

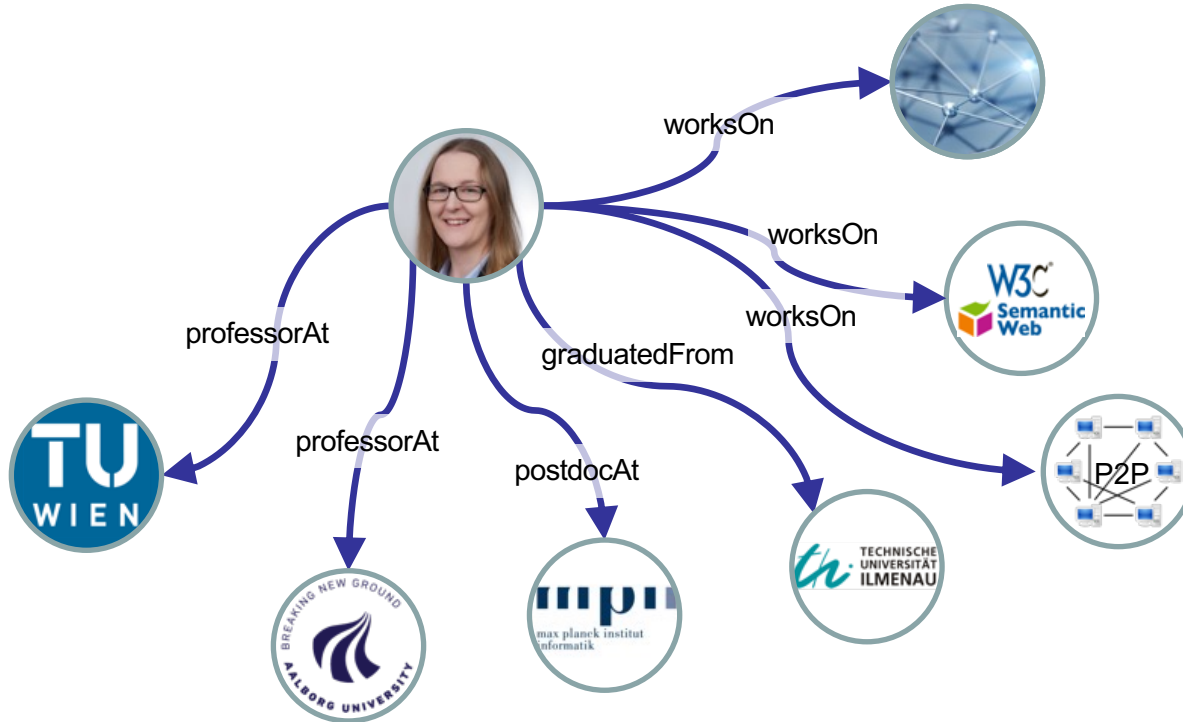
188.399 Introduction to Semantic Systems

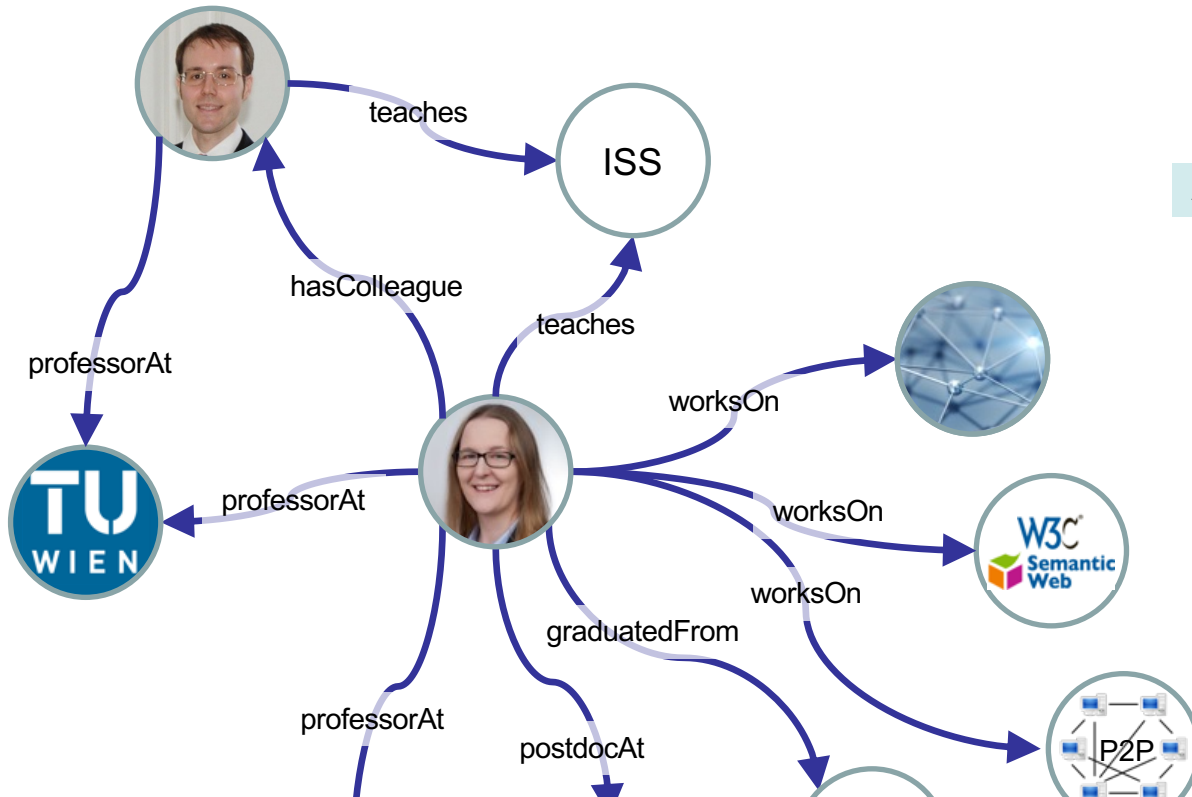
(2.0 VU / 3.0 ECTS)

2023W

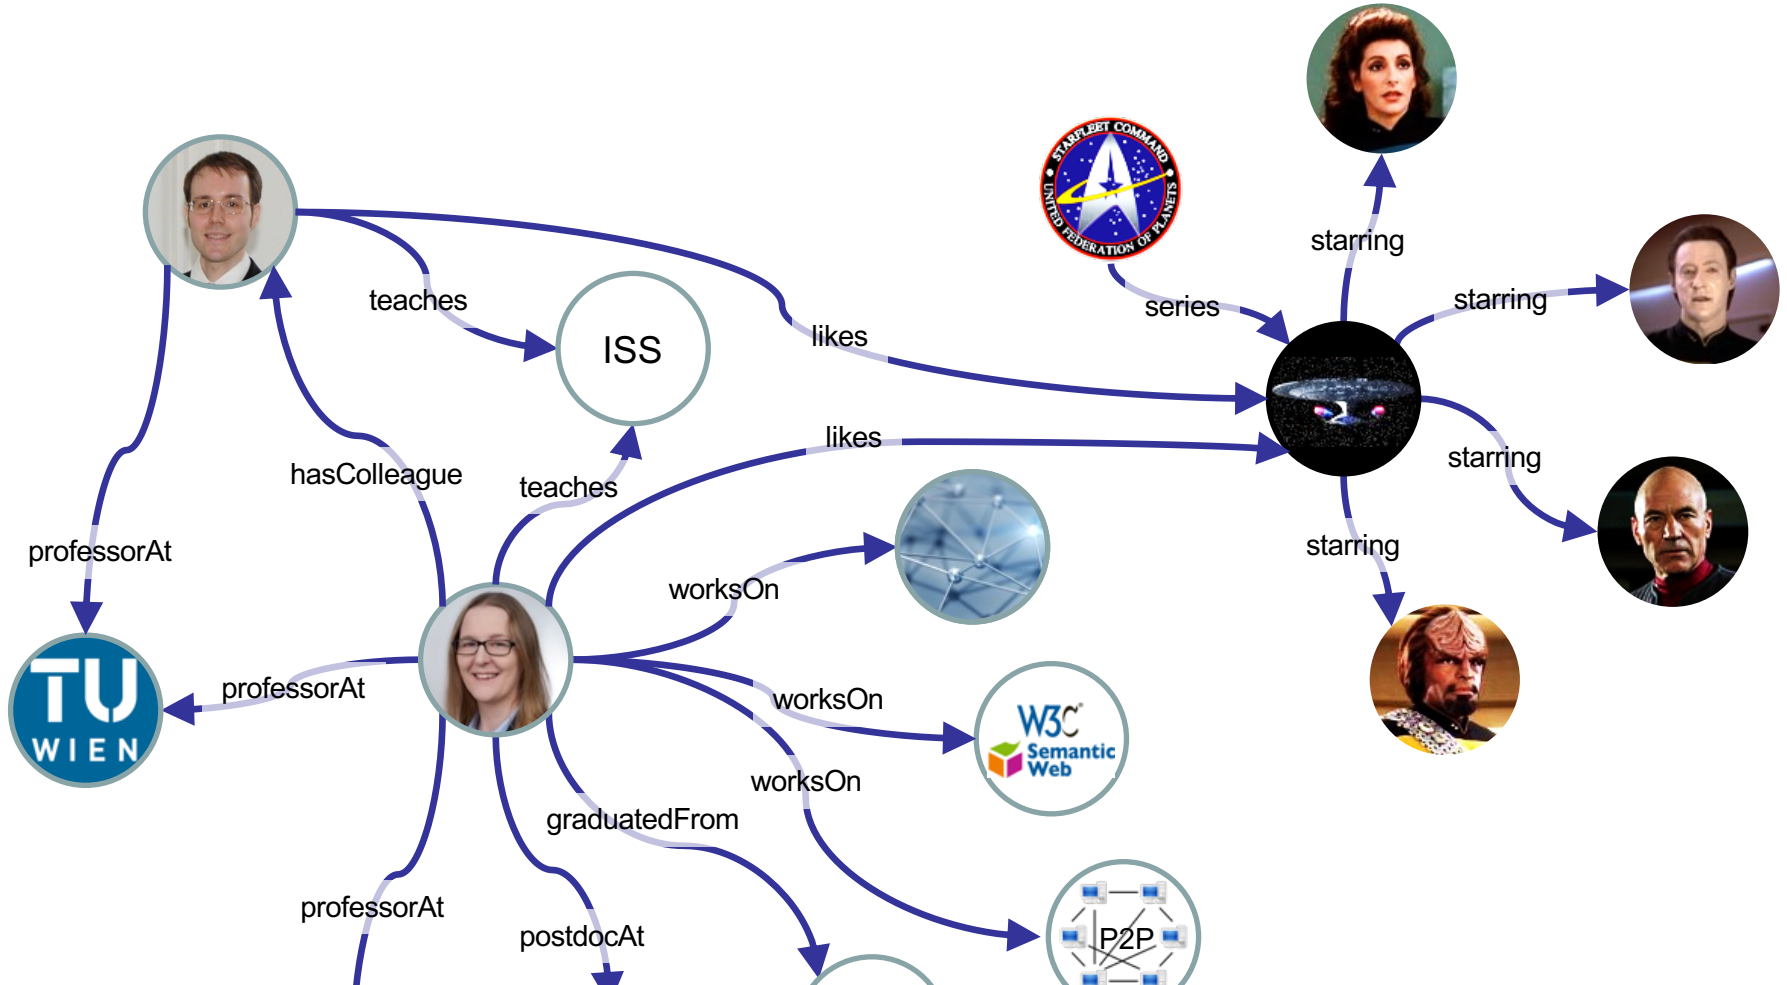
**Katja Hose
Emanuel Sallinger**

iss@dbai.tuwien.ac.at

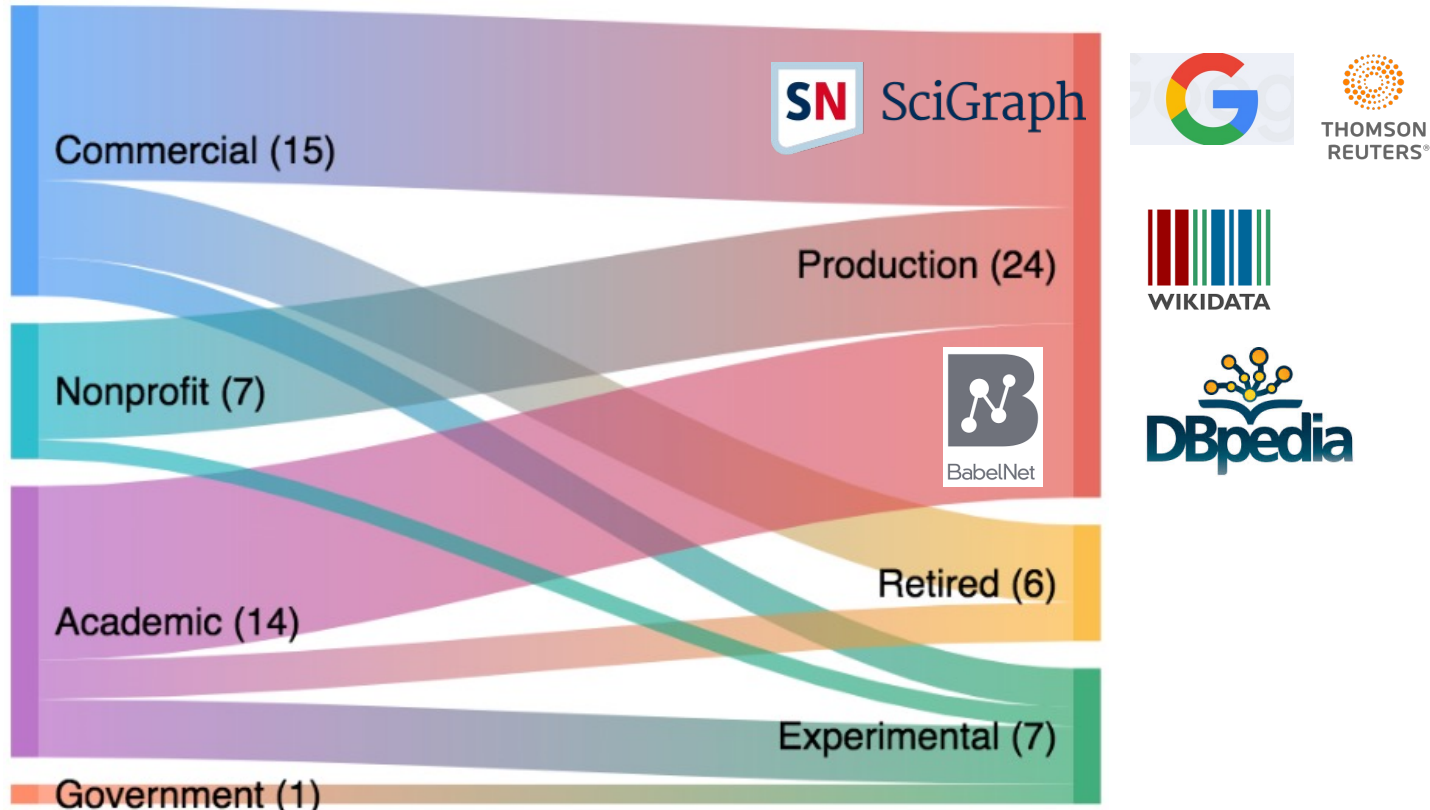




A knowledge graph is a directed labeled graph



Semantic systems: systems that make use of explicitly represented knowledge, through conceptual structures such as ontologies, taxonomies or **knowledge graphs**.



Knowledge graph proportions by publisher and status

What are knowledge graphs (KGs)?

- Definitions
- Examples

How to use semantic technologies and KGs?

- Ontology modeling
- Reasoning
- Representation languages
- KG management
- KG querying

What can I do with KGs?

- Use cases and applications

Lecture level: For Beginners (no prior knowledge required)

Knowledge graphs and semantic technologies with, without, and for AI



- Expressing meaning
- Knowledge representation
- Ontologies
- Agents
- Evolution of knowledge

Killer app: Agents

- Retrieve information about prescribed treatment/therapy
- Look for providers/specialists with good ratings in proximity
- Schedule times according to multiple busy schedules
- Share information with other agents and “negotiate”

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](#)

The entertainment system was belting out the Beatles' "We Can Work It Out" when the phone rang. When Pete answered, his phone turned the sound down by sending a message to all the other *local* devices that had a *volume control*. His sister, Lucy, was on the line from the doctor's office: "Mom needs to see a specialist and then has to have a series of physical therapy sessions. Biweekly or something. I'm going to have my agent set up the appointments." Pete immediately agreed to share the chauffeuring. At the



doctor's office, Lucy instructed her Semantic Web agent through her handheld Web browser. The agent promptly retrieved information about Mom's *prescribed treatment* from the doctor's agent, looked up several lists of *providers*, and checked for the ones *in-plan* for Mom's insurance within a *20-mile radius* of her home and with a *rating of excellent or very good* on trusted rating services. It then began trying to find a match between available *appointment times* (supplied by the agents of individual providers through their Web sites) and Pete's and Lucy's busy schedules. (The emphasized keywords indicate terms whose semantics, or meaning, were defined for the agent through the Semantic Web.)

In a few minutes the agent presented them with a plan. Pete didn't like it—University Hospital was all the way across town from Mom's place, and he'd be driving back in the middle of rush hour. He set his own agent to redo the search with stricter preferences about *location* and *time*. Lucy's agent, having *complete trust* in Pete's agent in the context of the present task, automatically assisted by supplying access certificates and shortcuts to the data it had already sorted through.

“A really important thing about data is:
the more things you have to connect
together, the more powerful it is.”



Google ✕ 🔍 Sign in

[🔍 All](#) [🖼️ Images](#) [📺 Videos](#) [📰 News](#) [📍 Maps](#) [⋮ More](#)

[Settings](#) [Tools](#)

About 463,000,000 results (0,73 seconds)

📰 Top stories

Stay, go, delay: Joe Biden is trapped and has no good choices in Afghanistan

The Guardian · 14 mins ago

Joe Biden won the presidency by making the most of his lucky breaks

The Washington Post · 6 hours ago

[→](#) [View all](#)

en.wikipedia.org › wiki › Joe_Biden

Joe Biden - Wikipedia

Joseph Robinette Biden Jr. (born November 20, 1942) is an / 46th president of the United States. A member of the Democr

Vice President: [Kamala Harris](#) Parents: Jo
President: [Barack Obama](#) Occupation

[Family of Joe Biden](#) · [Presidency of Joe Biden](#) · [Timeline of th](#)

www.whitehouse.gov › administration › president-biden

Joe Biden: The President | The White Hous

Joe Biden · [Joe Biden's Family](#) · [A Leader in the Senate](#) · [TI United States](#) · [A New Chapter.](#)

Microsoft Bing 🗨️ 🔍 🔄 🔍

ALL NEWS IMAGES VIDEOS

51,700,000 Results [Open links in new tab](#)



Joe Biden

[Quotes](#)

[Speeches](#)

[Career](#)

[Early life](#)

News about Joe Biden

[bing.com/news](#)



Donald Trump at CPAC: Ex-president expected to tear into ...

Trump is not expected to declare a 2024 presidential candidacy - he is likely to discuss ...

MSN · 21m



Joe Biden Is Covertly Defunding the Police by Matt Vespa

Townhall · 22h



Joe Biden's COVID relief bill with stimulus checks passes ...

Des Moines Register · 2d



Trump's unconditional retreat from Afghanistan leaves ...

USA Today · 2d

[See more news about joe biden](#)



[See all images](#)

Joe Biden

Joseph Robinette Biden Jr. is an American politician serving as the 46th president of the United States. A member of the Democratic Party, he served as the 47th vice president from 2009 to 2017 under Barack Obama and represented Delaware in the United States Senate from 1973 to 2009.



[Wikipedia](#)



[Facebook](#)



[Instagram](#)



[Official site](#)



[IMDb](#)

Born: 11/20/1942 (age 78) · Scranton, PA

Height: 1.83 m

Inauguration date: **January 20, 2021**

Transition plan: **Build Back Better**

December 9, 2021: Who is the chancellor of Germany?

Google

who is the chancellor of germany

Alle Bilder News Maps Videos Mehr


Suchfilter

Ungefähr 84.900.000 Ergebnisse (0,72 Sekunden)

Deutschland / Kanzler

Olaf Scholz

Seit 2021



Olaf Scholz
Bundeskanzler

Olaf Scholz ist ein deutscher Politiker und seit Dezember 2021 neunter Bundeskanzler der Bundesrepublik Deutschland. Zuvor war er von 2018 bis 2021 Stellvertreter der Bundesminister der Finanzen der Bundesrepublik Deutschland, vom 26. [Wikipedia](#)

Geboren: 14. Juni 1958 (Alter 63 Jahre)
Deutschland

Größe: 1,7 m

Ehepartnerin: Britta Ernst (verh. 1998)

Partei: Sozialdemokratische Partei Deutschlands

Ausbildung: Universität Hamburg (1979)

Wird auch oft gesucht

- Angela Merkel
- Armin Laschet
- Britta Ernst
- Annalena Baerbock
- Christoph Lindner
- Markus Söder
- Robert Habeck

Feedback geben

Microsoft Bing


who is the chancellor of germany

ALLE NEWS BILDER VIDEOS KARTEN SHOPPING

239.000 Ergebnisse Datum

Angela Merkel

Gemäß 4 Quellen



Much better known at home and abroad is **Angela Merkel**, the current Chancellor of Germany. She is serving her fourth term. Mrs. Merkel was elected in 2005 and will retire from politics at the end of her term in 2021. So, why is she better known? What are the roles of the German Chancellor versus the German President?

Published: 29 September,2021: 04:13 PM GST Updated: 29 September,2021: 04:43 PM GST Germany's outgoing Chancellor **Angela Merkel** has congratulated Finance Minister Olaf Scholz from the rival Social Democrats on his election victory, a government statement said Wednesday, as Merkel's conservatives remain in disarray after the vote.

[Role of German President vs. German Chancellor - W... walled-in-berlin.com](#)

[Merkel congratulates Scholz on election as Germany's ne... english.alarabiya.net](#)

Mehr anzeigen

Feedback

“

Knowledge graphs provide the structured data and factual knowledge that drive many products and make them more intelligent and "magical".

”

N. Noy, Y. Gao, A. Jain, A. Narayanan, A. Patterson, J. Taylor:
 Industry-scale knowledge graphs: lessons and challenges.
 Commun. ACM 62(8): 36-43 (2019)

Common characteristics of the knowledge graphs.


	Data model	Size of the graph	Development stage
Microsoft	The types of entities, relations, and attributes in the graph are defined in an ontology.	~2 billion primary entities, ~55 billion facts	Actively used in products
Google	Strongly typed entities, relations with domain and range inference	1 billion entities, 70 billion assertions	Actively used in products
Facebook	All of the attributes and relations are structured and strongly typed, and optionally indexed to enable efficient retrieval, search, and traversal.	~50 million primary entities, ~500 million assertions	Actively used in products
eBay	Entities and relation, well-structured and strongly typed	Expect around 100 million products, >1 billion triples	Early stages of development and deployment
IBM	Entities and relations with evidence information associated with them.	Various sizes. Proven on scales documents >100 million, relationships >5 billion, entities >100 million	Actively used in products and by clients




WIKIDATA
1.9B Facts


yago
select knowledge
210M Facts


DBpedia
52M Facts


linked life data
<http://linkedlifedata.com/sources.html>
6.7B Facts

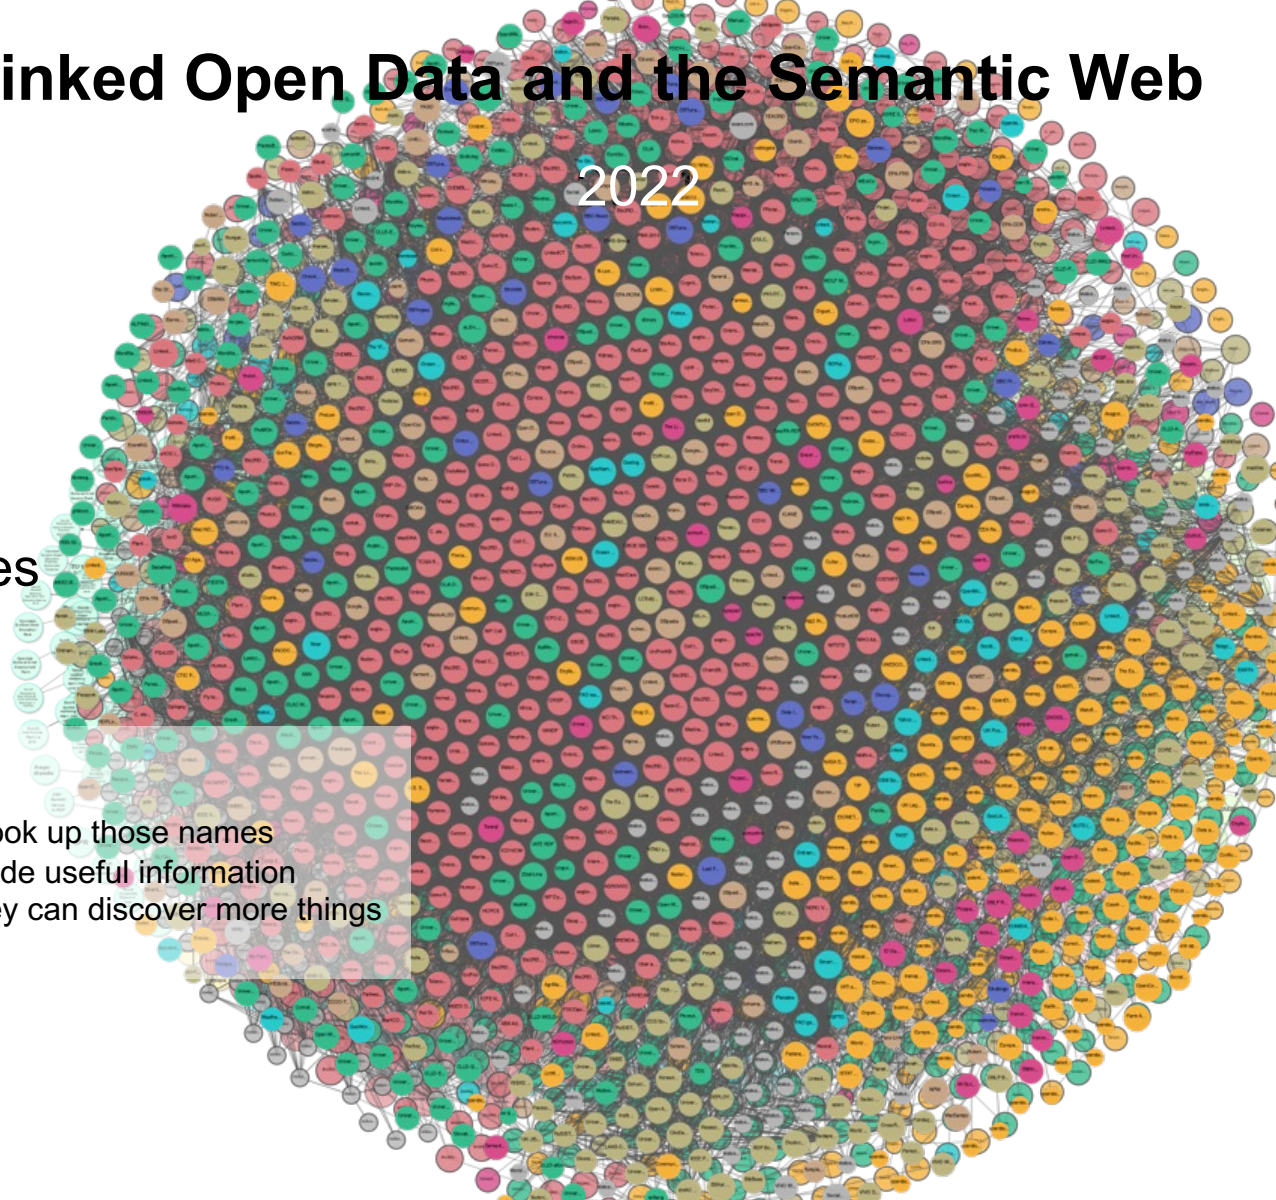
PubChem
<https://pubchem.ncbi.nlm.nih.gov/docs/rdf>
132B Facts

2022

- Constantly growing
- Ca. 1255 datasets
- Ca. 400 billion triples


Linked Data principles

1. Use URIs as names for things
2. Use HTTP URIs so that people can look up those names
3. When someone looks up a URI, provide useful information
4. Include links to other URIs so that they can discover more things



He Created the Web. Now He's Out to Remake the Digital World.

Tim Berners-Lee wants to put people in control of their personal data. He has technology and a start-up pursuing that goal. Can he succeed?



Tim Berners Lee's startup Inrupt releases Solid privacy platform for enterprises

Ron Miller @ron_miller / 6:01 AM GMT+1 • November 9, 2020



The screenshot shows a Financial Times article page. At the top, the 'FINANCIAL TIMES' logo is visible. Below it, a navigation bar includes 'HOME', 'WORLD', 'US', 'COMPANIES', 'TECH', 'MARKETS', 'CLIMATE', 'OPINION', 'WORK & CAREERS', 'LIFE & ARTS', and 'HOW TO SPEND IT'. The article is categorized under 'Technology' with a '+ Add to myFT' button. The main headline reads 'NHS signs up for Tim Berners-Lee pilot to reinvent web'. Below the headline, a sub-headline states 'BBC, NatWest and government of Flanders also test new Solid servers from Inrupt'. A large photograph of Tim Berners-Lee speaking into a microphone is featured. Below the photo, a caption reads: 'Tim Berners-Lee: "There is a feeling, a zeitgeist, that change is really overdue." © AFP/Getty Images'. At the bottom of the article preview, it says 'John Thornhill in London NOVEMBER 9 2020' and shows a comment icon with the number '118' and a share icon.



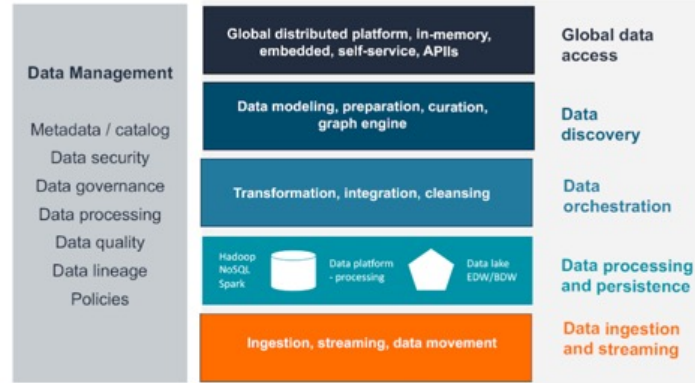


— — **In a data fabric, knowledge graphs** connect siloed data and simplify complicated data.

“A knowledge graph in a data fabric solves the data integration task and bridges the gap between disparate corporate datasets and the actual business requirements.

It allows business users to model, explore and discover how datasets are connected in a human readable model where there is little difference between the conceptual and the physical machine readable models.”

Source: Demystifying the Data Fabric | 9.17.20 | Gartner Group



Cloud ↑ Data sources ↑ On-premises

A knowledge graph built on semantics and graph data models fits into the data orchestration and discovery layers in a data fabric.

The Data Fabric is an architecture for modern data management that anticipates the need to connect data across the enterprise at speed and scale.

Knowledge graph is central to the success of any Data Fabric as an agile integration overlay that blends complex and diverse data into easy-to-consume data products for end-users, data scientists, and applications.





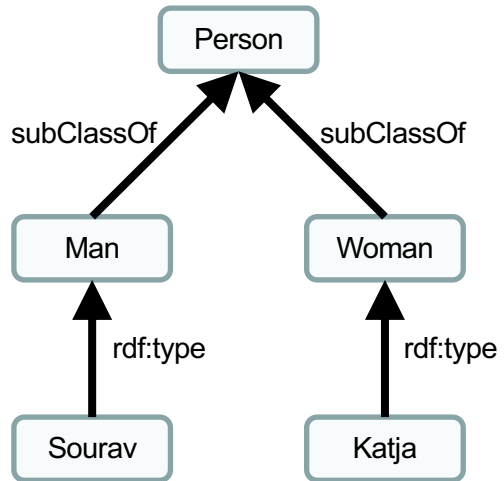
Pure Logic Reasoning

Pure Learning



Pure Logic Reasoning

Pure Learning


$$(Katja, \text{rdf:type}, Woman) \wedge (Woman, \text{subClassOf}, Person) \rightarrow (Katja, \text{rdf:type}, Person)$$



Pure Logic Reasoning

Pure Learning



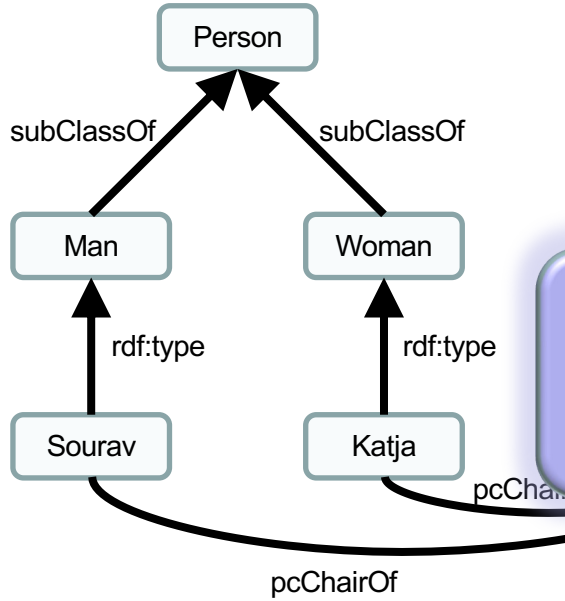
Person
Male

Person
Female


PC chairs opening EDBT 2023

role

conference



Who is who in the picture?
And how certain are you?




PC chairs opening EDBT 2023
role conference

EDBT 2023 has two PC chairs*
These are Katja and Sourav
Katja is a Woman, Sourav is a Man
Men and Women are Persons

The PC chairs of EDBT 2023 are (likely) on the picture
One is (likely) a male person, the other one (likely) female

* Partial Completeness Assumption (PCA)
Galárraga, Teflioudi, Hose, Suchanek. Fast Rule Mining in Ontological Knowledge Bases with AMIE+, VLDBJ 2015



What is knowledge?

Plato:


Knowledge is justified true belief.

Socrates:


True knowledge exists in knowing that you know nothing.




Data,
Information,
Knowledge,
Insight,
Wisdom
... and beyond

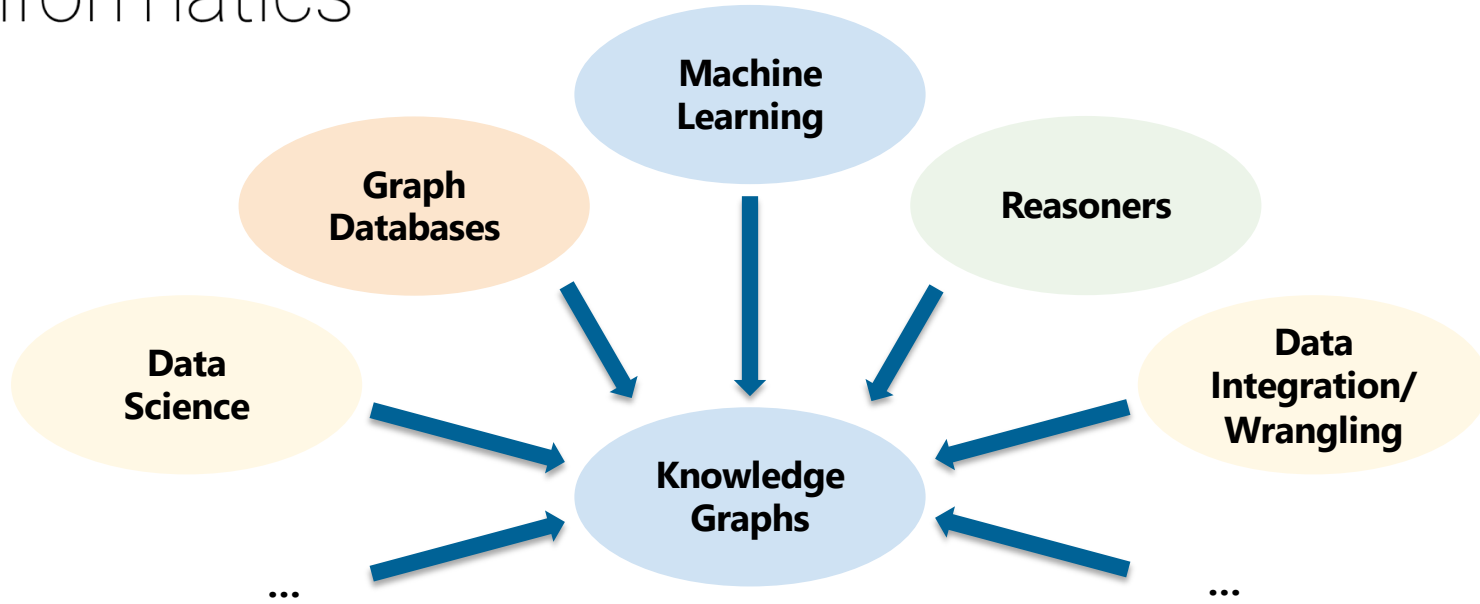



What is a Knowledge Graph?




What data scientists spend most of their time on







Source: <https://www.mkbergman.com/2244/a-common-sense-view-of-knowledge-graphs/>



Source: <https://www.mkbergman.com/2244/a-common-sense-view-of-knowledge-graphs/>

Google's Viewpoint

“A graph that understands real-world entities and their relationships to one another:
things, not strings”



“Introduction the Knowledge Graph: Things, Not Strings”. A. Singhal. Google Blog 2012.

Going Back to 1974

“A mathematical structure with **vertices** as *knowledge units* connected by **edges** that represent the *prerequisite relation*”

“On the Structure of the Teaching-learning Interactive Process”. E. Marchi and O. Miguel. I.J. Game Theory 1974.

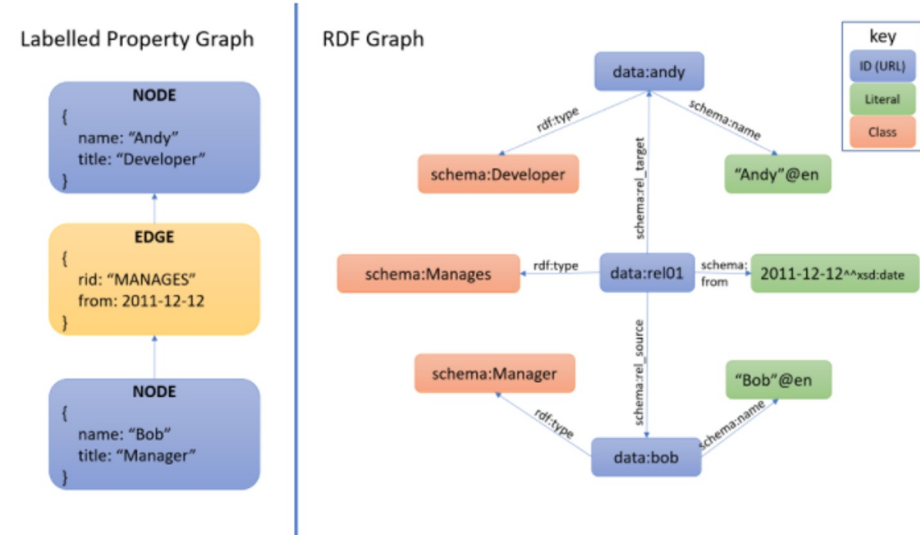
In this course

“A knowledge graph contains **semantically related** information encoded as **nodes and edges**.”

Katja Hose and Emanuel Sallinger, 2023


Enabling Technologies

- Two mainstream technologies (PG vs RDF)
- **Property Graphs (PG)**,
 - e.g., Neo4J, JanusGraph, OrientDB
 - *Programmer-friendly*
 - *Fast*
 - *Edge's attributes*
- **Resource Description Framework (RD)**
 - e.g., GraphDB, StarDog, Virtuoso
 - *Globally unique ID*
 - *Interoperability-at-scale*
 - *(shared) Schema-based*



Source: <https://terminusdb.com/blog/graph-fundamentals-part-2-labeled-property-graphs/>

Graph Databases become mainstream



188.399 Introduction to Semantic Systems

(2.0 VU / 3.0 ECTS)

2023W

Katja Hose

Emanuel Sallinger

iss@dbai.tuwien.ac.at

Learning Goals

After the lecture you will be able to:

- 1) explain the vision of semantic technologies
- 2) semantically model a domain of interest with ontologies
- 3) apply core semantic representation standards
- 4) manage and query knowledge graphs
- 5) build a knowledge graph by transforming (structured) data into a semantic representation

Course Organization and Grading

- 7 Lectures
- Three Individual Assignments, in total worth 60p
 - Submission of technical artifacts and report
- Written Exam (40p)
 - **19.01.2024**
 - You need to score at least 20p on the exam to pass the course
 - You may retake the exam (failed or already passed) on **15.03.2024**
 - **The later result counts!**

- Lectures:
 - 06.10.2023
 - 13.10.2023
 - 20.10.2023
 - 17.11.2023
 - 24.11.2023
 - 01.12.2023
 - 15.12.2023

Timeline of assignments

Assign ment	Assignment Topics	Available TUWEL	Submission Deadline
A1	Ontology modeling	20.10.2023	03.11.2023 (2 weeks)
A2	Knowledge Graph querying	17.11.2023	01.12.2023 (2 weeks)
A3	Programing Semantic Systems	01.12.2023	22.12.2023 (3 weeks)

Assignments need to be submitted in TUWEL by the specified deadline.

Assignments submitted after the deadline cannot be graded.

Course Grading

Grade = Assignments (max 60p) + Exam (max 40p)


Points	Evaluation	Grade
> 87	Sehr gut	1
86 –75	Gut	2
74 –61	Befriedigend	3
60 –50	Genügend	4
<50	Nicht Genügend	5

Recommended Literature

Book

Linked Data Management

Edited By [Andreas Harth](#), [Katja Hose](#), [Ralf Schenkel](#)




Book | [Open Access](#) | © 2020

Knowledge Graphs and Big Data Processing

[Home](#) > [Book](#)

Editors: [Valentina Janev](#), [Damien Graux](#), [Hajira Jabeen](#), [Emanuel Sallinger](#)




RESEARCH-ARTICLE

Knowledge Graphs



Authors: [Aidan Hogan](#), [Eva Blomqvist](#), [Michael Cochez](#), [Claudia D'amato](#), [Gerard De Melo](#), [Claudio Gutierrez](#), [Sabrina Kirrane](#), [José Emilio Labra Gayo](#), + 10 [Authors Info & Affiliations](#)

ACM Computing Surveys, Volume 54, Issue 4 • July 2021 • Article No.: 71, pp 1–37 • <https://doi.org/10.1145/3447772>



Next steps

- If you decide not to take the course, make sure to cancel your registration **before Oct 20, 2023 [!]**
- Register for the course in TUWEL to access all course materials. Link available in TISS.

Questions?

General Questions

- TUWEL Forum ... if the answer might be interesting for other students
- Answer questions from colleagues (preferably without posting the solution :-) ...

Specific Questions

- iss@dbai.tuwien.ac.at

What Is An Ontology?

Emanuel Sallinger
Katja Hose


What is an ontology?

- Ontology = ontos (being) + logos (word)
- Philosophy: Systematic explanation of Existence
 - Aristotle (400BC) attempts to establish universal categories for classifying everything that exists
- Merriam Webster dictionary:
 - A branch of metaphysics concerned with the nature and relations of being;
 - A particular theory about the nature of being and kinds of existents.
- Ontology - is both a science and its result


What is an ontology?

Studer(98):


Formal, explicit specification of a shared conceptualization.




