



PONTIFICIA UNIVERSIDAD CATOLICA DE CHILE
SCHOOL OF ENGINEERING

**A PRINCIPLED APPROACH TO
BRIDGING THE GAP BETWEEN RDF
DATA AND THEIR SCHEMAS**

GONZALO IGNACIO DÍAZ CÁCERES

Thesis submitted to the Office of Research and Graduate Studies
in partial fulfillment of the requirements for the degree of
Master of Science in Engineering

Advisor:

MARCELO ARENAS

Santiago de Chile, September 2013

© MMXIII, GONZALO I. DIAZ

© MMXIII, GONZALO I. DIAZ

Se autoriza la reproducción total o parcial, con fines académicos, por cualquier medio o procedimiento, incluyendo la cita bibliográfica que acredita al trabajo y a su autor.



PONTIFICIA UNIVERSIDAD CATOLICA DE CHILE
SCHOOL OF ENGINEERING

**A PRINCIPLED APPROACH TO
BRIDGING THE GAP BETWEEN RDF
DATA AND THEIR SCHEMAS**

GONZALO IGNACIO DÍAZ CÁCERES

Members of the Committee:

MARCELO ARENAS

JUAN REUTTER

JORGE PÉREZ

MARCELO GUARINI

Thesis submitted to the Office of Research and Graduate Studies
in partial fulfillment of the requirements for the degree of
Master of Science in Engineering

Santiago de Chile, September 2013

© MMXIII, GONZALO I. DIAZ

ACKNOWLEDGEMENTS

I thank the Becas CONICYT program of the Chilean Comisión Nacional de Investigación Científica y Tecnológica (CONICYT) and the School of Engineering at P.U.C. Chile for the financial support during my Masters program. To the members of my Graduate Committee, profs. Juan Reutter, Jorge Pérez, and Marcelo Guarini, I thank for their time in reviewing this work. I also thank my fellow postgraduate students for their companionship during these two years.

This work was done in collaboration with Anastasios Kementsietsidis, Kavitha Srinivas, and Achille Fokoue of IBM T.J. Watson Research Center. It was an honour to be able to work with them and I thank them for their guidance and for hosting me as an intern. Their excellence and passion for their work has been a great inspiration for me!

I feel extremely privileged to have been supervised by Marcelo Arenas, and I owe him an enormous debt of gratitude. Marcelo's patience, attention to detail, kindness and encouragement have been irreplaceable.

I would like to thank my family. Without Cecilia, Marco, and Ale, my education would not have been possible nor worth it. Finally, I thank Maca for always being by my side, for her love and support, and for sharing our adventure of living in U.K. together!

Contents

Acknowledgements	iv
List of Tables	vii
List of Figures	viii
Abstract	ix
Resumen	x
Chapter 1. Introduction	1
1.1. Summary of contributions	3
1.2. Related work	4
1.3. Organization of this document	6
Chapter 2. Preliminaries	7
2.1. A schema-oriented graph representation	7
2.2. Sample structuredness functions	8
2.2.1. The coverage function	8
2.2.2. The similarity function	9
2.2.3. The dependency functions	10
Chapter 3. A Language for Defining Structuredness Measures	11
3.1. Syntax of the language	11
3.2. Semantics of the language	12
Chapter 4. Sort Refinements and Signatures	15
Chapter 5. Formal Definition of the Decision Problem	19
Chapter 6. Reducing to Integer Linear Programming	21
6.1. Variable definitions	22

6.2.	Constraint definitions	23
6.3.	Implementation details	26
Chapter 7.	Experimental Results	28
7.1.	DBpedia Persons	29
7.1.1.	A highest θ sort refinement for $k = 2$	30
7.1.2.	A lowest k sort refinement for $\theta = 0.9$	31
7.1.3.	Dependency functions in DBpedia Persons	33
7.2.	For WordNet Nouns	35
7.2.1.	A highest θ sort refinement for $k = 2$	36
7.2.2.	A lowest k sort refinement for fixed θ	37
Chapter 8.	Conclusions	40
References	42
APPENDIX	44
APPENDIX A.	ADDITIONAL RESOURCES	45
A.1.	Proof of Theorem 5.0.3	45
A.1.1.	EXISTSSORTREFINEMENT(r) is in NP	45
A.1.2.	EXISTSSORTREFINEMENT(r_0) with $k = 3$ and $\theta = 1$ is NP-hard	46

List of Tables

7.1	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.	34
7.2	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.	35

List of Figures

2.1	Sample matrixes for datasets D_1 , D_2 and D_3	10
4.1	Signature view of the two real world datasets: (a) DBpedia Persons, and (b) WordNet Nouns. The datasets are depicted as horizontal tables, where the columns correspond to the properties of a dataset, the black regions correspond to data (non-null cells), and the white regions correspond to null cells. . .	18
7.1	DBpedia Persons split into $k = 2$ implicit sorts, using the structuredness functions (a) σ_{Cov} , (b) σ_{Sim} , and (c) σ_{Dep}	32
7.2	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}	33
7.3	WordNet Nouns split into $k = 2$ implicit sorts, using two different structuredness functions: (a) σ_{Cov} , and (b) σ_{Sim}	38
7.4	WordNet Nouns partitioned into the lowest k with a fixed threshold.	39

ABSTRACT

Although RDF graphs have schema information associated with them, in practice it is very common to find cases in which data do not fully conform to their schema. A prominent example of this is DBpedia, which is RDF data extracted from Wikipedia, a publicly editable source of information. In such situations, it becomes interesting to study the structural properties of the actual data, because the schema gives an incomplete description of the organization of a dataset.

In this work 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 or a database administrator 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 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 two real world datasets, DBpedia Persons and WordNet Nouns, and 4 different and intuitive rules, which gauge the structuredness in different ways. The rules give meaningful refinements of the datasets, showing that our language can be a powerful tool for understanding the structure of RDF data.

Keywords: RDF, Semantic Web

RESUMEN

Aunque grafos RDF tienen información de su esquema asociados a ellos, en la práctica es muy común encontrar situaciones en que los datos no se conforman totalmente a su esquema. Un ejemplo conspicuo es el de DBpedia, que son datos RDF extraídos desde Wikipedia, una fuente de información públicamente editable. En tales situaciones, se torna interesante estudiar las propiedades estructurales de los datos en sí, dado que el esquema de una descripción incompleta de la organización de una base de datos.

En este trabajo nos hemos acercado al estudio de la estructura de un grafo RDF desde primeros principios: proponemos un marco teórico para especificar funciones de estructura, que miden el grado de conformancia entre un grafo RDF y un esquema. En particular, primero se define un lenguaje formal para la especificación de funciones de estructura mediante expresiones que denominamos reglas. Este lenguaje permite a un usuario o a un administrador de una base de datos especificar una regla a la cual un grafo RDF puede conformarse de forma total o parcial. Luego, consideramos el problema de encontrar un refinamiento de un tipo (*sort*) mediante la partición de la base de datos en subconjuntos cuyo valor de estructura debe estar por sobre un valor crítico predeterminado. En particular, demostramos que el problema de decisión natural asociado a este problema de refinamiento es NP-completo, y diseñamos una reducción natural de este problema a una instancia de Programación Lineal Entera (ILP). Finalmente, ponemos esta solución ILP a prueba con dos bases de datos reales, Personas de DBpedia y Sustantivos de WordNet, y 4 reglas diferentes e intuitivas, que miden la estructura de formas diferentes. Las reglas resultan dar refinamientos razonables de las bases de datos, mostrando que nuestro lenguaje puede ser una herramienta potente para entender la estructura de los datos RDF.

Chapter 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 (Duan, Kementsietsidis, Srinivas, & Udreă, 2011). One example is the popular type of DBpedia persons (in this work, we will use the term *sort* as a synonym of *type*), which includes all the people having an entry in Wikipedia. According to the sort definition, each person in DBpedia has 8 properties, namely, a *name*, a *givenName*, a *surName*, a *birthDate*, a *birthPlace*, a *deathDate*, a *deathPlace*, and a *description*. There are currently 790,703 people 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, and only 323,368 have a birthplace, and for only 241,156 we have both pieces of information. There are approx. 40,000 people for whom we do not even know their last name. And when it comes to 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 that one needs to address is how to bridge the gap between these two worlds, the sorts and the respective data. In previous work (Duan et al., 2011), sorts are considered as being the unequivocal ground truth and methods are devised so as to make the data to 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 devise a sort refinement that better fits 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. 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 a 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 in this work several techniques enabling us to solve this problem in practice on real datasets as illustrated in our experimental evaluation section. Our first attack to 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 compressed representations 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 the compressed representation, 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 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.

Finally, 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 and characteristics of data.

