



HAL
open science

xR2RML: Relational and Non-Relational Databases to RDF Mapping Language

Franck Michel, Loïc Djimenou, Catherine Faron Zucker, Johan Montagnat

► **To cite this version:**

Franck Michel, Loïc Djimenou, Catherine Faron Zucker, Johan Montagnat. xR2RML: Relational and Non-Relational Databases to RDF Mapping Language. [Research Report] ISRN I3S/RR 2014-04-FR, CNRS. 2015. hal-01066663v3

HAL Id: hal-01066663

<https://hal.science/hal-01066663v3>

Submitted on 6 Jan 2015 (v3), last revised 26 Oct 2017 (v5)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



INFORMATIQUE, SIGNAUX ET SYSTÈMES
DE SOPHIA ANTIPOLIS
UMR7271

xR2RML: Relational and Non-Relational Databases to RDF Mapping Language

Franck Michel, Loïc Djimenou, Catherine Faron-Zucker, Johan Montagnat

Equipes Modalis/Wimmics

Rapport de Recherche
ISRN I3S/RR 2014-04-FR
Version 3

Oct. 2014 - 35 pages

Document history:

Version	Date	Description
v1	2014-09-22	Initial version
v2	2014-10-08	Fix spelling mistakes.
v3	2014-10-20	Convergence with RML language definitions, moving of reference formulation from language to processor environment, fix misc. minor mistakes.

1	Introduction.....	3
1.1	Document Conventions	3
1.2	Query Languages and Data Models	3
1.3	xR2RML mapping graphs and mapping documents.....	5
1.4	xR2RML processors.....	5
2	xR2RML Overview and Examples	6
2.1	Mapping CSV data	6
2.2	Mapping JSON data	6
2.3	Mapping XML data	7
2.4	Mapping data with mixed formats.....	8
2.5	Generating an RDF collection from a list of values	9
2.6	Generating an RDF container with a referencing object map	10
3	Language description.....	12
3.1	Mapping Logical Sources to RDF with Triples Maps.....	12
3.1.1	xR2RML Triples Map.....	12
3.1.2	Defining a Logical Source	12
3.1.3	xR2RML Triples Map Iteration Model	15
3.2	Creating RDF terms with Term Maps	16
3.2.1	xR2RML Term Maps	16
3.2.1.1	Constant-, Column-, Reference- and Template-valued Term Maps	17
3.2.1.2	Term Types of Term Maps	17
3.2.1.3	Nested Term Maps.....	17
3.2.2	Referencing data elements	19
3.2.2.1	Referencing simple data elements	19
3.2.2.2	Referencing data elements with mixed data formats	20
3.2.2.3	Production of multiple RDF terms.....	21
3.2.2.4	Production of RDF collections or containers	23
3.2.3	Parsing nested structured values	25
3.2.4	Multiple Mapping Strategies	27
3.2.5	Default Term Types	28
3.3	Reference relationships between logical sources.....	28
3.3.1	Reference relationship with structured values.....	30
3.3.2	Generating RDF collection/container with a referencing object map.....	31
3.3.3	Generating RDF collection/container with a referencing object map in the relational case.....	33
4	References	35

1 Introduction

This document describes xR2RML, a language for expressing customized mappings from various types of databases (XML, object-oriented, NoSQL) to RDF datasets.

xR2RML is an extension of the R2RML [1] mapping language, and relies on some properties of the RML mapping language [4][3]. R2RML addresses the mapping of relational databases to RDF. RML extends R2RML to address the mapping of heterogeneous data formats (CSV, XML, JSON) to RDF, but does not investigate the constraints that arise when dealing with different types of heterogeneous databases. xR2RML extends this scope to a wide range of non-relational databases. This document leverages the R2RML specification and mainly describes extensions. It also leverage the RML specification [2] (Draft, 17 Sept. 2014), either explicitly reusing RML properties when applicable, or by extending existing properties. Consequently, the reader should have a good understanding of both R2RML and RML before reading this document.

xR2RML is backward compatible with R2RML.

1.1 Document Conventions

In this document, examples assume the following namespace prefix bindings unless otherwise stated:

Prefix	IRI
rr:	http://www.w3.org/ns/r2rml#
rml:	http://semweb.mmlab.be/ns/rml#
ql:	http://semweb.mmlab.be/ns/ql#
xrr:	http://www.i3s.unice.fr/ns/xr2rml#
rdf:	http://www.w3.org/1999/02/22-rdf-syntax-ns#
rdfs:	http://www.w3.org/2000/01/rdf-schema#
xsd:	http://www.w3.org/2001/XMLSchema#
ex:	http://example.com/ns#

Vocabulary definitions are formatted in such grey boxes:

Definition

1.2 Query Languages and Data Models

R2RML is specifically focused on the translation of relational databases into RDF datasets. xR2RML extends this scope to non-relational databases. Although it is illusory to seek universal support of any database, our endeavor is to equally support most common relational and non-relational databases. In our approach, we more specifically analyze the requirements to support NoSQL and XML databases. Yet, xR2RML may not support all NoSQL databases, given the large variety of systems behind this word. Nevertheless, we argue that our work can be generalized to some other types of database, for instance object-oriented and directory (LDAP) databases.

Query languages:

Relational databases all support ANSI SQL (or at least a subset of it), and most XML databases support XQuery and XPath (which is subset of XQuery). By contrast, NoSQL is a catch-all term

referring to very diverse systems [6][5]. They have heterogeneous access methods ranging from low-level APIs to expressive query languages. Despite several propositions of common query language (N1QL¹, UnQL², SQL++ [8], ArangoDB Query Language³, CloudMdsQL[7]), no consensus has emerged yet, that would fit most NoSQL databases. Therefore, until a standard eventually arises, xR2RML must be agile enough to cope with various query languages and protocols.

Remark: R2RML relies on the ability of relational databases to support a declarative query language. xR2RML does the same assumption with regards to other types of databases, although this may be limitative. For instance, most NoSQL key-value stores provide simple key-based operations (such as put, get, delete) by means of APIs used in imperative programming languages, but they hardly provide a declarative query language. If such a system is to be mapped to RDF using xR2RML, an xR2RML processor should implement a mechanism to bridge this gap. In other words the xR2RML processor should define its own query language to be interpreted and compiled using an imperative programming language.

Data models:

Relational databases comply with a row-based model in which all rows have the same columns. NoSQL systems have heterogeneous data models (key-value store, document store, extensible column-store, graph store). Some of them also comply with the row-based model, such as extensible column-stores (also known as column family stores) with the difference that all rows do not necessarily share the same columns (BigTables, Cassandra...). Other databases in which data is formatted in JSON (document stores such as MongoDB, CouchDB...) or XML (native XM databases such as BaseX, eXistDB) can hardly be reduced to a row-based model. JSON or XML documents consist of structured (hierarchical) values representing collections or key-value associations:

- A JSON dictionary is an ordered association of keys and values, both keys and values may be of any type. A JSON array is an ordered collection of elements, it is a specific case of dictionary in which keys are implicit integer indexes: 0, 1, 2, etc.
- Similarly, a set of XML elements having the same parent element can be seen as an ordered association of keys (element names) and values (element values). A set of XML elements of the same type, having the same parent element, can be seen as an ordered collection. Besides, attributes of an XML element can be seen as a specific type of key-value association.

The model of structured values described above for JSON and XML can easily be applied to other databases. In an object-oriented model, an object can be approximated by as a key-value association: keys are attribute names while values are either a scalar, another object (composition or aggregation relationship), or a collection (depending on capabilities of the modeling language: list, map, etc). Similarly, an LDAP directory is organized as a tree. Each node has an identifier and a set of attributes represented as "name=value" that are nothing else than a key-value association. A set of attributes with the same name can be interpreted as either as a collection or a key-value association in which keys are not unique. Thus, xR2RML must be able to map data elements from rows as well as structured values (nested collections and key-value associations) to RDF terms.

Note: Below, we shall use the term "structured values" to refer to collections and key-value associations, whatever the representation syntax used.

¹ <http://www.couchbase.com/communities/n1ql>

² <http://unql.sqlite.org/index.html>

³ <http://docs.arangodb.org/Aql/README.html>

1.3 xR2RML mapping graphs and mapping documents

An **xR2RML mapping** defines a mapping from a database to RDF.

An xR2RML mapping is represented as an RDF graph called an **xR2RML mapping graph**.

An **xR2RML mapping document** is any document written in the Turtle RDF syntax that encodes an xR2RML mapping graph.

Any R2RML mapping graph is a valid xR2RML mapping graph.

1.4 xR2RML processors

An xR2RML processor is a system that, given an xR2RML mapping and an input database, provides access to the output RDF dataset.

An xR2RML processor has access to an execution environment consisting of:

- a connection to the input database,
- a reference formulation applicable to results of queries run against the connection,
- a base IRI used in resolving relative IRIs produced by the R2RML mapping.

The connection is used by the xR2RML processor to evaluate queries against the input database. It must be established with sufficient privileges for read access to all database elements (tables, views, documents, objects...) that are referenced in the xR2RML mapping. How the connection is established, or how users are authenticated against the database, is outside the scope of this document.

The reference formulation names a syntax used in triples maps to reference data elements within results of a query run against the database connection. It is the responsibility of an xR2RML processor developer to document supported reference formulations and how to configure one. Typically, this should include the following options: column name (applicable not only to row-based data models but also to any row-based query results), JSONPath, XPath or LDAP path expressions, etc. If the reference formulation is not provided, it defaults to column names, in order to ensure backward compatibility with R2RML.