Machine
readable &
interpretable



Concepts, properties,
functions, axioms
are explicitly defined

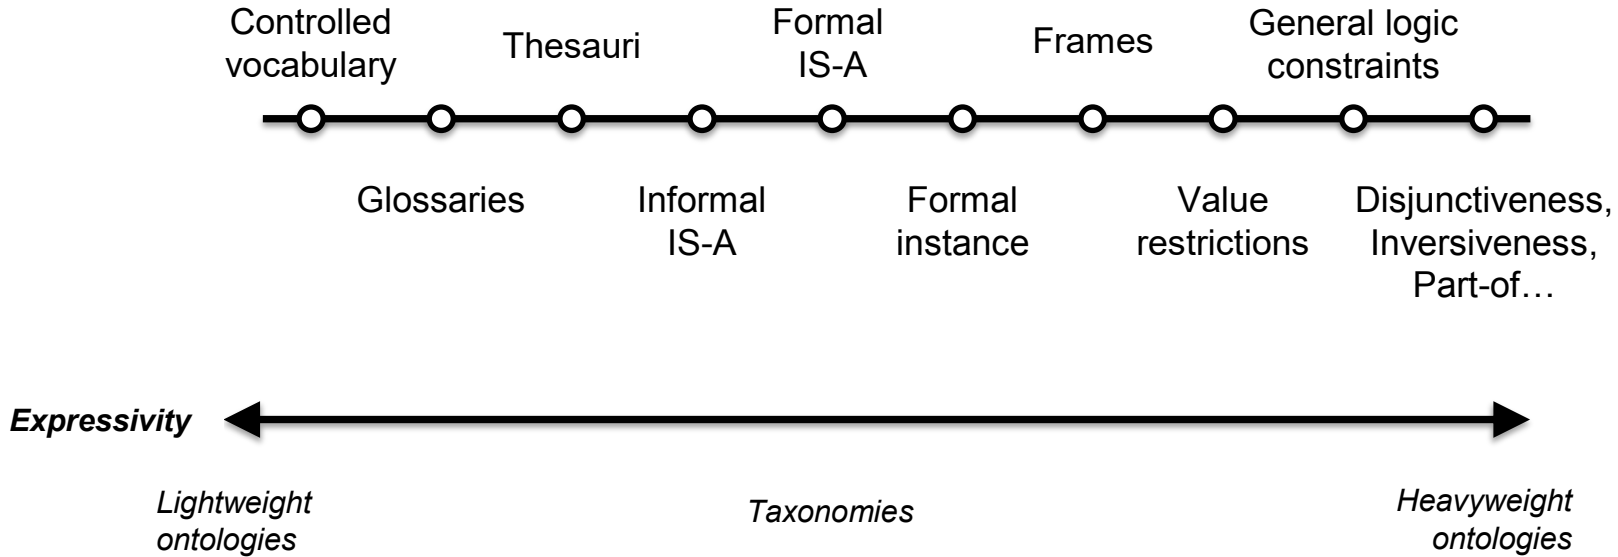


Shared by a
community




Abstract model
of some
phenomena
in the world

Ontology Types and Categories: Overview



Lightweight Ontologies

- **Controlled Vocabulary:** finite list of terms (e.g., a catalogue)
- **Glossary:**
finite list of terms + informal definition in natural language
- **Thesauri:**
controlled vocabulary + concepts are connected with relations
 - Equivalency (synonyms)
 - Hierarchies (subclasses, superclasses)
 - Homographs (Homonyms)
 - Associations (similar concepts)



A popular lightweight ontology

Friend of a Friend



FOAF Basics

- [Agent](#)
- [Person](#)
- [name](#)
- [nick](#)
- [title](#)
- [homepage](#)
- [mbox](#)
- [mbox_sha1sum](#)
- [img](#)
- [depiction](#) ([depicts](#))
- [surname](#)
- [family_name](#)
- [givenname](#)
- [firstName](#)

Personal Info

- [weblog](#)
- [knows](#)
- [interest](#)
- [currentProject](#)
- [pastProject](#)
- [plan](#)
- [based_near](#)
- [workplaceHomepage](#)
- [workInfoHomepage](#)
- [schoolHomepage](#)
- [topic_interest](#)
- [publications](#)
- [geekcode](#)
- [myersBriggs](#)
- [dnaChecksum](#)

Online Accounts / IM

- [OnlineAccount](#)
- [OnlineChatAccount](#)
- [OnlineEcommerceAccount](#)
- [OnlineGamingAccount](#)
- [holdsAccount](#)
- [accountServiceHomepage](#)
- [accountName](#)
- [icqChatID](#)
- [msnChatID](#)
- [aimChatID](#)
- [jabberID](#)
- [yahooChatID](#)

Projects and Groups

- [Project](#)
- [Organization](#)
- [Group](#)
- [member](#)
- [membershipClass](#)
- [fundedBy](#)
- [theme](#)

Documents and Images


- [Document](#)
- [Image](#)
- [PersonalProfileDocument](#)
- [topic](#) ([page](#))
- [primaryTopic](#)
- [tipjar](#)
- [sha1](#)
- [made](#) ([maker](#))
- [thumbnail](#)

Taxonomies

Definition of a hierarchical system of groups

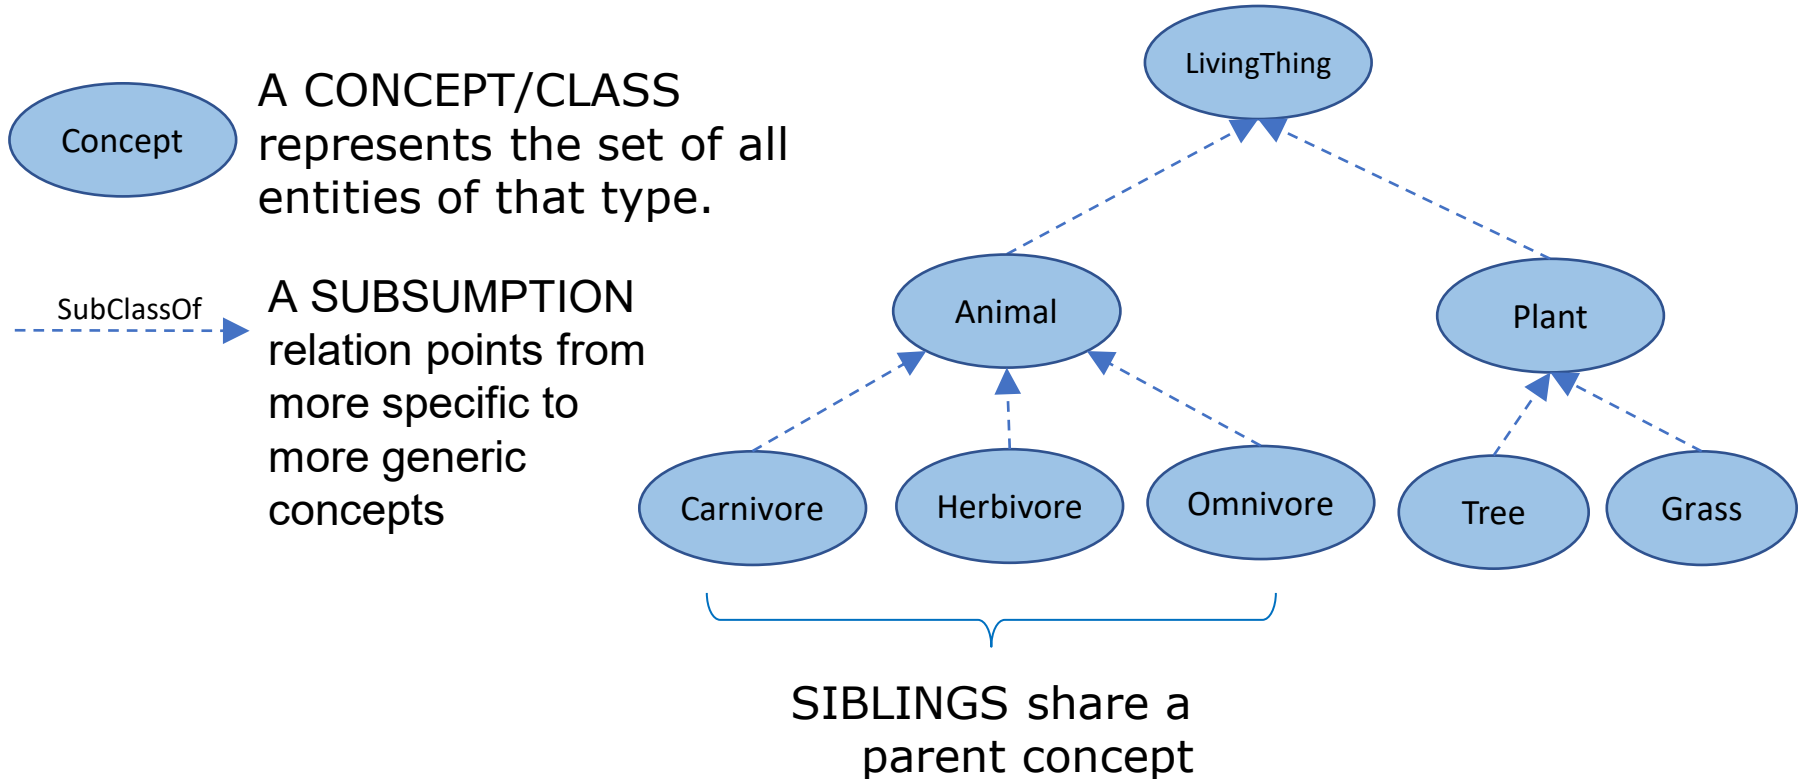
[from *τάξις* (*taxis*) = order, arrangement and *νόμος* (*nomos*) = law, science]

- **Informal IS-A-Hierarchy:** explicit hierarchy of classes
subclass relations are not strict (e.g. index of a library)
- **Formal IS-A-Hierarchy:**
explicit hierarchy of classes, subclass relations are strict
- **Formal instance:**
explicit class hierarchy, strict subclass relations +
instance-of relations are allowed



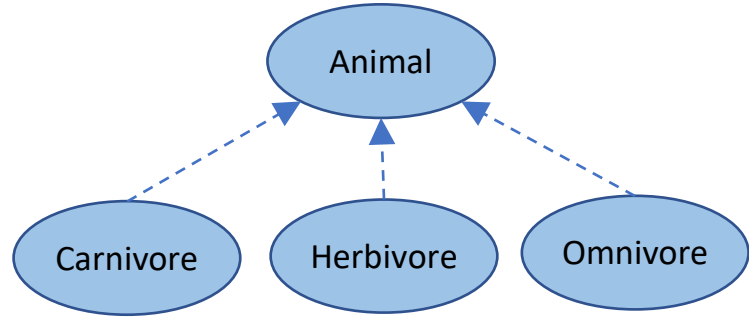
Taxonomy

Taxonomy: a controlled vocabulary organized into a hierarchical structure.



Subsumption/SubClass/Specialization Relation (in Formal IS-A Taxonomies)

Semantics (Meaning): If a *Herbivore* is a subClass of *Animal*, all instances of *Herbivore* are instances of *Animal*.



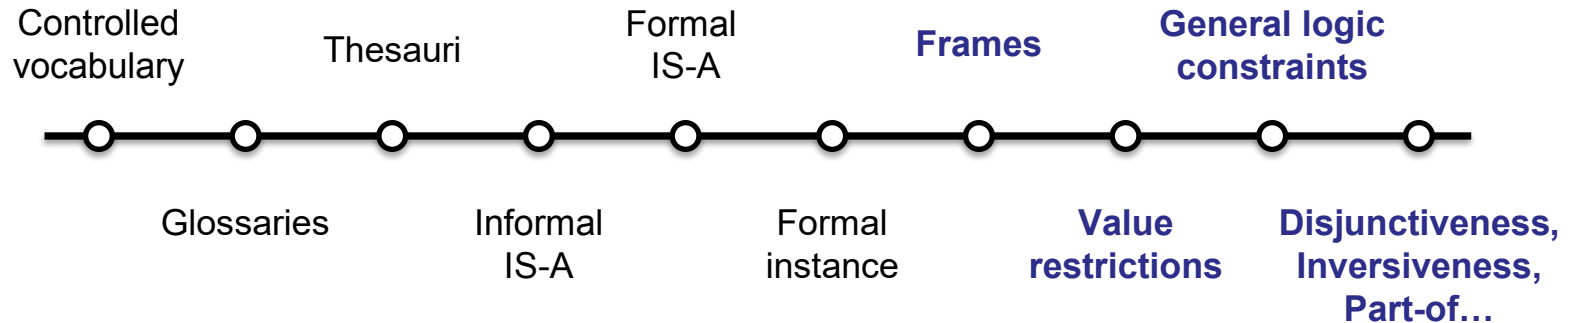
Common Mistake:

Subsumption is often incorrectly used to represent other types of relations such as part of (meronymy).



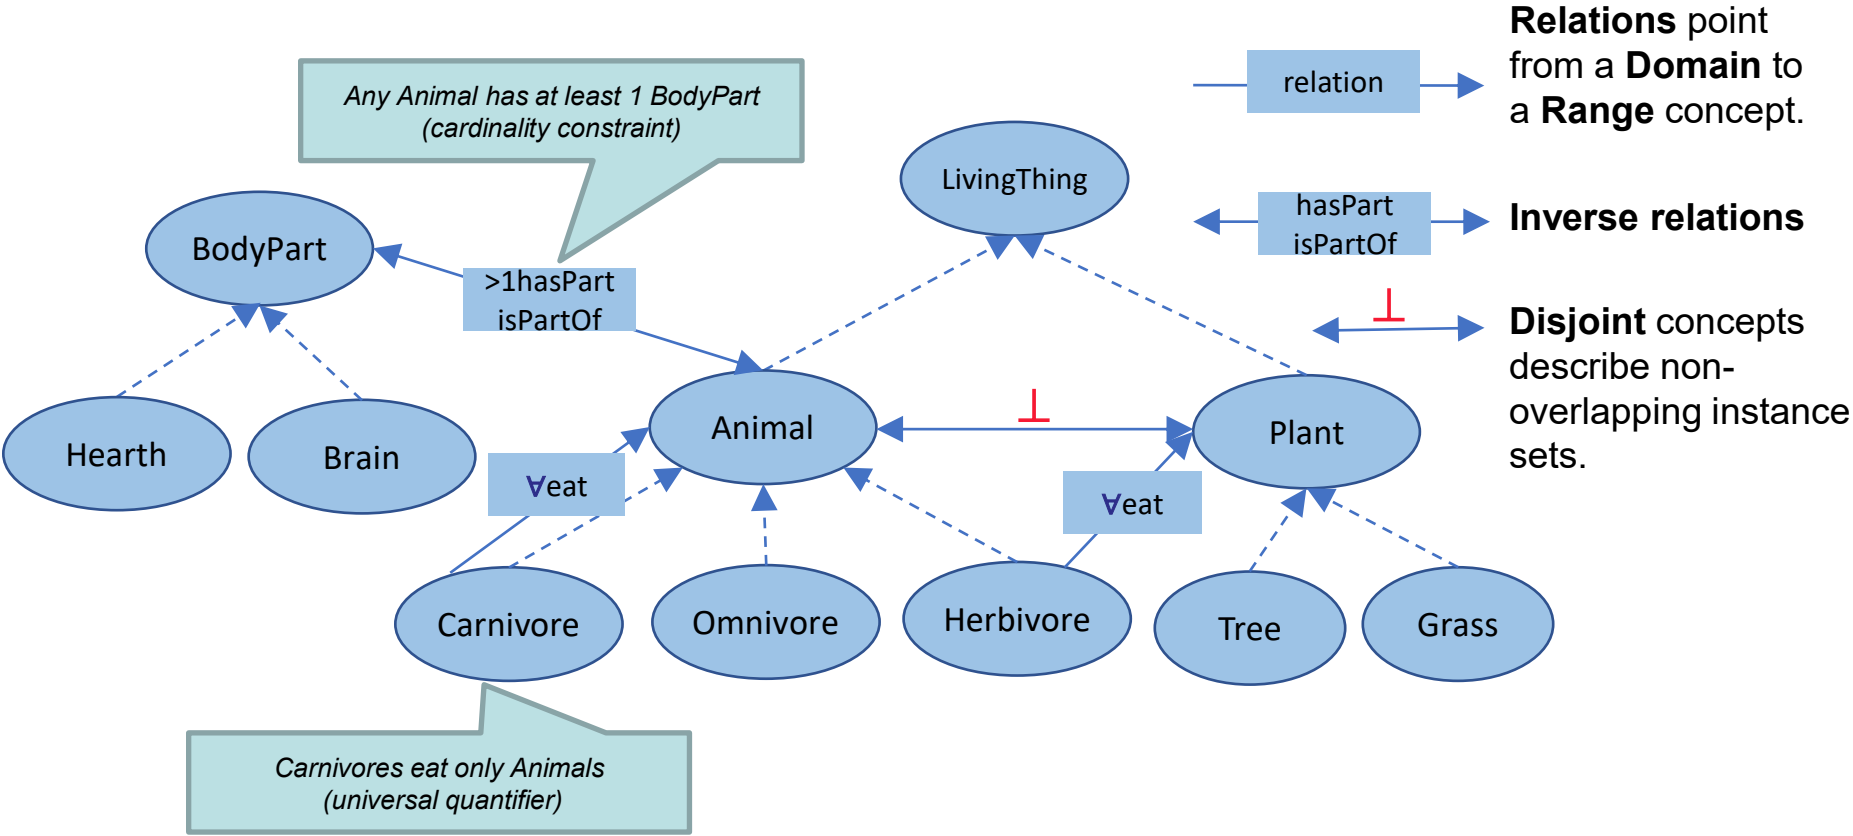
Heavyweight Ontologies

- Rigorous definition and organization of concepts
- Focus on *logic*, i.e., formally correct deductions and inferences
- Careful formal specification, need for consistency

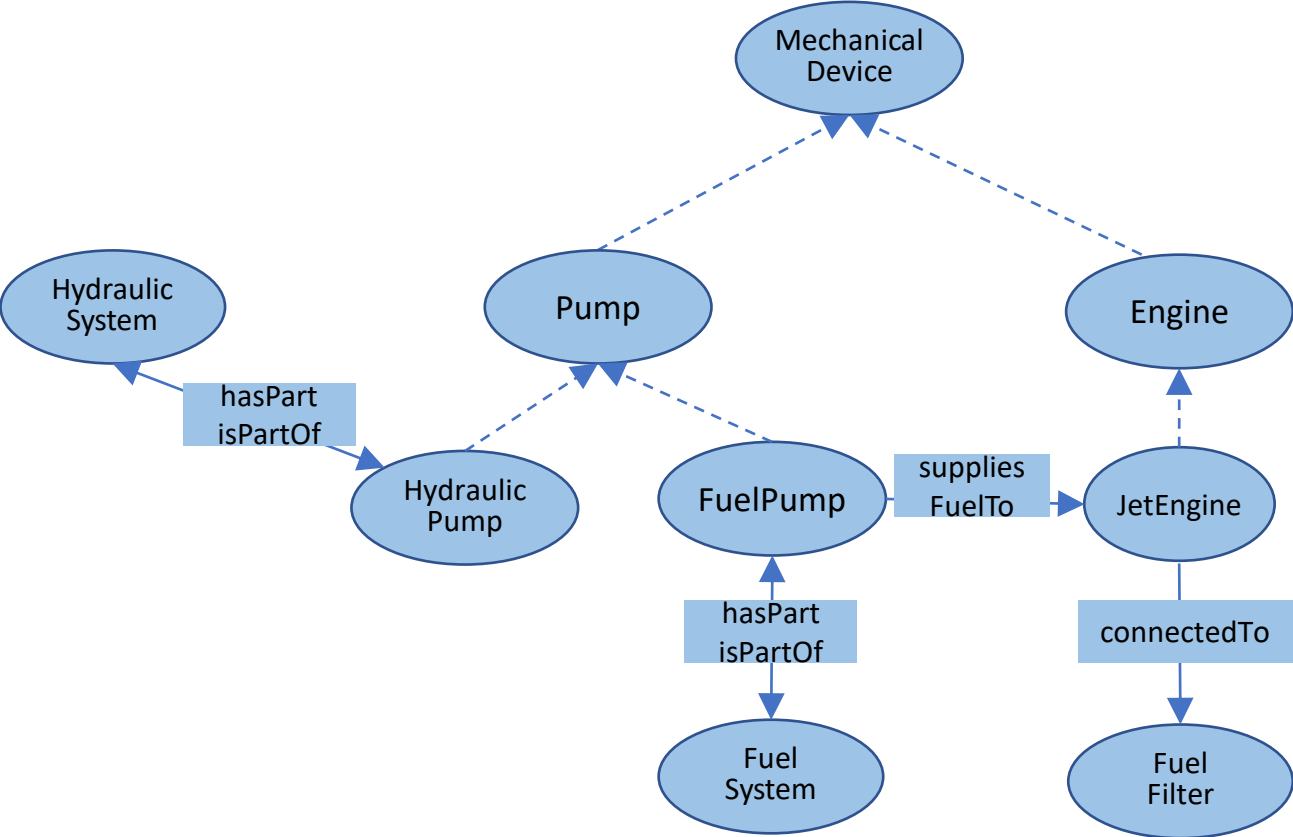


Ontology

Ontology: A taxonomy extended with other relations and further constraints.



Ontology- Mechanics




Instances/Entities



Metadata to describe the image:



- Tom is a *Carnivore*
- Jerry is an *Animal*
- Tom *eats* Jerry

Terminological Space



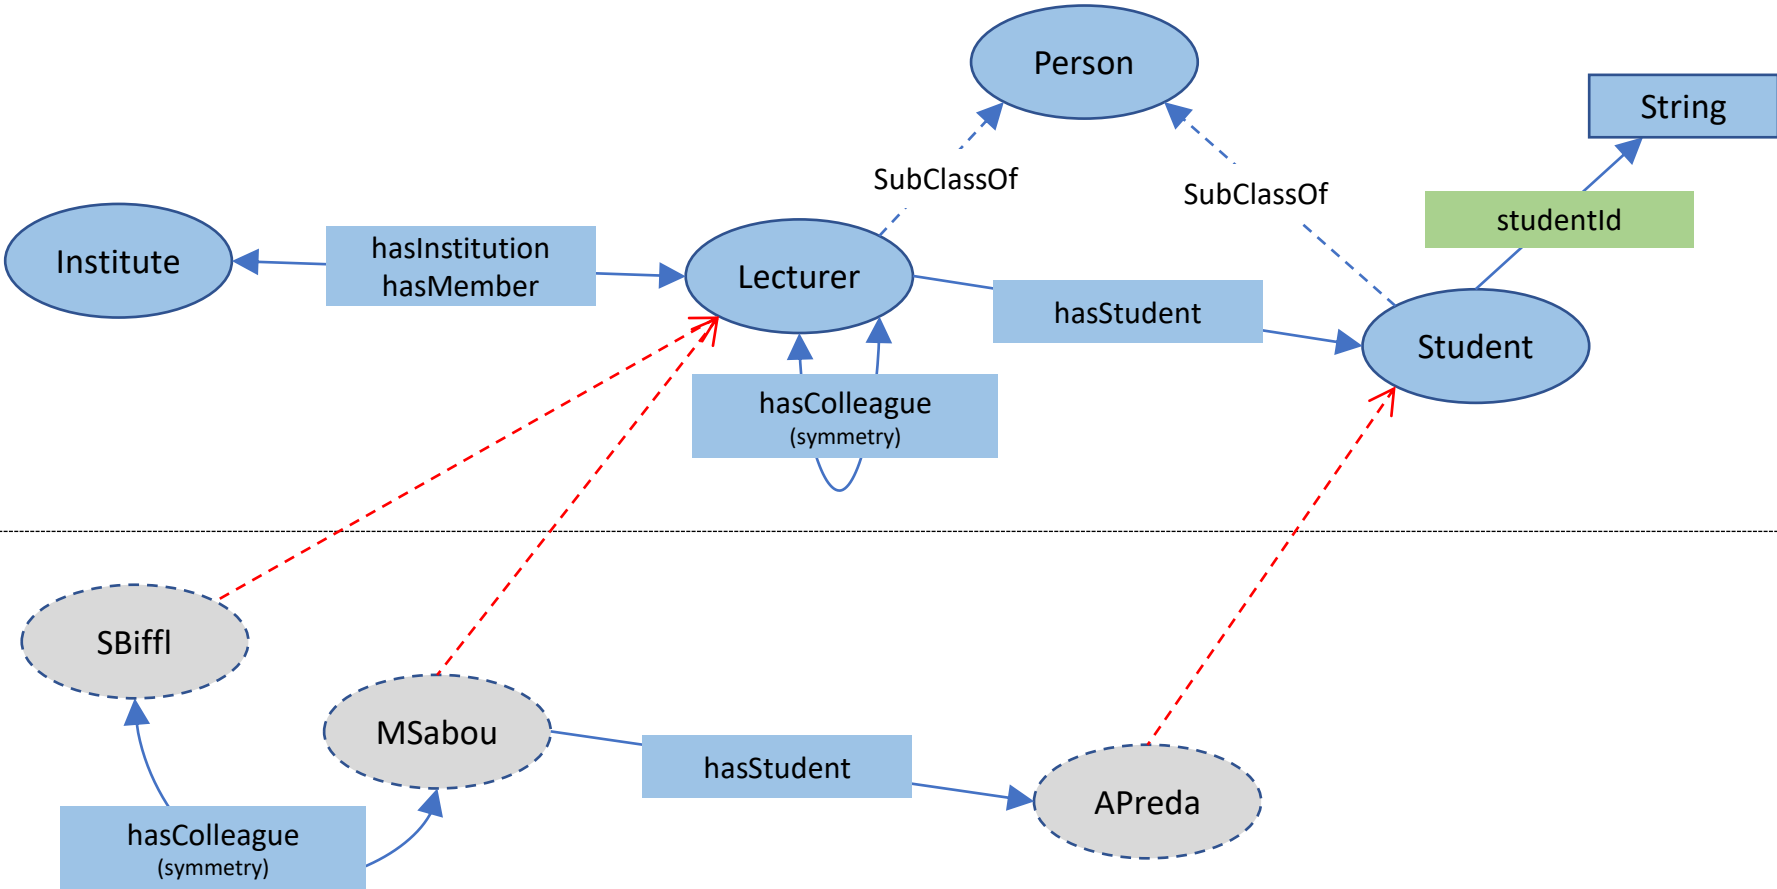
Instance represents an entity in the universe of discourse.

Instance Space



Types the instance with respect to a concept (the instance is of type concept).

More Examples ...



Recap:

What's inside an ontology?

- Concepts:
 - Denote the main concepts of the domain
 - *E.g.: Carnivore, Animal, Pump, Engine*
- Concept hierarchy
 - Denotes specialization/generalizations
 - *E.g.: Carnivore is a kind of Animal*
- Relations between classes
 - *E.g.: Carnivore eats Animal*
- Restrictions on relations (type, cardinality)
 - *E.g.: any Animal has at least one BodyPart*
- Instances:
 - Denote concrete entities in the domain
 - *E.g.: Tom, Jerry*

What Is An Ontology?

Emanuel Sallinger
Katja Hose

What Is An Ontology?

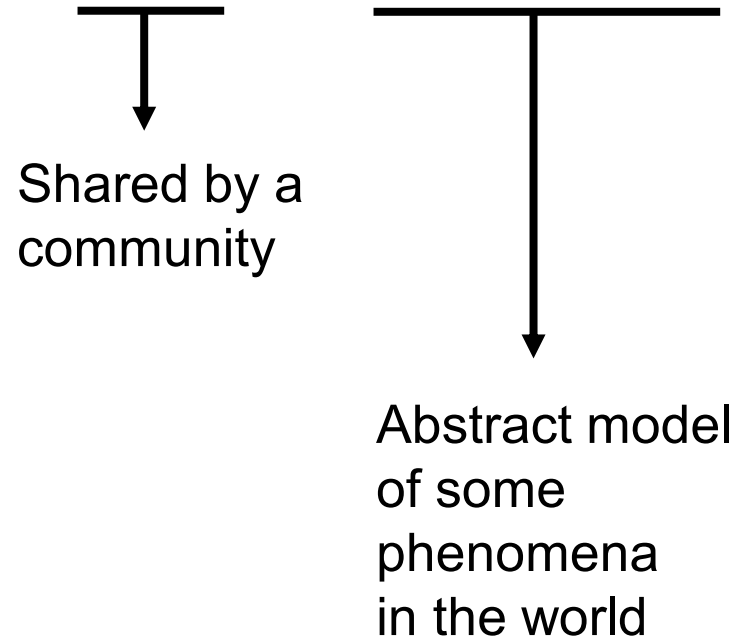
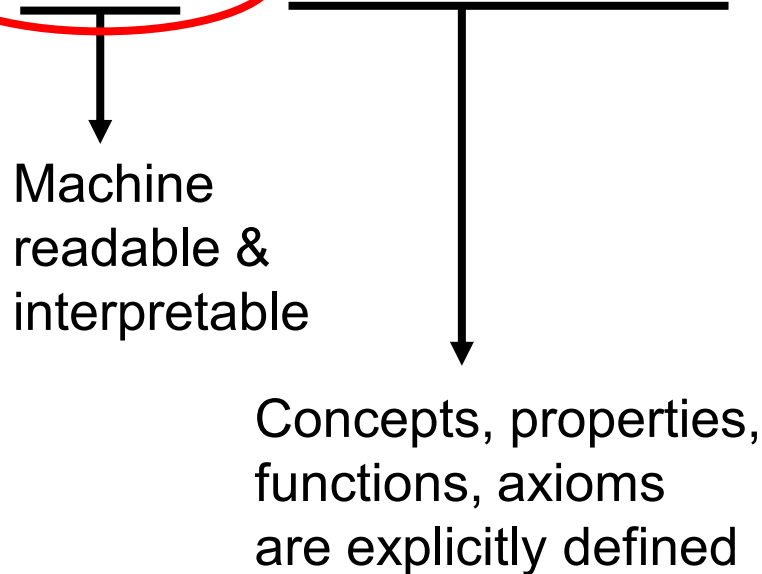
Description Logics

Emanuel Sallinger
Katja Hose

What is an ontology?

Studer(98):

Formal, explicit specification of a shared conceptualization.



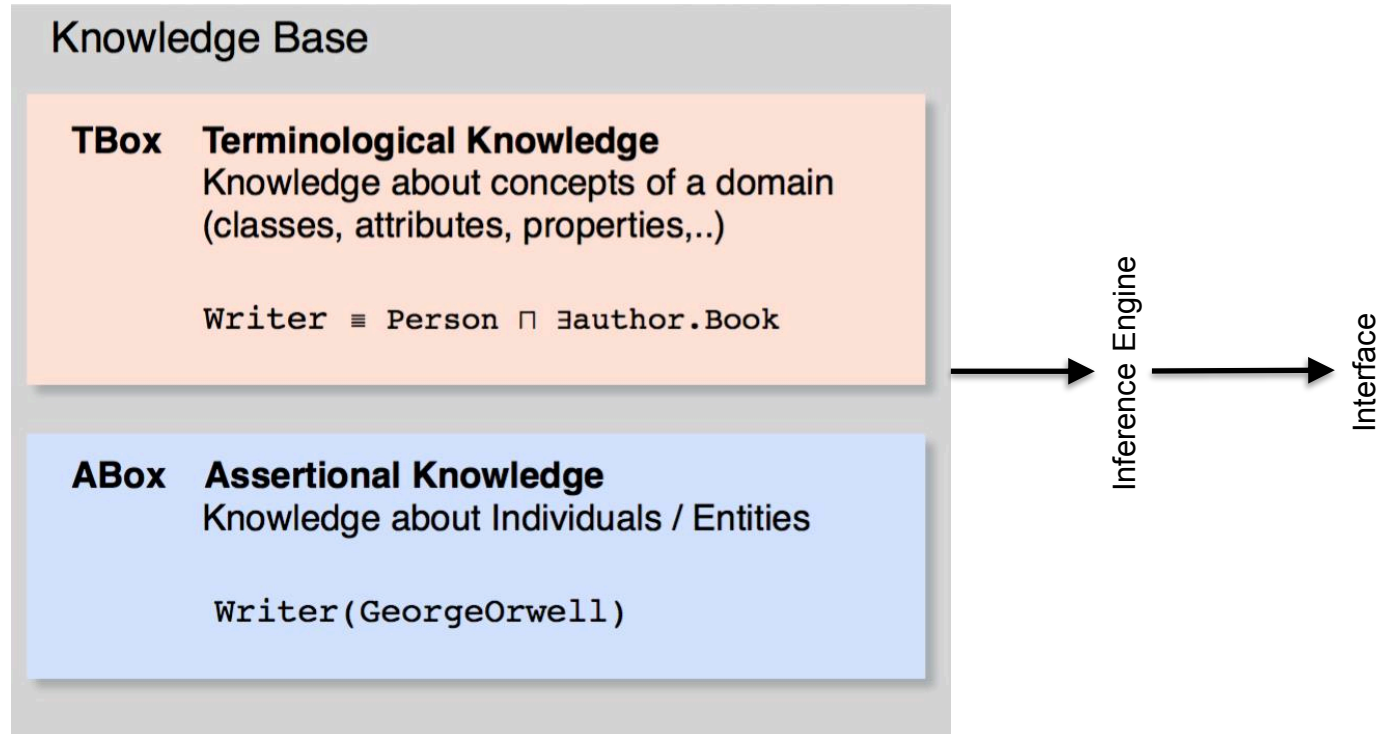
Logics as basis for Ontologies

- Ontologies should be “machine readable”
- ... but computers cannot read and make sense of images

- Formal logics are used as basis for representing knowledge
- A formal, logics-based language provides:
 - A syntax (which expressions are correct)
 - A semantics (the meaning of expressions)
 - A calculus (how to determine the meaning of expressions)
 - which allow computers to draw “logical” conclusions from the data (i.e., to reason with data)

- Current Semantic Web languages are based on Description Logics (a subset of First Order Logics)

General Description Logic (DL) Architecture



Atomic Types

- Concept names A, B, \dots
- Special concepts
 - \top Top (universal concept)
 - \perp Bottom concept
- Role names R, S, \dots

```
Writer(GeorgeOrwell)
Book(NineteenEightyFour)
Author(NineteenEightyFour, GeorgeOrwell)
```

Constructors

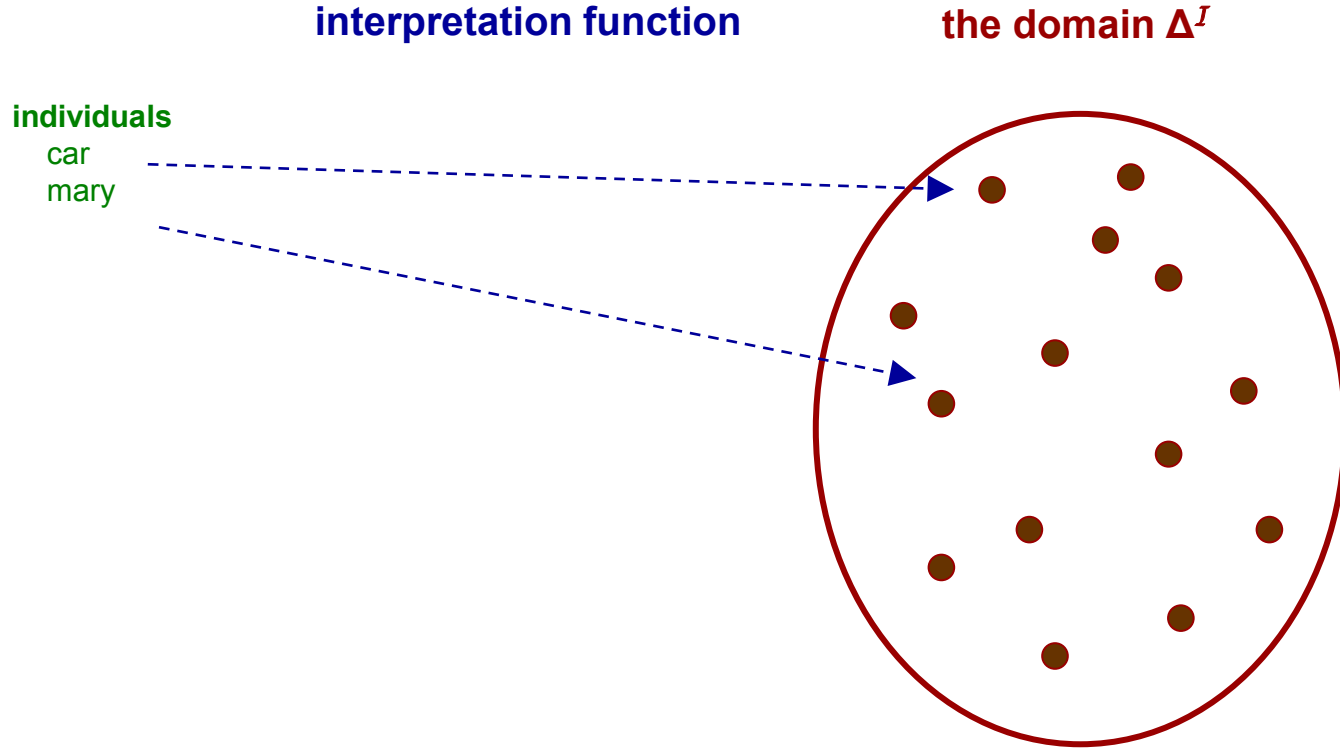
- Negation: $\neg C$
- Conjunction: $C \sqcap D$
- Disjunction: $C \sqcup D$
- Existential quantifier: $\exists R.C$
- Universal quantifier: $\forall R.C$

```
Novel  $\sqsubseteq$  Prose
Novel  $\sqsubseteq$  Book
Herbivore  $\sqsubseteq \neg$  Carnivore


PetOwner  $\sqsubseteq$  Person  $\sqcap \exists$  hasPet.Animal
Carnivore  $\sqsubseteq$  Animal  $\sqcap \forall$  eats.Animal
```

- Model theoretic semantics: An interpretation consists of
 - A domain of discourse (a collection of objects) Δ
 - Interpretative Functions I mapping
 - classes to sets of objects
 - properties to sets of pairs of objects
- In a DL, a class description is thus a characterization of the individuals that are members of that class


Description Logics Semantics (1/3)



Description Logics Semantics (2/3)



Description Logics Semantics (3/3)



What Is An Ontology?

Description Logics

Emanuel Sallinger
Katja Hose

How To Build an Ontology?

Ontology Engineering

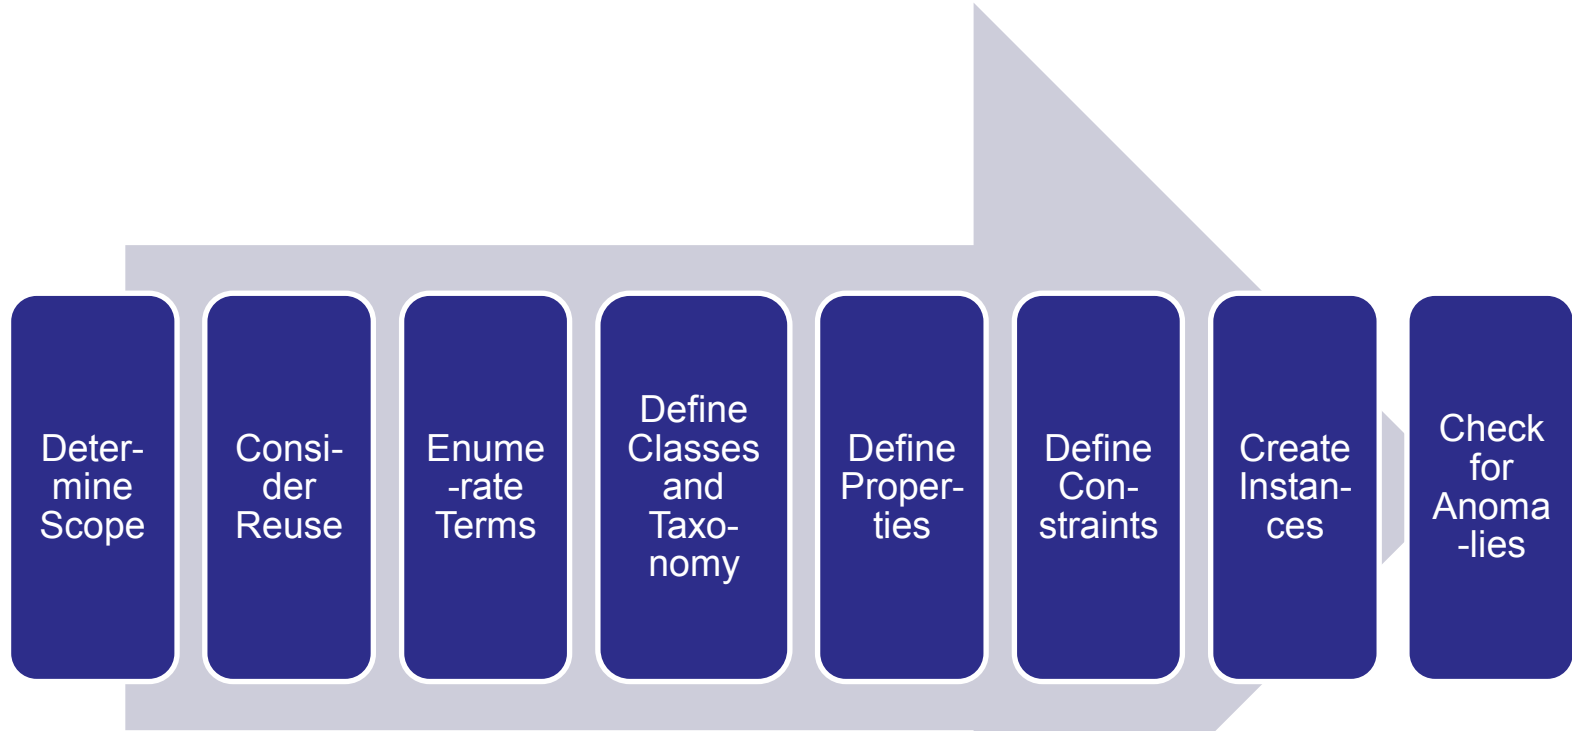
Emanuel Sallinger
Katja Hose

Ontology Engineering

- All activities involved in the development of ontologies
- Methodologies for:
 - specification and design of ontologies (**Ontology Design**)
 - comparison of ontologies (**Ontology Mapping**)
 - combination of ontologies (**Ontology Merging**)
 - retrieving ontologies from a set of information resources (**Ontology Learning**)
 - populating ontologies with individuals from information resources (**Ontology Population**)


Ontology Development 101

(Noy and McGuinness, 2000)




Not a linear process!

Determine Scope



- What is included in the ontology should be determined by:
 - the uses of the ontology
 - future extensions that are already anticipated
- Basic questions to be answered at this stage:
 - What is the **domain** that the ontology will cover?
 - For what are we going to **use** the ontology?
 - For what **types of questions** should the ontology provide answers?
 - Called “Competency Questions”
 - **Who** will use and maintain the ontology?


Determine Scope: Competency Questions



Competency questions for a wine recommender system:


- Which wine characteristics should I consider when choosing a wine?
- Is Bordeaux a red or white wine?
- Does Cabernet Sauvignon go well with seafood?
- What is the best choice of wine for grilled meat?
- Which characteristics of a wine affect its appropriateness for a dish?
- Does a bouquet or body of a specific wine change with vintage year?
- What were good vintages for Napa Zinfandel?

Consider Reuse




- Why reuse other ontologies?
 - to save ontology engineering effort
 - to interact with the tools that use other ontologies
 - to use ontologies that have been validated through use in applications
- Ontology libraries
 - Protégé ontology library (<http://protege.stanford.edu>)
 - NCBO Bioportal (<http://bioontology.org>)
 - <http://ontologydesignpatterns.org/wiki/Ontology:Main>

Enumerate Terms




- Write down a list of all the relevant terms that are expected to appear in the ontology
 - What are the terms we need to talk about?
 - Nouns form the basis for class names
 - What are the properties of these terms?
 - Verbs (or verb phrases) form the basis for property names
- ▶ **Examples: The Wine ontology**
 - wine, grape, winery, location,
 - wine color, wine body, wine flavor, sugar, content
 - white wine, red wine, Bordeaux wine

Define Classes and a Taxonomy

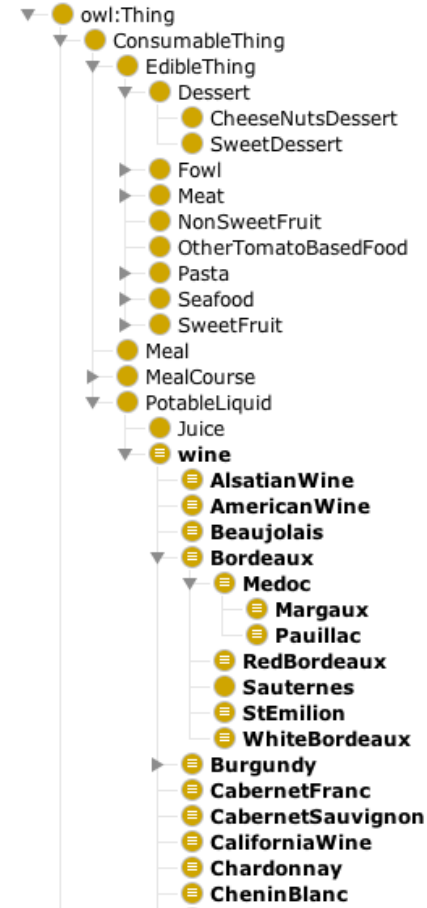


- A class is a concept in the domain
 - a class of **wines**
 - a class of **wineries**
 - a class of **red wines**
- A class is a collection of elements with similar properties
- Instances of classes
 - a glass of **California wine you'll have for lunch**


Define Classes and a Taxonomy



- Relevant terms must be organized in a taxonomic hierarchy
 - If A is a subclass of B, then every instance of A must also be an instance of B
- Example
 - *Red wines is a subclass of Wine*
 - *Every red wine is a wine*




Define Properties



- Often interleaved with the previous step
- Properties in a class definition:
 - describe attributes of instances of the class
 - each wine **has color**, **sugar content**, **producer**, etc.
 - describe relations to other classes:
 - each **Wine has maker Winery**
 - „parent class“ is the **domain** of property, e.g., **Wine**
 - „child class“ is the **range** of the property, e.g., **Winery**
- If A is a subclass of B, every property statement that holds for instances of B must also apply to instances of A
- Properties shared in a hierarchy should be specified in the highest possible class


Define Properties: Property Hierarchies



- Properties can be arranged in hierarchies too.



Define Constraints: Cardinality



- Cardinality restrictions
 - “each wine is made from at least one grape ”
 - “each wine has exactly one maker ”

```
Wine ⊆ >1madeFromGrape  
Wine ⊆ =1hasMaker
```

Description: wine


Equivalent To +

☰ Wine

SubClass Of +

- hasBody **exactly 1** owl:Thing
- hasColor **exactly 1** owl:Thing
- hasFlavor **exactly 1** owl:Thing
- hasMaker **exactly 1** owl:Thing
- hasMaker **only** Winery
- hasSugar **exactly 1** owl:Thing
- locatedIn **some** Region
- madeFromGrape **min 1** owl:Thing

Define Constraints: Required Values



- Universal restriction (all values from):
 - If a property is declared for an instance, all its values must be of a certain type
 - Wines can only be made in Wineries

Wherever there is a **hasMaker** property, all its values must be instances of **Winery**.
[There may be zero **hasMaker** properties.]

- Existential restriction (some values from):
 - There **must** be at least one property with that value
 - E.g., Wines are located in regions

There must be at least one **locatedIn** property whose value is **Region**.
[There must be at least one **locatedIn** property.]