1.1. Summary of contributions

Our key contributions in this work are fourfold:

- (i) We propose a language for specifying structuredness functions, which measure the degree to which an RDF graph conforms to a schema. This language has a boolean logic-like syntax and allows a user or a database administrator to specify a *rule* to which an RDF graph may fully or partially conform. The variables of a rule represent cells of the horizontal table corresponding to the RDF graph and the syntax of the language allows for specifying relationships between different subjects and properties of the RDF graph.
- (ii) 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. This constitutes the main theoretical result of our work. The NP-hardness of the decision problem we define is proven by means of a reduction from a known NP-complete problem (3-coloring of graphs, in this case).

- (iii) We provide a natural translation of an instance of the sort refinement problem into an ILP problem instance. This result allows us to tackle an otherwise difficult problem and the translation works for an *arbitrary rule* written in our language. This allows us to present a framework in which the description of rules is separated from the details of how to solve the sort refinement problem.
- (iv) Finally, we successfully test our ILP approach on two real world datasets and four different structuredness functions. The sort refinements produced by the commercial ILP solver give meaningful partitionings of the subjects of datasets studied.

1.2. 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 (Lee et al., 2013), "*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 (Duan et al., 2011) (Cov metric as per the notation of this work). 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 our algorithm may use overlapping sets of properties.

Similarly, (Ding, Wilkinson, Sayers, & Kuno, 2003) and (Levandovski & Mokbel, 2009) use frequent item set sequences (Agrawal & Srikant, 1995) 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 (Ding et al., 2003) and (Levandoski & Mokbel, 2009) 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 (Ding et al., 2003) and (Levandoski & Mokbel, 2009), various important parameters controlling the sort refinement approach are chosen in an ad-hoc manner (e.g., in (Ding et al., 2003) the minimum support used is chosen after manually inspecting partial results produced by an handful of minimum support values, and, in (Levandoski & Mokbel, 2009), it is explicit 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 (Völker & Niepert, 2011; d’Amato, Fanizzi, & Esposito, 2010; Lehmann, 2010; Delteil, Faron-Zucker, & Dieng, 2001; Maedche & Zacharias, 2002; Grimnes, Edwards, & Preece, n.d.). 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 (Vinh, Epps, & Bailey, 2010) 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 (Vinh et al., 2010).

1.3. Organization of this document

The remainder of the work is organized as follows. After a brief introduction of RDF data representation and a presentation of some examples of structuredness functions in Chapter 2, the syntax and the semantics of the language for specifying structuredness functions are formally defined in Chapter 3. In Chapter 4, we introduce the key concepts of signatures and sort refinements. After presenting the main computational complexity result of the sort refinement problem in Chapter 5, Chapter 6 describes the formulation of the problem as an ILP problem. In Chapter 7, we present the results of an experimental evaluation of our ILP based approach on two real world datasets (DBpedia Persons and WordNet Nouns). Finally, we conclude in Chapter 8.

Chapter 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 (Pan & Heflin, 2004), 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 $p \in P(D)$,

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

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:

DEFINITION 2.2.1. A *structuredness function* σ is a function which assigns to every RDF graph D a rational number $\sigma(D)$, such that $0 \leq \sigma(D) \leq 1$. \square

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 (Duan et al., 2011) 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 compressed graph data representation introduced in the previous section, the coverage metric of (Duan et al., 2011) can be expressed as follows: $\sigma_{\text{Cov}}(D) = (\sum_{sp} M(D)_{sp}) / |S(D)||P(D)|$. 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 2.1a. For this dataset, $\sigma_{\text{Cov}}(D_1) = 1$. Assume 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 2.1b. Then, the structuredness of $\sigma_{\text{Cov}}(D_2) \approx 0.5$ (for large N). This is because 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 $\text{total}(\varphi_1^{\text{Sim}}, D) = \{(s, s', p) \in S(D) \times S(D) \times P(D) \mid \varphi_1^{\text{Sim}} \text{ holds}\}$, and a set of favorable cases $\text{total}(\varphi_1^{\text{Sim}} \wedge \varphi_2^{\text{Sim}}, D) = \{(s, s', p) \in S(D) \times S(D) \times P(D) \mid \varphi_1^{\text{Sim}} \wedge \varphi_2^{\text{Sim}} \text{ holds}\}$ (note that $\text{total}(\varphi_1^{\text{Sim}} \wedge \varphi_2^{\text{Sim}}, D) \subseteq \text{total}(\varphi_1^{\text{Sim}}, D)$). 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 our example, 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 2.1c 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).

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

FIGURE 2.1. Sample matrixes for datasets D_1 , D_2 and D_3

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.

Chapter 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) combinations of these basic building components in the form of conjunctions, disjunctions and negations.

3.1. Syntax of the language

To define the syntax of the language, we need to introduce some terminology. From now on, assume that V is an infinite set of variables disjoint from U . We usually use c, c_1, c_2, \dots to denote the variables in 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 $(U \cup V)$. Then the set of terms in the language is defined as follows:

- $0, 1$, every $u \in U$ and every $c \in V$ is a term, and
- if $c \in V$, 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 $c \in V$ and $u \in U$, then $\text{val}(c) = 0$, $\text{val}(c) = 1$, $\text{prop}(c) = u$ and $\text{subj}(c) = u$ are formulas.
- If $c_1, c_2 \in V$, 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:

$$\varphi_1 \mapsto \varphi_2. \quad (3.1)$$

3.2. Semantics of the language

To define how rules of the form (3.1) 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 $\rho : V \rightarrow S(D) \times P(D)$ 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 $u \in U$, and $\rho(c) = (u, p)$.
- φ is the formula $\text{prop}(c) = u$, where $u \in U$, 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 (3.1). 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 .

Going back to the functions presented in Section 2.2, these 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 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 $\mathbf{p}_1, \mathbf{p}_2 \in \mathbf{U}$, we can also express the dependency measures. The $\sigma_{\text{Dep}}[\mathbf{p}_1, \mathbf{p}_2]$ structuredness measure can be expressed with the rule

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

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

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

Chapter 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 previous work (Duan et al., 2011), the data themselves were changed 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 $p, o \in U$) 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 $s, p_1, p_2, o_1, o_2 \in U$, 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 4.0.1. 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

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 $0 \leq \theta \leq 1$ (the threshold is required to be a rational number also, and this will be for compatibility with the reduction to the Integer Linear Programming instance).

DEFINITION 4.0.2. 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

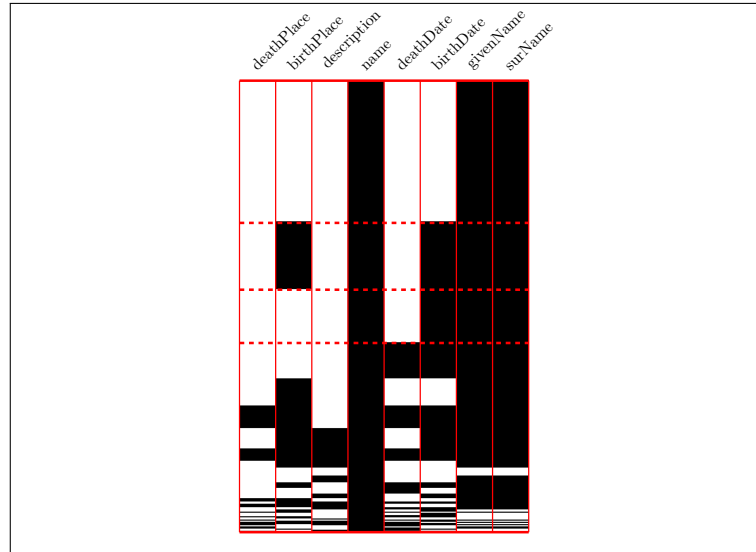
In the rest of this work, we will refer to the elements of the sort refinement (i.e. the elements of the partition of D) as implicit sorts.

We will now be concerned with discovering sort refinements in RDF data. We use the term signature in two ways: (i) to refer to the binary function described in Definition 4.0.1, and (ii) to refer to the set of all entities in an RDF graph D which share a common signature. In this second case, the *size* of a signature is the number of entities (or subjects) sharing that signature.

