

BOOTOX: Practical Mapping of RDBs to OWL 2

Ernesto Jiménez-Ruiz¹(✉), Evgeny Kharlamov¹, Dmitriy Zheleznyakov¹,
Ian Horrocks¹, Christoph Pinkel², Martin G. Skjæveland³, Evgenij
Thorstensen³, and Jose Mora⁴

¹ Department of Computer Science, University of Oxford, Oxford, UK
`ernesto.jimenez.ruiz@cs.ox.ac.uk`

² Fluid Operations AG, Walldorf, Germany

³ Department of Informatics, University of Oslo, Oslo, Norway

⁴ Sapienza Università di Roma, Rome, Italy

Abstract. Ontologies have recently become a popular mechanism to expose relational database (RDBs) due to their ability to describe the domain of data in terms of classes and properties that are clear to domain experts. Ontological terms are related to the schema of the underlying databases with the help of mappings, i.e., declarative specifications associating SQL queries to ontological terms. Developing appropriate ontologies and mappings for given RDBs is a challenging and time consuming task. In this work we present BOOTOX, a system that aims at facilitating ontology and mapping development by their automatic extraction (i.e., bootstrapping) from RDBs, and our experience with the use of BOOTOX in industrial and research contexts. BOOTOX has a number of advantages: it allows to control the OWL 2 profile of the output ontologies, bootstrap complex and provenance mappings, which are beyond the W3C direct mapping specification. Moreover, BOOTOX allows to import pre-existing ontologies via alignment.

1 Introduction

The Semantic Web community has actively investigated the problem of bridging the gap between relational databases and ontologies. One of the main targets behind this effort is to enable the access to the data stored in databases via Semantic Web technologies. An advantage of this approach is that ontologies provide a formal specification of the application domain that is close to the end-users' view of the domain, while databases are oriented towards an efficient data storage and retrieval and thus represent the data using structures that often not intuitive to end-users.

Manually building an ontology and connecting it to the data sources via mappings is, however, a costly process, especially for large and complex databases (a typical scenario in industry [20,22]). The cost of this manual process will typically be even more severe when dealing with multiple databases, e.g., in the context of accessing the Deep Web [16]. To aid this process, tools that can extract a preliminary ontology and mappings from database schemata play a critical role.

In the literature one can find a broad range of approaches to bootstrap an ontology and mappings from a relational database schema. The interested reader may have a look at the following surveys [34, 41]. These approaches can be classified with respect to different aspects such as: *(i)* level of automation (i.e., manual, semi-automatic, automatic), *(ii)* type of mappings (i.e., complex or direct mappings), *(iii)* language of the bootstrapped mappings and the ontology, *(iv)* reuse of external vocabularies (e.g., domain ontologies or thesauri), and *(v)* purpose (e.g., OBDA, constraint validation, database integration, ontology learning).

In this paper we present BOOTOX¹ (not) yet another ontology and mapping bootstrapper. Our main motivation to implement a new bootstrapper is to give more flexibility with respect to the classification aspects described above. Most of the existing approaches commit to concrete purposes or to a concrete ontology expressiveness. BOOTOX allows to define different “profiles” depending on the application scenario and the required Semantic Web technologies. For example, if the bootstrapped ontology is to be used in a so-called Ontology Based Data Access (OBDA) scenario where the ontology provides a virtual access layer to the data, OWL 2 QL will be chosen as the ontology language as it is required by the query rewriting engine.² Nevertheless, if the data is materialised, one could opt for other OWL 2 profiles depending on the used query answering engine. BOOTOX also allows to import domain ontologies, which will be integrated to the bootstrapped one via alignment [18] or directly mapped to the database. BOOTOX follows the W3C direct mappings directives to connect the ontological vocabulary to the relational database; moreover, it offers a suit of advanced techniques for bootstrapping of mapping that are beyond the direct ones. Furthermore, it extends the bootstrapped mappings with provenance information.

We have tested BOOTOX and compared it to several bootstrapping systems over a number of databases and use cases from industrial and research contexts, including the relational-to-ontology benchmark suite RODI [29].

The paper is organised as follows. Section 2 introduces the concepts behind semantic access to databases. Section 3 presents the problem of bootstrapping. Section 4 describes the techniques implemented in BOOTOX. Section 5 presents the scenarios where we deployed BOOTOX and the conducted experiments to evaluate the quality of the bootstrapping. In Section 6 we provide a summary of relevant related work, and we conclude in Section 7.

2 Exposing Relational Data Through Ontologies

The main idea behind exposing relational data via an ontology is to provide the user with access to the data store via the use of a domain specific vocabulary of classes, i.e., unary predicates, and properties, i.e., binary predicates, that the user is familiar with. This vocabulary is related to the database schema via view definitions, called *mappings*; thus, technical details of the database schema

¹ <http://www.cs.ox.ac.uk/isg/tools/BootOX/>

² The language underlying OWL 2 QL has the first-order (FO) rewritability property [8].

Well (<u>id</u> , name ^{NN} , type) Field (<u>id</u> , name, status ^{CK} , intersects_field ^{FK}) Operator (<u>id</u> , name ^{NN}) Operator_Field (operator ^{FK} , field ^{FK}) Wellbore (<u>id</u> , name ^{UQ} , content ^{CK} , depth ^{CK} , well ^{FK} , location ^{FK}) ExplorationWellbore (wellbore ^{FK} , seismic_location) DevelopmentWellbore (wellbore ^{FK} , production_facility)
--

Fig. 1. Relational Database Schema. Primary keys are underlined. FK stands for Foreign Key (the foreign key name indicates the target table), NN stands for Not Null, UQ stand for Unique and CK stands for Check

are hidden from end-users. Using an example from the oil industry domain, we now intuitively introduce the main concepts that are needed to understand the approach and refer the reader to [31] for details.

Relational Databases. In the relational data model, a database consists of a schema and instance, where the schema is a set of table names with corresponding attribute names and constraints, e.g., primary and foreign keys. The instance ‘populates’ tables of the schema with tuples by assigning values to the tables’ attributes. Figure 1 shows our running example, a fragment of a relational database based on the oil industry domain. For example, a Wellbore is given a name that is unique, has a content (e.g., GAS, OIL, or DRY), has a depth (e.g., 12,500 feet), belongs to a Well, and is located in a Field.

Ontologies. An ontology is usually referred to as a ‘conceptual model’ of (some aspect of) the world. It introduces the vocabulary of classes and properties that describe various aspects of the modelled domain. It also provides an explicit specification of the intended meaning of the vocabulary by describing the relationships between different vocabulary terms. These relationships are specified with special first-order formulae, also called axioms, over the vocabulary. An ontology can be thought of simply as a set of such axioms—i.e., a logical theory. Besides axioms, an ontology can contain ontological facts, which can be specified as first-order atoms with constants, but not variables. These constants are usually interpreted as (representations of) objects from the domain. Viewing an ontology as a logical theory opens up the possibility of using automated reasoning for different tasks, including checking an ontology’s internal consistency (checking whether the ontology does not entail ‘false’, i.e., if it is logically consistent), query answering, and other (non-)entailments. OWL 2 is a W3C standard ontology language and it is heavily used in the Semantic Web community. OWL 2 has profiles:³ EL, QL, and RL, that have different favourable computational properties.