Description: wine

Equivalent To +
Wine

SubClass Of +


- hasBody **exactly 1** owl:Thing
- hasColor **exactly 1** owl:Thing
- hasFlavor **exactly 1** owl:Thing
- hasMaker **exactly 1** owl:Thing
- hasMaker **only** Winery
- hasSugar **exactly 1** owl:Thing
- locatedIn **some** Region
- madeFromGrape **min 1** owl:Thing

Define Constraints: Property Characteristics



- Symmetry
 - If $P(x,y) \Rightarrow P(y,x)$
 - *E.g., marriedTo, friendOf, adjacentRegion*
- Transitivity
 - If $P(x,y)$ and $P(y,z) \Rightarrow P(x, z)$
 - *E.g., locatedIn, hasAncestor*
- Inverse properties
 - If $P(x,y) \Rightarrow \text{Inv}P(y,x)$ holds
 - *E.g., hasParent(John, Bob) \Rightarrow hasChild(Bob, John)*
- Functional values
 - can have only one (unique) value y for each instance x
 - i.e. there cannot be two distinct values y_1 and y_2 such that the pairs (x,y_1) and (x,y_2) are both instances of this property
 - *E.g., hasHusband*


Create Instances



- Create an instance of a class
 - The class becomes a direct type of the instance
 - Any superclass of the direct type is a type of the instance
- Assign property values for the instance frame
 - Property values should conform to the constraints
 - Knowledge-acquisition tools often check it
- Number of instances >> number of classes
- Populating an ontology with instances is usually not done manually
 - retrieved from legacy data sources
 - extracted automatically from a text corpus

See Lecture 6 on KG creation

Create Instances



Description: GaryFarrellMerlot ? || ▢ □ ✕ **Property assertions: GaryFarrellMerlot** || ▢ □ ✕

Types +
Merlot ? @ ✕ ○

Same Individual As +
Different Individuals +

Object property assertions +

hasFlavor	Moderate	? @ ✕ ○
hasSugar	Dry	? @ ✕ ○
hasMaker	GaryFarrell	? @ ✕ ○
locatedIn	SonomaRegion	? @ ✕ ○
hasBody	Medium	? @ ✕ ○

Description: MountadamRiesling ? || ▢ □ ✕ **Property assertions: MountadamRiesling** || ▢ □ ✕


Types +
DryRiesling ? @ ✕ ○

Same Individual As +
Different Individuals +

Object property assertions +

hasMaker	Mountadam	? @ ✕ ○
hasFlavor	Delicate	? @ ✕ ○
hasSugar	Dry	? @ ✕ ○
locatedIn	SouthAustraliaRegion	? @ ✕ ○
hasBody	Medium	? @ ✕ ○

Check for Anomalies

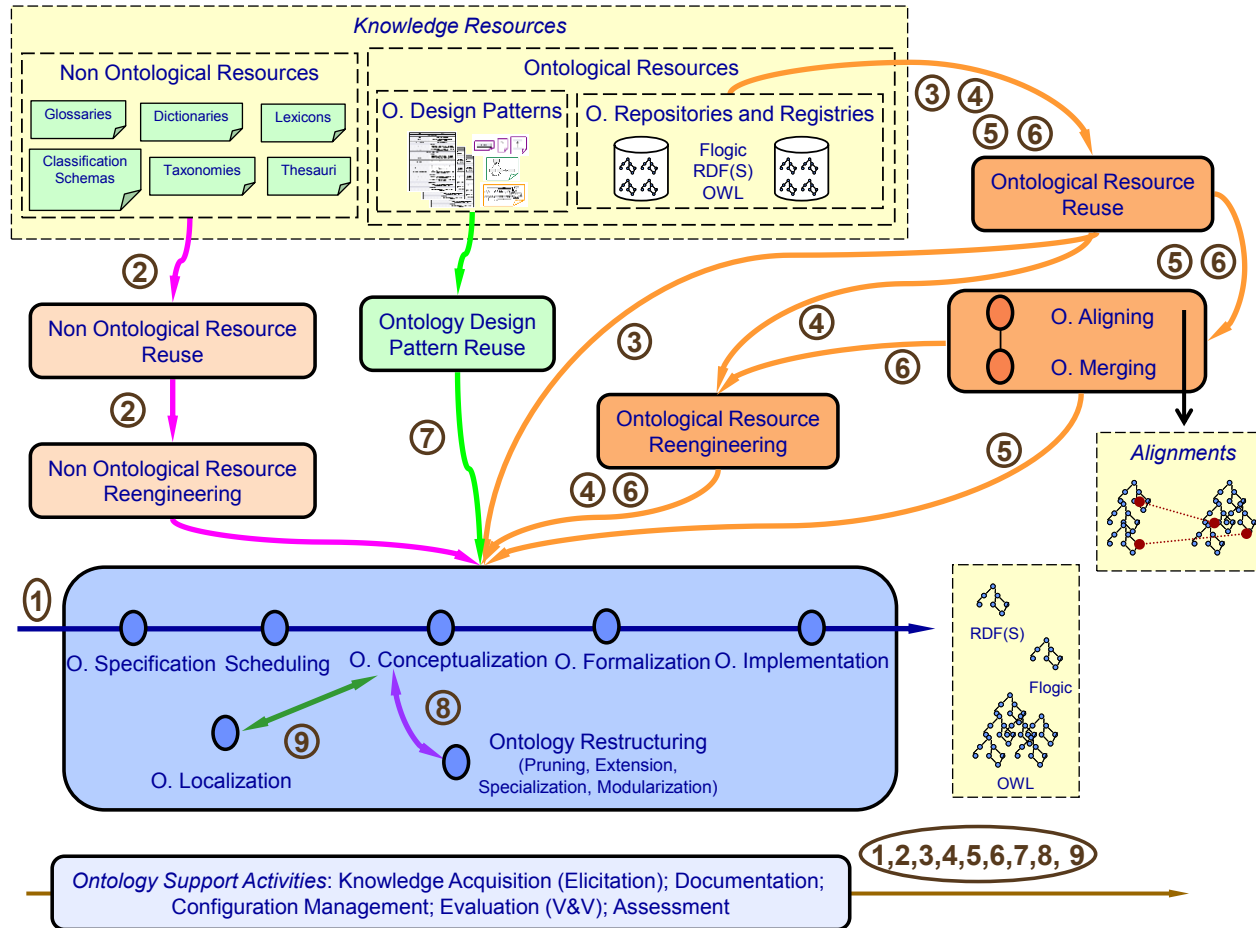


Exploits the formal semantics of ontologies to:

- 1) Validate ontology models
 - Consistency checking
 - Satisfiability
- 2) Derive additional knowledge
 - Classification

See "Reasoning" part of Lecture

NeOn Methodology - a much more complex picture



How To Build an Ontology?

Ontology Engineering

Emanuel Sallinger
Katja Hose

What Can I Do With An Ontology?

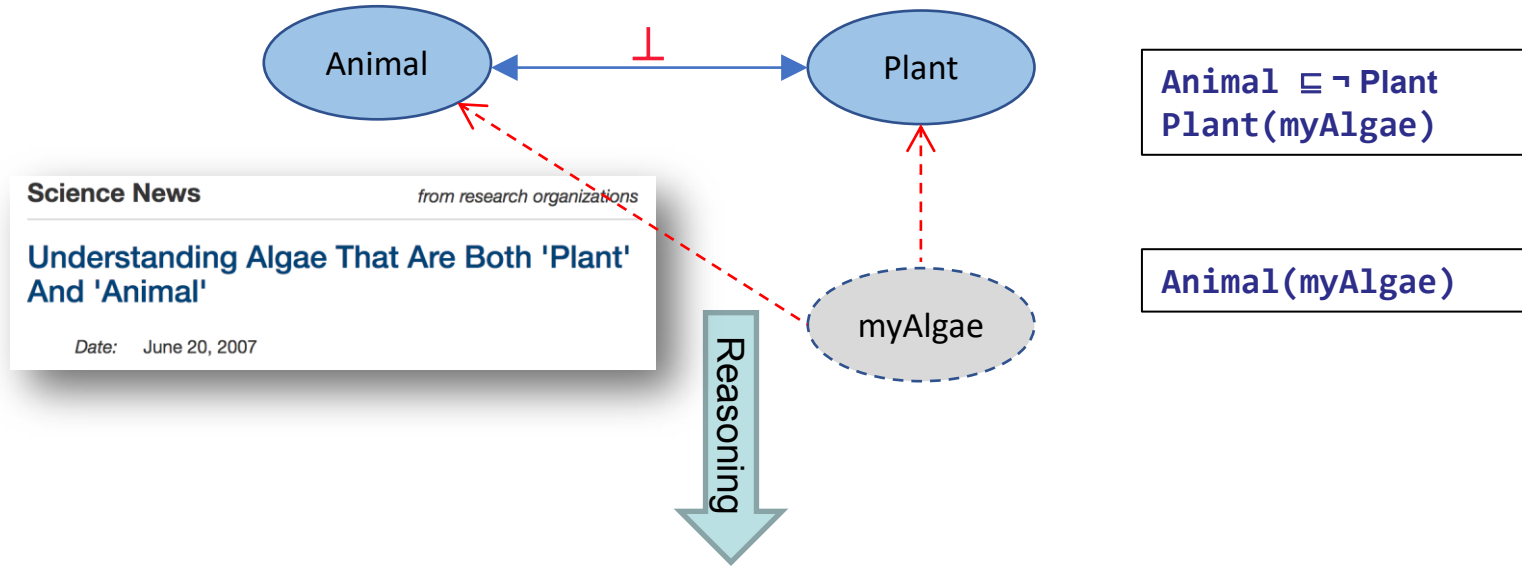
Reasoning

Emanuel Sallinger
Katja Hose

Reasoning with Ontologies

- Thanks to the formal nature of ontologies (their grounding in logic), **reasoning** can be used to **infer** new knowledge based on already declared information.
- Reasoning is performed with **reasoners (inference engines)**.
- Reasoning allows for, e.g.,:
 - **Consistency checking**: Checking the correctness of the ontology model (e.g., that the model described by the ontology is free of logical contradictions)
 - **Class inferences**: discovering relation between classes that were not explicitly declared.
 - **Instance inferences**: discovering class membership of instances.

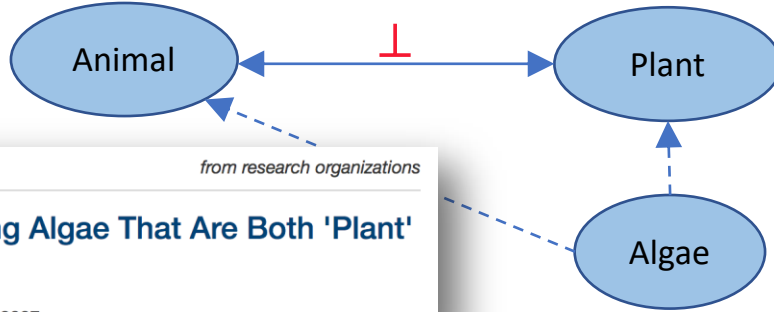
Consistency Checking



Inconsistent ontology: There are logical errors in the ontology

Satisfiability Checking

Are there classes that cannot possibly have any instances?



$Algae \sqsubseteq Plant$
 $Animal \sqsubseteq \neg Plant$
 $Algae \sqsubseteq Animal$

Science News *from research organizations*

Understanding Algae That Are Both 'Plant' And 'Animal'

Date: June 20, 2007

Unsatisfiable Concept: if there is a contradiction in the ontology that implies that the class cannot have any instances

Class hierarchy (inferred)
Class hierarchy

Class hierarchy: Algae

- owl:Thing
 - Plant
 - Algae
 - Animal
 - Algae

Annotations: Algae

Description: Algae

Equivalent To

SubClass Of

- Animal
- Plant

Reasoning

Class hierarchy (inferred)
Class hierarchy

Class hierarchy: Algae

- owl:Thing
 - Plant
 - Algae
 - Animal
 - Algae

Annotations: Algae

Description: Algae

Equivalent To

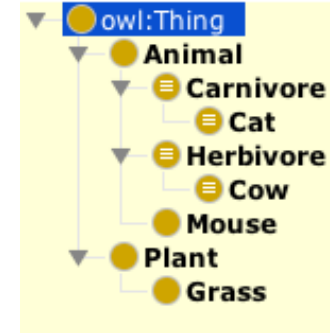

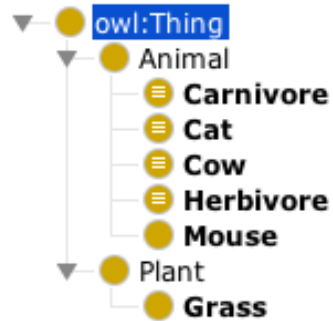
- owl:Nothing

SubClass Of

- Animal
- Plant

Class Inference

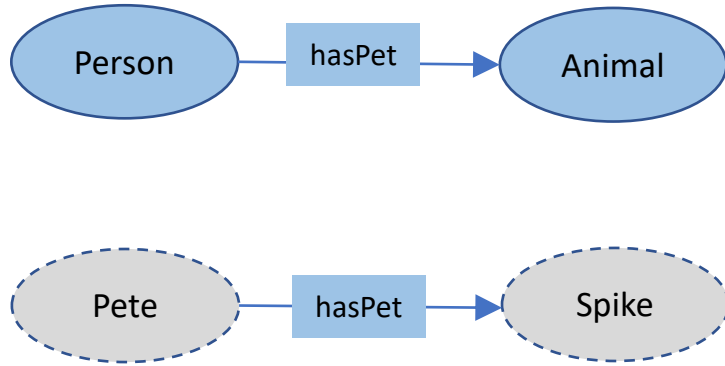
Discover new knowledge based on the given knowledge.




$\text{Carnivore} \sqsubseteq \text{Animal} \sqcap \forall \text{eats}.\text{Animal}$
 $\text{Herbivore} \sqsubseteq \text{Animal} \sqcap \forall \text{eats}.\text{Plant}$
 $\text{Cat} \sqsubseteq \text{Animal} \sqcap \forall \text{eats}.\text{Mouse}$
 $\text{Cow} \sqsubseteq \text{Animal} \sqcap \forall \text{eats}.\text{Grass}$
 $\text{Mouse} \sqsubseteq \text{Animal}$
 $\text{Grass} \sqsubseteq \text{Plant}$

$\text{Cat} \sqsubseteq \text{Carnivore}$
 $\text{Cow} \sqsubseteq \text{Herbivore}$

Instance Inference



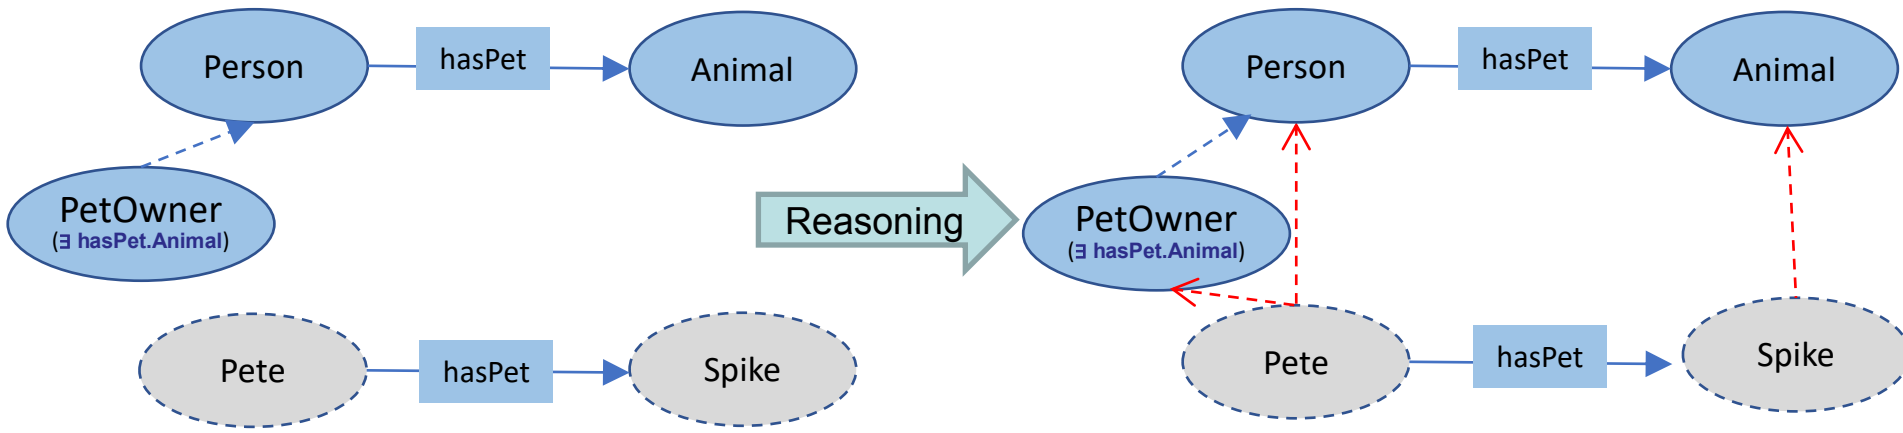
Reasoning



```
hasPet(Person, Animal)
Thing(Pete)
Thing(Spike)
hasPet(Pete, Spike)
```

```
Person(Pete)
Animal(Spike)
```

Instance Inference




hasPet(Person, Animal)
PetOwner \sqsubseteq Person $\sqcap \exists \text{hasPet}.\text{Animal}$

Thing(Pete)
Thing(Spike)
hasPet(Pete, Spike)

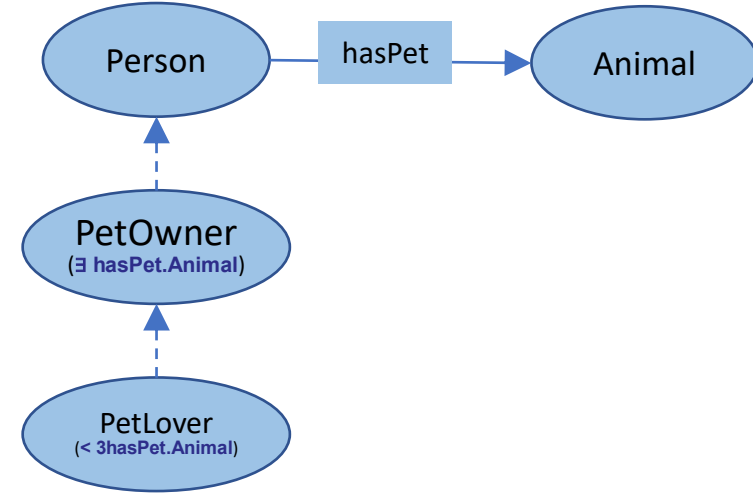
Person(Pete)
Animal(Spike)

PetOwner(Pete)

Class+Instance Inference



Reasoning



```
hasPet(Person, Animal)
PetOwner  $\sqsubseteq$  Person  $\sqcap$   $\exists \text{hasPet. Animal}$ 
PetLover  $\sqsubseteq$  Person  $\sqcap$   $\leq 3 \text{hasPet. Animal}$ 
```

```
Person(Pete)
Animal(Tom), Animal(Jerry), Animal(Spike)
hasPet(Pete, Spike)
hasPet(Pete, Tom)
hasPet(Pete, Jerry)
{Spike}  $\sqsubseteq$   $\neg\{\text{Tom, Jerry}\}$ 
```

```
PetOwner(Pete)
PetLover(Pete)
PetLover  $\sqsubseteq$  PetOwner
```


What Can I Do With An Ontology?

Reasoning

Emanuel Sallinger
Katja Hose

Ontology Editors in Practice

Emanuel Sallinger
Katja Hose

Practical Ontology Engineering

- To create your ontologies you can make use of:
 - WebVOWL (a light-weight, web-based editor)
 - Protégé (a complete ontology editor)
- These slides provide instructions to get started with these tools

Introduction to WebVOWL Editor: Creating a Class

1. Select one of the available class construct in the left panel
2. Double click anywhere on the screen
3. Edit the class name and label by double clicking on the screen OR
4. Use editing options from the right panel;
Selected Element -> IRI and Label.

The screenshot displays the WebVOWL Editor interface. On the left, a 'Default Element' panel lists various OWL constructs, with 'owl:Class' highlighted and a red circle '1' next to it. The central workspace shows a class hierarchy: 'Class1' is a subclass of 'Class2', and 'Class2' is connected to 'Class3' via a 'newObjectProperty'. 'Class3' is connected to 'Library' and 'Journal' via 'newDatatypeProperty'. A red circle '2' is on a 'newClass' button, and a red circle '3' is on a blue arrow icon. On the right, the 'Editing Options' panel shows fields for 'Title', 'IRI', 'Version', and 'Authors'. The 'Selected Element' section shows 'IRI: :Class14' (with a red circle '4'), 'Label: NewClass', and 'Type: owl:Class'. The bottom status bar includes a search field and icons for 'Ontology', 'Export', 'Filter', 'Modes', 'Reset', 'Pause', 'Options', and 'About'.

!Attention!

- Can be cumbersome
- Only limited editing functionality)

Introduction to WebVOWL Editor: Creating an Object Property

1. Select one of the available construct of Object Property in the left panel
2. Click on the domain (class) of the object property, and drag the arrow to the range (class) of the object property
3. Rename the property by double clicking the property name OR click on the property and edit it in the right panel.

The screenshot displays the WebVOWL Editor interface. On the left, a dark sidebar contains a 'Property: owl:objectPro...' section with a list of property types. The 'owl:objectProperty' type is highlighted with a red circle labeled '1'. The main workspace shows a class hierarchy with 'Class2' as a superclass of 'NewClass'. A new object property, 'newObjectProperty', is being created between 'Class2' and 'Class3'. A red circle labeled '2' is on the 'Class2' node, and another red circle labeled '3' is on the 'newObjectProperty' label. On the right, the 'Editing Options' panel is open, showing the 'Selected Element' section where the 'newObjectProperty' is selected with a red circle labeled '3'. The 'Type' is set to 'owl:ObjectProperty'. At the bottom, a search bar and navigation icons are visible.

Introduction to WebVOWL Editor: Creating an Datatype Property

1. Select one of the available construct of Datatype Property in the left panel
2. Click on the plus (+) button on the domain (class)
3. Rename the property by double clicking the property name OR click on the property and edit it in the right panel.
4. You can set the datatype by changing the value on the right panel

The screenshot displays the WebVOWL Editor interface. On the left, a 'Default Element' panel lists various constructs, with 'xsd:string' selected and marked with a red circle '1'. The central workspace shows a class hierarchy: Class1 is a subclass of Class2, and NewClass and Class3 are subclasses of Class2. Class3 has a red plus sign (+) on it, marked with a red circle '2'. A new datatype property, 'newDatatypeProperty', is being created on Class3, with its name highlighted in red and marked with a red circle '3'. The right panel, 'Editing Options', shows the configuration for the selected element, with the 'Type' set to 'owl:DatatypeProperty' and 'xsd:string' selected in the dropdown, both marked with red circles '3' and '4' respectively. The bottom of the interface features a search bar and navigation icons for 'Ontology', 'Export', 'Filter', 'Modes', 'Reset', 'Pause', 'Options', and 'About'.

Introduction to WebVOWL Editor: Export Function

- Export as TTL and SVG
 - Bottom menu -> export -> as TTL/SVG
 - JSON and URL exports don't seem to work at the moment

The screenshot displays the WebVOWL Editor interface. On the left, a sidebar lists various OWL elements such as 'owl:Class', 'owl:ObjectProperty', and 'rdfs:Literal'. The main workspace shows a graph with several classes and properties. A context menu is open over a red square on the graph, with options: 'Export as TTL', 'Export as JSON', 'Export as SVG', 'Export as URL: http://semantics.ic', and 'Copy'. The 'Export as TTL' and 'Export as SVG' options are highlighted. On the right, the 'Editing Options' panel is visible, showing fields for 'Title', 'IRI', 'Version', and 'Authors'. The 'Selected Element' section shows 'IRI: xsd:string', 'Label: string', 'Type: rdfs:Datatype', and 'Datatype: xsd:string'. The bottom status bar includes a search field and navigation icons for 'Ontology', 'Export', 'Filter', 'Modes', 'Reset', 'Pause', 'Options', and 'About'.

- Load your ontology file:
 - File -> Open

The screenshot displays the Protégé ontology editor interface. The main window title is 'movie (http://semantics.id/ns/example/movie/1.0.1)'. The menu bar includes File, Edit, View, Reasoner, Tools, Refactor, Window, Ontop, Mastro, and Help. The address bar shows the ontology IRI: 'http://semantics.id/ns/example/movie/1.0.1'. Below the address bar, there are tabs for 'Active Ontology', 'Entities', 'Object Properties', 'Data Properties', and 'Individuals by class'. The interface is divided into several panels:

- Ontology header:** Displays the 'Ontology IRI' as 'http://semantics.id/ns/example/movie' and the 'Ontology Version IRI' as 'http://semantics.id/ns/example/movie/1.0.1'.
- Annotations:** Shows a single annotation with the property 'rdfs:comment' and the value 'An ontology about movie and related concepts.'.
- Ontology metrics:** A table providing various counts for the ontology:

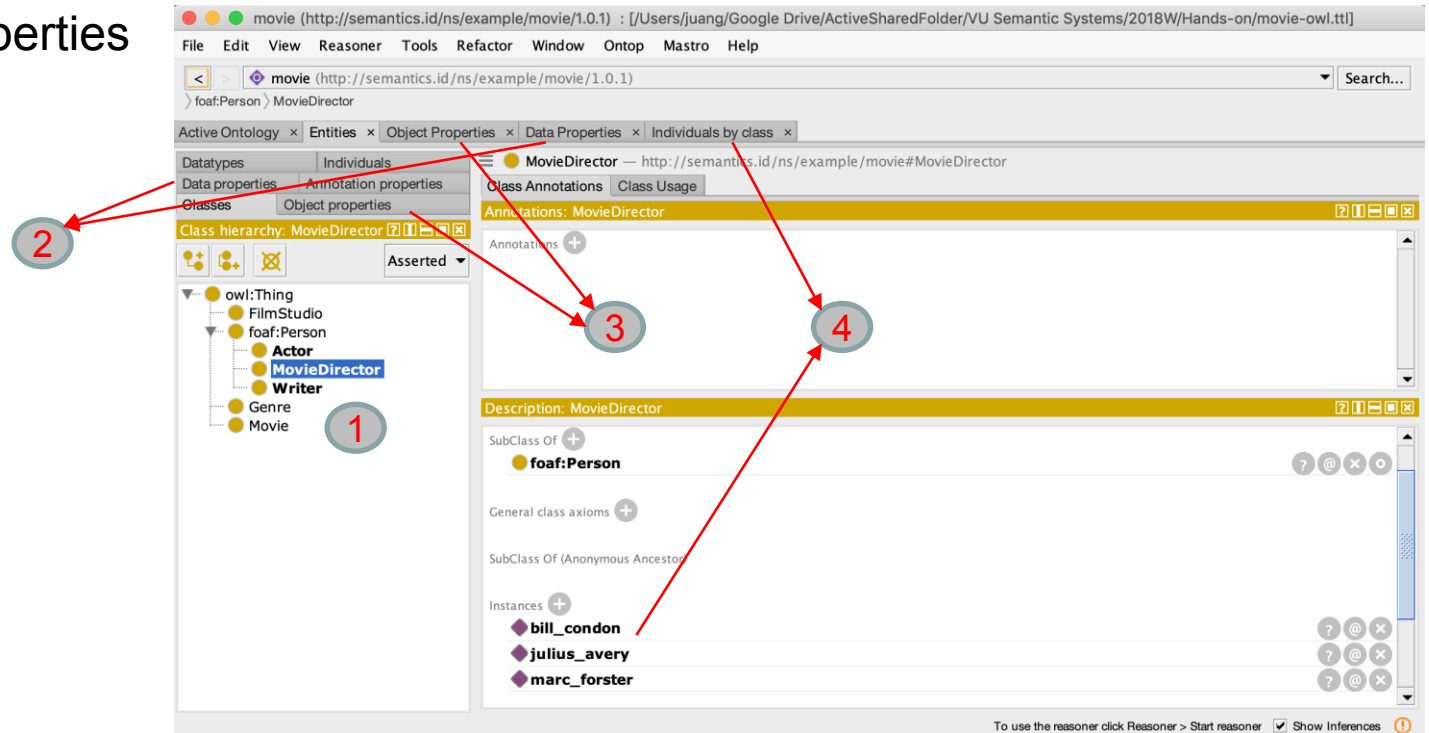
Metric	Count
Axiom	168
Logical axiom count	128
Declaration axioms count	24
Class count	7
Object property count	10
Data property count	5
Individual count	24
Annotation Property count	3
DL expressivity	ALHIF+ (D)
- Class axioms:** A table showing the number of axioms for different types:

Class axiom type	Count
SubClassOf	3
EquivalentClasses	0
DisjointClasses	0
GCI count	0
Hidden GCI Count	0
- Ontology imports:** A section for 'Imported ontologies' with sub-sections for 'Direct Imports' and 'Indirect Imports', both currently empty.

At the bottom right, there is a status bar with the text 'To use the reasoner click Reasoner > Start reasoner' and a checked checkbox for 'Show Inferences'.

Protégé Editor: Classes, Properties, and Instances

1. Classes
2. Datatype Properties
3. Object Properties
4. Instances



Protégé Editor: Customize your views

The screenshot shows the Protégé Editor interface with the following components:

- Menu Bar:** File, Edit, View, Reasoner, Tools, Refactor, **Window**, Ontop, Mastro, Help.
- Views Menu:** Open, showing options like Active Ontology, Entities, Classes, Object Properties, Data Properties, Annotation Properties, Individuals by class, OWLViz, DL Query, Individual Hierarchy Tab, **OntoGraf** (selected), SWRLTab, Ontop Mappings, Ontop SPARQL, Mastro, SQWRLTab, and SPARQL Query.
- Tabs Menu:** Open, showing options like Create new tab..., Delete custom tabs..., Import tab..., Export current tab..., Store current layout, Reset selected tab to default state, Capture view to clipboard..., Timestamp log / console, Show log..., Look & Feel, and Refresh user interface.
- Class Hierarchy:** A tree view showing the hierarchy: owl:Thing (parent) -> FilmStudio, foaf:Person (parent) -> Actor, MovieDirector (selected), Writer, Genre, Movie.
- Class Details:** The 'foaf:Person' class is selected, showing its general class axioms, subclass relationships (SubClass Of (Anonymous Ancestor)), and instances: bill_condon, julius_avery, and marc_forster.
- Workspace:** The main area displays the 'OntoGraf' view, showing a graph of the ontology with nodes and edges.
- Status Bar:** At the bottom, it says "To use the reasoner click Reasoner > Start reasoner" and "Show Inferences" (checked).

Explore ontology with OWLViz and OntoGraf

Protégé File Edit View Reasoner Tools Refactor Window Help

movie (http://semantics.id/ns/example/movie/1.0.1) : [/User
< > movie (http://semantics.id/ns/examp
af:Person > Actor

ve Ontology x Entities x Individuals by class x DL Query x

ss hierarchy Class hierarchy (inferred) Act

ss hierarchy: Actor Asserted

owl:Thing
● FilmStudio
● foaf:Person
● Actor
● MovieDirector
● Writer
● Genre
● Movie

Class An
Annotati

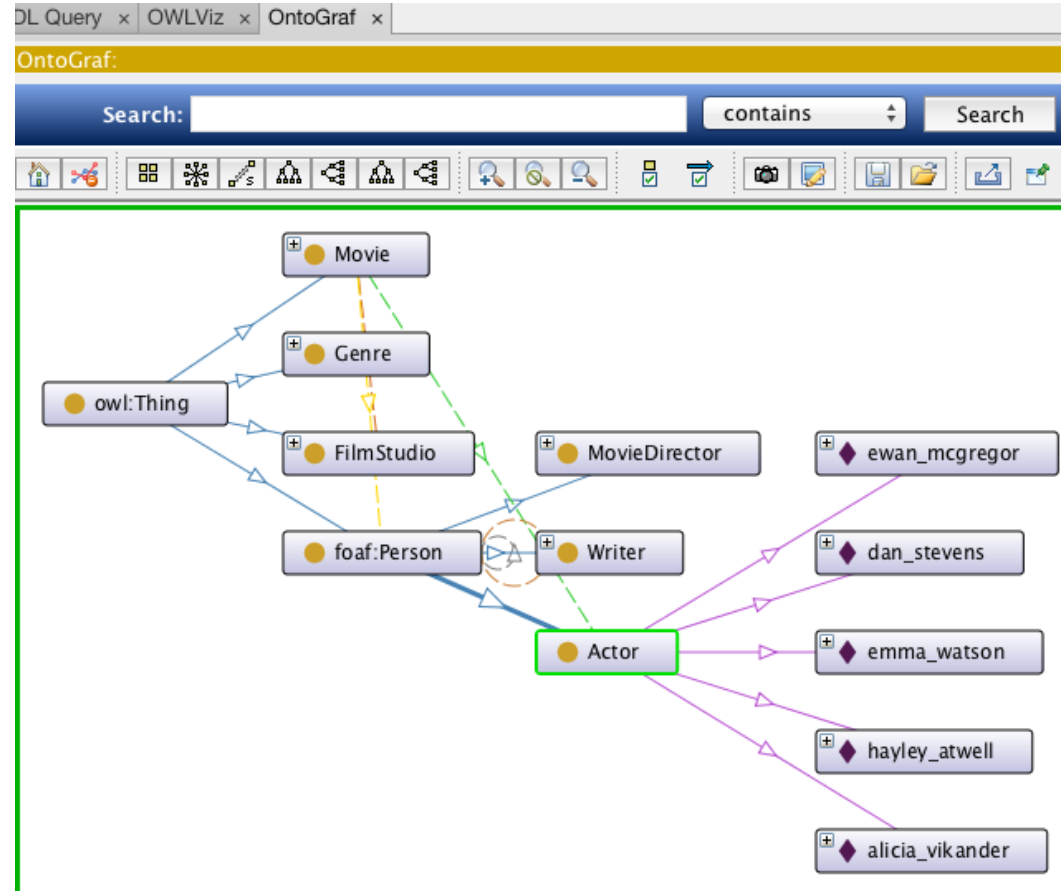
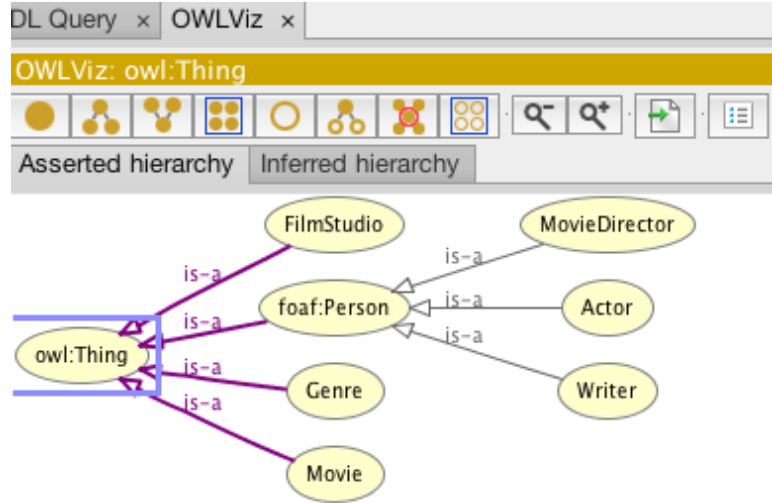
Annotation

Views
Tabs

Create new tab...
Delete custom tabs...
Import tab...
Export current tab...
Store current layout
Reset selected tab to default state
Capture view to clipboard...
Timestamp log / console
Show log...
Look & Feel
Refresh user interface

InnovationsLehrstuhl/Slides/P
✓ Active Ontology
✓ Entities
Classes
Object Properties
Data Properties
Annotation Properties
✓ Individuals by class
OWLviz
✓ DL Query
SWRLTab
OntoGraf
SQWRLTab
SPARQL Query

Explore ontology with OWLViz and OntoGraf



Run the Hermit Reasoner

Reasoner Tools Refactor Wind

- Start reasoner ⌘R
- Synchronize reasoner
- Stop reasoner
- Explain inconsistent ontology
- Configure...

✓ Hermit 1.3.8.413
None

Description: 'Son of a Gun' ? [] = [] ×

Types +

- Movie ? @ × ○

Same Individual As +

Different Individuals +

Property assertions: 'Son of a Gun' [] =

- hasFilmStudio 'Entertainment One' ? @ ×
- hasActor ewan_mcgregor ? @ ×
- hasGenre Drama ? @ ×
- hasGenre Action ? @ ×
- hasMovieDirector julius_avery ? @ ×
- hasActor alicia_vikander ? @ ×
- hasWriter julius_avery ? @ ×
- hasCrew writer_3 ?
- hasCrew julius_avery ?

Explanation

- owl:topObjectProperty
 - hasActor
 - hasCrew
 - hasMovieDirector
 - hasWriter
 - hasFilmStudio
 - hasGenre

Add one movie instance

Individuals by type: 'The_Shining'

- Actor (5)
 - ewan_mcgregor
 - dan_stevens
 - emma_watson
 - hayley_atwell
 - alicia_vikander
- FilmStudio (3)
- Genre (5)
 - Horror
 - Famil
 - Drama
 - Animation
 - Romance
- Movie (4)
 - 'The_Shining'
 - 'Christopher Robin'
 - 'Beauty and the Beast'
 - 'Son of a Gun'

Description: 'The_Shining'

Property assertions: 'The_Shining'

Types +

- Movie

Same Individual As +

Different Individuals +

Object property assertions +

Data property assertions +

Negative object property assertions +

'The_Shining'

hasGenre Horror

(Tip: Use CTRL+Space to auto-complete names)

Cancel OK

Cardinality constraint

Description: Movie

Equivalent To +

hasMovieDirector exactly 1 MovieDirector

SubClass Of +

General class axioms +

SubClass Of (Anonymous) +

Instances +

- 'Beauty and the Beast'
- 'Christopher Reeve'
- 'Son of a Gun'
- 'The_Shining'

Target for Key +

Disjoint With +

Disjoint Union Of +

Reasoning

Movie

Class expression editor | Object restriction creator

Restricted property

- owl:topObjectProperty
 - foaf:knows
 - hasActor
 - hasCrew
 - hasMovieDirector**
 - hasWriter
 - hasFilmStudio
 - hasGenre
 - ownsMovie
 - prequel
 - sequel

Restriction filler

- owl:Thing
 - FilmStudio
 - foaf:Person
 - Actor
 - MovieDirector**
 - MovieLover
 - MovieOwner
 - Writer
 - Genre
 - Movie

Restriction type

Exactly (exact cardinality) | Cardinality 1

Cancel | OK

movie can have exactly one director

Ontology Editors in Practice

Emanuel Sallinger
Katja Hose

188.399 Introduction to Semantic Systems

(2.0 VU / 3.0 ECTS)

2023W

**Katja Hose
Emanuel Sallinger**

iss@dbai.tuwien.ac.at



- RDF and Linked Data
- RDFS
- Ontology Representation and OWL

The Semantic Web



- Expressing meaning
- Knowledge representation
- Ontologies
- Agents
- Evolution of knowledge

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](#)

The entertainment system was belting out the Beatles' "We Can Work It Out" when the phone rang. When Pete answered, his phone turned the sound down by sending a message to all the other *local* devices that had a *volume control*. His sister, Lucy, was on the line from the doctor's office: "Mom needs to see a specialist and then has to have a series of physical therapy sessions. Biweekly or something. I'm going to have my agent set up the appointments." Pete immediately agreed to share the chauffeuring. At the doctor's office, Lucy instructed her Semantic Web agent through her handheld Web browser. The agent promptly retrieved information about Mom's *prescribed treatment* from the doctor's agent, looked up several lists of *providers*, and checked for the ones *in-plan* for Mom's insurance within a *20-mile radius* of her *home* and with a *rating of excellent* or *very good* on trusted rating services. It then began trying to find a match between available *appointment times* (supplied by the agents of individual providers through their Web sites) and Pete's and Lucy's busy schedules. (The emphasized keywords indicate terms whose semantics, or meaning, were defined for the agent through the Semantic Web.)



In a few minutes the agent presented them with a plan. Pete didn't like it—University Hospital was all the way across town from Mom's place, and he'd be driving back in the middle of rush hour. He set his own agent to redo the search with stricter preferences about *location* and *time*. Lucy's agent, having *complete trust* in Pete's agent in the context of the present task, automatically assisted by supplying access certificates and shortcuts to the data it had already sorted through.

Killer app: Agents

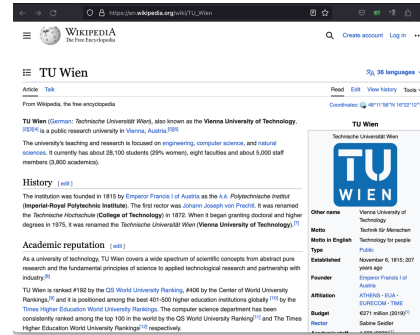
- Retrieve information about prescribed treatment/therapy
- Look for providers/specialists with good ratings in proximity
- Schedule times according to multiple busy schedules
- Share information with other agents and “negotiate”

“A really important thing about data is:
the more things you have to **connect**
together, the more powerful it is.”



The Web of Documents

TU Wien



Vienna



Austria

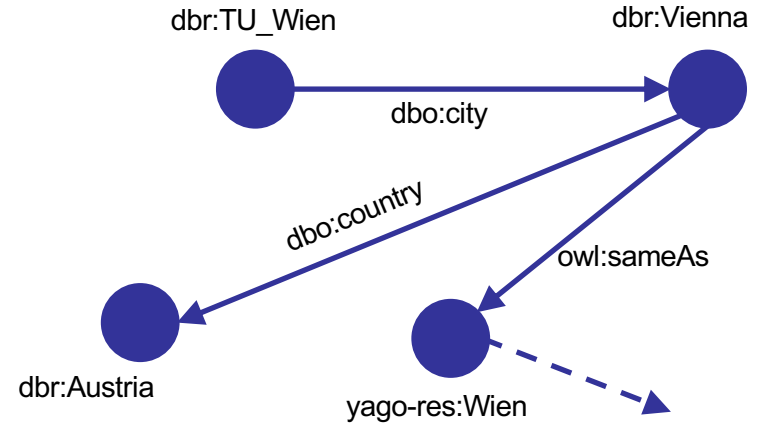


- Links between documents (websites, social media, etc.)
- Human-readable knowledge
- Normally accessed by humans via a Web browser
- Links are not qualified

The Web of Data

Linked Data


- Links between things/entities
- Machine-readable knowledge
- Normally accessed by machines
- Links are qualified
- A decentralized database



Web Technology Foundations



The Semantic Web is based on the Web



TCP/IP (Transfer Control Protocol/Internet Protocol):

protocols for data transfer via a network

HTTP (Hypertext Transfer Protocol):

- application protocol to transfer information (not just) on the WWW
- request/response protocol between client and server
- content negotiation (header lists acceptable media types etc.)

HTML (Hypertext Markup Language):

- language for the creation of web pages
- describes how it should be displayed by a web browser

XML (Extensible Markup Language):

- general purpose markup language for the creation of special markup languages
- used to facilitate data sharing among different systems

Basic principles

- “Everything has an identifier”
- “One name for one resource”

Uniform (different types of identifiers constructed according to a uniform schema)

Resource (whatever may be identified via URI)

(e.g., web pages, books, locations, persons, relations among objects, abstract concepts etc.)


Identifier (to distinguish one resource from another)

Examples:

```
ldap://[2001:db8::7]/c=GB?objectClass?one
mailto:John.Doe@example.com
news:comp.infosystems.www.servers.unix
tel:+1-816-555-1212
telnet://192.0.2.16:80/
urn:oasis:names:specification:docbook:dtd:xml:4.1.2
http://dbpedia.org/resource/Vienna
```

URI: Uniform Resource can be

- **URN (RFC 2141):** Uniform Resource Name
 - persistent identifier within a particular namespace
 - Refers to a resource without implying location or how to access it
 - e.g., urn:isbn:0-486-27557-4 refers to a specific edition of Shakespeare's Romeo and Juliet
- **URL (RFC 1738):**
 - Uniform Resource Locator
 - specifies (primary) access mechanism
 - may change during the lifetime of a resource
- **both URN+URL**

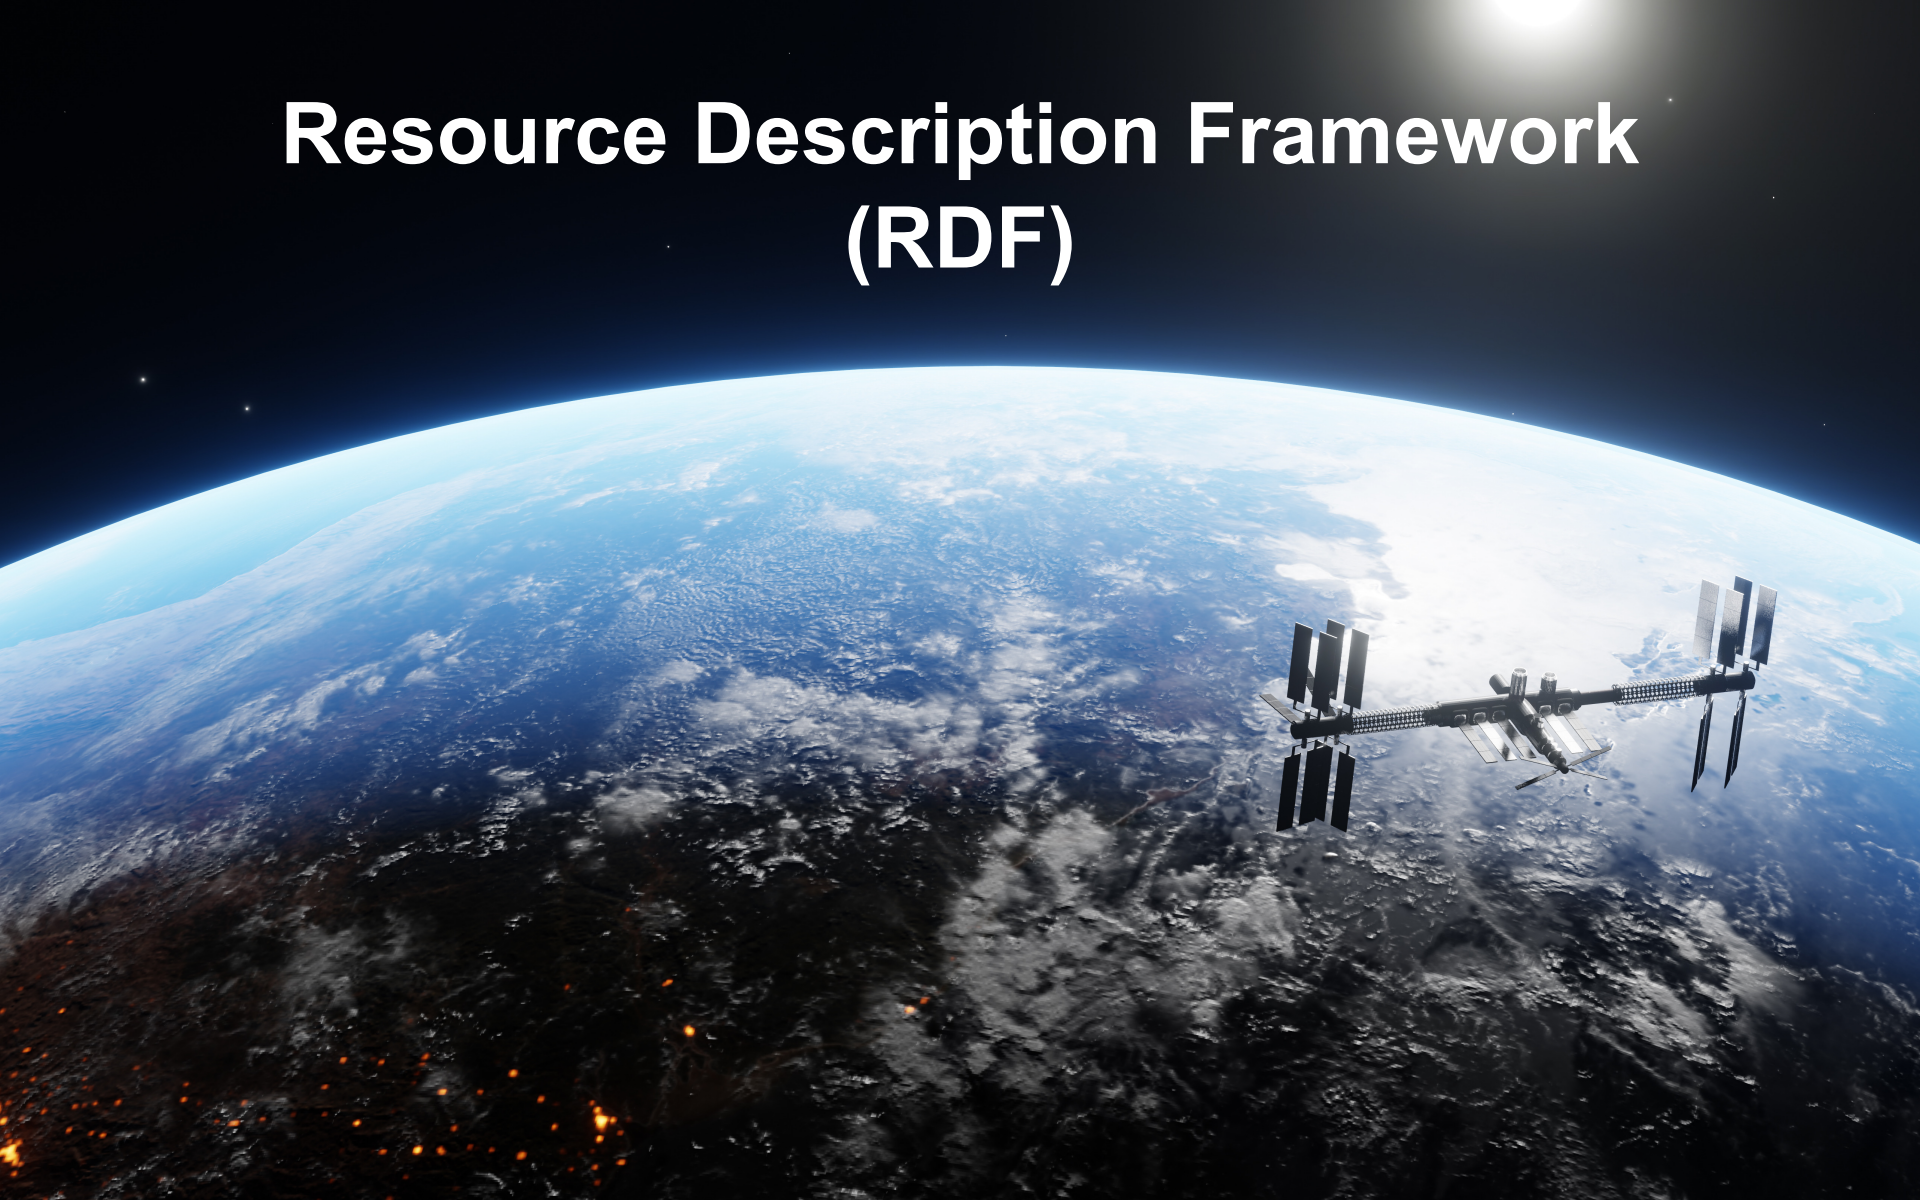


IRI (Internationalized Resource Identifier, RFC 3987) are URIs


- Character set of IRIs: Unicode (ISO 10646)
- Character set of URIs: US-ASCII

In the remainder of this presentation occurrences of the term URI can be replaced with IRI.