2 xR2RML Overview and Examples

This section gives a brief overview of the xR2RML mapping language, followed by simple examples of mapping various types of database to RDF.

An xR2RML mapping refers to logical sources to retrieve data from the input database. A logical source can be either an *xR2RML base table* (for input databases where tables or views exist, such as SQL views), or an *xR2RML view* representing the result of executing a query against the input database. An xR2RML processor is provided an xR2RML mapping description, a connection to the database and a reference formulation that specifies the syntax used to refer to data elements retrieved from the input database: this can be a column name in the case of row-based systems (RDB, extensible column-store), a JSONPath expression in case of a NoSQL document store, an XPath expression in case of an XML native database, an LDAP path in case of an LDAP directory, etc.

Each logical source is mapped to RDF using a triples map. As in R2RML, a triples map consists of a subject map that generates the subject of all RDF triples that will be generated from data elements, and multiple predicate-object maps that produce the predicate and object terms of triples.

2.1 Mapping CSV data

The input database in the example below consists of one CSV document. It is assumed that the xR2RML processor is provided a connection to that file, e.g. by means of a descriptor to a file on the local file system or a URL to locate the file on a web server. Data elements will be referenced using column names, i.e. the reference formulation passed to the xR2RML processor.

As a CSV file simply consists of a single unnamed table, the logical source can simply be left empty.

Input data	title, year, director Manhattan, 1979, Woody Allen Annie Hall, 1979, Woody Allen 2046, 2004, Wong Kar-wai In the Mood for Love, 2000, Wong Kar-wai
Mapping graph	<#CSVTriplesMap> rr:subjectMap [rr:template "http://example.org/movie/{title}";]; rr:predicateObjectMap [rr:predicate ex:directedBy; rr:objectMap [xrr:reference "director";];].
RDF triples produced	<http://example.org/movie/Manhattan> ex:directedBy "Woody Allen". <http://example.org/movie/Annie%20Hall> ex:directedBy "Woody Allen". <http://example.org/movie/2046> ex:directedBy "Wong Kar-wai". <http://example.org/movie/In%20the%20Mood%20for%20Love> ex:directedBy "Wong Kar-wai".

2.2 Mapping JSON data

The input database in the example below is a MongoDB database (document store). The query in the logical source retrieves one JSON document from collection "movies", that lists movie directors and movies they directed.

Without further instruction on how to parse the document, JSONPath expressions referring to data elements in the subject and object map will be evaluated against the whole document. For instance, a subject using expression "\$.directors.name" will return two terms, while an object map using expression "\$.directors.movies.*" will return four terms, one for each movie whatever its director. This will result in mixing up directors and movies. To avoid this, an `rml:iterator` property is added to the logical source, specifying that the triples map iteration should occur on each element of the array of directors.

References to data elements (`rr:template`, `xrr:reference`), as well as the iterator pattern, are expressed in JSONPath (i.e. the reference formulation, passed to the xR2RML processor along with the database connection).

Input data	<pre>{ "directors": [{ "name": "Wong Kar-wai", "movies": ["2046", "In the Mood for Love"] }, { "name": "Woody Allen", "movies": ["Manhattan", "Annie Hall"] }]}</pre>
Mapping graph	<pre><#Directors> xrr:logicalSource [xrr:query "db.movies.find({ directors: { \$exists: true } })"; rml:iterator "\$.directors.*";]; rr:subjectMap [rr:template "http://example.org/director/{\$.name}";]; rr:predicateObjectMap [rr:predicate ex:directed; rr:objectMap [xrr:reference "\$.movies.*";];].</pre>
RDF triples produced	<pre><http://example.org/director/Woody%20Allen> ex:directed "Manhattan". <http://example.org/director/Woody%20Allen> ex:directed "Annie Hall". <http://example.org/director/Wong%20Kar-wai> ex:directed "2046". <http://example.org/director/Wong%20Kar-wai> ex:directed "In the Mood for Love".</pre>

2.3 Mapping XML data

The example below is very similar to the previous one, using an XML database supporting XQuery. The query in the logical source retrieves "director" XML elements. As in the previous example, to avoid mixing up directors and movies, an `rml:iterator` property is added to the logical source, specifying that the triples map iteration should occur on each "director" XML element.

References to data elements (`rr:template`, `xrr:reference`), as well as the iterator pattern, use the XPath syntax (i.e. the reference formulation, passed to the xR2RML processor along with the database connection).

Input data	<pre><directors> <director name="Wong Kar-wai"> <movies> <movie>2046</movie> <movie>In the Mood for Love</movie> </movies> </director></pre>
------------	--

	<pre><director name="Woody Allen"> <movies> <movie>Manhattan</movie> <movie>Annie Hall</movie> </movies> </director> </directors></pre>
Mapping graph	<pre><#Directors> xrr:logicalSource [xrr:query "//directors/director"; rml:iterator "//director";]; rr:subjectMap [rr:template "http://example.org/director/{/director/@name}";]; rr:predicateObjectMap [rr:predicate ex:directed; rr:objectMap [xrr:reference "//movie";];].</pre>
RDF triples produced	<pre><http://example.org/director/Woody%20Allen> ex:directed "Manhattan". <http://example.org/director/Woody%20Allen> ex:directed "Annie Hall". <http://example.org/director/Wong%20Kar-wai> ex:directed "2046". <http://example.org/director/Wong%20Kar-wai> ex:directed "In the Mood for Love".</pre>

2.4 Mapping data with mixed formats

In some use cases, it is common to store values in a format which is not the native database format. For instance, an application designer may choose to embed JSON, CSV, or XML values in the cells of a relational database, for performance concerns or application design constraints.

xR2RML allows to reference data elements within such mixed contents with *mixed-syntax paths*. A path with mixed-syntax consists of the concatenation of several path expressions separated by the slash '/' character. Each individual path is enclosed in a syntax path constructor: Column(column-name), CSV(column-name), TSV(column-name), JSONPath(JSONPath-expression), XPath(XPath-expression).

In the example below, the logical source is a relational table with two columns. The second column, Movies, holds values formatted as JSON arrays. The xrr:reference property of the object map uses a mixed-syntax path: *Column(Movies)/JSONPath(\$.*)*. This expression selects values from column "Movies" and evaluates JSONPath expression "\$.*" against the values.

Input data	<p>Table DIRECTORS:</p> <table border="1"> <thead> <tr> <th>Name</th> <th>Movies</th> </tr> </thead> <tbody> <tr> <td>Wong Kar-wai</td> <td>["2046", "In the Mood for Love"]</td> </tr> <tr> <td>Woody Allen</td> <td>["Manhattan", "Annie Hall"]</td> </tr> </tbody> </table>	Name	Movies	Wong Kar-wai	["2046", "In the Mood for Love"]	Woody Allen	["Manhattan", "Annie Hall"]
Name	Movies						
Wong Kar-wai	["2046", "In the Mood for Love"]						
Woody Allen	["Manhattan", "Annie Hall"]						
Mapping graph	<pre><#Directors></pre>						

	<pre> rr:logicalTable [rr:tableName "DIRECTORS";]; rr:subjectMap [rr:template "http://example.org/director/{Name}";]; rr:predicateObjectMap [rr:predicate ex:directed; rr:objectMap [xrr:reference "Column(Movies)/JSONPath(\$.*)";];]. </pre>
RDF triples produced	<pre> <http://example.org/director/Woody%20Allen> ex:directed "Manhattan". <http://example.org/director/Woody%20Allen> ex:directed "Annie Hall". <http://example.org/director/Wong%20Kar-wai> ex:directed "2046". <http://example.org/director/Wong%20Kar-wai> ex:directed "In the Mood for Love". </pre>

2.5 Generating an RDF collection from a list of values

As illustrated by the previous example, several RDF terms can be generated by a term map during one triples map iteration step. While this can lead to the generation of several triples, this can also be used to generate hierarchical values in the form of RDF collections or containers.

The example below was already presented in section 2.2. In the object map we simply add an `rr:termType` property with value `xrr:RdfList`. All RDF terms produced by the object map during one triples map iteration step are then grouped as members of one term of type `rdf:List` (denoted as "(a b c...)" in Turtle syntax).

Finally, we want to add a language tag to the movie titles. The object map describes the generation of RDF lists, hence it would not make sense to add an `rr:language` property. To state properties about the members of the generated RDF list, we need a nested term map, introduced by the `xrr:nestedTermMap` property. A nested term map accepts the same properties as a term map, but it applies to members of RDF collection/container terms generated by its parent term map.

Input data	<pre> { "directors": [{ "name": "Wong Kar-wai", "movies": ["2046", "In the Mood for Love"] }, { "name": "Woody Allen", "movies": ["Manhattan", "Annie Hall"] }]} </pre>
Mapping graph	<pre> <#Directors> xrr:logicalSource [xrr:query "db.movies.find({ directors: { \$exists: true } })"; rml:iterator "\$.directors.*";]; rr:subjectMap [rr:template "http://example.org/director/{\$.name}";]; rr:predicateObjectMap [rr:predicate ex:directed; rr:objectMap [xrr:reference "\$.movies.*"; rr:termType xrr:RdfList; xrr:nestedTermMap [rr:language "en";]];]. </pre>

RDF triples produced	<pre><http://example.org/director/Woody%20Allen> ex:directed [a rdf:List; rdf:_1 "Manhattan"@en; rdf:_2 "Annie Hall"@en;] <http://example.org/director/<Wong%20Kar-wai> ex:directed ("2046"@en "In the Mood for Love"@en).</pre>
----------------------	---

2.6 Generating an RDF container with a referencing object map

The example below uses a referencing object map to describe a cross-reference relationship between logical resources. In addition, it generates an RDF bag from the result of the join condition in the referencing object map.

