with an RDF graph. Moreover, assume that 0, 1 do not belong to $(\mathbf{U} \cup \mathbf{V})$. Then the set of terms in the language is defined as follows: (i) 0, 1, every $u \in \mathbf{U}$ and every $c \in \mathbf{V}$ is a term, and (ii) if $c \in \mathbf{V}$, then $\text{val}(c)$, $\text{subj}(c)$ and $\text{prop}(c)$ are terms.

If c is a variable pointing to a particular cell in a matrix, then $\text{val}(c)$ represents the value of the cell, which must be either 0 or 1, $\text{subj}(c)$ denotes the row of the cell, which must be the subject of a triple in D , and $\text{prop}(c)$ denotes the column of the cell, which must be the property of a triple in D . Moreover, the set of formulas in the language is recursively defined as follows:

- If $c \in \mathbf{V}$ and $u \in \mathbf{U}$, then $\text{val}(c) = 0$, $\text{val}(c) = 1$, $\text{prop}(c) = u$ and $\text{subj}(c) = u$ are formulas.
- If $c_1, c_2 \in \mathbf{V}$, then $c_1 = c_2$, $\text{val}(c_1) = \text{val}(c_2)$, $\text{prop}(c_1) = \text{prop}(c_2)$ and $\text{subj}(c_1) = \text{subj}(c_2)$ are formulas.
- If φ_1 and φ_2 are formulas, then $(\neg\varphi_1)$, $(\varphi_1 \wedge \varphi_2)$, $(\varphi_1 \vee \varphi_2)$ are formulas.

If φ is a formula, then $\text{var}(\varphi)$ is the set consisting of all the variables mentioned in φ . With this notation, we can finally define the syntax of the rules in the language, which are used to define structuredness functions. Formally, if φ_1, φ_2 are formulas such that $\text{var}(\varphi_2) \subseteq \text{var}(\varphi_1)$, then the following is a rule:

$$\varphi_1 \mapsto \varphi_2. \quad (1)$$

3.2 Semantics of the language

To define how rules of the form (1) are evaluated, we need to define the notion of satisfaction of a formula. In the rest of this section, assume that D is an RDF graph and M is the $|S(D)| \times |P(D)|$ matrix associated with D . A partial function $\rho : \mathbf{V} \rightarrow S(D) \times P(D)$ is said to be a variable assignment for M , whose domain is denoted by $\text{dom}(\rho)$. Moreover, given a formula φ and a variable assignment ρ for M such that $\text{var}(\varphi) \subseteq \text{dom}(\rho)$, pair (M, ρ) is said to satisfy φ , denoted by $(M, \rho) \models \varphi$, if:

- φ is the formula $\text{val}(c) = i$, where $i = 0$ or $i = 1$, $\rho(c) = (s, p)$ and $M_{sp} = i$.
- φ is the formula $\text{subj}(c) = u$, where $u \in \mathbf{U}$, and $\rho(c) = (u, p)$.
- φ is the formula $\text{prop}(c) = u$, where $u \in \mathbf{U}$, and $\rho(c) = (s, u)$.
- φ is the formula $c_1 = c_2$, and $\rho(c_1) = \rho(c_2)$.
- φ is the formula $\text{val}(c_1) = \text{val}(c_2)$, $\rho(c_1) = (s_1, p_1)$, $\rho(c_2) = (s_2, p_2)$ and $M_{s_1 p_1} = M_{s_2 p_2}$.
- φ is the formula $\text{subj}(c_1) = \text{subj}(c_2)$, $\rho(c_1) = (s_1, p_1)$, $\rho(c_2) = (s_2, p_2)$ and $s_1 = s_2$.
- φ is the formula $\text{prop}(c_1) = \text{prop}(c_2)$, $\rho(c_1) = (s_1, p_1)$, $\rho(c_2) = (s_2, p_2)$ and $p_1 = p_2$.
- φ is the formula $(\neg\varphi_1)$ and $(M, \rho) \models \varphi_1$ does not hold.
- φ is the formula $(\varphi_1 \wedge \varphi_2)$, $(M, \rho) \models \varphi_1$ and $(M, \rho) \models \varphi_2$.
- φ is the formula $(\varphi_1 \vee \varphi_2)$, and $(M, \rho) \models \varphi_1$ or $(M, \rho) \models \varphi_2$.

Moreover, the set of satisfying assignments for a formula φ w.r.t. M , denoted by $\text{total}(\varphi, M)$, is defined as follows:

$$\{\rho \mid \rho \text{ is a variable assignment for } M \text{ such that} \\ \text{dom}(\rho) = \text{var}(\varphi) \text{ and } (M, \rho) \models \varphi\}.$$

We now have the necessary ingredients to define the semantics of rules. Assume that r is the rule (1). Then the structuredness function given by rule r is defined as a function σ_r that assigns to every matrix M the value

$$\sigma_r(M) = \frac{|\text{total}(\varphi_1 \wedge \varphi_2, M)|}{|\text{total}(\varphi_1, M)|}$$

if $|\text{total}(\varphi_1, M)| > 0$, and 1 otherwise (notice that $0 \leq \sigma_r(M) \leq 1$, as we assume that $\text{var}(\varphi_2) \subseteq \text{var}(\varphi_1)$). Thus, $\sigma_r(M)$ is defined as the probability that a variable assignment ρ satisfies φ_2 given that ρ satisfies φ_1 .

The functions presented in Section 2.2 can be expressed in our language as follows. The σ_{COV} structuredness measure can be expressed with the rule $c = c \mapsto \text{val}(c) = 1$. In this case, given a matrix M , $\text{total}(c = c, M)$ is the set of all cells of M and $\text{total}(c = c \wedge \text{val}(c) = 1, M)$ is the set of all cells of M containing a value 1 (which is represented by the condition $\text{val}(c) = 1$).

In some cases, it is desirable to compute a structuredness functions without considering some predicate (or set of predicates), which can be easily done in our language. For instance, a modified σ_{COV} structuredness measure which ignores a specific column called \mathbf{p} is defined by the following rule:

$$c = c \wedge \neg(\text{prop}(c) = \mathbf{p}) \mapsto \text{val}(c) = 1.$$

The σ_{SIM} structuredness measure can be expressed with the rule

$$\neg(c_1 = c_2) \wedge \text{prop}(c_1) = \text{prop}(c_2) \wedge \text{val}(c_1) = 1 \mapsto \\ \text{val}(c_2) = 1,$$

where $\neg(c_1 = c_2)$ considers two variables c_1 and c_2 that should point to different cells, and $\text{prop}(c_1) = \text{prop}(c_2)$ requires that the two variables range over the same property column, say property p . Taken together, the first two formulas iterate over all pairs of subjects for each property p . The last part of the formula $\text{val}(c_1) = 1$ requires that the value of the first cell be 1, i.e., the first subject actually has property p . If the consequence formula is satisfied, then the rule considers the cases where the value of the second cell is also 1, which translates to the second subject also having property p . Notice that this is exactly the definition of the function σ_{SIM} .