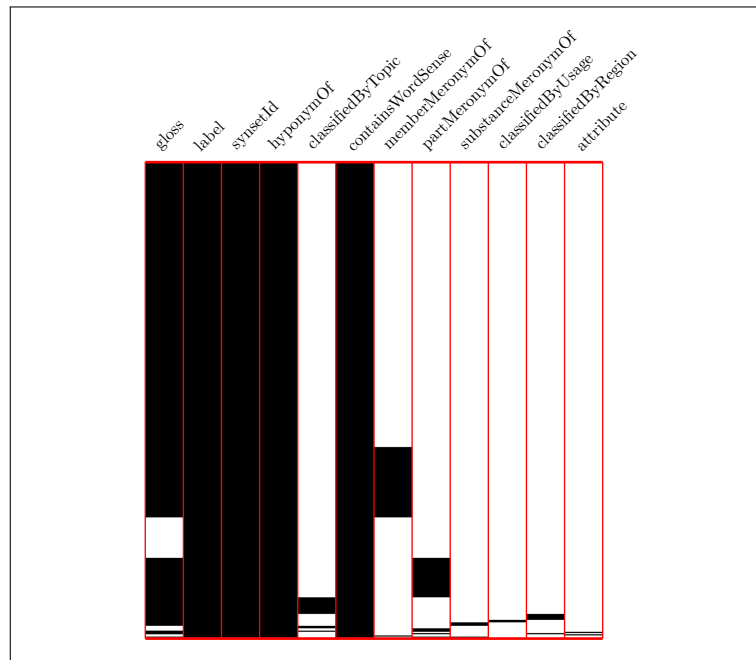
Figures 4.1a and 4.1b present a visual representation of an RDF graph's horizontal table. Every column represents a property and the rows have been grouped by signature,

in descending order of signature size. The first 3 signatures in figure 4.1a have been delimited with a dashed line, for clarity. The subsequent signatures 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. 4.1a) and WordNet Nouns (Fig. 4.1b) 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 for easy comparison the same number of columns, even if some columns are not present in a given subset of the sort refinement).



(A) The DBpedia Persons dataset has 790,703 subjects, 8 properties and 64 signatures. For this RDF graph, $\sigma_{Cov} = 0.54$, $\sigma_{Sim} = 0.77$.



(B) The WordNet Nouns dataset has 79,689 subjects, 12 properties and 53 signatures. For this RDF graph, $\sigma_{Cov} = 0.44$, $\sigma_{Sim} = 0.93$.

FIGURE 4.1. Signature view of the two real world datasets: (a) DBpedia Persons, and (b) WordNet Nouns. The datasets are depicted as horizontal tables, where the columns correspond to the properties of a dataset, the black regions correspond to data (non-null cells), and the white regions correspond to null cells.

Chapter 5. FORMAL DEFINITION OF THE DECISION PROBLEM

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

Problem: 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.

We pinpoint the complexity of the problem EXISTSSORTREFINEMENT(r) in the following theorem:

THEOREM 5.0.3.

- EXISTSSORTREFINEMENT(r) is in NP for every rule r .
- There is a rule r_0 for which 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 5.0.3 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 5.0.3 shows that there exists a (fixed) rule r_0 for which 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 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 corresponding RDF graph.

Although the rule r_0 will not be shown explicitly here, 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 in G . 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 (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. 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, 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 work, 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.

Chapter 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 will 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 will 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, the ILP instance we will define works in the following way: the integer variables will decide which signatures are to be included in which subsets, and they will keep track of which properties are used in each subset. Also, we will group variable assignments into objects we call *rough* variable assignments, which instead of assigning each variable to a subject and a property will assign each variable to a signature and a property. In this way, another set of variables will keep track of which rough assignments are valid in a given subset (i.e. the rough assignment mentions only signatures and properties which are present in the subset). With the previous, we will be 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 variables of the ILP instance. Recall that our goal when solving $\text{EXISTSSORTREFINEMENT}(r)$ is to find a σ_r -sort refinement of D with threshold θ that contains at most k implicit sorts.

All the variables used in the ILP instance will take only integer values. First, we define the set of signatures of D as $\Lambda(D) = \{\text{sig}(s, D) \mid s \in S(D)\}$, and for every $\mu \in \Lambda(D)$, we define 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 of the ILP instance, as they encode the sort refinement which is generated. 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 will consider a *rough* assignment of variables in φ to be a mapping of each variable to a signature and a property. We shall denote rough assignments with $\tau = ((\mu_1, p_1), \dots, (\mu_n, p_n)) \in (\Lambda(D) \times P(D))^n$. Then for each $i \in$

$\{1, \dots, k\}$ and each $\tau \in (\Lambda(D) \times P(D))^n$ 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:

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

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 set of inequalities that constrain the acceptable values of the previously defined variables.

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

$$0 \leq X_{i,\mu} \leq 1 \quad i \in \{1, \dots, k\} \text{ and } \mu \in \Lambda(D)$$

$$0 \leq U_{i,p} \leq 1 \quad i \in \{1, \dots, k\} \text{ and } p \in P(D)$$

$$0 \leq T_{i,\tau} \leq 1 \quad i \in \{1, \dots, k\} \text{ and } \tau \in (\Lambda(D) \times P(D))^n$$

- For every $\mu \in \Lambda(D)$, the following equation is used to indicate 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 $\mu(p) = 1$ ($p \in \text{supp}(\mu)$):

$$X_{i,\mu} \leq U_{i,p} \quad \text{if } p \in \text{supp}(\mu)$$

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

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.

From the definition of the ILP instance, it is easy to see that the following result holds.

PROPOSITION 6.2.1. 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 are enough to solve the decision problem, in practice the search space is too large to be manageable because of 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 will 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 therefore 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 in the case of DBpedia Persons) as large exponent values will 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.

Chapter 7. EXPERIMENTAL RESULTS

For our experiments, we consider two real datasets: DBpedia Persons and WordNet Nouns. With each, we consider 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 a 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.

A final comment on the general experimental settings is in order: 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 structuredness functions σ_{Cov} , σ_{Sim} , and σ_{Dep} (the last function is only used on DBpedia Persons). All experiments are done on a machine with two 2.3 GHz processors, each with 6 cores, and 64 GB of RAM. The ILP solver used is IBM ILOG CPLEX version 12.5.

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

$p_2 = \text{deathDate}$, for example, then the value of the function $\sigma_{\text{SymDep}}[\text{deathPlace}, \text{deathDate}]$ is 0.39. This specific choice of p_1 and 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 7.1a 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 experimental setup, each ILP instanced is solved in under 800 ms.

Figure 7.1b shows the results for the σ_{Sim} function. In this case, 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 the σ_{Sim} function in the first place). Also, notice that unlike the σ_{Cov} function, the cardinality of the generetic sorts from σ_{Sim} is more *balanced*. In this experimental setup each ILP instance is solved in under 2 minutes, except the infeasible instance (the last instance to be solved), which was completed in 2 hrs.

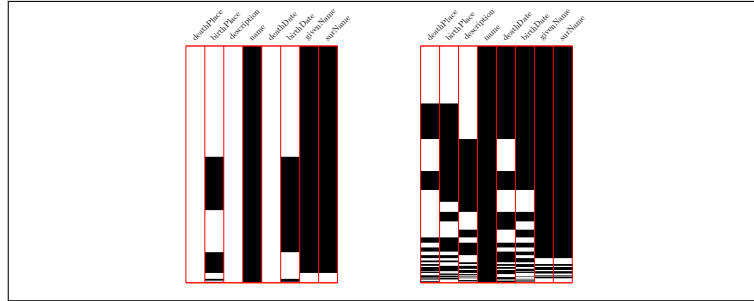
Finally, Figure 7.1c shows the results for $\sigma_{\text{SymDep}}[\text{deathPlace}, \text{deathDate}]$, a structuredness function in which we measure the probability that, if a subject has a `deathPlace` or a `deathDate`, it has both. In the resulting sort refinement, the second sort to the right has a high value of 0.82. From figure 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. This is because the absence of the said column eliminates all total cases (i.e. there are no assignments of variables in the rule that represents $\sigma_{\text{SymDep}}[\text{deathPlace}, \text{deathDate}]$ for which the antecedent is true, because it is never true that $\text{prop}(c_1) = \text{deathPlace}$). This setting is completed in under 1 minute.

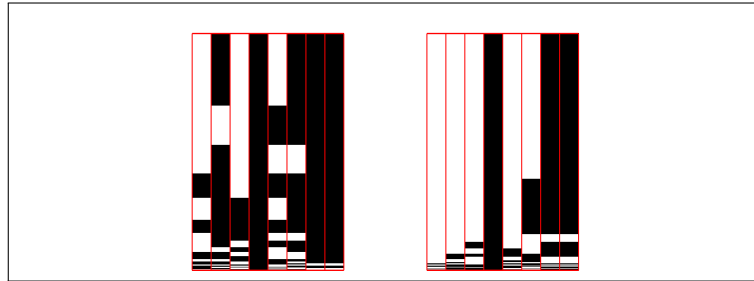
7.1.2. A lowest k sort refinement for $\theta = 0.9$

We now consider a fixed threshold $\theta = 0.9$. We seek the smallest sort refinement for DBpedia persons with this threshold. Figure 7.2a 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 setup was completed in a total of 30 minutes.