Triples map <#Movies> generates IRIs for the movies. The referencing object map in triples map <#Directors> uses IRI generated in triples map <#Movies> as the members of an RDF bag (property `rr:term:Type` `xrr:RdfBag`).

The join condition in triples map <#Directors> produces a result if a movie title (`rr:parent "$.title"`) matches at least one movie in the list of movies associated with each director (`rr:child "$.movies.*"`).

Input data	<pre>{ "directors": [{ "name": "Wong Kar-wai", "movies": ["2046", "In the Mood for Love"] }, { "name": "Woody Allen", "movies": ["Manhattan", "Annie Hall"] }]} { "movies": [{ "title": "Manhattan", "year": "1979" }, { "title": "Annie Hall", "year": "1977" }, { "title": "2046", "year": "2004" }, { "title": "In the Mood for Love", year: "2000"}]}</pre>
Mapping graph	<pre><#Movies> xrr:logicalSource [xrr:query "db.movies.find({ movies: { \$exists: true } })"; rml:iterator "\$.movies.*";]; rr:subjectMap [rr:template "http://example.org/movies/{\$.title}";]. <#Directors> xrr:logicalSource [xrr:query "db.movies.find({ directors: { \$exists: true } })"; rml:iterator "\$.directors.*";]; rr:subjectMap [rr:template "http://example.org/director/{\$.name}";]; rr:predicateObjectMap [rr:predicate ex:directed; rr:objectMap [rr:parentTriplesMap <#Movies>;</pre>

	<pre> rr:joinCondition [rr:child "\$.movies.*"; rr:parent "\$.title";]; rr:termType xrr:RdfBag;];]. </pre>
Generated RDF triples	<pre> <http://example.org/director/Woody%20Allen> ex:directed [a rdf:Bag; rdf:_1 <http://example.org/movie/Manhattan>; rdf:_1 <http://example.org/movie/Annie%20Hall>.]. <http://example.org/director/<Wong%20Kar-wai> ex:directed [a rdf:Bag; rdf:_1 <http://example.org/movie/2046>; rdf:_2 <http://example.org/movie/In%20the%20Mood%20for%20Love>.]. </pre>

3 Language description

3.1 Mapping Logical Sources to RDF with Triples Maps

3.1.1 xR2RML Triples Map

A triples map specifies a rule for translating data elements of a logical source to zero or more RDF triples. The RDF triples generated from one data element (row, document, set of XML elements, etc) in the logical source all share the same subject.

An xR2RML triples map extends R2RML triples map by referencing a logical source (next section) instead of a logical table. An xR2RML triples map is represented by a resource that references the following resources:

- It must have exactly one **xrr:logicalSource** property. Its object is a logical source that specifies a table or a query result to be mapped to triples.
- It must have exactly one subject map that specifies how to generate a subject for each data element of the logical source (row, document, set of XML elements, etc). A subject map may be specified in two ways:
 - using the `rr:subjectMap` property, whose value must be the subject map, or
 - using the constant shortcut property `rr:subject`.
- It may have zero or more `rr:predicateObjectMap` properties, whose values must be predicate-object maps. They specify pairs of predicate maps and object maps that, together with the subjects generated by the subject map, may form one or more RDF triples for each data element.

3.1.2 Defining a Logical Source

An R2RML logical table is a data set on which the triples map applies: this may be a relational table, SQL view, or the result of a valid SQL query. xR2RML intends to cope with a wide, non-limited scope of data sources, both relational and non-relational.

While relational databases have clearly identified commonalities (row-based data model, ANSI SQL compatibility), non-relational databases show much more heterogeneity. NoSQL, in particular, refers to a large variety of systems having specific data models (key-value, document, column, graph), and access methods (various protocols, APIs, query languages). Consequently, unless a standard NoSQL query language eventually arises, xR2RML must be agile enough to cope with various query languages and protocols in a transparent manner. Similarly, the way data elements are referenced within query results depends on the data model, query language and API capabilities. Thus, to be effective, xR2RML must remain free of any reference to specific expression syntaxes.

xR2RML vs. RML: In conformance with R2RML principles, database connection details (connection string, URL, login, etc) are out of the scope of xR2RML. These are part of the processor environment. In RML on the other hand, a logical source (`rml:LogicalSource`) points to a source containing the data to be mapped, denoted by property `rml:source`. The source may be a file on the local file system or retrieved from a web URL. This makes it difficult for xR2RML to extend RML's logical source concept. Consequently the xR2RML logical source does not extend `rml:LogicalSource`, but commonalities are addressed by using or extending some RML properties (`rml:referenceFormulation`, `rml:query`, `rml:iterator`).

xR2RML defines several extensions to describe an input database:

A **logical source** (property **xrr:logicalSource**) extends the R2RML concept of logical table (property `rr:logicalTable`) in the case of non-relational databases. A logical source is

the result of a query applied to the input database, to be mapped to RDF triples. A logical source is either an **xR2RML base table** or an **xR2RML view**.

An **xR2RML base table** is a logical source containing data from a table in the input database. It simply extends the concept of *R2RML table or view* to the context of tabular databases beyond relational databases (e.g. CSV, extensible column store). An xR2RML base table is represented by a resource that has exactly one **rr:tableName** property. Its object is a string literal representing the table name.

An **xR2RML view** is a logical source whose content is the result of executing a query against the input database. It is represented by a resource that has exactly one **xrr:query** property. Property xrr:query extends RML property rml:query. The object of property xrr:query is a string literal representing a valid expression with regards to the query language supported by the input database.

A logical source may have one data property **rml:iterator** that specifies the iteration pattern to apply on data retrieved from the input database (section 3.1.3). Its object is an expression written according to the syntax specified by the reference formulation (section 1.4). The rml:iterator property is ignored in the context of tabular result sets in which data is accessed by column names.

xR2RML vs. RML:

- RML introduced the reference formulation concept (property rml:referenceFormulation of a logical source) to name the syntax of data element references. These are ql:JSONPath for JSON data, ql:XPath for XML data, and rr:SQL2008 for relational databases. xR2RML adheres to R2RML's principle that database specificities are out of the scope of the mapping language. Moreover, it must remain free of explicit reference to specific syntaxes, as already mentioned. As a result, the concept of reference formulation is not part of the mapping language, instead it is a configuration parameter passed to an xR2RML processor along with the database connection (section 1.4).
- RML ontology defines property rml:query that subsumes properties rr:sqlQuery, rml:xmlQuery and rml:queryLanguage. However none of these properties are described nor used at the time of writing in the language specification and articles. We extend the rml:query property to allow room for RML authors to use it in their own way eventually.

xR2RML logical source and R2RML logical table definitions may equally be used in the case of a relational database. Examples:

R2RML logical table	Equivalent xR2RML logical source
[] rr:logicalTable [rr:tableName "SOME_TABLE".]	[] xrr:logicalSource [rr:tableName "SOME_TABLE";]
[] rr:logicalTable [rr:sqlQuery "SELECT NAME, DATE FROM MOVIES".]	[] xrr:logicalSource [xrr:query "SELECT NAME, DATE FROM MOVIES".]

The table below shows examples of xR2RML logical source definition with different flavors of input databases.

Type of database	Logical source definition
------------------	---------------------------

<p>Relational database</p>	<pre>[] rr:logicalTable [rr:sqlQuery "" SELECT TITLE FROM MOVIES WHERE YEAR > 1980 ORDER BY YEAR LIMIT 10"";];</pre>
<p>XML file. The xR2RML processor is provided with a file URL, e.g. <code>http://example.org/movies.xml</code>, and the XPath reference formulation. Therefore no further query is required (no <code>xrr:query</code> property). An iterator defines the pattern of XML elements to iterate on. XPath is used to refer to data elements within the XML data returned by the database.</p>	<pre>[] xrr:logicalSource [rml:iterator "//movie";];</pre>
<p>REST-based web service returning an XML stream based on parameters passed in the HTTP GET query string. The service URL (e.g. <code>http://example.org/service</code>) and XPath reference formulation are passed to the xR2RML processor by configuration, while the HTTP query string is provided by the <code>xrr:query</code> property.</p>	<pre>[] xrr:logicalSource [xrr:query "?minYear=1980&limit=10"; rml:iterator "//movie";];</pre>
<p>XML database supporting XQuery. XPath is used to describe the iterator and later on to refer to data elements within the XML data returned by the database.</p>	<pre>[] xrr:logicalSource [xrr:query ""for \$i in //movies/movie where \$i/year gt 1980 order by \$i/@title return \$i/@title""; rml:iterator "//movie";];</pre>
<p>JSON file. The xR2RML processor is provided with file URL and the JSONPath reference formulation. No further query is required (no <code>xrr:query</code> property). An iterator defines the pattern to iterate on.</p>	<pre>[] xrr:logicalSource [rml:iterator "\$.movies.*";];</pre>
<p>MongoDB database (document store). MongoDB Shell Query Language is used to search for documents in collection "movies". Then, JSONPath is used to refer to data elements within the JSON documents returned by the database.</p>	<pre>[] xrr:logicalSource [xrr:query ""db.movies.find({"year":{"\$gt: 1980}})"";];</pre>
<p>Cassandra (extensible column store) using Cassandra Query Language (CQL). Data element are referenced by column name (reference formulation passed to the xR2RML processor).</p>	<pre>[] xrr:logicalSource [xrr:query ""SELECT TITLE, YEAR FROM Movies LIMIT 10"";];</pre>
<p>AllegroGraph (RDF graph store) using SPARQL. The column name reference formulation is applied to a SPARQL result set: the result set is considered as a table in which columns are named after variable names.</p>	<pre>[] xrr:logicalSource [xrr:query ""select ?title ?year where { ?movie a ex:Movie; ex:name ? title; ex:year ?year. } filter (?year > "1980"^^xsd:integer) order by ?year limit 10"";];</pre>