³ OWL 2 profiles: <http://www.w3.org/TR/owl2-profiles/>

Mappings. Mappings declaratively define how ontological terms are related to terms occurring in the relational schema and are essentially view definitions of the following form that declare how to populate classes with objects—in OWL 2 objects are represented with Uniform Resource Identifiers (URIs)—and to populate properties with object-object and object-value pairs:

$$\begin{aligned} \text{Class}(f_o(x)) &\leftarrow \text{SQL}(x), \\ \text{objectProperty}(f_o(x), f_o(y)) &\leftarrow \text{SQL}(x, y), \\ \text{dataProperty}(f_o(x), f_v(y)) &\leftarrow \text{SQL}(x, y), \end{aligned}$$

where $\text{SQL}(x)$ and $\text{SQL}(x, y)$ are SQL queries with respectively one and two output variables, and f_o, f_v are functions that ‘cast’ values returned by the SQL queries into respectively objects, i.e. URIs, and values.⁴ Classes are populated with URIs $f_o(x)$ computed from the values x returned by $\text{SQL}(x)$. Properties can relate two objects, e.g., by stating that a wellbore has a particular location, or assigning a value to an object, e.g., stating that a field has a particular name (a string); and they are respectively populated with pairs of objects $f_o(x), f_o(y)$ or pairs of an object $f_o(x)$ and value $f_v(y)$ computed from the values x and y returned by the SQL query. For our running example, we may create the following mappings:

$$\begin{aligned} \text{Wellbore}(f_o(\text{id})) &\leftarrow \text{SELECT id FROM Wellbore}, \\ \text{hasLocation}(f_o(\text{id}), f_o(\text{location})) &\leftarrow \text{SELECT id, location FROM Wellbore}, \\ \text{hasContent}(f_o(\text{id}), f_v(\text{content})) &\leftarrow \text{SELECT id, content FROM Wellbore}. \end{aligned}$$

Query Answering. Consider a data access instance $(\mathcal{D}, \mathcal{V}, \mathcal{O}, \mathcal{M})$, where \mathcal{D} is an RDB, \mathcal{V} is an ontological vocabulary, \mathcal{O} is a set of ontological axioms over \mathcal{V} , and \mathcal{M} is a set of mappings between \mathcal{V} and \mathcal{D} . There are two approaches to answer a query Q over \mathcal{V} : (i) *materialisation*: ontological facts are materialised (i.e., classes and properties participating in mappings are populated with individuals by evaluating SQL queries participating in mappings) and this gives a set of ontological facts \mathcal{A} and then Q is evaluated over \mathcal{O} and \mathcal{A} with standard query-answering engines for ontologies, or (ii) *virtual*: Q should be first rewritten into an SQL query SQL using \mathcal{O} and \mathcal{M} and then SQL should be executed over \mathcal{D} .

In either case, the ontology and mappings are the backbone of query processing in this data access approach; thus, the problem of obtaining them has to be addressed in any implementation of such approach. The exact kind of ontology and mappings required for data access depends on the application scenario. For example, in a scenario where virtual query answering is required, the ontology should be in OWL 2 QL. In the case of materialisation the ontology should fall in the profile supported by the underlying systems, e.g., *EOLO* [42] can do query answering over OWL 2 EL ontologies, and thus the bootstrapping should output an EL ontology; a semantic faceted search system *SEMFACTET* [2–4, 15] relies on *RDFox* [28] that can do query answering over OWL 2 RL ontologies and thus the bootstrapper facilitating *SEMFACTET* should produce RL ontologies; *PAGODA* [47] does query answering over the whole OWL 2 and can work with any bootstrapper that produces OWL 2 ontologies.

⁴ These functions ensure coherent generation of URIs that respects primary and foreign keys.

3 Bootstrapping: Problem Definition

In this section we define the tasks for bootstrapping ontologies and mappings, or *assets*, from RDBs and quality metrics for these tasks. The tasks are the following:

- (i) *Semantic access instance generation*: Given a relational database \mathcal{D} , generate an instance $(\mathcal{D}, \mathcal{V}, \mathcal{O}, \mathcal{M})$. This task can be naturally divided into two sub-tasks.
 - *Vocabulary and Ontology generation*: Given \mathcal{D} , create a vocabulary \mathcal{V} and an ontology \mathcal{O} over \mathcal{V} .
 - *Mapping generation*: Given \mathcal{D} , \mathcal{V} , and \mathcal{O} create a set of mappings \mathcal{M} relating \mathcal{D} with \mathcal{V} .
- (ii) *Importing*: Given an instance $(\mathcal{D}, \mathcal{V}, \mathcal{O}_1, \mathcal{M})$ and an ontology \mathcal{O}_2 , return an instance $(\mathcal{D}, \mathcal{V}, \mathcal{O}, \mathcal{M})$, where \mathcal{O} is the alignment of \mathcal{O}_1 and \mathcal{O}_2 .

Task (ii) is important in applications where ontologies (partially) describing the domain of interest have been already created and users want to incorporate them into their semantic data access system.

The bootstrapping of the ontologies and the mappings enables a (preliminary) deployment of semantic data access system. However, the bootstrapped assets should meet some minimal requirements so that they can be usable in practice. We have identified the following *metrics* to measure the quality of generated assets:

- (1) *Ontology language*: compliance with standard ontology languages with well-defined semantics like OWL 2 and its profiles to enable the use of a wide-range of Semantic Web technologies. Note that the choice of the ontology language (e.g., one of OWL 2 profiles) also affects suitability of the different query answering engines.
- (2) *Mapping language*: compliance with standard directives like the W3C direct mapping specification⁵ and standard W3C mapping languages like R2RML.⁶
- (3) *Query coverage*: suitability of the ontology vocabulary to formulate the queries that the user is interested in.
- (4) *Query results*: accuracy of the query results obtained with the bootstrapped instance in the sense that the query answer should satisfy the user's expectations.

4 Bootstrapping Techniques in BOOTOX

In this section we present our bootstrapping techniques implemented in BOOTOX. Section 4.1 focuses on the techniques to generate the ontological vocabulary and the axioms over this vocabulary, where we give special attention

⁵ Direct mappings: <http://www.w3.org/TR/rdb-direct-mapping/>

⁶ R2RML language: <http://www.w3.org/TR/r2rml/>

to the concrete language of the generated ontology. In Section 4.2 we describe how the database terms are linked to the ontological vocabulary. Furthermore, we also present how mappings are enhanced with provenance information. Section 4.3 gives an overview of the incorporated ontology alignment techniques to import pre-existing ontologies. Finally, Section 4.4 summarises the conformance with the requirements presented in Section 3.

4.1 Vocabulary and Ontology Generation

The general rule to create ontological vocabulary from a relational schema would translate (i) each (non-binary) table into an OWL class; (ii) each attribute not involved in a foreign key into an OWL datatype property; (iii) each foreign key into an OWL object property. Generation of axioms, however, is usually a more involved process. BOOTOX follows Table 1 to automatically create the vocabulary and a set of OWL 2 axioms from the listed database features.⁷ One could opt for adding all the axioms associated with a feature or only a selection of them depending on the intended purpose. In the remainder of the section we will discuss the main considerations regarding bootstrapping that should be taken into account.