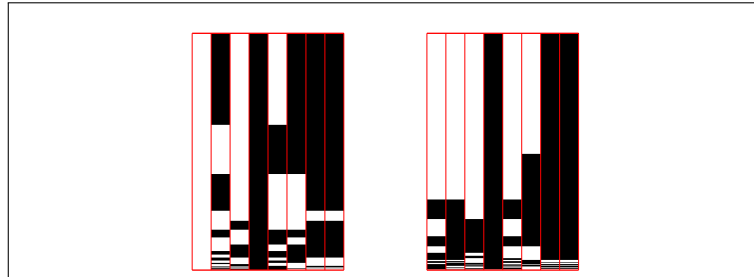
Figure 7.2b shows the result for σ_{Sim} , where the optimum value found is for $k = 4$. Again, the function is more lenient when properties appear for only a small amount of subjects (hence the smaller k). This is clearly evident in the first sort for this function, which



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




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




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

FIGURE 7.1. DBpedia Persons split into $k = 2$ implicit sorts, using the structuredness functions (a) σ_{Cov} , (b) σ_{Sim} , and (c) σ_{Dep} .

corresponds roughly to the second sort generated for the σ_{Cov} function (Fig. 7.2a) but also includes a small number of subjects with birth/death places/dates. This is also verified 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 experimental setup is clearly more difficult as the running time of individual ILP instances is apx. 8 hours.



(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 $\sigma_{\text{Cov}}(D_i) \geq 0.9$). 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 7.2. 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} .

7.1.3. Dependency functions in DBpedia Persons

We now turn our attention to the dependency functions. In terms of creating a new sort refinement using the function $\sigma_{\text{Dep}}[\mathbf{p}_1, \mathbf{p}_2]$, 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

	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 7.1. 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.

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 $\sigma_{\text{Dep}}[\mathbf{p}_1, \mathbf{p}_2]$ function, and we tabulate (in Table. 7.1) 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

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 7.2. 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.

(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 $\sigma_{\text{SymDep}}[p_1, p_2]$ function. In Table 7.2 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`, respectively.

7.2. For WordNet Nouns

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

NounSynset):

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

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 list of properties is 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. 4.1b); 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 the Sim rule.

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

As mentioned, the WordNet case proves to be very different from DBpedia Persons partly because in this dataset there are roughly 5 dominant signatures which represent a large portion of the subjects, and yet only use 8 of the 12 properties. This causes difficulties when partitioning the dataset into 2 sorts.

Figure 7.3a shows the result for σ_{Cov} . The most notable difference between both sorts is that the left sort mostly consists of subjects which have the memberMeronymOf property (the seventh 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 mostly due to the presence of many of signatures which represent very few subjects, and

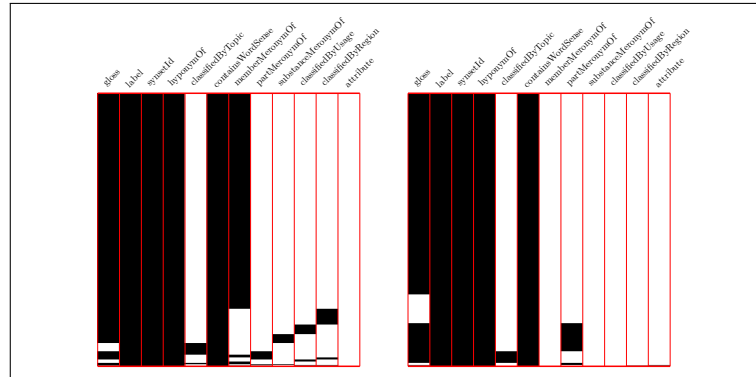
have different sets of properties. For this setup, all ILP instances were solved in under 1 second.

Figure 7.3b shows the result for σ_{Sim} . In this case, the clear 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, since the Sim function is not sensitive to their presence. Although the structuredness is high for this partitioning, the improvement is not large, since the original dataset is highly structured with respect to σ_{Sim} anyway. A discussion is in order with respect to the running times of this experiment. 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 seconds. 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 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 up to that moment.

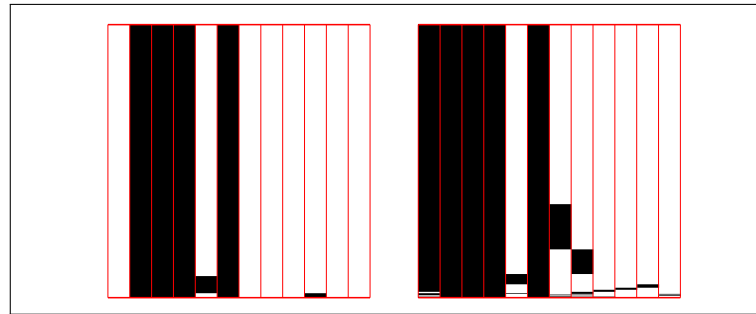
7.2.2. A lowest k sort refinement for fixed θ

As with the previous experimental setup, WordNet 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 the threshold is 0.9. For that reason, in this last case we fix the threshold at 0.98.

Figure 7.4a shows the first 10 sorts of the $k = 31$ solution for σ_{Cov} . The sheer amount of sorts needed is a 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



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




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

FIGURE 7.3. WordNet Nouns split into $k = 2$ implicit sorts, using two different structuredness functions: (a) σ_{Cov} , and (b) σ_{Sim} .

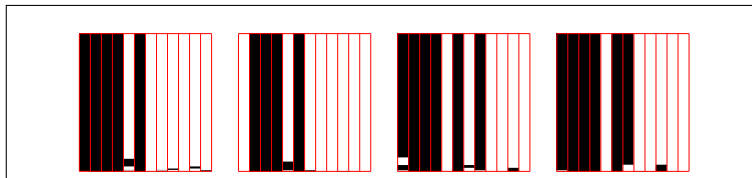
large number is another indication of the difficulty of partitioning a dataset with highly uniform entities.

Figure 7.4b 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. This setup was completed in apx. 15 minutes.

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 7.4. WordNet Nouns partitioned into the lowest k with a fixed threshold.

Chapter 8. CONCLUSIONS

The flexibility of the RDF language makes it a good candidate for storing data from diverse domains, as is the case of the Semantic Web vision. However, it is precisely this flexibility which permits the divergence of data and schema. As RDF and other similar languages proliferate, it will become increasingly important to have tools to study the structure of data and understand how they conform (or fail to conform) to their declared schemas.

In this work 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. Along with this language, we have proposed several intuitive rules which gauge the structuredness of an RDF graph in different ways. The language of structuredness rules can be used in a range of situations; we have considered the problem of discovering a partitioning of the entities of an RDF graph into subsets 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 Integer Linear Programming instance capable of solving this problem within reasonable time limits using commercially available ILP solvers. Most importantly, this ILP instance can be generated for an arbitrary rule expressed in our language.

We have used our framework to study two real world RDF datasets, namely DBpedia Persons and WordNet Nouns, the former ultimately 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 experimental results obtained were meaningful and intuitive. In the case of DBpedia Persons, the structuredness functions (defined via rules) which were used clearly indicate that the sort Person is too broad, and this can be inferred in two ways: (i) the initial RDF graph of DBpedia Persons has a low structuredness value under the structuredness functions, and

(ii) the sort refinement problem (via its ILP instance) was able to find a partitioning of DBpedia Persons which have higher structuredness and which are readily interpretable. For example, we may conclude that defining a subsort of Live People would allow for a more precise description of the data, since a subsort was created in which no subject has information about their deaths. Our framework, then, proves to be a powerful tool for producing a more accurate description of the structure of an RDF graph.

The obvious next goal is to better understand the expressiveness of structuredness rules (i.e. to explore which structuredness functions cannot be expressed in our language, or to find properties shared by functions which can be expressed in our language). Also, an extremely interesting goal would be to explore the existence of rules for which a high structuredness value can predict good performance for certain classes of queries. Since the evaluation of the structuredness value of an RDF graph under a certain rule is very efficiently computable, such a query performance predicting rule may open avenues of research on query optimization.

References

Agrawal, R., & Srikant, R. (1995). Mining sequential patterns. In *Proceedings of the eleventh international conference on data engineering* (pp. 3–14). Washington, DC, USA: IEEE Computer Society. Retrieved from <http://dl.acm.org/citation.cfm?id=645480.655281>

d’Amato, C., Fanizzi, N., & Esposito, F. (2010, April). Inductive learning for the semantic web: What does it buy? *Semant. web*, *1*(1,2), 53–59. Retrieved from <http://dl.acm.org/citation.cfm?id=2019445.2019452>

Delteil, A., Faron-Zucker, C., & Dieng, R. (2001). Learning Ontologies from RDF annotations. In A. Maedche, S. Staab, C. Nedellec, & E. H. Hovy (Eds.), *Ijcai 2001 workshop on ontology learning, proceedings of the second workshop on ontology learning ol 2001, seattle, usa, august 4, 2001 (held in conjunction with the 17th international conference on artificial intelligence ijcai 2001)* (Vol. 38). CEUR-WS.org.

