Virtual Knowledge Graphs — Applications and Tools

Guohui Xiao

University of Bergen & Ontopic

Huawei Ireland Research Centre

7 July 2022

Outline

1. A Brief History of Semantic Web & Knowledge Graphs
2. The Virtual Knowledge Graph (VKG) Framework
3. VKG Systems and Applications
4. Key Technologies in VKGs
5. A Demo for Data Integration through VKGs

1. A Brief History of Semantic Web & Knowledge Graphs

ONTOP IC

World Wide Web (1989)

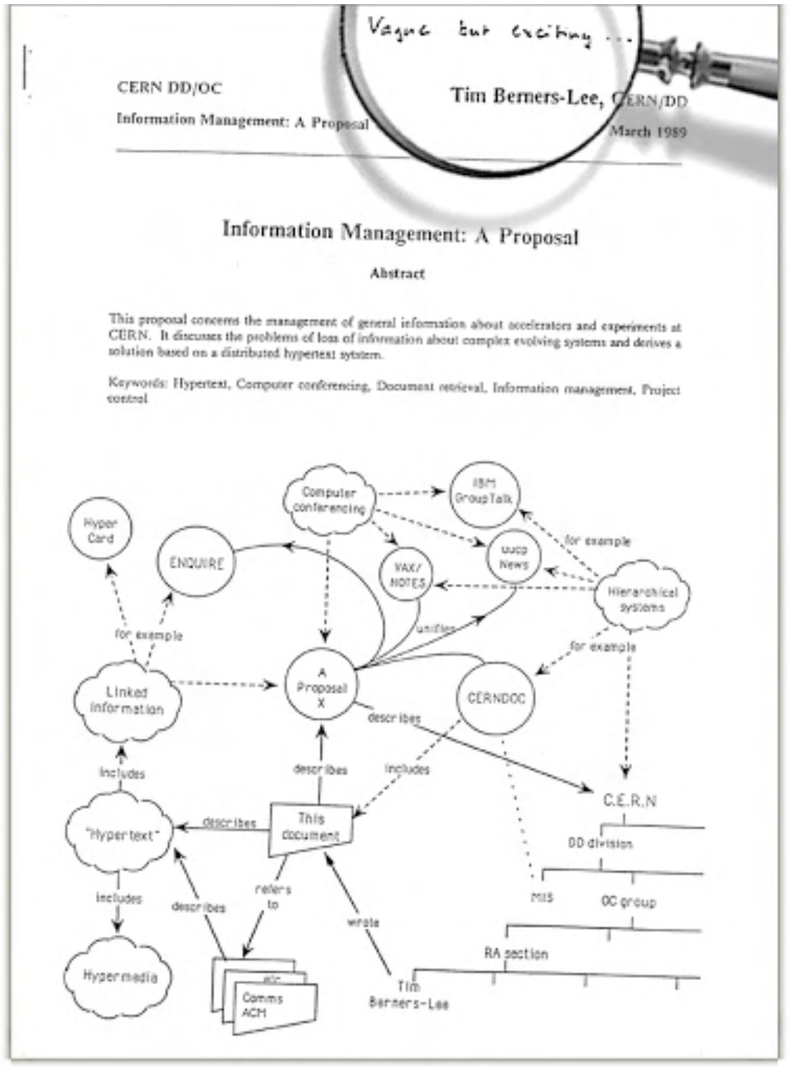

Vague but exciting ...

CERN DD/OC

Tim Berners-Lee, CERN/DD

Information Management: A Proposal

March 1989

Information Management: A Proposal

Abstract

This proposal concerns the management of general information about accelerators and experiments at CERN. It discusses the problems of loss of information about complex evolving systems and derives a solution based on a distributed hypertext system.

Keywords: Hypertext, Computer conferencing, Document retrieval, Information management, Project control

*… This is why a "web" of notes with links between them is far more **useful than a fixed hierarchical system**. … The system we need is like **a diagram of circles and arrows**, where circles and arrows can stand for anything.*

— Information management: A proposal. CERN. 1989

Links are the core of the Web.

WWW Inventor

Sir Tim Berners-Lee

Web of Documents

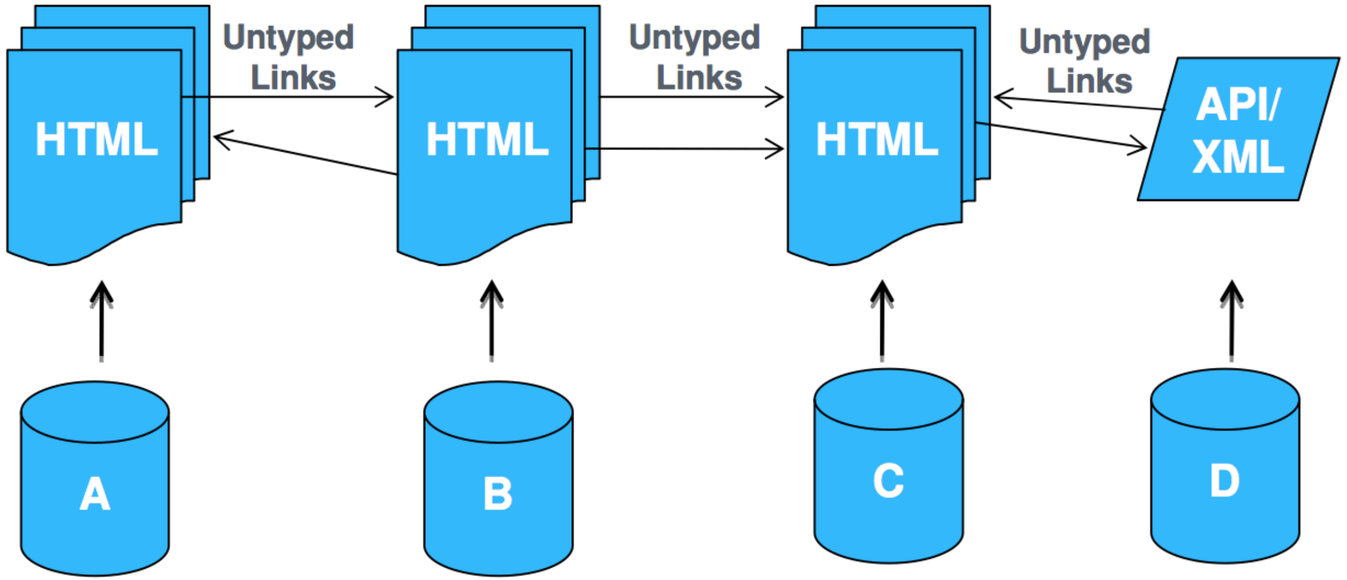

- Primary objects: **documents**
- Degree of structure in data: **low**
- Semantics of content: **implicit**
- Designed for: **human consumption**

Links between **documents**

Problem of the Web of Documents

- wealth of information
- … targeted at human users

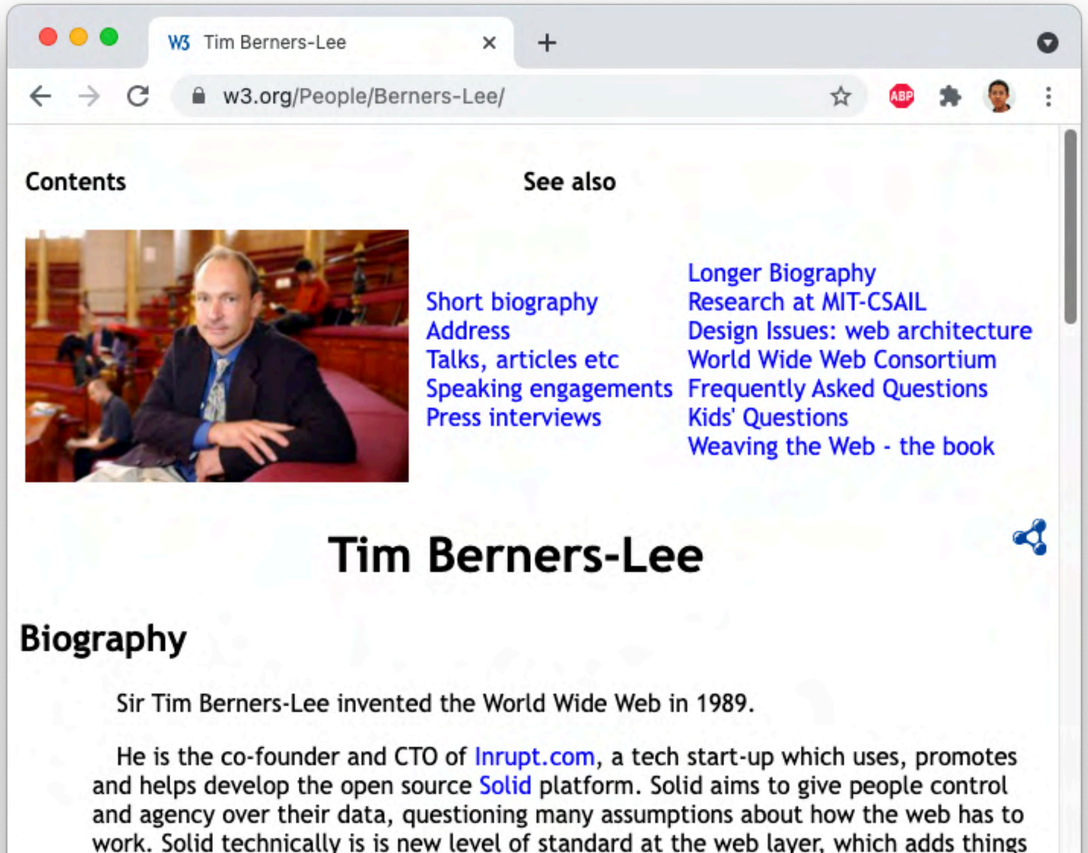


```
<?xml version="1.0" encoding="utf-8"?>
<!DOCTYPE html><html xmlns="http://www.w3.org/1999/xhtml">
  <head>
    <meta charset="utf-8" />
    <meta content="HTML Tidy for Mac OS X (vers 31 October 2006 - Apple Inc. build 1
      name="generator" />
    <title>Tim Berners-Lee</title>
    <link type="text/css" rel="stylesheet" href="general.css" />
    <link href="card.rdf" title="Contact" type="application/rdf+xml" rel="meta" />
    <!-- <link rel="meta" type="application/rdf+xml" title="FOAF" href="card.rdf"/>
    <link href="http://pip.verisignlabs.com/server" rel="openid.server" />
    <link href="http://timbl.pip.verisignlabs.com/" rel="openid.delegate" />
  </head>
  <body>
    <table summary="Nav" width="100%" cellpadding="2">
      <tbody>
        <tr>
          <td><b>Contents</b> </td>
          <th>
            <p align="left">See also</p>
          </th>
          <th>
            <p align="left"><br />
            </p>
          </th>
        </tr>
        <tr>
          <td> <img src="../../Press/Stock/Berners-Lee/2001-europaeum-eighth.jpg"
              alt="Tim BL"
              width="252"
              height="166" /><br />
            <!--
            <a href="https://www.battleforthenet.com/"> <img src="images/protest-loa
          </td>
          <td><a href="#Bio">Short biography</a><br />
            <a href="#Address">Address</a><br />
            <a href="#Talks">Talks, articles etc</a><br />
            <a href="#Speaking">Speaking engagements</a><br />
            <a href="#Press">Press interviews</a></td>
          <td><a href="Longer.html">Longer Biography</a><br />
```

Web of Documents

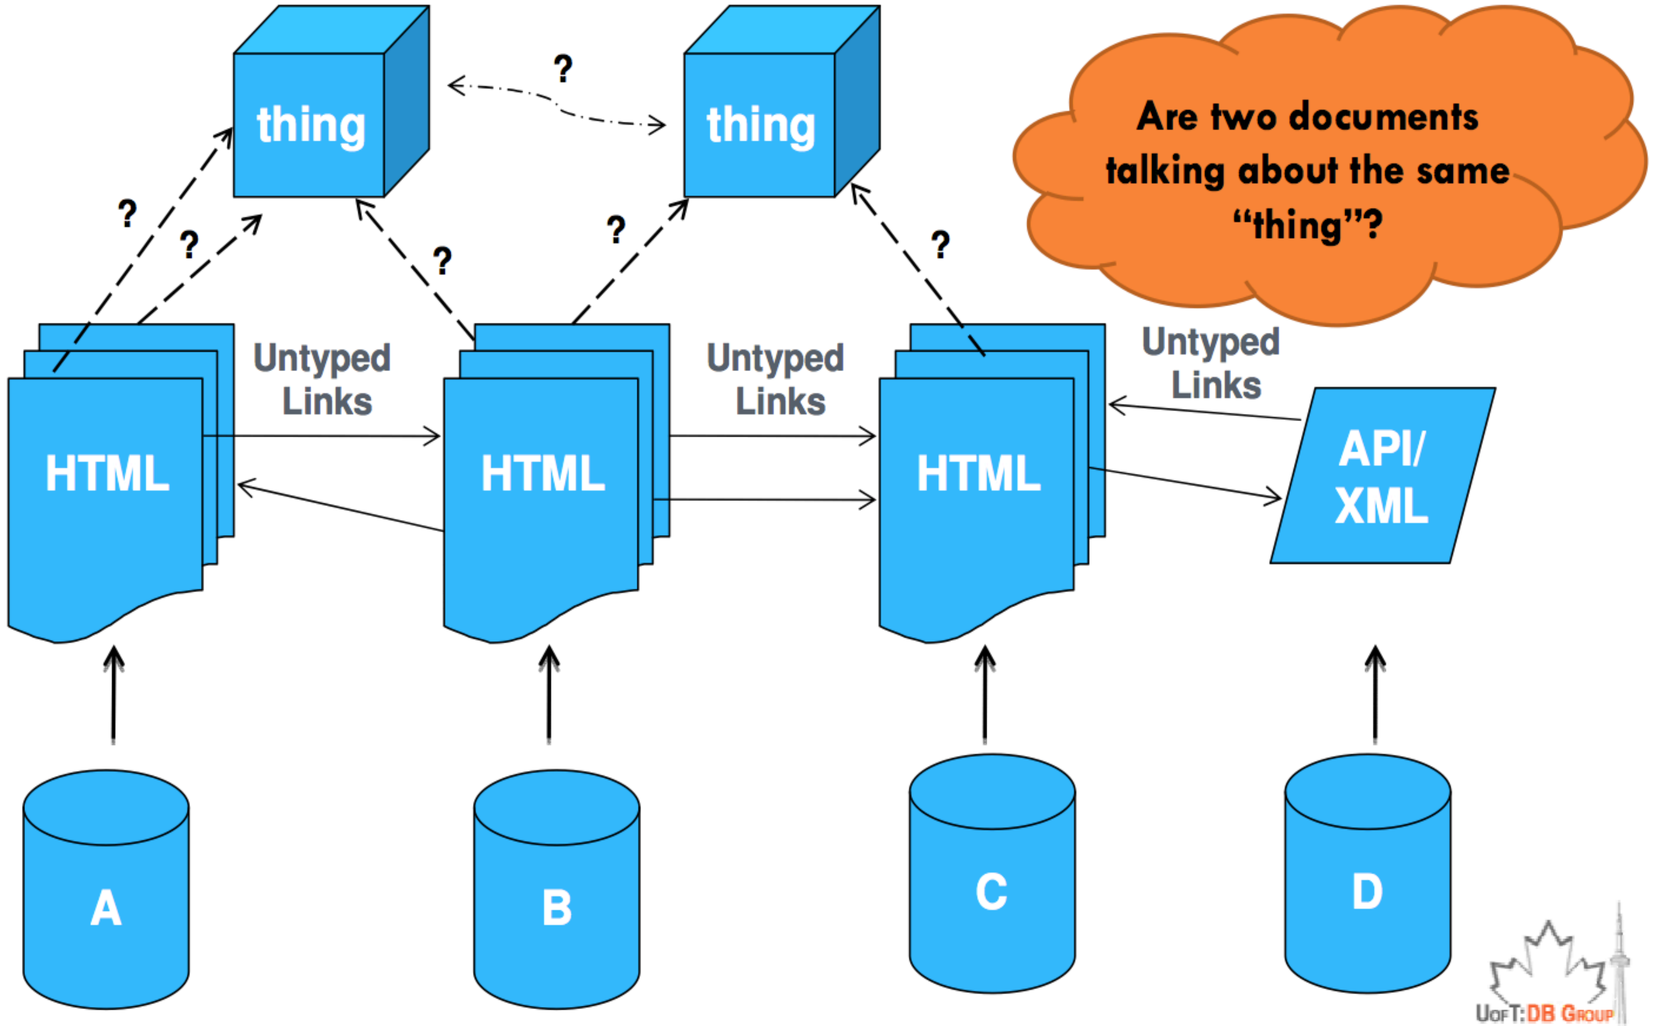

Problems of the Web of Documents

Approaches toward a solution:

① Ad hoc: Deployment of AI methods
(most notably NLP techniques) to evaluate
existing unstructured information on the Web

② A priori: Structure information on the Web
at authoring time
in a way facilitating later automated deployment

➔ Semantic Web

The Semantic Web (2001)

The Semantic Web — A new form of Web content that is meaningful to computers will unleash a revolution of new possibilities

By Tim Berners-Lee, James Hendler and Ora Lassila

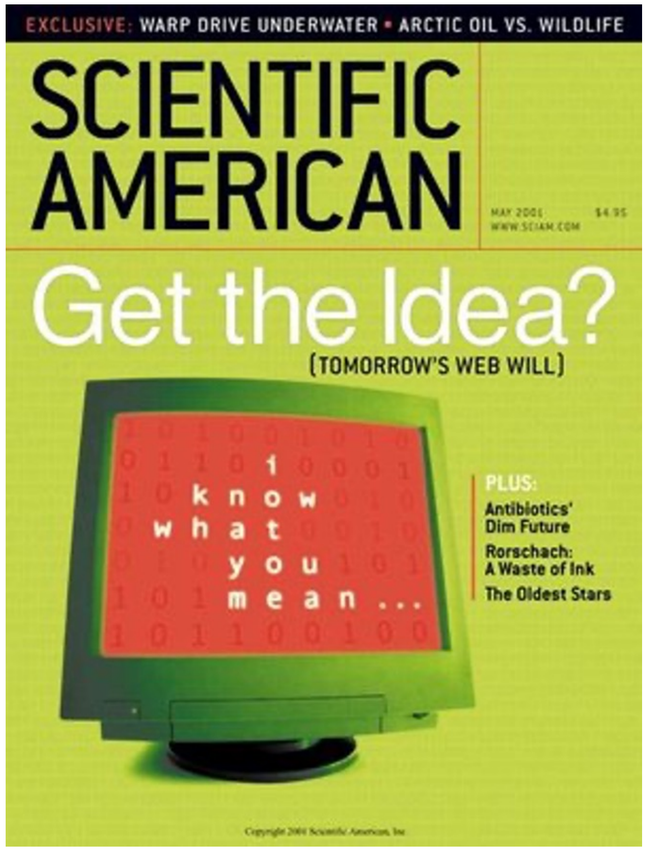

Key ideas:

- Expressing Meaning: things, but not strings
- Knowledge Representation: logic based, triple format
- Ontologies: expressive knowledge
- Agents: Machines should understand each other

Semantic Technologies

- **Not Only About the Web**
- The Semantic Web vision has generated technologies that are applied outside the web context including:
  - Retailing, supermarkets (see web pages of Best Buy)
  - Health care, medicine (see SNOMED ontology)
  - Enterprise intelligence (see solutions by Ontotext)
  - Research: Bio, Geo, Cultural heritage, etc.
  - Government (“smart cities”)
  - Software development (“ontology-driven software develoment”)

Semantic Web: W3C Standards

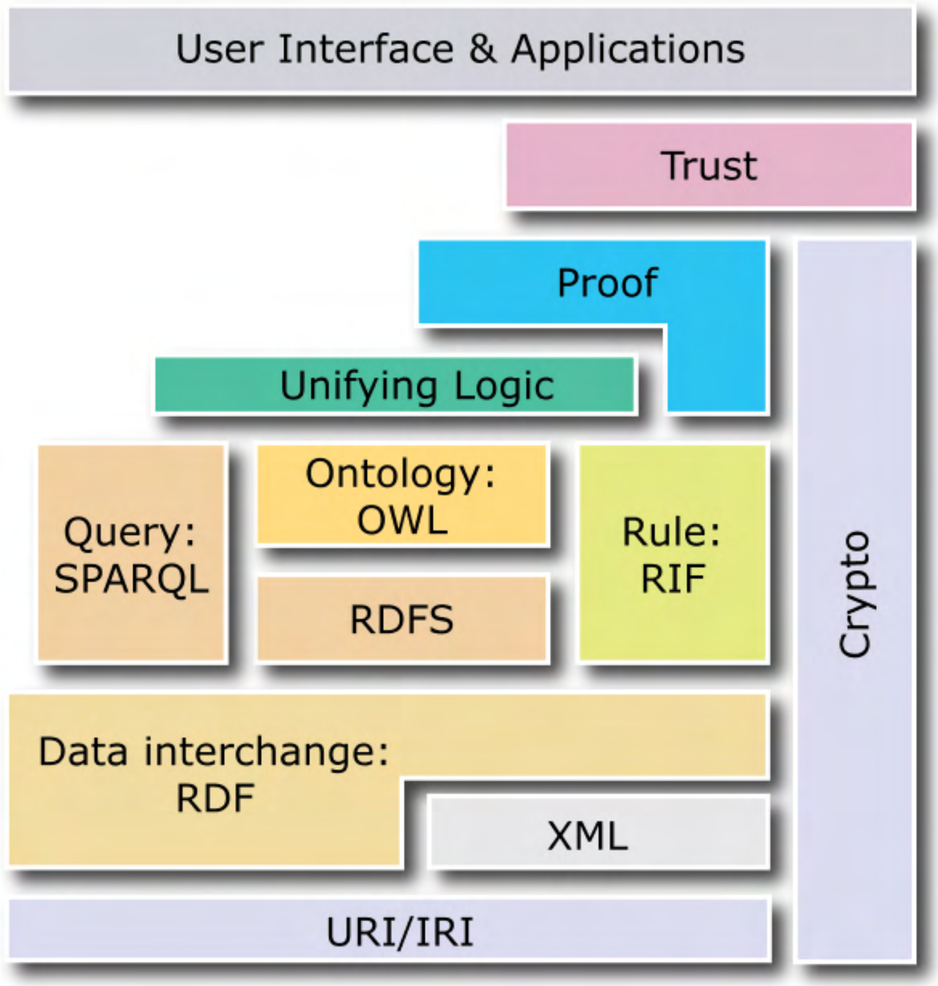

- 1994 First public presentation of the Semantic Web idea
- 1998 Start of standardization of data model (RDF) and a first ontology languages (RDFS) at W3C
- 2000 Start of large research projects about ontologies in the US and Europe
- 2002 Start of standardization of a new ontology language (OWL) based on research results
- 2004 Finalization of the standard for data (RDF) and ontology (OWL)
- 2008 Standardization of a query language (SPARQL)
- 2009 Extension of OWL to OWL 2.0
- 2010 Standard Rule Interchange Format (RIF)
- 2012 RDB2RDF Mapping Language (R2RML and Direct Mapping)
- 2013 SPARQL 1.1
- 2014 RDF 1.1

Google Knowledge Graph

“Introducing the Knowledge Graph: things, not strings” (2012)

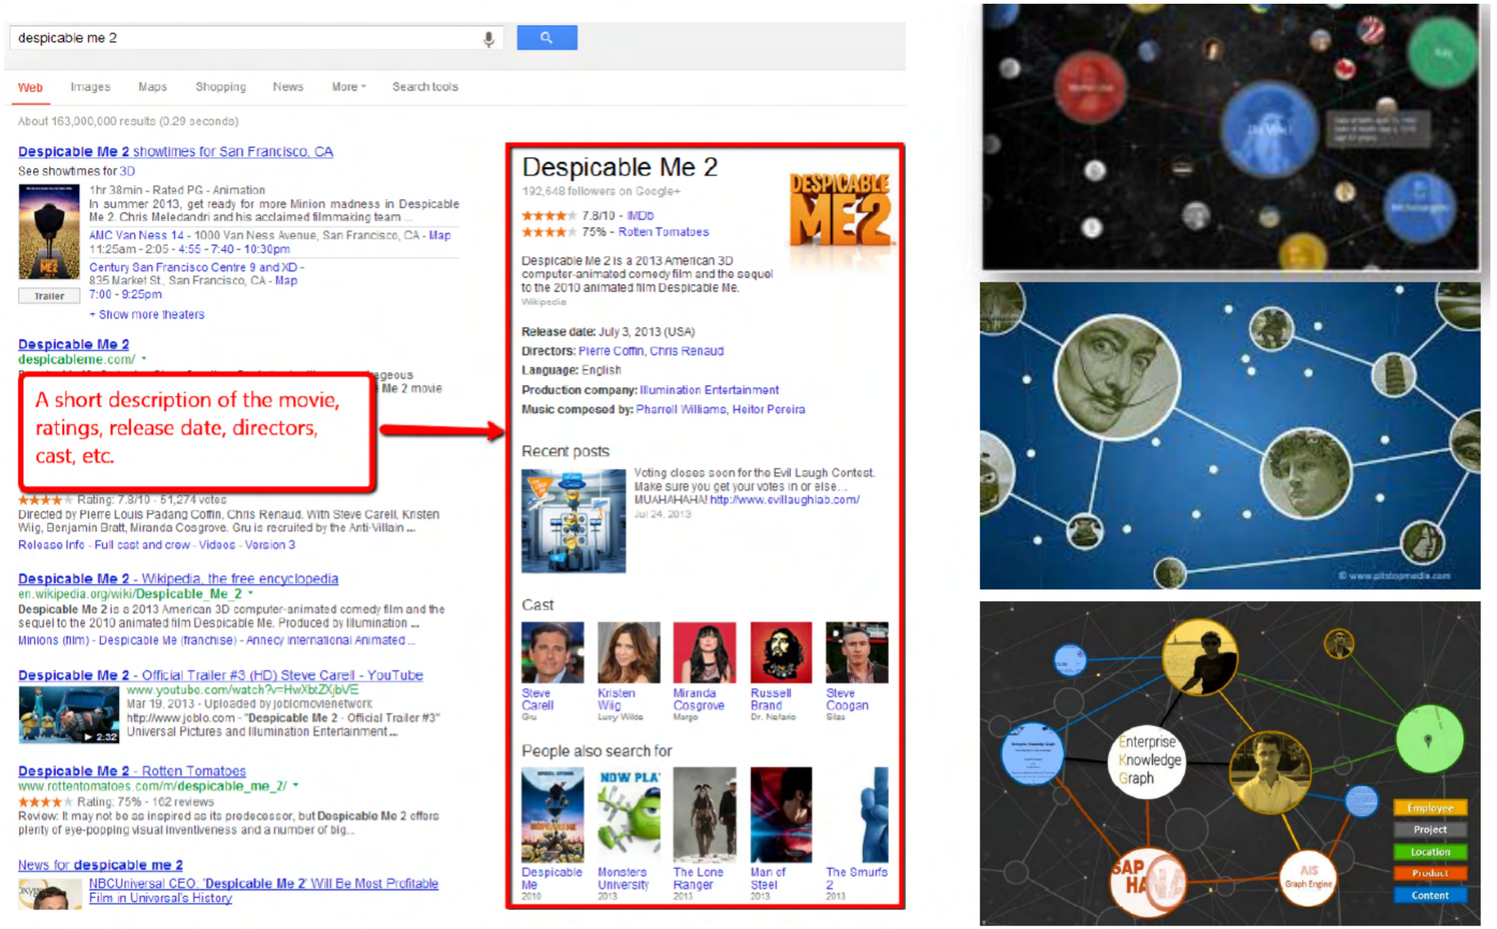

“Knowledge Graph” became a popular term for Semantic Web data.