Closed-world vs Open-world Semantics. Modelling database features in OWL 2 inevitably leads to different interpretations due to the Closed-world (CWA) and Open-world assumptions (OWA). In a relational database, every fact that occurs in the database is true, and all other facts are false (CWA). In an ontology, the truth value of some facts may be unknown, i.e., they are neither provably true nor provably false (OWA). Moreover, constraints have a very different meaning in databases and ontologies [27, 34]. In a database, constraints define valid database states, and updates that would violate these constraints are simply rejected. In an ontology, constraints are treated as material implications, and may lead to the derivation of “new” facts that are not explicitly stated in the data. Such derivations may produce unintended consequences,⁸ as we discuss next.

Table 1 presents a possible encoding of relational database features as OWL 2 axioms. The encoding, in general, does not lead to only one option, but several possible OWL 2 axioms. For example, considering the running example of Figure 1, for the data attribute *name* in table *Field* we apply encoding pattern (3) in Table 1, which proposes to add five different axioms. The general rule translates the data attribute *name* as a data property declaration axiom. In addition one may opt to add global restrictions (e.g., ‘*name* Domain: *Field*’)⁹ and/or local restrictions (e.g., ‘*Field* SubClassOf: *name* some *xsdstring*’). The use of global and/or local restrictions will lead to different interpretations and

⁷ When not stated the contrary in Table 1, a class C_T , an object property P_f , a data property R_a and a datatype d_t represent the ontology encoding of a table T , a foreign key fk , a data attribute a , and an SQL type t , respectively.

⁸ Note that the use of disjointness axioms may help in the detection of such “errors”.

⁹ For writing OWL 2 axioms we use the Manchester OWL Syntax [17].

Table 1. Encoding of relational database features as OWL 2 axioms. OWL 2 axioms are expressed in the Manchester OWL Syntax [17]. * Enumeration with only one literal

#	RDB feature	Ontology feature	OWL 2 axiom	OWL 2 profile		
				QL	RL	EL
(1)	Non-binary Relation / Table T	A class C_T for the non-binary table	Class: C_T	✓	✓	✓
(2)	Binary Relation or Many-to-Many Table referencing tables T_1 and T_2	A property P (and its inverse Q) associated to the classes C_{T_1} and C_{T_2} with local and/or global constraints	ObjectProperty: P	✓	✓	✓
			Q InverseOf: P	✓	✓	-
			P Domain: C_{T_1}	✓	✓	✓
			P Range: C_{T_2}	✓	✓	✓
			C_{T_1} SubClassOf: P some C_{T_2}	✓	-	✓
			C_{T_1} SubClassOf: P only C_{T_2}	-	✓	-
(3)	Data attribute in table T of (sql) type t	A property R_a associated to the class C_T and datatype d_t with local and/or global constraints.	DataProperty: R_a	✓	✓	✓
			R_a Domain: C_T	✓	✓	✓
			R_a Range: d_t	✓	✓	✓
			C_T SubClassOf: R_a some d_t	✓	-	✓
			C_T SubClassOf: R_a only d_t	-	✓	-
(4)	Foreign Key in table T_1 , referencing T_2 , no intersection with or strict subset of primary key	A property P_f associated to the classes C_{T_1} and C_{T_2} with local and/or global constraints	ObjectProperty: P_f	✓	✓	✓
			P_f Domain: C_{T_1}	✓	✓	✓
			P_f Range: C_{T_2}	✓	✓	✓
			C_{T_1} SubClassOf: P_f some C_{T_2}	✓	-	✓
			C_{T_1} SubClassOf: P_f only C_{T_2}	-	✓	-
(5)	Foreign Key is the primary key in T_1 , ref. T_2	Class C_{T_1} is subsumed by class C_{T_2}	C_{T_1} SubClassOf: C_{T_2}	✓	✓	✓
(6)	Foreign Key in table T referencing the same table	A property P_f associated to class C_T with a self-restriction. The property P_f may also be declared with several characteristics	C_T SubClassOf: P_f some C_T	✓	-	✓
			C_T SubClassOf: P_f some Self	-	-	✓
			P_f Characteristics: Transitive	-	✓	✓
			P_f Characteristics: Symmetric	✓	✓	-
			P_f Characteristics: Reflexive	✓	✓	✓
(7)	Primary Key or Unique constraint in table T_1 on a foreign key fk referencing T_2	Key axiom for class C_{T_1} and property P_f . P_f is associated to (local and/or global) cardinality constraints	C_{T_1} HasKey: P_f	-	✓	✓
			P_f Characteristics: Functional	-	✓	-
			P_f Characteristics: InverseFunctional	-	✓	-
			C_{T_1} SubClassOf: P_f exactly 1 C_{T_2}	-	-	-
			C_{T_1} SubClassOf: P_f max 1 C_{T_2}	-	✓	-
			C_{T_1} SubClassOf: P_f some C_{T_2}	✓	-	✓
(8)	Primary Key or Unique constraint on a data attribute a of (sql) type t in table T	Key axiom for class C_T and data property R_a . R_a is associated to (local and/or global) cardinality constraints.	C_T HasKey: R_a	-	✓	✓
			R_a Characteristics: Functional	-	✓	✓
			C_T SubClassOf: R_a exactly 1 d_t	-	-	-
			C_T SubClassOf: R_a max 1 d_t	-	✓	-
			C_T SubClassOf: R_a some d_t	✓	-	✓
(9)	Composed Primary Key in table T	Key axiom for the class C_T and the data and object properties involved in primary key	C_T HasKey: $R_1 \dots R_n P_1 \dots P_n$	-	✓	✓
(10)	Not Null Constraint on a data attribute in T , type t	Existential or cardinality restriction over R_a in C_T	C_T SubClassOf: R_a min 1 d_t	✓	-	-
			C_T SubClassOf: R_a some d_t	✓	-	✓
(11)	Not Null Constraint on a foreign key in T_1 , ref. T_2	Existential or cardinality restriction over P_f in C_T	C_{T_1} SubClassOf: P_f min 1 C_{T_2}	-	-	-
			C_{T_1} SubClassOf: P_f some C_{T_2}	✓	-	✓
(12)	Check Constraint on data attribute a of type t in table T listing possible values $v_1 \dots v_n$ ($n \geq 1$)	Enumeration of literals: in a restriction in class C_T and/or as a range of R_a . Alternatively, one could create subclasses for each of the values	C_T SubClassOf: R_a some $\{v_1 \dots\}$	-	-	✓*
			C_T SubClassOf: R_a only $\{v_1 \dots\}$	-	-	-
			C_T SubClassOf: R_a value v_1	-	-	✓
			R_a Range: $\{v_1 \dots\}$	-	-	✓*
			$C_{T_{v_i}}$ SubClassOf: C_T	✓	✓	✓
(13)	Check Constraint on attrib. a in table T restricting numerical range of t	Datatype restriction: in a class restriction in C_T and/or as a range of R_a	C_T SubClassOf: R_a some $d_t [> x]$	-	-	-
			C_T SubClassOf: R_a only $d_t [> x]$	-	-	-
			R_a Range: $d_t [> x]$	-	-	-
(14)	Several data attributes $a_1 \dots a_n$ in different tables $T_1 \dots T_n$ with the same name and type t	Group properties $R_1 \dots R_n$ under new superproperty R_a or merge $R_1 \dots R_n$ into new property R_a	R_i SubPropertyOf: R_a	✓	✓	✓
			R_a Domain: C_{T_1} or \dots or C_{T_n}	-	-	-
			C_{T_i} SubClassOf: R_a some d_t	✓	-	✓
			C_{T_i} SubClassOf: R_a only d_t	-	✓	-
(15)	T_1 and T_2 not involved in an inheritance relationship	Class C_{T_1} is disjoint with class C_{T_2}	C_{T_1} DisjointWith: C_{T_2}	✓	✓	✓