Ding, L., Wilkinson, K., Sayers, C., & Kuno, H. (2003). Application-specific schema design for storing large rdf datasets. In *In first intl workshop on practical and scalable semantic systems*.

Duan, S., Kementsietsidis, A., Srinivas, K., & Udrea, O. (2011). Apples and oranges: a comparison of rdf benchmarks and real rdf datasets. In *Sigmod conference* (p. 145-156).

Grimnes, G. A., Edwards, P., & Preece, A. (n.d.). Learning Meta-Descriptions of the FOAF Network. In *Proceedings of the Third International Semantic Web Conference (ISWC-04)* (p. 152-165). Hiroshima, Japan: Springer Verlag.

Lee, T. Y., Cheung, D. W., Chiu, J., Lee, S. D., Zhu, H., Yee, P., & Yuan, W. (2013). *Automating relational database schema design for very large semantic datasets* (Tech. Rep.). Department of Computer Science, University of Hong Kong.

Lehmann, J. (2010). *Learning owl class expressions*. Unpublished doctoral dissertation.

Levandoski, J. J., & Mokbel, M. F. (2009). Rdf data-centric storage. In *Proceedings of the 2009 IEEE international conference on web services* (pp. 911–918). Washington, DC, USA: IEEE Computer Society. Retrieved from <http://dx.doi.org/10.1109/ICWS.2009.49> doi: 10.1109/ICWS.2009.49

Maedche, A., & Zacharias, V. (2002). Clustering ontology-based metadata in the semantic web. In T. Elomaa, H. Mannila, & H. T. T. Toivonen (Eds.), *13th european conference on machine learning (ecml'02) 6th european conference on principles and practice of knowledge discovery in databases (pkdd'02)*. Helsinki, Finland.

Pan, Z., & Heflin, J. (2004). *Dldb: Extending relational databases to support semantic web queries* (Tech. Rep.). Department of Computer Science, Lehigh University.

Vinh, N. X., Epps, J., & Bailey, J. (2010, December). Information theoretic measures for clusterings comparison: Variants, properties, normalization and correction for chance. *J. Mach. Learn. Res.*, 9999, 2837–2854. Retrieved from <http://dl.acm.org/citation.cfm?id=1953011.1953024>

Völker, J., & Niepert, M. (2011). Statistical schema induction. In *Proceedings of the 8th extended semantic web conference on the semantic web: research and applications - volume part i* (pp. 124–138). Berlin, Heidelberg: Springer-Verlag. Retrieved from <http://dl.acm.org/citation.cfm?id=2008892.2008904>

APPENDIX

APPENDIX A. ADDITIONAL PROOFS

A.1. Proof of Theorem 5.0.3

Recall the definition of the decision problem, where r is a rule:

Problem: 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 a σ_r -sort refinement \mathcal{T} of D with threshold θ that contains at most k implicit sorts, and **false** otherwise.

We will now prove that EXISTSSORTREFINEMENT(r_0) is NP-complete for $\theta = 1$ and $k = 3$, where r_0 is the rule defined in equation A.1. For this, we will first prove that the general problem EXISTSSORTREFINEMENT(r) is in NP. Then we will show that fixing $\theta = 1$ and $k = 3$ and using rule r_0 yields a version of EXISTSSORTREFINEMENT(r) that is NP-hard.

A.1.1. EXISTSSORTREFINEMENT(r) is in NP

Given an RDF graph D , a rational number θ such that $0 \leq \theta \leq 1$, and a positive integer k , a Non-Deterministic Turing Machine (NDTM) must guess T_1, \dots, T_l with $l \leq k$. The NDTM must then verify that T_1, \dots, T_l form a sort refinement with threshold θ . For each T_i , $i \in [1, l]$, the structuredness can be determined in the following way: Consider that T_i contains S_i subjects (rows, if represented as a matrix) and P_i properties (columns). Let v be the number of variables in r_0 . Since rule r_0 is fixed, v is also fixed. There are $(S_i P_i)^v$ different possible assignments of the variables in r_0 , which is polynomial in the size of D .

For each assignment, the NDTM must check if (i) it satisfies the antecedent of r_0 and (ii) if it satisfies the antecedent and the consequent of r_0 , together. Since both checks are similar, we will focus on checking if the antecedent of r_0 is satisfied by an assignment ρ .

The antecedent of r_0 will consist of a number of equalities (or inequalities) bounded by the size of r_0 (which is itself fixed). If c_a and c_b are variables in r_0 , each equality will be of the form $x = y$, where x may be c_a , $\text{subj}(c_a)$, $\text{prop}(c_a)$, or $\text{val}(c_a)$, and y will be defined accordingly, obeying the syntax of the rules (defined previously). The NDTM can encode the assignment ρ in the following way: each of the v variables can be encoded using $\log(v)$ bits, and for each variable it must store the subject and property it is assigned to, using $\log(S_i)$ and $\log(P_i)$ bits, respectively. Since S_i and P_i are both bounded by the size $|D|$ of D , the assignment ρ can be encoded in space which is logarithmic in $|D|$.

To determine the value of a variable, the NDTM must find the subject-property pair in the encoded version of the RDF graph D . This lookup will take time at most linear in $|D|$. The comparison itself can be done polynomial time.

The entire verification process will take time which is polynomial in the size of the RDF graph D , although it depends non-trivially on the size of the rule r . Therefore, $\text{EXISTSSORTREFINEMENT}(r)$ is in NP.

A.1.2. $\text{EXISTSSORTREFINEMENT}(r_0)$ with $k = 3$ and $\theta = 1$ is NP-hard

To prove that $\text{EXISTSSORTREFINEMENT}(r_0, 1, 3)$ is NP-hard, we will use a reduction from 3-COLORABILITY, which is the following decision problem: given an undirected graph without loops (self-edges) $G = (V, E)$, decide if there exists a 3-coloring of G (i.e. a function $f : V \rightarrow \{1, 2, 3\}$ such that for all pairs of nodes $u, v \in V$, if $(u, v) \in E$ then $f(u) \neq f(v)$).

Consider a non-directed graph G with n nodes, defined by the $n \times n$ adjacency matrix A_G . We will construct an RDF graph D_G defined by its accompanying matrix $M_G = M(D_G)$ so that G is 3-colorable if and only if there exists a σ_{r_0} -sort refinement \mathcal{T} of D with threshold 1, consisting of at most 3 implicit sorts.

First, we construct the accompanying matrix M_G of D_G by blocks, in the following way:

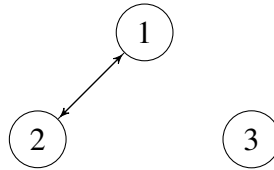
$$M_G = \begin{pmatrix} 0_{n \times 1} & 0_{n \times 1} & 1_{n \times 1} & D_{n \times n} & D_{n \times n} \\ 0_{n \times 1} & 1_{n \times 1} & 1_{n \times 1} & D_{n \times n} & D_{n \times n} \\ 1_{n \times 1} & 0_{n \times 1} & 1_{n \times 1} & D_{n \times n} & D_{n \times n} \\ 1_{n \times 1} & 1_{n \times 1} & 0_{n \times 1} & D_{n \times n} & \bar{A}_G \end{pmatrix}_{(4n) \times (2n+3)}$$

Here, $D_{n \times n}$ is a $n \times n$ unit matrix (i.e. with 1's in the diagonal and 0's everywhere else), $0_{n \times 1}$ is a single column of n zeroes, and $1_{n \times 1}$ is a single column of n ones. The block \bar{A}_G contains the complement of the adjacency matrix A_G ($\bar{A}_G[i, j] = 1 - A_G[i, j]$).

The first $3n$ rows will be referred to as the *upper section* of M_G , while the last n rows will be referred to as the *lower section* of M_G . The upper section consists of three sets of auxiliary rows, which are identical in every column except the first and second. The first 2 columns will be called **sp1** and **sp2**, respectively, the third column will be called the **idp** column, the next n columns will be referred to as the *left column set* and the last n rows as the *right column set*. In this way, the complemented adjacency matrix \bar{A}_G is contained in the right column set, in the lower section of M_G .

Every row of M_G represents a subject and every column represents a property, however, we will not explicitly mention the names of any subjects or properties, except for the first three properties: **sp1**, **sp2**, and **idp**.