3.1.3 xR2RML Triples Map Iteration Model

In R2RML, the row-based iteration implicitly occurs on a set of rows read from a logical table. xR2RML applies this principle to other row-based systems such as CSV/TSV files and extensible column stores, but also SPARQL result sets as already mentioned. In the context of non row-based databases, the model is implicitly extended to a **document-based iteration model**: a document is basically one result of a result set, whatever the form of such result. Typically, the document-based iteration applies to a set of JSON documents retrieved from a NoSQL document store, or a set of XML documents retrieved from an XML database. In the case of data sources that do not natively provide iterators over results, for instance a simple XML file or a web service returning an XML stream at once, then a single iteration occurs on the whole document retrieved. The latter case falls back to RML's standard behavior.

The document-based iteration model alone may not be sufficient to fulfill any iteration needs. For instance, a document store returns JSON documents. However it may be needed to iterate on selected entries of the document. This is illustrated by the following example.

The following JSON document is retrieved by a request to a MongoDB document store, it describes movie directors and respective movies.

```
{ "directors": [
  { "name": "Wong Kar-wai",
    "movies": ["2046", "In the Mood for Love"] },
  { "name": "Woody Allen",
    "movies": ["Manhattan", "Annie Hall"] }
]}
```

The following mapping graph is used:

```
<#Directors>
  xrr:logicalSource [
    xrr:query "db.movies.find( { directors: { $exists: true } } )";
  ];
  rr:subjectMap [
    rr:template "http://example.org/director/{$.name}";
  ];
  rr:predicateObjectMap [
    rr:predicate ex:directed;
    rr:objectMap [ xrr:reference "$.directors.*.movies.*"; ];
  ].
```

In this mapping, the subject map returns two terms (one per director) while the object map returns four terms (one per movie in the document). Consequently, triples are generated that mix up all directors and movies:

```
<http://example.org/director/Woody%20Allen> ex:directed "Manhattan".
<http://example.org/director/Woody%20Allen> ex:directed "Annie Hall".
<http://example.org/director/Woody%20Allen> ex:directed "2046".
<http://example.org/director/Woody%20Allen> ex:directed "In the Mood for Love".
<http://example.org/director/Wong%20Kar-wai> ex:directed "Manhattan".
<http://example.org/director/Wong%20Kar-wai> ex:directed "Annie Hall".
<http://example.org/director/Wong%20Kar-wai> ex:directed "2046".
```



```
<http://example.org/director/Wong%20Kar-wai> ex:directed "In the Mood for Love".
```

To deal with such cases, xR2RML relies on the concept of iterator defined in RML:

An **iterator** is defined within an xR2RML logical source by means of the **rml:iterator** property. It specifies the iteration pattern to apply to data retrieved from the input database.

The object of an rml:iterator property is a valid path expression written using the syntax specified by the reference formulation (section 1.4).

With the rml:iterator property, the previous example is modified as shown below:

```
<#Directors>
  xrr:logicalSource [
    xrr:query "db.movies.find( { directors: { $exists: true} } )";
    rml:iterator "$.directors.*";
  ];
  rr:subjectMap [
    rr:template "http://example.org/director/{$.name}";
  ];
  rr:predicateObjectMap [
    rr:predicate ex:directed;
    rr:objectMap [ xrr:reference "$.directors.*.movies.*"; ];
  ].
```

The rml:iterator property indicates that, within the document retrieved, the triples map iteration should be performed on each director element rather than on the whole document, thus producing the expected results:

```
<http://example.org/director/Woody%20Allen> ex:directed "Manhattan".
<http://example.org/director/Woody%20Allen> ex:directed "Annie Hall".
<http://example.org/director/Wong%20Kar-wai> ex:directed "2046".
<http://example.org/director/Wong%20Kar-wai> ex:directed "In the Mood for Love".
```

3.2 Creating RDF terms with Term Maps

3.2.1 xR2RML Term Maps

R2RML defines a term map as a function that generates RDF terms from a logical table row.

A term map is either a subject map, predicate map, object map or graph map.

A term map must be exactly one of the following:

- a constant-valued term map (property rr:constant)
- a column-valued term map (property rr:column)
- a template-valued term map (property rr:template).

R2RML treats all values from the input database as literals expressed in built-in data types. To deal with structured values such as collections or key-value associations, xR2RML term maps extend R2RML term maps so that structured values can be parsed, and data elements within structured values can be selected to build RDF terms. xR2RML extensions are described in this section.

3.2.1.1 Constant-, Column-, Reference- and Template-valued Term Maps

R2RML properties `rr:column` and `rr:template` reference columns of a logical table. xR2RML not only references columns but also any data element within structured values. xR2RML relies on the **rml:reference**, that extends property `rr:column`. Its object is a column name (in the case of an RDB, CSV/TSV file, extensible column store, SPARQL result set, etc), an XPath expression (in the case of XML data) and a JSONPath expression (in the case of JSON data). Furthermore, xR2RML introduces the possibility to reference data elements in data with mixed formats (§3.2.2.2). Thus, xR2RML extends property `rml:reference` with property **xrr:reference**. This leads to the following definition of an xR2RML term map. xR2RML changes to R2RML are highlighted.

A constant-valued term map is represented by a resource that has exactly one `rr:constant` property. A constant-valued term map always generates the same RDF term.

A column-valued term map has exactly one `rr:column` property. The value of the `rr:column` property is a valid column name.

A reference-valued term map has exactly one `xrr:reference` property. The value of the `xrr:reference` property is a valid reference to a data element, expressed using the syntax defined by the reference formulation (section 1.4).

A template-valued term map has exactly one `rr:template` property. The value of the `rr:template` property is a valid string template. A string template is a format string used to build strings from multiple components. It uses valid references to data elements by enclosing them in curly braces ("`{`" and "`}`"). Each reference is expressed using the syntax defined by the reference formulation (section 1.4).

3.2.1.2 Term Types of Term Maps

RDF terms generated by a term map have a term type (`rr:termType`) that may be one of the three R2RML term types: `rr:Literal`, `rr:BlankNode` or `rr:IRI`.

xR2RML extends the `rr:termType` property with four new values, hereafter referred to as **RDF collection or container term types**:

- A term map with **xrr:RdfList** as value of property `rr:termType` generates an RDF collection of type `rdf:List`;
- A term map with **xrr:RdfSeq**: as value of property `rr:termType` generates an RDF container of type `rdf:Seq`;
- A term map with **xrr:RdfBag**: as value of property `rr:termType` generates an RDF container of type `rdf:Bag`;
- A term map with **xrr:RdfAlt**: as value of property `rr:termType` generates an RDF container of type `rdf:Alt`.

3.2.1.3 Nested Term Maps

Structured data such as JSON or XML documents commonly have more than one level of nesting, resulting in tree-like values that may need to be parsed in depth to nest RDF collections and containers, e.g. to build an RDF collection of RDF collections.

An xR2RML term map may have an **xrr:nestedTermMap** property, whose range is the **xrr:NestedTermMap** class.

In a column-valued or reference-valued term map, the `xrr:nestedTermMap` property describes how to translate values produced by the `rr:column` or `xrr:reference` properties into RDF terms.

In a template-valued term map, the `xrr:nestedTermMap` property describes how to translate values produced by applying the template string to input values into RDF terms.

In a constant-valued term map, it is invalid to define a nested term map.

A **nested term map** may have the properties below:

- **xrr:reference** bears the same semantics as in the context of a term map. Its object is a valid path expression (possibly a mixed-syntax path, see §3.2.2.2). Evaluation of the path expression is performed against values retrieved by the parent of the nested term map. This parent may be a term map or a nested term map.
- **rr:template** bears the same semantics as in the context of a term map. References enclosed in capturing curly braces are valid path expressions (possibly mixed-syntax paths), they apply to values retrieved in the parent of the nested term map. This parent may be a term map or a nested term map.
- **xrr:nestedTermMap** is used to recursively parse any depth of nested structured values;
- **rr:termType** bears the same semantics as in the context of a term map;
- **rr:language** bears the same semantics as defined in R2RML;
- **rr:datatype** bears the same semantics as defined in R2RML.

A **simple nested term map** is a nested term map that has no `xrr:reference` nor `rr:template` property. A simple nested term map is used to qualify terms of an RDF collection or container generated by its parent term map or nested term map, i.e. assign them an optional term type, data type or language tag.

A **reference-valued nested term map** is a nested term map that has exactly one `xrr:reference` property.

A **template-valued nested term map** is a nested term map that has exactly one `rr:template` property.

xrr:nestedTermMap vs. rr:termType:

A nested term map *N* describes how to translate into RDF terms values produced by its parent *P*, *P* may be a term map or a nested term map.

If *P* has no `rr:termType` property, it simply returns values produced by *N*, therefore the term type of *P* is that of *N*.