Linked Data (Open Linked KGs) in 2017

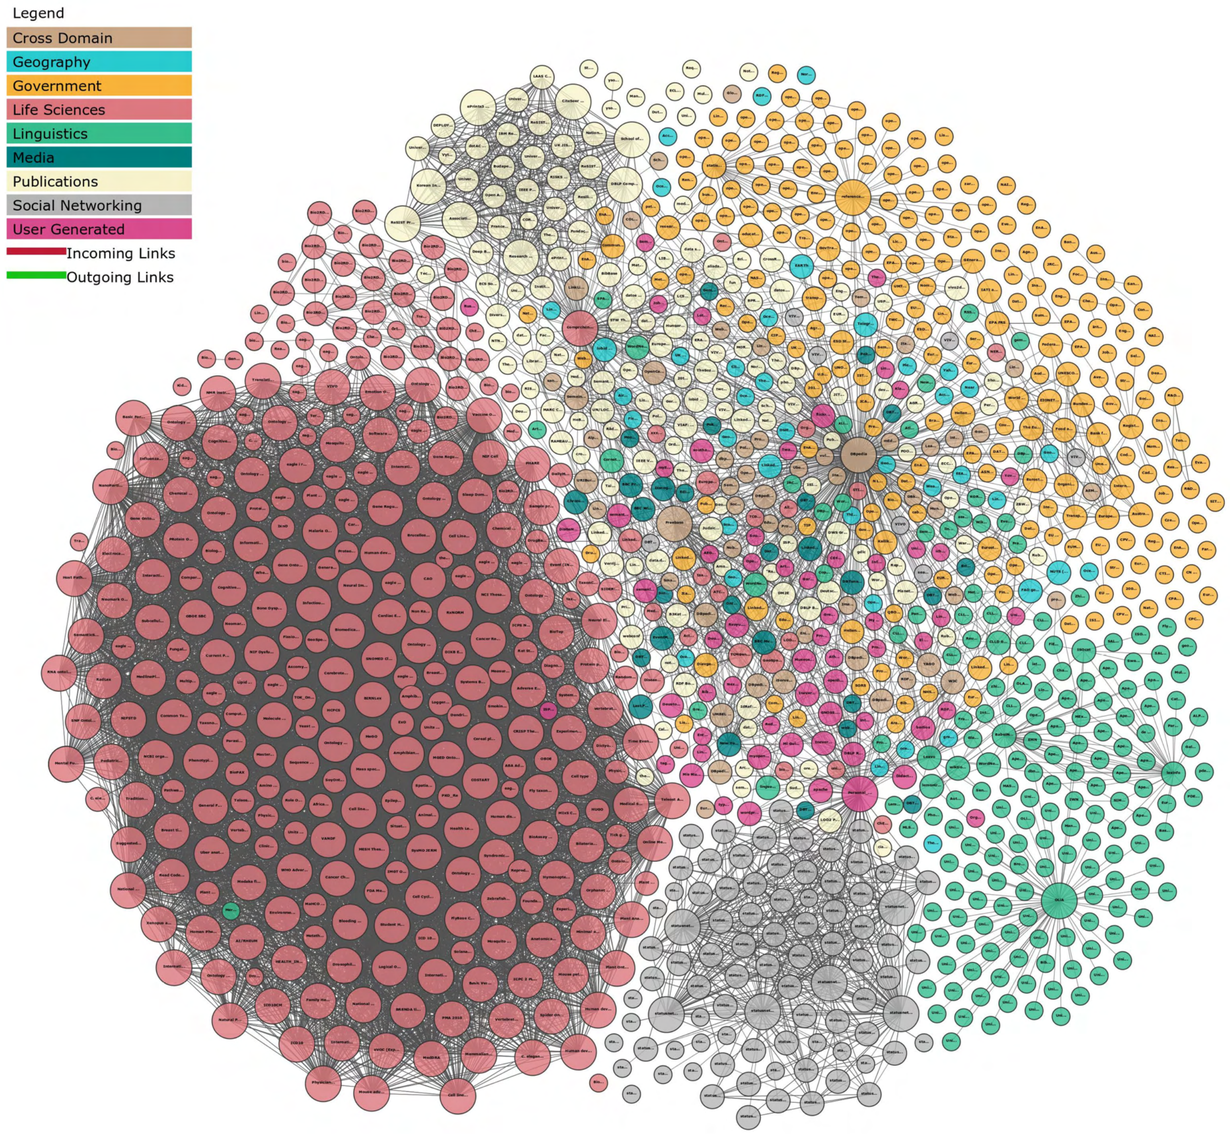

1163 datasets,
40 billion RDF triples

"Linking Open Data cloud diagram 2017, by Andrejs Abele, John P. McCrae, Paul Buitelaar, Anja Jentzsch and Richard Cyganiak. http://lod-cloud.net/"

Typical Knowledge Graph Projects

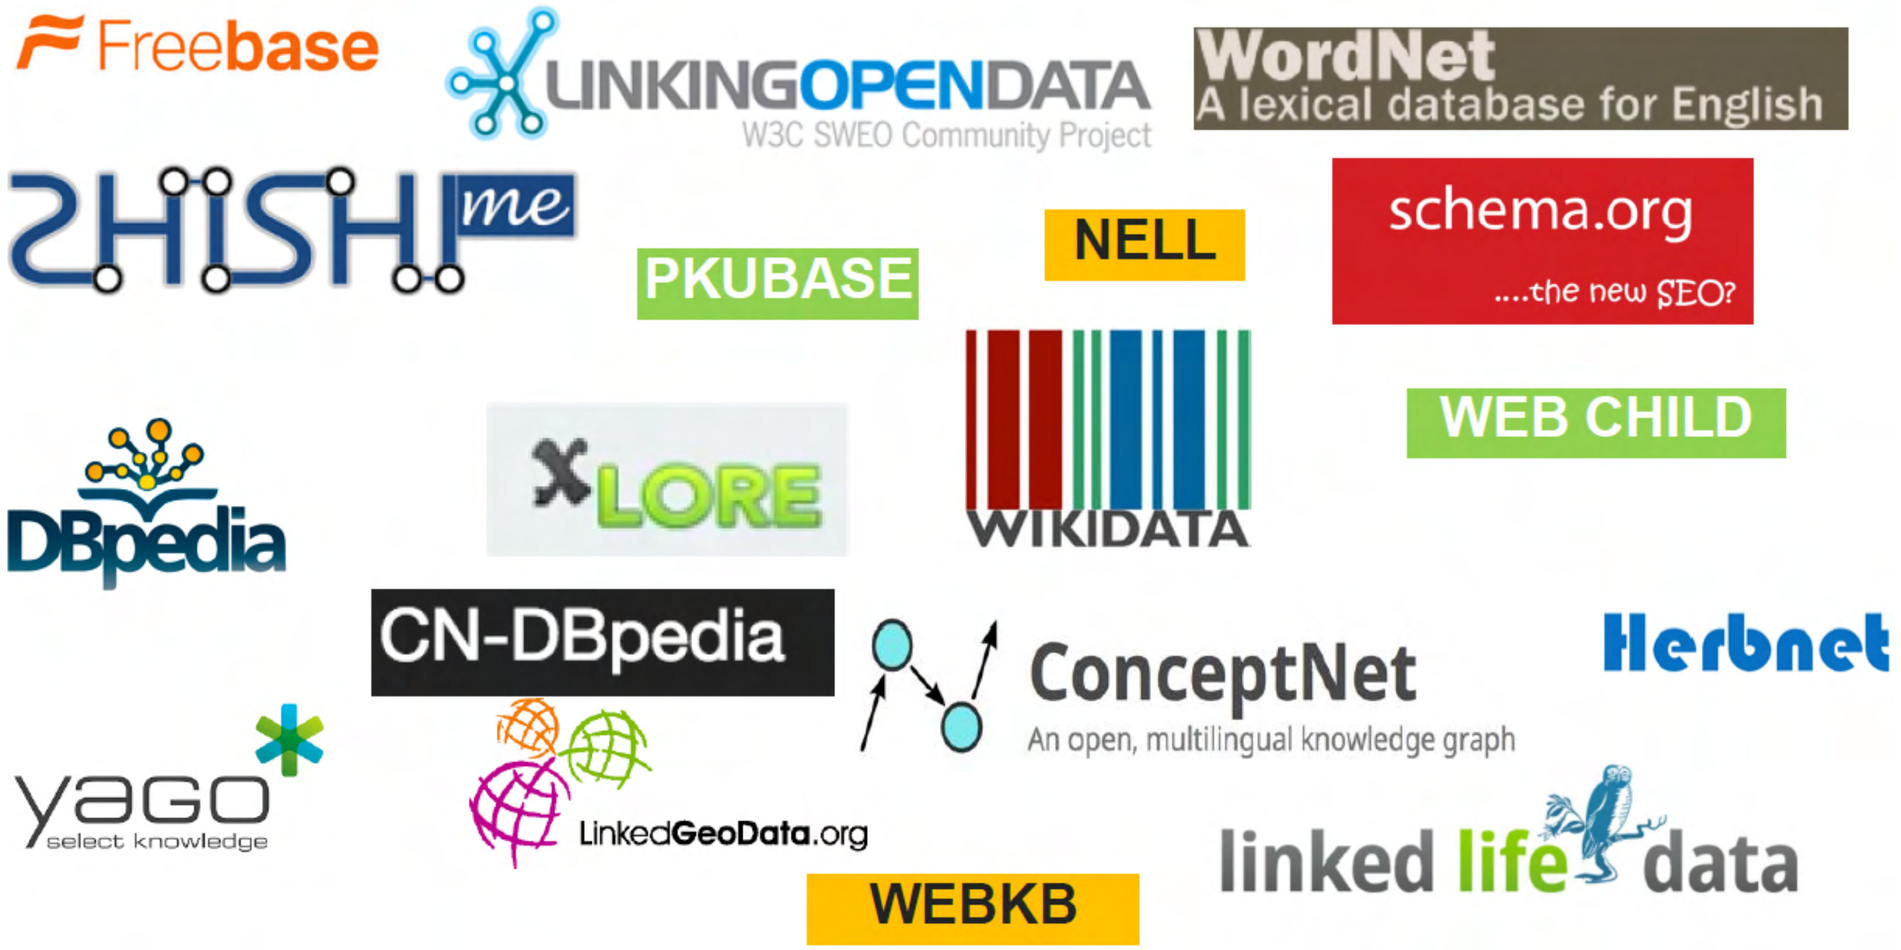

General Applications: Semantic Search

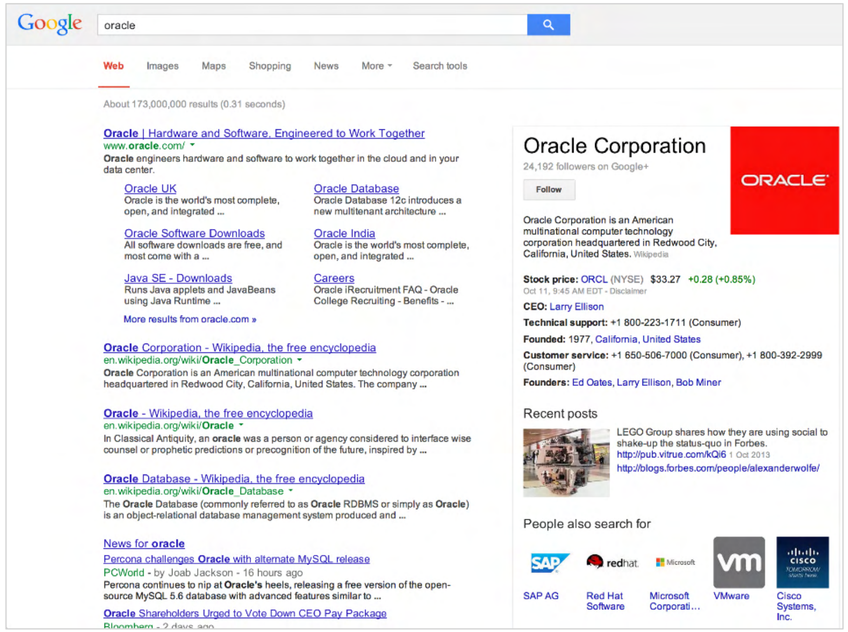

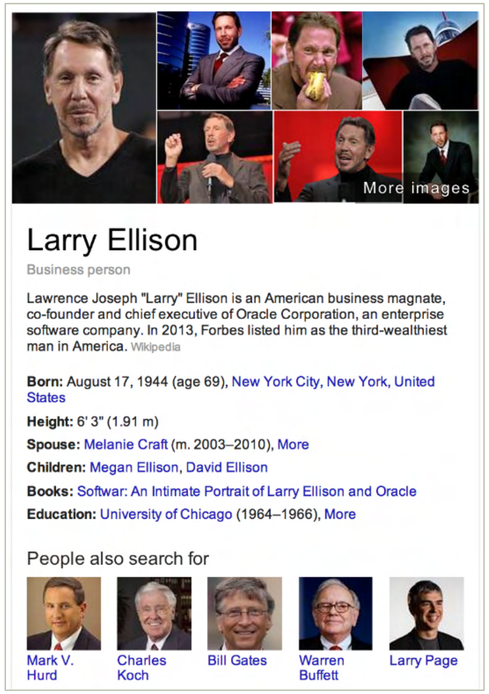

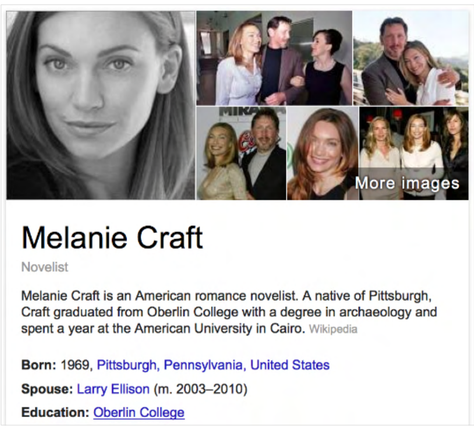

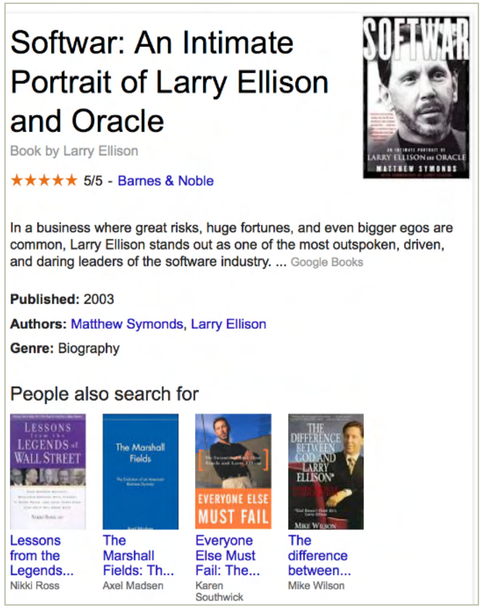

General Applications: Question Answering

General Application: Recommender system

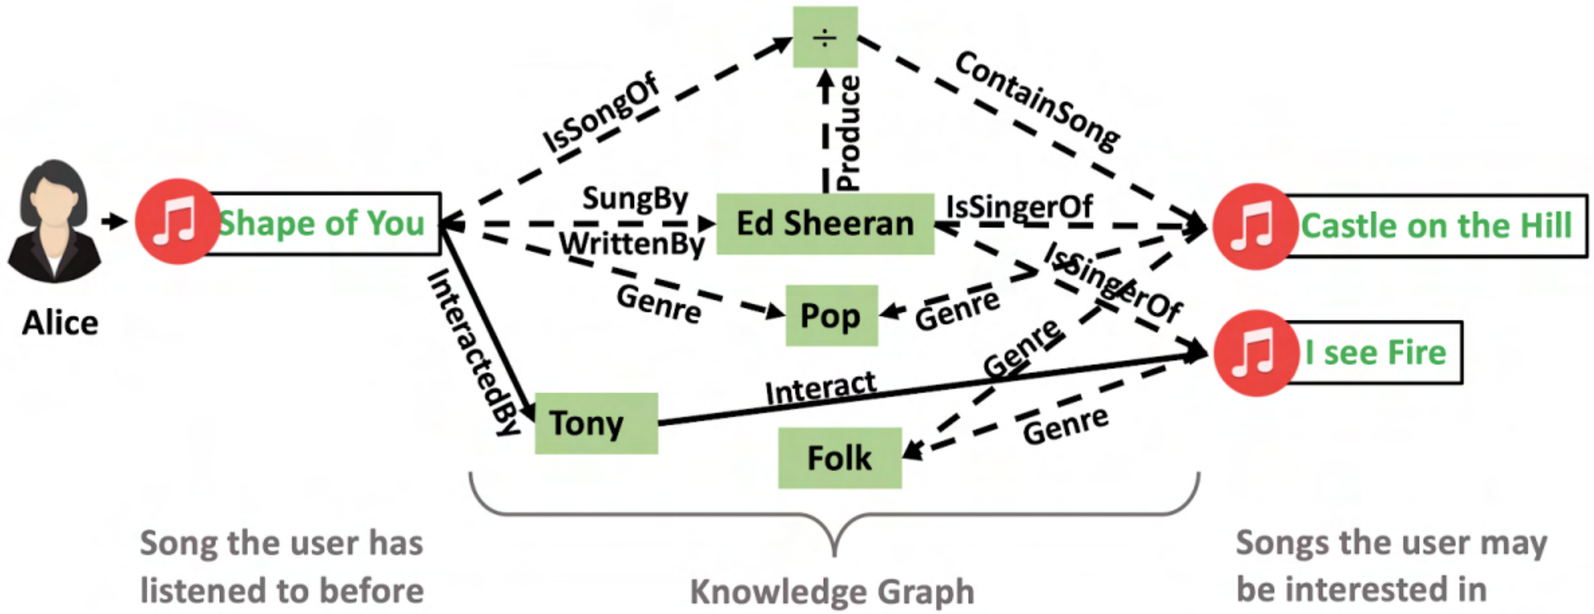

2. The Virtual Knowledge Graph Framework

Reality

- Knowledge Graphs are great, but how to get them?
- Many organization and companies are still relying on large relational databases to manage their internal data.
- It is not realistic to ask them completely migrate to knowledge graphs

Virtual Knowledge Graphs

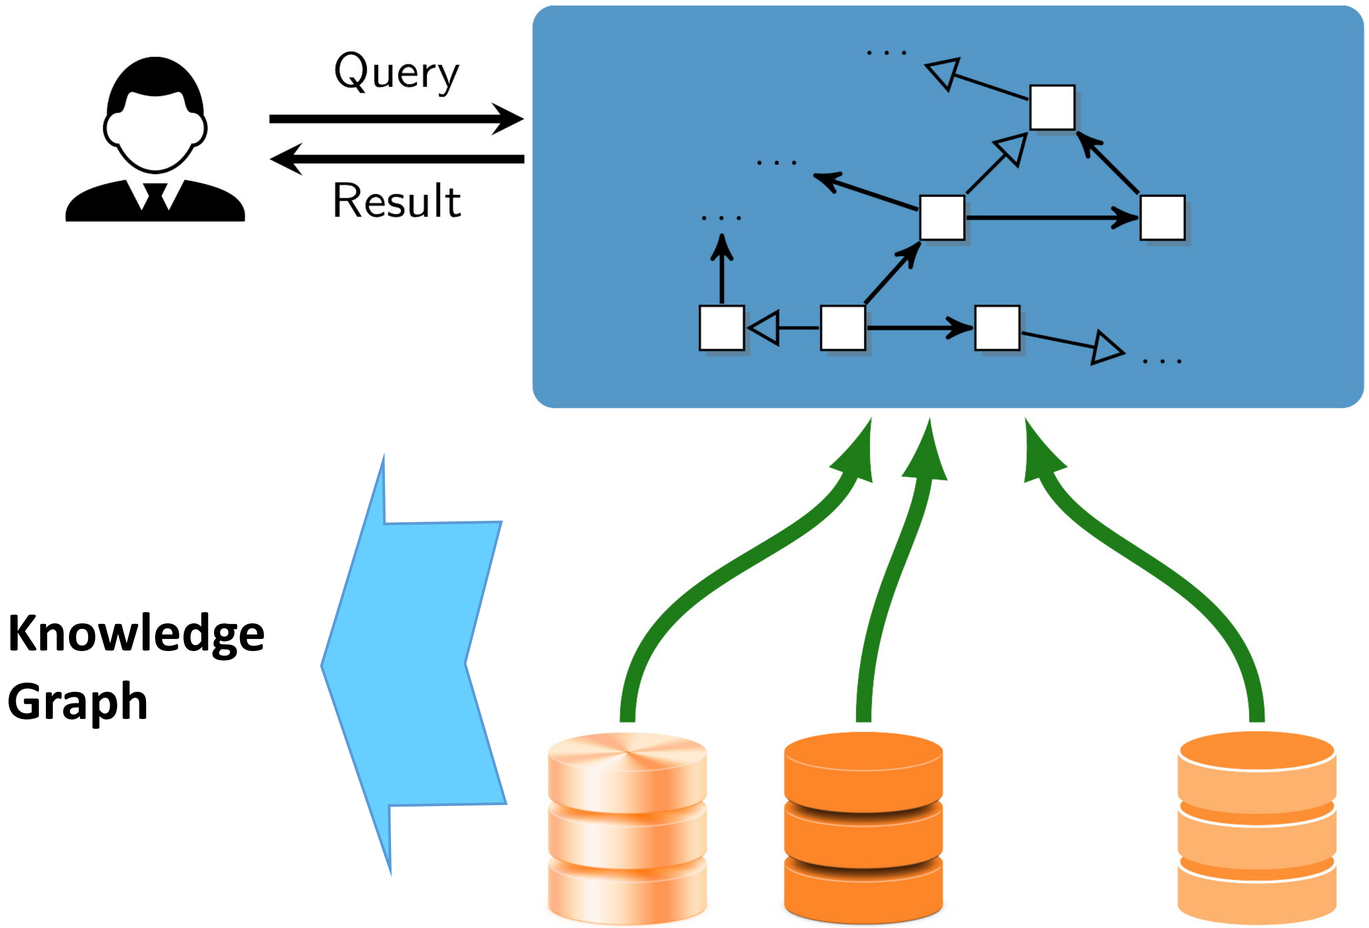

Ontology
provides global vocabulary and conceptual view

Mappings
how to populate the ontology from the data

Data Sources
external and heterogeneous

Why virtualization?