EXAMPLE A.1.1. Consider the following simple graph G , its adjacency matrix A_G and the complemented matrix \bar{A}_G :



$$A_G = \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \quad \bar{A}_G = \begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & 1 \\ 1 & 1 & 1 \end{pmatrix}$$

$$\left(\begin{array}{l}
\text{prop}(c_1) \neq \mathbf{sp1} \wedge \text{prop}(c_1) \neq \mathbf{sp2} \wedge \\
\text{prop}(c_2) \neq \mathbf{sp1} \wedge \text{prop}(c_2) \neq \mathbf{sp2} \wedge \\
\text{prop}(d_1) \neq \mathbf{sp1} \wedge \text{prop}(d_1) \neq \mathbf{sp2} \wedge \\
\text{prop}(d_2) \neq \mathbf{sp1} \wedge \text{prop}(d_2) \neq \mathbf{sp2} \wedge \\
\text{prop}(e) \neq \mathbf{sp1} \wedge \text{prop}(e) \neq \mathbf{sp2} \wedge \\
\text{prop}(f_1) \neq \mathbf{sp1} \wedge \text{prop}(f_1) \neq \mathbf{sp2} \wedge \\
\text{prop}(f_2) \neq \mathbf{sp1} \wedge \text{prop}(f_2) \neq \mathbf{sp2} \wedge \\
\text{prop}(x) = \mathbf{idp} \wedge \text{val}(x) = 1 \wedge \\
c_1 \neq x \wedge \text{subj}(c_1) = \text{subj}(x) \wedge \text{val}(c_1) = 1 \wedge \\
c_2 \neq x \wedge \text{subj}(c_2) = \text{subj}(x) \wedge \text{val}(c_2) = 1 \wedge \\
c_1 \neq c_2 \wedge \\
\text{prop}(y) = \mathbf{idp} \wedge \text{val}(y) = 0 \wedge \\
\text{subj}(d_1) = \text{subj}(y) \wedge \text{prop}(d_1) = \text{prop}(c_1) \wedge \\
\text{subj}(d_2) = \text{subj}(y) \wedge \text{prop}(d_2) = \text{prop}(c_2) \wedge \\
\text{prop}(z) = \mathbf{idp} \wedge \text{subj}(z) = \text{subj}(e) \wedge \\
\text{prop}(e) = \text{prop}(c_1) \wedge e \neq c_1 \wedge \text{val}(e) = 1 \wedge \\
\text{prop}(u) = \mathbf{idp} \wedge \text{val}(u) = 0 \wedge \\
\text{subj}(u) = \text{subj}(f_1) \wedge \text{prop}(f_1) = \text{prop}(c_1) \wedge \\
\text{subj}(u) = \text{subj}(f_2) \wedge \text{prop}(f_2) = \text{prop}(c_2) \wedge \\
\text{val}(f_1) = 1 \wedge \text{val}(f_2) = 1
\end{array} \right) \mapsto \begin{array}{l}
(\text{val}(d_1) = 1 \vee \text{val}(d_2) = 1) \wedge \\
\text{val}(z) = 0
\end{array}$$

(A.1)

Rule r_0 is constructed as follows:

(i) The subexpression

$$\begin{array}{l}
\text{prop}(c_1) \neq \mathbf{sp1} \wedge \text{prop}(c_1) \neq \mathbf{sp2} \wedge \\
\text{prop}(c_2) \neq \mathbf{sp1} \wedge \text{prop}(c_2) \neq \mathbf{sp2} \wedge \\
\text{prop}(d_1) \neq \mathbf{sp1} \wedge \text{prop}(d_1) \neq \mathbf{sp2} \wedge \\
\text{prop}(d_2) \neq \mathbf{sp1} \wedge \text{prop}(d_2) \neq \mathbf{sp2} \wedge \\
\text{prop}(e) \neq \mathbf{sp1} \wedge \text{prop}(e) \neq \mathbf{sp2} \wedge \\
\text{prop}(f_1) \neq \mathbf{sp1} \wedge \text{prop}(f_1) \neq \mathbf{sp2} \wedge \\
\text{prop}(f_2) \neq \mathbf{sp1} \wedge \text{prop}(f_2) \neq \mathbf{sp2}
\end{array}$$

ensures that no variables can be assigned to columns $\mathbf{sp1}$ or $\mathbf{sp2}$ (i.e. the set of total cases will only consider assignments where the variables are mapped to

cells whose columns are not **sp1** or **sp2**). It is not necessary to include variables x, y, z and u here, as their columns are fixed.

- (ii) The subexpression $\text{prop}(x) = \mathbf{idp} \wedge \text{val}(x) = 1$ forces variable x to point to a cell in the upper section of the **idp** column.
- (iii) The subexpression $c_1 \neq x \wedge \text{subj}(c_1) = \text{subj}(x) \wedge \text{val}(c_1) = 1 \wedge c_2 \neq x \wedge \text{subj}(c_2) = \text{subj}(x) \wedge \text{val}(c_2) = 1 \wedge c_1 \neq c_2$ defines variables c_1 and c_2 , which share their row with x . Both must point to an element of the diagonal of the $D_{n \times n}$ blocks (since their value must be 1) and since x, c_1 , and c_2 are all distinct, c_1 and c_2 must each point to a different $D_{n \times n}$ block. From here on, and without loss of generality, we will assume c_1 points to the left column set and c_2 points to the right column set.
- (iv) The subexpression $\text{prop}(y) = \mathbf{idp} \wedge \text{val}(y) = 0 \wedge \text{subj}(d_1) = \text{subj}(y) \wedge \text{prop}(d_1) = \text{prop}(c_1) \wedge \text{subj}(d_2) = \text{subj}(y) \wedge \text{prop}(d_2) = \text{prop}(c_2)$ defines variables y, d_1 , and d_2 . Variable y will point to the lower section of the **idp** column, and d_1 will point to a cell whose row and column are determined by y and c_1 , respectively. On the other hand, d_2 will point to a cell whose row and column are determined by y_2 and c_2 , respectively. With this configuration (and the assumption given in the previous item), d_1 will point to the $D_{n \times n}$ in the left column set, lower section, and d_2 will point to the \bar{A}_G block.
- (v) The subexpression $\text{prop}(z) = \mathbf{idp} \wedge \text{subj}(z) = \text{subj}(e) \wedge \text{prop}(e) = \text{prop}(c_1) \wedge e \neq c_1 \wedge \text{val}(e) = 1$ defines z , which points to the **idp** column, and e , whose position is fixed by the row of z and column of c_1 . Note that the rows of z and e are not fixed by the antecedent of the rule.
- (vi) The subexpression $\text{prop}(u) = \mathbf{idp} \wedge \text{val}(u) = 0 \wedge \text{subj}(u) = \text{subj}(f_1) \wedge \text{prop}(f_1) = \text{prop}(c_1) \wedge \text{subj}(u) = \text{subj}(f_2) \wedge \text{prop}(f_2) = \text{prop}(c_2)$ defines u , which points to the lower section of the **idp** column, f_1 is fixed by the row of u and the column of c_1 , and f_2 is analogous to f_1 .

- (vii) The subexpression $\text{val}(f_1) = 1 \wedge \text{val}(f_2) = 1$ ensures that the values of f_1 and f_2 are 1 (in an assignment that is to be included in the set of total cases).
- (viii) The subexpression $(\text{val}(d_1) = 1 \vee \text{val}(d_2) = 1) \wedge \text{val}(z) = 0$ constitutes the consequent of the rule and states the additional conditions that must be met by an assignment ρ for it to be included in the set of favorable cases.

EXAMPLE A.1.2. We can visualize the positioning of the variables assigned to cells in M_G :

$$\left(\begin{array}{c|c|c} & x & c_1 & c_2 \\ & z & e & \\ \hline & y & d_1 & d_2 \\ & u & f_1 & f_2 \end{array} \right)$$

The (x, c_1, c_2) trio must be located in the upper section as explained previously (items 2 and 3). The (y, d_1, d_2) trio and the (u, f_1, f_2) trio must be located in the lower section (items 4 and 6, respectively). The (z, e) duo is shown in the upper section, although it may be assigned to the lower section also (item 5).

□

Now that RDF graph D_G and rule r_0 have been defined, we will give an intuition as to how the existence of an implicit σ_{r_0} -sort refinement \mathcal{T} of D_G with threshold 1 implies that graph G is 3-colorable.

An implicit σ_{r_0} -sort T of D_G with threshold 1 can be understood as a subset of the rows of M_G , which themselves form an RDF graph which we will call D_T , with accompanying matrix $M_T = M(D_T)$. We must be careful, however, with the following condition: an implicit σ_{r_0} -sort T must be closed under signatures, that is, if a row is present in M_T , then all other identical rows must be included in M_T as well. This condition has been

made trivial with the creation of columns **sp1** and **sp2**, whose sole purpose is to ensure that there are no two identical rows in matrix M_G . Since each row of matrix M_G has its own unique signature, we have no need to discuss signatures for this reduction.