Resource Description Framework (RDF)



Standard information exchange is key



- **Resource Description Framework** (RDF, W3C standard)
- Basic concepts
 - A **resource** can basically be anything
 - e.g. persons, places, Web documents, abstract concepts
 - **Descriptions** of resources
 - Attributes
 - Relationships
 - The **framework** contains
 - A data model, and
 - Languages and syntaxes

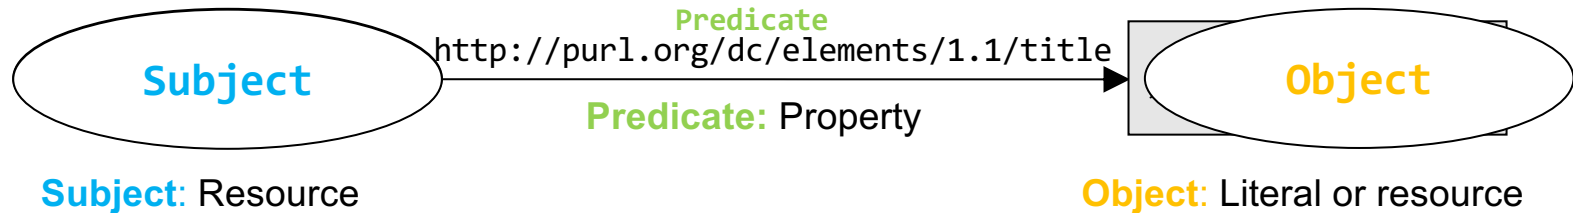
- Developed in the late 1990s
- Final W3C recommendation in 2004 (RDF 1.1)
- Graph-based model
- Formal semantics → machine-readable
- Knowledge expressed as a list of statements
- All statements follow the same simple schema (**RDF Triple**)
- Abstract model with various syntax notations and serialization formats:
 - Graph notation
 - RDF/XML (historically first serialization format)
 - N3/Turtle
 - N-Triples
 - JSON-LD






Basic RDF building block

- **Subject:** a resource, which may be identified with a URI
- **Predicate:** a URI-identified specification of the relationship between subject and object
- **Object:** a resource or literal to which the subject is related




Graph Notation Elements



Resource

Has a unique URI




Predicate

Connects two resources



Literal

- Only text, no URI
- Plain (lexical value and optional language tag) or typed

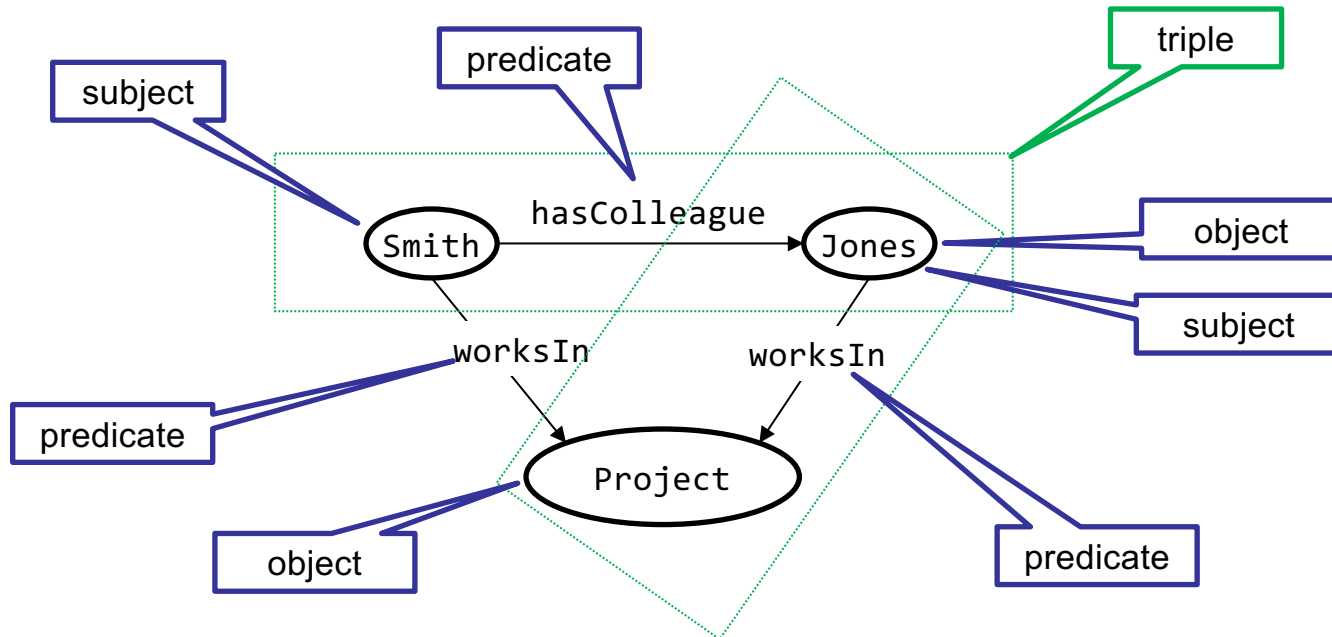


Blank Node


- No URI
- For unnamed objects

A set of RDF assertions in triples form a labeled directed graph with:

- **Resources:** the subjects and objects are nodes of the graph
- **Predicates:** each predicate becomes a label for an arc connecting the subject to the object.




! The object of one statement can be the subject of another.



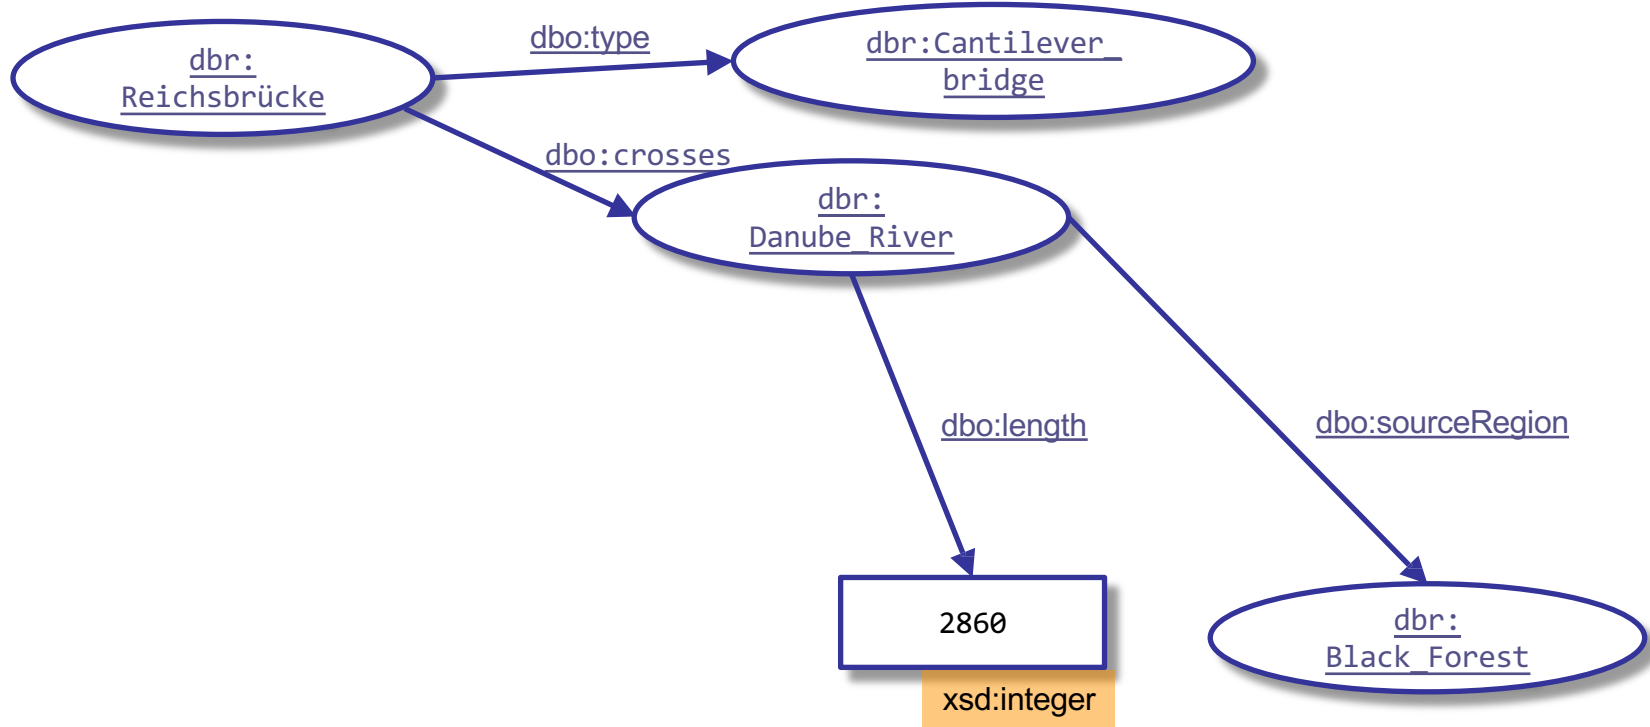
<http://dbpedia.org/resource/Reichsbrücke>

becomes

`dbr:Reichsbrücke`



Merging/Integrating Multiple Graphs



Plain literals have a **lexical form** (their lexical value) and optionally a **language tag**.

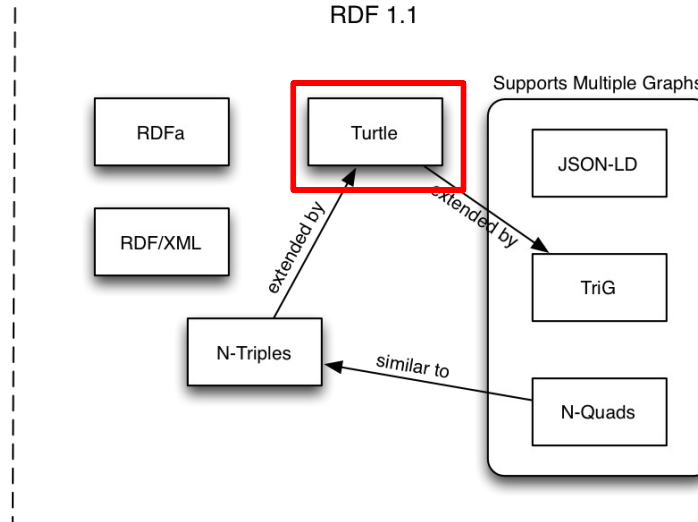
```
“27” , “Hello world”@en
```

- Formed by pairing a string with a URIref that identifies a particular datatype
- Indicates what datatype should be used to interpret a given literal
- RDF has no built-in set of datatypes of its own (except `rdf:XMLLiteral`)
- RDF datatype concepts are based on a conceptual framework from **XML Schema datatypes** that defines
 - the value space,
 - the lexical space and
 - the lexical-to-value mapping for a datatype (see RDF specifications)

```
"27"^^http://www.w3.org/2001/XMLSchema#integer
```


XML Schema datatypes

	Datatype
Core types	xsd:string
	xsd:boolean
	xsd:decimal
	xsd:integer
IEEE floating-point numbers	xsd:double
	xsd:float
Time and date	xsd:date
	xsd:time
	xsd:dateTime
	xsd:dateTimeStamp
Recurring and partial dates	xsd:gYear
	xsd:gMonth
	xsd:gDay
	xsd:gYearMonth
	xsd:gMonthDay
	xsd:duration
	xsd:yearMonthDuration
xsd:dayTimeDuration	
Limited-range integer numbers	xsd:byte
	xsd:short
	xsd:int
	xsd:long
	xsd:unsignedByte
	xsd:unsignedShort
	xsd:unsignedInt
	xsd:unsignedLong
	xsd:positiveInteger
	xsd:nonNegativeInteger
	xsd:negativeInteger
xsd:nonPositiveInteger	
Encoded binary data	xsd:hexBinary
	xsd:base64Binary
Miscellaneous XSD types	xsd:anyURI
	xsd:language
	xsd:normalizedString
	xsd:token
	xsd:NMTOKEN
	xsd:Name
xsd:NCName	



- Useful for inter-machine communication
- Every **Description** element describes a resource (referred to by an IRI)
- Every attribute or nested element inside a **Description** is a **property** of that Resource

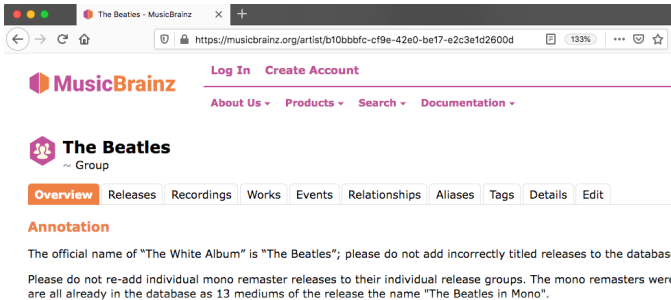
```
<rdf:Description rdf:about="http://musicbrainz.org/artist/b10bbbf-cf9e-42e0-be17-e2c3e1d2600d#_>
  <foaf:name>The Beatles</foaf:name>
  <owl:sameAs rdf:resource="http://dbpedia.org/resource/The_Beatles" />
</rdf:Description>
```

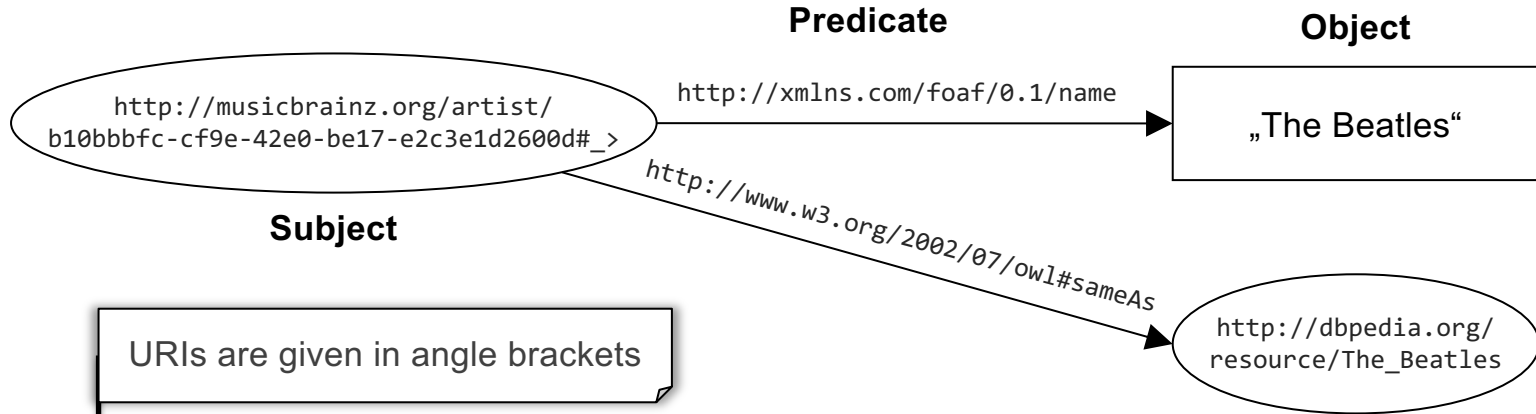
RDF/XML

Subject

Predicate

Object





URIs are given in angle brackets

Literals are given in quotes

```

<http://musicbrainz.org/artist/b10bbbfbc-cf9e-42e0-be17-e2c3e1d2600d# >
  <http://xmlns.com/foaf/0.1/name>
    "The Beatles".
<http://musicbrainz.org/artist/b10
  <http://www.w3.org/2002/07/owl#sameAs>
  <http://dbpedia.org/resource/The\_Beatles>.
  
```

Every statement is terminated with a full stop.

N-Triples

- Easy to read and write
- <subject> <predicate> <object>
- Every subject, predicate, object is identified with a URI
- Two shortcuts for several statements about the same subject
 - “;” introduce another predicate of the same subject
 - “,” introduce another object with the same predicate and subject

```

<#pat> <#child> <#al>, <#chaz>, <#mo>;
      <#age> 24;
      <#eyecolor> "blue".
  
```

N3

Interpretation:

Pat has the children al, chaz, mo;

Pat's age is 24;

Pat's eyecolor is blue

- "Terse RDF Triple Language"
- Subset of N3
- Many URIs share the same basis → use prefixes for namespace declarations

```
@prefix rdf:<http://www.w3.org/1999/02/22-rdf-syntax-ns#>.
@prefix rdfs:<http://www.w3.org/2000/01/rdf-schema#>.
@prefix owl:<http://www.w3.org/2002/07/owl#>.
@prefix mo:<http://purl.org/ontology/mo/>.
@prefix dbpedia:<http://dbpedia.org/resource/>.
```

Turtle

Note: You can find commonly used prefixes here: <http://prefix.cc>

- Unique base:

```
@base <http://musicbrainz.org/>.
```


Turtl

e

Has a simple shorthand for class membership (**a**):

```
@base <http://musicbrainz.org/>.
@prefix mo:<http://purl.org/ontology/mo/>.
<artist/b10bbbf-cf9e-42e0-be17-e2c3e1d2600d#_>
    a mo:MusicGroup.
```

Turtle



```
<http://musicbrainz.org/artist/b10bbbf-cf9e-42e0-be17-
e2c3e1d2600d#_>
    <http://www.w3.org/1999/02/22-rdf-syntax-ns#type>
    <http://purl.org/ontology/mo/MusicGroup>.
```

N-Triples


When multiple statements apply to **same subject (and predicate)** they can be abbreviated:

```
<artist/b10bbbfc-cf9e-42e0-be17-e2c3e1d2600d#_>
  rdfs:label "The Beatles" ;
  owl:sameAs dbpedia:The_Beatles,
  <http://www.bbc.co.uk/music/artists/
  b10bbbfc-cf9e-42e0-be17-e2c3e1d2600d#artist> .
```

Turtle

Same subject

Same subject and predicate



Turtle also provides a simple syntax for data types and language tags for literals:


```

<recording/5098d0a8-d3c3-424e-9367-1f2610724410#_>
  a mo:Signal;
  rdfs:label "All You Need Is Love" ;
  mo:duration "PT3M48S"^^xsd:duration .

dbpedia:The_Beatles dbpedia-owl:abstract
  "The Beatles were an English rock band formed (...) "@en,
  "The Beatles waren eine britische Rockband in den(...) "@de .
  
```

Turtle

Standard information exchange is key



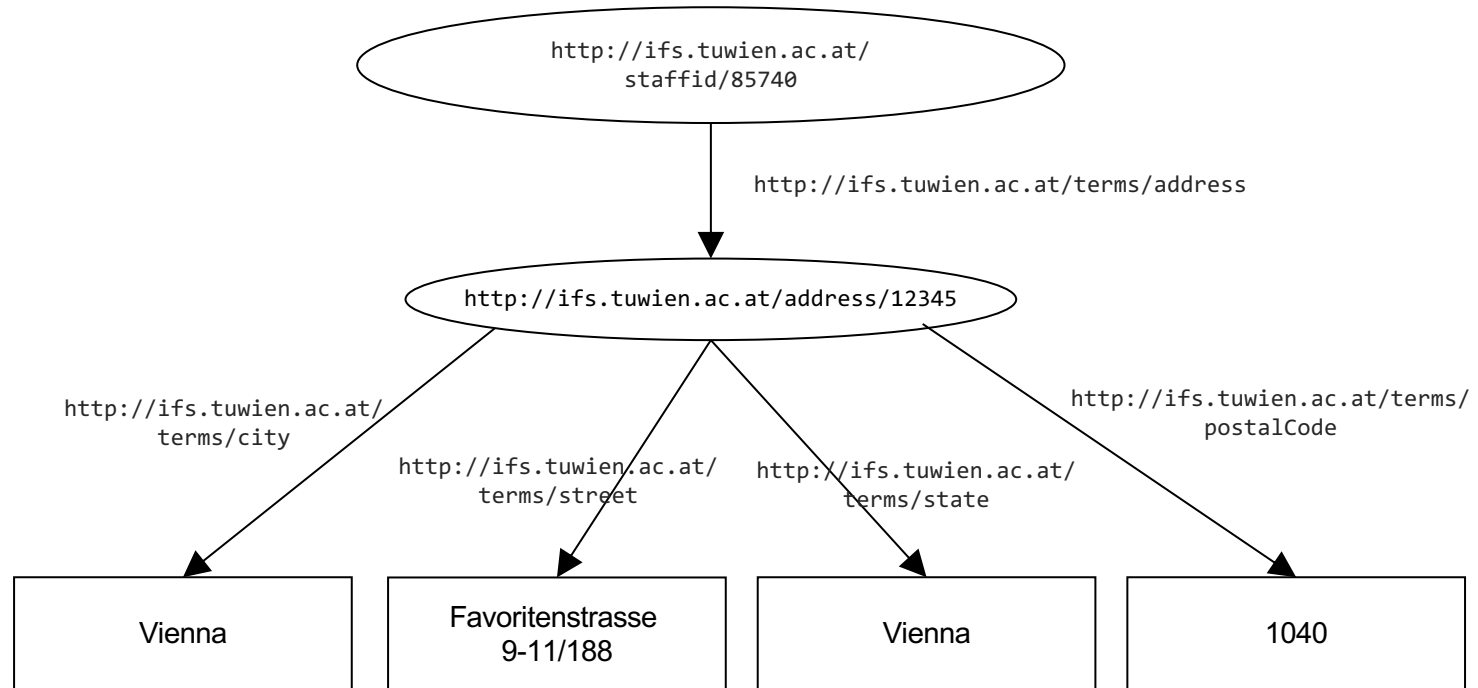
- Structured values
- Blank nodes
- Data structures (container, collection)
→ see appendix
- Reification

Consider the triple:

```
exstaff:85740
  exterms:address
    "Favoritenstrasse 9-11/188, 1040 Vienna" .
```

How do we represent address as a **structure** consisting of separate *street, city, state, and postal code* values? (N-ary relation)

→ Consider the **"aggregated thing" to be described** (like an address) **as a resource**, and then make statements about that new resource

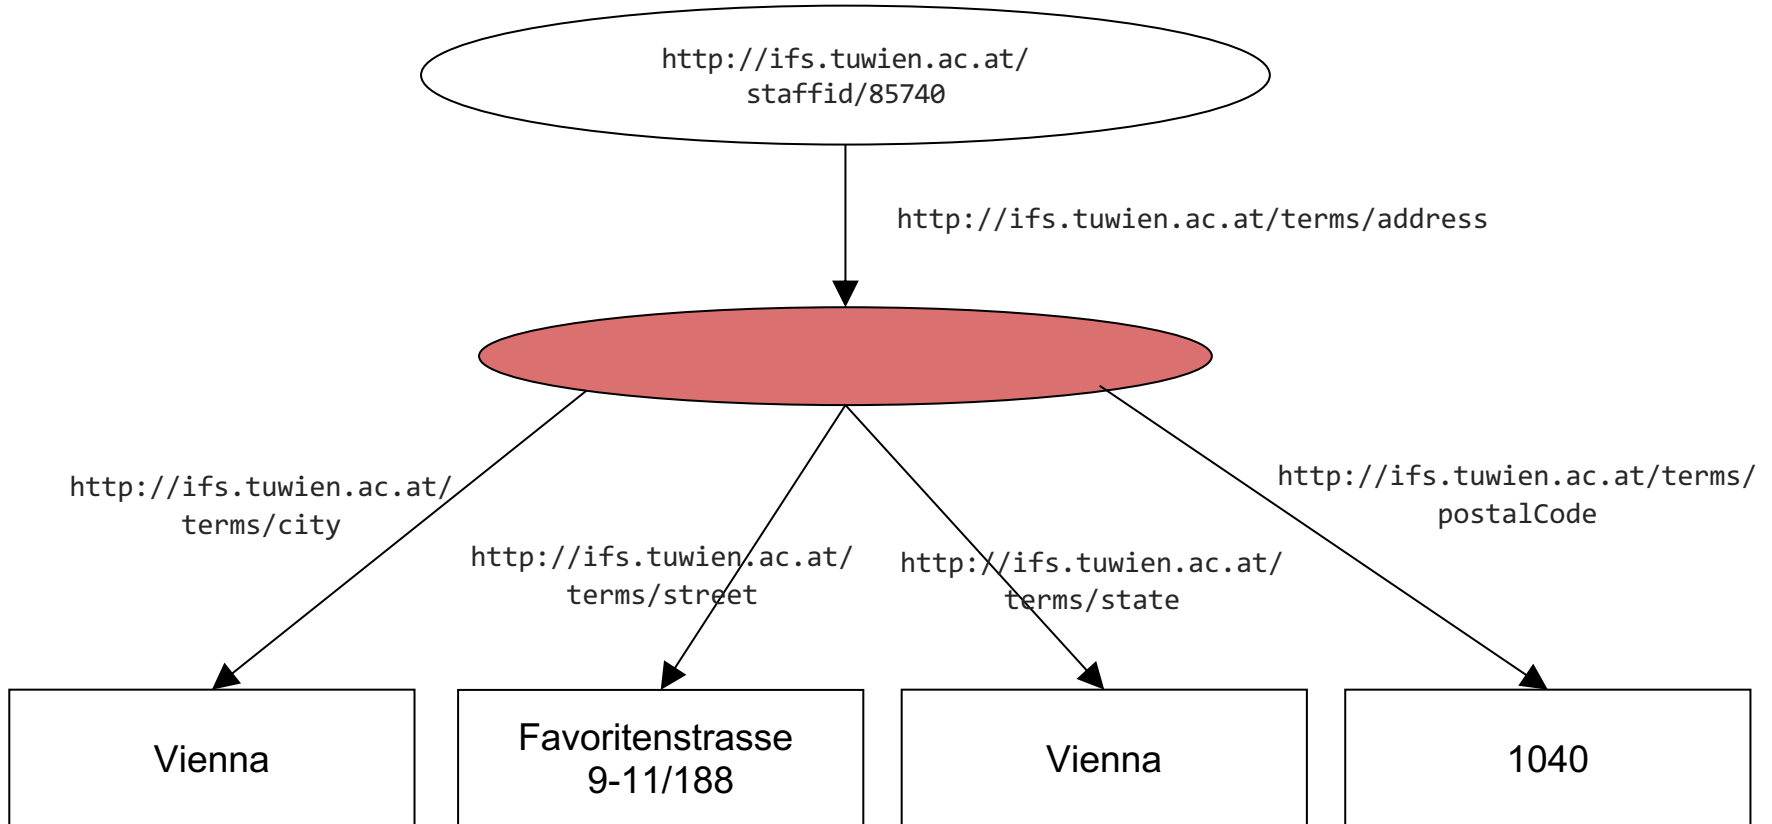


Or in triples notation:

```
exstaff:85740 exterms:address exaddressid:85740 .
exaddressid:85740 exterms:street "Favoritenstrasse 9-11/188" .
exaddressid:85740 exterms:city "Vienna" .
exaddressid:85740 exterms:state "Vienna" .
exaddressid:85740 exterms:postalCode "1040" .
```

```
exstaff:85740 exterms:address exaddressid:85740 .
exaddressid:85740 exterms:street "Favoritenstrasse 9-11/188" ;
                    exterms:city "Vienna" ;
                    exterms:state "Vienna" ;
                    exterms:postalCode "1040" .
```

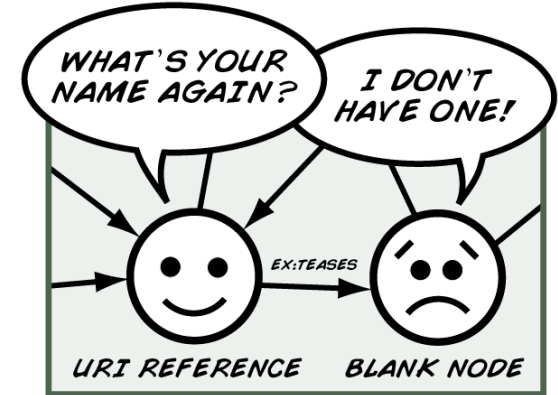
Turtle



- Represent unnamed resources
- Used widely, with varying intentions:
 - Make statements about resources that do not have URIs (but that are described in terms of relationships with other resources that do)
 - group related information
 - represent n-ary relationships

Issues:

- "Breaks" reasoning, because simple entailment in the presence of existentials is NP-complete [but: often not a big deal in practice*]
- Use in Linked Data discouraged [but: common in practice*]



<http://milicicvuk.com/blog/?p=4>

```
exstaff:85740 exterms:address ? .  
? exterms:street "1501 Grant Avenue" .  
? exterms:city "Bedford" .  
? exterms:state "Massachusetts" .  
? exterms:postalCode "01730" .
```

Turtle

Is this a valid notation for a blank node?
→ Menti


```
exstaff:85740  exterms:address _:johnaddress .  
_:johnaddress exterms:street "1501 Grant Avenue" .  
_:johnaddress exterms:city "Bedford" .  
_:johnaddress exterms:state "Massachusetts" .  
_:johnaddress exterms:postalCode "01730" .Turtle
```

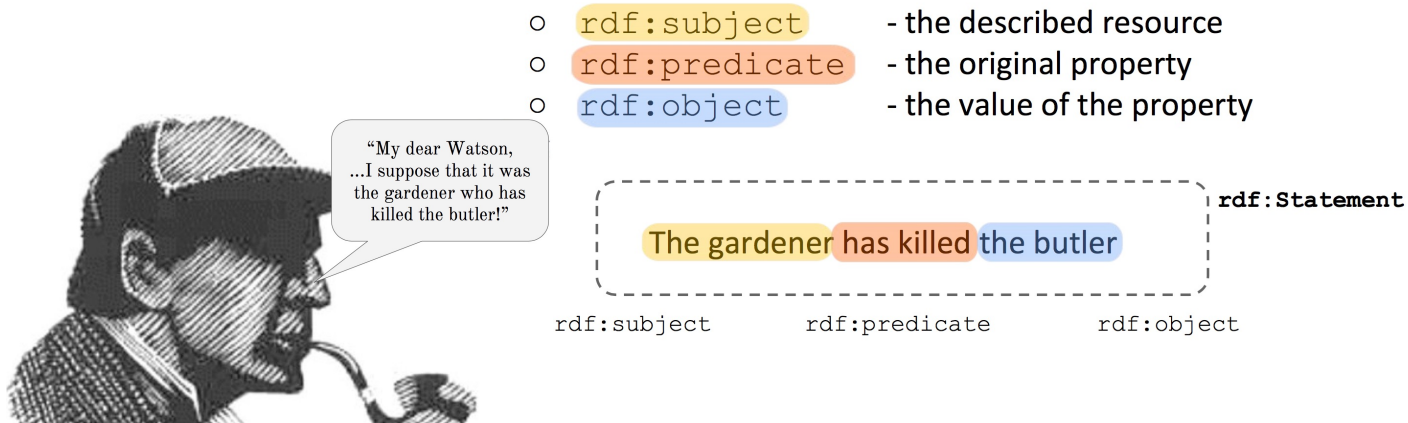
- Each distinct blank node must be given a different blank node identifier
- Blank node identifiers have significance only within a single graph
- Blank node identifiers may only appear as subjects or objects in triples (not predicates)

Used for making statements about statements to


- model data provenance
- formalize statements about reliability and trust
- define metadata about statements

Caveat:


- definition of infinite recursions and cycles



Sherlock Holmes supposes that the gardener has killed the butler



```
@prefix rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#> .
@prefix dbpedia <http://dbpedia.org/resource/> .
@prefix ex: <http://example.org/Crimestories#> .
```



```

@prefix dbpedia: <http://dbpedia.org/resource/> .
@prefix rdf:     <http://www.w3.org/1999/02/22-rdf-syntax-ns#> .
@prefix ex:      <http://example.org/Crimestories#> .


dbpedia:SherlockHolmes ex:supposes ex:StatementOfSherlock .
ex:StatementOfSherlock a rdf:Statement ;
  rdf:subject    ex:Gardener ;
  rdf:predicate ex:hasKilled ;
  rdf:object    ex:Butler .

```

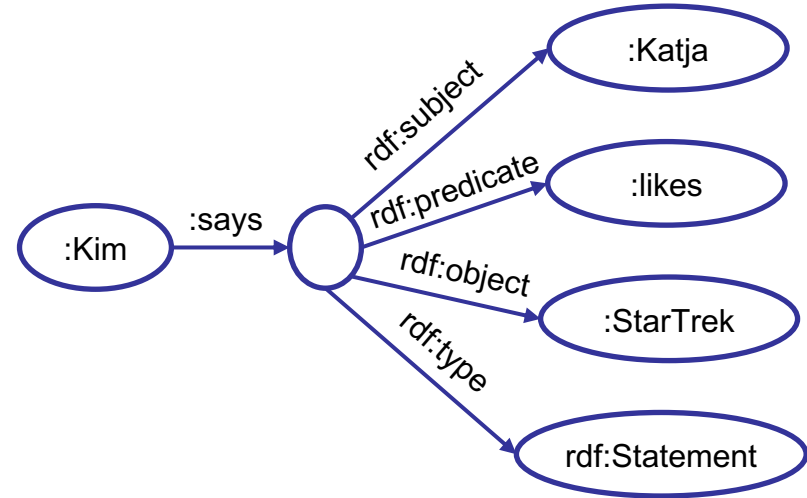
Turtle

Standard reification

Kim says



:Katja :likes :StarTrek




```

:Kim  :says      _:s1
_:s1  rdf:type   rdf:Statement
_:s1  rdf:subject :Katja
_:s1  rdf:predicate :likes
_:s1  rdf:object  :StarTrek
    
```

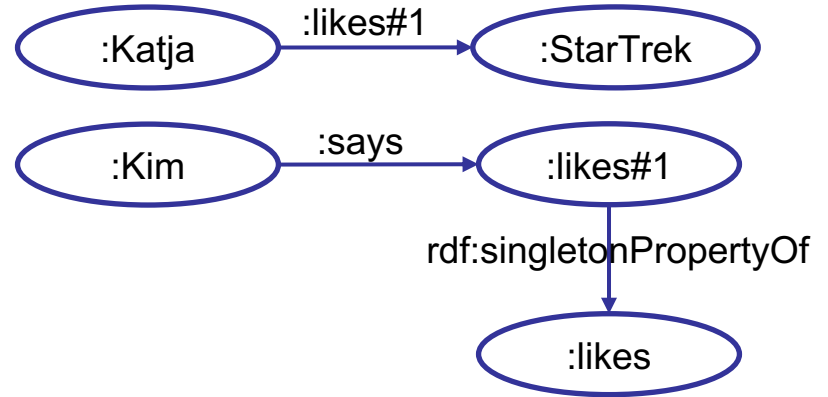
Pros: it works on any basic RDF/SPARQL engine

Cons: number of triples and query triple patterns is at least quadrupled

Singleton property



:Katja :likes :StarTrek




```
:Katja :likes#1 :StarTrek
:Kim :says :likes#1
:likes#1 rdf:singletonPropertyOf :likes
```


Pros: space efficient

Cons: it requires additional text indexes for predicates

Named graphs



:Katja :likes :StarTrek




:Katja :likes :StarTrek :graph#e
:Kim :says :graph#e :metadata

Pros: relatively space efficient


Cons: it requires efficient graph indexing – not always supported

RDF-star (RDF*)

Kim says



:Katja :likes :StarTrek




:Kim :says <<:Katja :likes :StarTrek>>


- Pros: nested triples treated as first-class citizens
- Cons: not supported by all systems

RDF-star (RDF*)

Kim says



:Katja :likes :StarTrek



:Kim :says <<:Katja :likes :StarTrek>>

W3C recommendation – not yet a standard

Currently not supported by all triple stores.

Supported (to a certain degree) by GraphDB, Stardog, Jena, Oxigraph

Details:

G. Abuoda, C. Aebeloe, D. Dell'Aglio, A. Keen, K. Hose:

StarBench: Benchmarking RDF-star Triplestores

In: QuWeDa icw. ISWC 2023

Classes


- **rdf:XMLLiteral** - the class of XML literal values
- **rdf:Property** - the class of properties
- **rdf:Statement** - the class of RDF statements
- **rdf:Alt**, **rdf:Bag**, **rdf:Seq** - containers of alternatives, unordered containers, and ordered containers*
- **rdf>List** - the class of RDF Lists
- **rdf:nil** - an instance of **rdf>List** representing the empty list

Properties


- **rdf:type** - an instance of **rdf:Property** used to state that a resource is an instance of a class
- **rdf:first** - the first item in the subject RDF list
- **rdf:rest** - the rest of the subject RDF list after **rdf:first**
- **rdf:value** - idiomatic property used for structured values
- **rdf:subject** - the subject of the subject RDF statement
- **rdf:predicate** - the predicate of the subject RDF statement
- **rdf:object** - the object of the subject RDF statement

* **rdfs:Container** is a super-class of the three

RDF Schema



Specifying conceptual models:
specifying **lightweight** ontologies with RDFS.



- On the Web “Anyone can say Anything about Any topic” – **AAA Slogan**
- **Non-unique Naming Assumption**
 - The same entity could be known by more than one name
 - E.g., PersonA can be the same instance as PersonB
 - Disjointness needs to be specified explicitly
- **Open World Assumption**
 - “Missing information is not evaluated as negative information!”
 - Assumes incomplete information by default, in this way, it allows for reasoning on incomplete models.
 - If fact f in a knowledge base is missing, $\neg f$ cannot be inferred
 - E.g., likes(PersonA, DrinkB)
 - CWA: likes(PersonA, DrinkC) – No
 - OWA: likes(PersonA, DrinkC) – don’t know; PersonA may also like other drinks...

- First version in April 1998, W3C recommendation in Feb. 2004
- Every RDF Schema document is an RDF document

Purpose:


1. Nominate

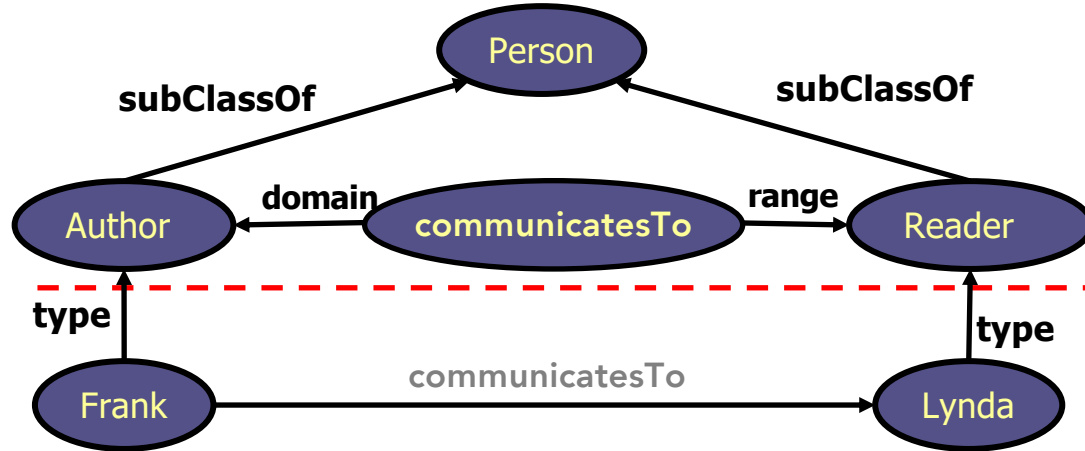
- the ‘types’ (i.e., classes) of things we might make assertions about, and
- the properties we might apply, as predicates in these assertions, to capture their relationships

2. Inference

Given a set of assertions, using these classes and properties, specify what should be inferred about assertions that are implicitly made

- Defines the **vocabulary** for RDF
- Organizes this vocabulary in a **typed hierarchy**
 - class, subClassOf, type
 - property, subPropertyOf
 - domain, range





```

<Person, rdf:type, rdfs:Class>
<Reader, rdfs:subClassOf, Person>
<Author, rdfs:subClassOf, Person>
<communicatesTo, rdf:type, rdfs:Property>
<communicatesTo, rdfs:domain, Author>
<communicatesTo, rdfs:range, Reader>
<Frank, communicatesTo, Lynda>

```

➔ `<Frank, rdf:type, Author>`
`<Lynda, rdf:type, Reader>`

rdfs:domain

- Restricts the scope of a property

```
mo:member rdfs:domain mo:MusicGroup .      Turtle
```

rdfs:range

- Used to restrict the type of values of a property

```
mo:member rdfs:range foaf:Agent .      Turtle
```

RDF and RDFS Vocabularies

RDF Classes:

`rdf:Property`, `rdf:Statement`,
`rdf:XMLLiteral`
`rdf:Seq`, `rdf:Bag`, `rdf:Alt`,
`rdf>List`

RDF Properties:

`rdf:type`, `rdf:subject`,
`rdf:predicate`, `rdf:object`,
`rdf:first`, `rdf:rest`, `rdf:_n`
`rdf:value`

RDF Resources:

`rdf:nil`

RDFS Classes

`rdfs:Resource`
`rdfs:Class`
`rdfs:Literal`
`rdfs:Datatype`
`rdfs:Container`
`rdfs:ContainerMembershipProperty`

RDFS Properties

`rdfs:domain`
`rdfs:range`
`rdfs:subPropertyOf`
`rdfs:subClassOf`
`rdfs:member`
`rdfs:seeAlso`
`rdfs:isDefinedBy`
`rdfs:comment`
`rdfs:label`

Classes:


- **rdfs:Resource** all resources within RDF are implicitly members of this class
- **rdfs:Class** declares a resource as a class (type or category) for other resources; instantiation of classes with `rdf:type`
- **rdfs:Literal** literal values such as strings and integers (plain or typed)
- **rdfs:Datatype** class of datatypes (both an instance of and a subclass of `rdfs:Class`; each instance of `rdfs:Datatype` is a subclass of `rdfs:Literal`)

Properties:

- **rdfs:subClassOf** allows to declare hierarchies of classes; properties of super class are inherited
- **rdfs:subPropertyOf** all resources related by one property are also related by another
- **rdfs:domain** defines the scope of a property (see slide)
- **rdfs:range** defines the range of values of a property (see slide)
- **rdfs:comment** used to provide a human-readable description of a resource
- **rdfs:seeAlso** used to indicate a resource that might provide additional information
- **rdfs:isDefinedBy** indicates a resource defining the subject source

- Allows basic inference
- Poor expressivity
- Does not allow complex expressions besides membership

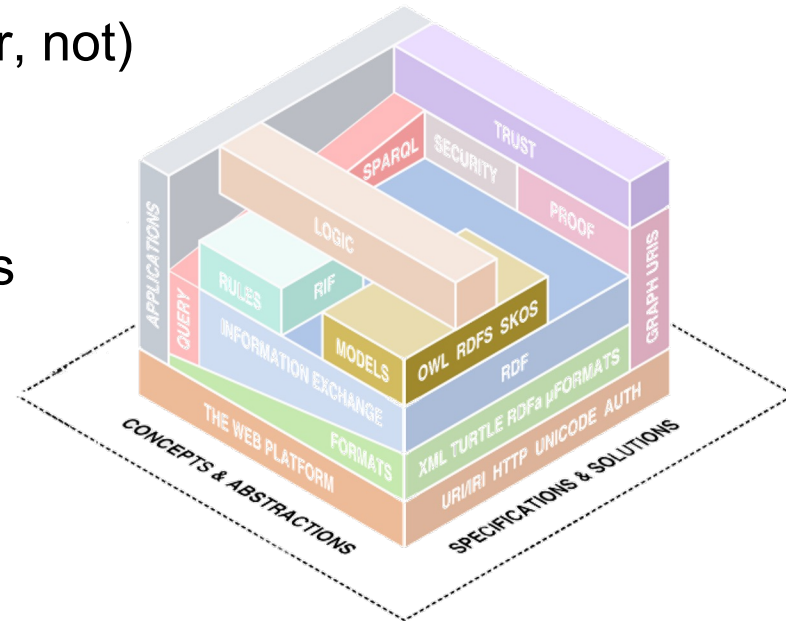
Specifying conceptual models:
specifying **heavyweight** ontologies with OWL.