Finally, for fixed $\mathbf{p}_1, \mathbf{p}_2 \in \mathbf{U}$, the dependency measures. $\sigma_{\text{DEP}}[\mathbf{p}_1, \mathbf{p}_2]$ can be expressed with the rule

$$\text{subj}(c_1) = \text{subj}(c_2) \wedge \text{prop}(c_1) = \mathbf{p}_1 \wedge \text{prop}(c_2) = \mathbf{p}_2 \\ \wedge \text{val}(c_1) = 1 \mapsto \text{val}(c_2) = 1,$$

while $\sigma_{\text{SYMDEP}}[\mathbf{p}_1, \mathbf{p}_2]$ can be expressed with the rule

$$\text{subj}(c_1) = \text{subj}(c_2) \wedge \text{prop}(c_1) = \mathbf{p}_1 \wedge \text{prop}(c_2) = \mathbf{p}_2 \\ \wedge (\text{val}(c_1) = 1 \vee \text{val}(c_2) = 1) \\ \mapsto \text{val}(c_1) = 1 \wedge \text{val}(c_2) = 1.$$

A variant of the dependency rule uses disjunction in the consequent and corresponds to the probability that a random subject s satisfies that: if s has \mathbf{p}_1 , then s also has \mathbf{p}_2 :

$$\text{subj}(c_1) = \text{subj}(c_2) \wedge \text{prop}(c_1) = \mathbf{p}_1 \wedge \text{prop}(c_2) = \mathbf{p}_2 \\ \mapsto \text{val}(c_1) = 0 \vee \text{val}(c_2) = 1.$$

4. SORT REFINEMENTS AND SIGNATURES

We can use the language from the previous section to define a structuredness measure for a dataset. If the value of the measure for the dataset is high, say 0.9 or even 1.0, then this is probably a positive indication for the current state of the data, and the measure computation can be repeated at a later stage, as the data change. Of most interest, however, is what happens if this value is relatively low, say, 0.5 or even 0.1? Then, we know that the dataset does not have the desired characteristic, as expressed by the measure, and the question is whether there is anything we can do about it. In particular, it is interesting to investigate if there is a way to convert the existing dataset into one whose measure is high.

In our previous work, we tried to change the data themselves to fit the measure, by introducing new triples or removing existing ones. The approach made sense in the context of benchmarking for which it was introduced, but in any practical setting one does not want to contaminate their data with dummy triples, or even worse lose real data by deleting triples just so that the data fit some desired measure. So a more pragmatic solution is to leave data as they are and try to figure out whether we can refine the sort that the data is supposed to fit, in an effort to improve structuredness.

To this end, we consider the situation in which one wishes to partition the dataset into k implicit sorts such that each implicit sort has a high structuredness (as defined by a rule in our language). For a certain subject $s_0 \in S(D)$ we are interested in keeping all triples of the form (s_0, p, o) (for some $p, o \in \mathbf{U}$) together. We refer to these triples collectively as the *entity* s_0 .

We define an *entity preserving* partition of size k of an RDF graph D to be a set of non-empty RDF graphs $\{D_1, \dots, D_k\}$ where (i) $D_i \subseteq D$ for every $i \in \{1, \dots, k\}$, (ii) $D_i \cap D_j = \emptyset$ for every $i, j \in \{1, \dots, k\}$ such that $i \neq j$, (iii) $\bigcup_{i=1}^k D_i = D$, and (iv) for all $s, p_1, p_2, o_1, o_2 \in \mathbf{U}$, we have that:

$$\text{if } (s, p_1, o_1) \in D_i \text{ and } (s, p_2, o_2) \in D_j, \text{ then } i = j.$$

While the first three items specify a partition of D , the last item indicates that for every entity s , we include the full entity in a sort.

A second consideration we shall make is concerned with the grouping of subjects which have the same properties in D . For this, we define the concept of signature:

DEFINITION 4.1. Given an RDF graph D and a subject $s \in S(D)$, the signature of s in D is a function $\text{sig}(s, D) : P(D) \rightarrow \{0, 1\}$, which assigns to every property $p \in P(D)$ a 1 if s has property p in D , and a 0 otherwise. \square

Similarly, a signature set is the set of all subjects in $S(D)$ which share the same signature. The *size* of a signature set is the number of subjects sharing that signature.

We are now ready to define our main objects of study. For the following definition, let D be a fixed RDF graph and θ be a rational number such that $0 \leq \theta \leq 1$ (θ is required to be a rational number for compatibility with the reduction to the Integer Linear Programming instance).

DEFINITION 4.2. Given a structuredness function σ , a σ -*sort refinement* \mathcal{T} of D with *threshold* θ is an entity preserving partition $\{D_1, \dots, D_n\}$ of D such that:

- i) $\sigma(D_i) \geq \theta$ for $i = 1, \dots, n$, and
- ii) each D_i ($1 \leq i \leq n$) is closed under signatures. That is, for every pair of subjects $s_1, s_2 \in S(D)$, if $\text{sig}(s_1, D) = \text{sig}(s_2, D)$ and $s_1 \in S(D_i)$, then $s_2 \in S(D_i)$. \square




Figure 2: DBpedia Persons has 790,703 subjects, 8 properties and 64 signature sets. Its $\sigma_{Cov} = 0.54$, while $\sigma_{Sim} = 0.77$.




Figure 3: WordNet Nouns has 79,689 subjects, 12 properties and 53 signature sets. Its $\sigma_{Cov} = 0.44$, while $\sigma_{Sim} = 0.93$.

In the rest of this paper, we will refer to the elements of the sort refinement (i.e. the elements of the partition of D) as implicit sorts.

The requirement that each implicit sort be closed under signatures is due to the fact that two subjects with equal signatures are structurally identical, and thus it would not be reasonable to place them in different subsets. This has the added benefit of reducing the difficulty of partitioning the dataset, as the basic units to be moved will be signature sets, and not individual entities.

In what follows we will be concerned with discovering sort refinements in RDF data.

Figures 2 and 3 present a visual representation of an RDF graph’s horizontal table. Every column represents a property and the rows have been grouped into signature sets, in descending order of signature set size. The first 3 signature sets in Figure 2 have been delimited with a dashed line, for clarity. The subsequent signature sets can be visually separated by searching for the change in pattern. The black zones represent data (i.e. non-null values) whereas the white regions represent null cells. The difference between DBpedia Persons (Fig. 2) and WordNet Nouns (Fig. 3) is immediately visible. DBpedia Persons is a relatively unstructured dataset, with only 3 clearly common properties: `name`, `givenName`, and `surName` (these three attributes are usually extractable directly from the URL of a Wikipedia article). On the other hand, WordNet Nouns has 5 clearly common properties, and the rest of the properties are relatively rare (very few subjects have them). The values of the structuredness functions show how they differ in judging the structuredness of an RDF graph.

We shall use this visual representation of the horizontal table of an RDF graph to present the results of the experimental settings. In this context, a sort refinement corresponds loosely to a partitioning of the rows of the horizontal table into subtables (in all figures for a given dataset, we depict the same number of columns for easy comparison, even if some columns are not present in a given implicit sort of the sort refinement).

5. FORMAL DEFINITION OF THE DECISION PROBLEM

Fix a rule r . The main problem that we address in this paper can be formalized as follows.

Problem: $\text{EXISTSSORTREFINEMENT}(r)$
Input: An RDF graph D , a rational number θ such that $0 \leq \theta \leq 1$, and a positive integer k .
Output: **true** if there exists an σ_r -sort refinement \mathcal{T} of D with threshold θ that contains at most k implicit sorts, and **false** otherwise.

In the following theorem, we pinpoint the complexity of the problem $\text{EXISTSSORTREFINEMENT}(r)$.

THEOREM 5.1.