If *P* has an `rr:termType` property with an RDF collection or container term type as object, then values produced by *N* will be assembled in an RDF collection or container.

Lastly, *P* should not have an `rr:termType` property with an R2RML term type (literal, blank node, IRI) or in other words, a nested term map cannot be used in the context of a term map or nested term map that has an R2RML term type (literal, IRI, blank node). Thus:

If a term map or nested term map has an `xrr:nestedTermMap` property, then it should have either no `rr:termType` property or an `rr:termType` property with an RDF collection or container term type. Formally:

?P is an `rr:TermMap` or `xrr:NestedTermMap`.

?P `xrr:nestedTermMap` ?N.

```
?P rr:termType ?tt.
⇒ ?tt is one of xrr:RdfList, xrr:RdfSeq, xrr:RdfBag or xrr:RdfAlt
```

A term map or nested term map with an RDF collection or container term type and no `xrr:nestedTermMap` property is assumed to have a **default `xrr:nestedTermMap` property** defined as follows:

- If the parent term map or nested term map is reference-valued:
`xrr:nestedTermMap [rr:termType rr:Literal];`
- If the parent term map or nested term map is template-valued:
`xrr:nestedTermMap [rr:termType rr:IRI];`

Finally, as defined in R2RML, properties `rr:language` and `rr:datatype` apply when generating literals only:

A term map or nested term map may have an `rr:language` or `rr:datatype` property only if its term type is `rr:Literal` (either stated by property `rr:termType` or inferred as a default value).

Nested term maps are exemplified in section 3.2.3.

3.2.2 Referencing data elements

3.2.2.1 Referencing simple data elements

The table below exemplifies the use of properties `rr:column`, `xrr:reference` and `rr:template` to reference simple data elements (i.e. with non-mixed data formats) from the logical source.

Logical source	Term map
Relational database: either <code>rr:column</code> or <code>xrr:reference</code> can be used to name a column.	<pre>[] rr:column "NAME". [] xrr:reference "NAME". [] rr:template "{NAME}".</pre>
Extensible column store: properties <code>xrr:reference</code> and <code>rr:template</code> reference data elements by column names.	<pre>[] xrr:reference "NAME". [] rr:template "{NAME}".</pre>
XML database supporting: properties <code>xrr:reference</code> and <code>rr:template</code> reference data elements by XPath expressions.	<pre>[] xrr:reference "//name". [] rr:template "{//name }".</pre>
NoSQL document store: <code>xrr:reference</code> and <code>rr:template</code> reference data elements using a valid JSONPath expression.	<pre>[] xrr:reference "\$.name". [] rr:template "{\$.name }".</pre>
RDF graph store accessed using a SPARQL SELECT query: <code>xrr:reference</code> and <code>rr:template</code> reference data elements by name of variable returned in the SPARQL result set.	<pre>[] xrr:reference "?name". [] rr:template "{?name }".</pre>

Remark: If a term map references a structured value but does not parse it using a nested term map, then generated RDF terms are string literals containing a simple serialization of structured values. Example:

Input data	{ "person": { "FirstName":"John", "LastName":"Smith" } }
Term map	[] rr:objectMap [xrr:reference "\$.person";];
Generated RDF term	The structured value matching the JSONPath expression "\$.person" is returned as a string literal in quotes: <code>'{ "FirstName":"John", "LastName":"Smith" }'</code>

3.2.2.2 Referencing data elements with mixed data formats

In some use cases, databases are commonly used to store values written in a data format that they cannot interpret. For instance, an application designer may choose to embed JSON, CSV, or XML values in the cells of a relational table, for performance concerns or application design constraints.

To reference data elements within such mixed contents, xR2RML allows a term map to reference data elements with **mixed-syntax paths**:

Properties xrr:reference and rr:template may use **mixed-syntax paths** to reference data elements across data in different formats. A mixed-syntax path consists of the concatenation of several path expressions separated by the slash '/' character. Each individual path is enclosed in a **syntax path constructor** naming the path syntax explicitly. Existing constructors are:

- **Column**(column-name): applies to row/column databases such as relational database and extensible column-store.
- **CSV**(column-name), **TSV**(column-name): applies to data formatted as comma-separated or tab-separated values. Column-name may be a 0-based column index, or an actual column name if a head line provides column names.
- **JSONPath**(JSONPath-expression): applies to any data formatted in JSON.
- **XPath**(XPath-expression): applies to any data formatted in XML.

In expressions enclosed in a syntax path constructor, special characters '/', '(', ')', '{' and '}' must be escaped with a '\'. Since, in Turtle syntax, the '\' character itself must be escaped with an additional '\', special characters are escaped with '\\'.

Example:

Input data	Relational table with one column: <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Name</th> </tr> </thead> <tbody> <tr> <td>{ "FirstName":"John", "LastName":"Smith" }</td> </tr> </tbody> </table>	Name	{ "FirstName":"John", "LastName":"Smith" }
Name			
{ "FirstName":"John", "LastName":"Smith" }			
Logical source definition and Term map	[] xrr:logicalSource [...]; ... rr:objectMap [xrr:reference " Column (Name)/ JSONPath (\$.FirstName)"; rr:language "en";];		
Generated RDF term	"John"@en		

From the example above, it can be noticed that (i) the leftmost syntax path constructor (Column) is consistent with the reference formulation (section 1.4), and (ii) data elements referenced by mixed-syntax path "Column(Name)/JSONPath(\$.FirstName)" are formatted in JSON, corresponding to the rightmost syntax path constructor. More generally:

The leftmost syntax path constructor of a mixed-syntax path must be consistent with the reference formulation (section 1.4).

- Constructors Column(), CSV() and TSV() apply with the column name reference formulation ,
- Constructor XPath() applies with the XPath reference formulation,
- Constructor JSONPath() applies with the JSONPath reference formulation.

The format of data retrieved by a mixed-syntax path is the format of the rightmost syntax path constructor.

3.2.2.3 Production of multiple RDF terms

In a row-based logical source, a column reference returns exactly one scalar value per triples map iteration step: the value of the cell identified by "column name" in the current row. Thus, an R2RML term map generates zero or one RDF term during each iteration step, ultimately a triples map generates zero or one triple during each iteration step.

Due to the tree-like nature of JSON and XML data formats, JSONPath and XPath expressions allow addressing not only literals but also structured values. Thus, using the xrr:reference or rr:template properties with a JSONPath or XPath expression may return multiple values during each triples map iteration step. Hence the introduction of the term map iteration.

A **term map iteration** is a process that occurs in a term map during each triples-map iteration step. Thus a reference-valued or template-valued term map generates zero to any number of RDF terms during each triples-map iteration step.

Examples:

Input data retrieved in one triples-map iteration step	{ "FirstNames": ["John", "Albert"], "LastName": "Smith" }	<person> <FirstNames> <item>John</item> <item>Albert</item> </FirstNames> <LastName>Smith</LastName> </person>
Term map	[] rr:objectMap [xrr:reference "\$.FirstNames.*";];	[] rr:objectMap [xrr:reference "/person/FirstNames/item";];
Generated RDF terms	"John" "Albert"	"John" "Albert"

The term map iteration applies identically in the context of mixed-syntax paths. Example:

Input data	<person> <name>John Smith</name> <items>[1,2,3]</items>
------------	---

	<pre></person></pre> <p>XML element "items" contains a value expressed as a JSON array.</p>
Term map	<pre>[] xrr:logicalSource [...] ... rr:objectMap [xrr:reference "XPath(\\person\\items)/JSONPath(\$.*)"; rr:datatype xsd:integer;]</pre> <p>The last expression of the mixed-syntax path, "JSONPath(\$.*)", indicates that (i) value "[1,2,3]" is formatted in JSON syntax, and (ii) it must be parsed as such using the "\$.*" JSONPath expression.</p>
Generated RDF terms	<pre>1^^xsd:integer 2^^xsd:integer 3^^xsd:integer</pre>

A template-valued term map may reference several data elements from the logical source, captured by curly braces ('{' and '}'). If at least one of the data elements referenced in a template string produces several terms, then the following applies:

A template-valued term map produces RDF terms by performing a Cartesian product between all values produced by all data elements referenced in the template.

Example:

Input data	<pre>{ "FirstNames": ["John", "Albert"], "LastName": "Smith" }</pre>
Term map	<pre>[] xrr:logicalSource [...]; rr:subjectMap [rr:template "http://example.org/{\$.FirstNames.*}/{\$.LastName}";];</pre>
Generated RDF terms	<p>The template performs a Cartesian product between "Smith" and ["John", "Albert"], resulting in two terms:</p> <pre><http://example.org/John/Smith> <http://example.org/Albert/Smith></pre>

Finally, below we define the behavior of a triples map in which one or several term maps generate multiple RDF terms during a single triples map iteration step:

During each iteration of an xR2RML triples map, triples are generated as the Cartesian product between RDF terms produced by the subject map and each predicate-object map. Predicate-object couples result of the Cartesian product between RDF terms produced by each predicate and object map.

Note that a graph map may also produce multiple terms, in which case triples are produced simultaneously in several target graphs.

xR2RML vs. RML: RML makes the restriction that a subject map should return zero or one value during each triples map iteration. In the case of xR2RML, we make no such restriction so that the Cartesian product be possibly applied between multiple subjects, multiple predicate-object couples, and multiple graph IRIs. Besides, RML does not describe how a template with several multi-valued references is processed. xR2RML states that the Cartesian product applies equally in this case, whether the template be used as a subject, predicate, object or graph map.