The problem of deciding if there exists an implicit σ_{r_0} -sort refinement \mathcal{T} of D_G with threshold 1 and with at most 3 sorts is the problem of partitioning the rows of M_G into at most 3 implicit σ_{r_0} -sorts T_1, T_2 , and T_3 such that each T_i satisfies $\sigma_{r_0}(T_i) = 1$ (i.e. the value of the structuredness of T_i is 1 when using r_0).

EXAMPLE A.1.3. Given matrix M_G of our working example, we show a possible partitioning of rows:

$$M_{T_1} = \left(\begin{array}{cc|c|c|cc} 0 & 0 & 1 & 1 & & 1 \\ 0 & 0 & 1 & & 1 & \\ 0 & 0 & 1 & & & 1 \\ \hline 1 & 1 & & 1 & 1 & 0 & 1 \\ 1 & 1 & & & 1 & 1 & 1 \end{array} \right), \quad M_{T_2} = \left(\begin{array}{cc|c|c|cc} 0 & 1 & 1 & 1 & & 1 \\ 0 & 1 & 1 & & 1 & \\ 0 & 1 & 1 & & 1 & \\ \hline 1 & 1 & & 1 & 0 & 1 & 1 \end{array} \right)$$

$$M_{T_3} = \left(\begin{array}{cc|c|c|cc} 1 & 0 & 1 & 1 & & 1 \\ 1 & 0 & 1 & & 1 & \\ 1 & 0 & 1 & & & 1 \end{array} \right)$$

Note that M_{T_1} includes two rows in the lower section which represent the set $\{1, 3\}$ of nodes of G and, analogously, M_{T_2} includes rows which represent the set $\{2\}$. In this case, each implicit σ_{r_0} -sort T_i has included one set of auxiliary rows from the upper section of matrix M_G . Also, the first sort, T_1 , has included the first and the third row from the lower section of M_G , while sort T_2 has included the second row from the lower section of M_G . Sort T_3 has not included any additional rows. We will later see that this choice of sorts represents a possible partitioning of the nodes of graph G into independent sets.

□

We will now see how an implicit σ_{r_0} -sort refinement \mathcal{T} of D_G with threshold 1 partitions the rows of M_G in such a way that each sort represents an independent set of graph G . Given an implicit sort T , we will explain how the rows from the lower section which

have been included represent nodes of graph G . These nodes form an independent set in G if and only if the structuredness of sort T under the structuredness function σ_{r_0} is 1.

For every pair of nodes included in T , r_0 must check that they are not connected, using the adjacency submatrix \bar{A}_G . The simplest way to do this would be to select a row from the lower section to represent the first node, a column from the right column set to represent the second node, and to check that the appropriate cell in \bar{A}_G is set to 1 (recall that \bar{A}_G is the complement of the adjacency matrix A_G of graph G). However, when building the implicit sort T for each row, all columns are included (i.e. full rows are included). The effect of this is that, while the rows included in M_T only represent a subset of the nodes in G , all nodes of G are represented in the columns of M_T . Rule r_0 must first ensure that only nodes included in T will be compared.

We now briefly review the role of the different variables in r_0 . The (z, e) pair serves to ensure that only one copy of the auxiliary rows is present in each implicit sort. To illustrate this, consider that, in a given assignment ρ , the triple (x, c_1, c_2) will occupy an auxiliary row. If this auxiliary row is duplicated in T , then it is possible to assign (z, e) to the duplicate auxiliary row. This assignment will satisfy the antecedent of r_0 , but will not satisfy the consequent, since $\text{val}(z) = 0$ will not hold. This will cause the structuredness of T to be less than 1 (see example).

EXAMPLE A.1.4. Using the same working example, consider the following subset T of D_G where two copies of an auxiliary row have been included (more precisely, the two copies of the auxiliary rows which are included are equal in every column except **sp1** and **sp2**):

$$M_T = \left(\begin{array}{cc|cc|c} 0 & 1 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & 1 & & 1 & 1 \\ 0 & 0 & 1 & & 1 & & 1 \\ \hline 1 & 1 & & 1 & 1 & 0 & 1 \\ 1 & 1 & & & 1 & 1 & 1 \end{array} \right),$$

A possible assignment ρ of the variables is shown as superscripts (note that only the variables relevant to the example are shown):

$$\left(\begin{array}{cc|cc|c} 0 & 1 & 1^z & 1^e & 1 \\ 0 & 0 & 1^x & 1^{c_1} & 1^{c_2} \\ 0 & 0 & 1 & & 1 \\ 0 & 0 & 1 & & 1 \\ \hline 1 & 1 & 0 & 1 & 1 & 0 & 1 \\ 1 & 1 & & & 1 & 1 & 1 \end{array} \right)$$

The assignment shown will cause the structuredness to be less than 1. This is because the assignment satisfies the antecedent but not the consequent, as $\text{val}(z) = 0$ is false. This example illustrates how variables (z, e) serve to ensure only one set of auxiliary rows can be included in a sort.

□

Next, the (u, f_1, f_2) triple ensures that c_1 (and, by association, c_2) is assigned to a column of M_T which represents a node of G that is also represented by a row of M_T . To understand this better, recall that a sort T of D_G is represented by its matrix M_T , which is itself built by selecting a subset of the rows of M_G . When including a row from the lower section of M_G in M_T , we are also selecting the appropriate node of G . In this way, the rows of the lower section of M_T represent a certain subset of the nodes in G . This is not true for the columns in M_T , which are all included indiscriminately. We will use variables (u, f_1, f_2) to only consider only the columns which represent the appropriate nodes.

EXAMPLE A.1.5. Consider matrix M_{T_1} as was defined in example A.1.3:

$$M_{T_1} = \left(\begin{array}{cc|c|c|c} 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 & 1 \\ \hline 1 & 1 & 0 & 1 & 1 \\ 1 & 1 & 0 & 1 & 1 \end{array} \right)$$

Because of the inclusion of the first and third rows from the lower section of matrix M_G , sort T_1 represents the subset $V_1 = \{1, 3\}$ of the nodes of G .

Consider the following assignment ρ of the variables in r_0 (shown as superscripts):

$$\left(\begin{array}{cc|c|c|c} 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & 1^x & 1^{c_1} & 1^{c_2} \\ 0 & 0 & 1 & 1 & 1 \\ \hline 1 & 1 & 0 & 1 & 1 \\ 1 & 1 & 0^u & 0^{f_1} & 1^{f_2} \end{array} \right)$$

Here, c_1 is assigned to the column which represents node 2 of G . This is undesired, since node 2 is not present in V_1 . This assignment will not be included in the set of total cases because $\text{val}(f_1) = 0$. Furthermore, it is not possible to build an assignment which assigns c_1 to the column representing node 2 and which also satisfies the antecedent, since there is no 1 valued cell to be found in the lower section of that column.

In contrast, the following assignment does satisfy the antecedent:

$$\left(\begin{array}{cc|c|c|c} 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & 1^x & 1^{c_1} & 1^{c_2} \\ \hline 1 & 1 & 0 & 1 & 1 \\ 1 & 1 & 0^u & 1^{f_1} & 1^{f_2} \end{array} \right)$$

As a final comment, note that there are two possibilities for discarding a variable assignment ρ : (i) it may be that the variable assignment does not satisfy the antecedent of rule r_0 , in which case it will never be counted (as is the case in this example), or (ii)

it may be that a variable assignment, if valid in a subgraph, will cause the structuredness to be less than 1, as it does satisfy the antecedent but not the consequent (as is the case in example A.1.4).

□

Finally, the trio (y, d_1, d_2) serve to assure that the node represented by the row of y and the node represented by the column of c_1 (and c_2) are not connected in G . The y variable is free to be assigned to any cell whose row is in the lower section of M_T and whose column is **idp** (i.e. any cell representing a node which is included in T). Also, d_1 and d_2 are already assigned cells whose columns represent a node in N_T (possibly different to the node represented by the row of y). While one of d_1, d_2 point to the lower section, left column set, the other of the two points to the adjacency submatrix. In the consequent of r_0 , it would be enough to ask that the d_i variable which points to the adjacency matrix point to a 1 valued cell. However, it is not possible to distinguish d_1 from d_2 , therefore, both are checked in a symmetrical fashion, with the subexpression $\text{val}(d_1) = 1 \vee \text{val}(d_2) = 1$. Momentarily, let us assume that d_1 points to the left column set and d_2 points to \bar{A}_G and let n be the node represented by y and m be the node represented by the column of c_1 . In this case, if $\text{val}(d_1) = 1$, then $n = m$ and d_2 will necessarily point to the cell (n, n) of \bar{A}_G , which will always contain a 1 (recall that G does not have self-edges). If $\text{val}(d_1) = 0$, then d_2 will point to the cell (n, m) of \bar{A}_G . This last cell must contain a 1 for the assignment to satisfy the consequent (meaning n and m are not connected).