potentially to different logical consequences when combined with data facts. For example, the fact ‘*operator987* Facts: *name* Statoil’ in combination with the domain axiom for *Field* and *name* will lead to the undesired consequence ‘*operator987* Types: *Field*’. In this case, if the property *name* is to be used in different contexts, using only local restrictions seems to be more appropriate.

Profiling the Ontology. BOOTOX takes into account the concrete language (i.e., one of OWL 2 profiles) that is generated from the relational database features, which will enable the use of different Semantic Web technologies (e.g., query answering engines). As mentioned above, database features can be encoded using different axioms which may lead to different OWL 2 profiles. For example, primary keys and unique constraints (patterns (7) and (8) in Table 1) can be modelled with the OWL 2 construct *HasKey* (e.g., ‘*Well* *HasKey*: *id*’), which is supported in OWL 2 RL and OWL 2 EL, but must be avoided if the target ontology language is OWL 2 QL. Alternatively (or in addition) one could use global and local cardinality restrictions (e.g., ‘*id* *Characteristics*: *Functional*’ and ‘*Well* *SubClassOf*: *id* exactly 1 *xsdint*’, respectively), which leads to similar issues since not all profiles support them. For example, none of the OWL 2 profiles support exact cardinalities. If a profile does not support a type of axiom, an approximation is used instead (e.g., ‘*Well* *SubClassOf*: *id* max 1 *xsdint*’ that is supported in OWL 2 RL, or ‘*Well* *SubClassOf*: *id* some *xsdint*’ that is allowed in both OWL 2 QL and EL profiles). However, a representative approximation or alternative may not be always available for a database feature, as for the *datatype restrictions* (pattern (13) in Table 1), which is not supported by any of the profiles. BOOTOX, when the required ontology output is one of the OWL 2 profiles, keeps this knowledge as OWL 2 annotation axioms (e.g., ‘*depth* *Annotations*: *annmax_value*“12,500”’). These axioms have no logical impact on the ontology, but they can potentially be used in user interfaces to guide the formulation of queries (e.g., [38–40]).

Selection of Ontology Axioms. BOOTOX, in the default setting, encodes the database features with all axioms suitable for a required OWL 2 profile. The user can optionally select the kind of preferred axioms (i.e., local or global restrictions). Additionally, the suggested OWL 2 axioms may require an extra validation step to accurately represent a database feature. This is the case of pattern (6) in Table 1 where the generated object property can be declared with different characteristics which may not necessarily hold in the database. For example, in our running example from Figure 1, the foreign key *intersects_field* in table *Field* represents a self-reference to the same table. In this case, the foreign key can be encoded as a *reflexive* (a field intersects with itself) and *symmetric* object property; *transitivity*, however, does not necessarily apply for this property.

Mapping SQL Datatypes to OWL 2. There is a clear mapping between SQL datatypes to XML schema datatypes. However, there are some XML schema datatypes that are not built-in datatypes in OWL 2. For example, *xsddate* is not supported in OWL 2. Furthermore, OWL 2 profiles also include specific

restrictions on the supported datatypes. For example, *xsdboolean* or *xsddouble* are not supported in OWL 2 QL nor OWL 2 EL. BOOTOX addresses this issue by mapping SQL datatypes, whose XML counterparts are not supported in OWL 2 (resp., OWL 2 profile), to the OWL 2 built-in datatype *rdfsLiteral*. This datatype denotes the union of all datatypes (i.e., top datatype). Note that a more intuitive approximation of unsupported datatypes (e.g., *xsddouble* to *xsddecimal*) is not possible since the value spaces of primitive datatypes are disjoint.¹⁰

4.2 Mapping Generation

Generation of Direct Mappings. BOOTOX follows the W3C direct mapping guidelines to generate mappings from a relational database to an ontological vocabulary. BOOTOX relies on the R2RML language to produce direct mappings, which are particular cases of the mappings that can be expressed in R2RML. Intuitively, an R2RML mapping allows to map any valid SQL query or view (i.e., logical table) into a target vocabulary.

The W3C direct mapping specification does not introduce specific restrictions on the used target ontological vocabulary, that is, it only requires to reference the ontological vocabulary via its URI. Hence, the generated mappings may reference the vocabulary of any given ontology. BOOTOX allows this behaviour and it can produce mappings for a vocabulary generated from the relational database (see Section 4.1) or for the vocabulary of an external domain ontology.

Generation of Mappings beyond Direct Ones. Besides direct mappings described above, that involve only projection operator in the SQL part, BOOTOX can help user to build *complex* R2RML mappings, that may also include selection and/or join operators. There are three general approaches that we use to find candidate mappings:

- (i) In the direct mapping approach, tables are mapped to classes. To generalise this, consider a table T and a table T' that has a foreign key constraint referencing T or that can be joined. By taking the left join of T and T' we obtain a subset of the tuples in T . If this subset of tuples is sufficiently big and sufficiently different from T itself, we obtain a subclass of the class T was mapped to. Furthermore, by letting this process to continue recursively, we obtain rich class hierarchies. Here, as in direct mappings, the many-to-many exception applies.
- (ii) Another way to discover subclasses in the database is to look for clusters of tuples in a table. To this end, we can consider a tuple to be a vector of numerical values (assigning distinct numbers to each value), and look for sets of tuples whose distance from each other is sufficiently small.
- (iii) Finally, one can also consider subsets of the attributes of a table, particularly when tables are not known to in standard normal forms, such as BCNF. In such tables, there will usually be tuples with repeating values in

¹⁰ <http://www.w3.org/TR/swbp-xsch-datatypes/>

attributes. Such sets of tuples can, again, be considered subclasses of the class the table itself was mapped to.

Note that the generation of these complex mappings heavily relies on interaction with the users since these techniques allow to generate only the SQL queries for mappings, while they do not offer the names of classes and properties defined by these queries; the names should be introduced by the users.

Generation of Provenance Mappings. BOOTOX automatically extends direct mappings with meta-information about provenance.¹¹ The provenance meta-information is modelled in the mapping assertion, adding, for instance, the source database from which the information is extracted, and more granular information, like table and column identifiers. Currently, provenance in BOOTOX comes into three different granularity levels, whose convenience will depend on the intended use of this information:

- (i) *URI level:* each generated URI is associated with provenance information. This is especially interesting when different database elements are merged into the same URI (e.g., see pattern (14) in Table 1).
- (ii) *Triple level:* triples are annotated with provenance via RDF reification (i.e., three new triples are added for each triple).
- (iii) *Graph level:* in this granularity, provenance is attached to RDF named graphs that group triples with similar provenance characteristics. For example, one could group triples generated by automatically generated mappings or triples derived from the same data source.