- The data stays in the sources and is only accessed at query time
- There is no need to construct a large and potentially costly materialization pipeline
- The data is always fresh wrt the latest updates at the sources
- One can rely on the existing data infrastructure and expertise
- Better supports an incremental approach to integration

Components of the VKG framework

- We consider now the main components that make up the VKG framework, and the languages used to specify them.

- The W3C has standardized languages suitable for VKGs:
- Knowledge graph: expressed in RDF [W3C Rec. 2014] (v1.1)
- Ontology O: expressed in OWL2QL [W3C Rec. 2012]
- MappingM: expressed in R2RML [W3C Rec. 2012]
- Query: expressed in SPARQL [W3C Rec. 2013] (v1.1)

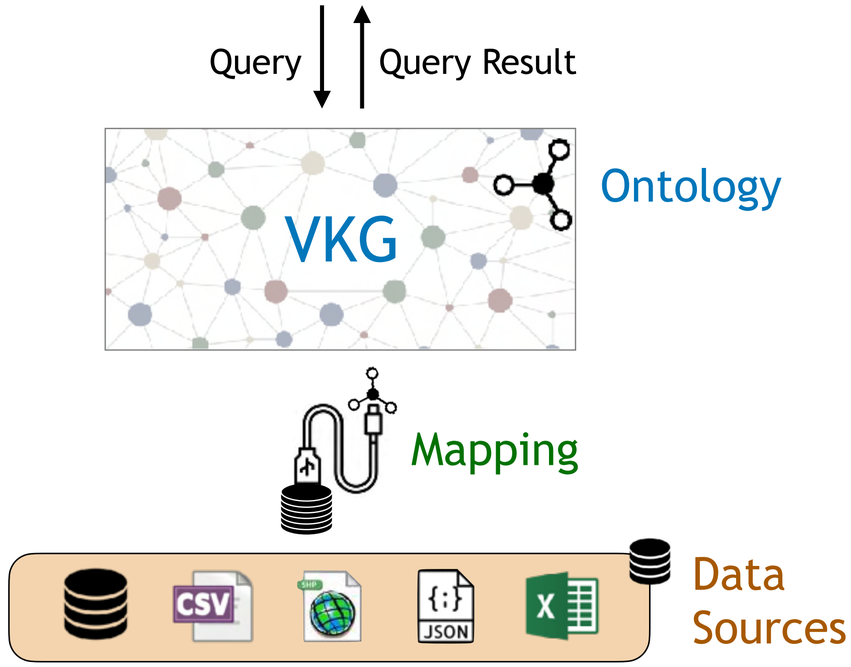

Use of mappings

In the VKG framework, the mapping encodes how the data in the sources should be used to create the Virtual Knowledge Graph, which is formulated in the vocabulary of the ontology.

VKG defined from the mapping and the data.

- Queries are answered with respect to the ontology and the data of the VKG.
- The data of the VKG is not materialized (it is virtual!).
- Instead, the information in the ontology and the mapping is used to translate queries over the ontology into queries formulated over the sources.

Note: The graph is always up to date wrt the data sources.