EXAMPLE A.1.6. Using the same matrix M_{T_1} once again, consider the following assignment ρ of the variables in r_0 (shown as superscripts):

$$\left(\begin{array}{cc|c|c|c|c|c} 0 & 0 & 1 & 1 & & 1 & \\ 0 & 0 & 1 & & 1 & & 1 \\ 0 & 0 & 1^x & & 1^{c_1} & & 1^{c_2} \\ \hline 1 & 1 & 0^y & 1 & 0^{d_1} & 1 & 0 & 1^{d_2} \\ 1 & 1 & 0 & & 1 & 1 & 1 & 1 \end{array} \right)$$

Here, the row of y represents node 1 and the column of c_1 represents node 3. This assignment satisfies the antecedent of r_0 . As such, both nodes are in V_1 , so we must now check if they are connected in G . This is done with the subexpression $\text{val}(d_1) = 1 \vee \text{val}(d_2) = 1$ of the consequent of r_0 . Here, the consequent is also satisfied by ρ , since $\text{val}(d_2) = 1$, which tells us that nodes 1 and 3 are not connected in G .

□

We will now prove that graph G is 3-colorable if and only if RDF graph D_G has a σ_{r_0} -sort refinement \mathcal{T} with threshold 1 consisting of at most 3 implicit sorts.

A.1.2.1. G is 3-colorable if and only if the sort refinement exists

If the graph $G = (V, E)$ is 3-colorable, then let $f : V \rightarrow \{1, 2, 3\}$ be the coloring function, with the property that for every pair of nodes $u, v \in V$, if $(u, v) \in E$ then $f(v) \neq f(u)$. This property can also be expressed with the adjacency matrix A_G : for every pair $i, j \in \{1, \dots, n\}$, if $A_G[i, j] = 1$ then $f(i) \neq f(j)$. Let n be the number of nodes in G .

The RDF graph D_G , constructed as was previously specified, has three sets of n auxiliary rows (upper section) and one set of n rows (lower section). The right column set of the lower section contains the complemented adjacency matrix \bar{A}_G .

We construct a sort refinement with three sorts, T_1, T_2 and T_3 , as follows. The first set of auxiliary rows is assigned to T_1 , the second set of auxiliary rows is assigned to T_2 and the third set of auxiliary rows is assigned to T_3 . The last set of rows range from $3n + 1$ to $3n + n$. For every row $3n + i$ with $i \in [1, n]$, we assign this row to the sort $T_{f(i)}$. That is, the row $3n + i$ is assigned to the sort given by the color of the corresponding node in G .

We now argue that the sort refinement $\mathcal{T} = \{T_1, T_2, T_3\}$ has threshold 1 using rule r_0 . For this, consider a assignment ρ of the variables $(x, c_1, c_2, y_1, d_1, d_2, z, e, u, f_1, f_2)$ in r_0 to the cells in the sort T_1 . To be included in the set of total cases, the assignment must satisfy the antecedent of r_0 . We will now consider the restrictions produced by this

fact (without loss of generality, we assume that the rows of M_{T_1} are ordered in the same fashion as presented previously).

Considering the subexpression $\text{prop}(x) = \mathbf{idp}' \wedge \text{val}(x) = 1$ of r_0 , let $\rho(x) = (i_x, 3)$, where $i_x \in [1, n]$. That is, the variable x is assigned to a cell in row i_x and column 3 (recall that column \mathbf{idp} is the third column).

With the subexpression $c_1 \neq x \wedge \text{subj}(c_1) = \text{subj}(x) \wedge \text{val}(c_1) = 1 \wedge c_2 \neq x \wedge \text{subj}(c_2) = \text{subj}(x) \wedge \text{val}(c_2) = 1 \wedge c_1 \neq c_2$ we can restrict the assignment of c_1 and c_2 . Without loss of generality, assume that:

$$\rho(c_1) = (i_x, 3 + i_x) \quad \text{and} \quad \rho(c_2) = (i_x, 3 + n + i_x)$$

Both c_1 and c_2 are assigned to the same row as x and furthermore, since the auxiliary rows contain diagonal matrices ($D_{n \times n}$), the row i_x restricts the possible columns to only two, given by the cells with value 1.

Given that the subexpression $\text{prop}(y) = \mathbf{idp}' \wedge \text{val}(y) = 0 \wedge \text{subj}(d_1) = \text{subj}(y) \wedge \text{prop}(d_1) = \text{prop}(c_1) \wedge \text{subj}(d_2) = \text{subj}(y) \wedge \text{prop}(d_2) = \text{prop}(c_2)$ holds, we have:

$$\rho(y) = (n + i_y, 3), \quad \rho(d_1) = (n + i_y, 3 + i_x), \quad \text{and} \quad \rho(d_2) = (n + i_y, 3 + n + i_x).$$

where i_y ranges from 1 to the number of nodes with color 1.

Next, the subexpression $\text{prop}(z) = \mathbf{idp}' \wedge \text{subj}(z) = \text{subj}(e) \wedge \text{prop}(e) = \text{prop}(c_1) \wedge e \neq c_1 \wedge \text{val}(e) = 1$ sets:

$$\rho(e) = (n + \alpha_{i_x}, 3 + i_x) \quad \text{and} \quad \rho(z) = (n + \alpha_{i_x}, 3).$$

Since the value of the cell assigned to e must be 1, it must share columns with c_1 , and $e \neq c_1$, we must assign e to a cell of the diagonal of the lower section, left column set. However, note that this block is incomplete: only the rows corresponding to nodes in color 1 have been included. If there is no node which provides a value 1 on column $3 + i_x$, then this assignment cannot be considered as a total case. We therefore assume that this node

is present. In that case, we define $n + \alpha_{i_x}$ to be the row at which node i_x has been placed. Furthermore, if this subset is empty (i.e. no nodes of G have been included) then there will be no total cases. In this case, by definition the structuredness is assigned value 1.

The last subexpression of the antecedent is $\text{prop}(u) = \text{'idp'} \wedge \text{val}(u) = 0 \wedge \text{subj}(u) = \text{subj}(f_1) \wedge \text{prop}(f_1) = \text{prop}(c_1) \wedge \text{subj}(u) = \text{subj}(f_2) \wedge \text{prop}(f_2) = \text{prop}(c_2) \wedge \text{val}(f_1) = 1 \wedge \text{val}(f_2) = 1$. Since the variables u , f_1 , and f_2 do not appear in the consequent, this subexpression acts only as a restriction to the total cases. It allows us to write:

$$\rho(u) = (n + i_u, 3), \quad \rho(f_1) = (n + i_u, 3 + i_x), \quad \text{and} \quad \rho(f_2) = (n + i_u, 3 + n + i_x)$$

This restriction has the effect of only considering, in the set of total cases, the assignments where column $3 + n + i_x$ (and column $3 + i_x$ also) refers to a node of G which has been included in color 1. This is because the value of the cell assigned to f_1 must be 1. Variable f_2 will be restricted to the cell which represents the edge (i_x, i_x) of the graph G , which, since there are no self-edges, will be 1 (when complemented).

We now turn our attention to the consequent of r_0 . Given an assignment ρ of the variables in r_0 as seen before, we shall show that it will also satisfy the consequent.

Consider the subexpression $\text{val}(d_1) = 1 \vee \text{val}(d_2) = 1$. Variable d_2 will be assigned to a cell of the complemented adjacency matrix. Since column $3 + n + i_x$ of d_2 is fixed, d_2 will point to a cell which corresponds to the edge (i_y, i'_x) of G , where we define i'_x to be the node represented by row i_x (recall that the index i_x actually ranges from 1 to the number of nodes included in the subset). If it is the case that $i'_x = i_y$, then $\text{val}(d_1) = 1$ will hold. Now, we know node i'_x is included in color 1. Since color 1 is an independent set, the complemented adjacency must contain a 1 in the corresponding cell. Therefore, the value of the cell assigned to d_2 will be 1.

The subexpression $\text{val}(z) = 0$ will be true because we have included exactly one copy of the n auxiliary rows in T_1 . More precisely, since the column of e is fixed and $e \neq c_1$, e must point to the only other cell in that column which has value 1, for the value of z to

be 0. If any auxiliary row had been included twice, then z could be assigned to a cell with value 1.

We have shown that every assignment which satisfies the antecedent (is a total case) also satisfies the consequent (is a favorable case). Therefore, the structuredness value for color 1 is 1. The same reasoning can be applied to T_2 and T_3 . Therefore, the sort refinement constructed has threshold 1.

The proof of the other direction (if the sort refinement exists then G is 3-colorable) is analogous and will not be shown explicitly here.