- $\text{EXISTSSORTREFINEMENT}(r)$ is in NP for every rule r .
- There is a rule r_0 for which $\text{EXISTSSORTREFINEMENT}(r_0)$ is NP-complete. Moreover, this result holds even if we fix $k = 3$ and $\theta = 1$. \square

The first part of Theorem 5.1 is a corollary of the fact that one can efficiently check if a sort refinement is an entity preserving partition of an RDF graph and has the correct threshold, as for every (fixed) rule r , function σ_r can be computed in polynomial time. The second statement in Theorem 5.1 shows that there exists a (fixed) rule r_0 for which $\text{EXISTSSORTREFINEMENT}(r_0)$ is NP-hard, even if the structuredness threshold θ and the maximum amount of implicit sorts k are fixed. The proof of this part of the theorem relies on a reduction from the graph 3-coloring problem to $\text{EXISTSSORTREFINEMENT}(r_0)$ with $\theta = 1$ and $k = 3$. In this reduction, a graph G (the input to the 3-coloring problem) is used to construct an RDF graph D_G in such a way that a partition of the nodes of G can be represented by an entity preserving partitioning of the RDF graph. Although the rule r_0 will not be shown explicitly, it is designed to calculate the probability that 2 subjects in a subset of the entity preserving partitioning of D_G represent 2 nodes of G which are not adjacent. This probability will be 1 only when said subset represents an independent set of G . Therefore, setting the threshold $\theta = 1$ ensures that each subset of D_G will represent an independent set of G . Finally, setting $k = 3$ ensures that at most 3 subsets will be generated. If the graph G is 3-colorable, then it will be possible to generate the sort refinement of D_G in which each subset represents an independent set of G , and thus will have a structuredness value of 1. Conversely, if there is a sort refinement of at most 3 subsets, then it is possible to partition the nodes of G into 3 or less independent sets, and thus, is 3-colorable.

Note that the fixed rule r_0 used in the reduction does not contain statements of the form $\text{subj}(c) = a$ (where a is a constant URI), although it does use statements of the form $\text{prop}(c) = a$ and other equalities. It is natural to exclude rules which mention specific subjects, as the structuredness of an RDF graph should not depend on the presence of a particular subject, but rather on the general uniformity of all entities in the RDF graph.

The decision problem presented in this section is theoretically intractable, which immediately reduces the prospects of finding reasonable algorithms for its solution. Nevertheless, the inclusion of the problem in NP points us to three NP-complete problems for which much work has been done to produce efficient solvers: the travelling salesman problem, the boolean satisfiability problem, and the integer linear programming problem.

An algorithm for our problem must choose a subset for each signature set, producing a series of decisions which could in principle be expressed as boolean variables, suggesting the boolean satisfiability problem. However, for a candidate sort refinement the function σ_r must be computed for every subset, requiring non-trivial arithmetics which cannot be naturally formulated as a boolean formula. Instead, and as one of the key contributions of this paper,

we have successfully expressed the previous decision problem in a natural way as an instance of Integer Linear Programming. It is to this reduction that we turn to in the next section.

6. REDUCING TO INTEGER LINEAR PROGRAMMING

We start by describing the general structure of the Integer Linear Programming (ILP) instance which, given a fixed rule r , solves the problem $\text{EXISTSSORTREFINEMENT}(r)$. Given an RDF graph D , a rational number θ such that $0 \leq \theta \leq 1$ and a positive integer k , we define in this section an instance of integer linear programming, which can be represented as a pair $(A_{(D,k,\theta)}, \vec{b}_{(D,k,\theta)})$, where $A_{(D,k,\theta)}$ is a matrix of integer values, $\vec{b}_{(D,k,\theta)}$ is a vector of integer values, and the problem is to find a vector \vec{d} of integer values (i.e. the values assigned to the variables of the system of equations) such that $A_{(D,k,\theta)}\vec{d} \leq \vec{b}_{(D,k,\theta)}$. Moreover, we prove that $(D, k, \theta) \in \text{EXISTSSORTREFINEMENT}(r)$ if and only if the instance $(A_{(D,k,\theta)}, \vec{b}_{(D,k,\theta)})$ has a solution.

Intuitively, our ILP instance works as follows: the integer variables decide which signature sets are to be included in which subsets, and they keep track of which properties are used in each subset. Also, we group variable assignments into objects we call *rough* variable assignments, which instead of assigning each variable to a subject and a property, they assign each variable to a *signature set* and a property. In this way, another set of variables keeps track of which rough assignments are valid in a given subset (i.e. the rough assignment mentions only signature sets and properties which are present in the subset). With the previous, we are able to count the total and favorable cases of the rule for each subset.

For the following, fix a rule $r = \varphi_1 \mapsto \varphi_2$ and assume that $\text{var}(\varphi_1) = \{c_1, \dots, c_n\}$ (recall that $\text{var}(\varphi_2) \subseteq \text{var}(\varphi_1)$). Also, fix a rational number $\theta \in [0, 1]$, a positive integer k , and an RDF graph D , with the matrix $M = M(D)$.

6.1 Variable definitions

We begin by defining the ILP instance variables. Recall that our goal when solving $\text{EXISTSSORTREFINEMENT}(r)$ is to find a σ_r -sort refinement of D with threshold θ and at most k implicit sorts.

All the variables used in the ILP instance take only integer values. To introduce these variables, we begin by defining the set of signatures of D as $\Lambda(D) = \{\text{sig}(s, D) \mid s \in S(D)\}$, and for every $\mu \in \Lambda(D)$, by defining the support of μ , denoted by $\text{supp}(\mu)$, as the set $\{p \in P(D) \mid \mu(p) = 1\}$. Then for each $i \in \{1, \dots, k\}$ and each $\mu \in \Lambda(D)$, we define the variable:

$$X_{i,\mu} = \begin{cases} 1 & \text{if signature } \mu \text{ is placed in implicit sort } i \\ 0 & \text{otherwise.} \end{cases}$$

These are the primary variables, as they encode the generated sort refinement. Notice that it could be the case that for some $i \in \{1, \dots, k\}$ value 0 is assigned to every variable $X_{i,\mu}$ ($\mu \in \Lambda(D)$), in which case we have that the i -th implicit sort is empty.

For each $i \in \{1, \dots, k\}$ and each $p \in P(D)$ define the variable:

$$U_{i,p} = \begin{cases} 1 & \text{if implicit sort } i \text{ uses property } p \\ 0 & \text{otherwise.} \end{cases}$$

Each variable $U_{i,p}$ is used to indicate whether the i -th implicit sort uses property p , that is, if implicit sort i includes a signature $\mu \in \Lambda(D)$ such that $\mu(p) = 1$ ($p \in \text{supp}(\mu)$).

For the last set of variables, we consider a *rough* assignment of variables in φ_1 to be a mapping of each variable to a signature

and a property. Recall that $\text{var}(\varphi_1) = \{c_1, \dots, c_n\}$. Then we denote rough assignments with $\tau = ((\mu_1, p_1), \dots, (\mu_n, p_n)) \in (\Lambda(D) \times P(D))^n$, and for each $i \in \{1, \dots, k\}$ and each $\tau \in (\Lambda(D) \times P(D))^n$, we define the variable:

$$T_{i,\tau} = \begin{cases} 1 & \text{if } \tau \text{ is consistent in the } i\text{-th implicit sort} \\ 0 & \text{otherwise.} \end{cases}$$

The rough assignment $\tau = ((\mu_1, p_1), \dots, (\mu_n, p_n))$ is consistent in the i -th implicit sort if it only mentions signatures and properties which are present in it, that is, if for each $j \in \{1, \dots, n\}$ we have that μ_j is included in the i -th implicit sort and said implicit sort uses p_j .

6.2 Constraint definitions

Define function $\text{count}(\varphi, \tau, M)$ to be the number of variable assignments for rule r which are restricted by the rough assignment τ and which satisfy the formula φ . Formally, if $\tau = ((\mu_1, p_1), \dots, (\mu_n, p_n))$, then $\text{count}(\varphi, \tau, M)$ is defined as the cardinality of the following set:

$$\{\rho \mid \rho \text{ is a variable assignment for } D \text{ s.t. } \text{dom}(\rho) = \text{var}(\varphi), \\ (M, \rho) \models \varphi \text{ and for every } i \in \{1, \dots, n\}, \\ \text{if } \rho(c_i) = (s, p) \text{ then } \text{sig}(s, D) = \mu_i \text{ and } p = p_i\}.$$

Note that the value of $\text{count}(\varphi, \tau, M)$ is calculated offline and is used as a constant in the ILP instance. We now present the inequalities that constrain the acceptable values of the defined variables.