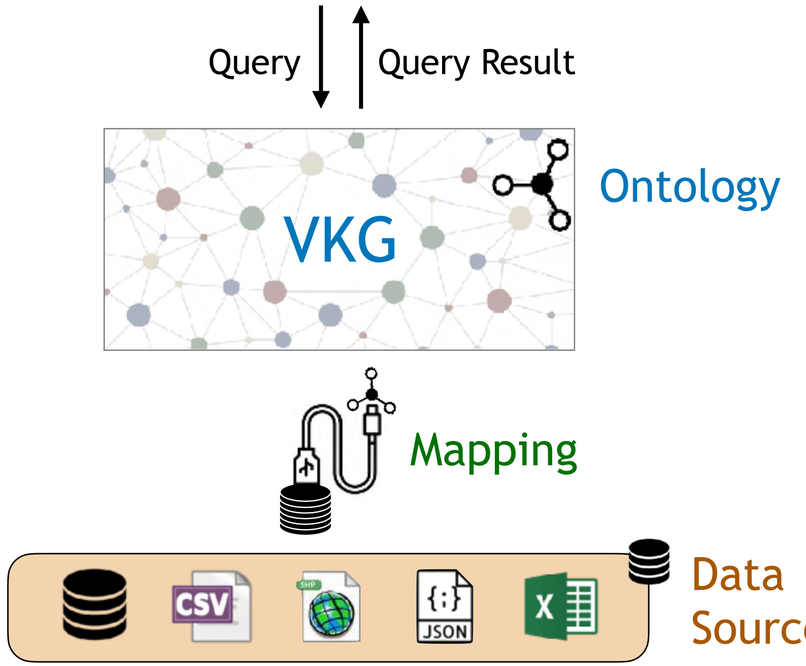

VKG Approach: Query answering by rewriting

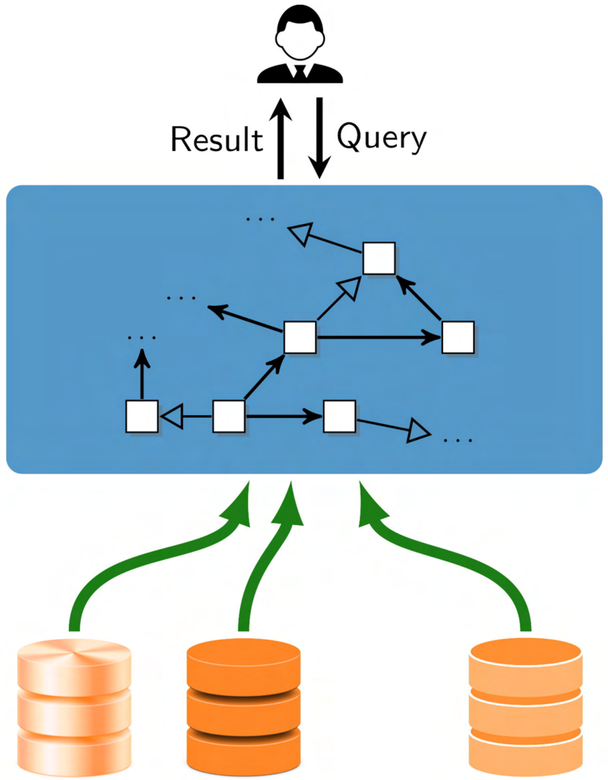

VKG Approach: Query answering by rewriting

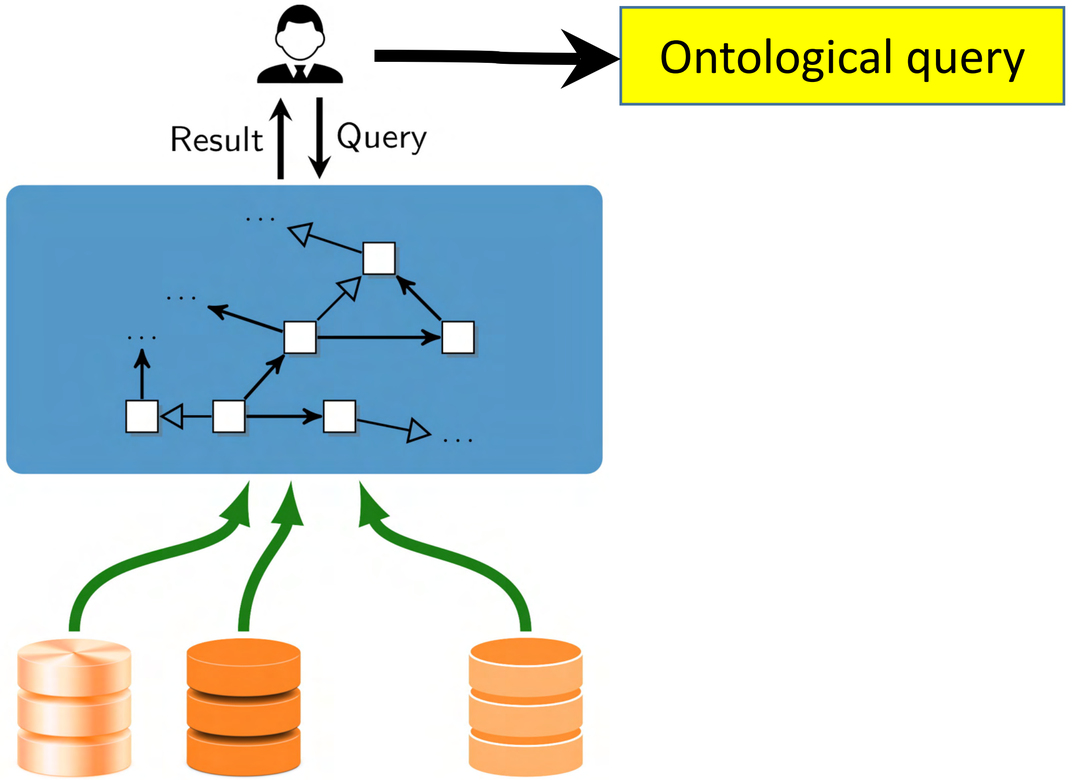

VKG Approach: Query answering by rewriting

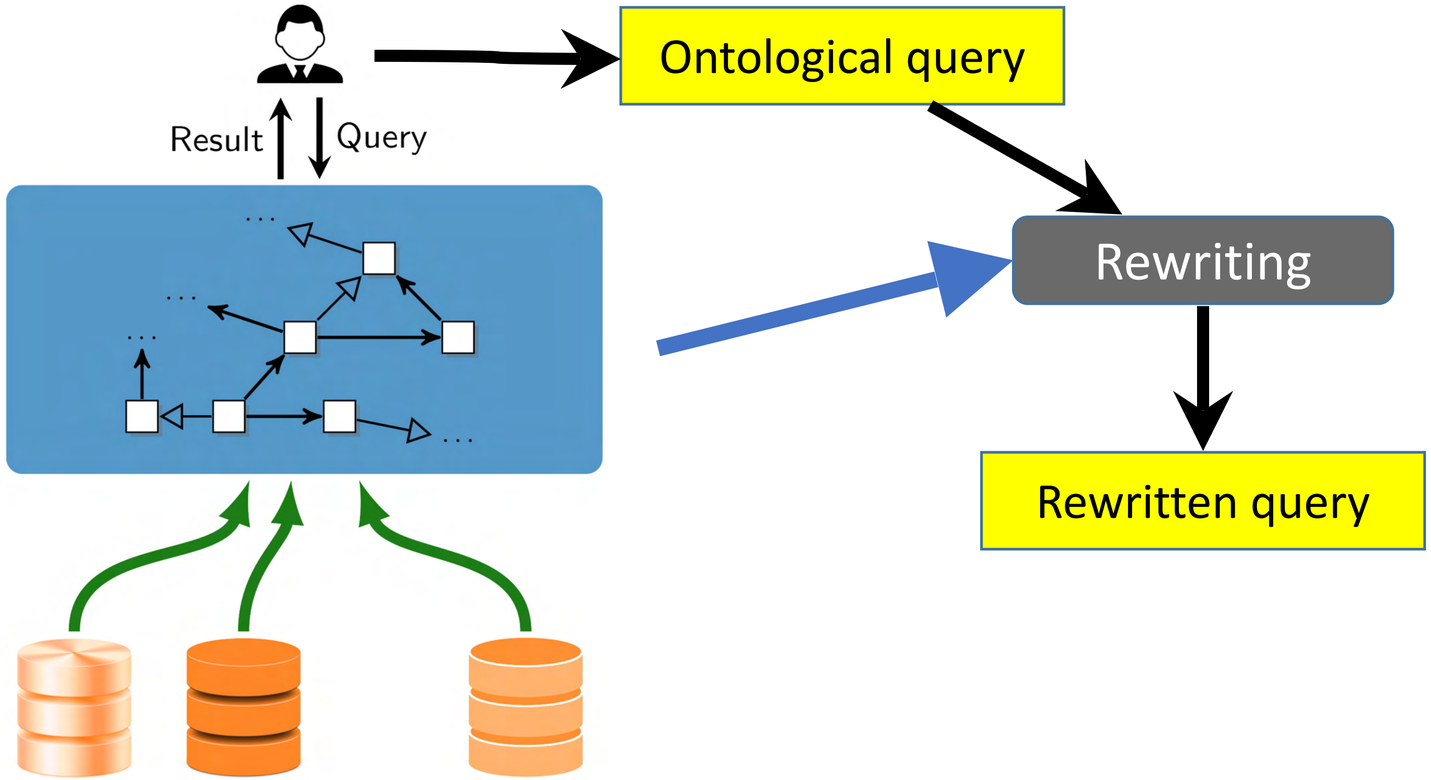

VKG Approach: Query answering by rewriting

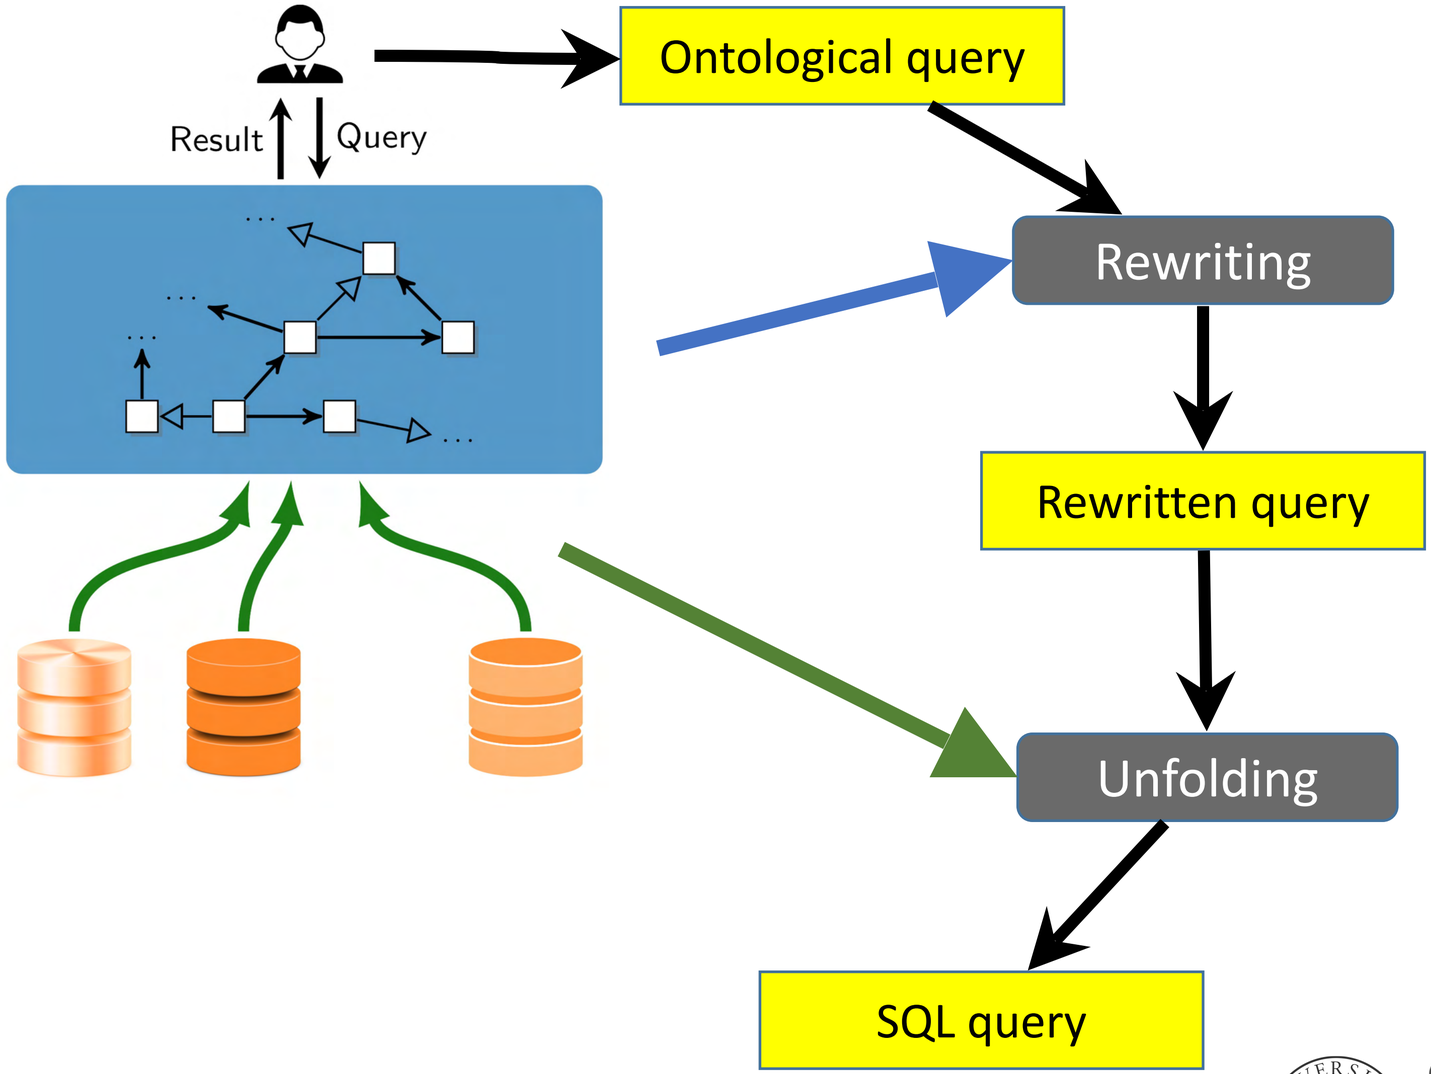

VKG Approach: Query answering by rewriting

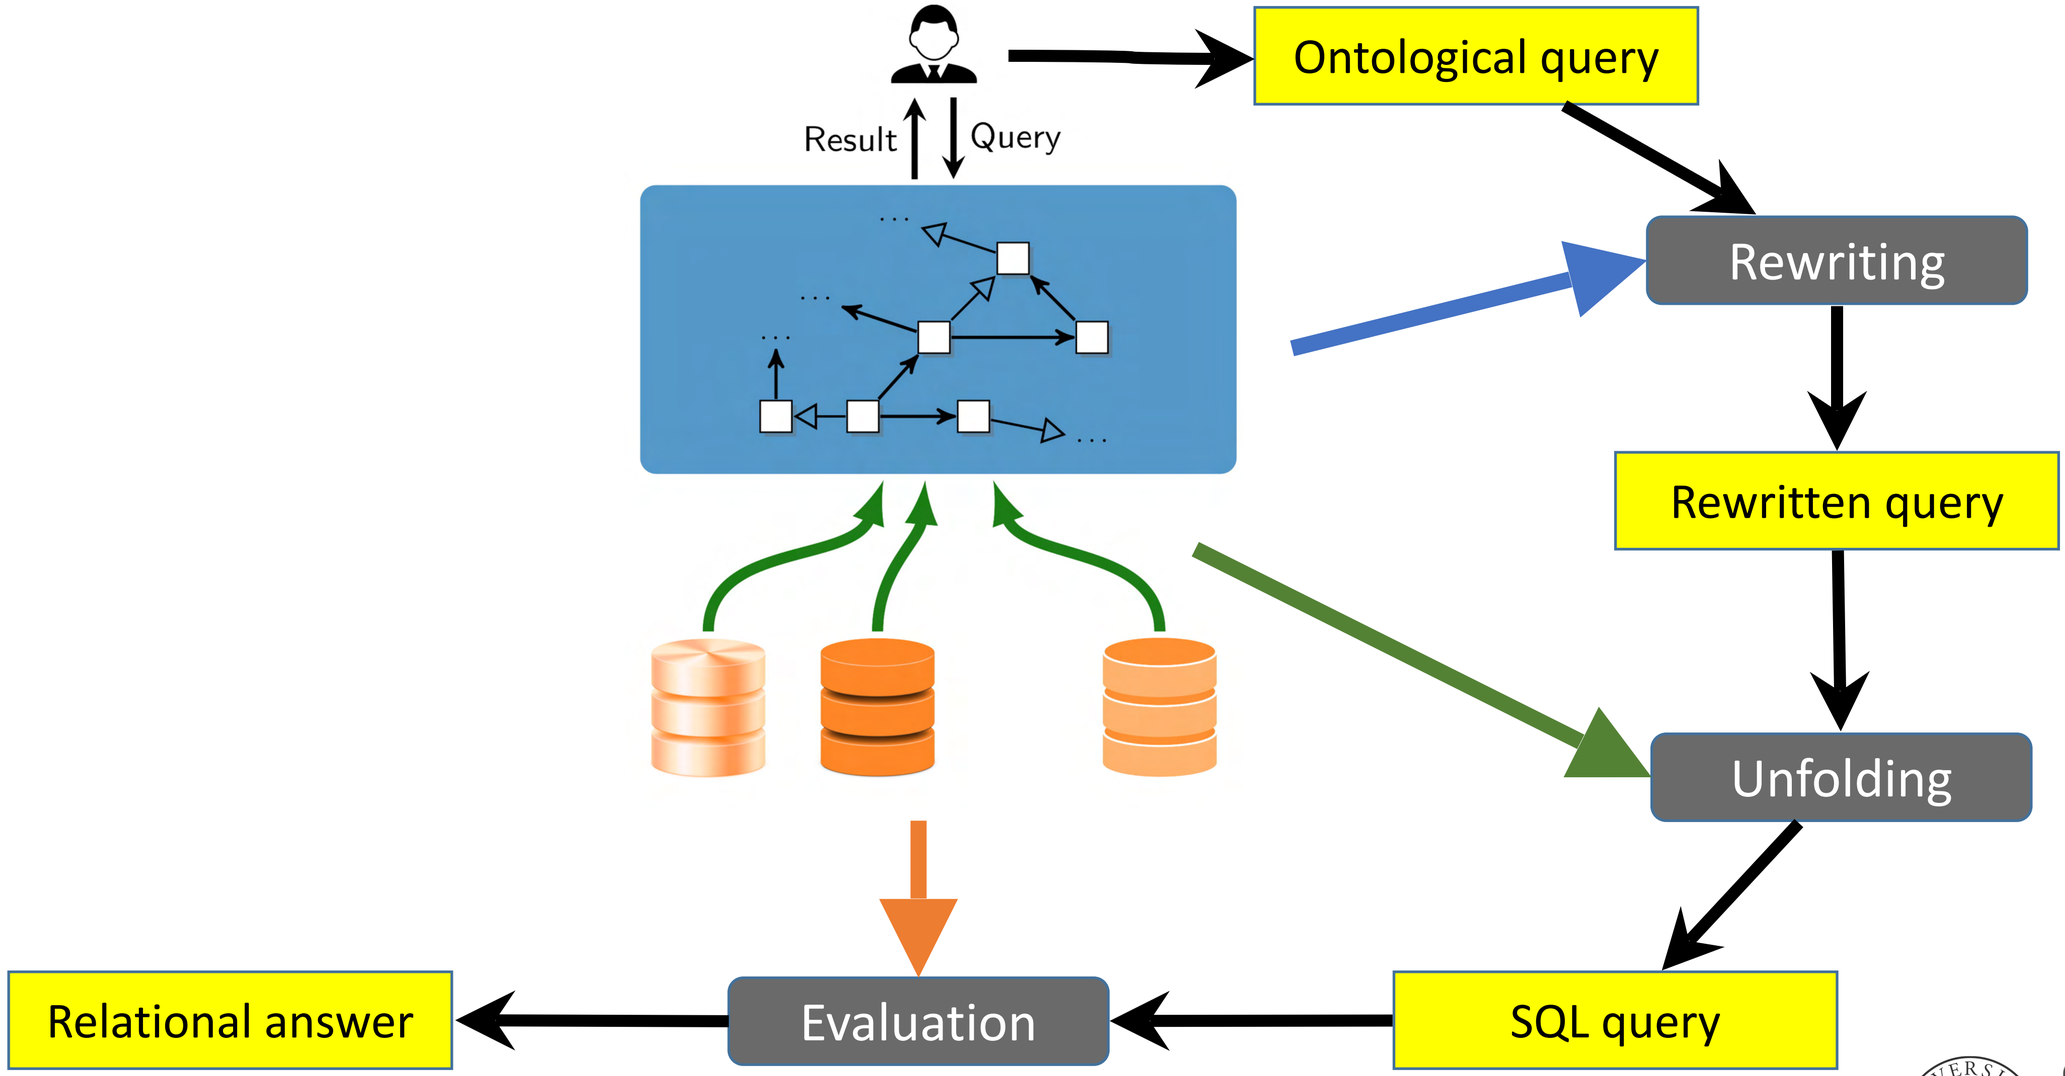

VKG Approach: Query answering by rewriting

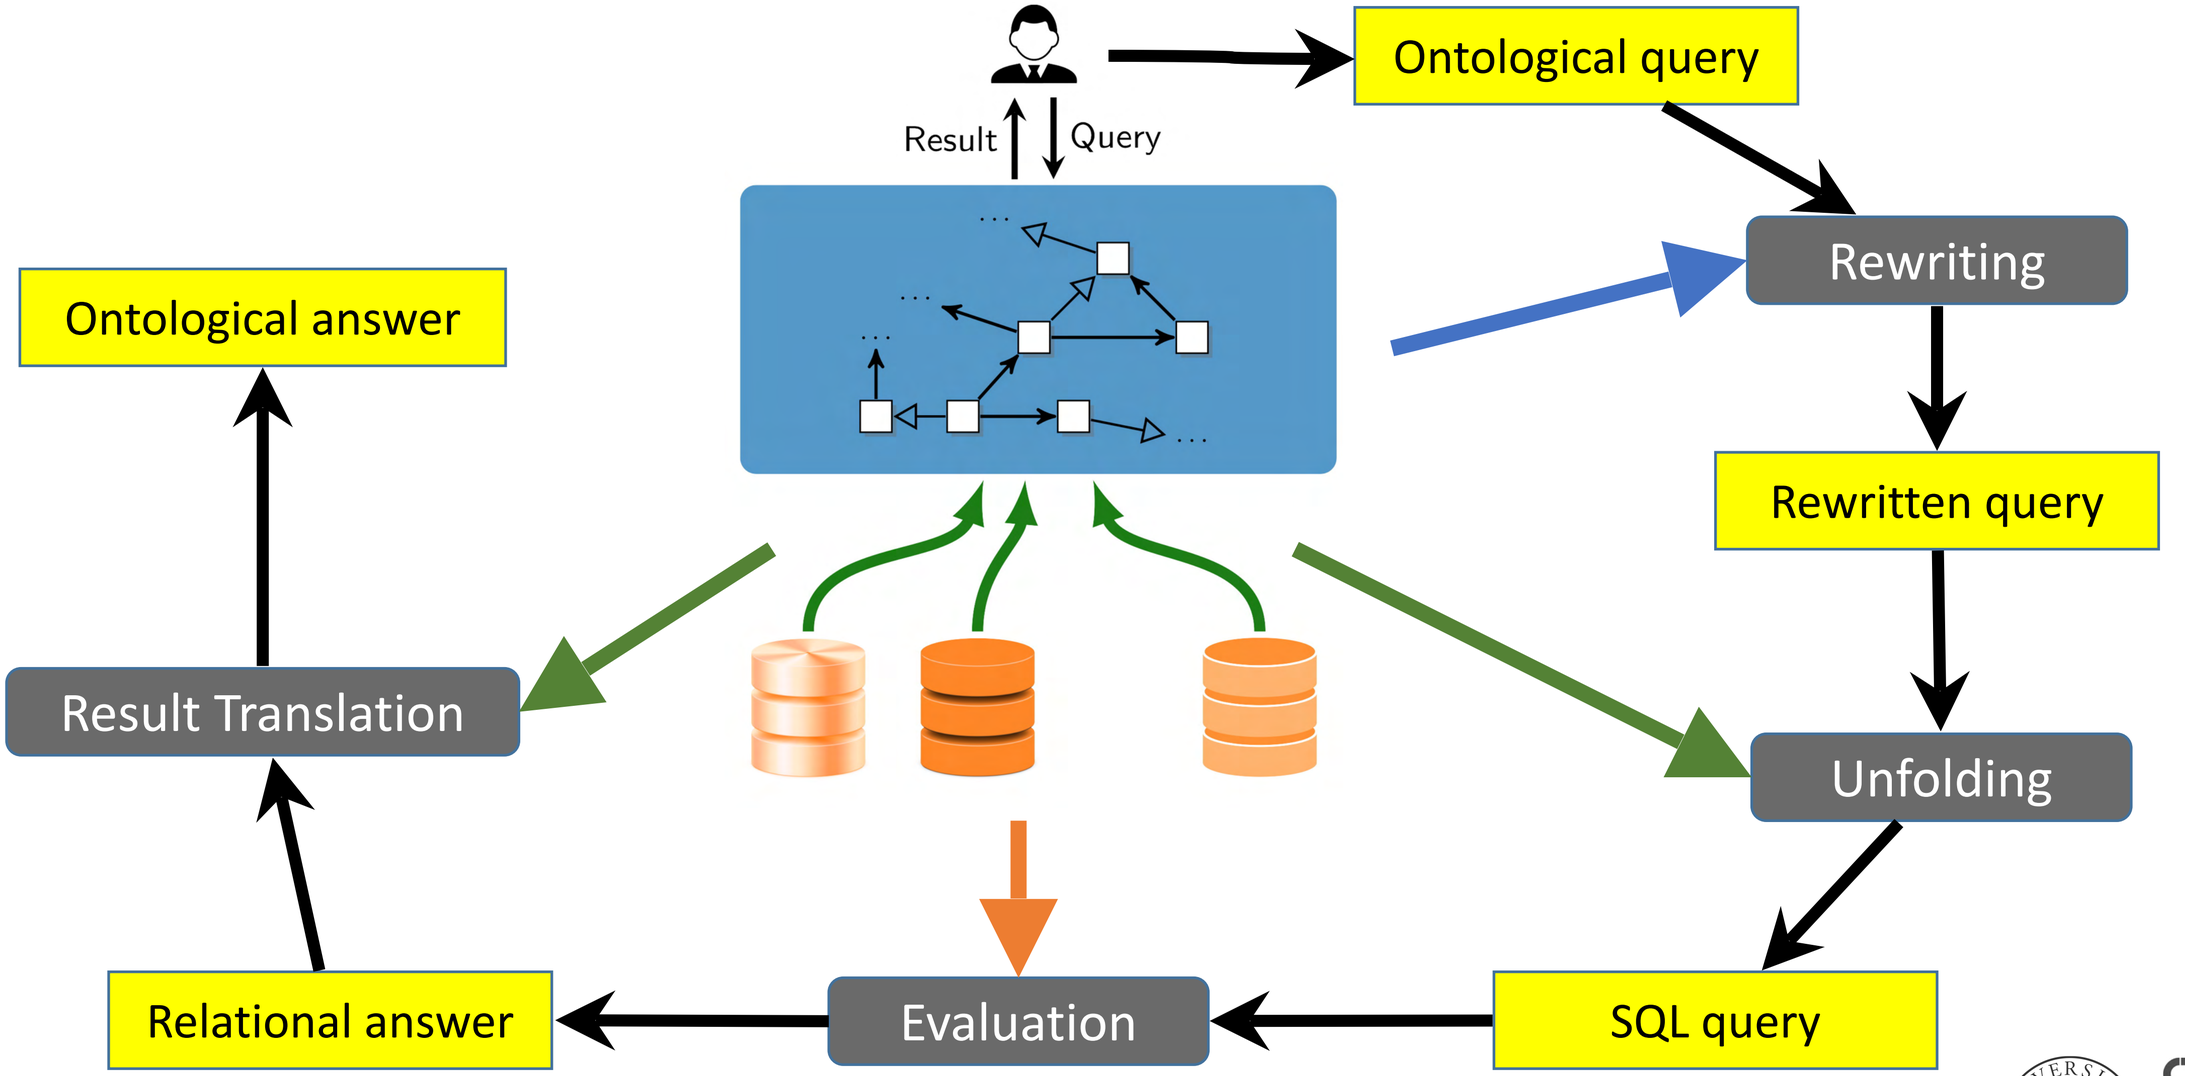

VKG Approach: Query answering by rewriting

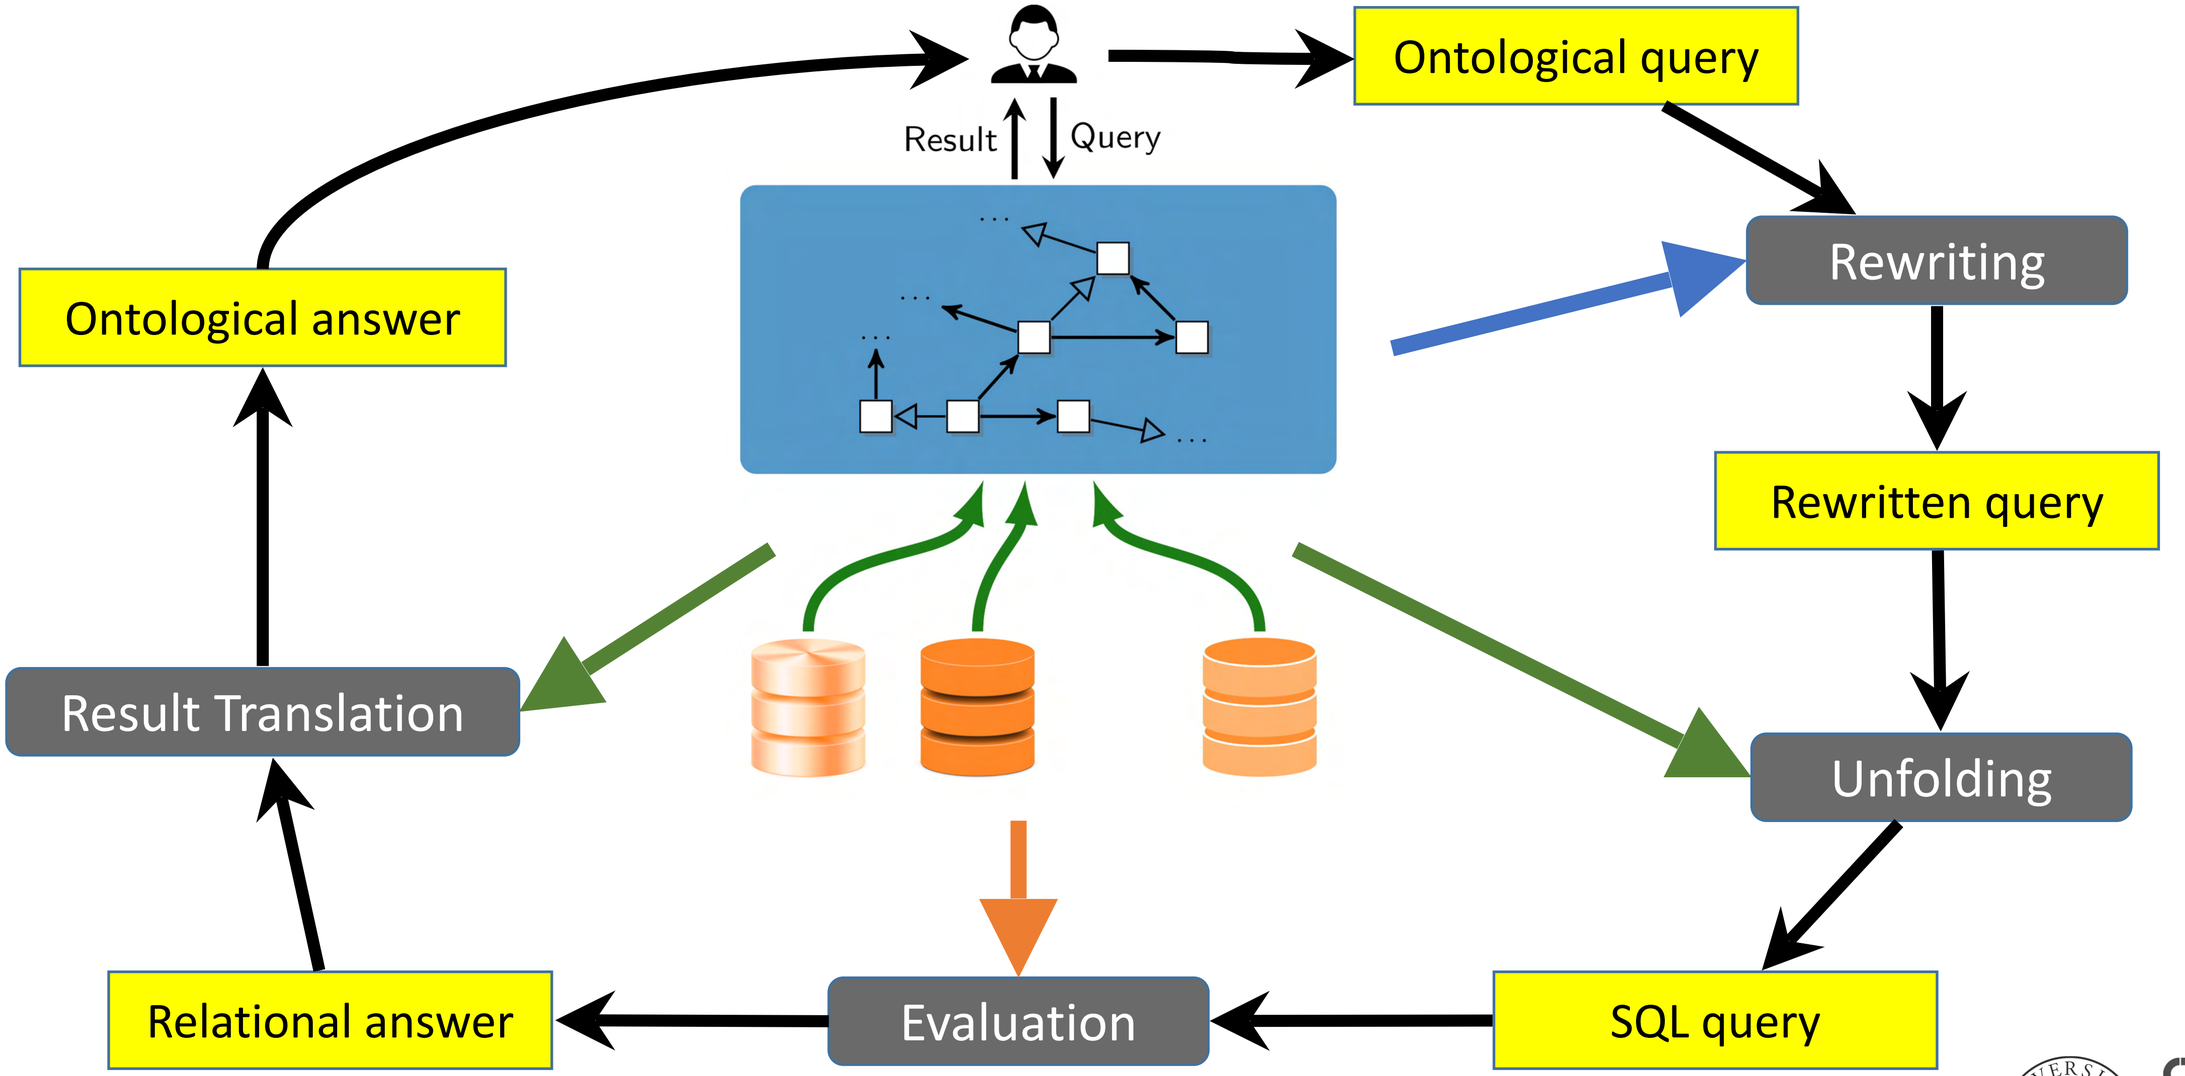

3. VKG Systems and Applications

VKG systems

- Mastro, Sapienza Università di Roma & OBDA systems SRL, Italy
- Morph, Technical University of Madrid, Spain
- Ontop, Free Uni. of Bozen-Bolzano, Italy & Ontopic SRL, Italy
- Stardog, Stardog Union, US
  - powered by Ontop in the first version, later own implementation
- GraphDB, OntoText, US & EU
  - powered by Ontop
- Allegro Graph, Franz Inc., US
  - powered by Ontop
- Oracle Spatial and Graph, US

The VKG system Ontop

- State-of-the-art VKG system developed at the Free University of Bozen-Bolzano.
- Compliant with all relevant Semantic Web standards:

  RDF, RDFS, OWL 2 QL, R2RML, and SPARQL
- Supports all major relational DBs:

  Oracle, DB2, MS SQL Server, Postgres, MySQL, Teiid, Dremio, Denodo, etc.
- Open-source and released under Apache 2 license.
- Development: started in 2009
  - Already well established (e.g., ~5000 downloads / year, +200 members in mailing list)
  - Recently: major release of v4, with several improvements and new features.

Developer Community

unibz

ON TOP IC

Birkbeck UNIVERSITY OF LONDON

UNIVERSITAS BERGENSIS

UiO : University of Oslo

HELLENIC REPUBLIC
National and Kapodistrian University of Athens

UNIVERSITÄT LEIPZIG

ontotext

POLITECNICO MILANO 1863

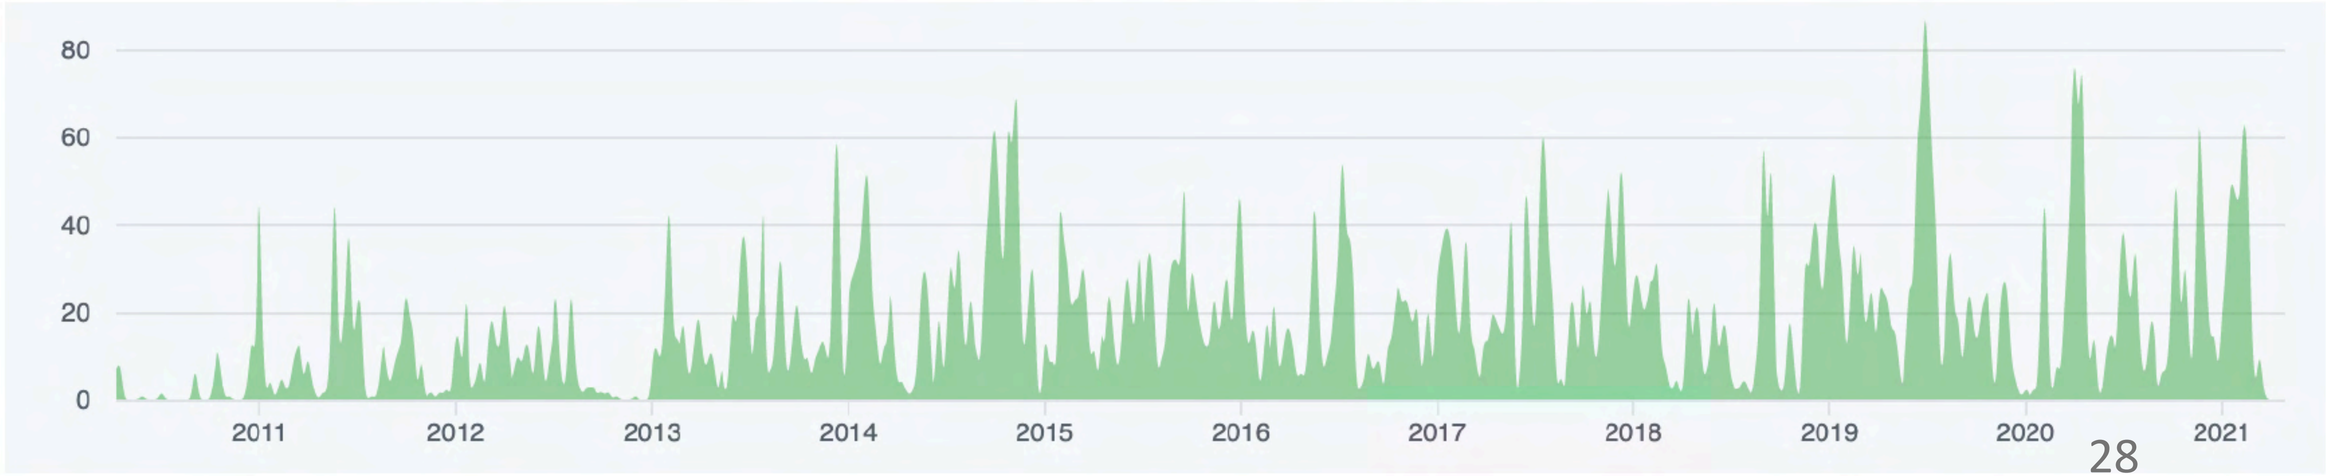

ON TOP IC

Ontop Downloads

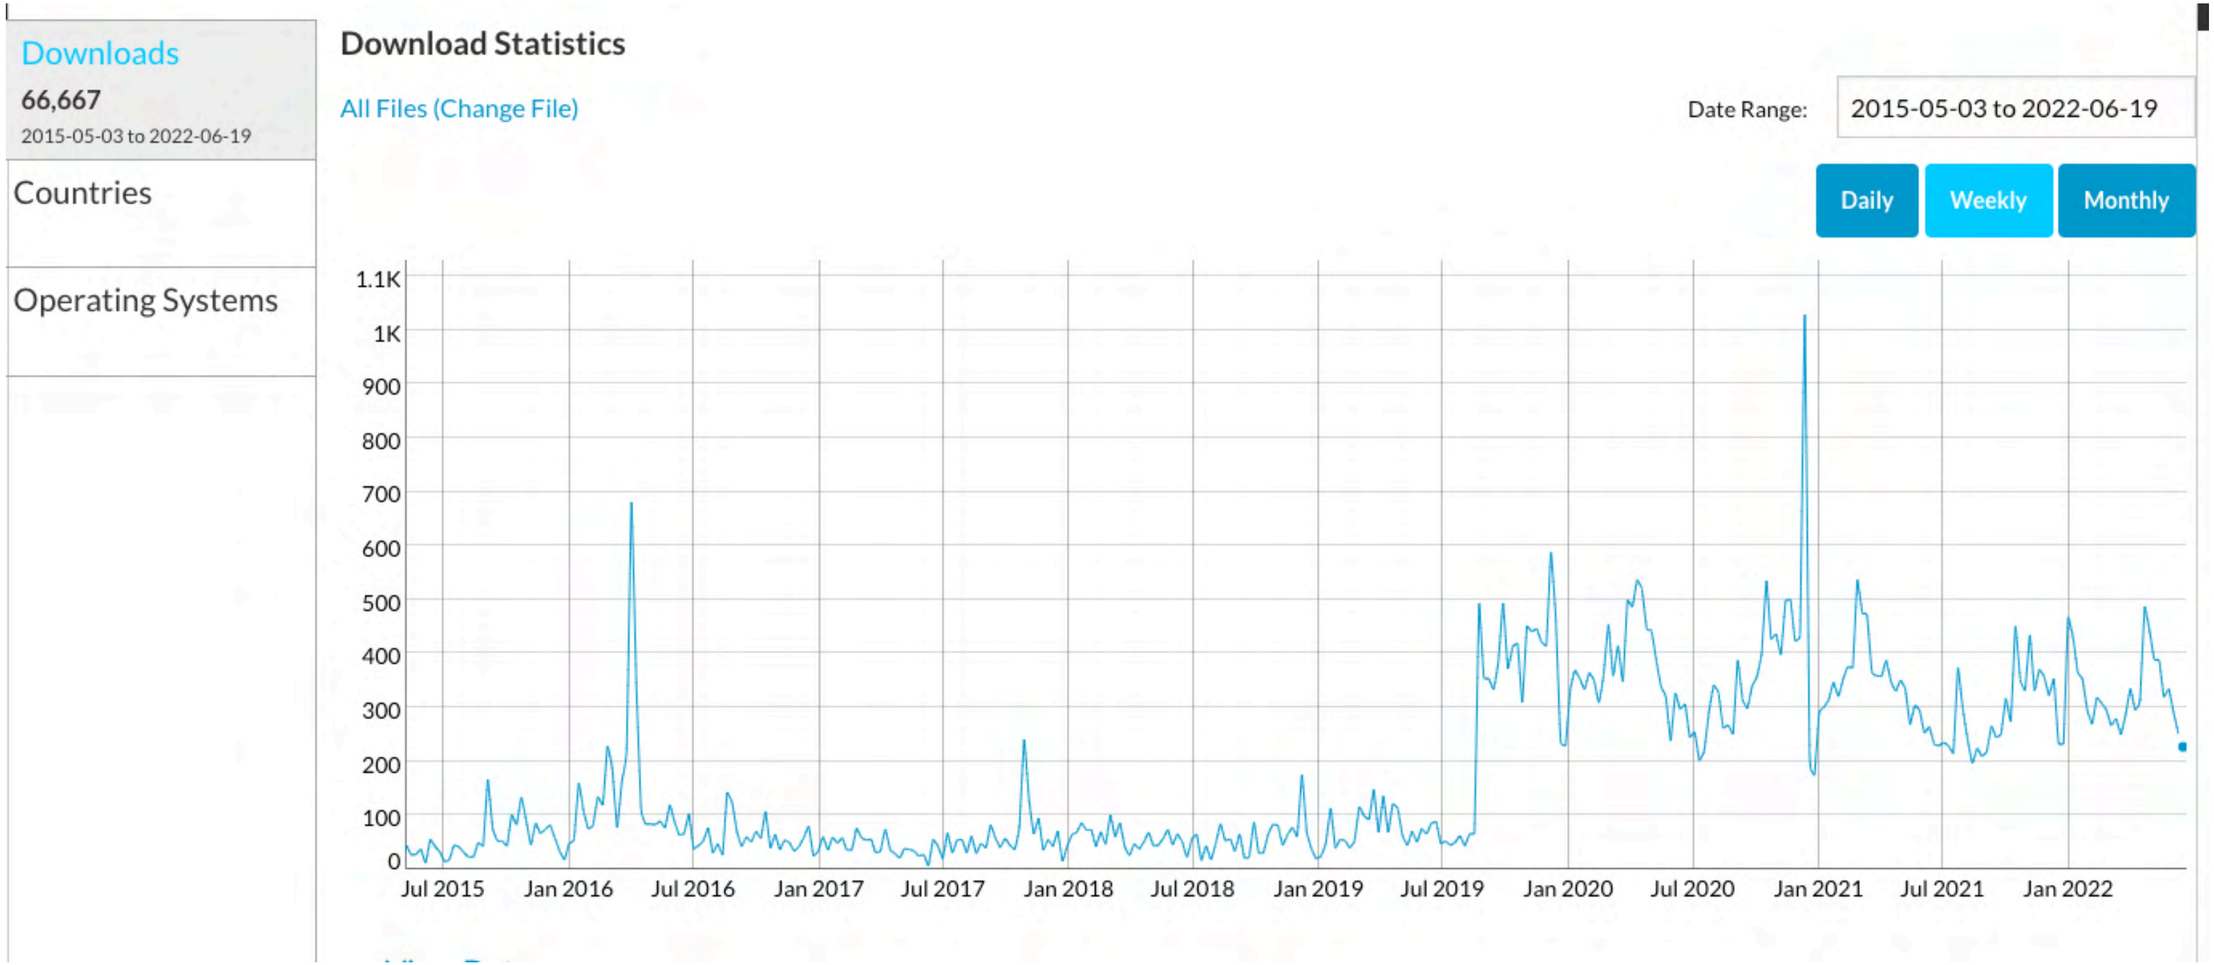

Downloads
66,667
2015-05-09 to 2022-06-19

Countries
Top: **US**, at 14%

Operating Systems
Top: **Other**, at 73%

Download Statistics

All Files (Change File)