In the example below, during one triples map iteration step, the subject map produces two RDF terms `<http://example.org/company/Dell>` and `<http://example.org/company/Asus>`, while the object map produces two literals "Laptop" and "Desktop". A Cartesian product between the two subjects and the two objects results in the production of four triples:

Input data: one row retrieved from a relational table, values are formatted in JSON in column "cos", and XML in column "products"	<table border="1"> <thead> <tr> <th data-bbox="411 645 719 712">cos</th> <th data-bbox="719 645 1209 712">products</th> </tr> </thead> <tbody> <tr> <td data-bbox="411 712 719 869">["Dell", "Asus"]</td> <td data-bbox="719 712 1209 869"> <pre><list> <product>Laptop</product> <product>Desktop</product> </list></pre> </td> </tr> </tbody> </table>	cos	products	["Dell", "Asus"]	<pre><list> <product>Laptop</product> <product>Desktop</product> </list></pre>
cos	products				
["Dell", "Asus"]	<pre><list> <product>Laptop</product> <product>Desktop</product> </list></pre>				
Mapping graph	<pre>[] xrr:logicalSource [...]; rr:subjectMap [rr:template "http://example.org/{Column(cos)}/JSONPath(\$.*)";]; rr:predicateObjectMap [rr:predicate ex:produces; rr:objectMap [xrr:reference "Column(products)/XPath(\\list*)";];];</pre>				
Generated triples	<pre><http://example.org/Dell> ex:produces "Laptop". <http://example.org/Dell> ex:produces "Desktop". <http://example.org/Asus> ex:produces "Laptop". <http://example.org/Asus> ex:produces "Desktop".</pre>				

3.2.2.4 Production of RDF collections or containers

A term map with an RDF collection or container term type generates one RDF term during each triples map iteration step. The elements of the collection or container are the RDF terms produced by the term map, whether using property `rr:column`, `xrr:reference` or `rr:template`. In the example below, the triples map generates one triple per iteration step, the object of this triple is an RDF bag (itself consisting of several triples):

Input data: JSON document retrieved in a single iteration step	<pre><company "name"="Dell"> <products> <product>Laptop</product> <product>Desktop</product> </products> </company></pre>
Mapping graph	<pre>[] xrr:logicalSource [...];</pre>

	<pre>rr:subjectMap [rr:template "http://example.org/{/company/@name}";]; rr:predicateObjectMap [rr:predicate ex:builds; rr:objectMap [xrr:reference "//company/products/*"; rr:termType xrr:RdfBag;];];</pre>
Generated triples	<pre><http://example.org/Dell> ex:builds [a rdf:Bag; rdf:_1 "Laptop"; rdf:_2 "Desktop" .]</pre>

Unlike RDF terms of type IRI or blank node, RDF terms of type RDF collection or container cannot be used as subject or predicate of an RDF triple, nor as a graph IRI. Consequently:

A term map with term type `xrr:RdfList`, `xrr:RdfSeq`, `xrr:RdfBag` or `xrr:RdfAlt` is an object map (hence it cannot be a subject map, predicate map nor graph map).

Formally:

?X an `rr:TermMap`.

?X `rr:termType` ?tt.

?tt is one of `xrr:RdfList`, `xrr:RdfSeq`, `xrr:RdfBag` or `xrr:RdfAlt`

⇒ ?X an `rr:ObjectMap`.

A nested term map (property `xrr:nestedTermMap`) can be used to specify a term type, language tag or data type of members of an RDF collection or container. The example below illustrates the usage of a nested term map to generate an RDF collection of IRIs (first example), and an RDF sequence of data-typed literals (second example):

Input data	{ "key1": ["url1", "url2"] }	{ "key1": [10, 20] }
Term map	<pre>[] rr:objectMap [xrr:reference "\$.key1.*"; rr:termType xrr:RdfList; xrr:nestedTermMap [rr:termType rr:IRI;];];</pre>	<pre>[] rr:objectMap [xrr:reference "\$.key1.*"; rr:termType xrr:RdfSeq; xrr:nestedTermMap [rr:termType rr:Literal; rr:datatype xsd:integer;];];</pre>
Generated RDF terms	In Turtle abbreviated notation: (<url1> <url2>)	<pre>[a rdf:Seq; rdf:_1 10^^xsd:integer; rdf:_2 20^^xsd:integer.];</pre>

In a template-valued term map, the `xrr:nestedTermMap` property applies to values resulting from the application of the template string to the input values. In the first example below, term type `rr:IRI` applies to the result of the template string. The same principle applies in the second example with term type `rr:Literal` and datatype `xsd:string`.

Input data	{ "FirstNames": ["John", "Albert"], "LastName": "Smith" }	{ "FirstNames": ["John", "Albert"], "LastName": "Smith" }
Term map	[rr:objectMap [rr:template "http://example.org/ {\$.FirstNames.*}/{\$.LastName}"; rr:termType xrr:RdfList; xrr:nestedTermMap [rr:termType rr:IRI;];];	[rr:objectMap [rr:template "{\$.FirstNames.*} {\$.LastName}"; rr:termType xrr:RdfList; xrr:nestedTermMap [rr:termType rr:Literal; rr:datatype xsd:string;];];
Generated RDF terms	(<http://example.org/John/Smith> <http://example.org/Albert/Smith>)	("John Smith"^^xsd:string "Albert Smith"^^xsd:string)

3.2.3 Parsing nested structured values

The example below illustrates the use of a nested term map to (i) parse nested structured values ("teams" are collections of "team" elements, which are collections of "member" elements) and (ii) translate those nested structured values into RDF terms of class `rdf:List`.

Input data	<teams> <team> <member>John</member> <member>Paul</member> </team> <team> <member>Cathy</member> <member>Ed</member> </team> </teams>
Term map	[rr:objectMap [xrr:reference "/teams/team"; xrr:nestedTermMap [xrr:reference "/member"; rr:termType xrr:RdfList;];]; The first <code>xrr:reference</code> property (" <code>/teams/team</code> ") selects "team" elements from the XML input, each "team" element being the root of an XML tree whose descendants are "member" elements. The second <code>xrr:reference</code> property (" <code>/member</code> "), within the <code>xrr:nestedTermMap</code> property, is evaluated sequentially against the results of the parent reference expression. Thus, the <code>xrr:RdfList</code> term type successively applies to "member" elements of the first team, then to "member" elements of the second team. Finally the term map generates two RDF collections, one per team element.
Generated	("John" "Paul")

RDF terms	<p>("Cathy" "Ed")</p> <p>Note: the object map has no <code>rr:termType</code> property, therefore its term type is that of its nested term type, that is <code>xrr:RdfList</code>.</p>
-----------	--

The subsequent example generates one RDF sequence of nested RDF collections. Elements of the inner RDF collections are typed as `rr:Literal` and assigned a language tag using a second nested `xrr:nestedTermMap` property.

Input data	{ "teams": [["John", "Paul"], ["Cathy", "Ed"]] }
Term map	<pre>[rr:objectMap [xrr:reference "\$.teams.*"; rr:termType xrr:RdfSeq; # represent "teams" as an rdf:Seq # Describe the elements of the RDF sequence xrr:nestedTermMap [rr:template "Player {\$.*.}"; rr:termType xrr:RdfList; # represent each team as an rdf:List # Type members of each team as literals with language "en" # using a simple nested term map xrr:nestedTermMap [rr:termType rr:Literal; rr:language "en";];];];]</pre>
Generated RDF terms	<pre>[a rdf:Seq; rdf:_1 ("Player John"@en "Player Paul"@en); rdf:_2 ("Player Cathy"@en "Player Ed"@en);]</pre>

As already mentioned, in a template-valued term map, property `xrr:nestedTermMap` applies to values resulting from the application of the template string to input values. Thus, defining a nested term map in a template-valued term map suggests that the template produces a valid expression with regards to the current data format, that, in turn, is interpreted against a path expression provided by an `xrr:reference` or `rr:template` property.

For instance, applying the template string:

```
'\{ "first": "{FirstNames}", "last": "{LastName}" \}'
```

would produce a string formatted as a JSON dictionary, like:

```
{ "first": "John", "last": "Smith" }
```

This use case is illustrated in the example below:

Input data	<pre>{ "FirstNames": ["John", "Albert"], "LastName": "Smith" }</pre>
Term map	<pre>[rr:objectMap [rr:template '\{ "first": "{\$.FirstNames.*}", "last": "{\$.LastName}" \}';];]</pre>

	<pre>xrr:nestedTermMap [xrr:reference "\$.*"; rr:termType xrr:RdfList;];]</pre>
Generated RDF terms	<pre>("John" "Smith") ("Albert" "Smith")</pre> <p>Two values are generated by applying the template string, those values are formatted as JSON arrays: <pre>{ "first": "John", "last": "Smith" } { "first": "Albert", "last": "Smith" }</pre></p> <p>The <code>xrr:nestedTermMap</code> property instructs to parse those values using the JSONPath expression <code>"\$.*"</code> (property <code>xrr:reference</code>), and generates an RDF collection (<code>rdf:List</code>) for each of them.</p>

Note: this use case may seem rather awkward and probably of little use, but insofar as it is consistent with the xR2RML language definition, we think it should be considered as valid.

3.2.4 Multiple Mapping Strategies

The flexibility offered by nested term maps allows the same mapping to be written using various strategies: path expressions of properties `xrr:reference` and `rr:template` can be split in several levels of term map and nested term map.