- The following inequalities specify the obvious lower and upper bounds of all variables:

$$\begin{aligned} 0 \leq X_{i,\mu} \leq 1 & \quad i \in \{1, \dots, k\} \text{ and } \mu \in \Lambda(D) \\ 0 \leq U_{i,p} \leq 1 & \quad i \in \{1, \dots, k\} \text{ and } p \in P(D) \\ 0 \leq T_{i,\tau} \leq 1 & \quad i \in \{1, \dots, k\} \text{ and } \tau \in (\Lambda(D) \times P(D))^n \end{aligned}$$

- For every $\mu \in \Lambda(D)$, the following equation indicates that the signature μ must be assigned to exactly one implicit sort:

$$\sum_{i=1}^k X_{i,\mu} = 1.$$

- For every $i \in \{1, \dots, k\}$ and $p \in P(D)$, we include the following equations to ensure that $U_{i,p}$ is assigned to 1 if and only if the i -th implicit sort includes a signature $\mu \in \Lambda(D)$ such that $\mu(p) = 1$ ($p \in \text{supp}(\mu)$):

$$\begin{aligned} X_{i,\mu} & \leq U_{i,p} & \text{if } p \in \text{supp}(\mu) \\ U_{i,p} & \leq \sum_{\mu' \in \Lambda(D) : p \in \text{supp}(\mu')} X_{i,\mu'} \end{aligned}$$

The first equation indicates that if signature μ has been assigned to the i -th implicit sort and $p \in \text{supp}(\mu)$, then p is one of the properties that must be considered when computing σ_r in this implicit sort. The second equation indicates that if p is used in the computation of σ_r in the i -th implicit sort, then this implicit sort must include a signature $\mu' \in \Lambda(D)$ such that $p \in \text{supp}(\mu')$.

- For $i \in \{1, \dots, k\}$, and $\tau = ((\mu_1, p_1), \dots, (\mu_n, p_n)) \in (\Lambda(D) \times P(D))^n$, recall that $T_{i,\tau} = 1$ if and only if for every $j \in \{1, \dots, n\}$, it holds that $X_{i,\mu_j} = 1$ and $U_{i,p_j} = 1$. This is expressed as integer linear equations as follows:

$$\sum_{j=1}^n (X_{i,\mu_j} + U_{i,p_j}) \leq T_{i,\tau} + 2 \cdot n - 1$$

$$2 \cdot n \cdot T_{i,\tau} \leq \sum_{j=1}^n (X_{i,\mu_j} + U_{i,p_j})$$

The first equation indicates that if the signatures μ_1, \dots, μ_n are all included in the i -th implicit sort (each variable X_{i,μ_j} is assigned value 1), and said implicit sort uses the properties p_1, \dots, p_n (each variable U_{i,p_j} is assigned value 1), then τ is a valid combination when computing favorable and total cases (variable $T_{i,\tau}$ has to be assigned value 1). Notice that if any of the variables $X_{1,\mu_1}, U_{1,p_1}, \dots, X_{n,\mu_n}, U_{n,p_n}$ is assigned value 0 in the first equation, then $\sum_{j=1}^n (X_{i,\mu_j} + U_{i,p_j}) \leq 2 \cdot n - 1$ and, therefore, no restriction is imposed on $T_{i,\tau}$ by this equation, as we already have that $0 \leq T_{i,\tau}$. The second equation indicates that if variable $T_{i,\tau}$ is assigned value 1, meaning that τ is considered to be a valid combination when computing σ_r over the i -th implicit sort, then each signature mentioned in τ must be included in this implicit sort (each variable X_{i,μ_j} has to be assigned value 1), and each property mentioned in τ is used in this implicit sort (each variable U_{i,p_j} has to be assigned value 1).

- Finally, assuming that $\theta = \theta_1/\theta_2$, where θ_1, θ_2 are natural numbers, we include the following equation for each $i \in \{1, \dots, k\}$:

$$\theta_2 \cdot \left(\sum_{\tau \in (\Lambda(D) \times P(D))^n} \text{count}(\varphi_1 \wedge \varphi_2, \tau, M) \cdot T_{i,\tau} \right)$$

$$\geq \theta_1 \cdot \left(\sum_{\tau \in (\Lambda(D) \times P(D))^n} \text{count}(\varphi_1, \tau, M) \cdot T_{i,\tau} \right)$$

To compute the numbers of favorable and total cases for σ_r over the i -th implicit sort, we consider each rough assignment τ in turn. The term $\sum_{\tau \in (\Lambda(D) \times P(D))^n} \text{count}(\varphi_1 \wedge \varphi_2, \tau, M) \cdot T_{i,\tau}$ evaluates to the amount of favorable cases (i.e. variable assignments which satisfy the antecedent and the consequent of the rule), while the term $\sum_{\tau \in (\Lambda(D) \times P(D))^n} \text{count}(\varphi_1, \tau, M) \cdot T_{i,\tau}$ evaluates to the number of total cases (i.e. variable assignments which satisfy the antecedent of the rule). Consider the former term as an example: for each rough variable assignment τ , if τ is a valid combination in the i -th implicit sort, then the amount of variable assignments which are compatible with τ and which satisfy the full rule are added.

It is now easy to see that the following result holds.

PROPOSITION 6.1. There exists a σ_r -sort refinement of D with threshold θ that contains at most k implicit sorts if and only if the instance of ILP defined in this section has a solution. \square

6.3 Implementation details

Although the previously defined constraints suffice to solve the decision problem, in practice the search space is still too large due to the presence of sets of solutions which are *equivalent*, in the sense that the variables describe the same partitioning of the input RDF graph D . More precisely, if there is a solution of the ILP instance where for each $i \in \{1, \dots, k\}$, $\mu \in \Lambda(D)$, $p \in P(D)$, and

$\tau \in (\Lambda(D) \times P(D))^n$, $X_{i,\mu} = a_{i,\mu}$, $U_{i,p} = b_{i,p}$, and $T_{i,\tau} = c_{i,\tau}$, then for any permutation (l_1, \dots, l_k) of $(1, \dots, k)$, the following is also a solution: $X_{i,\mu} = a_{l_i,\mu}$, $U_{i,p} = b_{l_i,p}$, and $T_{i,\tau} = c_{l_i,\tau}$.

In order to break the symmetry between these equivalent solutions, we define the following hash function for the i -th implicit sort. For this, consider $\ell = |\Lambda(D)|$ and consider any (fixed) ordering μ_1, \dots, μ_ℓ of the signatures in $\Lambda(D)$. Then:

$$\text{hash}(i) = \sum_{j=0}^{\ell} 2^j X_{i,\mu_j}$$

With the previous hash function defined, the following constraint is added, for $i = 1, \dots, k - 1$:

$$\text{hash}(i) \leq \text{hash}(i + 1).$$

The hash function as defined above uniquely identifies a subset of signatures, and thus the previous constraints eliminate the presence of multiple solutions due to permutations of the i index. Care must be taken, however, if the amount of signatures in the RDF graph is large (64 for DBpedia Persons) as large exponent values cause numerical instability in commercial ILP solvers. This issue may be addressed on a case by case basis. One alternative is to limit the maximum exponent in the term 2^j , which has the drawback of increasing the amount of collisions of the hash function, and therefore permitting the existence of more equivalent solutions.