Date Range: 2015-05-09 to 2022-06-19

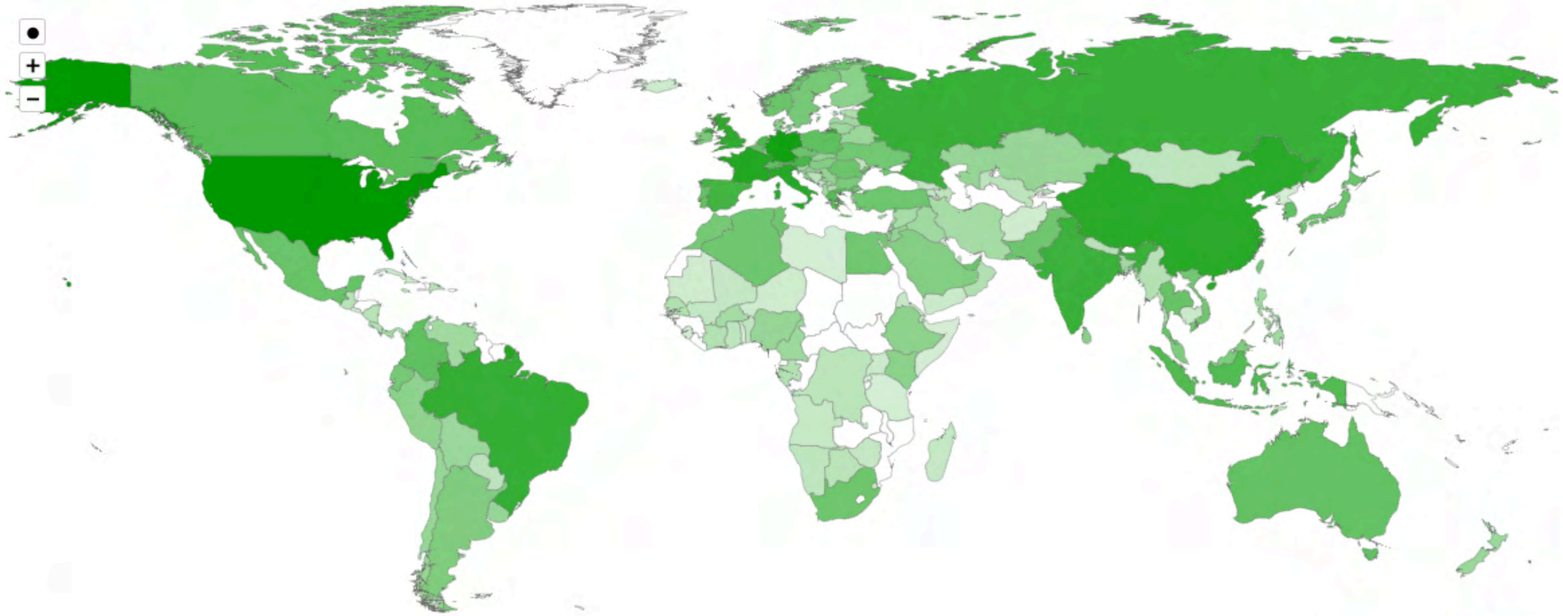

OS downloads as: Percent

	Country	Android	BSD	Linux	Macintosh	Unknown	Windows	Total
1.	United States	0%	0%	2%	5%	85%	9%	9,653
2.	Germany	0%	0%	6%	2%	75%	16%	5,958
3.	France	0%	0%	6%	3%	82%	9%	4,182
4.	Italy	0%	0%	8%	14%	53%	25%	3,808
5.	China	0%	0%	2%	4%	28%	66%	3,653
6.	United Kingdom	0%	0%	3%	7%	72%	18%	3,126
7.	Brazil	0%	0%	4%	3%	69%	24%	2,695
8.	Russia	0%	0%	3%	1%	81%	16%	2,401

ONTOPIC

- https://ontopic.ai
- First spin-off of the Free University of Bozen-Bolzano.
- Incorporated in April 2019.
- Product: *Ontopic Studio* based on the Ontop engine.
- Delivers services around Ontop and the Ontopic Suite.
- Consultancy for VKG-based data integration projects.

Selected Use Cases

- Oil & Gas: Statoil [Kharlamov et al. 2017a]
- Turbine Diagnoses: Siemens [Kharlamov et al. 2017b]
- Log Extraction in Process Mining [Calvanese, Kalayci, et al. 2017]
- Cultural heritage [Calvanese, Liuzzo, et al. 2016]
- Maritime security [Brüggemann et al. 2016]
- Manufacturing [Petersen et al. 2017]
- Health care: electronic health records [Rahimi et al. 2014]
- Public debt: the Italian Ministry of Economy and Finance [Antonioli et al. 2014]
- Smart cities: IBM Ireland [Lopez et al. 2015]
- Knowledge resource management: Huawei & Ontopic
- and many more … (see our survey [Xiao et al . 2019])

VKG over LinStar in Huawei

Goal: build a virtual knowledge graph integrating structured data in a proprietary platform (LinStar) and the results of information extraction from related semi-structured data.

- Structured data is provided by a relational database.
- Semi-structured data consisting of text with little (if any) structure or markup, such as natural language text, HTML documents, PDF files.
- Information extraction (IE) aims at extracting structured information from semi-structured data, possibly leveraging natural language processing (NLP) techniques.

Motivations:

- Provide an unambiguous formalization of the knowledge in the platform, to ease exploitation.
- Provide an integrated, queryable, up-to-date view over all available information.
- Enable more advanced services, such as intelligent search and intelligent recommendation.

High-Level architecture of VKG over semi-structured data

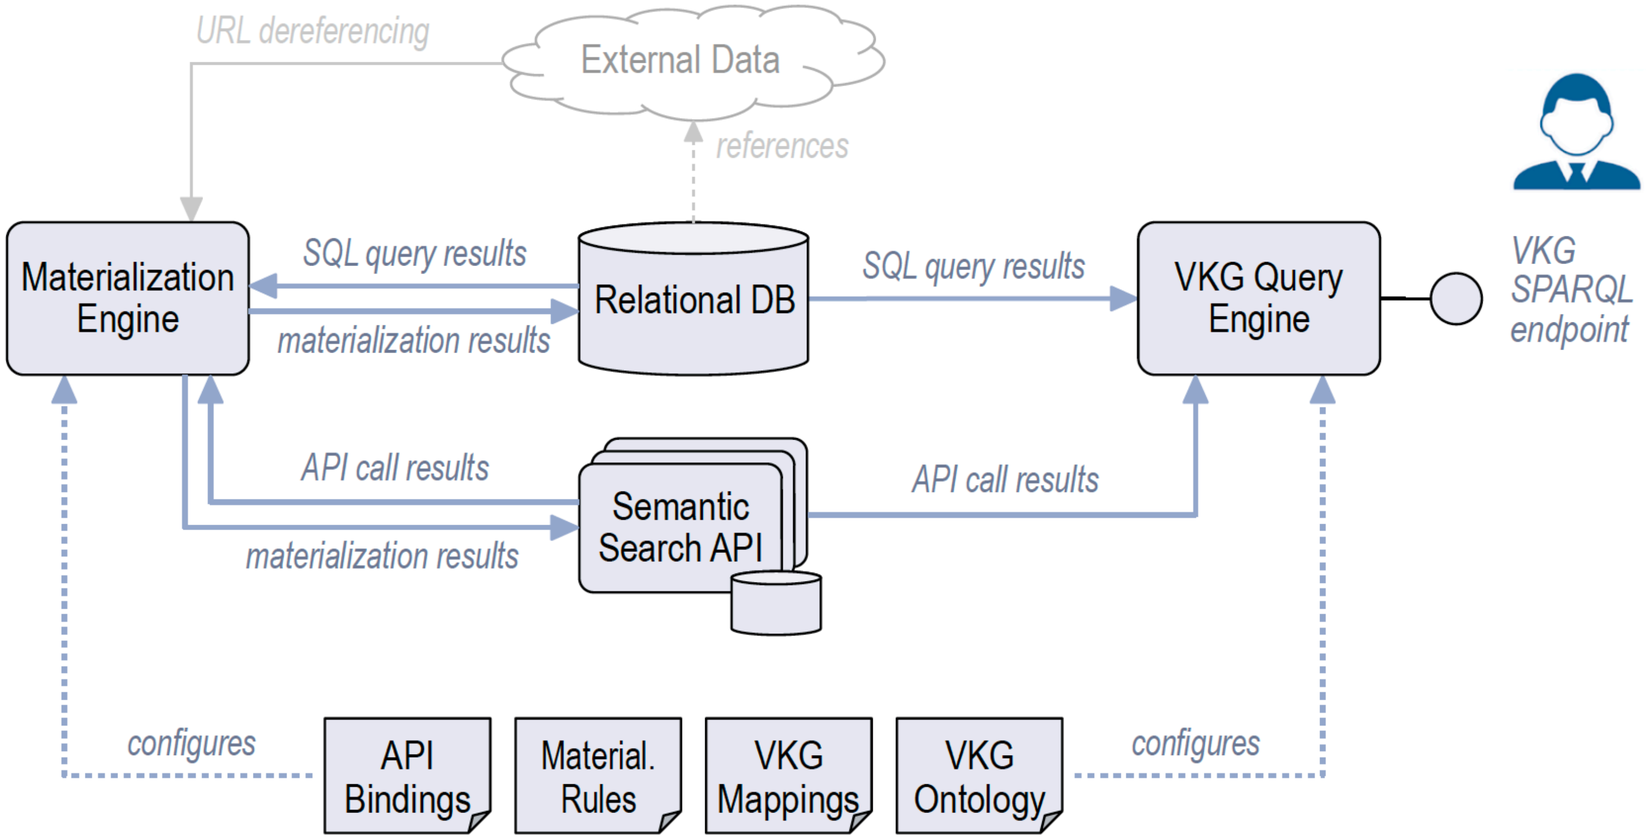

3. Key Technologies in Virtual Knowledge Graphs

Some Knowledge Graph technologies

- RDF Graph Language
- SPARQL Query Language
- OWL Ontology Language
- R2RML Mapping Language

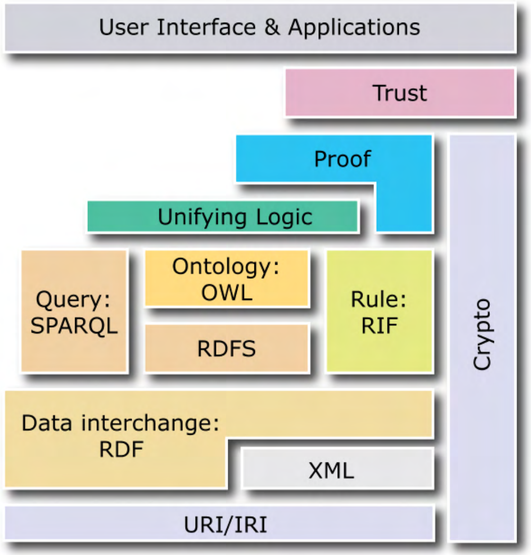

RDF — Data is represented as a graph

The graph consists of a set of subject-predicate-object triples.

Relational DB vs. graph data — Example

Addresses

id integer	istat integer	frac_code integer	strt_code integer	label_it character varying (254)	label_de character varying (254)	num character varying	geom geometry
79520478	21008	0	8280	Via Leonardo Da Vinci 1/F	Leonardo-Da-Vinci-Strasse 1/F	1/F	01010000C...

Municipalities

gem_id integer	istat_code integer	name_i character varying (254)	name_d character varying (254)	area numeric	geom geometry
8	21008	Bolzano	Bozen	52337186.505577500000000	...

Streets

comistat integer	fraistat integer	ascot_wege integer	desc_i character varying (254)	desc_d character varying (254)	geom geometry
21008	0	8280	VIA LEONARDO DA VINCI	LEONARDO-DA-VINCI-STRASSE	01050000C...