This meta-information can be later used for a wide range of purposes. For example it can be used to provide a richer query answering interface, to help in the debugging of a data access instance (i.e., identifying faulty ontology axioms or mappings), or to discard some sources based on a variety of criteria, e.g., licenses, access privileges, costs, etc.

4.3 Importing

It is increasingly often the case that a high quality ontology of (parts of) the domain already exists and captures the domain experts vocabulary better than the directly mapped ontology. When such an ontology is available, BOOTOX allows importing it to extend the automatically generated ontology.¹² To this end, BOOTOX integrates the ontology alignment system LogMap [18, 19], that aligns two ontologies \mathcal{O}_1 and \mathcal{O}_2 by deriving OWL 2 equivalence and sub-class(property) axioms between the entities from \mathcal{O}_1 's and \mathcal{O}_2 's vocabularies using the lexical characteristics of the terms and the structure of the ontologies. BOOTOX gives special care to avoid introducing unwanted consequences that may lead to unexpected answers. Thus, LogMap will discard alignment axioms that would lead to

¹¹ Based on the W3C recommendation PROV-O <http://www.w3.org/TR/prov-o/>.

¹² If the imported ontology is outside the desired target profile it should first be approximated using off-the-shelf semantic or syntactic approximation techniques (e.g., [11]).

inconsistencies, or faulty consequences like ‘*Well* SubClassOf: *WellBore*’. This is based on novel techniques to avoid violations of the so called consistency and conservativity principles (the interested reader please refer to [37]).

4.4 Conformance with Problem Requirements

BOOTOX covers the two bootstrapping tasks defined in Section 3 and fully meets the Metrics (1) and (2) regarding the ontology and mapping language.

Furthermore, as shown in Section 5, it provides reasonable results in practice with respect to Metrics (3) and (4). *Query coverage*, i.e., Metric (3), is enhanced thanks to the importing facility. Regarding *query results*, i.e., Metric (4), note that we supply the system with *provenance* capabilities, which are useful for analysis. In the case when query results contain unexpected answers, provenance will help the user to “trace” the source of these answers, thus helping to gain a better understanding of possible issues with the ontology or mappings.

5 BOOTOX at Work

BOOTOX has been tested with a number of databases and scenarios (e.g., virtual and materialised data access). In this section we present results with respect to the quality of the bootstrapped assets (i.e., vocabulary, ontology, and mappings). We have evaluated (i) the ability of formulating queries with the bootstrapped vocabulary in an industrial scenario (Section 5.1), and (ii) the ability of (enabling the) answering of queries with the bootstrapped ontology and mappings in a controlled scenario (Section 5.2).

We compared BOOTOX¹³ to three other bootstrapping systems: IncMap [30], MIRROR [26], and -ontop- [32]. IncMap is designed to directly map a relational database to a target ontology, but is focused on a semi-automatic, incremental/interactive approach rather than direct automated bootstrappings. Both -ontop- and MIRROR follow an approach that is similar to the one employed in BOOTOX, with respect to the generation of a semantic data access instance (i.e., vocabulary, ontology, and mappings).¹⁴

5.1 Query Coverage in the EU Optique Project

In this section we assess the quality of the bootstrapped ontology vocabulary to enable the formulation of queries. To this end we used the terms from the query catalogs available in the industrial scenario provided by the EU Optique project. We looked at how many terms from the catalogs were *covered* by the bootstrapped ontological vocabulary and then did manual verification of the quality of the coverage.

¹³ Note that we have evaluated BOOTOX with its automatic setting, that is, with the functionality to generate complex mappings turned off.

¹⁴ -ontop- generates only vocabulary, that is, an ontology containing only declaration axioms.


The figure displays three main components of the BOOTOX interface:

- Main Menu (Top Left):** A grid of icons for various bootstrapping and analysis tools: Logical Bootstrapping, Step-by-Step Bootstrapping, Import-Align Ontology, Import Layer Ontology, Approximate Ontology, Statistics Bootstrapping, Integrated Bootstrapping, Provenance Bootstrapping, Statistics Provenance, Bootstrapping for VQS, Extracted Annotations, and Coverage Query Catalog.
- Integrated Ontology and Mapping Bootstrapper (Top Right):** A form for configuring bootstrapping. It includes:
 - Available schemata:** A table with columns for Table name and Primary key. Example: `access_object` with primary key `access_object_s`.
 - Provenance mappings:** Checkboxes for `At URI level` and `At triple level`.
 - Extended features for the bootstrapped ontology:** Checkboxes for `Extract annotations for QP1`, `Infer subclass relationships`, and `Merge properties with same name and range`.
 - Import domain ontology:** A text input field for an ontology URL.
 - OWL 2 profile:** A table showing `OWL 2 DL` with `687` Axioms and `97` Classes.
 - OWL 2 QL approximation:** Checkboxes for `Via ontology alignment`, `Via direct connection to the database`, and `Perform approximation of the bootstrapped ontology`.
- Coverage of Query Catalogs (Bottom Left):** A section showing the selection of an ontology and a query catalog. It includes two pie charts:
 - Coverage of class terms:** A pie chart showing the distribution of class terms, with a legend for `Not covered`, `Covered (0.0, 1.0)`, and `Covered (0.0, 1.0)`.
 - Coverage of property terms:** A pie chart showing the distribution of property terms, with a legend for `Not covered`, `Covered (0.0, 1.0)`, and `Covered (0.0, 1.0)`.


Fig. 2. BOOTOX interfaces in the Optique platform, from left to right, top to bottom: the main BOOTOX menu, the bootstrapping form, and the query catalog coverage

Industrial Scenario. The EU Optique project aims at facilitating scalable end-user access to big data in the oil and gas industry [14]. Optique advocates for an OBDA approach where the ontology provides a virtual access to the data and the mappings connect the ontology with the data source. The project is focused around two demanding use cases provided by Optique industry partners Siemens and Statoil. BOOTOX has been deployed in Siemens [20] and Statoil [22] as part of the “Ontology and mapping management module” provided by Optique’s platform [21]. Furthermore, BOOTOX has already shown to save time, in the creation of the initial OBDA assets, to the IT experts of both Statoil and Siemens. Note that the Optique scenario requires BOOTOX to bootstrap an OWL 2 QL as it is required by the query rewriting engine in order to rewrite the queries formulated over the ontology into queries on the database using reasoning [8].

BOOTOX in the Optique Platform. Figure 2 shows an overview of the BOOTOX related interfaces. The main menu presents to the user the available options currently supported in BOOTOX: automatic bootstrapper, guided bootstrapper, ontology alignment, provenance bootstrapping, bootstrapping related statistics, etc. The screenshot on the right shows the integrated bootstrapping form in BOOTOX where the user can select from a number of options including the database schema, imported ontology, provenance, OWL 2 QL approximation. The bottom screenshot shows the coverage of a selected query catalog by the vocabulary of a bootstrapped ontology.