7. EXPERIMENTAL RESULTS

For our first experiments, we consider two real datasets—DBpedia Persons and WordNet Nouns—and two settings:

- **A highest θ sort refinement for $k = 2$:** This setup can be used to obtain an intuitive understanding of the dataset at hand. We fix $k = 2$ to force at most 2 implicit sorts.
- **A lowest k sort refinement for $\theta = 0.9$:** As a complementary approach, we specify $\theta = 0.9$ as the threshold, and we search for the lowest k such that an sort refinement with threshold θ and k implicit sorts exists. This approach allows a user to refine their current sort by discovering sub-sorts. In some cases the structuredness of the original dataset under some structuredness function is higher than 0.9, in which case we increase the threshold to a higher value.

In the first case the search for the optimum value of θ is done in the following way: starting from the initial structuredness value $\theta = \sigma_r(D)$ (for which a solution is guaranteed) and for values of θ incremented in steps of 0.01, an ILP instance is generated with $k = 2$ and the current value of θ . If a solution is found by the ILP solver, then said solution is stored. If the ILP instance is found to be infeasible, then the last stored solution is used (this is the solution with the highest threshold). This sequential search is preferred over a binary search because the latter will generate more infeasible ILP instances on average, and it has proven to be much slower to find an instance infeasible than to find a solution to a feasible instance. A similar strategy is used for the second case (the search for the lowest k), with the following difference: for some setups it is more efficient to search downwards, starting from $k = |\Lambda(D)|$ (i.e. as many available sorts as signatures in the dataset), and yet for others it is preferable to search upwards starting from $k = 1$, thus dealing with a series of infeasible ILP instances, before discovering the first value of k such that a solution is found. Which of the two directions is to be used has been decided on a case by case basis.

The amount of variables and constraints in each ILP instance depends on the amount of variables of the rules, on the degrees

of freedom given to the variables in the rules (e.g. the two variables in $\sigma_{\text{Dep}}[\mathbf{p}_1, \mathbf{p}_2]$ lose a degree of freedom when considering the restriction $\text{subj}(c_1) = \text{subj}(c_2)$ in the antecedent), and on the characteristics of the dataset. Here, the enormous reduction in size offered by the signature representation of a dataset has proven crucial for the efficiency of solving the ILP instances.

The previous two settings are applied both to the DBpedia Persons and WordNet Nouns datasets. Furthermore, they are repeated for the σ_{Cov} , σ_{Sim} , and σ_{Dep} functions (the last function is only used on DBpedia Persons). All experiments are conducted on dual 2.3GHz processor machine (with 6 cores per processor), and 64GB of RAM. The ILP solver used is IBM ILOG CPLEX version 12.5.

In section 7.3 we study how our solution scales with larger and more complex datasets by extracting a representative sample of explicit sorts from the knowledge base YAGO and solving a highest θ sort refinement for fixed k for each explicit sort. Finally, in section 7.4 we challenge the solution to recover two different explicit sorts from a mixed dataset, providing a practical test of the results.

7.1 DBpedia Persons

DBpedia corresponds to RDF data extracted from Wikipedia. DBpedia Persons refers to the following subgraph (where Person is a shorthand for `http://xmlns.com/foaf/0.1/Person`):

$$D_{\text{DBpedia Persons}} = \{(s, p, o) \in D_{\text{DBpedia}} \mid (s, \text{type}, \text{Person}) \in D_{\text{DBpedia}}\}.$$


This dataset is 534 MB in size, and contains 4,504,173 triples, 790,703 subjects, and 8 properties (excluding the type property). It consists of 64 signatures, requiring only 3 KB of storage. The list of properties is as follows: `deathPlace`, `birthPlace`, `description`, `name`, `deathDate`, `birthDate`, `givenName`, and `surName` (note that these names are abbreviated versions of the full URIs).

For this sort, $\sigma_{\text{Cov}} = 0.54$, and $\sigma_{\text{Sim}} = 0.77$. We are also interested in studying the dependency functions for different properties \mathbf{p}_1 and \mathbf{p}_2 . If $\mathbf{p}_1 = \text{deathPlace}$ and $\mathbf{p}_2 = \text{deathDate}$, for example, then the value of the function $\sigma_{\text{SymDep}}[\text{deathPlace}, \text{deathDate}]$ is 0.39. This specific choice of \mathbf{p}_1 and \mathbf{p}_2 is especially interesting because it might be tempting to predict that a death date and a death place are equally obtainable for a person. However, the value 0.39 reveals the contrary. The generally low values for the three structuredness functions discussed make DBpedia Persons interesting to study.


7.1.1 A highest θ sort refinement for $k = 2$

We set $k = 2$ in order to find a two-sort sort refinement with the best threshold θ . Figure 4a shows the result for the σ_{Cov} function. The left sort, which is also the largest (having 528,593 subjects) has a very clear characteristic: no subject has a `deathDate` or a `deathPlace`, i.e. it represents the sort for people that are alive! Note that without encoding any explicit schema semantics in the generic rule of σ_{Cov} , our ILP formulation is able to discover a very intuitive decomposition of the initial sort. In the next section, we show that this is the case even if we consider larger values of k . In this experiment, each ILP instance is solved in under 800 ms.

Figure 4b shows the results for the σ_{Sim} function. Here, the second sort accumulates subjects for which very little data is known (other than a person’s name). Notice that whereas Cov has excluded the columns `deathPlace`, `description`, and `deathDate` from its first sort, Sim does not for its second sort, since it does not *penalize* the largely missing properties in these columns (which was what motivated us to introduce σ_{Sim} in the first place). Also, notice that unlike the σ_{Cov} function, the cardinality of the generated sorts from σ_{Sim} is more *balanced*. In this experiment, each



(a) Using the σ_{Cov} function, the left sort has 528,593 subjects and 8 signatures, $\sigma_{\text{Cov}} = 0.73$, and $\sigma_{\text{Sim}} = 0.85$. The right sort has 262,110 subjects and 56 signatures, $\sigma_{\text{Cov}} = 0.71$, and $\sigma_{\text{Sim}} = 0.78$.



(b) Using the σ_{Sim} function, the left sort has 387,297 subjects and 37 signatures, $\sigma_{\text{Cov}} = 0.67$, and $\sigma_{\text{Sim}} = 0.82$. The right sort has 403,406 subjects and 27 signatures, $\sigma_{\text{Cov}} = 0.42$ and $\sigma_{\text{Sim}} = 0.85$.




(c) Using the $\sigma_1 = \sigma_{\text{SymDep}}[\text{deathPlace}, \text{deathDate}]$ function, the left sort has 305,610 subjects and 25 signatures, $\sigma_{\text{Cov}} = 0.66$, $\sigma_{\text{Sim}} = 0.80$, and $\sigma_1 = 1.0$. The right sort has 485,093 subjects and 39 signatures, $\sigma_{\text{Cov}} = 0.52$, $\sigma_{\text{Sim}} = 0.78$, and $\sigma_1 = 0.82$.

Figure 4: DBpedia Persons split into $k = 2$ implicit sorts, using the structuredness functions (a) σ_{Cov} , (b) σ_{Sim} , and (c) σ_{Dep} .


ILP instance is solved in under 2 minutes, except the infeasible (last) instance which was completed in 2 hrs.

Finally, Figure 4c shows the results for $\sigma_{\text{SymDep}}[\text{deathPlace}, \text{deathDate}]$, a structuredness function in which we measure the probability that, if a subject has a `deathPlace` or a `deathDate`, it has both. In the resulting sort refinement, the second sort to the right has a high value of 0.82. It is easy to see that indeed our ILP solution does the right thing. In the sort on the right, the `deathDate` and `deathPlace` columns look almost identical which implies that indeed whenever a subject has one property it also has the other. As far as the sort on the left is concerned, this includes all subjects that do not have a `deathPlace` column. This causes the sort to have a structuredness value of 1.0 for $\sigma_{\text{SymDep}}[\text{deathPlace}, \text{deathDate}]$ since the rule is trivially satisfied; the absence of said column eliminates all total cases (i.e. there are no assignments in the rule that represents $\sigma_{\text{SymDep}}[\text{deathPlace}, \text{deathDate}]$ for which the antecedent is true, as it is never true that $\text{prop}(c_1) = \text{deathPlace}$). This setting is completed in under 1 minute.