Graph data is easy to integrate

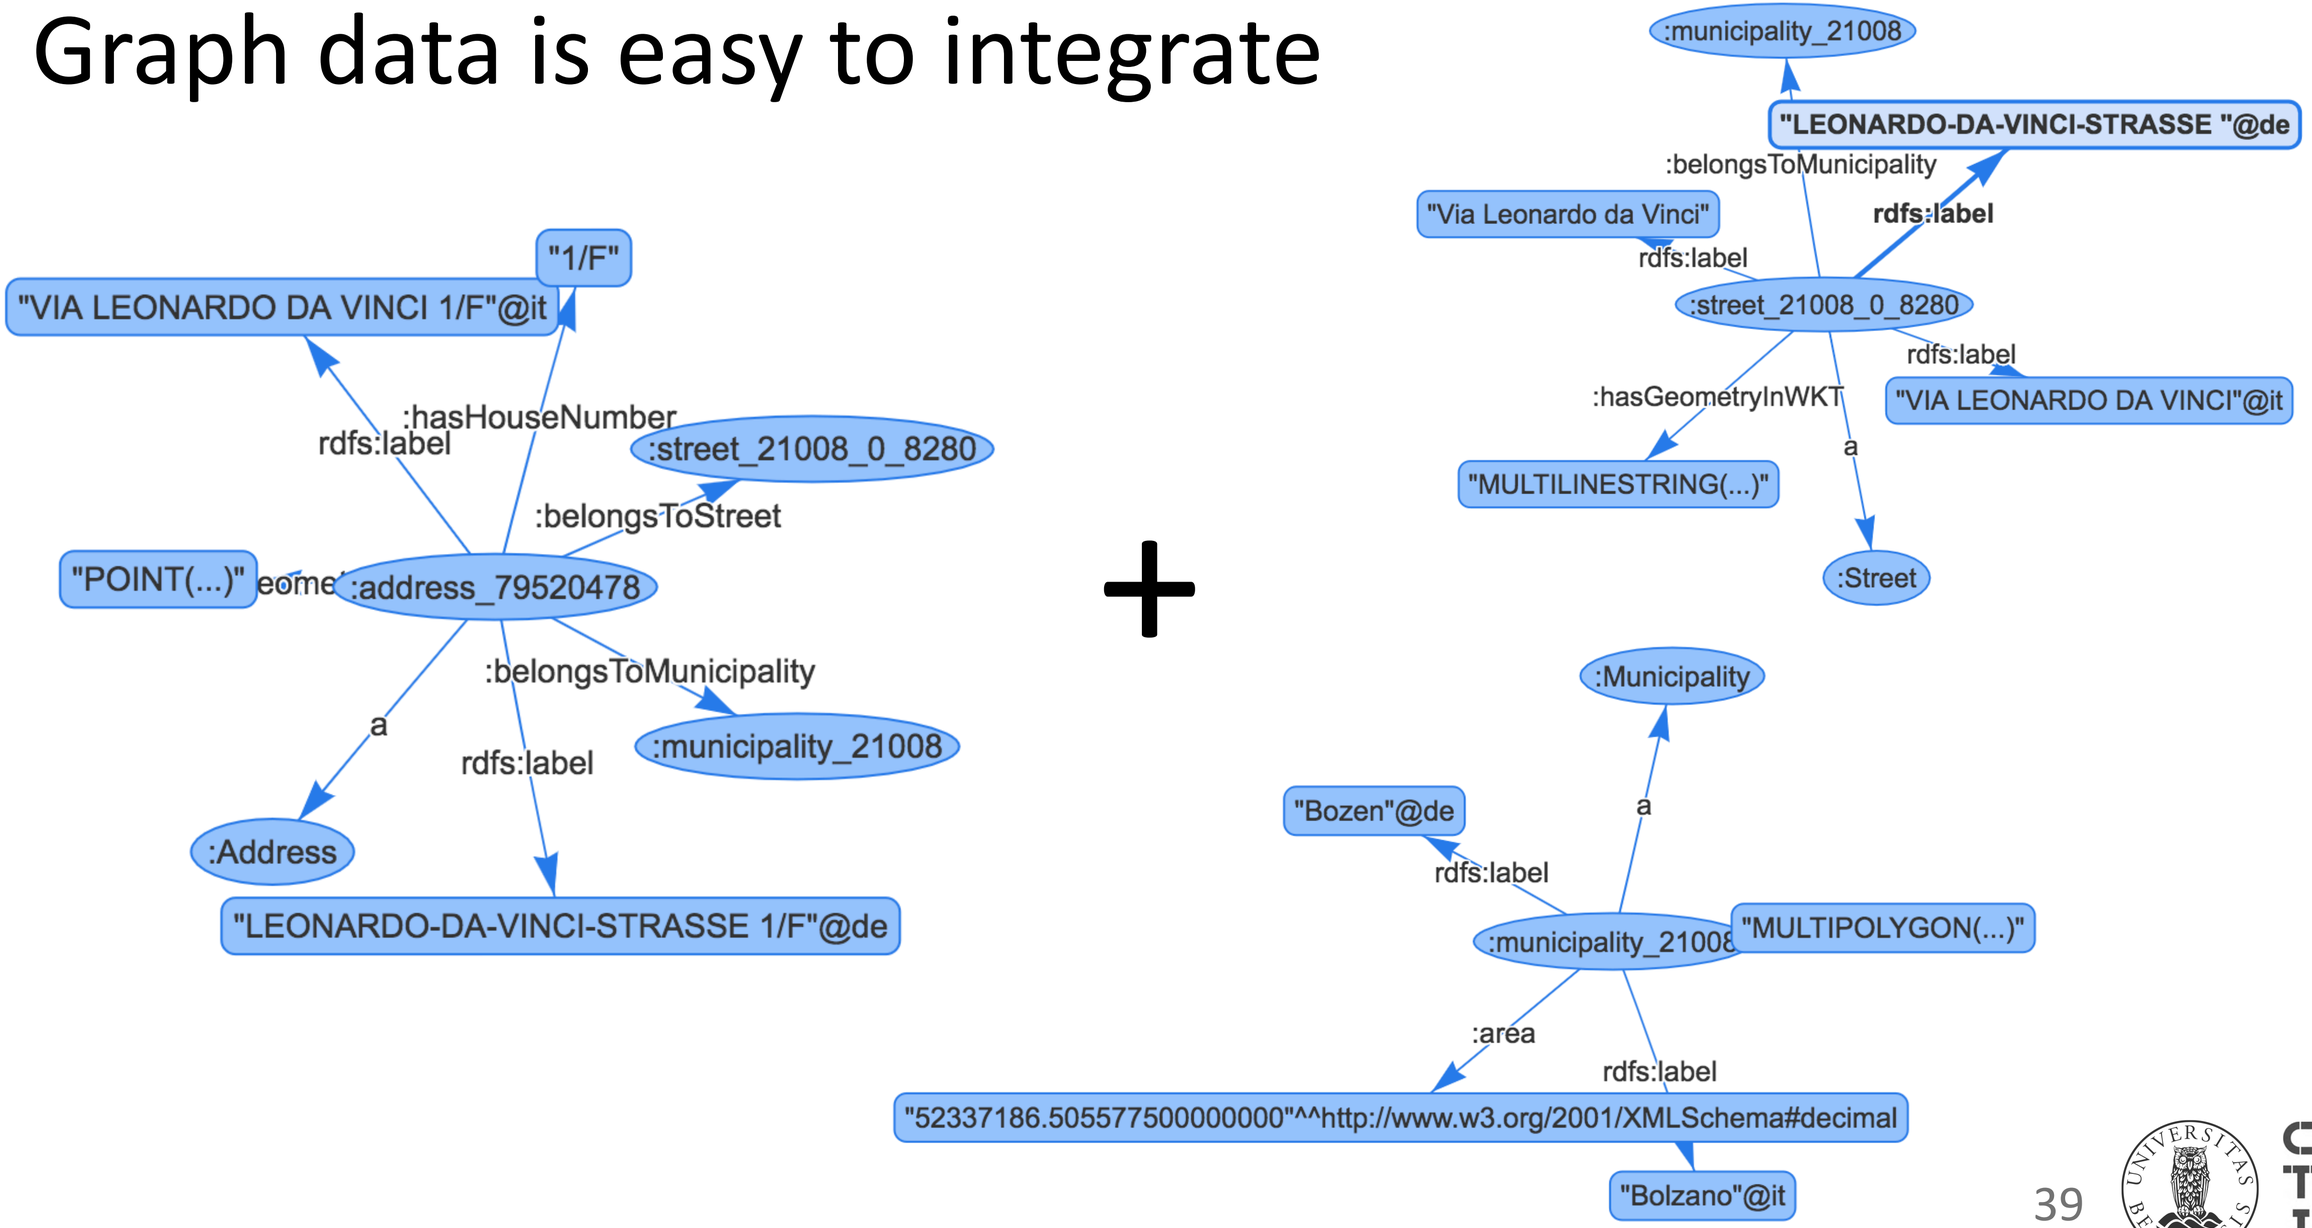

=

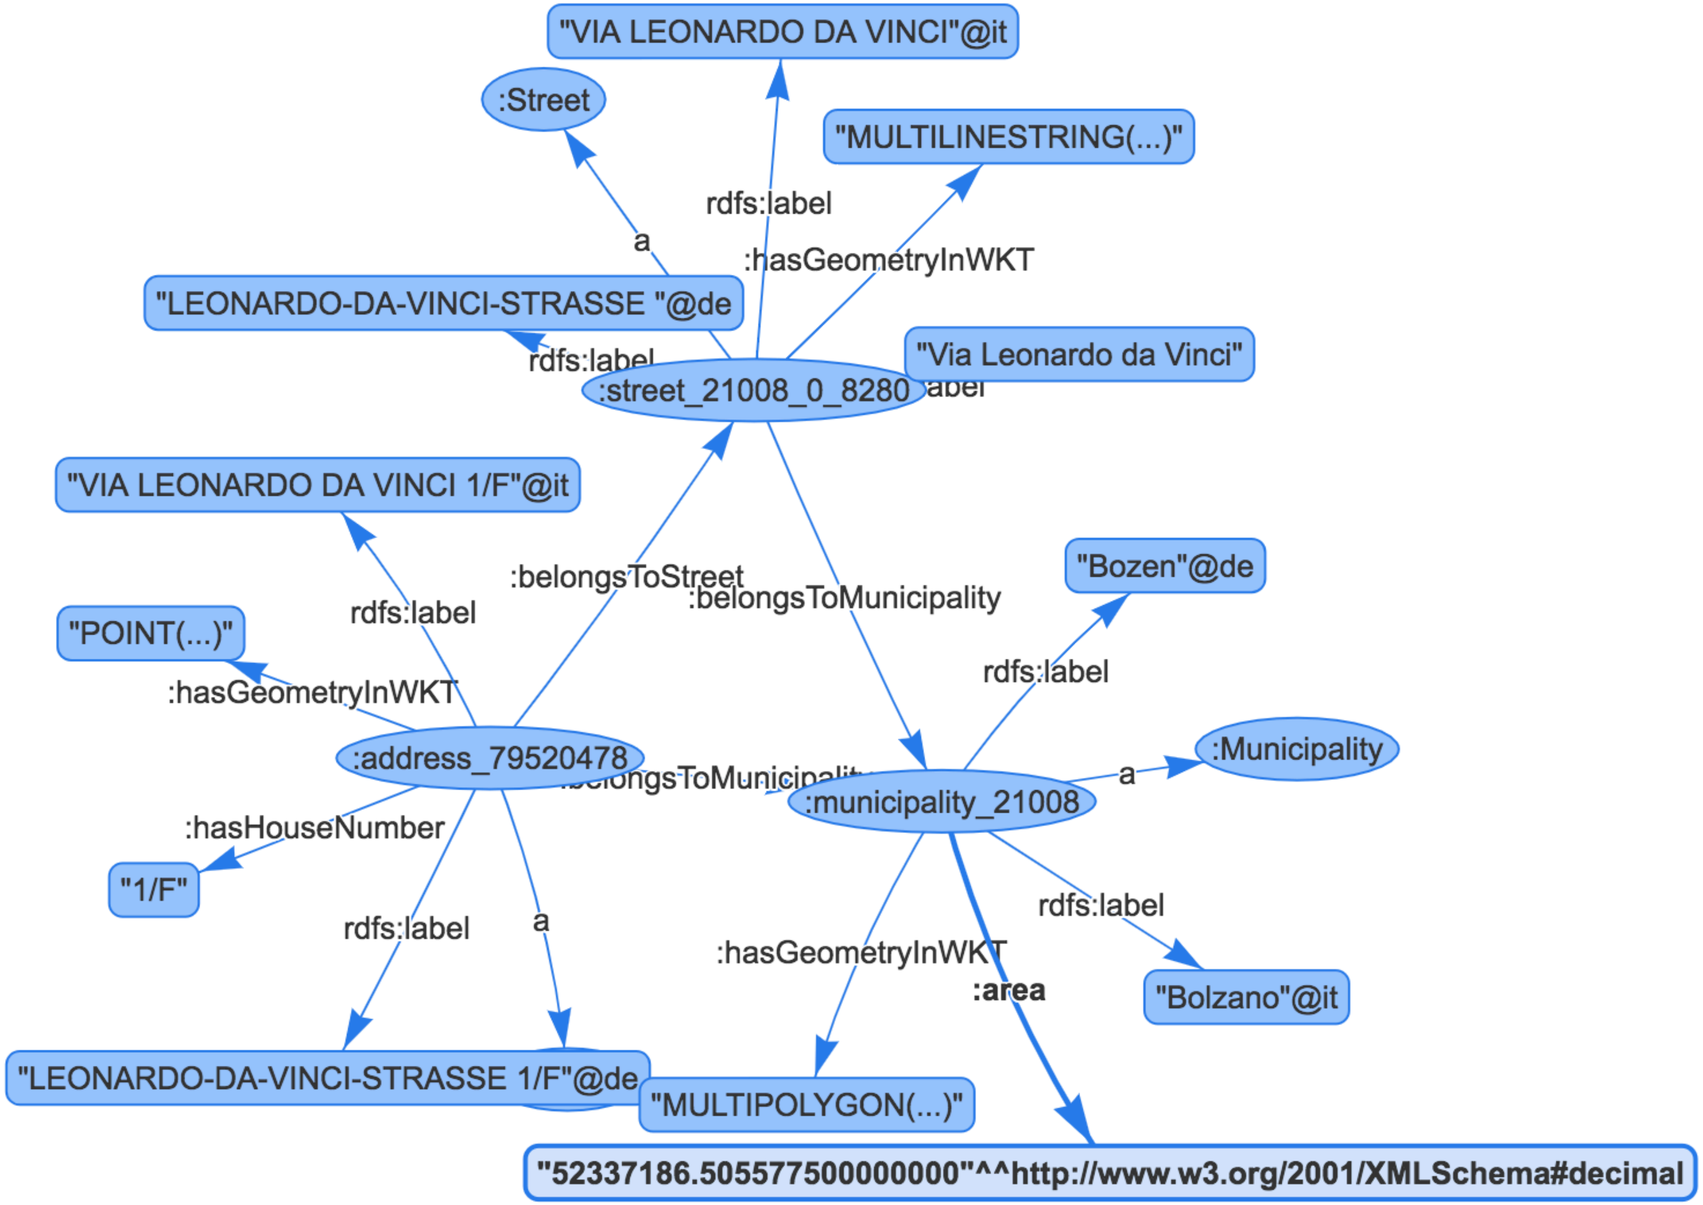
"VIA LEONARDO DA VINCI"@it
:Street
"MULTILINESTRING(...)"
rdfs:label
a
:hasGeometryInWKT
"LEONARDO-DA-VINCI-STRASSE "@de
rdfs:label
"Via Leonardo da Vinci"
:street_21008_0_8280
"VIA LEONARDO DA VINCI 1/F"@it
:belongsToStreet
:belongsToMunicipality
"Bozen"@de
rdfs:label
"POINT(...)"
rdfs:label
:hasGeometryInWKT
:address_79520478
:municipality_21008
a
:Municipality
:hasHouseNumber
"1/F"
rdfs:label
a
:hasGeometryInWKT
:area
rdfs:label
"Bolzano"@it
"LEONARDO-DA-VINCI-STRASSE 1/F"@de
"MULTIPOLYGON(...)"
"52337186.505577500000000"^^http://www.w3.org/2001/XMLSchema#decimal

The SPARQL query language

- Is the standard query language for RDF data. [W3C Rec. 2008, 2013]
- Core query mechanism is based on graph matching:

```
SELECT ?a ?t
WHERE { ?a rdf:type Actor .
        ?a playsIn ?m .
        ?m rdf:type Movie .
        ?m title ?t .
      }
```

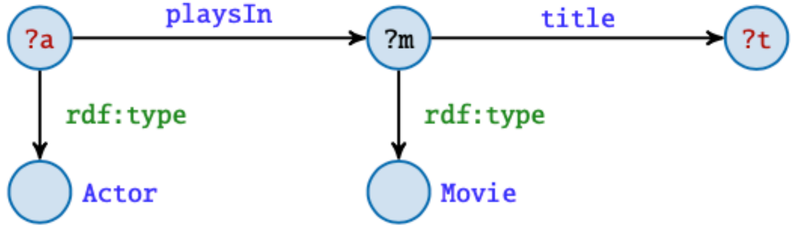

- Has many additional features:
  - **`UNION`**: matches one of alternative graph patterns
  - **`OPTIONAL`**: produces a match even when part of the pattern is missing
  - complex **`FILTER`** conditions
  - **`GROUP BY`**: to express aggregations
  - **`MINUS`**: to remove possible solutions
  - property paths (regular expressions)

What is an ontology?

- An ontology conceptualizes a domain of interest in terms of concepts/classes, (binary) relations, and their properties.
- It typically organizes the concepts in a hierarchical structure.

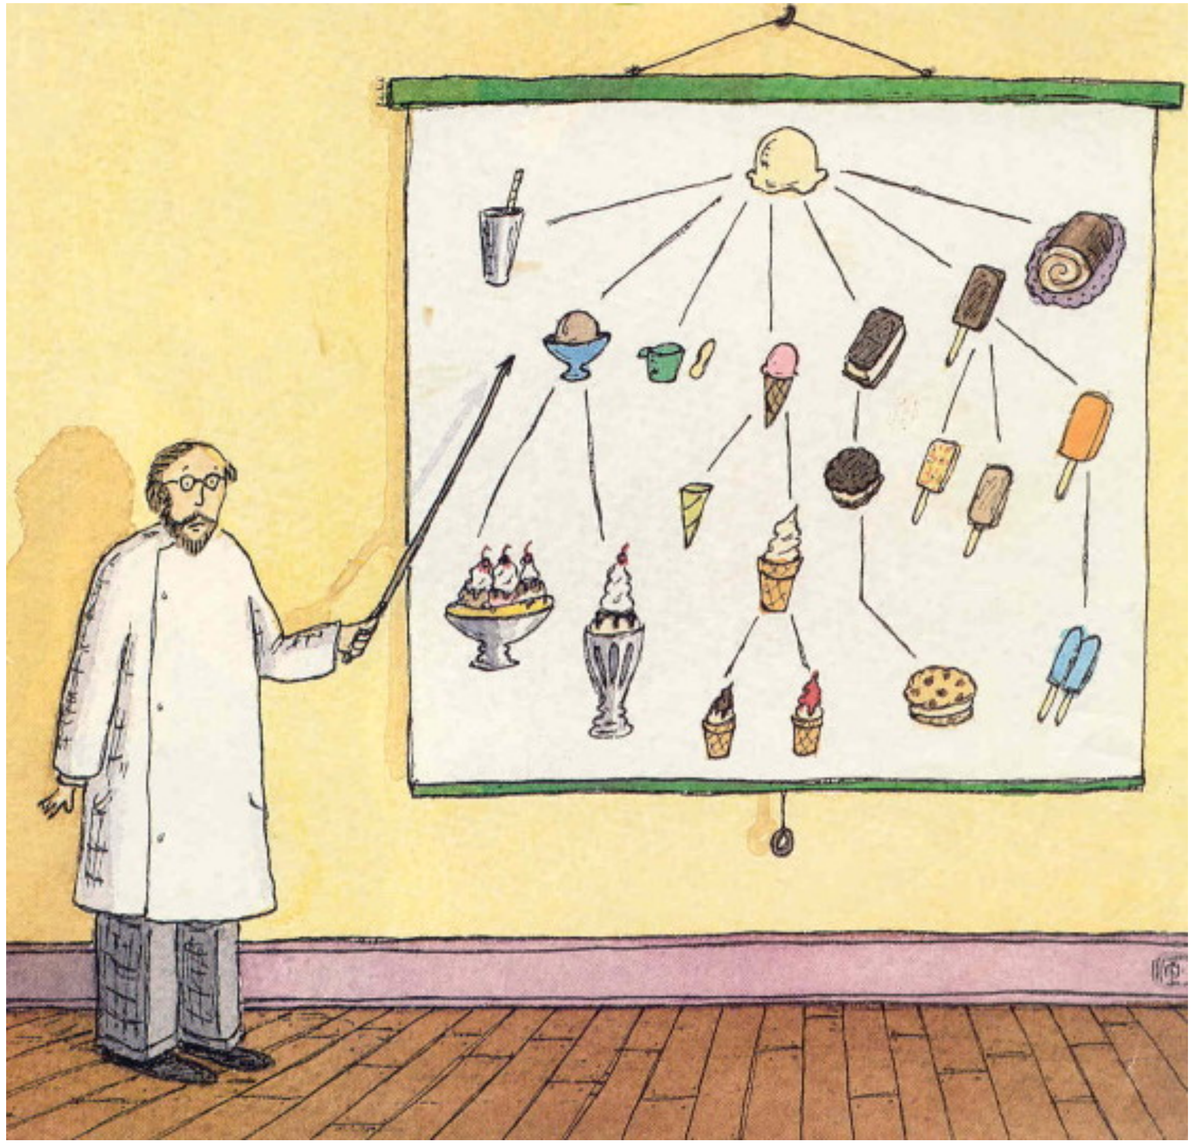

What is an ontology?

- An ontology conceptualizes a domain of interest in terms of concepts/classes, (binary) relations, and their properties.
- It typically organizes the concepts in a hierarchical structure.
- Ontologies are often represented as graphs. Sometimes they are also considered as Knowledge Graphs.

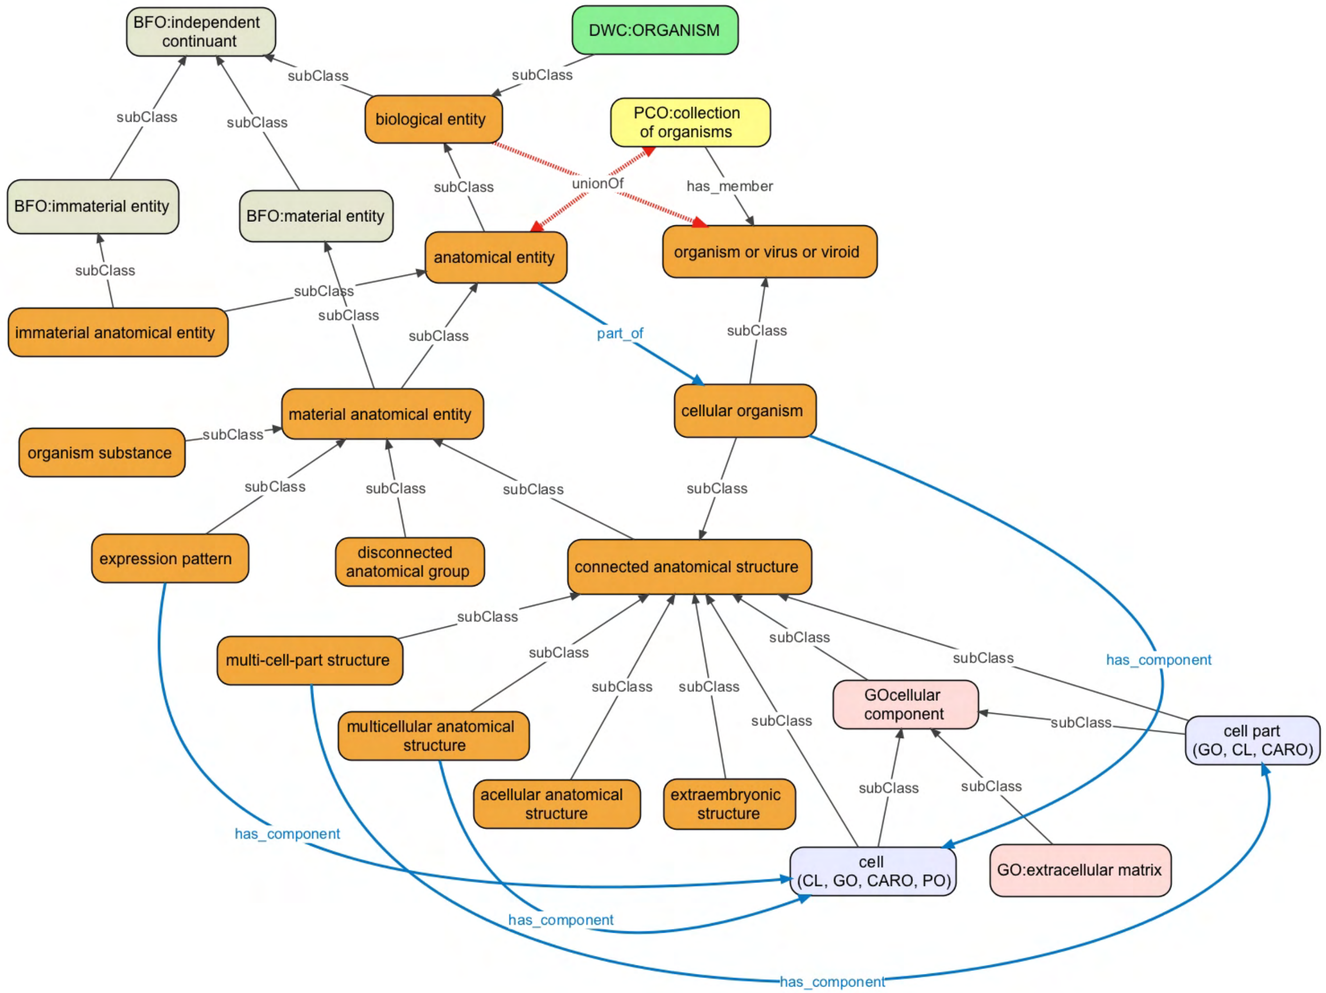

What is an ontology?

- An ontology conceptualizes a domain of interest in terms of concepts/classes, (binary) relations, and their properties.
- It typically organizes the concepts in a hierarchical structure.
- Ontologies are often represented as graphs. Sometimes they are also considered as Knowledge Graphs.
- However, an ontology is actually a logical theory, expressed in a suitable fragment of first-order logic.

$$\forall x.\, \text{Pressure}(x) \rightarrow \text{Measurement}(x)$$
$$\forall x.\, \text{Porosity}(x) \rightarrow \text{Measurement}(x)$$
$$\forall x.\, \text{Permeability}(x) \rightarrow \text{Measurement}(x)$$
$$\forall x.\, \text{Temperature}(x) \rightarrow \text{Measurement}(x)$$
$$\forall x.\, \text{Pressure}(x) \rightarrow \neg\text{Porosity}(x) \wedge \neg\text{Permeability}(x) \wedge \neg\text{Temperature}(x)$$
$$\forall x.\, \text{Porosity}(x) \rightarrow \neg\text{Permeability}(x) \wedge \neg\text{Temperature}(x)$$
$$\forall x.\, \text{Permeability}(x) \rightarrow \neg\text{Temperature}(x)$$

$$\forall x.\, \text{HydrostaticPressure}(x) \rightarrow \text{Pressure}(x)$$
$$\forall x.\, \text{FormationPressure}(x) \rightarrow \text{Pressure}(x)$$
$$\forall x.\, \text{PorePressure}(x) \rightarrow \text{Pressure}(x)$$
$$\forall x.\, \text{HydrostaticPressure}(x) \rightarrow \neg\text{FormationPressure}(x) \wedge \neg\text{PorePressure}(x)$$
$$\forall x.\, \text{FormationPressure}(x) \rightarrow \neg\text{PorePressure}(x)$$

$$\forall x, y.\, \text{hasFormationPressure}(x, y) \rightarrow \text{Wellbore}(x) \wedge \text{FormationPressure}(y)$$
$$\forall x, y.\, \text{hasDepth}(x, y) \rightarrow \text{FormationPressure}(x) \wedge \text{Depth}(y)$$
$$\forall x.\, \text{FormationPressure}(x) \rightarrow \exists y.\, \text{hasDepth}(x, y)$$

$$\forall x, y.\, \text{hasFormationPressure}(x, y) \rightarrow \text{hasMeasurement}(x, y)$$

$$\forall x, y.\, \text{completionDate}_{\text{Wellbore}}(x, y) \rightarrow \text{Wellbore}(x) \wedge \text{xsd:dateTime}(y)$$
$$\forall x.\, \text{Wellbore}(x) \rightarrow (\sharp\{y \mid \text{completionDate}_{\text{Wellbore}}(x, y)\} \leq 1)$$
$$\forall x, y.\, \text{wellboreTrack}_{\text{Wellbore}}(x, y) \rightarrow \text{Wellbore}(x) \wedge \text{xsd:string}(y)$$
$$\forall x.\, \text{Wellbore}(x) \rightarrow (\sharp\{y \mid \text{wellboreTrack}_{\text{Wellbore}}(x, y)\} \leq 1)$$

$$\forall x, y.\, \text{hasCoreSample}(x, y) \rightarrow \text{Core}(x) \wedge \text{CoreSample}(y)$$
$$\forall x.\, \text{CoreSample}(x) \rightarrow \exists y.\, \text{hasCoreSample}(y, x) \wedge \text{Core}(y)$$
$$\ldots$$

What is an ontology?

- An ontology conceptualizes a domain of interest in terms of concepts/classes, (binary) relations, and their properties.
- It typically organizes the concepts in a hierarchical structure.
- Ontologies are often represented as graphs.
- However, an ontology is actually a logical theory, expressed in a suitable fragment of first-order logic, or better, in description logics, the theoretical foundation of the OWL standard.

$$
\begin{aligned}
\text{Pressure} &\sqsubseteq \text{Measurement}\\
\text{Porosity} &\sqsubseteq \text{Measurement}\\
\text{Permeability} &\sqsubseteq \text{Measurement}\\
\text{Temperature} &\sqsubseteq \text{Measurement}\\
\text{Pressure} &\sqsubseteq \neg\text{Porosity} \sqcap \neg\text{Permeability} \sqcap \neg\text{Temperature}\\
\text{Porosity} &\sqsubseteq \neg\text{Permeability} \sqcap \neg\text{Temperature}\\
\text{Permeability} &\sqsubseteq \neg\text{Temperature}\\
\\
\text{HydrostaticPressure} &\sqsubseteq \text{Pressure}\\
\text{FormationPressure} &\sqsubseteq \text{Pressure}\\
\text{PorePressure} &\sqsubseteq \text{Pressure}\\
\text{HydrostaticPressure} &\sqsubseteq \neg\text{FormationPressure} \sqcap \neg\text{PorePressure}\\
\text{FormationPressure} &\sqsubseteq \neg\text{PorePressure}\\
\\
\exists\text{hasFormationPressure} &\sqsubseteq \text{Wellbore}\\
\exists\text{hasFormationPressure}^{-} &\sqsubseteq \text{FormationPressure}\\
\exists\text{hasDepth} &\sqsubseteq \text{FormationPressure}\\
\exists\text{hasDepth}^{-} &\sqsubseteq \text{Depth}\\
\text{FormationPressure} &\sqsubseteq \exists\text{hasDepth}\\
\\
\text{hasFormationPressure} &\sqsubseteq \text{hasMeasurement}\\
\\
\exists\text{completionDate} &\sqsubseteq \text{Wellbore}\\
\exists\text{completionDate}^{-} &\sqsubseteq \text{xsd:dateTime}\\
\text{Wellbore} &\sqsubseteq (\leq 1\ \text{completionDate})\\
\exists\text{wellboreTrack} &\sqsubseteq \text{Wellbore}\\
&\ldots
\end{aligned}
$$

Reasoning in Ontology

Class hierarchy: rdfs:subClassOf
Example: :MovieActor rdfs:subClassOf :Actor .
Inference: <person/2> rdf:type :MovieActor . => <person/2> rdf:type :Actor .

Domain of properties: rdfs:domain
Example: :playsIn rdfs:domain :MovieActor .
Inference: <person/2> :playsIn <movie/3> . => <person/2> rdf:type :MovieActor .

Range of properties: rdfs:range
Example: :playsIn rdfs:range :Movie .
Inference: <person/2> :playsIn <movie/3> . => <movie/3> rdf:type :Movie .

Reasoning in Ontology

Class disjointness: owl:disjointWith
Example: :Actor owl:disjointWith :Movie .
Inference: <person/2> rdf:type :Actor . <person/2> rdf:type :Movie .
=> RDF graph inconsistent with the ontology

Inverse properties: owl:inverseOf
Example: :actsIn owl:inverseOf :hasActor .
Inference: <person/2> :actsIn <movie/3> . => <movie/3> :hasActor <person/2> .

Property hierarchy
Property disjointness
Mandatory participation
…

Mapping Language

The mapping consists of a set of assertions of the form

SQL Query ⇝ Class membership assertion

SQL Query ⇝ Property membership assertion

Intuition behind the mapping

The answers returned by the SQL Query in the left-hand side are used to create the objects (and values) that populate the Class / Property in the right-hand side.

Note: The mapping contains also a mechanism to transform values retrieved from the database into objects of the VKG (thus solving the so-called impedance mismatch).

Why mappings?

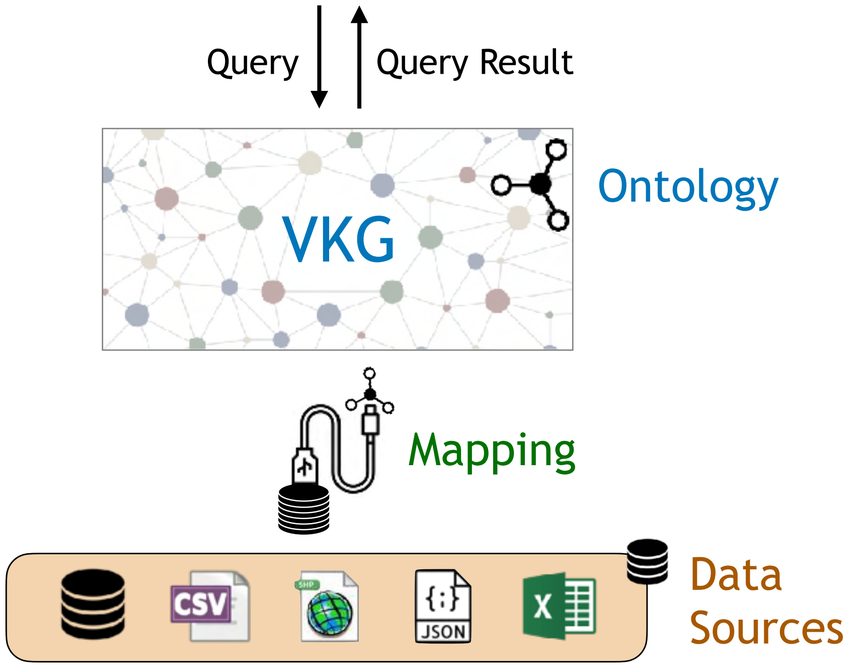

The traditional approach to data integration relies on mediators, which are specified through complex code.