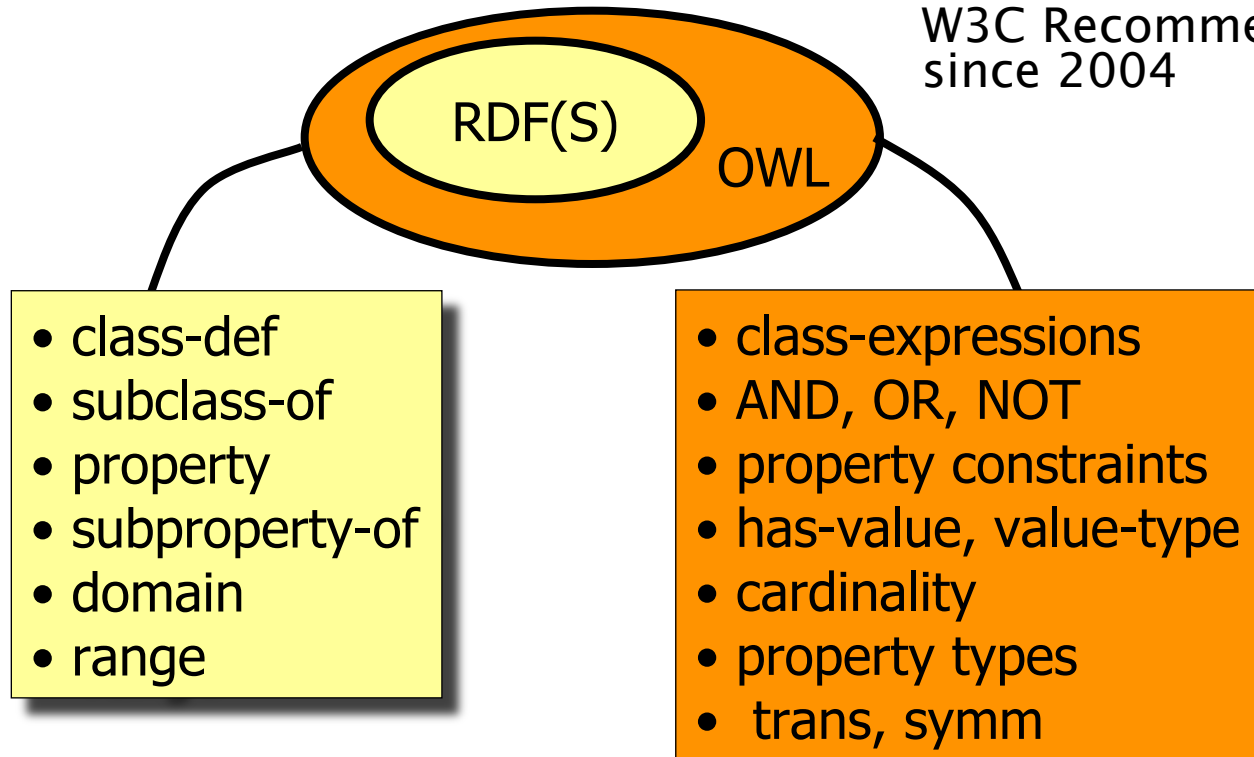
- W3C Recommendation for the representation of ontologies
- Fully-fledged knowledge representation language for the Web
- Provides a wider range of ontological constructs and avoids some of the potential confusion in RDFS
- Formal foundations: Description Logics (DL)
 - well-defined semantics
 - Well-understood properties (e.g., computational complexity)
 - allows to verify consistency of defined knowledge
 - inferencing allows to make implicit knowledge explicit
 - many tools available
- OWL ontologies can be exchanged as RDF documents

OWL extends RDF Schema to a fully-fledged knowledge representation language for the Web

- Logical expressions (and, or, not)
- (In)equality
- Local properties
- Required/optional properties
- Required values
- Enumerated classes
- Symmetry, inverse



W3C Recommendation
since 2004



- Aka. “type”, “concept”, “category”
- Sets of individuals belong to a class
- Class membership is dependent on logical description, not name
- There exist two predefined classes
 - `owl:Thing` (class that contains all individuals) \top
 - `owl:Nothing` (empty class) \perp

Definition of a class:

```
:Book a owl:Class .
```

- Definition of individuals via class membership

```
:NineteenEightyfour a :Book .
```

```
Book(NineteenEightyfour)
```

- Individuals may be (and are likely to be) a member of multiple classes
- Individuals can also be defined without direct class membership as named entity

```
:FrankSmith a owl:NamedIndividual .
```

- Used to relate individuals, binary in OWL
- Two property variants:

1. ObjectProperties

```
:author a owl:ObjectProperty .
```

2. DatatypeProperties

```
:publicationYear a owl:DatatypeProperty .
```

Domain and Range of object properties

```
:author a owl:ObjectProperty ;  
    rdfs:domain          :Book ;  
    rdfs:range :Writer .
```

Domain: $\exists \text{author.T} \sqsubseteq \text{Book}$

Range: $\text{T} \sqsubseteq \forall \text{author.Writer}$

Domain and range of datatype properties

```
:publicationYear a owl:DatatypeProperty ;  
    rdfs:domain      :Book ;  
    rdfs:range xsd:integer .
```

```
∃publicationYear.T ⊆ Book  
T ⊆ ∀publicationYear.Integer
```

Many XML datatypes can be used

```
:Book a owl:Class .
:Writer a owl:Class .
:GeorgeOrwell a :Writer .

:author a owl:ObjectProperty ;
        rdfs:domain :Book;
        rdfs:range  :Writer .

:publicationYear a owl:DatatypeProperty ;
        rdfs:domain :Book;
        rdfs:range  xsd:integer .

:NineteenEightyFour a :Book ;
        :author   :GeorgeOrwell ;
        :publicationYear 1948 .
```

Tbox and **Abox** do not have to be separated in OWL!

```
:Novel a owl:Class ;  
      rdfs:subClassOf :Book .  
  
:Book a owl:Class ;  
      rdfs:subClassOf :Work .  
  
:Work a owl:Class .
```

Novel \sqsubseteq Book
Book \sqsubseteq Work

(Via inference, it can be entailed that Novel is
also a subclass of Work)

```
:Book a owl:Class .
:Writer a owl:Class .
:Novel a owl:Class ;
        rdfs:subClassOf :Book .
:Poet a owl:Class ;
        rdfs:subClassOf :Writer .
:Book owl:disjointWith :Writer.
```

Novel \sqsubseteq Book
Poet \sqsubseteq Writer
Book \sqcap Writer $\sqsubseteq \perp$

(Via inference, it can be entailed that Novel and Poet are also disjoint classes)


```
:Writer a owl:Class .  
:Author a owl:Class .  
:Poet a owl:Class ;  
      rdfs:subClassOf :Writer .  
  
:Writer owl:equivalentClass :Author .
```

Poet \sqsubseteq Writer
Writer \equiv Author

(Via inference it can be entailed that Poet is also an Author)

```
:NineteenEightyfour a :Novel ;  
    :author :GeorgeOrwell ;  
    :publicationYear 1948;  
    owl:sameAs :ARX012345 .  
:Novel a owl:Class ;  
    rdfs:subClassOf :Book .  
:Book a owl:Class.
```

Via inference it can be entailed that ARX012345 is a Book

Differences of Individuals via owl:differentFrom:

```
:ARX012345 a :Novel ;  
    owl:differentFrom :ARX0123456 .
```

Overview of the basic constructs:

1. Classes:

```
:Book a owl:Class .  
:Novel a owl:Class ;  
    rdfs:subClassOf :Book  
:Book owl:disjointWith :Writer .  
:Writer owl:equivalentClass :Author .
```

2. Properties:

```
:author a owl:ObjectProperty;  
    rdfs:domain :Book;  
    rdfs:range :Writer .  
  
:publicationYear a owl:DatatypeProperty;  
    rdfs:domain : Book;  
    rdfs:range xsd:integer .
```

3. Individuals:

```
:NineteenEightyfour a :Book .  
:NineteenEightyfour owl:sameAs :ARX012345 .
```

Overview of the basic constructs:

1. Classes:

```
:Book a owl:Class .
:Novel a owl:Class ;
    rdfs:subClassOf :Book
:Book owl:disjointWith :Writer .
```

Further Reading:
 OWL Complex Classes in Appendix 2

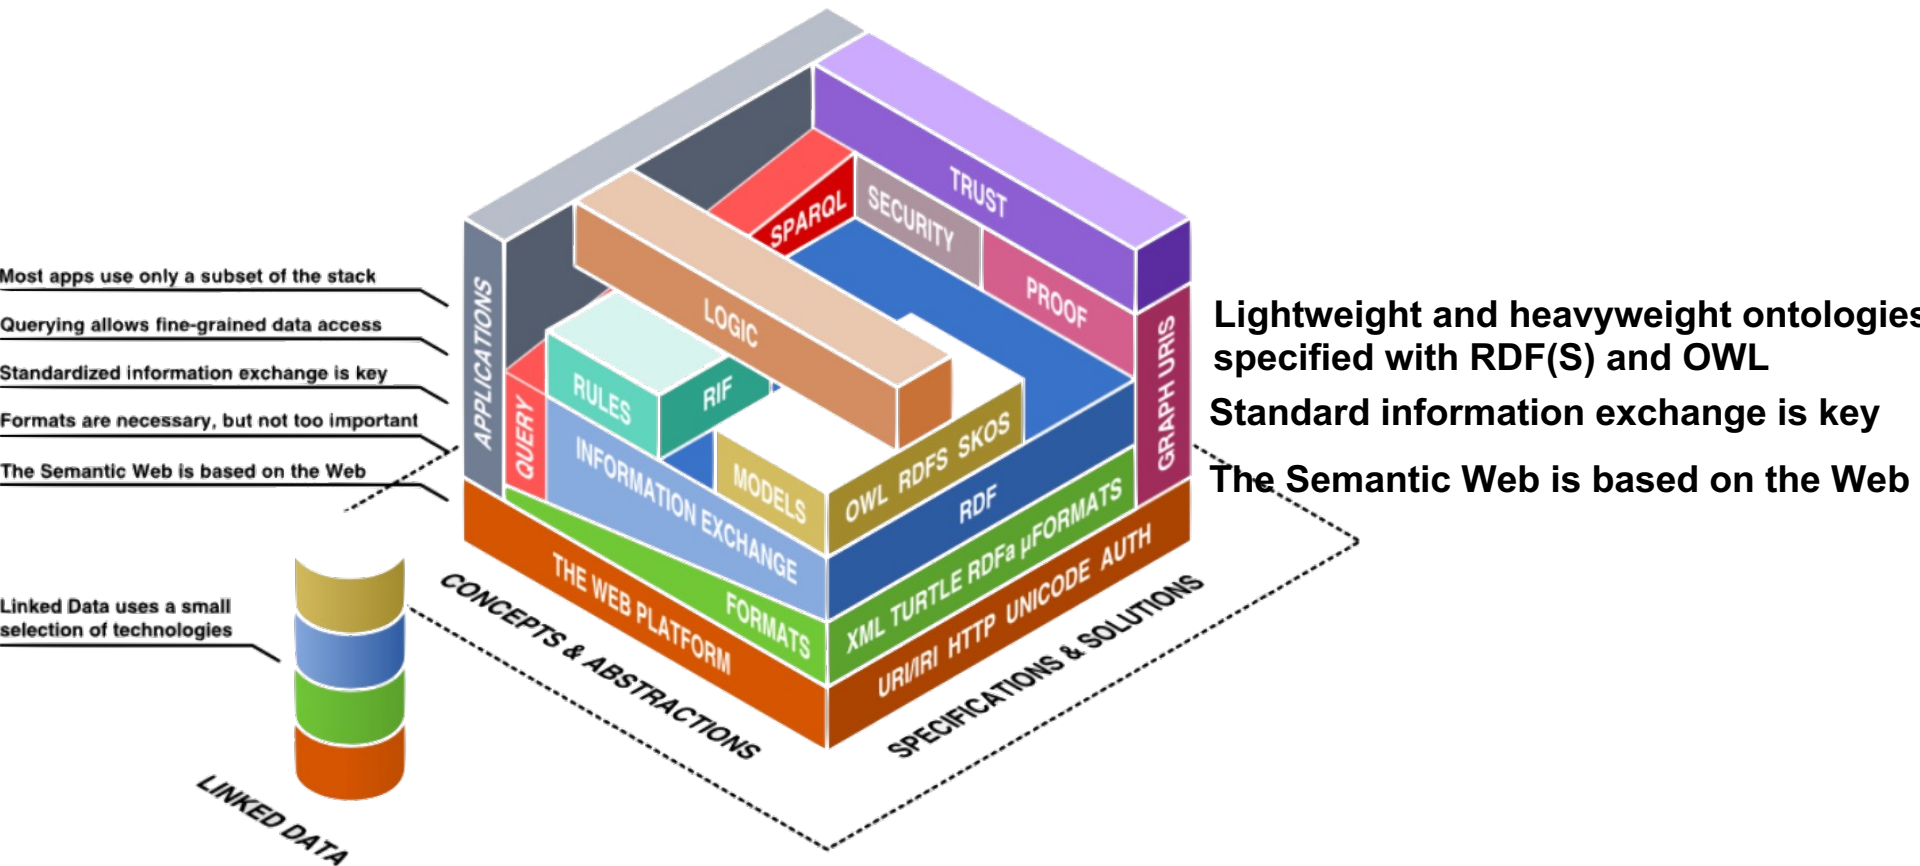
2. Prop

```
.publicationYear a owl:DatatypeProperty ;
    rdfs:domain : Book;
    rdfs:range xsd:integer .
```

3. Individuals:

```
:NineteenEightyfour a :Book .
:NineteenEightyfour owl:sameAs :ARX012345 .
```

Summary: The Semantic Web Technology Stack



Every layer can access the functionality of the layers below

- RDF and Linked Data
- RDFS
- Ontology Representation and OWL

Read Appendix 2 in preparation for the assignment!

Appendix 1

RDF



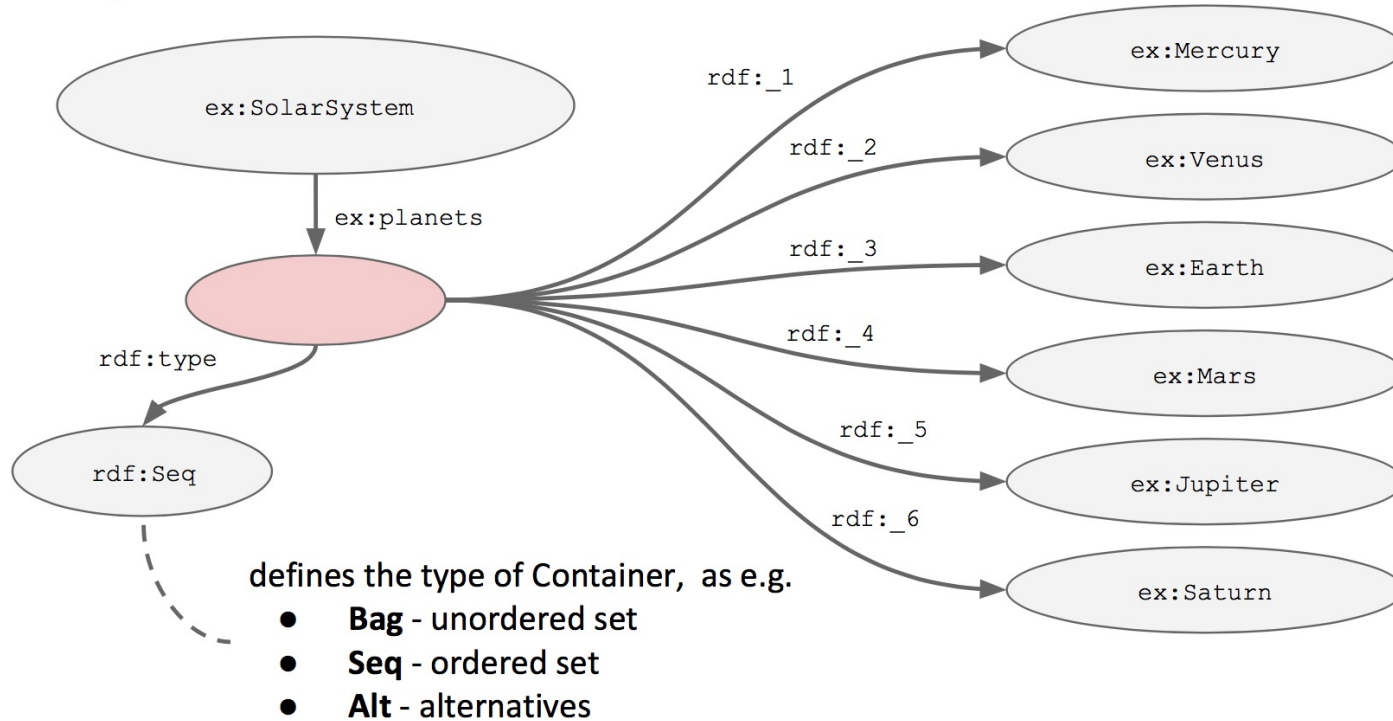
```
exstaff:85740  exterms:address _:johnaddress .
_:johnaddress  exterms:street "1501 Grant Avenue" .
_:johnaddress  exterms:city "Bedford" .
_:johnaddress  exterms:state "Massachusetts" .
_:johnaddress  exterms:postalCode "01730" .Turtle
```

- Each distinct blank node must be given a different blank node identifier
- Blank node identifiers have significance only within a single graph
- Blank node identifiers may only appear as subjects or objects in triples (not predicates)

- General Data structure to enumerate resources or literals
- Only shortcuts without additional semantic expressivity (syntactic sugar)

Distinguish:

- **Container:**
open list, i.e. extension (new entries) possible
- **Collections:**
closed list, i.e. no extension possible



```

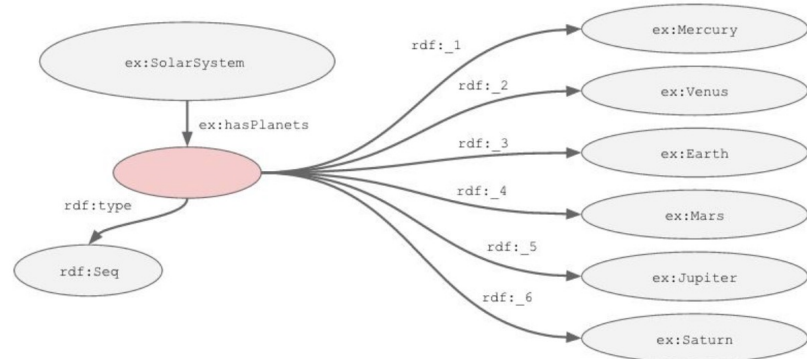
@prefix rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#> .
@prefix ex: <http://example.org/test#> .

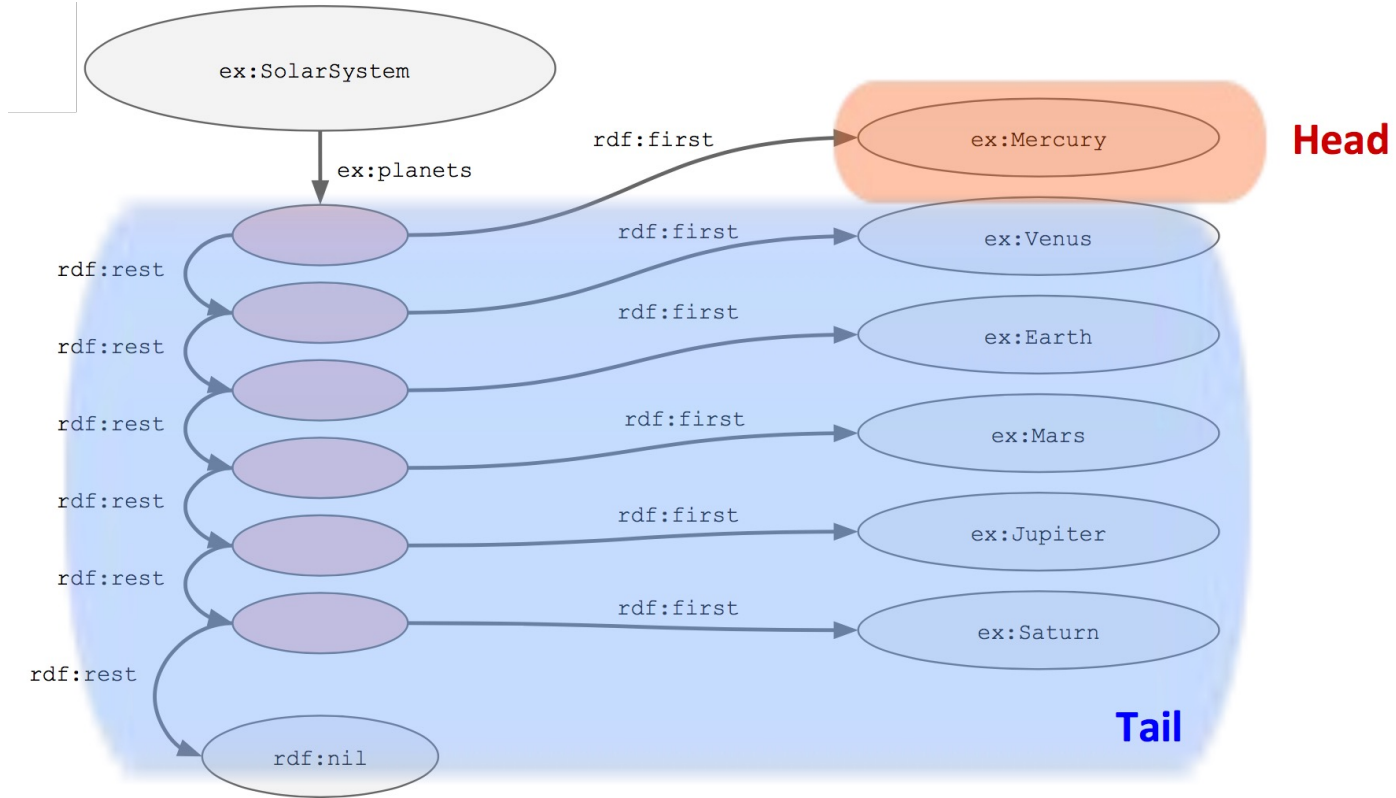
ex:SolarSystem ex:planets [
  a rdf:Seq ;
  rdf:_1 ex:Mercury ;
  rdf:_2 ex:Venus ;
  rdf:_3 ex:Earth ;
  rdf:_4 ex:Mars ;
  rdf:_5 ex:Jupiter ;
  rdf:_6 ex:Saturn
] .

```

Turtle


= **rdf:type** rdf:Seq





```
@prefix rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#> .
@prefix ex: <http://example.org/test#> .

ex:SolarSystem ex:planets [
  rdf:first ex:Mercury ; rdf:rest [
    rdf:first ex:Venus ; rdf:rest [
      rdf:first ex:Earth ; rdf:rest [
        rdf:first ex:Mars ; rdf:rest [
          rdf:first ex:Jupiter ; rdf:rest [
            rdf:first ex:Saturn ;
            rdf:rest rdf:nil
          ]
        ]
      ]
    ]
  ]
]
]
]
]
] .
```



```
@prefix rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#> .
@prefix ex: <http://example.org/test#> .

ex:SolarSystem ex:planets (
  ex:Mercury ex:Venus ex:Earth ex:Mars ex:Jupiter ex:Saturn
) .
```

Turtle

Turtle provides a nice syntax for RDF collections: simply put elements into parentheses

Appendix 2

OWL Complex Classes and Property Restrictions



```
:Novel a owl:Class ;  
:AnimalFarm a :Novel .  
:NineteenEightyfour a :Novel .  
:NovelsInStore a owl:Class ;  
    owl:oneOf  
    (:AnimalFarm  
     :NineteenEightyfour) .
```

`NovelsInStore \sqsubseteq {AnimalFarm, NineteenEightyfour}`

This says that there are only two novels available
in the store.

- Are used to describe complex classes via properties

- Restrictions on values:
 - owl:hasValue
 - owl:allValuesFrom
 - owl:someValuesFrom

- Restrictions on cardinality
 - owl:cardinality
 - owl:minCardinality
 - owl:maxCardinality


```
:Writer a owl:Class .  
:Book a owl:Class .  
:GeorgeOrwell a :Person .  
:author a owl:ObjectProperty ;  
  rdfs:domain :Book ;
```

```
:OrwellsBooks a owl:Class ;  
  rdfs:subClassOf  
    [ a owl:Restriction ;  
      owl:onProperty :author ;  
      owl:hasValue :GeorgeOrwell  
    ] .
```

OrwellsBooks \sqsubseteq author.(GeorgeOrwell)

- Class OrwellsBooks is described via fixed value assignment (=constant) of the individual GeorgeOrwell to the property author

Fixes all instances of a specific class C as allowed range for a property P \rightarrow (Universal Quantification)
 $\forall P.C$

```

:Poetry a owl:Class ;
      rdfs:subClassOf
        [ a owl:Restriction ;
          owl:onProperty :author ;
          owl:allValuesFrom :Poet
        ] .
  
```

Poetry \sqsubseteq \forall author.Poet

```

:Poet a owl:Class .
      rdfs:subClassOf :Writer .
:Poetry a owl:Class rdfs:subClassOf :Book .
:author a owl:ObjectProperty ;
      rdfs:domain :Book ;
      rdfs:range :Writer .
  
```

Universal restriction over an object property expression: contains those individuals that are connected through an object property expression **only** to instances of a given class expression.

- for a given property, universal restrictions *do not specify the existence* of a relationship. They merely state that if a relationship exists for the property then it must be to individuals that are members of a specific class.
- may be denoted by the universal quantifier (\forall). In OWL `owl:allValuesFrom`.
- It can be read as **only** or **no other than**
 - `VegetarianPizza hasTopping only VegetarianTopings`
 - `VegetarianPizza hasTopping no other than VegetarianTopings`

```
<owl:Class rdf:about="VegetarianPizza">
  <rdfs:subClassOf>
    <owl:Restriction>
      <owl:onProperty rdf:resource="hasTopping"/>
      <owl:allValuesFrom rdf:resource="VegetarianTopping"/>
    </owl:Restriction>
  </rdfs:subClassOf>
</owl:Class>
```

How about an individual that has no “hasTopping” relation? Will it be also considered a “Vegetarian” Pizza?

“trivial satisfaction of universal restriction”

Describes that there **must exist** an individual for property P and fixes its range to class C →
(existential quantification) $\exists P.C$

```

:Reader a owl:Class ;
    rdfs:subClassOf
        [ a owl:Restriction ;
          owl:onProperty :reads ;
          owl:someValuesFrom :Book
        ] .
  
```

Reader $\sqsubseteq \exists \text{reads}.\text{Book}$

```

:Reader a owl:Class .
:Book a owl:Class .
:reads a owl:ObjectProperty ;
    rdfs:domain :Person ;
    rdfs:range :ReadingMaterial .
  
```

An existential restriction describes a class of individuals that have *at least one (some)* relationship along a specified property to an individual that is a member of a specified class.

- To be a member of the class, the individual *must have at least one relation* along the given property to individuals that are members of a specific class (necessary condition).
- may be denoted by the existential quantifier (\exists). In OWL `owl:someValuesFrom`.
- It can be read as **some** or **at least one**
 - `Pizza hasBase some PizzaBase`
 - `Pizza hasBase at least one PizzaBase`

```
<owl:Class rdf:about="Pizza">
  <rdfs:subClassOf>
    <owl:Restriction>
      <owl:onProperty rdf:resource="hasBase"/>
      <owl:someValuesFrom rdf:resource="PizzaBase"/>
    </owl:Restriction>
  </rdfs:subClassOf>
</owl:Class>
```

A common pattern is combining universal and existential restrictions to define classes that have *at least one* relation of a property and all values of the property range over individuals of a certain class.

```
<owl:Class rdf:about="VegetarianPizza">
  <rdfs:subClassOf>
    <owl:Restriction>
      <owl:onProperty rdf:resource="hasTopping"/>
      <owl:someValuesFrom rdf:resource="VegetarianTopping"/>
    </owl:Restriction>
  </rdfs:subClassOf>
  <rdfs:subClassOf>
    <owl:Restriction>
      <owl:onProperty rdf:resource="hasTopping"/>
      <owl:allValuesFrom rdf:resource="VegetarianTopping"/>
    </owl:Restriction>
  </rdfs:subClassOf>
</owl:Class>
```

- `owl:cardinality` restricts to an exact number
- `owl:minCardinality`, `owl:maxCardinality` restricts to upper / lower bounds

```

:Tetralogy a owl:Class ;
  rdfs:subClassOf
    [ a owl:Restriction ;
      owl:onProperty :hasVolumes ;
      owl:cardinality 4
    ] .

```

Tetralogy \sqsubseteq (=4)hasVolumes

```

:hasVolumes a owl:DatatypeProperty ;
  rdfs:domain :Thing ;
  rdfs:range xsd:integer .

```

- Property hierarchies can be created via specializations:
rdfs:subPropertyOf
- Inverse properties are defined via owl:inverseOf
- Identical properties are defined via owl:equivalentProperty


```

:isMadeOf a owl:ObjectProperty ;
          rdfs:subPropertyOf consistsOf .


:reads a owl:ObjectProperty ;
       owl:inverseOf isReadBy .
  
```

```


isMadeOf ⊆ consistsOf
reads- ≡ isReadBy
  
```


- **owl:TransitiveProperty**
e.g.: if $\text{isPartOf}(a,b)$ and $\text{isPartOf}(b,c)$
then it holds that $\text{isPartOf}(a,c)$




- **owl:SymmetricProperty**
e.g.: if $\text{isNeighborOf}(a,b)$,
then it holds that $\text{isNeighborOf}(b,a)$



- **owl:FunctionalProperty**
e.g.: if $\text{hasMother}(a,b)$ and $\text{hasMother}(a,c)$,
then it holds that $b=c$




- **owl:InverseFunctionalProperty**
e.g.: if $\text{isBirthMotherOf}(b,a)$ and $\text{isBirthMotherOf}(c,a)$
then it holds that $b=c$



owl:AsymmetricProperty


e.g.: if it holds that `isLeftOf(a,b)`

then it is not possible that also `isLeftOf(b,a)`



owl:ReflexiveProperty

e.g.: `isRelatedTo(x,x)`



owl:IrreflexiveProperty

e.g.: if `isParentOf(x,y)`

then $x \neq y$

```
:isPublishedBefore a owl:ObjectProperty ;  
    a owl:TransitiveProperty ;  
    rdfs:domain owl:Book ;  
    rdfs:range owl:Book .  
  
:AnimalFarm a owl:Book ;  
    :isPublishedBefore :NineteenEightyfour .  
  
:BraveNewWorld a :Book ;  
    : isPublishedBefore :AnimalFarm .
```

Via inference it can be entailed that BraveNewWorld has also been published before NineteenEightyfour .

Two properties R and S are **disjunctive**, if two individuals x,y are never related via both properties

```
:hasParent a owl:ObjectProperty ;  
  owl:propertyDisjointWith :hasChild .
```

Shortcut for several disjunctive properties:

```
[] rdf:type owl:AllDisjointProperties ;  
  owl:members ( :hasParent :hasChild :hasGrandchild ) .
```

Basic constructs:**1. Classes:**

```
:Book a owl:Class .
:Novel a owl:Class ;
    rdfs:subClassOf :Book
:Book owl:disjointWith :Writer .
:Writer owl:equivalentClass :Author .
```

2. Properties:

```
:author a owl:ObjectProperty;
    rdfs:domain :Book;
    rdfs:range :Writer .

:publicationYear a owl:DatatypeProperty;
    rdfs:domain :Book;
    rdfs:range xsd:integer .
```

3. Individuals:

```
:NineteenEightyfour a :Book .
:NineteenEightyfour owl:sameAs :ARX012345 .
```

- **Complex classes** (logical class constructors, property restrictions)
- **Property relationships**

188.399 Introduction to Semantic Systems

(2.0 VU / 3.0 ECTS)

2023W

**Katja Hose
Emanuel Sallinger**

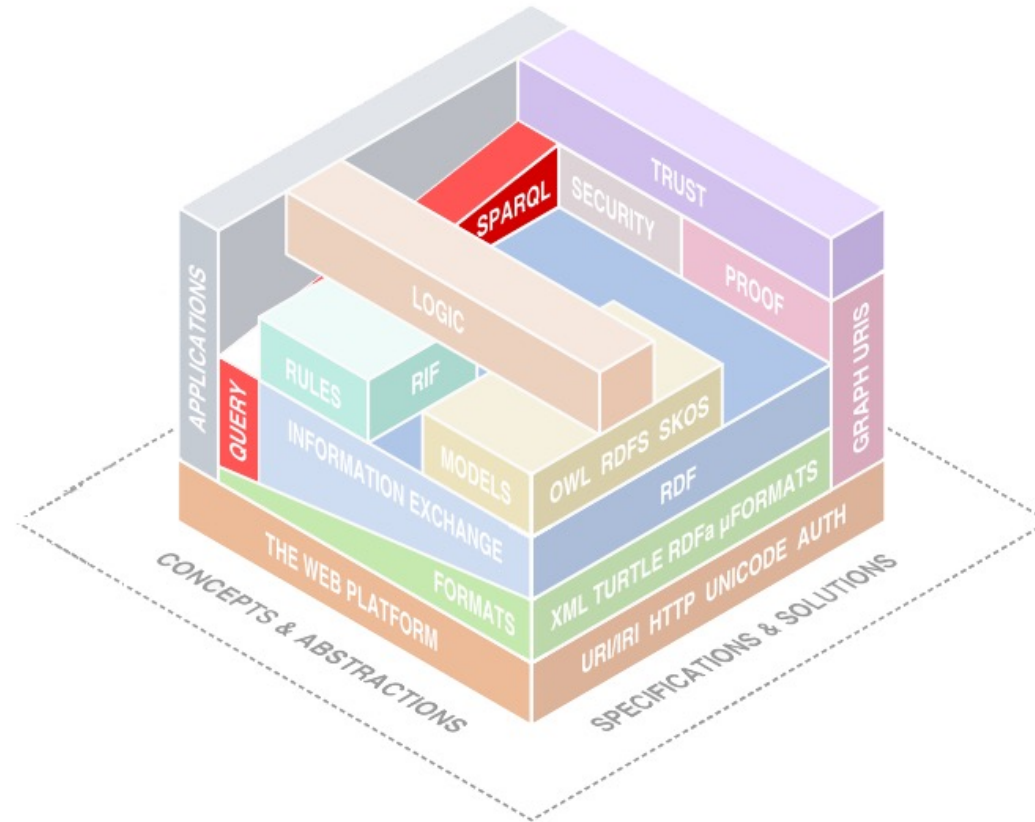
iss@dbai.tuwien.ac.at



- Querying Knowledge Graphs with SPARQL
 - Basic SPARQL
 - Advanced SPARQL
- Updating Knowledge Graphs with SPARQL

SPARQL





- Basics
 - **Introductory notions**
 - The SELECT query
 - Other query types
- Advanced notions
 - CONSTRUCT queries
 - Filter & Order
 - Aggregate & Group
 - Union
 - Subqueries
 - Federated queries
 - Entailment regimes

- SPARQL = **SPARQL Protocol and RDF Query Language**
 - Semantic query language
 - Developed by W3C RDF Data Access Working group (DAWG)
 - January 2008: SPARQL 1.0
 - March 2013: SPARQL 1.1
 - <https://www.w3.org/TR/sparql11-query/>

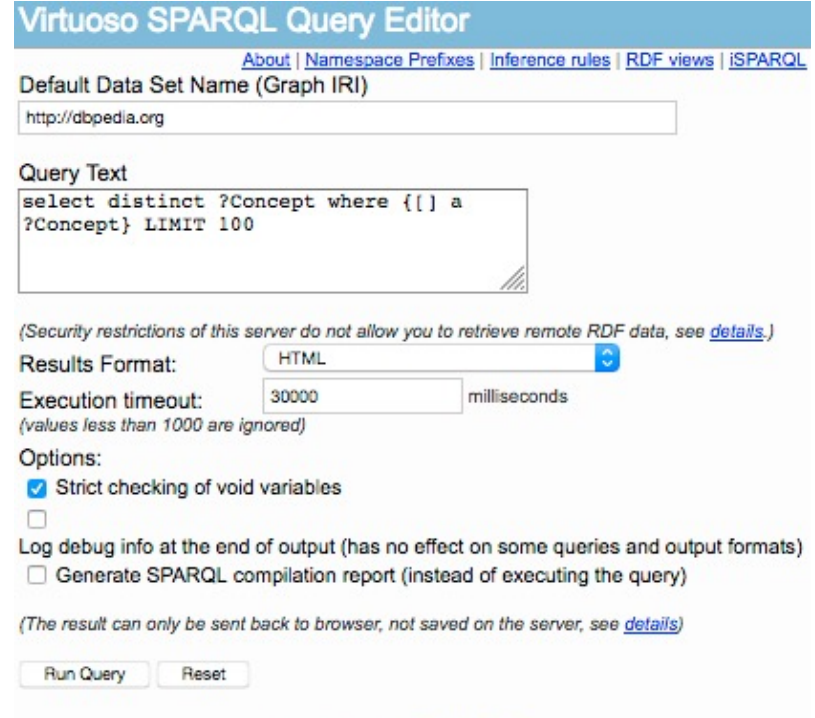
- SPARQL lets us
 - Retrieve and manipulate data stored in RDF
 - Explore data by querying unknown relationships
 - Perform complex joins of disparate databases in a single, simple query
 - Transform RDF data from one vocabulary to another

- Triple stores
 - OpenLink Virtuoso (<https://virtuoso.openlinksw.com>)
 - Apache Jena ARQ (<http://jena.apache.org>)
 - Many more: <https://www.w3.org/wiki/SparqlImplementations>

- **RDF graph:** Set of RDF assertions, manipulated as a labeled directed graph.
- **RDF dataset:** set of RDF triples; comprised of
 - one default graph
 - zero or more named graphs
- **SPARQL protocol client:** HTTP client that sends requests for SPARQL Protocol operations (queries or updates)
- **SPARQL protocol service:** HTTP server that services requests for SPARQL Protocol operations
- **SPARQL endpoint:** The URI at which a SPARQL Protocol service listens for requests from SPARQL clients

SPARQL Fundamentals

<https://dbpedia.org/sparql>



The screenshot shows the Virtuoso SPARQL Query Editor interface. At the top, there is a blue header with the text "Virtuoso SPARQL Query Editor". Below the header, there are several navigation links: "About", "Namespace Prefixes", "Inference rules", "RDF views", and "iSPARQL". The main content area is divided into several sections. The first section is "Default Data Set Name (Graph IRI)", which contains a text input field with the value "http://dbpedia.org". The second section is "Query Text", which contains a text area with the query: "select distinct ?Concept where {{}} a ?Concept} LIMIT 100". Below the query text, there is a warning message: "(Security restrictions of this server do not allow you to retrieve remote RDF data, see [details](#).)". The third section is "Results Format:", which has a dropdown menu set to "HTML". The fourth section is "Execution timeout:", which has a text input field with the value "30000" and the unit "milliseconds". Below this, there is a note: "(values less than 1000 are ignored)". The fifth section is "Options:", which contains three checkboxes: "Strict checking of void variables" (checked), "Log debug info at the end of output (has no effect on some queries and output formats)" (unchecked), and "Generate SPARQL compilation report (instead of executing the query)" (unchecked). At the bottom of the interface, there are two buttons: "Run Query" and "Reset".

Copyright © 2018 [OpenLink Software](#)
Virtuoso version 07.20.3229 on Linux (i686-generic-linux-glibc25-64), Single Server Edition

RDF triple


```
dbpedia:The_Beatles foaf:name "The Beatles" .
```



RDF triple pattern



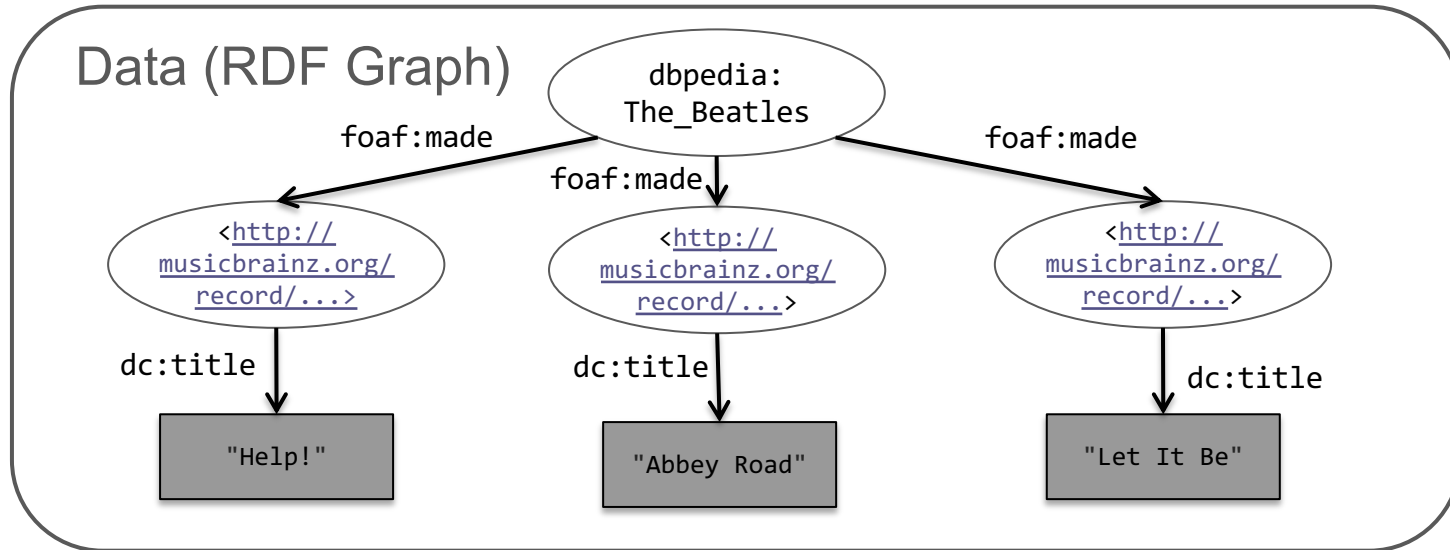
```
dbpedia:The_Beatles foaf:made ?album .
```

Just like triples, but any of their parts can be replaced by a variable.





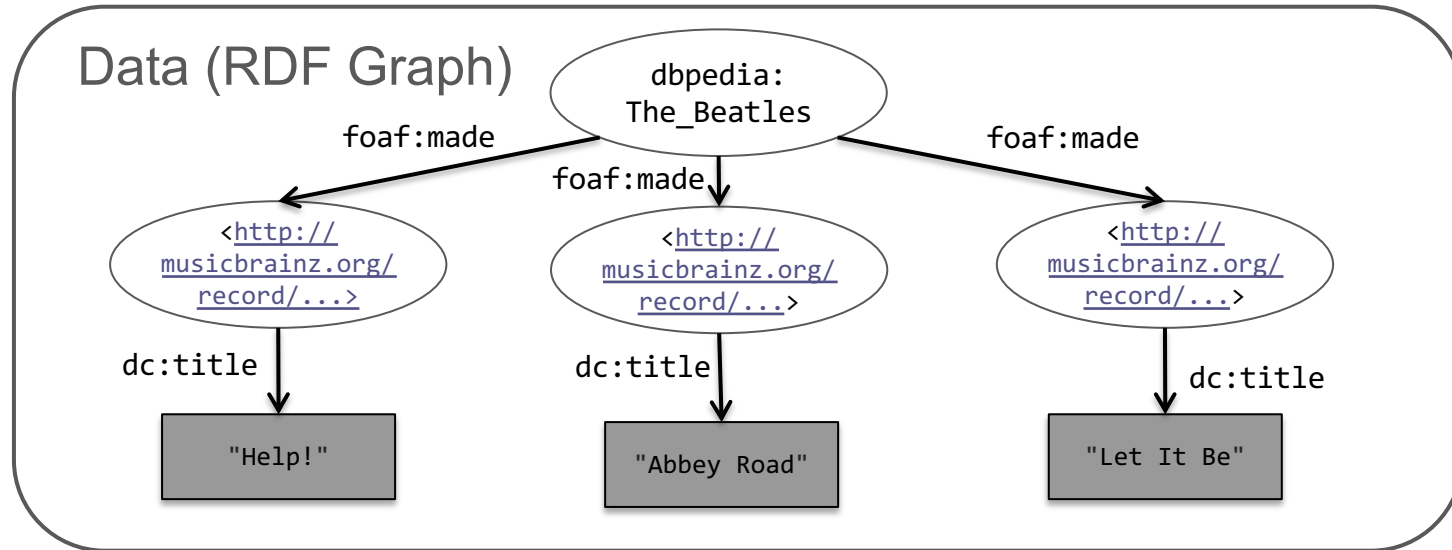
Other triple pattern examples:

```
?album mo:track ?track .  
?s ?p ?o .
```




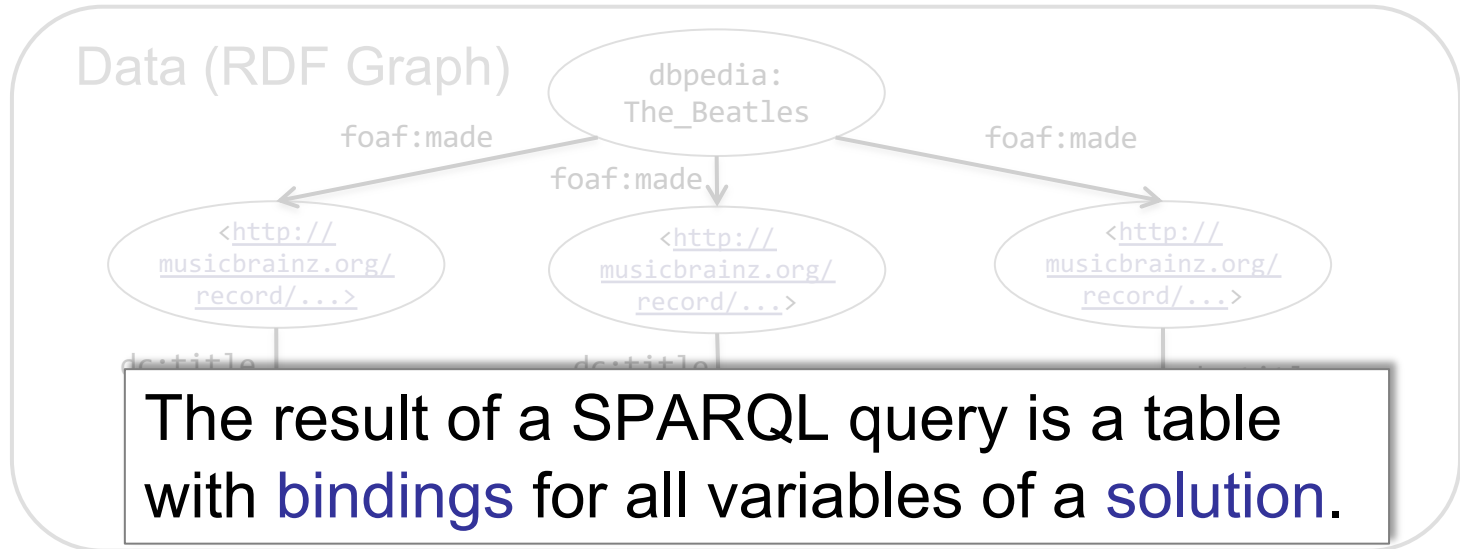
Results

What are the results?
→ Menti




Results

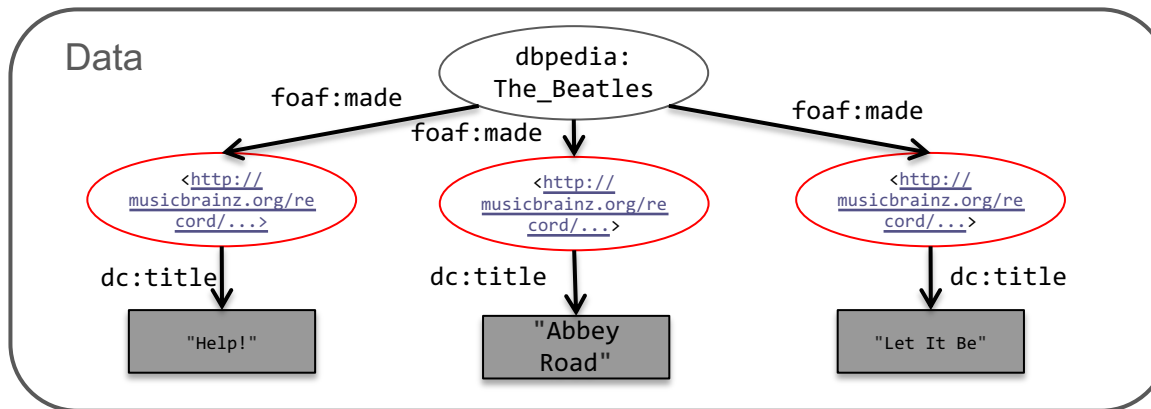
?album
http://musicbrainz.org/...
http://musicbrainz.org/...
http://musicbrainz.org/...