7.1.2 A lowest k sort refinement for $\theta = 0.9$



(a) DBpedia Persons split into $k = 9$ implicit sorts, using the σ_{Cov} function. The threshold of this sort refinement is $\theta = 0.9$, (i.e. every sort D_i has $\sigma_{Cov}(D_i) \geq 0.9$). The sizes of the sorts range from 260,585 subjects (the second sort) to 10,748 subjects (the seventh sort).



(b) A $k = 4$ σ_{Sim} -sort refinement with threshold 0.9 for DBpedia Persons. The sizes of the sorts range from 292,880 subjects (the first sort) to 87,117 subjects (the third sort).

Figure 5: DBpedia Persons split into the lowest k such that the threshold is $\theta = 0.9$, using the structuredness functions (a) σ_{Cov} , and (b) σ_{Sim} .

We now consider a fixed threshold $\theta = 0.9$. We seek the smallest sort refinement for DBpedia persons with this threshold. Figure 5a shows the result for σ_{Cov} , where the optimum value found is for $k = 9$. As in the previous setting, the Cov function shows a clear tendency to produce sorts which do not use all the columns (i.e. sorts which exclude certain properties). People that are alive can now be found in the first, second, third, fourth, and sixth sorts. The first sort considers living people who have a description (and not even a birth place or date). The second sort shows living people who are even missing the description field. The third sort considers living people who have a description and a birth date or a birth place (or both). The fourth sort considers living people with a birth place or birth date but no description. Finally, the sixth sort considers living people with a birth place only. It is easy to see that similarly dead people are separated into different sorts, based on the properties that are known for them. The eighth sort is particularly interesting since it contains people for which we mostly have all the properties. This whole experiment was completed in 30 minutes.

Figure 5b shows the result for σ_{Sim} , where the optimum value found is for $k = 4$. Again, the function is more lenient when properties appear for only a small amount of subjects (hence the smaller k). This is evident in the first sort for this function, which corresponds roughly to the second sort generated for the σ_{Cov} function (Fig. 5a) but also includes a few subjects with birth/death places/dates. This is confirmed by the relative sizes of the two sorts, with the sort for σ_{Cov} having 260,585 subjects, while the sort for σ_{Sim} having 292,880 subjects. This experiment is more difficult as the running time of individual ILP instances is apx. 8 hrs.

7.1.3 Dependency functions in DBpedia Persons

We now turn our attention to the dependency functions. In terms

	dP	bP	dD	bD
deathPlace	1.0	.93	.82	.77
birthPlace	.26	1.0	.27	.75
deathDate	.43	.50	1.0	.89
birthDate	.17	.57	.37	1.0

Table 1: DBpedia Persons structuredness according to σ_{Dep} with different combinations of parameters \mathbf{p}_1 and \mathbf{p}_2 . The property names are abbreviated in the column headers.

of creating a new sort refinement using the function $\sigma_{Dep}[\mathbf{p}_1, \mathbf{p}_2]$, for any constants $\mathbf{p}_1, \mathbf{p}_2 \in \mathcal{U}$, we can generate a sort refinement with $\theta = 1.0$ for $k = 2$, consisting of the following two sorts: (i) all entities which do not have \mathbf{p}_1 , and (ii) all entities which do have \mathbf{p}_2 . The sort (i) will have structuredness 1.0 because there are no assignments that satisfy the antecedent (no assignments satisfy $\text{prop}(c_2) = \mathbf{p}_1$), and sort (ii) has structuredness 1.0 because every assignment which satisfies the antecedent will also satisfy the consequent ($\text{val}(c_2) = 1$ because all entities have \mathbf{p}_2). On the other hand, σ_{SymDep} with constants $\mathbf{p}_1, \mathbf{p}_2 \in \mathcal{U}$ can generate an sort refinement with $\theta = 1.0$ for $k = 3$, consisting of the following three sorts: (i) entities which have \mathbf{p}_1 but not \mathbf{p}_2 , (ii) entities which have \mathbf{p}_2 but not \mathbf{p}_1 , and (iii) entities which have both \mathbf{p}_1 and \mathbf{p}_2 or have neither. The first two sorts will not have any total cases, and for the third sort every total case is also a favorable case.

The dependency functions, as shown, are not very well suited to the task of finding the lowest k such that the threshold θ is met, which is why these functions were not included in the previous results. The dependency functions are useful, however, for characterizing an RDF graph or a sort refinement which was generated with a different structuredness function, such as σ_{Cov} or σ_{Sim} , since they can help analyze the relationship between the properties in an RDF graph. To illustrate, we consider the $\sigma_{Dep}[\mathbf{p}_1, \mathbf{p}_2]$ function, and we tabulate (in Table. 1) the structuredness value of DBpedia Persons when replacing the parameters \mathbf{p}_1 and \mathbf{p}_2 by all possible combinations of deathPlace, birthPlace, deathDate, and birthDate. Recall that σ_{Dep} with parameters $\mathbf{p}_1 = \text{deathPlace}$ and $\mathbf{p}_2 = \text{birthPlace}$ measures the probability that a subject which has deathPlace also has birthPlace.

The table reveals a very surprising aspect of the dataset. Namely, the first row shows high structuredness values when $\mathbf{p}_1 = \text{deathPlace}$. This implies that if we somehow know the deathPlace for a particular person, there is a very high probability that we also know all the other properties for her. Or, to put it another way, knowing the death place of a person implies that we know a lot about the person. This is also an indication that it is somehow the hardest fact to acquire, or the fact that is least known among persons in DBpedia. Notice that none of the other rows have a similar characteristic. For example, in the second row we see that given the birthPlace of a person there is a small chance (0.27) that we know her deathDate. Similarly, given the deathDate of a person there is only a small chance (0.43) that we know the deathPlace.

We can do a similar analysis with the $\sigma_{SymDep}[\mathbf{p}_1, \mathbf{p}_2]$ function. In Table 2 we show the pairs of properties with the highest and lowest values of σ_{SymDep} . Given that the name property in DBpedia persons is the only property that every subject has, one would expect that the most correlated pair of properties would include name. Surprisingly, this is not the case. Properties givenName and surName are actually the most correlated properties, probably stemming from the fact that these two properties are extracted from the same source. The least correlated properties all involve deathPlace and the properties of name, givenName and surName.

7.2 For WordNet Nouns

P1	P2	σ_{SymDep}
givenName	surName	1.0
name	givenName	.95
name	surName	.95
name	birthDate	.53
...
description	givenName	.14
deathPlace	name	.11
deathPlace	givenName	.11
deathPlace	surName	.11

Table 2: A ranking of DBpedia Persons structuredness according to σ_{SymDep} with different combinations of the 8 properties in $P(D_{\text{DBpedia Persons}})$. Only the highest and lowest entries are shown.

WordNet is a lexical database for the english language. WordNet Nouns refers to the following subgraph (where Noun stands for <http://www.w3.org/2006/03/wn/wn20/schema/NounSynset>):

$$D_{\text{WordNet Nouns}} = \{(s, p, o) \in D_{\text{WordNet}} \mid (s, \text{type, Noun}) \in D_{\text{WordNet}}\}.$$

This dataset is 101 MB in size, and contains 416,338 triples, 79,689 subjects, and 12 properties (excluding the type property). Its signature representation consists of 53 signatures, stored in 3 KB. The properties are the following: gloss, label, synsetId, hyponymOf, classifiedByTopic, containsWordSense, memberMeronymOf, partMeronymOf, substanceMeronymOf, classifiedByUsage, classifiedByRegion, and attribute.