Mappings, instead:

- Provide a declarative specification, and not code.
- Are easier to understand, and hence to design and to maintain.
- Support an incremental approach to integration.
- Are machine processable, hence can be used for query optimization.

5. A Demo for Data Integration using VKGs

Demo: Using Ontop for data integration

```
$ git clone \
https://github.com/ontopic-vkg/destination-tutorial \
--config core.autocrlf=input # important for Windows
$ cd destination-tutorial
$ docker-compose -f docker-compose.solution.yml up
```

1. Check the database in DBeaver
2. Open vkg/dest-solution.ttl in Protégé
3. Open http://localhost:8080 in the browser

Demo Database

Three data sources:

- Sources 1 & 2: hotels and rooms in different areas
- Source 3: weather information

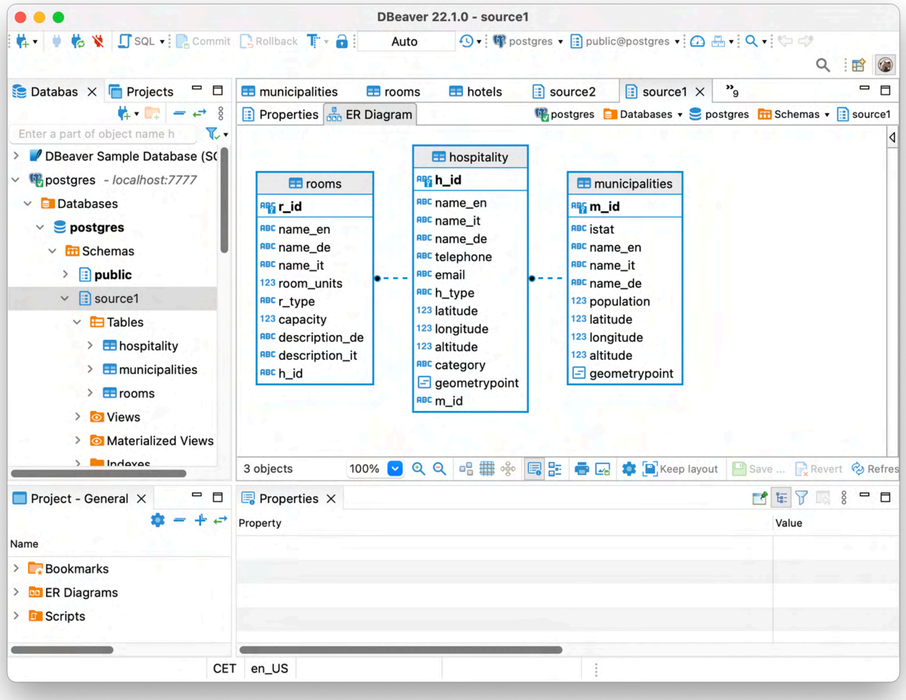

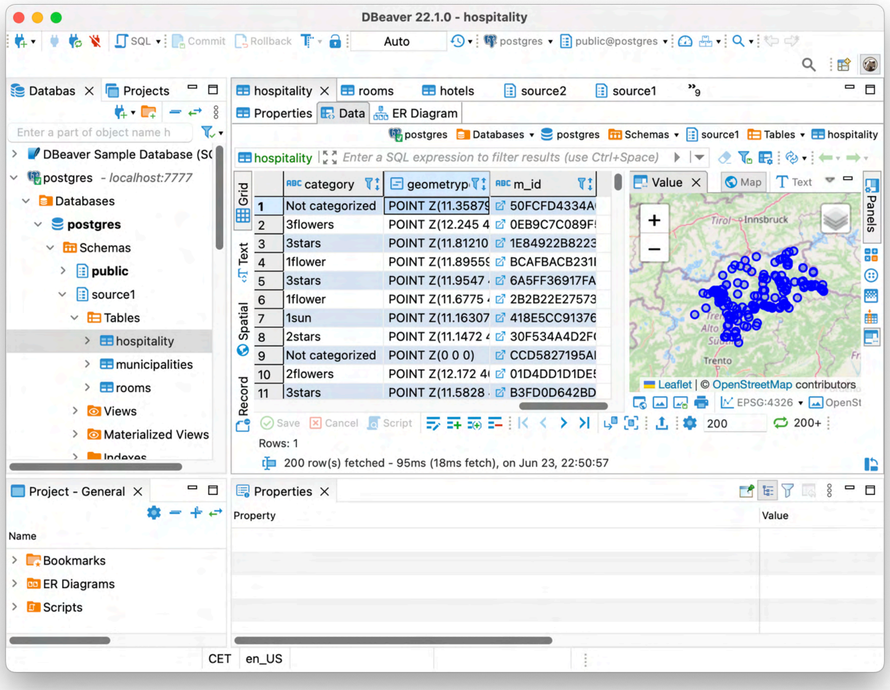

Ontology for tables in the schemas source1 and source2

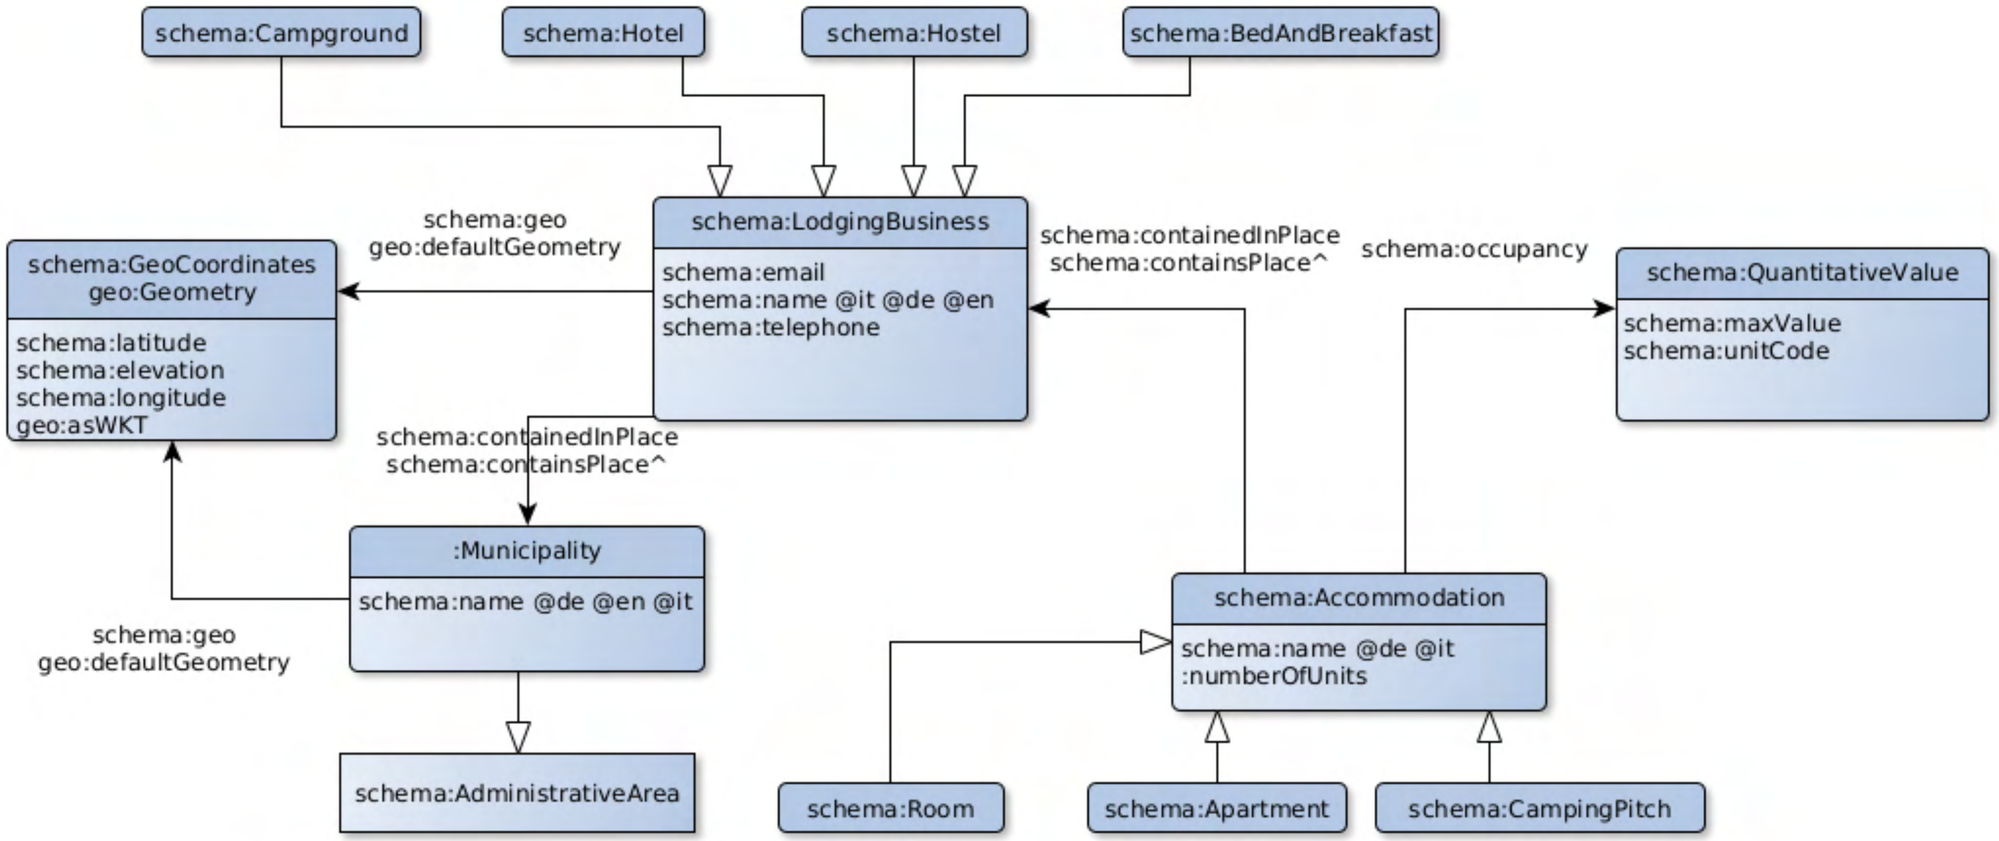

Ontology for tables in the schema source3

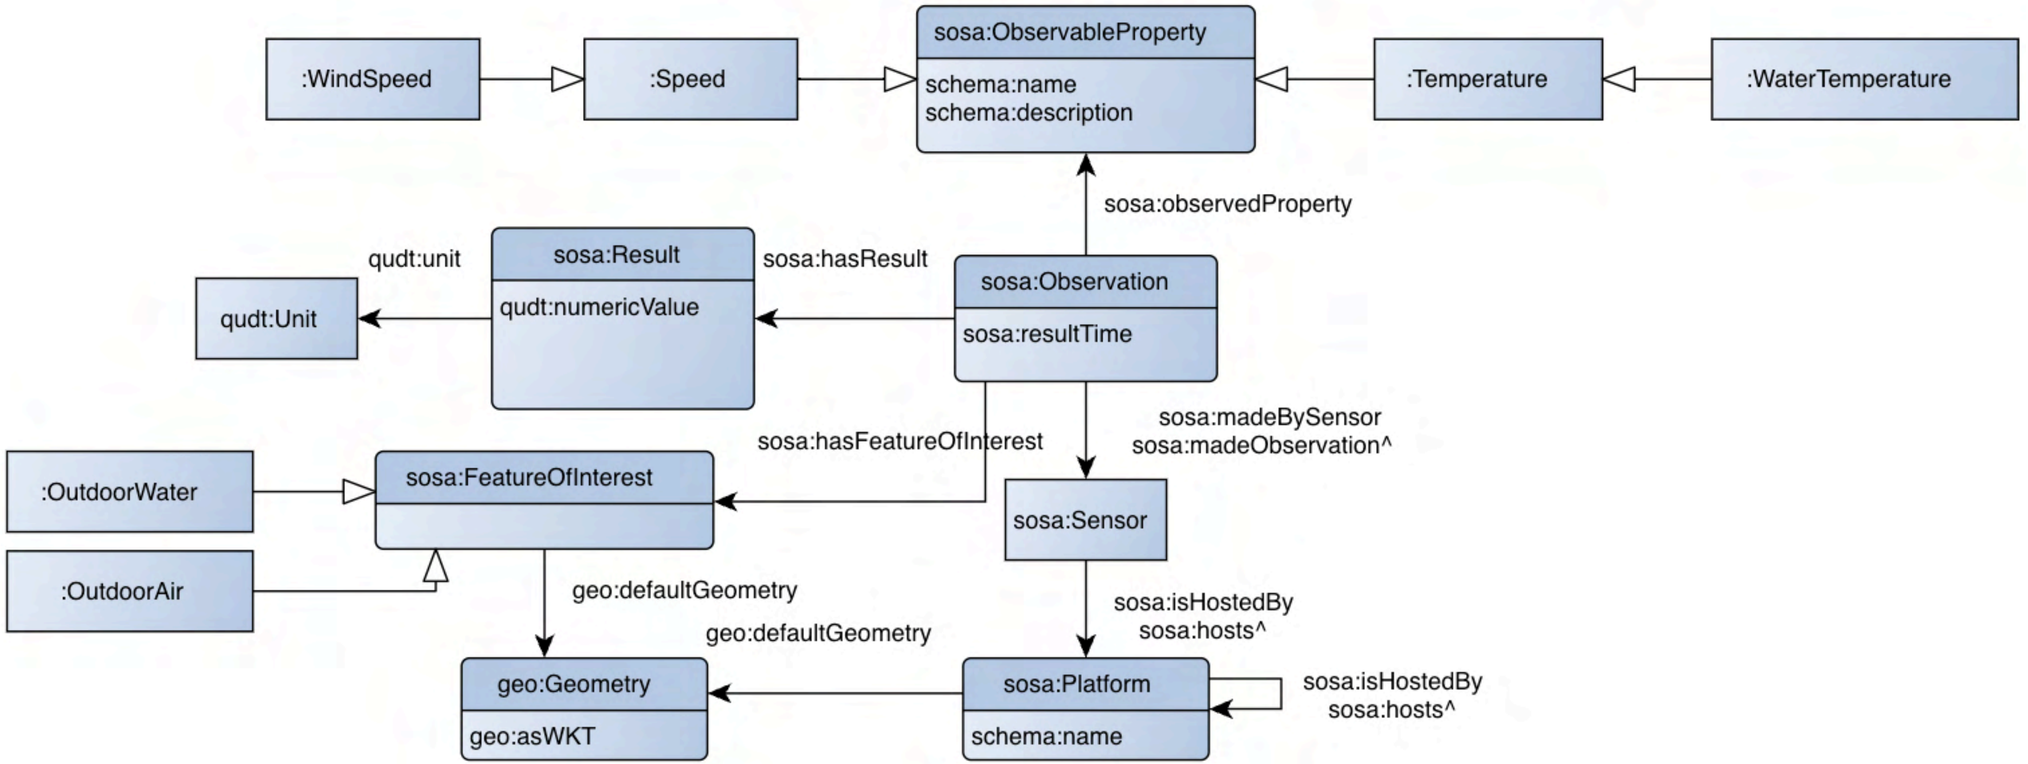

Protege Editor with Ontop Plugin

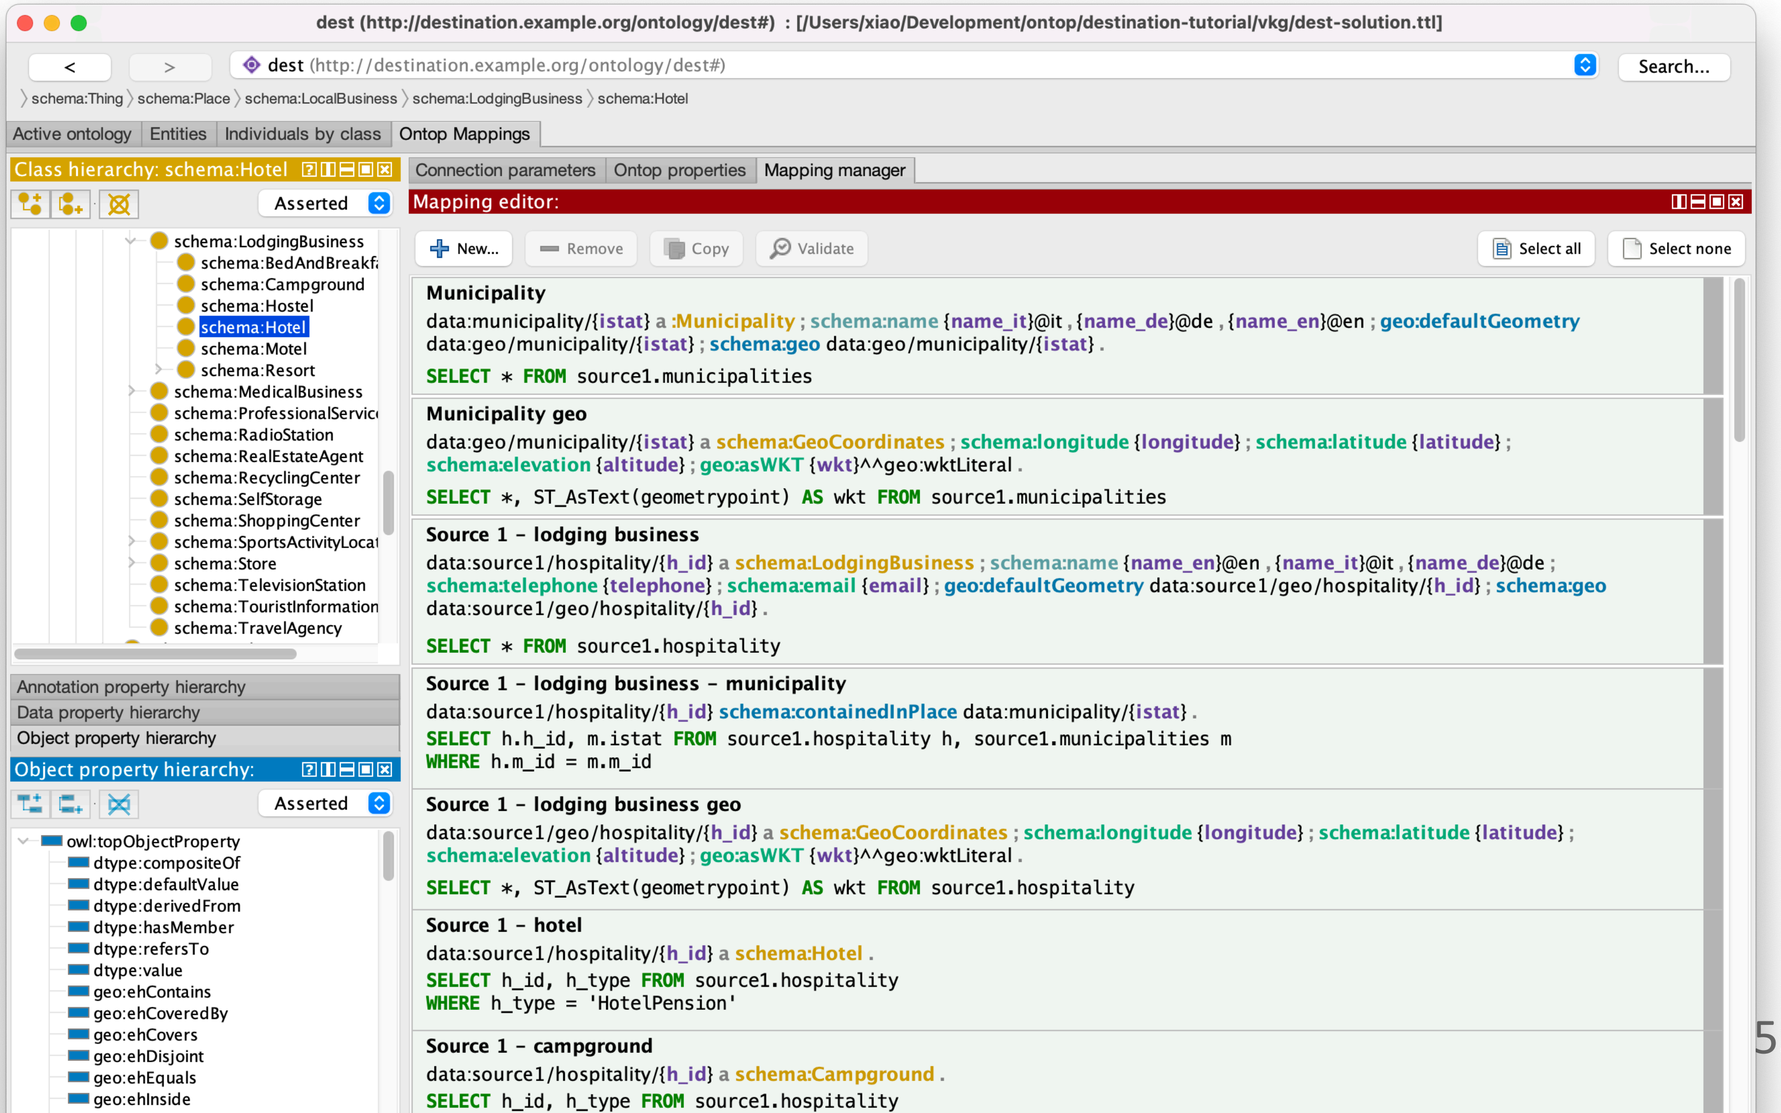

Ontopic Studio

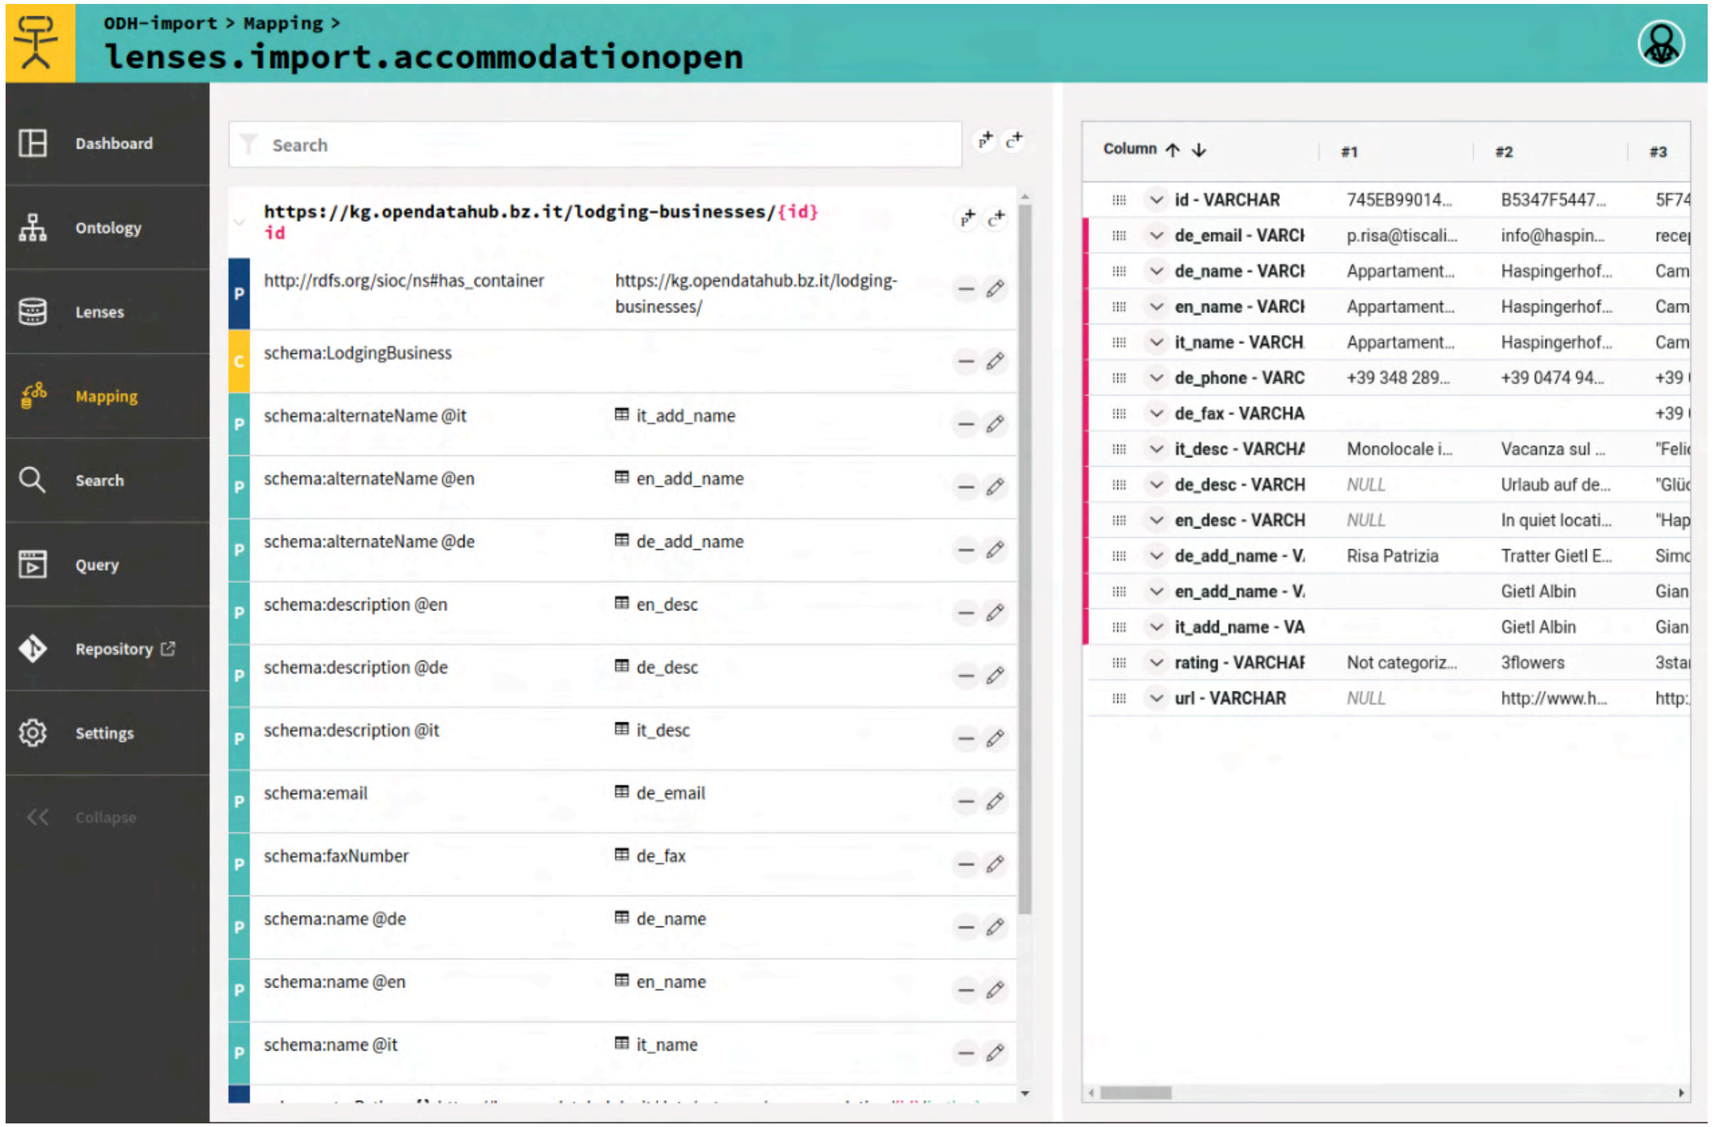

SPARQL Query Answering

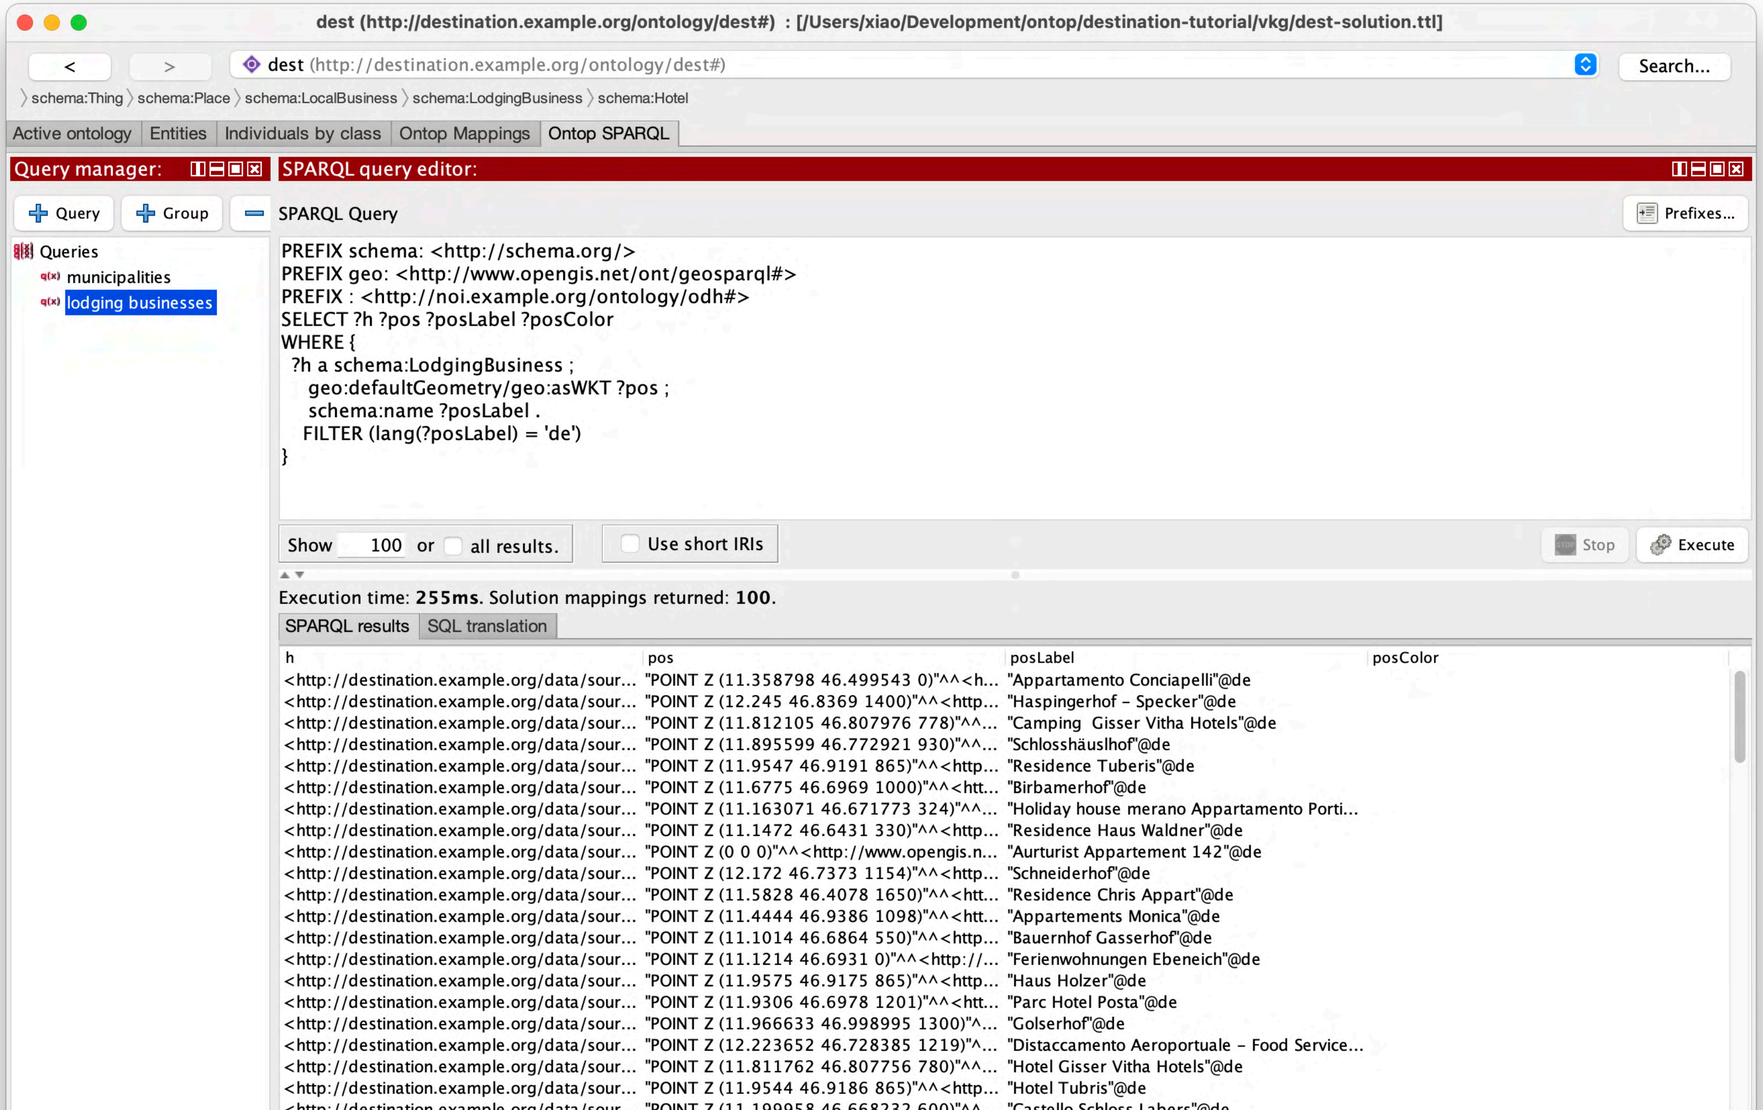