(a) Statoil use case.



(b) Siemens use case.

Fig. 3. Coverage of terms from the Optique query catalog with terms from ontologies. Inner pie charts show coverage by lexical confidence using I-SUB [43]: ■ in [0.9, 1.0], ■ in [0.8, 0.9), □ in [0.6, 0.8). Outer pie charts represent the manual verification of the quality of the terms with a coverage above 0.6: ■ *true positive*, ■ *semi-true positive*, ■ *false positive*. Note that *semi-true positives* are not clear-cut cases where the ontology term has a broader or narrower meaning with respect to the query term

Coverage of Statoil Query Catalog. The query catalog at Statoil currently includes 60 queries that contain representative information needs from Statoil geologists. Most of the data needed by Statoil geologists is stored in the Exploration and Production Data Store (EPDS), Statoil’s corporate data store for exploration and production data and interpretations. The *NPD FactPages* [36] ontology is relevant to the domain of EPDS.

We bootstrapped ontological vocabulary from the relevant parts of EPDS using BOOTOX, MIRROR and -ontop-. The generated vocabulary, on average,

contained more than 3,000 classes, 3,500 object properties, and 42,000 datatype properties. Note that IncMap relies on the vocabulary of the available domain ontology. BOOTOX, unlike -ontop- and MIRROR, includes a built-in ontology alignment system which allows to import the vocabulary of the domain ontologies into the bootstrapped ontology.

The results of the query catalog coverage are summarised in Figure 3a. The first column represent the coverage of the bootstrapped ontologies computed by BOOTOX (without importing the domain ontologies), -ontop- and MIRROR. Since all three systems rely on the direct mapping directives, the bootstrapped vocabulary is, apart from minor differences, basically the same. The middle columns represent the coverage of the vocabulary of the domain ontology, which is equal to the coverage of IncMap. The third columns shows the coverage results achieved by the ontology bootstrapped by BOOTOX including importing. For example, 44% of the classes in the query catalog has a good lexical intersection (greater or equal 0.8) with terms of the ontology bootstrapped by BOOTOX; furthermore, 29% of the classes are fully covered (i.e., true positives).

Coverage of Siemens Query Catalog. The data in the Siemens use-case is stored in several databases with different schemata. Although the schemata are not specially large, the size of the data is in the order of hundreds of terabytes, e.g., there is about 15 GB of data associated to a single turbine, and it currently grows with the average rate of 30 GB per day [20]. As for the Statoil use case, we extracted the relevant terms of the query catalog, bootstrapped an ontology from one of the Siemens databases using BOOTOX, MIRROR and -ontop-. Additionally, BOOTOX performed alignment with two Siemens ontologies about diagnostic procedures and turbines. Finally, IncMap relied on the vocabulary of these ontologies. The results of the coverage are summarised in Figure 3b.

Quality Assessment. The experiments show that the bootstrapped ontologies without importing had a higher coverage than the domain ontologies in isolation, e.g., 39% of query class coverage against 27% in the Statoil use case. These results suggest that there is an adequate number of table and column names with potentially adequate semantic relations with the terms that domain experts at Statoil and Siemens have in mind when they access data, and thus, the ontology vocabulary computed by the ontology bootstrappers is indeed relevant to query formulation. Nevertheless, the domain ontologies naturally complement the vocabulary obtained from the database and hence BOOTOX is able to bootstrap and ontology with better coverage over the query catalog than the ones generated by -ontop- and MIRROR. For example, 48% of the classes (resp., 90%) in the Statoil (resp., Siemens) catalog are fully or partially covered in the bootstrapped ontology computed by BOOTOX.

Table 2. RODI results of all tests per scenario. Values in cells are the success scores

Scenario	IncMap	MIRROR	-ontop-	BootOX
Adjusted naming				
CMT	0.5	0.28	0.39	0.39
CONFERENCE	0.26	0.27	0.37	0.37
SIGKDD	0.21	0.3	0.45	0.45
Cleaned hierarchies				
CMT	0.44	0.17	0.28	0.28
CONFERENCE	0.16	0.23	0.3	0.3
SIGKDD	0.11	0.11	0.16	0.16
Combined case				
SIGKDD	0.05	0.11	0.16	0.16
Missing FKs				
CONFERENCE	0.03	0.17	-	0.17
Denormalised				
CMT	0.22	0.22	0.28	0.28

5.2 Query Answering in a Controlled Scenario

In this section we assess the quality of the bootstrapped ontology and mappings to enable the answering of queries in a controlled scenario.¹⁵ To this end we ran experiments with a recently released relational-to-ontology benchmark suite, RODI [29], comparing BOOTOX to IncMap, -ontop- and MIRROR. RODI is designed to test relational-to-ontology mappings end-to-end: it provides an input database and a target ontology and requests complete mappings or mapped data to query. RODI is based on *scenarios*, with each scenario comprising several query tests. While RODI is extensible and can run scenarios in different application domains, it ships with a set of *default scenarios* in the conference domain that are designed to test a wide range of fundamental relational-to-ontology mapping challenges in a controlled fashion. The effectiveness of mappings is then judged by a score that mainly represents the number of query tests that return expected results on mapped data.

IncMap is designed to automatically map the target ontology directly to the input database, while BOOTOX approached this task in two steps: first, it bootstrapped an intermediate ontology and mappings from the database. Then, it aligned this intermediate, bootstrapped ontology to the target ontology as provided by the benchmark. As mentioned in Section 5.1, neither -ontop- nor MIRROR include a built-in ontology alignment system to support the importing of the target ontology provided by the benchmark. In order to be able to evaluate these systems with RODI, we aligned the generated ontologies by -ontop- and MIRROR with the target ontology using the LogMap system in a similar setup to the one used in BOOTOX.

Scenarios. RODI default scenarios are a selection of benchmark scenarios set in the conference domain, based on three different conference ontologies: CMT,

¹⁵ Note that assessing the quality of the bootstrapped ontology and mappings in an open scenario like Optique requires a huge involvement of domain experts, thus we leave it for future work.

CONFERENCE, and SIGKDD¹⁶. For each ontology, the benchmark provides a set of database instances that are to be mapped to the ontology vocabulary. While all databases are modelled to contain the same data that the ontologies might contain, they deviate largely in how close they are to their corresponding ontologies in terms of structure and modelling patterns. For each ontology, there is a number of mapping challenges that can be tested: (i) *Adjusted naming*. Here the identifier names in the databases are syntactically changed in comparison with the names in the ontology. (ii) *Cleaned hierarchies*. In this scenarios, the databases are remodelled from ground to follow widely used relational database design patterns. Most significantly, this includes cases where abstract parent classes have no corresponding table in the database, several sibling classes are jointly represented in a single table, etc. (iii) *Combined case* mixes changes made in scenarios with adjusted naming and cleaned hierarchies. (iv) *Missing FKs* represents the case of a database with no explicit referential constraints at all. (v) In the *denormalised* case, the database contains a few denormalised tables.