For this sort, $\sigma_{\text{Cov}} = 0.44$, and $\sigma_{\text{Sim}} = 0.93$. There is a significant difference in the structuredness of WordNet Nouns as measured by the two functions. This difference is clearly visible in the signature view of this dataset (fig. 3); the presence of nearly empty properties (i.e. properties which relatively few subjects have) is highly penalized by the Cov rule, though mostly ignored by Sim.

7.2.1 A highest θ sort refinement for $k = 2$


As mentioned, the WordNet case proves to be very different from the previous, partly because this dataset has roughly 5 dominant signatures representing a most subjects, and yet only using 8 of the 12 properties, causing difficulties when partitioning the dataset.

Figure 6a shows the result for σ_{Cov} . The largest difference between both sorts is that the left sort mostly consists of subjects which have the memberMeronymOf property (the 7th property). The improvement in the structuredness of these two sorts is very small in comparison to the original dataset (from 0.44 to 0.55), suggesting that $k = 2$ is not enough to discriminate sub-sorts in this dataset, and with this rule. This is due to the presence of many of signatures which represent very few subjects, and have different sets of properties. Here, all ILP instances were solved in under 1 s.

Figure 6b shows the result for σ_{Sim} . The difference between the two sorts is gloss, which is absent in the left sort. The placement of the smaller signatures does not seem to follow any pattern, as the Sim function is not sensitive to their presence. Although the structuredness is high in this case, the improvement is small, since the original dataset is highly structured with respect to σ_{Sim} anyway. A discussion is in order with respect to the running times. Recall that the ILP instances are solved for increasing values of θ (the increment being 0.01). For all values of θ lower than 0.95 each ILP instance is solved in less than 5 s. For the value $\theta = 0.95$ however (the first value for which there is no solution), after 75 hours of running time, the ILP solver was not able to find a solution or prove the system infeasible. Although there is an enormous asymmetry between the ease of finding a solution and the difficulty of proving



(a) Using the σ_{Cov} function, the left sort has 14,938 subjects and 35 signatures, $\sigma_{\text{Cov}} = 0.55$, $\sigma_{\text{Sim}} = 0.93$. The right sort has 64,751 subjects and 18 signatures, $\sigma_{\text{Cov}} = 0.56$, $\sigma_{\text{Sim}} = 0.95$.



(b) Using the σ_{Sim} function, the left sort has 7,311 subjects and 13 signatures, $\sigma_{\text{Cov}} = 0.34$, and $\sigma_{\text{Sim}} = 0.98$. The right sort has 72,378 subjects and 40 signatures, $\sigma_{\text{Cov}} = 0.45$, and $\sigma_{\text{Sim}} = 0.94$.

Figure 6: WordNet Nouns split into $k = 2$ implicit sorts, using two different structuredness functions: (a) σ_{Cov} , and (b) σ_{Sim} .

an instance infeasible, in every instance a higher threshold solution is found, in which case it is reasonable to let the user specify a maximum running time and keep the best solution found so far.


7.2.2 A lowest k sort refinement for fixed θ

As with the previous experimental setup, Nouns proves more difficult to solve. For the σ_{Cov} we set the usual threshold of 0.9, however, since the structuredness value of Wordnet Nouns under the σ_{Sim} function is 0.93 originally, this exercise would be trivial if θ is 0.9. For that reason, in this last case we fix the threshold at 0.98.


Figure 7a shows the first 10 sorts of the $k = 31$ solution for σ_{Cov} . The sheer amount of sorts needed is an indication that WordNet Nouns already represents a highly structured sort. The sorts in many cases correspond to individual signatures, which are the smallest sets of identically structured entities. In general, it is probably not of interest for a user or database administrator to be presented with an sort refinement with so many sorts. This setup was the longest running, at an average 7 hours running time per ILP instance. This large number is another indication of the difficulty of partitioning a dataset with highly uniform entities.

Figure 7b shows the solution for σ_{Sim} , which is for $k = 4$. As with the $k = 2$ case, there is a sort which does not include the gloss property. The general pattern of this sort refinement, however, is that the four largest signatures are each placed in their own sort. Beyond that, the presence of the smaller signatures does not greatly affect the structuredness value (runtime: apx. 15 min.).

It is to be expected that a highly structured RDF graph like WordNet Nouns will not be a prime candidate for discovering refinements of the sort, which is confirmed by these experiments.



(a) WordNet Nouns split into $k = 31$ implicit sorts, using the σ_{Cov} function and a threshold of $\theta = 0.9$. Only the first 12 sorts are shown here.



(b) WordNet Nouns split into $k = 4$ implicit sorts, using the σ_{Sim} function. The threshold of this sort refinement is $\theta = 0.98$. The sizes of the sorts range from 52,880 subjects (the third sort) to 7,037 subjects (the first sort).

Figure 7: WordNet Nouns partitioned into the lowest k with a fixed threshold.

7.3 Scalability Analysis

To study how the ILP-based solution scales, we turn to the knowledge base YAGO. This proves to be a useful dataset for this study, as it contains apx. 380,000 explicit sorts extracted from several sources, including Wikipedia. From YAGO, a sample of apx. 500 sorts is randomly selected (as most explicit sorts in YAGO are small, we manually included larger sorts in the sample). The sample contains sorts with sizes ranging from ~ 100 to $\sim 10^5$ subjects, from 1 to ~ 350 signatures, and from ~ 10 to ~ 40 properties.


As can be seen in the histograms of figure 8, there exist larger and more complex explicit sorts in YAGO than those sampled (e.g. sorts with $\sim 20,000$ signatures or ~ 80 properties). However, the same histograms show that 99.9% of YAGO explicit sorts have under 350 signatures, and 99.8% have under 40 properties, and thus the sample used is representative of the whole. We also note that, in a practical setting, the amount of properties can be effectively reduced *with the language itself* by designing a rule which restricts the structuredness function to a set of predefined properties.

For each sort in the sample, a highest θ sort refinement for $k = 2$ is solved, and the total time taken to solve the ILP instances is recorded (we may solve multiple ILP instances, each for larger values of θ , until an infeasible instance is found). See figure 8.


Due to lack of space, a plot for the runtime as a function of the number of subjects in a sort is not shown, but the result is as expected: the runtime of the ILP-based solution does not depend on the amount of subjects in an explicit sort; the difficulty of the problem depends on the structure of the dataset, and not its size.

In contrast, figures 8a and 8b both reveal clear dependencies of the runtime on the number of signatures and properties in a sort.

In the first case, the best polynomial fit shows that runtime is $\mathcal{O}(s^{2.5})$, where s is the number of signatures (the best fit is obtained by minimising the sum of squared errors for all data points). This result limits the feasibility of using the ILP-based solution for more complicated explicit sorts. However, sorts with more than 350



(a) Left: A plot of the runtime as a function of the number of signatures in a sort. The best fit is shown, and corresponds to $f(s) \approx x^{2.53}$ ($R^2 = 0.72$). Right: A logarithmic histogram of the number of signatures in YAGO explicit sorts. Note that 99.9% of sorts have less than 350 signatures.



(b) Left: A logarithmic plot of the runtime as a function of the number of properties in a sort. The best fit is shown, and corresponds to $f(p) \approx e^{0.28p}$ ($R^2 = 0.61$). Right: A histogram of the number of properties in YAGO explicit sorts. Note that 99.8% of sorts have less than 40 properties.

Figure 8: For a sample of YAGO explicit sorts, a highest θ sort refinement for $k = 2$ is solved.

signatures are very rare. In the second case, the best exponential fit shows that runtime is $\mathcal{O}(e^{0.28p})$, where p is the number of properties. Although an exponential dependency can be worrisome, the situation is similar to that for signatures.