SPARQL Query Answering

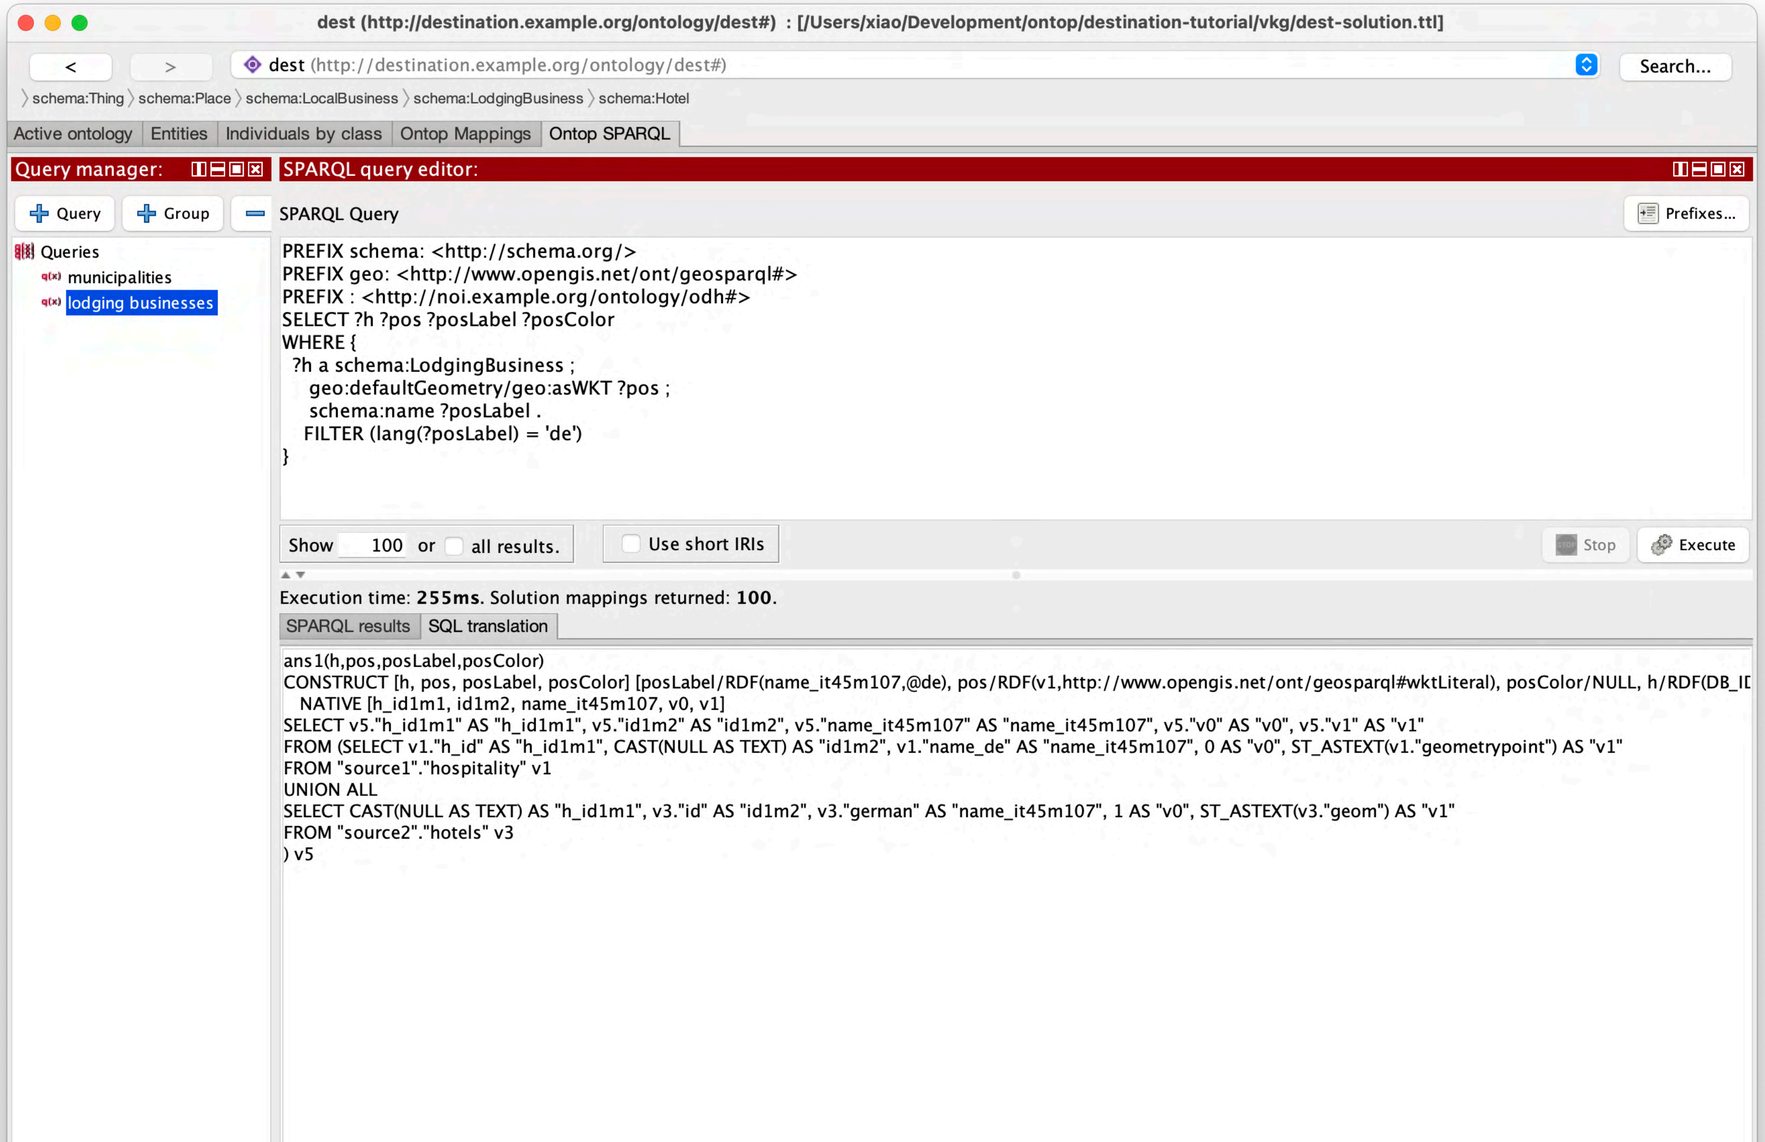

SPARQL Endpoint

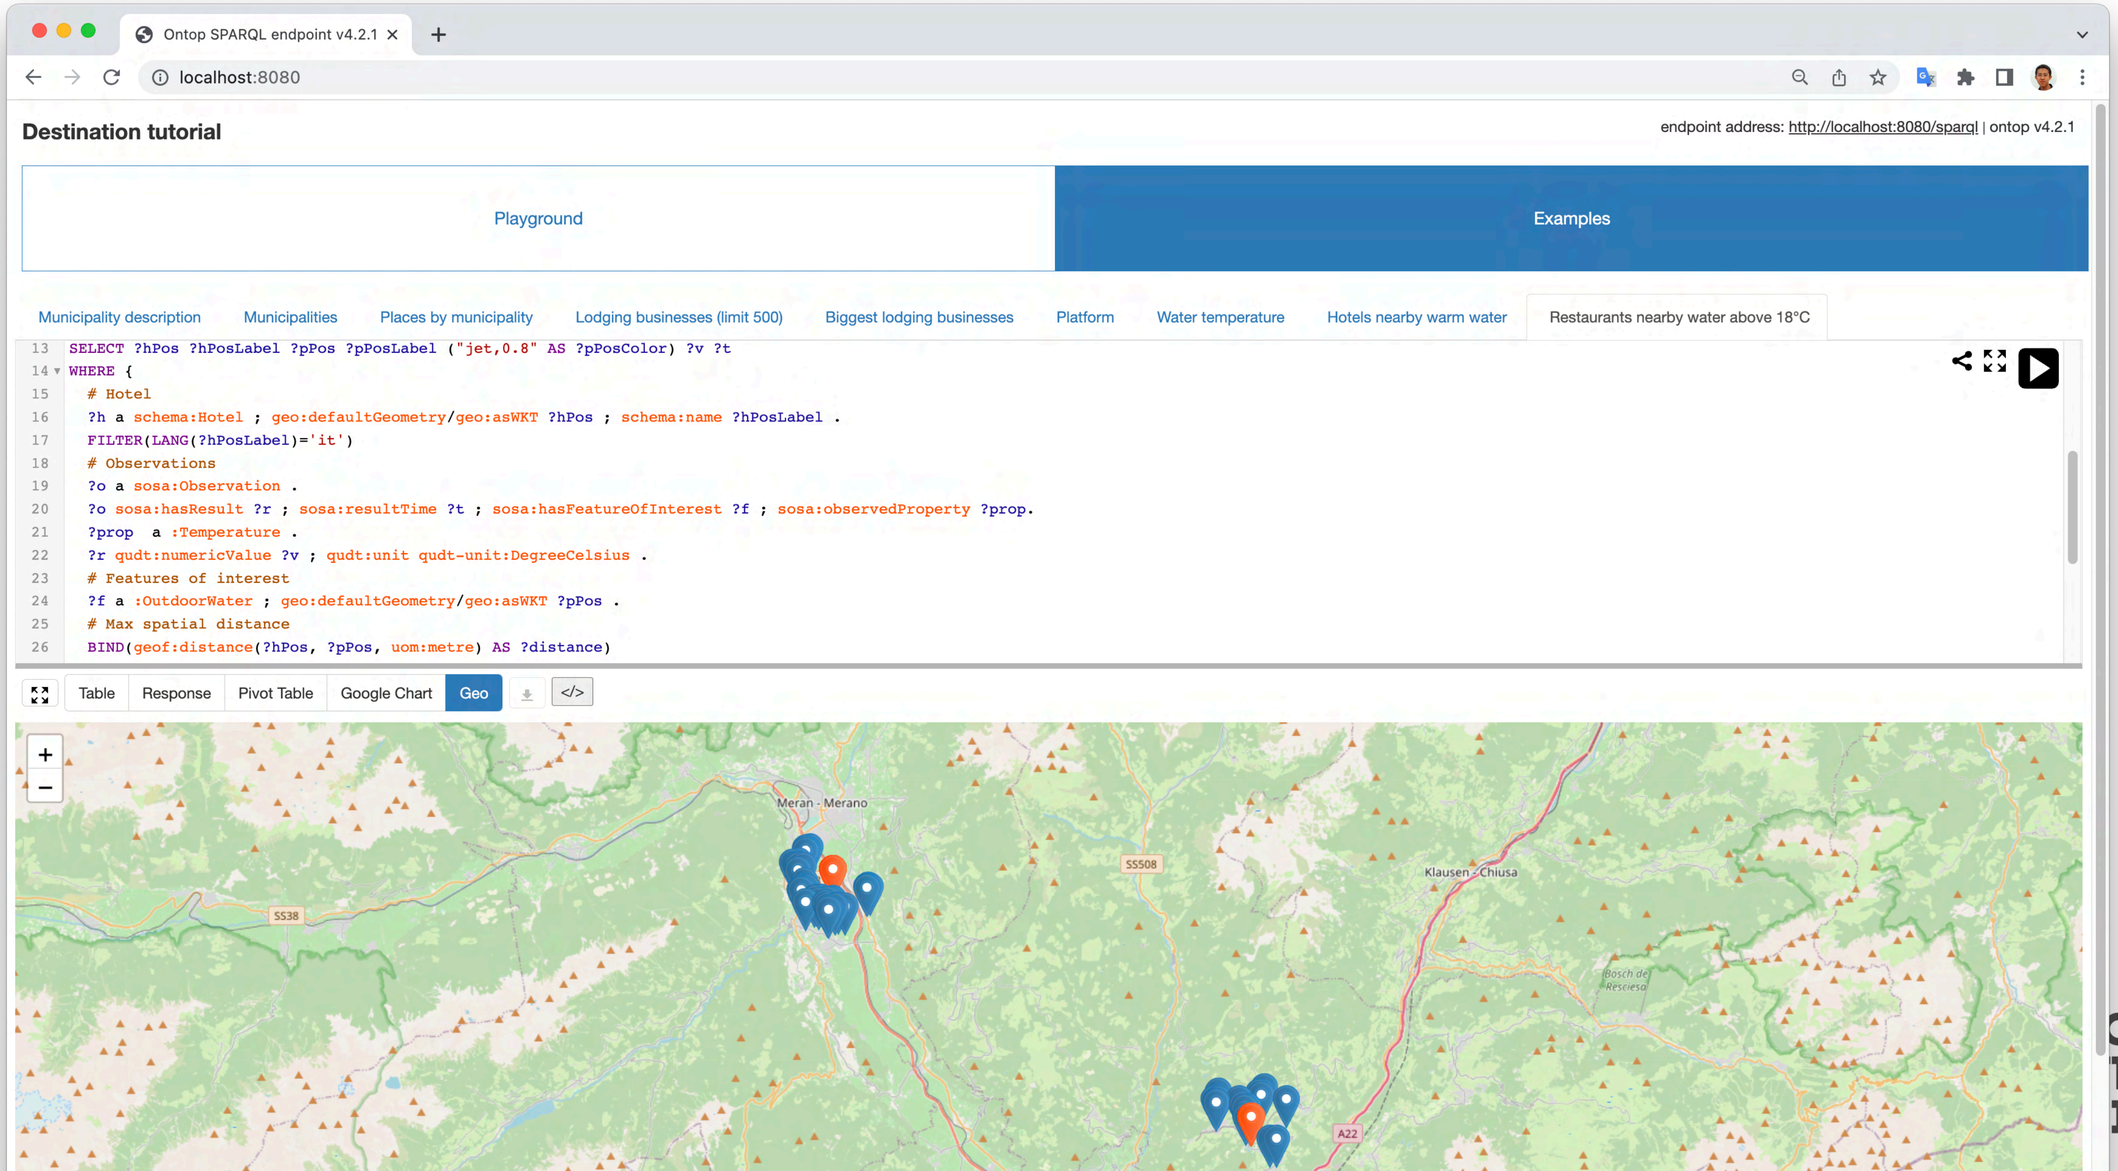

Take-home messages

- KG is a very simple but powerful data structure
- KGs give semantics to the data: things, not strings
- KGs have been applied in many applications
- VKGs are by now a mature technology to address the challenges related to data access and integration.
- It has been well-investigated and applied in many different scenarios mostly for the case of relational data sources.

Further readings

1. Tim Berners-Lee. **Information management: A proposal**. Technical report, 1989.
2. Tim Berners-Lee, James Hendler, and Ora Lassila. **The semantic web: A new form of web content that is meaningful to computers will unleash a revolution of new possibilities**. Scientific American, 2001.
3. A. Hogan, et al. **Knowledge graphs**. ACM Computer Survey, 54(4):71:1–71:37, 2021.
4. G. Xiao, et al. **Ontology-based data access: A survey**. IJCAI, 2018.
5. G. Xiao, Linfang Ding, Benjamin Cogrel, and Diego Calvanese. **Virtual knowledge graphs: An overview of systems and use cases**. Data Intelligence, 1:201–223, 2019.

Thank you for your attention

Questions?

Guohui Xiao – guohui.xiao@ontopic.ai