For instance, both term maps below produce equivalent results. In the first case (left), the JSONPath expression (`$.teams.*.*`) retrieves all team members at once. In the second case (right), teams are retrieved first (`$.teams.*`), then the `xrr:nestedTermMap` property runs a second JSONPath evaluation to retrieve and datatype team members.

Input data	<pre>{ "teams": [["John", "Paul"] , ["Cathy", "Ed"]] }</pre>	
Term maps	<pre>[] xrr:logicalSource [...]; rr:objectMap [xrr:reference "\$.teams.*.*"; rr:datatype xsd:string;];</pre>	<pre>[] xrr:logicalSource [...]; rr:objectMap [xrr:reference "\$.teams.*"; xrr:nestedTermMap [xrr:reference "\$.*"; rr:datatype xsd:string;];];</pre>
Generated RDF terms	<pre>"John"^^xsd:string "Paul"^^xsd:string "Cathy"^^xsd:string "Ed"^^xsd:string</pre>	

It is likely that the first case will be more efficient as only one JSONPath evaluation is performed, whereas in the second case two JSONPath evaluations are performed in sequence.

Similarly, the example below shows how a mixed-syntax path can be split into a term map and a nested term map:

<pre>[] xrr:logicalSource [...]; rr:objectMap [xrr:reference "Column(col)/XPath(\\Vperson\\Vname)"; rr:datatype xsd:string;];</pre>	<pre>[] xrr:logicalSource [...]; rr:objectMap [rr:column "col"; xrr:nestedTermMap [xrr:reference "XPath(\\Vperson\\Vname)"; rr:datatype xsd:string;];];</pre>
--	--

Both mappings are likely to be equally efficient, as both evaluations (column selection and XPath expression evaluation) need to be done anyway.

3.2.5 Default Term Types

This section is an adaptation of section 7.4 (<http://www.w3.org/TR/r2rml/#termtype>) of the R2RML specification. xR2RML additions to R2RML are highlighted.

If the term map has an optional `rr:termType` property then its term type is the value of that property. The value MUST be one of the following options:

- If the term map is a subject map: `rr:IRI` or `rr:BlankNode`
- If the term map is a predicate map: `rr:IRI`
- If the term map is an object map: `rr:IRI`, `rr:BlankNode`, `rr:Literal`, `rdf:List`, `rdf:Seq`, `rdf:Bag`, `rdf:Alt`.
- If the term map is a graph map: `rr:IRI`.

If the term map does not have an `rr:termType` property, then its term type is:

- `rr:Literal`, if it is an object map and at least one of the following conditions is true:
 - It is a column-based term map.
 - It has an `rr:language` property (and thus a specified language tag).
 - It has an `rr:datatype` property (and thus a specified datatype).
 - It does not have an `rr:language` property and it has a nested term map that has an `rr:language` property.
 - It does not have an `rr:datatype` property and it has a nested term map that has an `rr:datatype` property.
- the term type of the value of its nested term map.
- `rr:IRI`, otherwise.

A corollary of this definition is that the `xrr:nestedTermMap` property may be used in a subject map, predicate map or graph map only if it produces IRIs. Consequently:

A term map with an `xrr:nestedTermMap` property may be a subject map or graph map only if (i) it does not have an `rr:termType` property and (ii) its nested term map has an `rr:termType` property with object `rr:IRI` or `rr:BlankNode`.
 A term map with an `xrr:nestedTermMap` property may be a predicate map only if (i) it does not have an `rr:termType` property and (ii) its nested term map property has an `rr:termType` property with object `rr:IRI`.

3.3 Reference relationships between logical sources

The following definitions are an adaptation of R2RML specification section 8 (<http://www.w3.org/TR/r2rml/#foreign-key>). xR2RML additions to R2RML are highlighted.

A referencing object map allows using the subjects of another triples map as the objects

generated by a predicate-object map. Since both triples maps may be based on different logical sources, this may require a join between the logical sources.

A referencing object map resource has exactly one `rr:parentTriplesMap` property (its value is a triples map), and optional `rr:joinCondition` properties. A join condition has exactly one `rr:child` property and one `rr:parent` property. The `rr:child` property references the join condition's child data element, the `rr:parent` property references the join condition's parent data element. Data element references are valid path expressions with regards to the reference formulation (section 1.4), possibly using mixed-syntax paths.

A referencing object map may have an `rr:termType` property with an RDF collection or container term type (see further details in §3.3.2).

The child query of a referencing object map is the query or source name of the logical source of the triples map containing the referencing object map.

The parent query of a referencing object map is the query or source name of the logical source of the referencing object map's parent triples map.

Properties `rr:child` and `rr:parent` use valid path expressions to reference data elements. As described in §3.2.2.3, such path expressions may produce multiple terms. Consequently, the equivalent joint query of a referencing object map must take into account the fact that child and parent references be multi-valued. More precisely, a join between two multi-valued references should be satisfied if at least one data element of the first reference matches one data element of the second reference.

The joint query of a referencing object map is defined below using SQL syntax (`SELECT... FROM... AS... WHERE`) and first order logic for the description of `WHERE` conditions:

If a referencing object map has no join condition, its joint query is:

```
SELECT * FROM ({child-query}) AS tmp
```

If a referencing object map has at least one join condition, its joint query is:

```
SELECT * FROM ({child-query}) AS child, ({parent-query}) AS parent
WHERE  $\exists c1 \in \text{eval}(\text{child}, \{\text{child-ref1}\}), \exists p1 \in \text{eval}(\text{parent}, \{\text{parent-ref1}\}), c1 = p1$ 
AND  $\exists c2 \in \text{eval}(\text{child}, \{\text{child-ref2}\}), \exists p2 \in \text{eval}(\text{parent}, \{\text{parent-ref2}\}), c2 = p2$ 
AND ...
```

where `{child-ref1}` and `{parent-ref1}` are the child reference and parent reference of the first join condition and so on, and `eval(source, {ref})` is the result of the evaluation of expression "{ref}" on data "source".

Note: when applied to a relational database, in which child and parent references are single-valued, this definition can be simplified into the R2RML joint query definition:

```
SELECT * FROM ({child-query}) AS child, ({parent-query}) AS parent
WHERE child.{child-ref1} = parent.{parent-ref1}
AND child.{child-ref2} = parent.{parent-ref2}
AND ...
```

3.3.1 Reference relationship with structured values

The relational database example below models the relation between medical doctors and the studies for which they are investigators. Column "Doctor.studies" contains JSON arrays of which elements are references (similar to foreign keys) to column "Study.study_id".

Input data	<p>Table Study</p> <table border="1" data-bbox="424 448 812 607"> <thead> <tr> <th>study_id</th> <th>study_name</th> </tr> </thead> <tbody> <tr> <td>1</td> <td>study1</td> </tr> <tr> <td>2</td> <td>study2</td> </tr> <tr> <td>3</td> <td>study3</td> </tr> </tbody> </table> <p>Table Doctor</p> <table border="1" data-bbox="424 674 979 792"> <thead> <tr> <th>doc_id</th> <th>doc_name</th> <th>studies</th> </tr> </thead> <tbody> <tr> <td>1</td> <td>D1</td> <td>[1,2]</td> </tr> <tr> <td>2</td> <td>D2</td> <td>[3]</td> </tr> </tbody> </table>	study_id	study_name	1	study1	2	study2	3	study3	doc_id	doc_name	studies	1	D1	[1,2]	2	D2	[3]
study_id	study_name																	
1	study1																	
2	study2																	
3	study3																	
doc_id	doc_name	studies																
1	D1	[1,2]																
2	D2	[3]																
Mapping graph	<pre data-bbox="424 857 1257 1534"> <#Study> rr:logicalTable [rr:tableName "Study"]; rr:subjectMap [rr:template "http://example.org/study/{study_name}";]. <#Doctor> rr:logicalTable [rr:tableName "Doctor"]; rr:subjectMap [rr:template "http://example.org/doc/{doc_name}";]; rr:predicateObjectMap [rr:predicate ex:investigatorOf; rr:objectMap [rr:parentTriplesMap <#Study>; rr:joinCondition [rr:parent "study_id"; rr:child "Column(studies)/JSONPath(\$.*)";];];]. </pre> <p data-bbox="424 1568 1455 1657">The rr:child property uses a mixed-syntax path specifying that the data retrieved is formatted in JSON, and that each element of this structured value is considered in the join operation.</p>																	
Generated triples	<pre data-bbox="424 1697 1257 1881"> <http://example.org/doc/D1> ex:investigatorOf <http://example.org/study/study1> . <http://example.org/doc/D1> ex:investigatorOf <http://example.org/study/study2> . <http://example.org/doc/D2> ex:investigatorOf <http://example.org/study/study3> . </pre> <p data-bbox="424 1915 1455 1982">According to the equivalent joint query definition, the joint query is as follows ("child" and "parent" notations have been removed for readability):</p> <pre data-bbox="518 1993 933 2027"> SELECT * FROM Doctor, Study </pre>																	