7.4 Semantic Correctness

A final practical question is whether our ILP-based solution can reproduce existing differences between sorts. To address this, we have chosen two explicit sorts from YAGO: Drug Companies and Sultans. All triples whose subject is declared to be of sort Drug Company or Sultan are included in one mixed dataset. We then solve a highest θ sort refinement for fixed $k = 2$ problem on this mixed dataset and compare the resulting sort refinement with the actual separation between Drug Companies and Sultans.

	Is Drug Company	Is Sultan
Classified as Drug Company	27	17
Classified as Sultan	0	23

In order to describe the quality of these results, we interpret this experiment as the binary classification problem of Drug Companies. Thus, Drug Companies become the positive cases, whereas Sultans become negative cases. With this, the resulting accuracy of the result is 74.6%, the precision is 61.4%, and the recall is 100%. These results can be improved by considering a modified Cov rule which ignores properties defined in the syntax of RDF (`type`, `sameAs`, `subClassOf`, and `label`), which can be achieved by adding four conjuncts of the form $\text{prop}(c) \neq u$ to the antecedent of Cov. This results in 82.1% accuracy, 69.2% precision, and 100% recall.

It must be noted, however, that this experiment relies on the assumption that the initial explicit sorts are well differentiated to begin with, which is precisely the assumption we doubt in this paper.

8. RELATED WORK

Our work is most related to efforts that mine RDF data to discover frequently co-occurring property sets that can be stored in

separate so-called ‘property tables’. A recent example of this approach is exemplified in [7], “*Attribute Clustering by Table Load*” where the authors consider the problem of partitioning the properties of an RDF graph into *clusters*. Each cluster defines the columns of a property table in which each row will represent a subject. A cluster is valid insofar as the *table load factor* remains above a threshold. The table load factor Lee et. al. defined is equivalent to the coverage value defined in [5] (Cov metric as per the notation of this paper). Their approach, however, differs from ours in the following way: while they seek to partition the properties of an RDF graph for the purpose of generating property tables, we seek to discover sets of subjects which, when considered together as an RDF graph, result in a highly structured relational database. The sub-sorts generated by us may use overlapping sets of properties.

Similarly, [4] and [9] use frequent item set sequences [1] data mining techniques to discover, in a RDF dataset, properties that are frequently defined together for the same subject (e.g., first name, last name, address, etc.). Such properties represent good candidates to be stored together in a property table. Although the goal of [4] and [9] is to improve performance by designing a customized database schema to store a RDF dataset, a property table can also be viewed as a refined sort whose set of instances consists of all resources specifying at least one of the properties of the table. In [4] and [9], various important parameters controlling the sort refinement approach are chosen in an ad-hoc manner (e.g., in [4] the minimum support used is chosen after manually inspecting partial results produced by a handful of minimum support values, and, in [9], it is explicitly specified by the user); whereas, in our approach, key parameters (e.g., k and θ) are selected in a principled way to reach an optimal value of a user defined structuredness metric.

Other than the property tables area, our work can be positioned in the broader context of inductive methods to acquiring or refining schema-level knowledge for RDF data [13, 2, 8, 3, 10, 6]. Prior works have typically relied on statistical or logic programming approaches to discover ontological relations between sorts and properties. However, to the best of our knowledge, our work presents the first principled approach to refine the sort by altering the assignment of resources to a refined set of sorts in order to improve some user defined measure of structuredness.

In the area of knowledge discovery in general, the work by Yao [12] offers a nice overview of several information-theoretic measures for knowledge discovery, including, attribute entropy and mutual information. A common characteristic of all these measures is that they focus on the particular values of attributes (in our case, predicates) and attempt to discover relationships between values of the same attribute, or relationships between values of different attributes. As is obvious from Section 3, our work focuses on discovering relationships between entities (and their respective schemas) and therefore we are only interested in the presence (or absence) of predicates for particular attributes for a given entity, therefore ignoring the concrete values stored there. Hence the our measures are orthogonal to those discussed by Yao [12].

9. CONCLUSIONS

We have presented a framework within which it is possible to study the structuredness of RDF graphs using measures which are tailored to the needs of the user or database administrator. This framework includes a formal language for expressing structuredness rules which associate a structuredness value to each RDF graph. We then consider the problem of discovering a partitioning of the entities of an RDF graph into subset which have high structuredness with respect to a specific structuredness function chosen by the user. Although this problem is intractable in general, we define

an ILP instance capable of solving this problem within reasonable time limits using commercially available ILP solvers.

We have used our framework to study two real world RDF datasets, namely DBpedia Persons and WordNet Nouns, the former depending on a publicly editable web source and therefore containing data which does not clearly conform to its schema, and the latter corresponding to a highly uniform set of dictionary entries. In both cases the results obtained were meaningful and intuitive. We also studied the scalability of the ILP-based solution on a sample of explicit sorts extracted from the knowledge base YAGO, showing that our solution is practical for all but a small minority of existing sorts.

The obvious next goal is to better understand the expressiveness of structuredness rules, and to explore subsets of our language with possibly lower computational complexity. In particular, the NP-hardness of the decision problem has been proven for a rule which uses disjunction in the consequent; if this were to be disallowed, it might be possible to lower the complexity of the problem. Finally, it would also be interesting to explore the existence of rules for which a high structuredness value can predict good performance for certain classes of queries.

10. REFERENCES

- [1] R. Agrawal and R. Srikant. Mining sequential patterns. In *ICDE*, pages 3–14, 1995.
- [2] C. d’Amato, N. Fanizzi, and F. Esposito. Inductive learning for the semantic web: What does it buy? *Semant. web*, 1(1,2):53–59, Apr. 2010.
- [3] A. Delleil, C. Faron-Zucker, and R. Dieng. Learning Ontologies from RDF annotations. In *IJCAI 2001 Workshop on Ontology Learning*, volume 38 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2001.
- [4] L. Ding, K. Wilkinson, C. Sayers, and H. Kuno. Application-specific schema design for storing large rdf datasets. In *First Intl Workshop On Practical And Scalable Semantic Systems*, 2003.
- [5] S. Duan, A. Kementsietsidis, K. Srinivas, and O. Udrea. Apples and oranges: a comparison of rdf benchmarks and real rdf datasets. In *SIGMOD*, pages 145–156, 2011.
- [6] G. A. Grimmes, P. Edwards, and A. Preece. Learning Meta-Descriptions of the FOAF Network. In *Proceedings of the Third International Semantic Web Conference (ISWC-04)*, LNCS, pages 152–165, Hiroshima, Japan.
- [7] T. Y. Lee, D. W. Cheung, J. Chiu, S. D. Lee, H. Zhu, P. Yee, and W. Yuan. Automating relational database schema design for very large semantic datasets. Technical report, Dept of Computer Science, Univ. of Hong Kong, 2013.
- [8] J. Lehmann. *Learning OWL class expressions*. PhD thesis, 2010.
- [9] J. J. Levandoski and M. F. Mokbel. Rdf data-centric storage. In *ICWS*, pages 911–918, 2009.
- [10] A. Maedche and V. Zacharias. Clustering ontology-based metadata in the semantic web. In *PKDD*, pages 348–360, 2002.
- [11] Z. Pan and J. Heflin. Dldb: Extending relational databases to support semantic web queries. Technical report, Department of Computer Science, Lehigh University, 2004.
- [12] N. X. Vinh, J. Epps, and J. Bailey. Information theoretic measures for clusterings comparison: Variants, properties, normalization and correction for chance. *J. Mach. Learn. Res.*, 9999:2837–2854, December 2010.
- [13] J. Völker and M. Niepert. Statistical schema induction. In *ESWC*, pages 124–138, 2011.