Results. Results show that the BOOTOX comes out in top position for seven out of nine tested scenarios (see Table 2). This shows that BOOTOX is well suited for meeting the requirements of an end-to-end scenario. Note that the results between BOOTOX and -ontop- may look very similar; however, currently RODI only provides the percentage of correctly answered queries (i.e., BOOTOX and -ontop- fail in the same queries). RODI is being extended to include a more fine grained evaluation in terms of Precision and Recall in order to take into account partially answered queries, which may reveal more differences among the evaluated systems. On a more generic note, however, results also demonstrate that none of the tested systems to date manages to solve relational-to-ontology mapping challenges with a score above 0.5. This confirms the need for specialised relational-to-ontology bootstrapping systems such as BOOTOX, which build the foundation for better solutions.

Provenance. When we evaluated BOOTOX with RODI without the use of provenance mappings, the results were slightly worse than the ones in Table 2: in three scenarios we detected unexpected answers. Then we made use of the provenance functionality of BOOTOX to analyse the source of these unexpected answers. As an outcome of such an analysis we identified and fixed the faulty mappings, hence improving the results for those scenarios.

6 Related Work

The implementation of BOOTOX has been motivated by the fact that existing ontology and mapping bootstrappers provide limited or no support for the bootstrapping tasks and quality requirements described in Section 3 (see, e.g., [34,41]). Most of the state of the art bootstrappers fail to conform with

¹⁶ <http://oei.ontologymatching.org/2015/conference/>

the ontology and mapping language standards, or they do not provide profiling capabilities for the output ontology. Moreover, to the best of our knowledge they do not provide bootstrapping of complex mappings. For historical reasons (i.e., OWL was not yet defined), former systems used RDFS and F-Logic axioms (e.g., [5,44]). Other systems have also used DLR-Lite based languages (e.g., [25]) and extensions based on SWRL (e.g., [13,24]). Regarding mapping generation, before R2RML became a W3C recommendation, system typically relied on their own native language to define mappings (e.g., D2RQ [7], Mastro [10]). To the best of our knowledge, currently only IncMap [30], MIRROR [26], -ontop- [32], and Ultrawrap [33,35] produce mappings in the R2RML language.

Among the systems using OWL or OWL 2 as the ontology language, only BOOTOX put special attention to the target ontology expressiveness. BOOTOX allows to output different ontology axioms to conform to the required OWL 2 profile as discussed in Section 4.1. Many bootstrapping systems typically use exact or min cardinality restrictions which fall outside the three OWL 2 profiles (e.g., [1,46]). Furthermore, other systems, like [6], produce an ontology that falls into OWL 2 Full due to the use of the *InverseFunctional* characteristic in both data and object properties. Finally, MIRROR, -ontop-, and Ultrawrap are conformant to the OWL 2 QL, but they do not support profiling to the other sublanguages of OWL 2.

As BOOTOX, systems like Automapper [13], Relational.OWL [23] and ROSEX [12] complement the automatically generated ontology with links to domain ontologies. However, none of these systems apply logic-based techniques to assess the consequences of such links to domain ontologies.

Special mention require the approaches in [9,36]. These approaches use (semi-automatic) ontology learning techniques to exploit the data and discover interesting patterns that can be included to enrich the ontology. Currently, BOOTOX only relies on a (fully-automatic) schema-driven generation of the ontology and mappings.

In the literature one can also find several approaches to overcome the OWA problem when dealing with data-centric applications (e.g., [27,45]). These approaches typically extend the semantics of OWL 2. The integration of these approaches in a bootstrapping scenario is, however, still an open problem.

7 Conclusions and Future Work

We presented BOOTOX, an automatic ontology and mapping bootstrapper. To the best of our knowledge, BOOTOX is the only bootstrapper that *(i)* profiles the concrete language of the output OWL ontology, *(ii)* puts special attention to the datatype conversion, *(iii)* enhances the direct mappings with provenance meta-information, *(iv)* bootstraps a range of complex mappings, and *(v)* includes a built-in logic-based ontology alignment system. Furthermore, we tested BOOTOX in a number of databases and test cases involving materialised or virtual data access. The evaluation suggests that automatic techniques to bootstrap an initial ontology and mappings are suitable to be used in practice.

We see bootstrapping as the first step towards the creation of a fully-fledged semantic data access system. Bootstrapped assets are by no means perfect, and thus they should be post-processed, validated, and extended. Our evaluation in both industrial (i.e., Optique) and research contexts has also served to guide the extension of BOOTOX with semi-automatic techniques. For example, while working with Statoil's EPDS database, we found that the discovery of implicit constraints represents a critical feature since EPDS has very few constraints, e.g., many tables are materialised views without specified primary or foreign keys. In the close future we aim at extending our evaluation to understand the limits of semi-automatic ontology and mapping bootstrappers to enable the answering of the queries in open scenarios like Optique.

Acknowledgments. This work was partially funded by the EU project Optique (FP7-ICT-318338), the Research Council of Norway through the project DOIL (RCN project n. 213115), and the EPSRC projects MaSI³, Score!, and DBOnto.

References

1. Alalwan, N., et al.: Generating OWL ontology for database integration. In: SEMAPRO (2009)
2. Arenas, M., et al.: Enabling faceted search over OWL 2 with SemFacet. In: OWLED (2014)
3. Arenas, M., et al.: Faceted search over ontology-enhanced RDF data. In: CIKM (2014)
4. Arenas, M., et al.: SemFacet: semantic faceted search over yago. In: WWW (2014)
5. Astrova, I.: Reverse engineering of relational databases to ontologies. In: Bussler, C.J., Davies, J., Fensel, D., Studer, R. (eds.) ESWS 2004. LNCS, vol. 3053, pp. 327–341. Springer, Heidelberg (2004)
6. Astrova, I.: Rules for mapping SQL relational databases to OWL ontologies. In: MTSR (2007)
7. Bizer, C., et al.: D2RQ-treating non-RDF databases as virtual RDF graphs. In: ISWC Posters (2004)
8. Calvanese, D., et al.: Tractable Reasoning and Efficient Query Answering in Description Logics: The DL-Lite Family. English. JAR **39**(3) (2007)
9. Cerbah, F., et al.: Perspectives in ontology learning. In: *Ontology Learning from Databases: Some Efficient Methods to Discover Semantic Patterns in Data*. AKA / IOS Press. Serie (2012)
10. Civili, C., et al.: Mastro studio: managing ontology-based data access applications. In: PVLDB, vol. 6, no. 12 (2013)
11. Console, M., Mora, J., Rosati, R., Santarelli, V., Savo, D.F.: Effective computation of maximal sound approximations of description logic ontologies. In: Mika, P., Tudorache, T., Bernstein, A., Welty, C., Knoblock, C., Vrandečić, D., Groth, P., Noy, N., Janowicz, K., Goble, C. (eds.) ISWC 2014, Part II. LNCS, vol. 8797, pp. 164–179. Springer, Heidelberg (2014)
12. Curino, C., Orsi, G., Panigati, E., Tanca, L.: Accessing and documenting relational databases through OWL ontologies. In: Andreasen, T., Yager, R.R., Bulskov, H., Christiansen, H., Larsen, H.L. (eds.) FQAS 2009. LNCS, vol. 5822, pp. 431–442. Springer, Heidelberg (2009)