- Key idea: Pattern matching
- **Queries** (graph patterns):
 - describe subgraphs of an RDF graph
 - Roughly correspond to RDF graphs specified in Turtle syntax but containing variables (prefixed by **?**)



- **Results:** Subgraphs matching the graph pattern
→ or rather a table with the variables of those subgraphs



```
# Prefix declarations, for abbreviating URIs
PREFIX dbpedia: http://dbpedia.org/resource/
...

# Result clause, identifying what information to return from the query
SELECT/ASK/CONSTRUCT ...

# Dataset definitions, stating what RDF graph(s) are being queried
FROM <http://musicbrainz.org/20130302>

# Query pattern, specifying what to query for in the underlying dataset
WHERE {
    ...
}

# Query modifiers, slicing, ordering, rearranging query results
ORDER BY ...
```

- Basics
 - Introductory notions
 - **The SELECT query**
 - Other query types
- Advanced notions
 - CONSTRUCT queries
 - Filter & Order
 - Aggregate & Group
 - Union
 - Subqueries
 - Federated queries
 - Entailment regimes

```
PREFIX dbpedia: <http://dbpedia.org/resource/>
```

```
PREFIX foaf: <http://xmlns.com/foaf/0.1/>
```

```
SELECT ?album
```

```
FROM <http://musicbrainz.org/20130302>
```

```
WHERE {
```

```
    dbpedia:The_Beatles foaf:made ?album .
```

```
}
```

Query form:

- **SELECT** retrieves variables and their bindings as a table
- Other query forms: **ASK**, **DESCRIBE** or **CONSTRUCT**

?album
< http://musicbrainz.org/20130302 >
< http://musicbrainz.org/20130302 >
< http://musicbrainz.org/20130302 >

```
PREFIX dbpedia: <http://dbpedia.org/resource/>
PREFIX foaf: <http://xmlns.com/foaf/0.1/>

SELECT ?album
FROM <http://musicbrainz.org/20130302>
WHERE {
    dbpedia:The_Beatles foaf:made ?album .
}
```

Data set specification

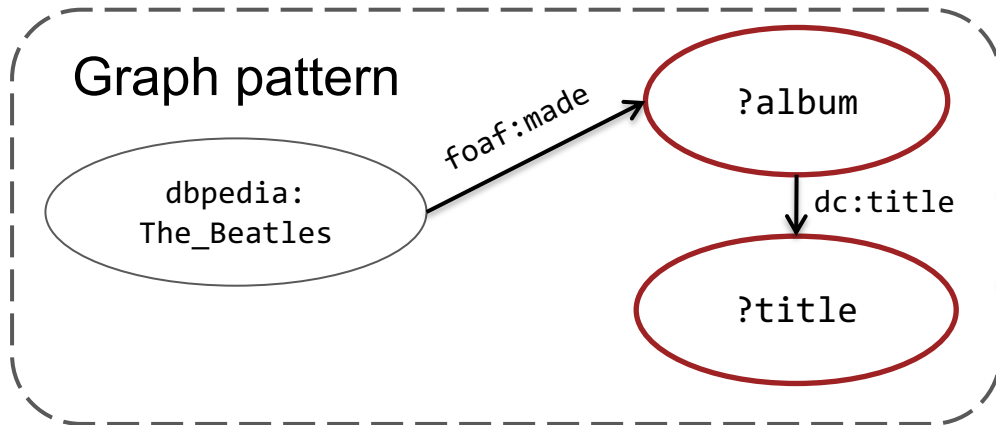

- This clause is optional
- **FROM** or **FROM NAMED**
- Indicates the sources for the data against which to find matches

```
PREFIX dbpedia: <http://dbpedia.org/resource/>
PREFIX foaf: <http://xmlns.com/foaf/0.1/>

SELECT ?album
FROM <http://musicbrainz.org/20130302>
WHERE {
    dbpedia:The_Beatles foaf:made ?album .
}
```


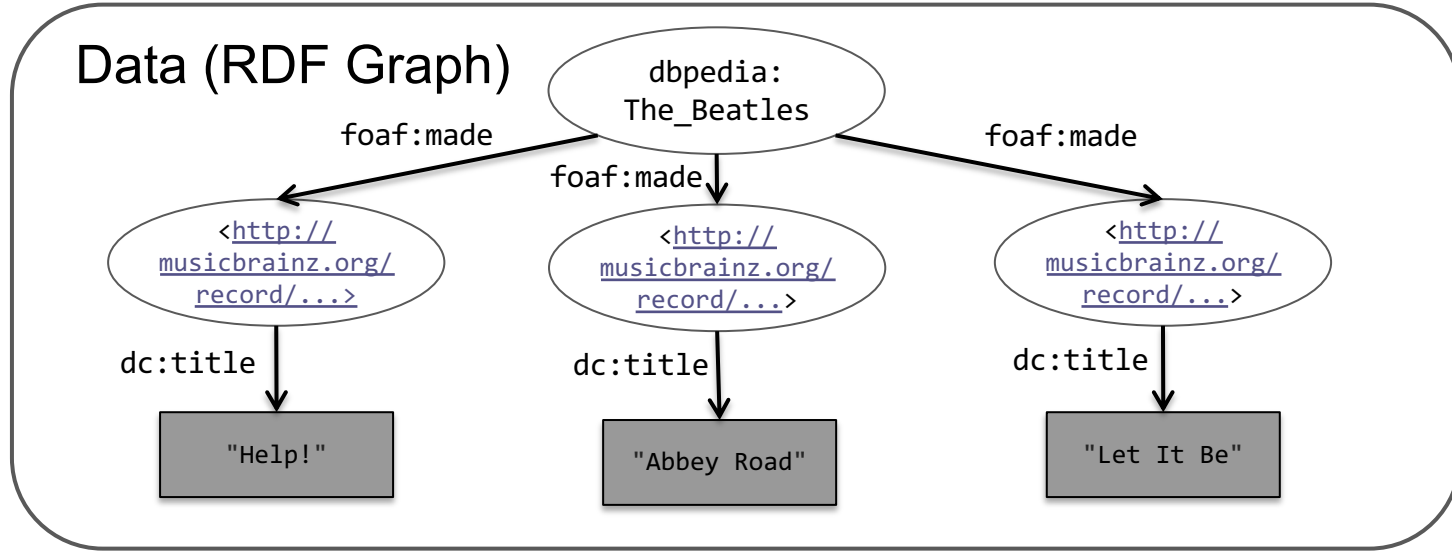
Query pattern

- Generalizes Turtle with variables and keywords
 - The final “.” after last triple pattern is optional
- One or multiple triple patterns between “{}”
 - together describing a graph pattern to be matched against RDF triples specified in the FROM clause



Results

What are the results?
→ Menti



Results

?album	?title
<http://...>	"Help!"
<http://...>	"Abbey Road"
<http://...>	"Let It Be"


```
PREFIX dbpedia, foaf, dc, mo

SELECT *
FROM <http://musicbrainz.org/20130302>
WHERE {
    dbpedia:The_Beatles foaf:made ?album .
    ?album dc:title ?title .
}
ORDER BY ?title
```

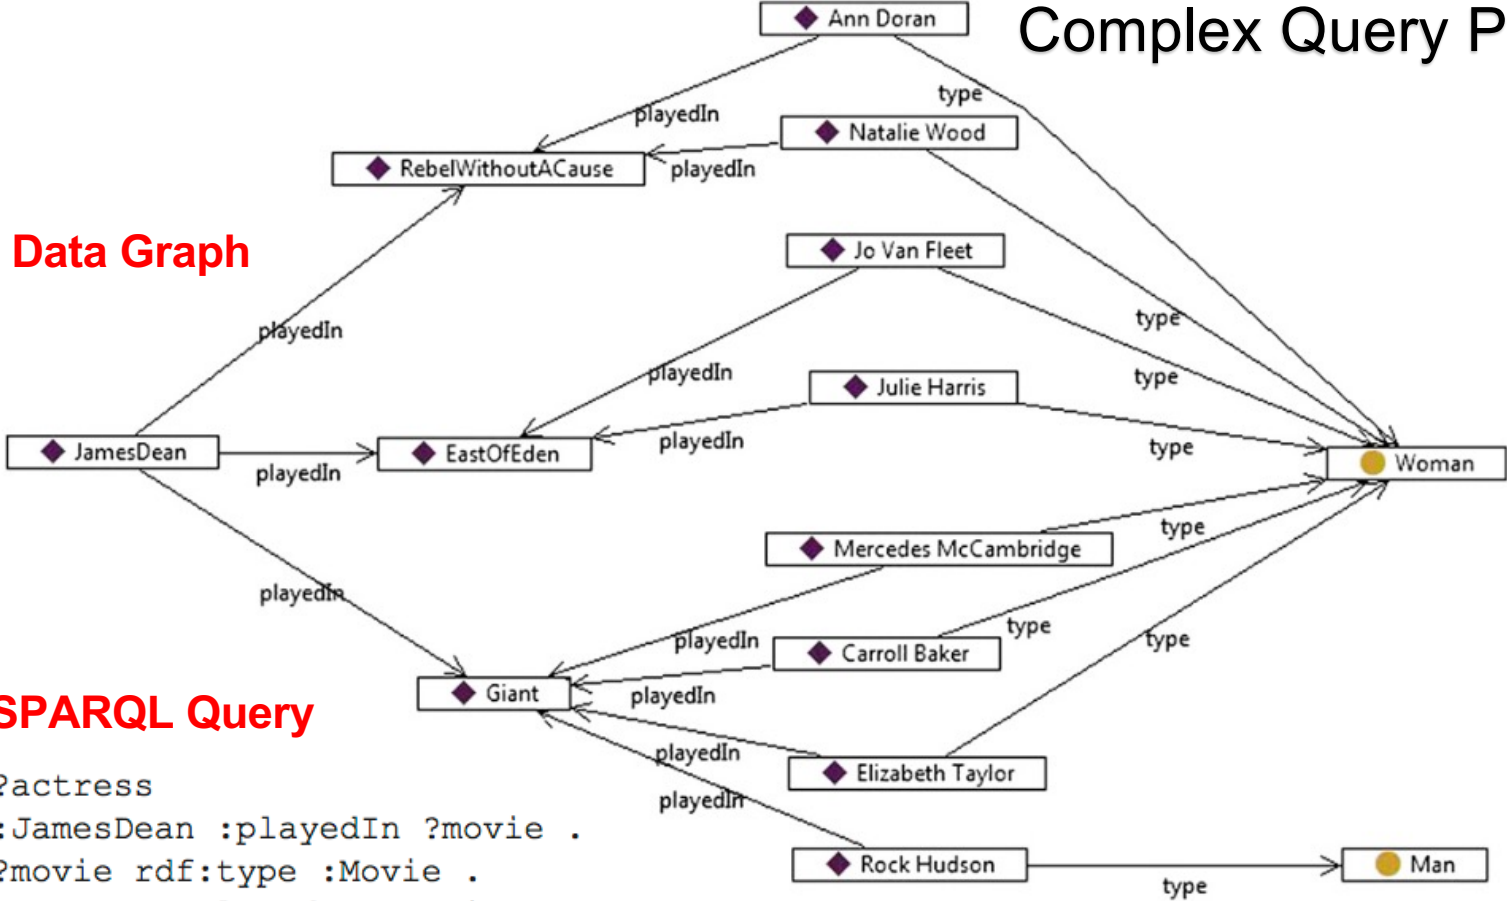
Solution modifier:

- Modify the result set
- **ORDER BY**, **LIMIT** or **OFFSET**
- **GROUP BY**
- **SELECT ***

?album	?title
< http://... >	"Abbey Road"
< http://... >	"Help!"
< http://... >	"Let It Be"

Complex Query Patterns

Data Graph




SPARQL Query

```
SELECT ?actress
WHERE {
  :JamesDean :playedIn ?movie .
  ?movie rdf:type :Movie .
  ?actress :playedIn ?movie .
  ?actress rdf:type :Woman }

```

Query Pattern



- Basics
 - Introductory notions
 - The SELECT query
 - **Other query types**
- Advanced notions
 - CONSTRUCT queries
 - Filter & Order
 - Aggregate & Group
 - Union
 - Subqueries
 - Federated queries
 - Entailment regimes

SPARQL supports the following query types

- **SELECT**
returns variables and their bindings directly
- **ASK**
tests whether or not a query pattern has a solution
Returns true or false
- **DESCRIBE**
returns a single RDF graph containing RDF data about a resource
- **CONSTRUCT**
returns a single RDF graph specified by a graph template

ASK - tests whether or not a query pattern has a solution. Returns yes/no

Query: *Is Paul McCartney a member of 'The Beatles'?*

```
PREFIX dbpedia: <http://dbpedia.org/resource/>
```

```
PREFIX mo: <http://purl.org/ontology/mo/>
```

```
ASK WHERE { dbpedia:The_Beatles mo:member dbpedia:Paul_McCartney. }
```

In your own words:
What does this query mean?

Results

true

ASK - tests whether or not a query pattern has a solution. Returns yes/no

Query: *Is Elvis Presley a member of 'The Beatles'?*

```
PREFIX dbpedia: <http://dbpedia.org/resource/>
```

```
PREFIX mo: <http://purl.org/ontology/mo/>
```

```
ASK WHERE { dbpedia:The_Beatles mo:member dbpedia:Elvis_Presley. }
```

In your own words:
What does this query mean?

Results

false

ASK - tests whether or not a query pattern has a solution. Returns yes/no

Query: *Is Elizabeth Taylor dead?*

```
ASK WHERE { :ElizabethTaylor :diedOn ?any . }
```

Results

false

(no triple matches this pattern)

Query: *Were any of the actors of "Giant" born after 1950?*

```
ASK WHERE { ?any :playedIn :Giant .  
             ?any :bornOn ?birthday .  
             FILTER (?birthday > "1950-01-01"^^xsd:date) }
```

Results

false

*(everyone playing in Giant
was born before 1950)*

Takes the resources within the solution and provides information about them as RDF statements. They can be identified by:

- Specifying **explicit IRIs**

```
PREFIX dbpedia: <http://dbpedia.org/resource/>
DESCRIBE dbpedia:Paul_McCartney
```

- **Bindings of variables** in the WHERE clause

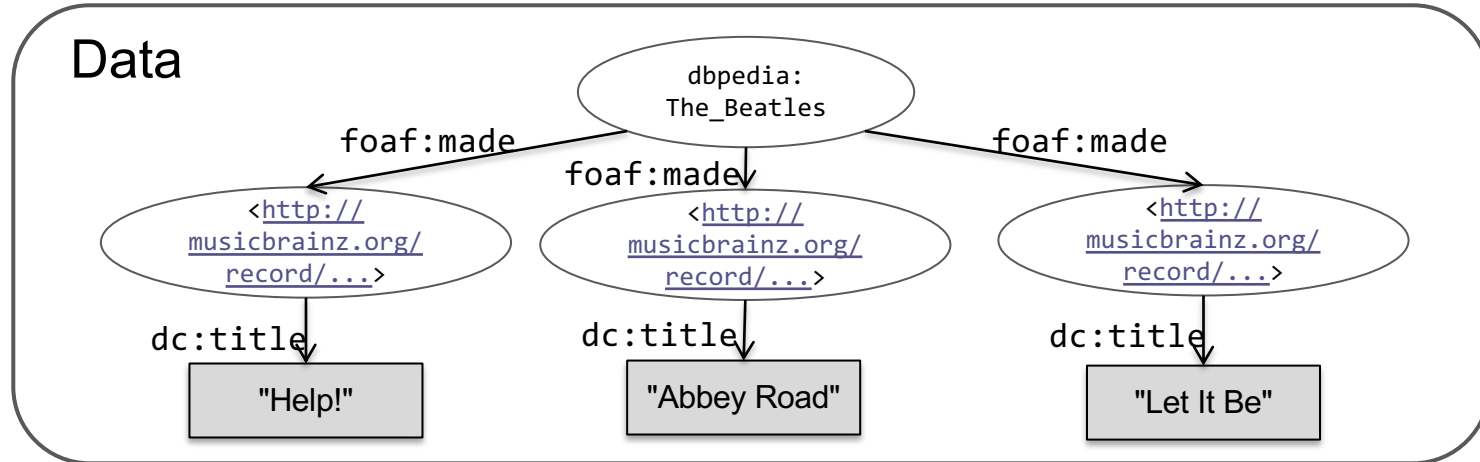
```
PREFIX dbpedia: <http://dbpedia.org/resource/>
PREFIX mo: <http://purl.org/ontology/mo/>

DESCRIBE ?member
WHERE { dbpedia:The_Beatles mo:member ?member . }
```

Very useful for exploration when you are confronted with a new dataset that you don't know the structure of.

- Basics
 - Introductory notions
 - The SELECT query form
 - Other query forms
- Advanced notions
 - **CONSTRUCT queries**
 - Filter & Order
 - Aggregate & Group
 - Union
 - Subqueries
 - Federated queries
 - Entailment regimes

Returns RDF statements created from variable bindings

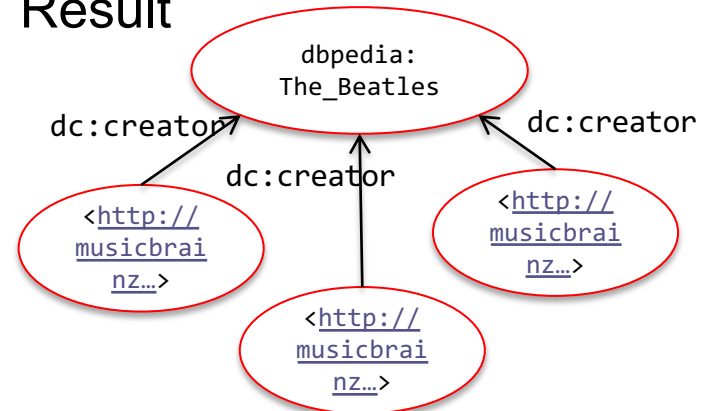


Query

```

CONSTRUCT {
  ?album dc:creator dbpedia:The_Beatles .}
WHERE {
  dbpedia:The_Beatles foaf:made ?album .}
  
```

Result



Query: *Create the dc:creator descriptions for albums and their tracks recorded by 'The Beatles'.*

```
PREFIX dbpedia, foaf, mo, dc ...
CONSTRUCT {
  ?album dc:creator dbpedia:The_Beatles .
  ?track dc:creator dbpedia:The_Beatles .}
WHERE {
  dbpedia:The_Beatles foaf:made ?album .
  ?album mo:track ?track .}
```

CONSTRUCT queries allow phrasing rules of the type
 “IF this is found THEN generate that”

```

CONSTRUCT {?q1 :hasSibling ?q2} WHERE {?q1 :hasBrother ?q2}
CONSTRUCT {?q1 :hasSibling ?q2} WHERE {?q1 :hasSister ?q2}
CONSTRUCT {?q1 :hasParent ?q2} WHERE {?q1 :hasFather ?q2}
CONSTRUCT {?q1 :hasParent ?q2} WHERE {?q1 :hasMother ?q2}
  
```

```

CONSTRUCT {?q1 :hasUncle ?q2}
WHERE {?q2 :hasSibling ?parent .
        ?q2 a :Man .
        ?q1 :hasParent ?parent }
  
```

- Basics
 - Introductory notions
 - The SELECT query form
 - Other query forms
- Advanced notions
 - CONSTRUCT queries
 - **Filter & Order**
 - Aggregate & Group
 - Union
 - Subqueries
 - Federated queries
 - Entailment regimes

- Different types of **filters and functions** may be used
- Filters use comparison and logical operators
- Positioned **within** the WHERE clause

Query: *Retrieve the albums and tracks recorded by ‘The Beatles’, where the duration of the song is more than 300 secs. and no longer than 400 secs.*

```
PREFIX dbpedia, foaf, dc, mo

SELECT ?album_name ?track_title ?duration
WHERE {
    dbpedia:The_Beatles foaf:made ?album .
    ?album dc:title ?album_name ;
           mo:track ?track .
    ?track dc:title ?track_title ;
           mo:duration ?duration;
    FILTER (?duration > 300000 && ?duration < 400000)
}
```

Filter expressions as regular expressions over strings

Query: *Create the `dc:creator` descriptions of the albums recorded by 'The Beatles' whose title contains the word "love".*

```
PREFIX dbpedia, foaf, dc

CONSTRUCT {?album dc:creator dbpedia:The_Beatles .}

WHERE {
  dbpedia:The_Beatles foaf:made ?album .
  ?album dc:title ?album_name ;
  FILTER (REGEX(?album_name, ".*love.*", i))
}
```

- Elimination of duplicates with **DISTINCT** and **REDUCED**
- **REDUCED** does not guarantee elimination of duplicates

Query: *Retrieve the name of the albums recorded by 'The Beatles' which have at least two different songs.*

```

PREFIX dbpedia, foaf, dc, mo

SELECT MODIFIER ?album_name

WHERE {
  dbpedia:The_Beatles foaf:made ?album .
  ?album dc:title ?album_name ;
  mo:track ?track1 .
  mo:track ?track2 .
  FILTER (?track1 != ?track2)
}
  
```

Results

MODIFIER=
DISTINCT

?album
"Revolver"
"Sessions"
"Abbey Road"

MODIFIER=
REDUCED

?album
"Revolver"
"Revolver"
"Revolver"
"Sessions"
"Abbey Road"
"Abbey Road"

It is possible to combine the query with solution modifiers (**ORDER BY**, **LIMIT**, **OFFSET**) to get **subsets of results**:

- **ORDER BY (DESC/ASC)**: result set is shown in decreasing/ascending order
- **LIMIT**: limits the number of results returned by a query

Query *Create the dc:creator descriptions for the 10 most recent albums and their tracks recorded by 'The Beatles'.*

```

PREFIX dbpedia, foaf, mo, dc;

CONSTRUCT {
  ?album dc:creator dbpedia:The_Beatles .
  ?track dc:creator dbpedia:The_Beatles .}
WHERE {
  dbpedia:The_Beatles foaf:made ?album .
  ?album mo:track ?track ;
         dc:date ?date .
} ORDER BY DESC(?date)
LIMIT 10
  
```

- Basics
 - Introductory notions
 - The SELECT query form
 - Other query forms
- Advanced notions
 - CONSTRUCT queries
 - Filter & Order
 - **Aggregate & Group**
 - Union
 - Subqueries
 - Federated queries
 - Entailment regimes

Aggregates

- Calculate aggregate values: COUNT, SUM, MIN, MAX, AVG, GROUP_CONCAT and SAMPLE
- In combination with the GROUP BY operator
- Prune at group level (cf. FILTER) using HAVING

Company	Amount	Year
ACME	\$1250	2010
PRIME	\$3000	2009
ABC	\$2500	2009
ABC	\$2800	2010
PRIME	\$1950	2010
ACME	\$2500	2009
ACME	\$3100	2010
ABC	\$1500	2009
ACME	\$1250	2009
PRIME	\$2350	2009
PRIME	\$1850	2010

Q: How high are the total sales?

```

SELECT  (SUM(?val) AS ?total)
        WHERE {
            ?s a co:Sale .
            ?s co:amount ?val .
        }
```

```

:row1 a :Sale .
:row1 :company :ACME .
:row1 :amount 1250 .
:row1 :year 2010 .
```

?total
24050

DATA

Company	Amount	Year
ACME	\$1250	2010
PRIME	\$3000	2009
ABC	\$2500	2009
ABC	\$2800	2010
PRIME	\$1950	2010
ACME	\$2500	2009
ACME	\$3100	2010
ABC	\$1500	2009
ACME	\$1250	2009
PRIME	\$2350	2009
PRIME	\$1850	2010

Solution Modifier: **GROUP BY**

```
:row1 a :Sale .
:row1 :company :ACME .
:row1 :amount 1250 .
:row1 :year 2010 .
```

Q: How high are the sales per year?

```
SELECT ?year (SUM(?val) AS ?total)
WHERE {
  ?s a co:Sale .
  ?s co:amount ?val .
  ?s co:year ?year .
}
GROUP BY ?year
```

?year	?total
2009	13100
210	10950

Q: Which companies generated a yearly sales higher than 5K and in which year?

```
SELECT ?year ?company (SUM(?val) AS ?total)
WHERE {
  ?s a co:Sale .
  ?s co:amount ?val .
  ?s co:year ?year .
  ?s co:company ?company .
}
GROUP BY ?year ?company
HAVING (?total > 5000)
```

HAVING is similar to FILTER but applies to the variables of the groups and is **outside** the graph pattern

?year	?company	?total
2009	PRIME	5350

Query: *Retrieve albums recorded by 'The Beatles' that have a duration higher than a given value.*

```
PREFIX dbpedia, foaf, mo

SELECT ?album (SUM(?track_duration) AS ?album_duration)
  WHERE {
    dbpedia:The_Beatles foaf:made ?album .
    ?album mo:track ?track .
    ?track mo:duration ?track_duration .
  } GROUP BY ?album
  HAVING (?album_duration > 3600000)
```

- Basics
 - Introductory notions
 - The SELECT query form
 - Other query forms
- Advanced notions
 - CONSTRUCT queries
 - Filter & Order
 - Aggregate & Group
 - **Union**
 - Subqueries
 - Federated queries
 - Entailment regimes

- Allows the specification of alternatives (disjunctions)
- UNION = the set union of all solution bindings of the connected patterns

Query: *Which actors played in 'Giant' or 'Rebel without a cause' or both?*

```
PREFIX mo

SELECT ?actor
WHERE {
    {?actor mo:playedIn mo:Giant .}
    UNION
    {?actor mo:playedIn mo:RebelWithoutACause .}
}
```

actor
Ann Doran
Carroll Baker
Elizabeth Taylor
James Dean
James Dean
Jim Backus
Mercedes McCambridge
Natalie Wood
Rock Hudson
Sal Mineo
Sal Mineo

- Basics
 - Introductory notions
 - The SELECT query form
 - Other query forms
- Advanced notions
 - CONSTRUCT queries
 - Filter & Order
 - Aggregate & Group
 - Union
 - **Subqueries**
 - Federated queries
 - Entailment regimes


```
PREFIX co
```

```
SELECT ?company
```

```
WHERE
```

```
{ SELECT ?company (SUM(?val) AS ?total09)
```

```
WHERE {
```

```
  ?s a co:Sale .
```

```
  ?s co:amount ?val .
```

```
  ?s co:year 2009 .
```

```
  ?s co:company ?company .}
```

```
GROUP BY ?company } .
```

```
{ SELECT ?company (SUM(?val) AS ?total10)
```

```
WHERE {
```

```
  ?s a co:Sale .
```

```
  ?s co:amount ?val .
```

```
  ?s co:year 2010 .
```

```
  ?s co:company ?company .}
```

```
GROUP BY ?company } .
```

```
FILTER (?total10 >?total09) .
```

```
}
```

Aggregate values can be "named" using variables (keyword **AS**)

Query:

Select companies for which the total sales value in 2010 is larger than the total sales value in 2009.

?company
ACME

CONSTRUCT and **GROUP BY** can be combined

```

PREFIX co

CONSTRUCT {?company a co:PreferredCustomer .
           ?company co:totalSales ?total .}

WHERE
{ SELECT ?year ?company (SUM(?val) AS ?total)
  WHERE {
    ?s a co:Sale .
    ?s co:amount ?val .
    ?s co:year ?year .
    ?s co:company ?company .
  }
  GROUP BY ?year ?company
  HAVING (?total > 5000)
}
    
```

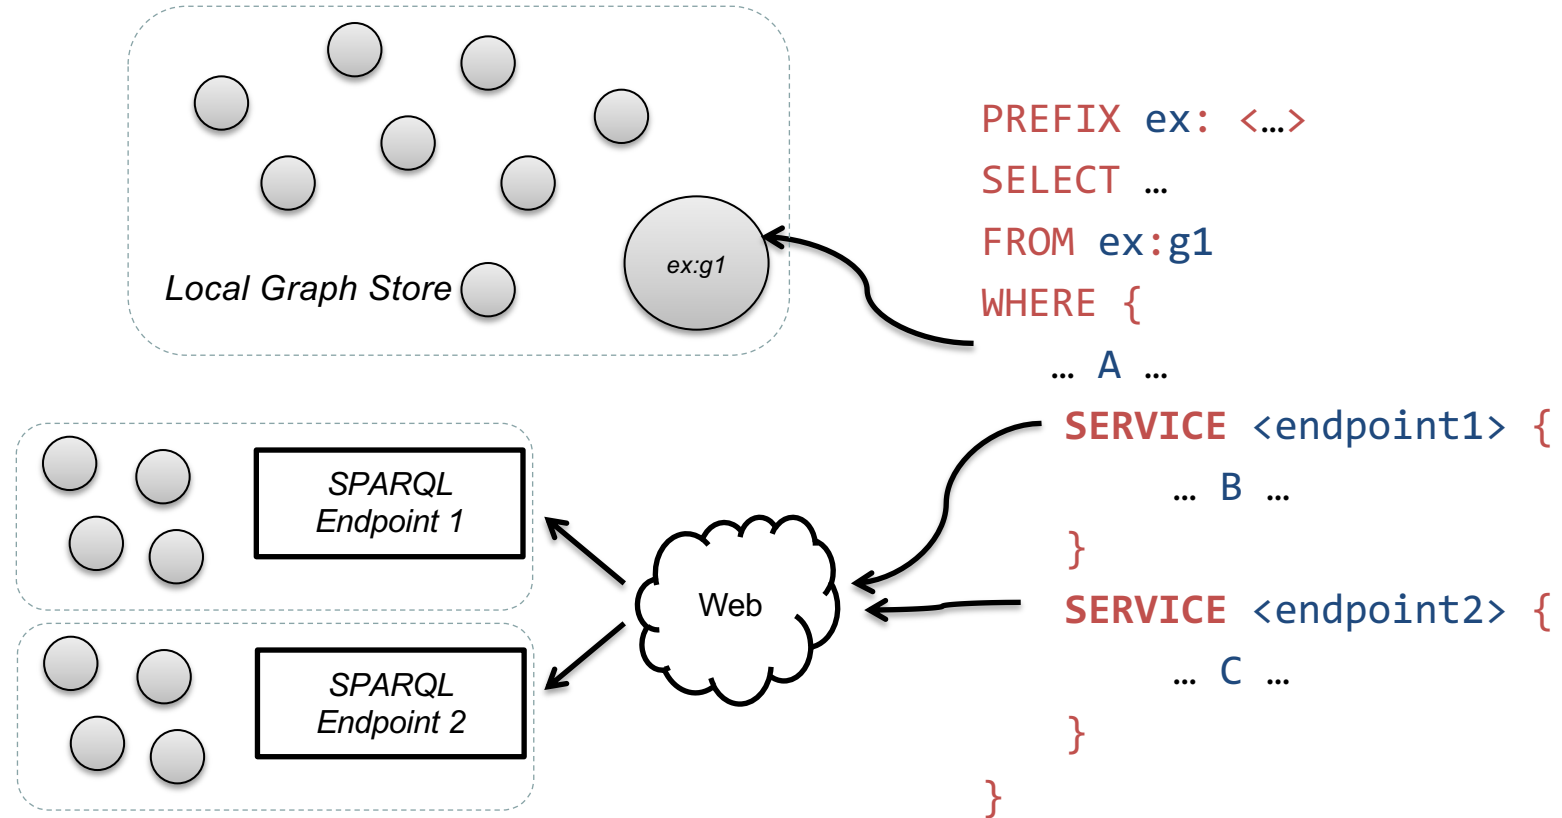
Query:

Companies for which the total sales value in any year is larger than 5000 are PreferredCustomers.

```

:PRIME a :PreferredCustomer .
:PRIME :totalSales 5350.00 .
    
```

- Basics
 - Introductory notions
 - The SELECT query form
 - Other query forms
- Advanced notions
 - CONSTRUCT queries
 - Filter & Order
 - Aggregate & Group
 - Union
 - Subqueries
 - **Federated queries**
 - Entailment regimes



Find the birth dates of all of the actors in *Star Trek: The Motion Picture*

PREFIX movie: <http://data.linkedmdb.org/resource/movie/>

PREFIX dbpedia: <http://dbpedia.org/ontology/>

PREFIX foaf: <http://xmlns.com/foaf/0.1/>

SELECT ?actor_name ?birth_date

FROM <http://dig.csail.mit.edu/2008/webdav/timbl/foaf.rdf> # placeholder graph

WHERE {

{

SERVICE <http://data.linkedmdb.org/sparql> {

<http://data.linkedmdb.org/resource/film/675> movie:actor ?actor .

?actor movie:actor_name ?actor_name

}

BIND(STRLANG(?actor_name, "en") AS ?actor_name_en)

}

SERVICE <http://dbpedia.org/sparql> {

?actor2 a foaf:Person ;

foaf:name ?actor_name_en ;

dbpedia:birthDate ?birth_date

}

}

actor_name	birth_date
"William Shatner"	"1931-03-22" ^^<http://www.w3.org/2001/XMLSchema#date>
"Leonard Nimoy"	"1931-03-26" ^^<http://www.w3.org/2001/XMLSchema#date>
"DeForest Kelley"	"1920-01-20" ^^<http://www.w3.org/2001/XMLSchema#date>
"James Doohan"	"1920-03-03" ^^<http://www.w3.org/2001/XMLSchema#date>
"Nichelle Nichols"	"1932-12-28" ^^<http://www.w3.org/2001/XMLSchema#date>
"Stephen Collins"	"1947-10-01" ^^<http://www.w3.org/2001/XMLSchema#date>
"George Takei"	"1937-04-20" ^^<http://www.w3.org/2001/XMLSchema#date>

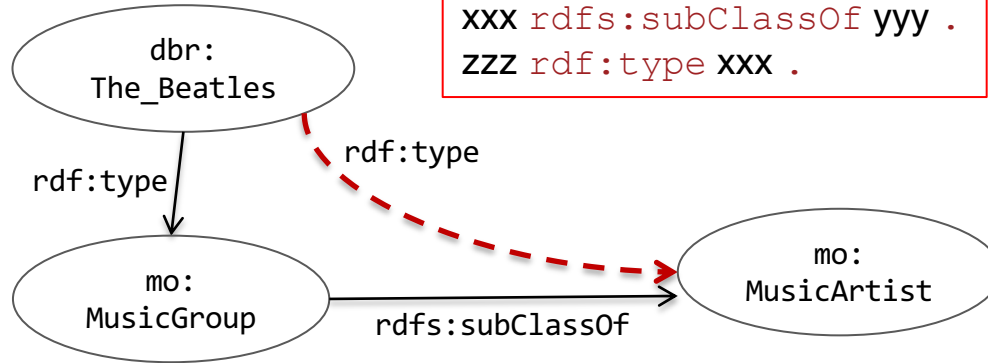
- Basics:
 - Introductory notions
 - The SELECT query form
 - Other query forms
- Advanced notions
 - CONSTRUCT queries
 - Filter & Order
 - Aggregate & Group
 - Union
 - Subqueries
 - Federated queries
 - **Entailment regimes**

- A triple store can use the knowledge encoded in the schema to infer new facts

Schema: `mo:MusicGroup rdfs:subClassOf mo:MusicArtist .`

Inferred facts: `mo:MusicGroup a rdfs:Class .`
`mo:MusicArtist a rdfs:Class .`

- Whether a triple store supports reasoning, influences the outcome of querying



`xxx rdfs:subClassOf yyy .`
`zzz rdf:type xxx .`

`zzz rdf:type yyy .`

Schema:

`mo:MusicGroup rdfs:subClassOf mo:MusicArtist .`

Query:

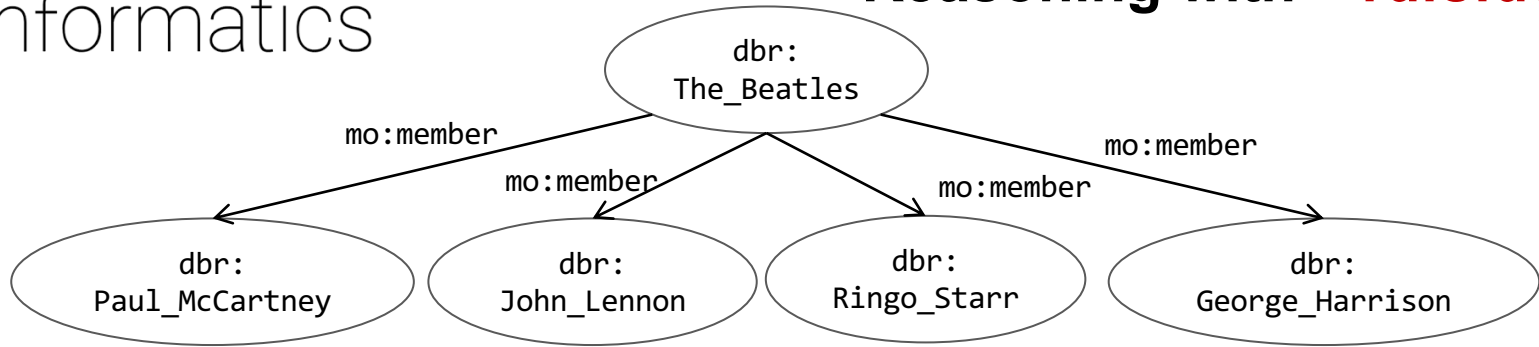
`SELECT ?x`
`WHERE {?x a mo:MusicArtist.}`

Result set **with inference:**

?x
dbr:The_Beatles ...

Result set **without inference:**

?x



Schema:

```
mo:member rdfs:domain mo:MusicGroup .
```

Query:

```
SELECT ?x
WHERE {?x a mo:MusicGroup.}
```

Result set **with inference:**

?x
dbr:The_Beatles ...


Result set without interference:

?x

```
aaa rdfs:domain xxx .
yyy aaa zzz .
```



```
yyy rdf:type xxx .
```



Schema:

```
dbo:bandMember rdfs:range foaf:Musician .
```

Query:

```
SELECT ?x
WHERE {?x a foaf:Musician .}
```

Result set **with inference**:

?x
dbr:Paul_McCartney
dbr:John_Lennon
dbr:Ringo_Starr
dbr:George_Harrison ...

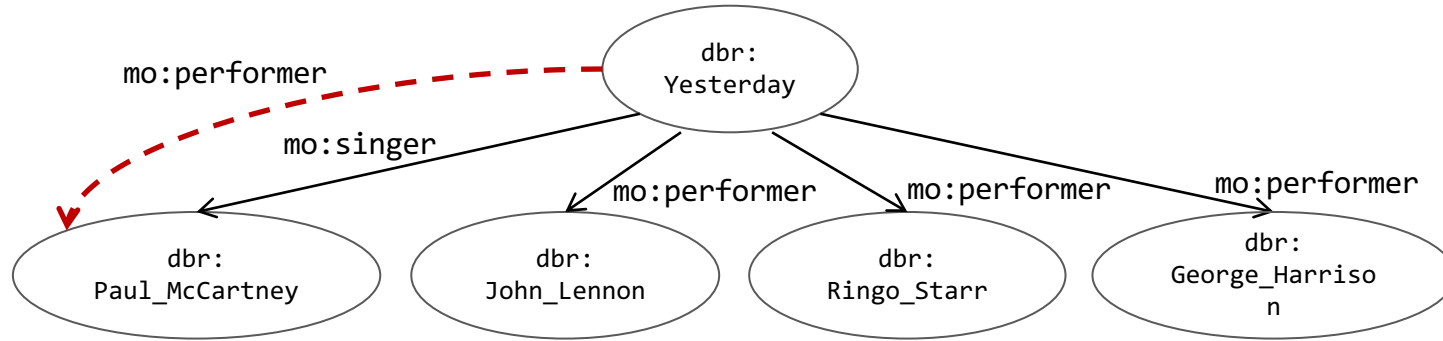
Result set **without inference**:

?x

```
aaa rdfs:range xxx .
yyy aaa zzz .
```



```
zzz rdf:type xxx .
```



Schema:

Query:

```
mo:singer rdfs:subPropertyOf mo:performer .
```

```
SELECT ?x
WHERE {dbr:Yesterday mo:performer ?x .}
```

Result set with inference:

?x
dbr:John_Lennon
dbr:Ringo_Starr
dbr:George_Harrison
dbr:Paul_McCartney

Result set without inference:

?x
dbr:John_Lennon
dbr:Ringo_Starr
dbr:George_Harrison

```
aaa rdfs:subPropertyOf bbb .
yyy aaa zzz .
```



```
yyy bbb zzz .
```

	If S contains:	then S RDFS entails recognizing D:
<i>rdfs1</i>	any IRI aaa in D	aaa <code>rdf:type rdfs:Datatype .</code>
<i>rdfs2</i>	aaa <code>rdfs:domain xxx .</code> yyy aaa zzz .	yyy <code>rdf:type xxx .</code>
<i>rdfs3</i>	aaa <code>rdfs:range xxx .</code> yyy aaa zzz .	zzz <code>rdf:type xxx .</code>
<i>rdfs4a</i>	xxx aaa yyy .	xxx <code>rdf:type rdfs:Resource .</code>
<i>rdfs4b</i>	xxx aaa yyy .	yyy <code>rdf:type rdfs:Resource .</code>
<i>rdfs5</i>	xxx <code>rdfs:subPropertyOf yyy .</code> yyy <code>rdfs:subPropertyOf zzz .</code>	xxx <code>rdfs:subPropertyOf zzz .</code>
<i>rdfs6</i>	xxx <code>rdf:type rdf:Property .</code>	xxx <code>rdfs:subPropertyOf xxx .</code>
<i>rdfs7</i>	aaa <code>rdfs:subPropertyOf bbb .</code> xxx aaa yyy .	xxx bbb yyy .
<i>rdfs8</i>	xxx <code>rdf:type rdfs:Class .</code>	xxx <code>rdfs:subClassOf rdfs:Resource .</code>
<i>rdfs9</i>	xxx <code>rdfs:subClassOf yyy .</code> zzz <code>rdf:type xxx .</code>	zzz <code>rdf:type yyy .</code>
<i>rdfs10</i>	xxx <code>rdf:type rdfs:Class .</code>	xxx <code>rdfs:subClassOf xxx .</code>
<i>rdfs11</i>	xxx <code>rdfs:subClassOf yyy .</code> yyy <code>rdfs:subClassOf zzz .</code>	xxx <code>rdfs:subClassOf zzz .</code>
<i>rdfs12</i>	xxx <code>rdf:type rdfs:ContainerMembershipProperty .</code>	xxx <code>rdfs:subPropertyOf rdfs:member .</code>
<i>rdfs13</i>	xxx <code>rdf:type rdfs:Datatype .</code>	xxx <code>rdfs:subClassOf rdfs:Literal .</code>

S.. subgraph

D.. Set of IRIs identifying datatypes



Updating RDF with SPARQL 1.1

- SPARQL 1.0 only allows data access (querying)
- SPARQL 1.1 introduces:

Query extensions

Aggregates, subqueries, negation, expressions in the **SELECT** clause, property paths, assignment, expanded set of functions and operators

Updates

Insert, Delete, Delete/Insert
Create, Load, Clear, Drop
Copy, Move, Add

Federation extension

Service, values, service variables (informative)

SPARQL 1.1 provides data update operations:

- **INSERT (DATA)**: adds triples, given inline in the request, into a graph
- **DELETE** data: removes triples, given inline in the request, if the respective graphs contains those
- **DELETE/INSERT** data: uses in parallel **INSERT** and **DELETE**

INSERT + GP that specifies a set of RDF triples, given bindings for any variables that are included in the pattern

Insert "Peter Best" as a formerMember of dbr:The_Beatles

```

PREFIX dbr: <http://dbpedia.org/resource/>
PREFIX dbo: <http://dbpedia.org/ontology/>
PREFIX foaf: <http://xmlns.com/foaf/0.1/>

INSERT { dbr:The_Beatles dbr:formerMember ?x }
WHERE { dbr:The_Beatles dbr:currentMember ?x .
        ?x foaf:name "Peter Best"
      }
  
```


Insert the following albums recorded by The Beatles into the graph
<http://musicbrainz.org/20130209-004702>

```
PREFIX dc: <http://purl.org/dc/elements/1.1/>
```

```
PREFIX foaf: <http://xmlns.com/foaf/0.1/>
```

```
INSERT DATA {
```

```
    GRAPH { <http://musicbrainz.org/20130302>
```

```
    <http://musicbrainz.org/artist/b10bbbfc-cf9e-42e0-be17-e2c3e1d2600d>
```

```
      foaf:made <http://musicbrainz.org/release/3a685770-7326-34fc-9f18-e5f5626f3dc5>
```

```
    ,
```

```
      <http://musicbrainz.org/release/cb6f8798-d51e-4fa5-a4d1-2c0602bfe1b6> .
```

```
    <http://musicbrainz.org/release/3a685770-7326-34fc-9f18-e5f5626f3dc5>
```

```
      dc:title "Please Please Me".
```

```
    <http://musicbrainz.org/release/cb6f8798-d51e-4fa5-a4d1-2c0602bfe1b6>
```

```
      dc:title "Something New". } }
```

Delete all the information about the album "Casualties" of The Beatles.

```
PREFIX dc: <http://purl.org/dc/elements/1.1/>
```

```
PREFIX foaf: <http://xmlns.com/foaf/0.1/>
```

```
DELETE { ?album ?predicate ?object . }
```

```
WHERE {
```

```
http://musicbrainz.org/artist/b10bbbfc-cf9e-42e0-be17-e2c3e1d2600d foaf:made  
    ?album .
```

```
?album dc:title "Casualties".
```

```
?album ?predicate ?object .}
```

Delete the status of 'Peter Best' as current member of "The Beatles" and **insert** his status as former member of the band.

```
PREFIX dbr, dbo, foaf ...
```

```
DELETE { dbr:The_Beatles dbo:currentMember ?x . }
```

```
INSERT { GRAPH <http://musicbrainz.org/20130209-004702> {
  dbr:The_Beatles dbo:formerBandMember ?x .}
```

```
}
```

```
WHERE {
```

```
  dbr:The_Beatles dbo:currentMember ?x .
```

```
  ?x foaf:name "Peter Best" .}
```

SPARQL 1.1 provides graph update operations:

- **CREATE** creates an empty graph in the Graph Store
- **LOAD** reads the content of a document into a graph in the Graph Store
- **CLEAR** removes all triples in one or more graphs
- **DROP** removes the graph from the Graph Store
- Other operations: **COPY**, **MOVE**, **ADD**

- Creates a new named graph
- Can be used with the **DEFAULT** and **ALL** keywords

```
CREATE GRAPH <http://musicbrainz.org/20130302>
```

- **LOAD**: An RDF graph can be loaded from a URL

```
LOAD <http://xmlns.com/foaf/spec/20100809.rdf>
```

```
LOAD <http://xmlns.com/foaf/spec/20100809.rdf>  
    INTO <http://xmlns.com/foaf/0.1/>
```

Named graph

- **LOAD** can be used with the **SILENT** keyword (suppresses errors)

```
CLEAR GRAPH <http://musicbrainz.org/20130302>
```

- **CLEAR** removes all triples in the graph (emptied, but not deleted!)
- The graph(s) can be specified with the following keywords:
DEFAULT, **NAMED**, **ALL**, **GRAPH**
- Can be used with the **SILENT** keyword

```
DROP GRAPH <http://musicbrainz.org/20130302>
```

- **DROP** removes the given graph from the graph store, including its content
- Can be used with the **DEFAULT** and **ALL** keywords

- COPY ... TO ...

```
COPY GRAPH <http://musicbrainz.org/20130302>  
TO GRAPH <http://musicbrainz.org/20130303>
```

- MOVE ... TO ...

```
MOVE GRAPH <http://musicbrainz.org/temp>  
TO GRAPH <http://musicbrainz.org/20130303>
```

- ADD ... TO ...

```
ADD GRAPH <http://musicbrainz.org/20130302>  
TO GRAPH <http://musicbrainz.org/20130303>
```

- SPARQL basics
 - Introductory notions
 - The SELECT query
 - Other query types
- Advanced SPARQL
 - CONSTRUCT queries
 - Filter & Order
 - Aggregate & Group
 - Union
 - Subqueries
 - Federated queries
 - Entailment regimes
- Updates with SPARQL

**188.399 Introduction to
Semantic Systems**
(2.0 VU / 3.0 ECTS)
2023W

**Katja Hose
Emanuel Sallinger**

iss@dbai.tuwien.ac.at

- Basics
 - Introductory notions
 - The SELECT query form
 - Other query forms
- Advanced notions
 - CONSTRUCT queries
 - Filter & Order
 - Aggregate & Group
 - Union
 - Subqueries
 - **Optional**
 - Federated queries
 - Entailment regimes

PREFIX ex: <http://ex.com/>

PREFIX foaf: <http://xmlns.com/foaf/0.1/>

SELECT ?s ?o ?c

WHERE {

 ?s foaf:name ?o.