	<p>WHERE $\exists c \in \text{eval}(\text{Doctor}, \text{Column}(\text{studies})/\text{JSONPath}(\\$.*))$, $\exists p \in \text{Study.study_id}$, $c = p$</p> <p>where $\text{eval}(\text{Doctor}, \text{Column}(\text{studies})/\text{JSONPath}(\\$.*))$ represents the evaluation of mixed-syntax path "Column(studies)/JSONPath(\$.*)" on table Doctor.</p> <p>Since Study.study_id is single-valued, we can rewrite the query as:</p> <p>SELECT * FROM Doctor, Study WHERE $\exists c \in \text{Doctor.Column}(\text{studies})/\text{JSONPath}(\\$.*)$, $c = \text{Studies.study_id}$</p> <p>The join query results in this table:</p> <table border="1"> <thead> <tr> <th>doc_id</th> <th>doc_name</th> <th>studies</th> <th>study_id</th> <th>study_name</th> </tr> </thead> <tbody> <tr> <td>1</td> <td>D1</td> <td>[1,2]</td> <td>1</td> <td>study1</td> </tr> <tr> <td>1</td> <td>D1</td> <td>[1,2]</td> <td>2</td> <td>study2</td> </tr> <tr> <td>2</td> <td>D2</td> <td>[3]</td> <td>3</td> <td>study3</td> </tr> </tbody> </table>	doc_id	doc_name	studies	study_id	study_name	1	D1	[1,2]	1	study1	1	D1	[1,2]	2	study2	2	D2	[3]	3	study3
doc_id	doc_name	studies	study_id	study_name																	
1	D1	[1,2]	1	study1																	
1	D1	[1,2]	2	study2																	
2	D2	[3]	3	study3																	

3.3.2 Generating RDF collection/container with a referencing object map

In R2RML, referencing object maps do not have an `rr:termType` property as they should only produce RDF terms of type `rr:IRI`. In xR2RML however, the result of a joint query may be translated into an RDF collection or container using property `rr:termType`. The `rr:termType` has a specific semantics here: it groups joint query results by subjects of the generated triples, i.e. by child reference, and renders all objects in the same grouping as an RDF collection or container.

If a referencing object map has no `rr:termType` property, then its term type is `rr:IRI` (compliant with the R2RML definition).

A referencing object map may have an `rr:termType` property with an RDF collection or container term type (`xrr:RdfList`, `xrr:RdfSeq`, `xrr:RdfBag` or `xrr:RdfAlt`). In that case, members of the collection or container are necessarily of type `rr:IRI`.

In a referencing object map with an RDF collection or container term type, results of the joint query are grouped by value of the child reference. The objects of the triples map are grouped in a single object of type RDF collection or container, as instructed by the `rr:termType` property.

In the example below the referencing object map has an `rr:termType` property with value `xrr:RdfList`:

Input data	<p>JSON documents retrieved by the query in the <#Study> triples map:</p> <pre>{ "study_id":1, "study_name":"study1" } { "study_id":2, "study_name":"study2" } { "study_id":3, "study_name":"study3" }</pre> <p>JSON documents retrieved by the query in the <#Doctor> triples map:</p>
------------	---

	<pre>{ "doc_name":"D1", "studies": [1,2] } { "doc_name":"D2", "studies": [2,3] }</pre>
Mapping graph	<p>Below, queries to retrieve Studies and Doctors are referred to as <i><Study query></i> and <i><Doctor query></i>.</p> <pre><#Doctor> xrr:logicalSource [xrr:query "<Doctor query>";]; rr:subjectMap [rr:template "http://example.org/doc/{\$.doc_name}";]. <#Study> xrr:logicalSource [xrr:query "<Study query>";]; rr:subjectMap [rr:template "http://example.org/study/{\$.study_name}";]; rr:predicateObjectMap [rr:predicate ex:hasInvestigators; rr:objectMap [rr:parentTriplesMap <#Doctor>; rr:joinCondition [rr:child "\$.study_id"; rr:parent "\$.studies.*";]; rr:termType xrr:RdfList;];].</pre>
Generated RDF triples	<pre><http://example.org/study/study1> ex:hasInvestigators (<http://example.org/doc/D1>). <http://example.org/study/study2> ex:hasInvestigators (<http://example.org/doc/D1> <http://example.org/doc/D2>). <http://example.org/study/study3> ex:hasInvestigators (<http://example.org/doc/D2>).</pre> <p>Explanation: according to the equivalent joint query definition, the joint query is as follows:</p> <pre>SELECT * FROM (<Study query>) as child, (<Doctor query>) as parent WHERE ∃ p ∈ eval(parent, \$.studies.*), p = eval(child, \$.study_id)</pre> <p>where <i>eval(parent, \$.studies.*)</i> represents the evaluation of path "\$.studies.*" on the result of the parent query, and <i>eval(child, \$.study_id)</i> represents the evaluation of path "\$.study_id" on the result of the child query.</p> <p>The equivalent joint query results in the following documents:</p> <pre>{ "study_id":1, "study_name":"study1", "doc_name":"D1", "studies": [1,2] } { "study_id":2, "study_name":"study2", "doc_name":"D1", "studies": [1,2] } { "study_id":2, "study_name":"study2", "doc_name":"D2", "studies": [2,3] } { "study_id":3, "study_name":"study3", "doc_name":"D2", "studies": [2,3] }</pre> <p>Then, term type <code>xrr:RdfList</code> instructs to group results by child reference, i.e. by</p>

"study_id".

3.3.3 Generating RDF collection/container with a referencing object map in the relational case

An interesting consequence of using the `rr:termType` in a referencing object map is the ability, in the case of a relational database with standard SQL values, to build an RDF collection or container reflecting a one-to-many relation. In the example below, foreign key `Study.doctor` relates each study to its investigator in a many-to-one relation (several studies may have the same investigator). Considered the other way round, it can be seen as a one-to-many relation (one doctor investigates several studies). The mapping graph describes the generation of each doctor along with the list of studies he/she investigates.

Input data	<p>Table Study</p> <table border="1"> <thead> <tr> <th>study_id</th> <th>study_name</th> <th>doctor</th> </tr> </thead> <tbody> <tr> <td>1</td> <td>study1</td> <td>1</td> </tr> <tr> <td>2</td> <td>study2</td> <td>1</td> </tr> <tr> <td>3</td> <td>study3</td> <td>2</td> </tr> </tbody> </table> <p>Table Doctor</p> <table border="1"> <thead> <tr> <th>doc_id</th> <th>doc_name</th> </tr> </thead> <tbody> <tr> <td>1</td> <td>D1</td> </tr> <tr> <td>2</td> <td>D2</td> </tr> </tbody> </table>	study_id	study_name	doctor	1	study1	1	2	study2	1	3	study3	2	doc_id	doc_name	1	D1	2	D2
study_id	study_name	doctor																	
1	study1	1																	
2	study2	1																	
3	study3	2																	
doc_id	doc_name																		
1	D1																		
2	D2																		
Mapping graph	<pre> <#Study> rr:logicalTable [rr:tableName "Study"]; rr:subjectMap [rr:template "http://example.org/study/{study_name}";]. <#Doctor> rr:logicalTable [rr:tableName "Doctor"]; rr:subjectMap [rr:template "http://example.org/doc/{doc_name}";]. rr:predicateObjectMap [rr:predicate ex:investigatesStudies; rr:objectMap [rr:parentTriplesMap <#Study>; rr:joinCondition [rr:child "doc_id"; rr:parent "doctor";]; rr:termType xrr:RdfList;];]. </pre>																		
Generated RDF triples	<pre> <http://example.org/doc/D1> ex: investigatesStudies (<http://example.org/study/study1> <http://example.org/study/study2>) . <http://example.org/doc/D2> ex: investigatesStudies </pre>																		

(<<http://example.org/study/study3>>) .

The equivalent joint query results in this table:

doc_id	doc_name	study_id	study_name	doctor
1	D1	1	study1	1
1	D1	2	study2	1
2	D2	3	study3	2

Results are grouped by child reference, i.e. by column "doc_id" to generate RDF lists.

4 References

- [1] S. Das, S. Sundara, R. Cyganiak, R2RML: RDB to RDF Mapping Language, (2012).
- [2] A. Dimou, M.V. Sande, RDF Mapping Language (RML) Unofficial Draft 17 September 2014, (2014).
- [3] A. Dimou, M. Vander Sande, P. Colpaert, E. Mannens, R. Van de Walle, Extending R2RML to a source-independent mapping language for RDF, in: Workshop Proceedings, 12th International Semantic Web Conference Posters & Demos, Sydney, Australia, 2013: pp. 237–240.
- [4] A. Dimou, M. Vander Sande, P. Colpaert, R. Verborgh, E. Mannens, R. Van de Walle, RML: A Generic Language for Integrated RDF Mappings of Heterogeneous Data, in: Proceedings of the 7th Workshop on Linked Data on the Web (LDOW2014), Seoul, Korea, 2014.
- [5] S.K. Gajendran, A Survey on NoSQL Databases (technical report), 2013.
- [6] R. Hecht, S. Jablonski, NoSQL Evaluation: A Use Case Oriented Survey, in: Proceedings of the 2011 International Conference on Cloud and Service Computing, IEEE Computer Society, Washington, DC, USA, 2011: pp. 336–341.
- [7] B. Kolev, P. Valduriez, R. Jimenez-Peris, N. Martínez-Bazan, J. Pereira, CloudMdsQL: Querying Heterogeneous Cloud Data Stores with a Common Language.pdf, in: Proceeding of the BDA 2014 Conference, Autrans, France, 2014.
- [8] K.W. Ong, Y. Papakonstantinou, R. Vernoux, The SQL++ Unifying Semi-structured Query Language, and an Expressiveness Benchmark of SQL-on-Hadoop, NoSQL and NewSQL Databases (submitted), CoRR. abs/1405.3631 (2014).