13. Fisher, M., et al.: Use of OWL and SWRL for semantic relational database translation. In: OWLED (2008)
14. Giese, M., et al.: Optique — Zooming In on Big Data Access. *Computer* **48**(3) (2015)
15. Grau, B.C., et al.: Querying life science ontologies with SemFacet. In: SWAT4LS (2014)
16. He, B., et al.: Accessing the Deep Web. *Commun. ACM* **50**(5) (2007)
17. Horridge, M., et al.: The manchester OWL syntax. In: OWLED (2006)
18. Jiménez-Ruiz, E., et al.: Large-scale interactive ontology matching: algorithms and implementation. In: ECAI (2012)
19. Jiménez-Ruiz, E., Cuenca Grau, B.: LogMap: logic-based and scalable ontology matching. In: Aroyo, L., Welty, C., Alani, H., Taylor, J., Bernstein, A., Kagal, L., Noy, N., Blomqvist, E. (eds.) ISWC 2011, Part I. LNCS, vol. 7031, pp. 273–288. Springer, Heidelberg (2011)
20. Kharlamov, E., Solomakhina, N., Özçep, Ö.L., Zheleznyakov, D., Hubauer, T., Lamparter, S., Roshchin, M., Soyly, A., Watson, S.: How semantic technologies can enhance data access at siemens energy. In: Mika, P., Tudorache, T., Bernstein, A., Welty, C., Knoblock, C., Vrandečić, D., Groth, P., Noy, N., Janowicz, K., Goble, C. (eds.) ISWC 2014, Part I. LNCS, vol. 8796, pp. 601–619. Springer, Heidelberg (2014)
21. Kharlamov, E., Jiménez-Ruiz, E., Zheleznyakov, D., Bilidas, D., Giese, M., Haase, P., Horrocks, I., Kllapi, H., Koubarakis, M., Özçep, Ö., Rodríguez-Muro, M., Rosati, R., Schmidt, M., Schlatte, R., Soyly, A., Waaler, A.: Optique: towards OBDA systems for industry. In: Cimiano, P., Fernández, M., Lopez, V., Schlobach, S., Völker, J. (eds.) ESWC 2013. LNCS, vol. 7955, pp. 125–140. Springer, Heidelberg (2013)
22. Kharlamov, E., et al.: Ontology based access to exploration data at statoil. In: ISWC (2015)
23. de Laborda, C.P., et al.: Database to semantic web mapping using RDF query languages. In: ER (2006)
24. Levshin, D.V.: Mapping Relational Databases to the SemanticWeb with Original Meaning. *Int. J. Software and Informatics* **4**(1) (2010)
25. Lubyte, L., Tessaris, S.: Automatic extraction of ontologies wrapping relational data sources. In: Bhowmick, S.S., Küng, J., Wagner, R. (eds.) DEXA 2009. LNCS, vol. 5690, pp. 128–142. Springer, Heidelberg (2009)
26. de Medeiros, L.F., Priyatna, F., Corcho, O.: MIRROR: automatic R2RML mapping generation from relational databases. In: Cimiano, P., Frasincar, F., Houben, G.-J., Schwabe, D. (eds.) ICWE 2015. LNCS, vol. 9114, pp. 326–343. Springer, Heidelberg (2015)
27. Motik, B., et al.: Bridging the gap between OWL and relational databases. In: WWW (2007)
28. Motik, B., et al.: Parallel materialisation of datalog programs in centralised, main-memory RDF systems. In: AAAI (2014)
29. Pinkel, C., Binnig, C., Jiménez-Ruiz, E., May, W., Ritze, D., Skjæveland, M.G., Solimando, A., Kharlamov, E.: RODI: a benchmark for automatic mapping generation in relational-to-ontology data integration. In: Gandon, F., Sabou, M., Sack, H., d’Amato, C., Cudré-Mauroux, P., Zimmermann, A. (eds.) ESWC 2015. LNCS, vol. 9088, pp. 21–37. Springer, Heidelberg (2015)
30. Pinkel, C., et al.: IncMap: pay as you go matching of relational schemata to OWL ontologies. In: OM (2013)

31. Poggi, A., et al.: Linking Data to Ontologies. *J. Data Semantics* **X** (2008)
32. Rodriguez-Muro, M., et al.: Efficient SPARQL-to-SQL with R2RML Mappings. *J. Web Sem.* (to appear, 2015)
33. Sequeda, J., et al.: On directly mapping relational databases to RDF and OWL. In: *WWW* (2012)
34. Sequeda, J., et al.: Survey of directly mapping SQL databases to the semantic web. *Knowledge Eng. Review* **26**(4) (2011)
35. Sequeda, J., et al.: Ultrawrap: SPARQL Execution on Relational Data. *J. Web Sem.* **22** (2013)
36. Skjæveland, M.G., et al.: Engineering Ontology-Based Access to Real-World Data Sources. *J. Web Sem.* (to appear, 2015)
37. Solimando, A., Jiménez-Ruiz, E., Guerrini, G.: Detecting and correcting conservativity principle violations in ontology-to-ontology mappings. In: Mika, P., Tudorache, T., Bernstein, A., Welty, C., Knoblock, C., Vrandečić, D., Groth, P., Noy, N., Janowicz, K., Goble, C. (eds.) *ISWC 2014, Part II*. LNCS, vol. 8797, pp. 1–16. Springer, Heidelberg (2014)
38. Soyly, A., et al.: Experiencing optiqueVQS: a multi-paradigm and ontology-based visual query system for end users. In: *Univ. Access in the Inform. Society* (2015)
39. Soyly, A., et al.: OptiqueVQS: towards an ontology-based visual query system for big data. In: *MEDES* (2013)
40. Soyly, A., Giese, M., Jimenez-Ruiz, E., Kharlamov, E., Zheleznyakov, D., Horrocks, I.: Towards exploiting query history for adaptive ontology-based visual query formulation. In: Closs, S., Studer, R., Garoufallou, E., Sicilia, M.-A. (eds.) *MTSR 2014*. CCIS, vol. 478, pp. 107–119. Springer, Heidelberg (2014)
41. Spanos, D.-E., et al.: Bringing Relational Databases into the Semantic Web: A Survey. *Semantic Web* **3**(2) (2012)
42. Stefanoni, G., et al.: Answering conjunctive queries over EL knowledge bases with transitive and reflexive roles. In: *AAAI* (2015)
43. Stoilos, G., Stamou, G., Kollias, S.D.: A string metric for ontology alignment. In: Gil, Y., Motta, E., Benjamins, V.R., Musen, M.A. (eds.) *ISWC 2005*. LNCS, vol. 3729, pp. 624–637. Springer, Heidelberg (2005)
44. Stojanovic, L., et al.: Migrating data-intensive web sites into the semantic web. In: *SAC* (2002)
45. Tao, J., et al.: Integrity constraints in OWL. In: *AAAI* (2010)
46. Tirmizi, S.H., Sequeda, J., Miranker, D.P.: Translating SQL applications to the semantic web. In: Bhowmick, S.S., Küng, J., Wagner, R. (eds.) *DEXA 2008*. LNCS, vol. 5181, pp. 450–464. Springer, Heidelberg (2008)
47. Zhou, Y., et al.: Pay-as-you-go OWL query answering using a triple store. In: *AAAI* (2014)