 OPTIONAL { ?s ex:country ?c }.

}

ex:Anzhelika%20Sidorova rdf:type foaf:Person.
 ex:Sandi%20Morris rdf:type foaf:Person.
 ex:Katerina%20Stefanidi rdf:type foaf:Person.
 ex:Holly%20Bradshaw rdf:type foaf:Person.
 ex:Alysha%20Newman rdf:type foaf:Person.
 ex:Angelica%20Bengtsson rdf:type foaf:Person.

ex:Anzhelika%20Sidorova ex:score "4.95"^^xsd:decimal.
 ex:Sandi%20Morris ex:score "4.90"^^xsd:decimal.
 ex:Katerina%20Stefanidi ex:score "4.85"^^xsd:decimal.
 ex:Holly%20Bradshaw ex:score "4.80"^^xsd:decimal.
 ex:Alysha%20Newman ex:score "4.80"^^xsd:decimal.
 ex:Angelica%20Bengtsson ex:score "4.80"^^xsd:decimal.

ex:Anzhelika%20Sidorova foaf:name "Anzhelika Sidorova"@en.
 ex:Sandi%20Morris foaf:name "Sandi Morris"@en.
 ex:Katerina%20Stefanidi foaf:name "Katerina Stefanidi"@en.
 ex:Holly%20Bradshaw foaf:name "Holly Bradshaw"@en.
 ex:Alysha%20Newman foaf:name "Alysha Newman"@en.
 ex:Angelica%20Bengtsson foaf:name "Angelica Bengtsson"@en.

ex:Anzhelika%20Sidorova ex:country <http://ex.com/RU>.
~~ex:Sandi%20Morris ex:country <http://ex.com/US>.~~
 ex:Katerina%20Stefanidi ex:country <http://ex.com/EL>.
 ex:Holly%20Bradshaw ex:country <http://ex.com/UK>.
 ex:Alysha%20Newman ex:country <http://ex.com/CA>.
 ex:Angelica%20Bengtsson ex:country <http://ex.com/SE>.

PREFIX ex: <http://ex.com/>

PREFIX foaf: <http://xmlns.com/foaf/0.1/>

SELECT ?s ?o ?c

WHERE {

 ?s foaf:name ?o.

 OPTIONAL { ?s ex:country ?c }.

}

ex:Anzhelika%20Sidorova rdf:type foaf:Person.
 ex:Sandi%20Morris rdf:type foaf:Person.
 ex:Katerina%20Stefanidi rdf:type foaf:Person.
 ex:Holly%20Bradshaw rdf:type foaf:Person.
 ex:Alysha%20Newman rdf:type foaf:Person.
 ex:Angelica%20Bengtsson rdf:type foaf:Person.

ex:Anzhelika%20Sidorova ex:score "4.95"^^xsd:decimal.
 ex:Sandi%20Morris ex:score "4.90"^^xsd:decimal.
 ex:Katerina%20Stefanidi ex:score "4.85"^^xsd:decimal.
 ex:Holly%20Bradshaw ex:score "4.80"^^xsd:decimal.
 ex:Alysha%20Newman ex:score "4.80"^^xsd:decimal.
 ex:Angelica%20Bengtsson ex:score "4.80"^^xsd:decimal.

ex:Anzhelika%20Sidorova foaf:name "Anzhelika Sidorova"@en.
 ex:Sandi%20Morris foaf:name "Sandi Morris"@en.
 ex:Katerina%20Stefanidi foaf:name "Katerina Stefanidi"@en.
 ex:Holly%20Bradshaw foaf:name "Holly Bradshaw"@en.
 ex:Alysha%20Newman foaf:name "Alysha Newman"@en.
 ex:Angelica%20Bengtsson foaf:name "Angelica Bengtsson"@en.

ex:Anzhelika%20Sidorova ex:country <http://ex.com/RU>.
~~ex:Sandi%20Morris ex:country <http://ex.com/US>.~~
 ex:Katerina%20Stefanidi ex:country <http://ex.com/EL>.
 ex:Holly%20Bradshaw ex:country <http://ex.com/UK>.
 ex:Alysha%20Newman ex:country <http://ex.com/CA>.
 ex:Angelica%20Bengtsson ex:country <http://ex.com/SE>.

?s	?o	?c
ex:Alysha%20Newman	"Alysha Newman"@en	ex:CA
ex:Angelica%20Bengtsson	"Angelica Bengtsson"@en	ex:SE
ex:Anzhelika%20Sidorova	"Anzhelika Sidorova"@en	ex:RU
ex:Holly%20Bradshaw	"Holly Bradshaw"@en	ex:UK
ex:Katerina%20Stefanidi	"Katerina Stefanidi"@en	ex:EL
ex:Sandi%20Morris	"Sandi Morris"@en	

PREFIX ex: <http://ex.com/>

PREFIX foaf: <http://xmlns.com/foaf/0.1/>

SELECT ?s ?o ?c

WHERE {

 ?s foaf:name ?o.

 ?s ex:country ?c .

}

ex:Anzhelika%20Sidorova rdf:type foaf:Person.
 ex:Sandi%20Morris rdf:type foaf:Person.
 ex:Katerina%20Stefanidi rdf:type foaf:Person.
 ex:Holly%20Bradshaw rdf:type foaf:Person.
 ex:Alysha%20Newman rdf:type foaf:Person.
 ex:Angelica%20Bengtsson rdf:type foaf:Person.

ex:Anzhelika%20Sidorova ex:score "4.95"^^xsd:decimal.
 ex:Sandi%20Morris ex:score "4.90"^^xsd:decimal.
 ex:Katerina%20Stefanidi ex:score "4.85"^^xsd:decimal.
 ex:Holly%20Bradshaw ex:score "4.80"^^xsd:decimal.
 ex:Alysha%20Newman ex:score "4.80"^^xsd:decimal.
 ex:Angelica%20Bengtsson ex:score "4.80"^^xsd:decimal.

ex:Anzhelika%20Sidorova foaf:name "Anzhelika Sidorova"@en.
 ex:Sandi%20Morris foaf:name "Sandi Morris"@en.
 ex:Katerina%20Stefanidi foaf:name "Katerina Stefanidi"@en.
 ex:Holly%20Bradshaw foaf:name "Holly Bradshaw"@en.
 ex:Alysha%20Newman foaf:name "Alysha Newman"@en.
 ex:Angelica%20Bengtsson foaf:name "Angelica Bengtsson"@en.

ex:Anzhelika%20Sidorova ex:country <http://ex.com/RU>.
~~ex:Sandi%20Morris ex:country <http://ex.com/US>.~~
 ex:Katerina%20Stefanidi ex:country <http://ex.com/EL>.
 ex:Holly%20Bradshaw ex:country <http://ex.com/UK>.
 ex:Alysha%20Newman ex:country <http://ex.com/CA>.
 ex:Angelica%20Bengtsson ex:country <http://ex.com/SE>.

?s	?o	?c
ex:Alysha%20Newman	"Alysha Newman"@en	ex:CA
ex:Angelica%20Bengtsson	"Angelica Bengtsson"@en	ex:SE
ex:Anzhelika%20Sidorova	"Anzhelika Sidorova"@en	ex:RU
ex:Holly%20Bradshaw	"Holly Bradshaw"@en	ex:UK
ex:Katerina%20Stefanidi	"Katerina Stefanidi"@en	ex:EL
ex:Sandi%20Morris	"Sandi Morris"@en	

PREFIX ex: <http://ex.com/>

PREFIX foaf: <http://xmlns.com/foaf/0.1/>

SELECT ?s ?o ?c

WHERE {

 ?s foaf:name ?o.

 ?s ex:country ?c .

}

ex:Anzhelika%20Sidorova rdf:type foaf:Person.
 ex:Sandi%20Morris rdf:type foaf:Person.
 ex:Katerina%20Stefanidi rdf:type foaf:Person.
 ex:Holly%20Bradshaw rdf:type foaf:Person.
 ex:Alysha%20Newman rdf:type foaf:Person.
 ex:Angelica%20Bengtsson rdf:type foaf:Person.

ex:Anzhelika%20Sidorova ex:score "4.95"^^xsd:decimal.
 ex:Sandi%20Morris ex:score "4.90"^^xsd:decimal.
 ex:Katerina%20Stefanidi ex:score "4.85"^^xsd:decimal.
 ex:Holly%20Bradshaw ex:score "4.80"^^xsd:decimal.
 ex:Alysha%20Newman ex:score "4.80"^^xsd:decimal.
 ex:Angelica%20Bengtsson ex:score "4.80"^^xsd:decimal.

ex:Anzhelika%20Sidorova foaf:name "Anzhelika Sidorova"@en.
 ex:Sandi%20Morris foaf:name "Sandi Morris"@en.
 ex:Katerina%20Stefanidi foaf:name "Katerina Stefanidi"@en.
 ex:Holly%20Bradshaw foaf:name "Holly Bradshaw"@en.
 ex:Alysha%20Newman foaf:name "Alysha Newman"@en.
 ex:Angelica%20Bengtsson foaf:name "Angelica Bengtsson"@en.

ex:Anzhelika%20Sidorova ex:country <http://ex.com/RU>.
~~ex:Sandi%20Morris ex:country <http://ex.com/US>.~~
 ex:Katerina%20Stefanidi ex:country <http://ex.com/EL>.
 ex:Holly%20Bradshaw ex:country <http://ex.com/UK>.
 ex:Alysha%20Newman ex:country <http://ex.com/CA>.
 ex:Angelica%20Bengtsson ex:country <http://ex.com/SE>.

?s	?o	?c
ex:Alysha%20Newman	"Alysha Newman"@en	ex:CA
ex:Angelica%20Bengtsson	"Angelica Bengtsson"@en	ex:SE
ex:Anzhelika%20Sidorova	"Anzhelika Sidorova"@en	ex:RU
ex:Holly%20Bradshaw	"Holly Bradshaw"@en	ex:UK
ex:Katerina%20Stefanidi	"Katerina Stefanidi"@en	ex:EL

PREFIX ex: <http://ex.com/>

PREFIX foaf: <http://xmlns.com/foaf/0.1/>

SELECT ?s ?o ?c ?sc

WHERE { ?s foaf:name ?o.

OPTIONAL { ?s ex:country ?c }.

OPTIONAL { ?s ex:score ?sc }.

}

ex:Anzhelika%20Sidorova rdf:type foaf:Person.
 ex:Sandi%20Morris rdf:type foaf:Person.
 ex:Katerina%20Stefanidi rdf:type foaf:Person.
 ex:Holly%20Bradshaw rdf:type foaf:Person.
 ex:Alysha%20Newman rdf:type foaf:Person.
 ex:Angelica%20Bengtsson rdf:type foaf:Person.

ex:Anzhelika%20Sidorova ex:score "4.95"^^xsd:decimal.
 ex:Sandi%20Morris ex:score "4.90"^^xsd:decimal.
 ex:Katerina%20Stefanidi ex:score "4.85"^^xsd:decimal.
 ex:Holly%20Bradshaw ex:score "4.80"^^xsd:decimal.
 ex:Alysha%20Newman ex:score "4.80"^^xsd:decimal.
 ex:Angelica%20Bengtsson ex:score "4.80"^^xsd:decimal.

ex:Anzhelika%20Sidorova foaf:name "Anzhelika Sidorova"@en.
 ex:Sandi%20Morris foaf:name "Sandi Morris"@en.
 ex:Katerina%20Stefanidi foaf:name "Katerina Stefanidi"@en.
 ex:Holly%20Bradshaw foaf:name "Holly Bradshaw"@en.
 ex:Alysha%20Newman foaf:name "Alysha Newman"@en.
 ex:Angelica%20Bengtsson foaf:name "Angelica Bengtsson"@en.

ex:Anzhelika%20Sidorova ex:country <http://ex.com/RU>.
~~ex:Sandi%20Morris ex:country <http://ex.com/US>.~~
 ex:Katerina%20Stefanidi ex:country <http://ex.com/EL>.
 ex:Holly%20Bradshaw ex:country <http://ex.com/UK>.
 ex:Alysha%20Newman ex:country <http://ex.com/CA>.
 ex:Angelica%20Bengtsson ex:country <http://ex.com/SE>.

PREFIX ex: <http://ex.com/>

PREFIX foaf: <http://xmlns.com/foaf/0.1/>

SELECT ?s ?o ?c ?sc

WHERE { ?s foaf:name ?o.

OPTIONAL { ?s ex:country ?c }.

OPTIONAL { ?s ex:score ?sc }.

}

ex:Anzhelika%20Sidorova rdf:type foaf:Person.
 ex:Sandi%20Morris rdf:type foaf:Person.
 ex:Katerina%20Stefanidi rdf:type foaf:Person.
 ex:Holly%20Bradshaw rdf:type foaf:Person.
 ex:Alysha%20Newman rdf:type foaf:Person.
 ex:Angelica%20Bengtsson rdf:type foaf:Person.

ex:Anzhelika%20Sidorova ex:score "4.95"^^xsd:decimal.
 ex:Sandi%20Morris ex:score "4.90"^^xsd:decimal.
 ex:Katerina%20Stefanidi ex:score "4.85"^^xsd:decimal.
 ex:Holly%20Bradshaw ex:score "4.80"^^xsd:decimal.
 ex:Alysha%20Newman ex:score "4.80"^^xsd:decimal.
 ex:Angelica%20Bengtsson ex:score "4.80"^^xsd:decimal.

ex:Anzhelika%20Sidorova foaf:name "Anzhelika Sidorova"
 ex:Sandi%20Morris foaf:name "Sandi Morris"@en.
 ex:Katerina%20Stefanidi foaf:name "Katerina Stefanidi"
 ex:Holly%20Bradshaw foaf:name "Holly Bradshaw"@en.
 ex:Alysha%20Newman foaf:name "Alysha Newman"@en.
 ex:Angelica%20Bengtsson foaf:name "Angelica Bengtsson"

ex:Anzhelika%20Sidorova ex:country <http://ex.com/RU>
~~ex:Sandi%20Morris ex:country <http://ex.com/US>~~
 ex:Katerina%20Stefanidi ex:country <http://ex.com/EL>
 ex:Holly%20Bradshaw ex:country <http://ex.com/UK>
 ex:Alysha%20Newman ex:country <http://ex.com/CA>
 ex:Angelica%20Bengtsson ex:country <http://ex.com/SE>

?s	?o	?c	?sc
ex:Anzhelika%20Sidorova	"Anzhelika Sidorova"@en	ex:RU	"4.95"^^xsd:decimal
ex:Sandi%20Morris	"Sandi Morris"@en		"4.90"^^xsd:decimal
ex:Katerina%20Stefanidi	"Katerina Stefanidi"@en	ex:EL	"4.85"^^xsd:decimal
ex:Holly%20Bradshaw	"Holly Bradshaw"@en	ex:UK	"4.80"^^xsd:decimal
ex:Alysha%20Newman	"Alysha Newman"@en	ex:CA	"4.80"^^xsd:decimal
ex:Angelica%20Bengtsson	"Angelica Bengtsson"@en	ex:SE	"4.80"^^xsd:decimal

PREFIX ex: <http://ex.com/>

PREFIX foaf: <http://xmlns.com/foaf/0.1/>

SELECT ?s ?o ?c ?sc

WHERE { ?s foaf:name ?o.

OPTIONAL { ?s ex:country ?c .

?s ex:score ?sc }.

}

ex:Anzhelika%20Sidorova rdf:type foaf:Person.
 ex:Sandi%20Morris rdf:type foaf:Person.
 ex:Katerina%20Stefanidi rdf:type foaf:Person.
 ex:Holly%20Bradshaw rdf:type foaf:Person.
 ex:Alysha%20Newman rdf:type foaf:Person.
 ex:Angelica%20Bengtsson rdf:type foaf:Person.

ex:Anzhelika%20Sidorova ex:score "4.95"^^xsd:decimal.
 ex:Sandi%20Morris ex:score "4.90"^^xsd:decimal.
 ex:Katerina%20Stefanidi ex:score "4.85"^^xsd:decimal.
 ex:Holly%20Bradshaw ex:score "4.80"^^xsd:decimal.
 ex:Alysha%20Newman ex:score "4.80"^^xsd:decimal.
 ex:Angelica%20Bengtsson ex:score "4.80"^^xsd:decimal.

?s	?o	?c	?sc
ex:Anzhelika%20Sidorova	"Anzhelika Sidorova"@en	ex:RU	"4.95"^^xsd:decimal
ex:Sandi%20Morris	"Sandi Morris"@en		"4.90"^^xsd:decimal
ex:Katerina%20Stefanidi	"Katerina Stefanidi"@en	ex:EL	"4.85"^^xsd:decimal
ex:Holly%20Bradshaw	"Holly Bradshaw"@en	ex:UK	"4.80"^^xsd:decimal
ex:Alysha%20Newman	"Alysha Newman"@en	ex:CA	"4.80"^^xsd:decimal
ex:Angelica%20Bengtsson	"Angelica Bengtsson"@en	ex:SE	"4.80"^^xsd:decimal

ex:Anzhelika%20Sidorova foaf:name "Anzhelika Sidorova"
 ex:Sandi%20Morris foaf:name "Sandi Morris"@en.
 ex:Katerina%20Stefanidi foaf:name "Katerina Stefanidi"
 ex:Holly%20Bradshaw foaf:name "Holly Bradshaw"@
 ex:Alysha%20Newman foaf:name "Alysha Newman"@
 ex:Angelica%20Bengtsson foaf:name "Angelica Bengtsson"

ex:Anzhelika%20Sidorova ex:country <http://ex.com/RU>
~~ex:Sandi%20Morris ex:country <http://ex.com/US>~~
 ex:Katerina%20Stefanidi ex:country <http://ex.com/EL>
 ex:Holly%20Bradshaw ex:country <http://ex.com/UK>
 ex:Alysha%20Newman ex:country <http://ex.com/CA>
 ex:Angelica%20Bengtsson ex:country <http://ex.com/SE>.

PREFIX ex: <http://ex.com/>

PREFIX foaf: <http://xmlns.com/foaf/0.1/>

SELECT ?s ?o ?c ?sc

WHERE { ?s foaf:name ?o.

OPTIONAL { ?s ex:country ?c .

?s ex:score ?sc }.

}

ex:Anzhelika%20Sidorova rdf:type foaf:Person.
 ex:Sandi%20Morris rdf:type foaf:Person.
 ex:Katerina%20Stefanidi rdf:type foaf:Person.
 ex:Holly%20Bradshaw rdf:type foaf:Person.
 ex:Alysha%20Newman rdf:type foaf:Person.
 ex:Angelica%20Bengtsson rdf:type foaf:Person.

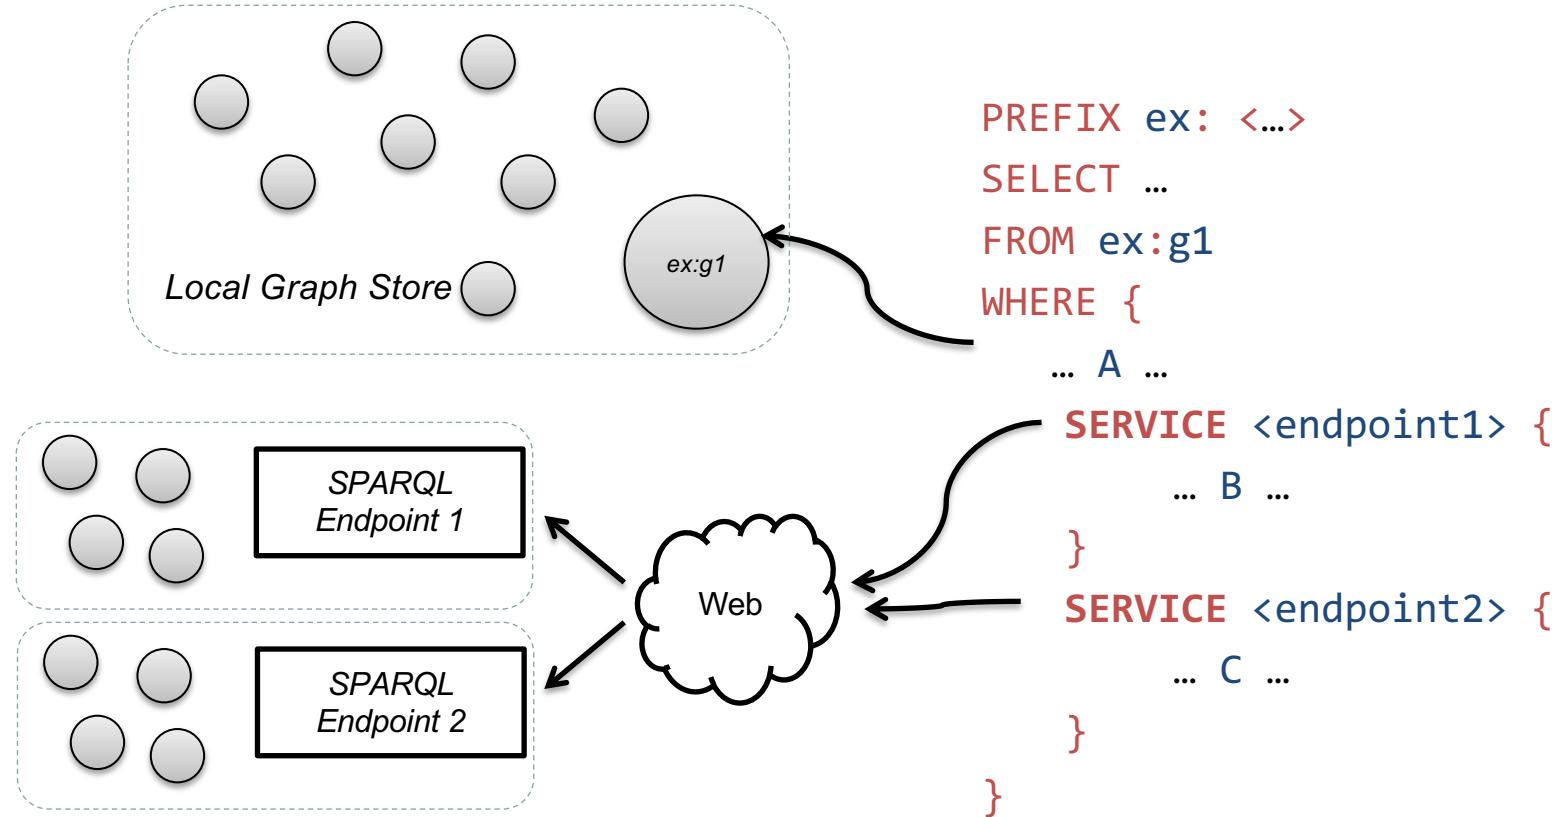
ex:Anzhelika%20Sidorova ex:score "4.95"^^xsd:decimal.
 ex:Sandi%20Morris ex:score "4.90"^^xsd:decimal.
 ex:Katerina%20Stefanidi ex:score "4.85"^^xsd:decimal.
 ex:Holly%20Bradshaw ex:score "4.80"^^xsd:decimal.
 ex:Alysha%20Newman ex:score "4.80"^^xsd:decimal.
 ex:Angelica%20Bengtsson ex:score "4.80"^^xsd:decimal.

?s	?o	?c	?sc
ex:Anzhelika%20Sidorova	"Anzhelika Sidorova"@en	ex:RU	"4.95"^^xsd:decimal
ex:Sandi%20Morris	"Sandi Morris"@en		
ex:Katerina%20Stefanidi	"Katerina Stefanidi"@en	ex:EL	"4.85"^^xsd:decimal
ex:Holly%20Bradshaw	"Holly Bradshaw"@en	ex:UK	"4.80"^^xsd:decimal
ex:Alysha%20Newman	"Alysha Newman"@en	ex:CA	"4.80"^^xsd:decimal
ex:Angelica%20Bengtsson	"Angelica Bengtsson"@en	ex:SE	"4.80"^^xsd:decimal

ex:Anzhelika%20Sidorova foaf:name "Anzhelika Sidorova"
 ex:Sandi%20Morris foaf:name "Sandi Morris"@en.
 ex:Katerina%20Stefanidi foaf:name "Katerina Stefanidi"
 ex:Holly%20Bradshaw foaf:name "Holly Bradshaw"@
 ex:Alysha%20Newman foaf:name "Alysha Newman"@
 ex:Angelica%20Bengtsson foaf:name "Angelica Bengtsson"

ex:Anzhelika%20Sidorova ex:country <http://ex.com/RU>
~~ex:Sandi%20Morris ex:country <http://ex.com/US>~~
 ex:Katerina%20Stefanidi ex:country <http://ex.com/EL>
 ex:Holly%20Bradshaw ex:country <http://ex.com/UK>
 ex:Alysha%20Newman ex:country <http://ex.com/CA>
 ex:Angelica%20Bengtsson ex:country <http://ex.com/SE>.

- Basics
 - Introductory notions
 - The SELECT query form
 - Other query forms
- Advanced notions
 - CONSTRUCT queries
 - Filter & Order
 - Aggregate & Group
 - Union
 - Subqueries
 - Optional
 - **Federated queries**
 - Entailment regimes



Find the birth dates of all of the actors in *Star Trek: The Motion Picture*

PREFIX movie: <http://data.linkedmdb.org/resource/movie/>

PREFIX dbpedia: <http://dbpedia.org/ontology/>

PREFIX foaf: <http://xmlns.com/foaf/0.1/>

SELECT ?actor_name ?birth_date

FROM <http://dig.csail.mit.edu/2008/webdav/timbl/foaf.rdf> # placeholder graph

WHERE {

```
{
  SERVICE <http://data.linkedmdb.org/sparql> {
    <http://data.linkedmdb.org/resource/film/675> movie:actor ?actor .
    ?actor movie:actor_name ?actor_name
  }
  BIND(STRLANG(?actor_name, "en") AS ?actor_name_en)
}
```

```
  SERVICE <http://dbpedia.org/sparql> {
    ?actor2 a foaf:Person ;
      foaf:name ?actor_name_en ;
      dbpedia:birthDate ?birth_date
  }
}
```

?actor_name	? birth_date
"William Shatner"	"1931-03-22" ^^<http://www.w3.org/2001/XMLSchema#date>
"Leonard Nimoy"	"1931-03-26" ^^<http://www.w3.org/2001/XMLSchema#date>
"DeForest Kelley"	"1920-01-20" ^^<http://www.w3.org/2001/XMLSchema#date>
"James Doohan"	"1920-03-03" ^^<http://www.w3.org/2001/XMLSchema#date>
"Nichelle Nichols"	"1932-12-28" ^^<http://www.w3.org/2001/XMLSchema#date>
"Stephen Collins"	"1947-10-01" ^^<http://www.w3.org/2001/XMLSchema#date>
"George Takei"	"1937-04-20" ^^<http://www.w3.org/2001/XMLSchema#date>


- Basics:
 - Introductory notions
 - The SELECT query form
 - Other query forms
- Advanced notions
 - CONSTRUCT queries
 - Filter & Order
 - Aggregate & Group
 - Union
 - Subqueries
 - Optional
 - Federated queries
 - **Entailment regimes**

- A triple store can use the knowledge encoded in the schema to infer new facts

Schema: `mo:MusicGroup rdfs:subClassOf mo:MusicArtist .`

Inferred facts: `mo:MusicGroup a rdfs:Class .`
`mo:MusicArtist a rdfs:Class .`

- Whether a triple store supports reasoning, influences the outcome of querying



```
xxx rdfs:subClassOf yyy .
zzz rdf:type xxx .
```

```
zzz rdf:type yyy .
```

Schema:

Query:

```
mo:MusicGroup rdfs:subClassOf mo:MusicArtist .
```

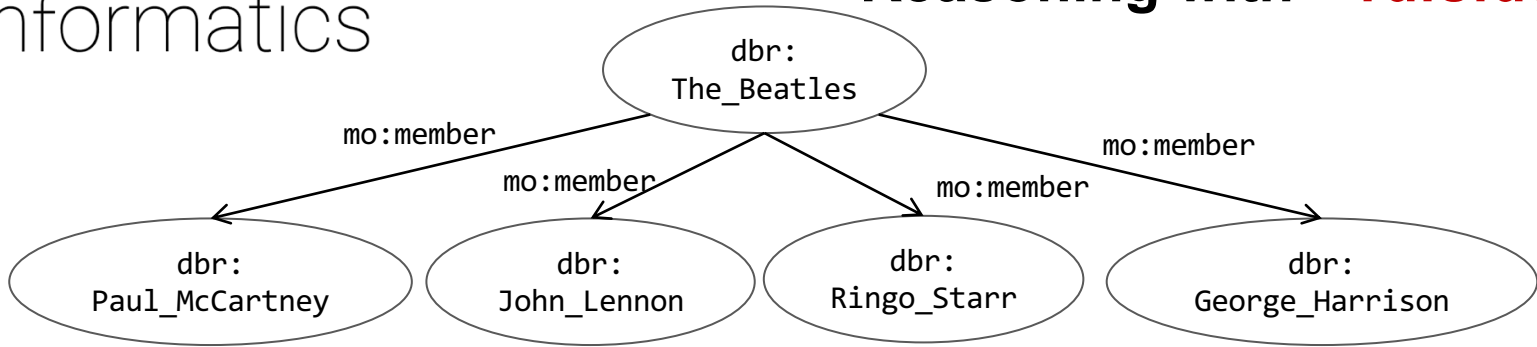
```
SELECT ?x
WHERE {?x a mo:MusicArtist.}
```

Result set **with inference:**

?x
dbr:The_Beatles ...

Result set **without inference:**

?x



Schema:

```
mo:member rdfs:domain mo:MusicGroup .
```

Query:

```
SELECT ?x
WHERE {?x a mo:MusicGroup.}
```

Result set **with inference**:

?x
dbr:The_Beatles ...


Result set without interference:

?x

```
aaa rdfs:domain xxx .
yyy aaa zzz .
```



```
yyy rdf:type xxx .
```



Schema:

```
dbo:bandMember rdfs:range foaf:Musician .
```

Query:

```
SELECT ?x
WHERE {?x a foaf:Musician .}
```

Result set with inference:

?x
dbr:Paul_McCartney
dbr:John_Lennon
dbr:Ringo_Starr
dbr:George_Harrison ...

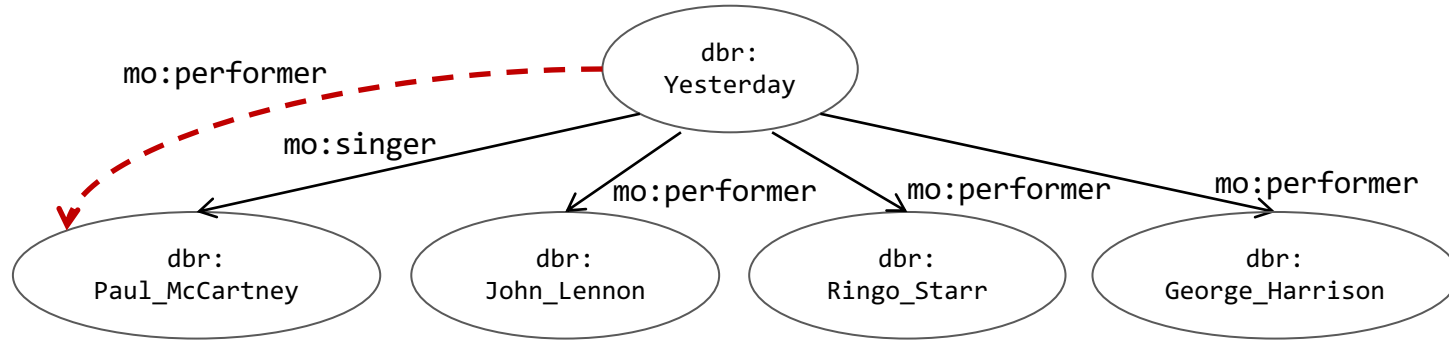
Result set without inference:

?x

```
aaa rdfs:range xxx .
yyy aaa zzz .
```



```
zzz rdf:type xxx .
```



Schema:

Query:

```
mo:singer rdfs:subPropertyOf mo:performer .
```

```
SELECT ?x
WHERE {dbr:Yesterday mo:performer ?x .}
```

Result set **with inference**:

?x
dbr:John_Lennon
dbr:Ringo_Starr
dbr:George_Harrison
dbr:Paul_McCartney

Result set without inference:

?x
dbr:John_Lennon
dbr:Ringo_Starr
dbr:George_Harrison

```
aaa rdfs:subPropertyOf bbb .
yyy aaa zzz .
```



```
yyy bbb zzz .
```

	If S contains:	then S RDFS entails recognizing D:
<i>rdfs1</i>	any IRI aaa in D	aaa <code>rdf:type rdfs:Datatype .</code>
<i>rdfs2</i>	aaa <code>rdfs:domain xxx .</code> yyy aaa zzz .	yyy <code>rdf:type xxx .</code>
<i>rdfs3</i>	aaa <code>rdfs:range xxx .</code> yyy aaa zzz .	zzz <code>rdf:type xxx .</code>
<i>rdfs4a</i>	xxx aaa yyy .	xxx <code>rdf:type rdfs:Resource .</code>
<i>rdfs4b</i>	xxx aaa yyy .	yyy <code>rdf:type rdfs:Resource .</code>
<i>rdfs5</i>	xxx <code>rdfs:subPropertyOf yyy .</code> yyy <code>rdfs:subPropertyOf zzz .</code>	xxx <code>rdfs:subPropertyOf zzz .</code>
<i>rdfs6</i>	xxx <code>rdf:type rdf:Property .</code>	xxx <code>rdfs:subPropertyOf xxx .</code>
<i>rdfs7</i>	aaa <code>rdfs:subPropertyOf bbb .</code> xxx aaa yyy .	xxx bbb yyy .
<i>rdfs8</i>	xxx <code>rdf:type rdfs:Class .</code>	xxx <code>rdfs:subClassOf rdfs:Resource .</code>
<i>rdfs9</i>	xxx <code>rdfs:subClassOf yyy .</code> zzz <code>rdf:type xxx .</code>	zzz <code>rdf:type yyy .</code>
<i>rdfs10</i>	xxx <code>rdf:type rdfs:Class .</code>	xxx <code>rdfs:subClassOf xxx .</code>
<i>rdfs11</i>	xxx <code>rdfs:subClassOf yyy .</code> yyy <code>rdfs:subClassOf zzz .</code>	xxx <code>rdfs:subClassOf zzz .</code>
<i>rdfs12</i>	xxx <code>rdf:type rdfs:ContainerMembershipProperty .</code>	xxx <code>rdfs:subPropertyOf rdfs:member .</code>
<i>rdfs13</i>	xxx <code>rdf:type rdfs:Datatype .</code>	xxx <code>rdfs:subClassOf rdfs:Literal .</code>

S.. subgraph

D.. Set of IRIs identifying datatypes

Updating RDF with SPARQL 1.1



- SPARQL 1.0 only allows data access (querying)
- SPARQL 1.1 introduces:

Query extensions

Aggregates, subqueries, negation, expressions in the **SELECT** clause, property paths, assignment, expanded set of functions and operators

Updates

Insert, Delete, Delete/Insert

**Create, Load, Clear, Drop
Copy, Move, Add**

Federation extension

Service, values, service variables (informative)

SPARQL 1.1 provides data update operations

- **INSERT (DATA)**: adds triples, given inline in the request, into a graph
- **DELETE (DATA)**: removes triples, given inline in the request, if the respective graphs contains those
- **DELETE/INSERT (DATA)**: uses in parallel **INSERT** and **DELETE**

INSERT + graph pattern specifying a set of RDF triples, given bindings for any variables that are included in the pattern

Insert "Peter Best" as a formerMember of dbr:The_Beatles

```

PREFIX dbr: <http://dbpedia.org/resource/>
PREFIX dbo: <http://dbpedia.org/ontology/>
PREFIX foaf: <http://xmlns.com/foaf/0.1/>

INSERT { dbr:The_Beatles dbr:formerMember ?x }
WHERE { dbr:The_Beatles dbr:currentMember ?x .
        ?x foaf:name "Peter Best"
      }
  
```


Insert the following albums recorded by The Beatles into the graph
<http://musicbrainz.org/20130209-004702>

```
PREFIX dc: <http://purl.org/dc/elements/1.1/>
```

```
PREFIX foaf: <http://xmlns.com/foaf/0.1/>
```

```
INSERT DATA {
```

```
    GRAPH { <http://musicbrainz.org/20130302>
```

```
    <http://musicbrainz.org/artist/b10bbbfc-cf9e-42e0-be17-e2c3e1d2600d>
```

```
        foaf:made <http://musicbrainz.org/release/3a685770-7326-34fc-9f18-e5f5626f3dc5> ,
```

```
        <http://musicbrainz.org/release/cb6f8798-d51e-4fa5-a4d1-2c0602bfe1b6> .
```

```
    <http://musicbrainz.org/release/3a685770-7326-34fc-9f18-e5f5626f3dc5>
```

```
        dc:title "Please Please Me".
```

```
    <http://musicbrainz.org/release/cb6f8798-d51e-4fa5-a4d1-2c0602bfe1b6>
```

```
        dc:title "Something New". } }
```

Delete all the information about the album "Casualities" of The Beatles.

```
PREFIX dc: <http://purl.org/dc/elements/1.1/>
```

```
PREFIX foaf: <http://xmlns.com/foaf/0.1/>
```

```
DELETE { ?album ?predicate ?object . }
```

```
WHERE {
```

```
http://musicbrainz.org/artist/b10bbbfbc-cf9e-42e0-be17-e2c3e1d2600d foaf:made  
    ?album .
```

```
?album dc:title "Casualities".
```

```
?album ?predicate ?object .}
```

Delete the status of 'Peter Best' as current member of "The Beatles" and **insert** his status as former member of the band.

```
PREFIX dbr, dbo, foaf ...
```

```
DELETE { dbr:The_Beatles dbo:currentMember ?x . }
```

```
INSERT { GRAPH <http://musicbrainz.org/20130209-004702> {
  dbr:The_Beatles dbo:formerBandMember ?x .}
```

```
}
```

```
WHERE {
```

```
  dbr:The_Beatles dbo:currentMember ?x .
```

```
  ?x foaf:name "Peter Best" .}
```

SPARQL 1.1 provides graph update operations

- **CREATE** creates an empty graph in the Graph Store
- **LOAD** reads the content of a document into a graph in the Graph Store
- **CLEAR** removes all triples in one or more graphs
- **DROP** removes the graph from the Graph Store
- Other operations: **COPY**, **MOVE**, **ADD**

- Creates a new named graph
- Can be used with the **DEFAULT** and **ALL** keywords

```
CREATE GRAPH <http://musicbrainz.org/20130302>
```

- **LOAD**: An RDF graph can be loaded from a URL

```
LOAD <http://xmlns.com/foaf/spec/20100809.rdf>
```

```
LOAD <http://xmlns.com/foaf/spec/20100809.rdf>  
    INTO <http://xmlns.com/foaf/0.1/>
```

Named graph

- **LOAD** can be used with the **SILENT** keyword (suppresses errors)

```
CLEAR GRAPH <http://musicbrainz.org/20130302>
```

- **CLEAR** removes all triples from the graph (emptied, but not deleted!)
- The graph(s) can be specified with the following keywords **DEFAULT**, **NAMED**, **ALL**, **GRAPH**
- Can be used with the **SILENT** keyword

```
DROP GRAPH <http://musicbrainz.org/20130302>
```

- **DROP** removes the given graph from the graph store, including its content
- Can be used with the **DEFAULT** and **ALL** keywords

- COPY ... TO ...

```
COPY GRAPH <http://musicbrainz.org/20130302>  
TO GRAPH <http://musicbrainz.org/20130303>
```

- MOVE ... TO ...

```
MOVE GRAPH <http://musicbrainz.org/temp>  
TO GRAPH <http://musicbrainz.org/20130303>
```

- ADD ... TO ...

```
ADD GRAPH <http://musicbrainz.org/20130302>  
TO GRAPH <http://musicbrainz.org/20130303>
```

- SPARQL basics
 - Introductory notions
 - The SELECT query
 - Other query types
- Advanced SPARQL
 - CONSTRUCT queries
 - Filter & Order
 - Aggregate & Group
 - Union
 - Subqueries
 - Optional
 - Federated queries
 - Entailment regimes
- Updates with SPARQL

188.399 Introduction to Semantic Systems

(2.0 VU / 3.0 ECTS)

2023W

**Katja Hose
Emanuel Sallinger**

iss@dbai.tuwien.ac.at



- SHACL – Shapes Constraint Language

RDF Quality and SHACL



Language for validating RDF graphs against a set of **conditions**


W3C recommendation since July 2017: <https://www.w3.org/TR/shacl/>

data graph: RDF graph to be validated against a shapes graph

shapes graph: RDF graphs with **conditions** provided as shapes and other constructs
descriptions of the data graphs that do satisfy these conditions

Language for validating RDF graphs against a set of conditions

W3C recommendation since July 2017



Shape graph

```
ex:PersonShape a sh:NodeShape ;  
sh:targetClass ex:Person ;  
sh:property [  
  a sh:PropertyShape;  
  sh:path ex:ssn;  
  sh:minCount 1 ] .
```

sh:minCount specifies the minimum number of value nodes that satisfy the condition.

Data graph

```
ex:Alice rdf:type ex:Person .  
ex:Alice ex:ssn "987-65-432A" .
```

```
ex:Bob rdf:type ex:Person .  
ex:Bob ex:ssn "123-45-6789" .  
ex:Bob ex:ssn "124-35-6789" .
```


```
ex:Calvin rdf:type ex:Person .  
ex:Calvin ex:birthDate "1971-07-07"^^xsd:date .  
ex:Calvin ex:worksFor ex:UntypedCompany .
```

- **Targets** specify what RDF graph nodes are validated against a shape
- During validation, we refer to the targets as **focus nodes**


Targets in SHACL Core

Value	Description
targetNode	directly specified nodes
targetClass	All instances of a class
targetSubjectOf	All subjects of a specific property
targetObjectOf	All objects of a specific property

Data graph



Apply targets to select focus nodes



Target definitions in a shape graph

ex:PersonShape sh:targetNode ex:Alice .

ex:PersonShape sh:targetClass ex:Person .

ex:PersonShape sh:targetSubjectsOf ex:ssn .

ex:PersonShape sh:targetObjectsOf ex:worksFor .

Data Graph

ex:Alice rdf:type ex:Person .

ex:Alice ex:ssn "987-65-432A".

ex:Bob rdf:type ex:Person .

ex:Bob ex:ssn "123-45-6789" .

ex:Bob ex:ssn "124-35-6789" .

ex:Calvin rdf:type ex:Person .

ex:Calvin ex:birthDate "1971-07-07"^^xsd:date .

ex:Calvin ex:worksFor ex:UntypedCompany .

Shape Graph

ex:PersonShape rdf:type sh:NodeShape .

ex:PersonShape sh:targetNode **ex:Alice** .

ex:PersonShape sh:targetClass ex:Person .

ex:PersonShape sh:targetSubjectsOf ex:ssn .

ex:PersonShape sh:targetObjectsOf ex:worksFor .

Data Graph

ex:Alice rdf:type ex:Person .

ex:Alice ex:ssn "987-65-432A" .

ex:Bob rdf:type ex:Person .

ex:Bob ex:ssn "123-45-6789" .

ex:Bob ex:ssn "124-35-6789" .

ex:Calvin rdf:type ex:Person .

ex:Calvin ex:birthDate "1971-07-07"^^xsd:date .

ex:Calvin ex:worksFor ex:UntypedCompany .

Shape Graph

ex:PersonShape rdf:type sh:NodeShape .

ex:PersonShape sh:targetNode **ex:Alice** .

ex:PersonShape rdf:type sh:NodeShape .

ex:PersonShape sh:targetClass ex:Person .

ex:PersonShape sh:targetSubjectsOf ex:ssn .

ex:PersonShape sh:targetObjectsOf ex:worksFor .

Data Graph

ex:Alice rdf:type ex:Person .

ex:Alice ex:ssn "987-65-432A" .

ex:Bob rdf:type ex:Person .

ex:Bob ex:ssn "123-45-6789" .

ex:Bob ex:ssn "124-35-6789" .

ex:Calvin rdf:type ex:Person .

ex:Calvin ex:birthDate "1971-07-07"^^xsd:date .

ex:Calvin ex:worksFor ex:UntypedCompany .

Shape Graph

ex:PersonShape rdf:type sh:NodeShape .

ex:PersonShape sh:targetNode **ex:Alice** .

ex:PersonShape rdf:type sh:NodeShape .

ex:PersonShape sh:targetClass **ex:Person** .

ex:PersonShape rdf:type sh:NodeShape .

ex:PersonShape sh:targetSubjectsOf ex:ssn .

ex:PersonShape sh:targetObjectsOf ex:worksFor .

Data Graph

ex:Alice rdf:type ex:Person .

ex:Alice ex:ssn "987-65-432A" .

ex:Bob rdf:type ex:Person .

ex:Bob ex:ssn "123-45-6789" .

ex:Bob ex:ssn "124-35-6789" .

ex:Calvin rdf:type ex:Person .

ex:Calvin ex:birthDate "1971-07-07"^^xsd:date .

ex:Calvin ex:worksFor ex:UntypedCompany .

Shape Graph

ex:PersonShape rdf:type sh:NodeShape .

ex:PersonShape sh:targetNode **ex:Alice** .

ex:PersonShape rdf:type sh:NodeShape .

ex:PersonShape sh:targetClass **ex:Person** .

ex:PersonShape rdf:type sh:NodeShape .

ex:PersonShape sh:targetSubjectsOf **ex:ssn** .

ex:PersonShape rdf:type sh:NodeShape .

ex:PersonShape sh:targetObjectsOf ex:worksFor .

Data Graph

ex:Alice rdf:type ex:Person .

ex:Alice ex:ssn "987-65-432A" .

ex:Bob rdf:type ex:Person .

ex:Bob ex:ssn "123-45-6789" .

ex:Bob ex:ssn "124-35-6789" .

ex:Calvin rdf:type ex:Person .

ex:Calvin ex:birthDate "1971-07-07"^^xsd:date .

ex:Calvin ex:worksFor ex:UntypedCompany .

Shape Graph

ex:PersonShape rdf:type sh:NodeShape .

ex:PersonShape sh:targetNode **ex:Alice** .

ex:PersonShape rdf:type sh:NodeShape .

ex:PersonShape sh:targetClass **ex:Person** .

ex:PersonShape rdf:type sh:NodeShape .

ex:PersonShape sh:targetSubjectsOf **ex:ssn** .

ex:PersonShape rdf:type sh:NodeShape .

ex:PersonShape sh:targetObjectsOf **ex:worksFor** .

Data Graph

ex:Alice rdf:type ex:Person .

ex:Alice ex:ssn "987-65-432A" .

ex:Bob rdf:type ex:Person .

ex:Bob ex:ssn "123-45-6789" .

ex:Bob ex:ssn "124-35-6789" .

ex:Calvin rdf:type ex:Person .

ex:Calvin ex:birthDate "1971-07-07"^^xsd:date .

ex:Calvin ex:worksFor **ex:UntypedCompany** .

Node Shapes specify constraints that need to be met by the focus nodes

ex:PersonShape rdf:type sh:NodeShape

Property Shapes specify constraints on the values of particular properties/paths of the focus nodes

ex:PersonPropertyShape a sh:PropertyShape

Shape Graph

```
ex:PersonShape rdf:type sh:NodeShape .  
ex:PersonShape sh:targetClass ex:Person .  
ex:PersonShape sh:property [  
  rdf:type sh:PropertyShape ;  
  sh:path ex:ssn ] .
```

What does this shape check?

Data Graph

```
ex:Alice rdf:type ex:Person .  
ex:Alice ex:ssn "987-65-432A".  
  
ex:Bob rdf:type ex:Person .  
ex:Bob ex:ssn "123-45-6789" .  
ex:Bob ex:ssn "124-35-6789" .
```

```
ex:Calvin rdf:type ex:Person .  
ex:Calvin ex:birthDate "1971-07-07"^^xsd:date .  
ex:Calvin ex:worksFor ex:UntypedCompany .
```

Shape Graph

```
ex:PersonShape rdf:type sh:NodeShape .
ex:PersonShape sh:targetClass ex:Person .
ex:PersonShape sh:property [
  rdf:type sh:PropertyShape ;
  sh:path ex:ssn ] .
```

How can we extend this example to check that `ex:Person` can have at most 1 value for the property `ex:ssn`?

Data Graph

```
ex:Alice rdf:type ex:Person .
ex:Alice ex:ssn "987-65-432A" .
```

```
ex:Bob rdf:type ex:Person .
ex:Bob ex:ssn "123-45-6789" .
ex:Bob ex:ssn "124-35-6789" .
```

```
ex:Calvin rdf:type ex:Person .
ex:Calvin ex:birthDate "1971-07-07"^^xsd:date .
ex:Calvin ex:worksFor ex:UntypedCompany .
```


Shape Graph

```
ex:PersonShape rdf:type sh:NodeShape .
ex:PersonShape sh:targetClass ex:Person .
ex:PersonShape sh:property [
  rdf:type sh:PropertyShape ;
  sh:path ex:ssn ;
  sh:maxCount 1 ].
```

Data Graph

```
ex:Alice rdf:type ex:Person .
ex:Alice ex:ssn "987-65-432A".


ex:Bob rdf:type ex:Person .
ex:Bob ex:ssn "123-45-6789" .
ex:Bob ex:ssn "124-35-6789" .
```

```
ex:Calvin rdf:type ex:Person .
ex:Calvin ex:birthDate "1971-07-07"^^xsd:date .
ex:Calvin ex:worksFor ex:UntypedCompany .
```


Type	Constraint
Cardinality	minCount, maxCount
Values	node, in, hasValue
Value range	minInclusive, maxInclusive, minExclusive, maxExclusive
String-based	minLength, maxLength, pattern, languageIn, uniqueLang
Logical constraints	not, and, or, xone
Closed shapes	closed, ignoredProperties
Property pair constraints	equals, disjoint, lessThan, lessThanOrEquals
Non-validating constraints	name, description, order, group, defaultValue
...	...

Online survey


- Common practice to create shapes manually
- Automatic extraction often has scalability issues
- Automatic extraction results in many spurious shapes




(a) Area of occupation




(b) Size of graph (Number of triples)




(c) Size of ontology (Number of classes)



(d) Distinct # of properties in the graph



(e) Number of generated shapes



(f) Scope of the generated shapes

- SHACL constraints can help improve the quality of knowledge graphs
- Shapes can be used to define a broad range of constraints
- Validation reports can be used to increase the quality of an RDF graph

Knowledge Graph Creation

Emanuel Sallinger
Katja Hose

Option 1: Template-based on-the-fly generation

Typical web application architecture:

- Data resides in back-end databases
- Front-end translates entries via templates into HTML, JSON,.. representations

Use, e.g., JSON-LD (cf. RDF notations) to impose semantics on JSON documents by providing a context that identifies properties and structural elements

→ can be interpreted as plain JSON document or as RDF graphs and triples

```
{
  "@context": "http://schema.org/",
  "@id": "http://dbpedia.org/resource/Tim_Berners-Lee",
  "givenName": "Tim",
  "knows": [{
    "@id": "http://dbpedia.org/resource/Ted_Nelson",
    "givenName": "Ted"
  }]
}
```

PREFIX dbr: <http://dbpedia.org/resource/>
PREFIX schema: <http://schema.org>

=

```
dbr:Ted_Nelson      schema:givenName  "Ted".
dbr:Tim_Berners-Lee schema:givenName  "Tim";
                   schema:knows    dbr:Ted_Nelson.
```

- But:** not always sufficient..
- Remodeling data to RDF on the fly is tricky
 - Does not create links between resources
 - Link to external concepts?
 - Reuse of external identifiers?

RDF Notations: JSON-LD

- JavaScript Object Notation for Linked Data
- W3C recommendation
- Imposes structure on JSON with @context annotations (often external vocabularies)

```
{ "@context": {  
  "name": "http://xmlns.com/foaf/0.1/name",  
  "homepage": {  
    "@id": "http://xmlns.com/foaf/0.1/workplaceHomepage",  
    "@type": "@id"  
  },  
  "Person": "http://xmlns.com/foaf/0.1/Person"  
},  
"@id": "http://me.example.com",  
"@type": "Person",  
"name": "John Smith",  
"homepage": "http://www.example.com/"  
}
```

Option 2: Mapping and Transformation Tools




Tabular data




Databases



Structured Data



Text



Knowledge Graph Creation

Tabular Data to RDF: OpenRefine

Emanuel Sallinger
Katja Hose

Mapping and Transformation Tools



Tabular data



Example Tabular Data



```
The Beatles, 250 million
Elvis Presley, 203.3 million
Michael Jackson, 157.4 million
Madonna, 160.1 million
Led Zeppelin, 135.5 million
Queen, 90.5 million
```



```
{
  "artist": {
    "class": "artist",
    "name": "The Beatles"
  },
  "rank": 1,
  "value": 250 million
},
...
```

HTML tables

Artist	Country of origin	Period active	Release-year of first charted record	Total certified units (from available markets)^[Notes]
The Beatles	United Kingdom	1960–1970 ^[4]	1962 ^[4]	Total available certified units: 250 million [show]
Elvis Presley	United States	1954–1977 ^[28]	1954 ^[28]	Total available certified units: 203.3 million [show]
Michael Jackson ^[Note 2]	United States	1964–2009 ^[32]	1971 ^[32]	Total available certified units: 157.4 million [show]
Madonna	United States	1979–present ^[44]	1982 ^[44]	Total available certified units: 160.1 million [show]
Led Zeppelin	United Kingdom	1968–1980 ^[50]	1969 ^[50]	Total available certified units: 135.5 million [show]
Queen	United Kingdom	1971–present ^[53]	1973 ^[53]	Total available certified units: 90.5 million [show]

http://en.wikipedia.org/wiki/List_of_best-selling_music_artists

Extraction tools: OpenRefine



- Tool for **messy data**
- May look like a spreadsheet program, but
 - NOT for entering new data one cell at a time
 - NOT for accounting
- Key feature: **apply transformations** over many existing cells in bulk to
 - clean up data
 - integrate data from multiple sources
 - perform entity reconciliation (against SPARQL endpoint)
 - convert data (e.g., RDF plugin)
 - ... *among other things*

Using OpenRefine for KG Population




1. Import messy input data, transform it into a table, and clean it



	B	C	E	F	G	I	J
1	Series Title	Released Year	Runtime	Genre	IMDB Rating	Meta score	Director
2	The Shawshank Redemption	1994	142 min	Drama	9.3	80	Frank Darabont
3	The Godfather	1972	175 min	Crime, Drama	9.2	100	Francis Ford Coppola
4	The Dark Knight	2008	152 min	Action, Crime, Drama	9	84	Christopher Nolan
5	The Godfather: Part II	1974	202 min	Crime, Drama	9	90	Francis Ford Coppola
6	12 Angry Men	1957	96 min	Crime, Drama	9	96	Sidney Lumet

2. (optional) Apply entity reconciliation to interlink with existing data sets.



3. Define the structure of the RDF output.
4. Export the data into some RDF syntax.

```
<http://example-instance.org/The%20Shawshank%20Redemption> a mv:Movie;  
    mv:hasName "The Shawshank Redemption";  
    mv:hasReleaseDate "1994-01-01T00:00Z";  
    mv:hasGenre "Drama";  
    mv:hasMovieDirector <http://example-instance.org/Frank%20Darabont> .  
  
<http://example-instance.org/Frank%20Darabont> a mv:MovieDirector;  
    mv:hasName "Frank Darabont" .
```

Step 1 - Data Transformation

Typical steps:

- Group and explore data items
- Delete columns or rows based on filter condition
- Split columns into several columns based on condition
- Modify messy data items with **GREL**, a powerful expression language
- Replay steps from a previous Refine project

	Series_Title	Released_Year	Runtime	Genre
1.	The Shawshank Redemption	1994	142 min	Drama
2.	The Godfather	1972	175 min	Crime, Drama
3.	The Dark Knight	2008	152 min <small>edit</small>	Action, Crime, Drama
4.	The Godfather: Part II	1974	202 min	Crime, Drama

Transform column "Released_Year" to date and execute String processing to get value of the "Runtime" as integer

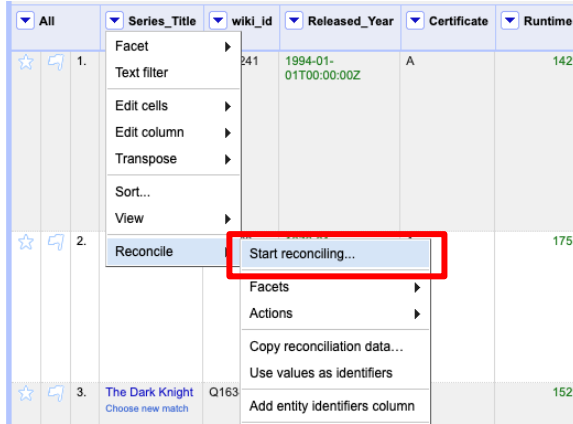
	Series_Title	Released_Year	Runtime	Genre
1.	The Shawshank Redemption	1994-01-01T00:00:00Z	142	Drama
2.	The Godfather	1972-01-01T00:00:00Z	175	Crime, Drama
3.	The Dark Knight	2008-01-01T00:00:00Z	152	Action, Crime, Drama
4.	The Godfather: Part II	1974-01-01T00:00:00Z	202	Crime, Drama

Split column "Genre" -> to create several "Genres"

	Series_Title	Released_Year	Runtime	Genre 1	Genre 2	Genre 3
1.	The Shawshank Redemption	1994-01-01T00:00:00Z	142	Drama		
2.	The Godfather	1972-01-01T00:00:00Z	175	Crime	Drama	
3.	The Dark Knight	2008-01-01T00:00:00Z	152	Action	Crime	Drama
4.	The Godfather: Part II	1974-01-01T00:00:00Z	202	Crime	Drama	

GREL = Google Refine Expression Language

Step 2 - Entity Reconciliation



All	Series_Title	wiki_id	Released_Year	Certificate	Runtime
1.	Facet	241	1994-01-01T00:00:00Z	A	142
2.	Reconcile				175
3.	The Dark Knight	Q163			152

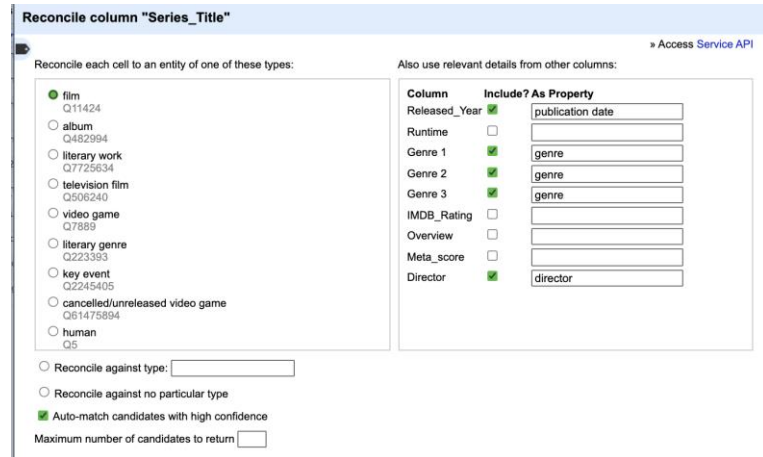
Reconciliation: process of matching text names to the corresponding semantic entities in another knowledge graph.

Typical steps:

- Define a reconciliation service (e.g. Wikidata)
- Select specific types to reconcile against (e.g., Film)
- Start reconciling a column against the service



	Series_Title	Wiki_ID	Released_Year
1.	The Shawshank Redemption Choose new match	Q172241	1994-01-01T00:00:00Z
2.	The Godfather Choose new match	Q47703	1972-01-01T00:00:00Z
3.	The Dark Knight Choose new match	Q163872	2008-01-01T00:00:00Z
4.	The Godfather Part II Choose new match	Q184768	1974-01-01T00:00:00Z



Reconcile column "Series_Title"

Reconcile each cell to an entity of one of these types:

- film (Q11424)
- album (Q482994)
- literary work (Q7725634)
- television film (Q506240)
- video game (Q7889)
- literary genre (Q223363)
- key event (Q2245405)
- cancelled/unreleased video game (Q61475894)
- human (Q5)

Reconcile against type:

Reconcile against no particular type

Auto-match candidates with high confidence

Maximum number of candidates to return

Also use relevant details from other columns:

Column	Include? As Property
Released_Year	<input checked="" type="checkbox"/> publication date
Runtime	<input type="checkbox"/>
Genre 1	<input checked="" type="checkbox"/> genre
Genre 2	<input checked="" type="checkbox"/> genre
Genre 3	<input checked="" type="checkbox"/> genre
IMDB_Rating	<input type="checkbox"/>
Overview	<input type="checkbox"/>
Meta_score	<input type="checkbox"/>
Director	<input checked="" type="checkbox"/> director

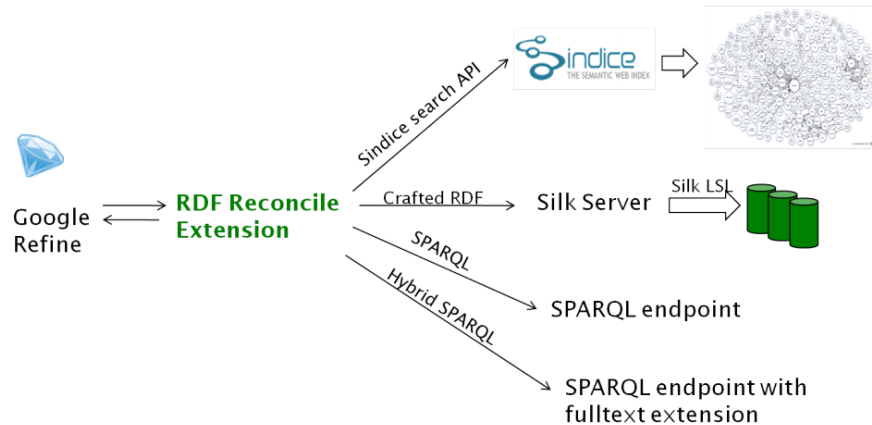
OpenRefine Reconciliation Extension



Several reconciliation services can be accessed from OpenRefine


39.	The Pianist	2002-01-01T00:00:00Z
	<input checked="" type="checkbox"/> <input checked="" type="checkbox"/> The Pianist (100)	
	<input checked="" type="checkbox"/> <input checked="" type="checkbox"/> The Pianist (100)	
	<input checked="" type="checkbox"/> <input checked="" type="checkbox"/> The Pianist (100)	
	<input checked="" type="checkbox"/> <input checked="" type="checkbox"/> The Passionate Pianist (67)	
	<input checked="" type="checkbox"/> <input checked="" type="checkbox"/> Create new item	
	Search for match	
40.	Gladiator	2000-01-01T00:00:00Z
	<input checked="" type="checkbox"/> <input checked="" type="checkbox"/> Gladiator (100)	
	<input checked="" type="checkbox"/> <input checked="" type="checkbox"/> The Gladiator (100)	
	<input checked="" type="checkbox"/> <input checked="" type="checkbox"/> Gladiator (100)	
	<input checked="" type="checkbox"/> <input checked="" type="checkbox"/> The Gladiator (100)	
	<input checked="" type="checkbox"/> <input checked="" type="checkbox"/> Gladiators Seven (90)	
	<input checked="" type="checkbox"/> <input checked="" type="checkbox"/> Gladiators of Rome (67)	
	<input checked="" type="checkbox"/> <input checked="" type="checkbox"/> Create new item	
	Search for match	
41.	American History X	1998-01-01T00:00:00Z
	Choose new match	
42.	The Usual Suspects	1995-01-01T00:00:00Z
	Choose new match	
43.	Léon: The Professional	1994-01-01T00:00:00Z
	Choose new match	
44.	The Lion King	1994-01-01T00:00:00Z
	<input checked="" type="checkbox"/> <input checked="" type="checkbox"/> The Lion King (100)	
	<input checked="" type="checkbox"/> <input checked="" type="checkbox"/> The Lion King (100)	
	<input checked="" type="checkbox"/> <input checked="" type="checkbox"/> The Lion King II: Simba's Pride (93)	
	<input checked="" type="checkbox"/> <input checked="" type="checkbox"/> The Lion King 1½ (93)	
	<input checked="" type="checkbox"/> <input checked="" type="checkbox"/> Leo the Lion: King of the Jungle (59)	
	<input checked="" type="checkbox"/> <input checked="" type="checkbox"/> Create new item	
	Search for match	

OpenCorporates service reconciles company names against legal entities database:
<http://vimeo.com/17924204>



How to Generate RDF?

	Series_Title	Released_Year	Runtime	Genre 1	Genre 2	Genre 3
1.	The Shawshank Redemption	1994-01-01T00:00:00Z	142	Drama		
2.	The Godfather	1972-01-01T00:00:00Z	175	Crime	Drama	
3.	The Dark Knight	2008-01-01T00:00:00Z	152	Action	Crime	Drama
4.	The Godfather: Part II	1974-01-01T00:00:00Z	202	Crime	Drama	



RDF Refine:

- Extension of OpenRefine to support RDF Generation
- Originally developed by INSIGHT – Galway
- Now further developed by open source communities

RDF Refine facilitate RDF export with the help of RDF skeletons

- Define the vocabulary and graph structure of the RDF serialization
- Export to Turtle or RDF/XML

Step 3 - Define RDF Skeletons

An RDF skeleton defines the structure of the RDF triples that are exported

The screenshot displays the 'RDF Schema alignment' interface. At the top right, there are buttons for 'Open...', 'Export', and 'Help'. Below them is a menu with 'Extensions: RDF - Wikidata' and options: 'Edit RDF skeleton...' (highlighted with a red box), 'Reset RDF skeleton...', and 'Add reconciliation service'. The main area is titled 'RDF skeleton' and 'RDF Preview'. It shows a 'Base URI: http://example-instance.org/ Edit'. Under 'Available prefixes:', there is a list: 'rdf owl mv rdfs foaf + Add Manage' (the entire list is highlighted with a red box). The 'RDF skeleton' section shows a tree structure of nodes and properties. A red box highlights the 'Director URI' node, which has a sub-property 'mv:MovieDirector' with a value 'mv:hasName Cell' and an 'Add property' link. Other nodes include 'Series_Title URI', 'Released_Year Cell', 'Genre 1 Cell', 'Genre 2 Cell', 'Genre 3 Cell', and 'Overview Cell'. The 'Director URI' node also has a sub-property 'Director' with a value 'Cell' and an 'Add property' link.

Step 4 - Export RDF Graph Data

RDF Schema alignment

The RDF schema alignment skeleton below specifies how the RDF data that will get generated from your grid-shaped data. The cells in each record of your data will get placed into nodes within the skeleton. Configure the skeleton by specifying which column to substitute into which node.

Base URI: <http://example-instance.org/> [Edit](#)

RDF skeleton | **RDF Preview**

Available prefixes: `rdf owl mv rdfs foaf` [+ Add](#) [Manage](#)

Series_Title URI [X](#) [>](#) mv:hasName → Series_Title Cell

[X](#) mv:Movie [>](#) mv:hasReleaseDate → Released_Year Cell

[Add type](#) [>](#) mv:hasGenre → Genre 1 Cell

[>](#) mv:hasGenre → Genre 2 Cell

[>](#) mv:hasGenre → Genre 3 Cell


[>](#) rdfs:comment → Overview Cell

[>](#) mv:hasMovieDirector → Director URI Director

mv:MovieDirector [>](#) mv:hasName Cell

[Add type](#) [>](#) Add property

[>](#) owl:sameAs → Wiki_ID URI ...



RDF skeleton | **RDF Preview**

This is a sample `Turtle` representation of (up-to) the *first 10 rows*

```
@prefix rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#> .
@prefix owl: <http://www.w3.org/2002/07/owl#> .
@prefix mv: <http://semantics.id/ns/example/movie#> .
@prefix rdfs: <http://www.w3.org/2000/01/rdf-schema#> .
@prefix foaf: <http://xmlns.com/foaf/0.1/> .

<http://example-instance.org/The%20Shawshank%20Redemption> a mv:Movie;
mv:hasName "The Shawshank Redemption";
mv:hasReleaseDate "1994-01-01T00:00Z";
mv:hasGenre "Drama";
rdfs:comment "Two imprisoned men bond over a number of years, finding solace and eventual redemption through acts of common decency." .
<http://example-instance.org/Frank%20Darabont> a mv:MovieDirector;
mv:hasName "Frank Darabont" .
<http://example-instance.org/The%20Shawshank%20Redemption> mv:hasMovieDirector <http://example-instance.org/Frank%20Darabont> .
<http://example-instance.org/The%20Godfather> a mv:Movie;
mv:hasName "The Godfather";
mv:hasReleaseDate "1972-01-01T00:00Z";
mv:hasGenre "Crime","Drama";
rdfs:comment "An organized crime dynasty's aging patriarch transfers control of his clandestine empire to his reluctant son." .
<http://example-instance.org/Francis%20Ford%20Coppola> a mv:MovieDirector;
mv:hasName "Francis Ford Coppola" .
<http://example-instance.org/The%20Godfather> mv:hasMovieDirector <http://example-instance.org/Francis%20Ford%20Coppola> .
<http://example-instance.org/The%20Dark%20Knight> a mv:Movie;
```

Open... | Export ▾ | Help

OpenRefine project archive to file

Tab-separated value

Comma-separated value

HTML table

Excel (.xls)

Excel 2007+ (.xlsx)

ODF spreadsheet

Custom tabular exporter...

SQL Exporter...

Templating...

RDF as RDF/XML

RDF as Turtle

OpenRefine project archive to Google Drive...

Google Sheets

Wikidata edits...

QuickStatements file

Wikidata schema

Further Pointers

- Open Refine: <https://openrefine.org/>
- RDF Refine extension:
<https://github.com/stkenny/grefine-rdf-extension>
- An example project with RDF skeletons:
https://bit.ly/iss_openrefine_example

Knowledge Graph Creation


Relational Data to RDF: R2RML

Emanuel Sallinger
Katja Hose

Extracting Relational Data: R2RML



Databases



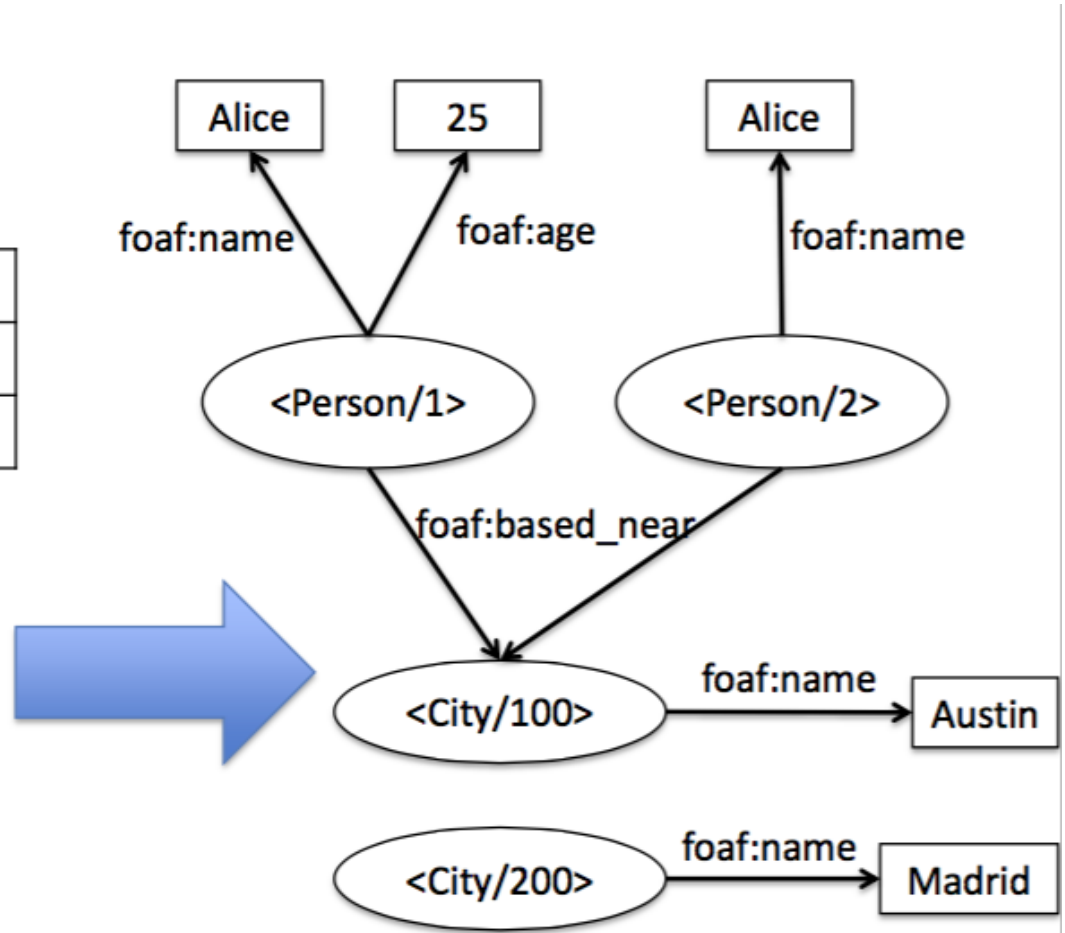
Goal: Transform Databases to RDF

Person

ID	NAME	AGE	CID
1	Alice	25	100
2	Bob	NULL	100


City

CID	NAME
100	Austin
200	Madrid



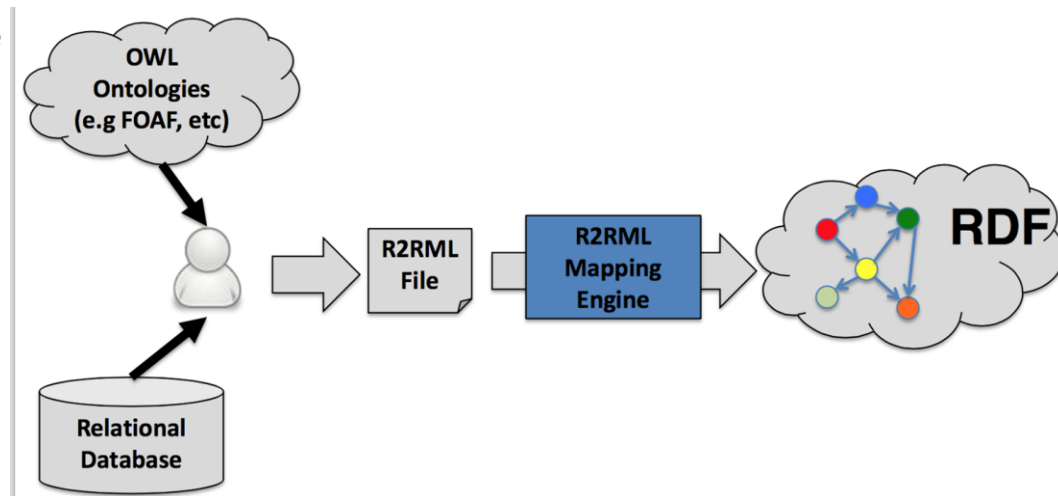
Direct mapping: directly exposes data as RDF

- Does not allow for vocabulary mapping
- No interlinking (unless URIs used in relational data)



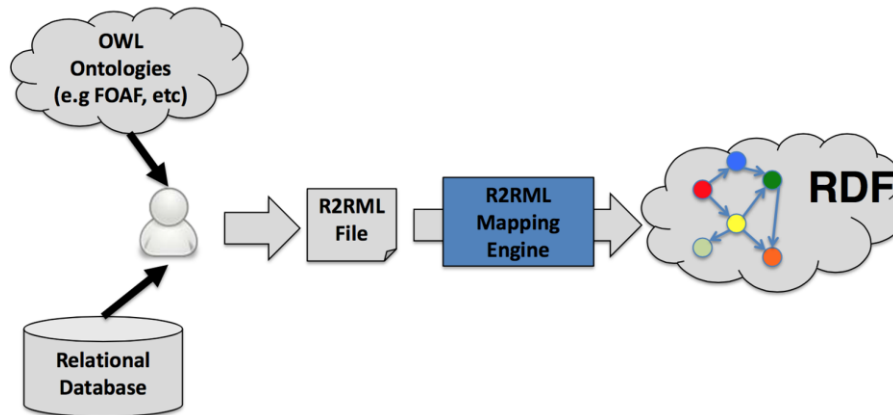
R2RML, the RDB to RDF mapping language

- Allows vocabulary mapping (subject, predicate and object maps with class options)
- Allows interlinking – URIs can be constructed



R2RML Mapping Language – Basic Structure

- Each mapping called Triples Map “`rr:TriplesMap`”
- Consist of
 - The definition of the data source “`rr:LogicalTable`”
 - One rule to generate the subject of the triples “`rr:SubjectMap`”
 - Zero or more rules to generate the predicate and object of the triples “`rr:PredicateObjectMap`”



Direct Mapping as R2RML

```
@prefix rr: <http://www.w3.org/ns/r2rml#> .
```

```
<TriplesMap1>
```

```
  a rr:TriplesMap;
```

```
  rr:logicalTable [ rr:tableName "Person"];
```

```
  rr:subjectMap [
```

```
    rr:template "http://www.ex.com/Person/ID={ID}";
```

```
    rr:class <http://www.ex.com/Person>
```

```
  ];
```

```
  rr:predicateObjectMap [
```

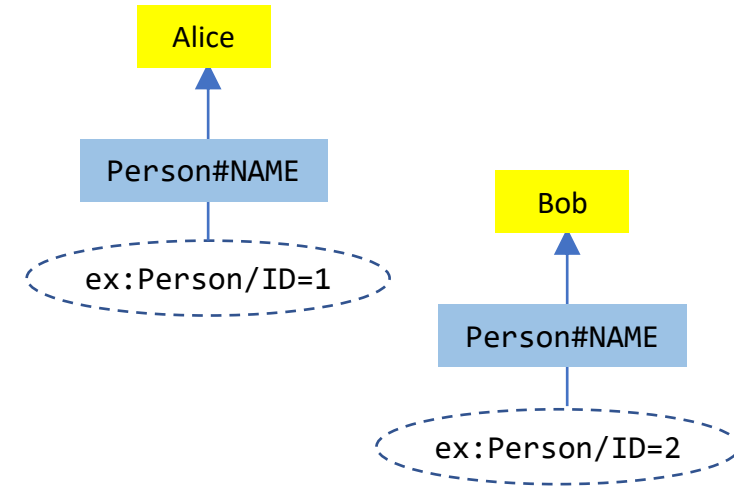
```
    rr:predicate <http://www.ex.com/Person#NAME> ;
```

```
    rr:objectMap [rr:column "NAME" ]
```

```
  ].
```

Person

ID	NAME	AGE	CID
1	Alice	25	100
2	Bob	NULL	100



Direct Mapping as R2RML

```
@prefix rr: <http://www.w3.org/ns/r2rml#> .
```

```
<TriplesMap1>
```

```
  a rr:TriplesMap;
```

What is being mapped?

```
  rr:logicalTable [ rr:tableName "Person" ];
```

How to generate the subject URI?

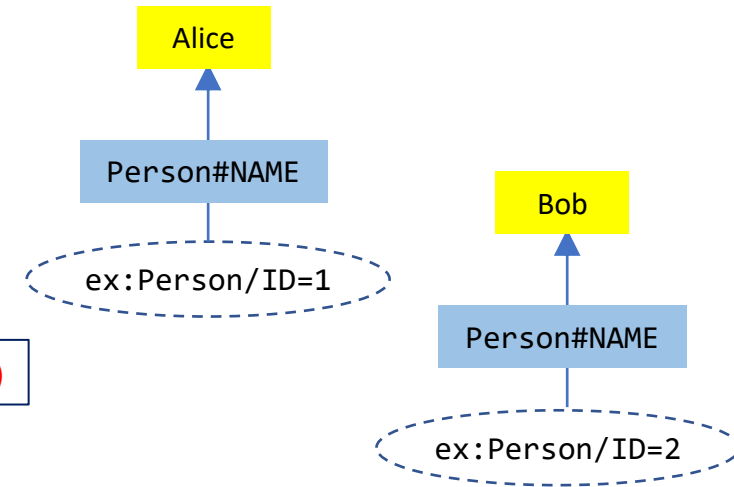
```
  rr:subjectMap [
    rr:template "http://www.ex.com/Person/ID={ID}";
    rr:class <http://www.ex.com/Person>
  ];
```

How to generate the predicate and object? (optional)

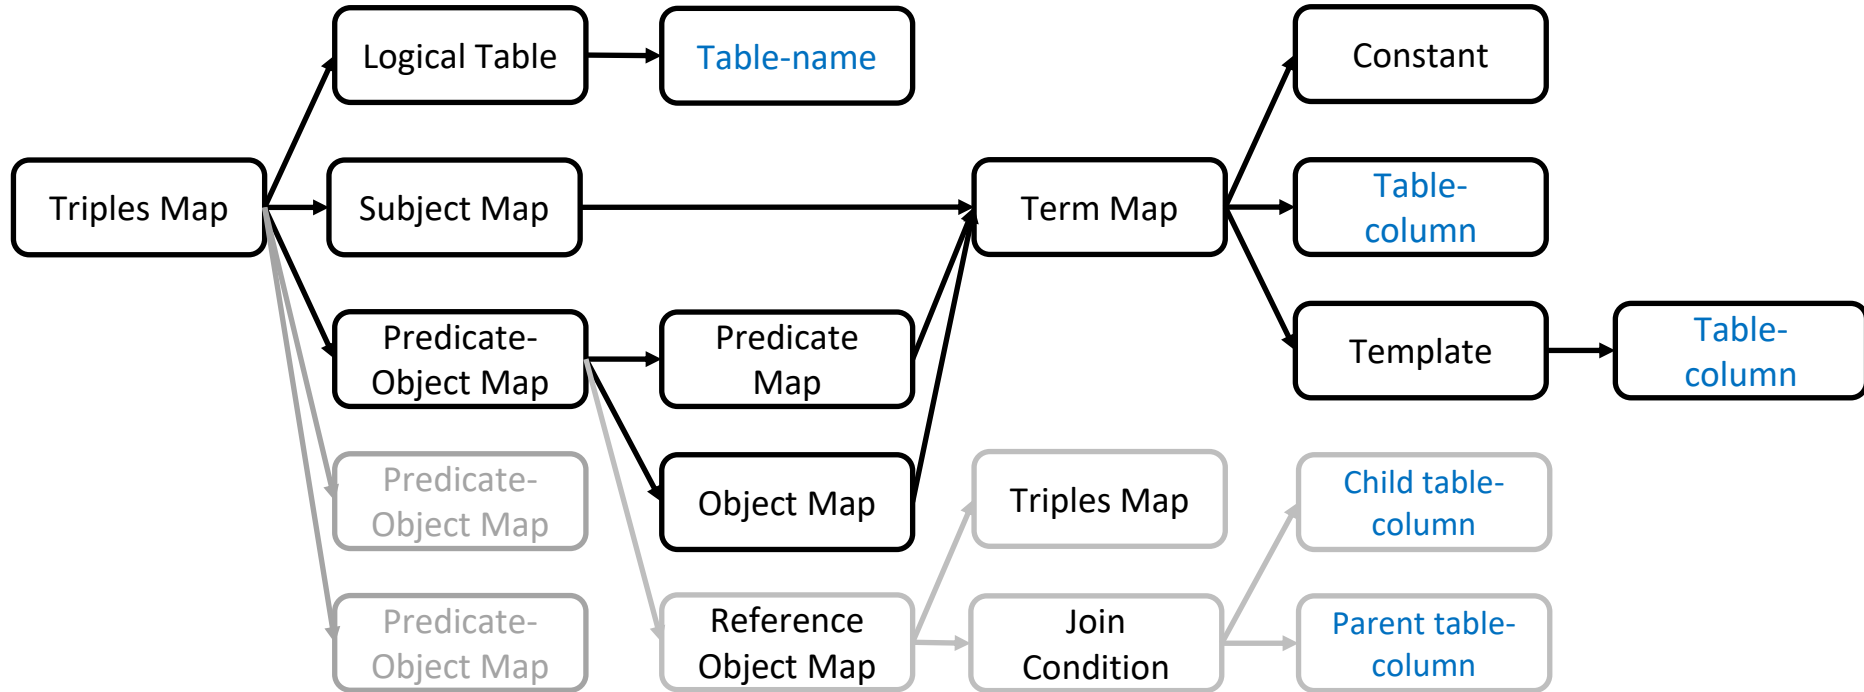
```
  rr:predicateObjectMap [
    rr:predicate <http://www.ex.com/Person#NAME> ;
    rr:objectMap [ rr:column "NAME" ]
  ]
```

Person

ID	NAME	AGE	CID
1	Alice	25	100
2	Bob	NULL	100



R2RML Mapping Definition



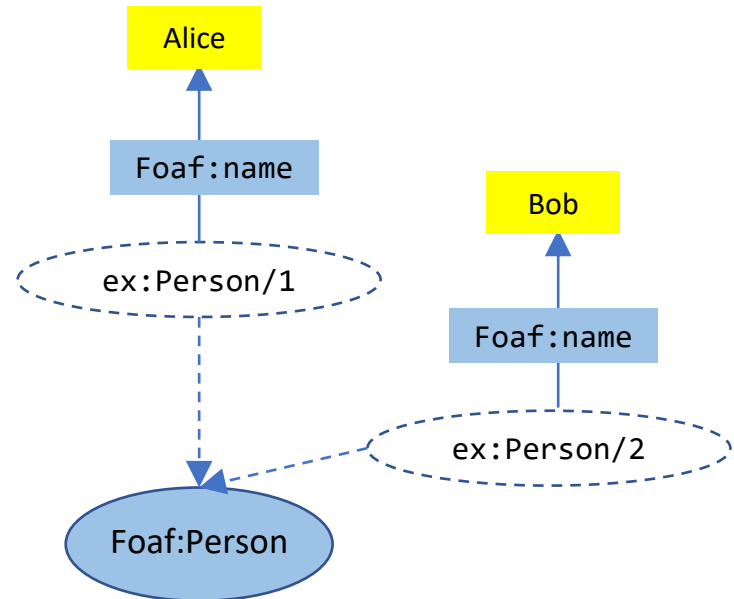
Customizing R2RML Mapping

```
@prefix rr: <http://www.w3.org/ns/r2rml#> .  
@prefix foaf: <http://xmlns.com/foaf/0.1/> .
```

```
<TriplesMap1> a rr:TriplesMap;  
  rr:logicalTable [ rr:tableName "Person" ];  
  
  rr:subjectMap [  
    rr:template "http://www.ex.com/Person/{ID}";  
    rr:class    foaf:Person  
  ];  
  
  rr:predicateObjectMap [  
    rr:predicate foaf:name;  
    rr:objectMap [ rr:column "NAME" ]  
  ];
```

Person

ID	NAME	AGE	CID
1	Alice	25	100
2	Bob	NULL	100



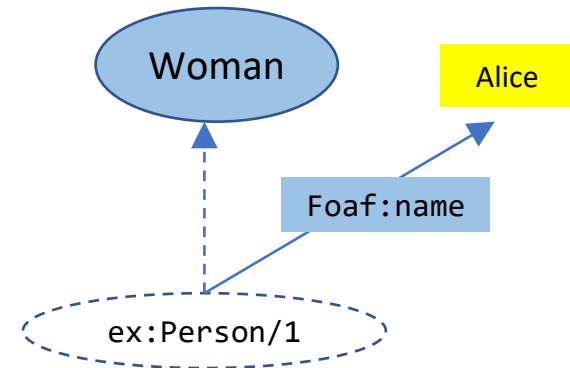
Customizing R2RML Source

```
@prefix rr: <http://www.w3.org/ns/r2rml#> .  
@prefix foaf: <http://xmlns.com/foaf/0.1/> .
```

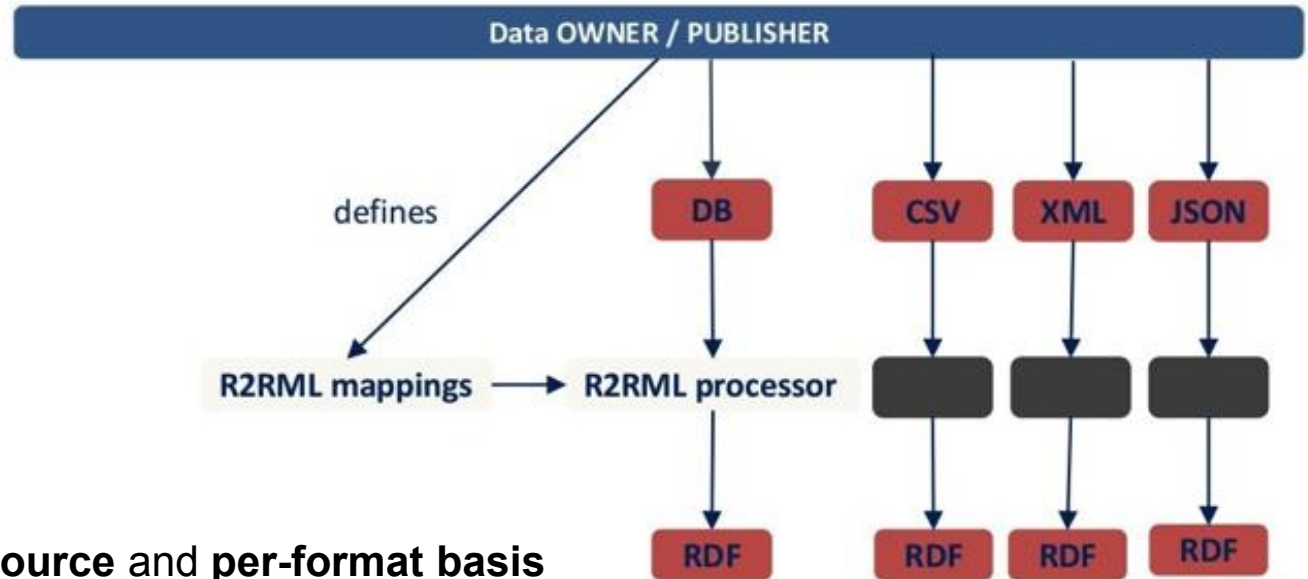
```
<TriplesMap1>  
  a rr:TriplesMap;  
  rr:logicalTable [ rr:sqlQuery  
    """ SELECT ID, NAME  
    FROM Person WHERE gender = "F" """ ];  
  rr:subjectMap [  
    rr:template "http://www.ex.com/Person/{ID}";  
    rr:class <http://www.ex.com/Woman>  
  ];  
  rr:predicateObjectMap [  
    rr:predicate foaf:name;  
    rr:objectMap [rr:column "NAME" ]  
  ] .
```

Person

ID	NAME	GENDER
1	Alice	F
2	Bob	M



R2RML is good, but **how about other formats?**



- Mapping data **per-source** and **per-format** basis
- Mappings are **tied to implementations** and not **interoperable** across different implementations
- No **uniform way to describe mappings** of heterogeneous resources
 - **Mapping definitions can't be reused** for data in the same or different formats

Knowledge Graph Creation

Structured Data to RDF: RML

Emanuel Sallinger
Katja Hose

Extracting Structured Data: RML



Structured Data



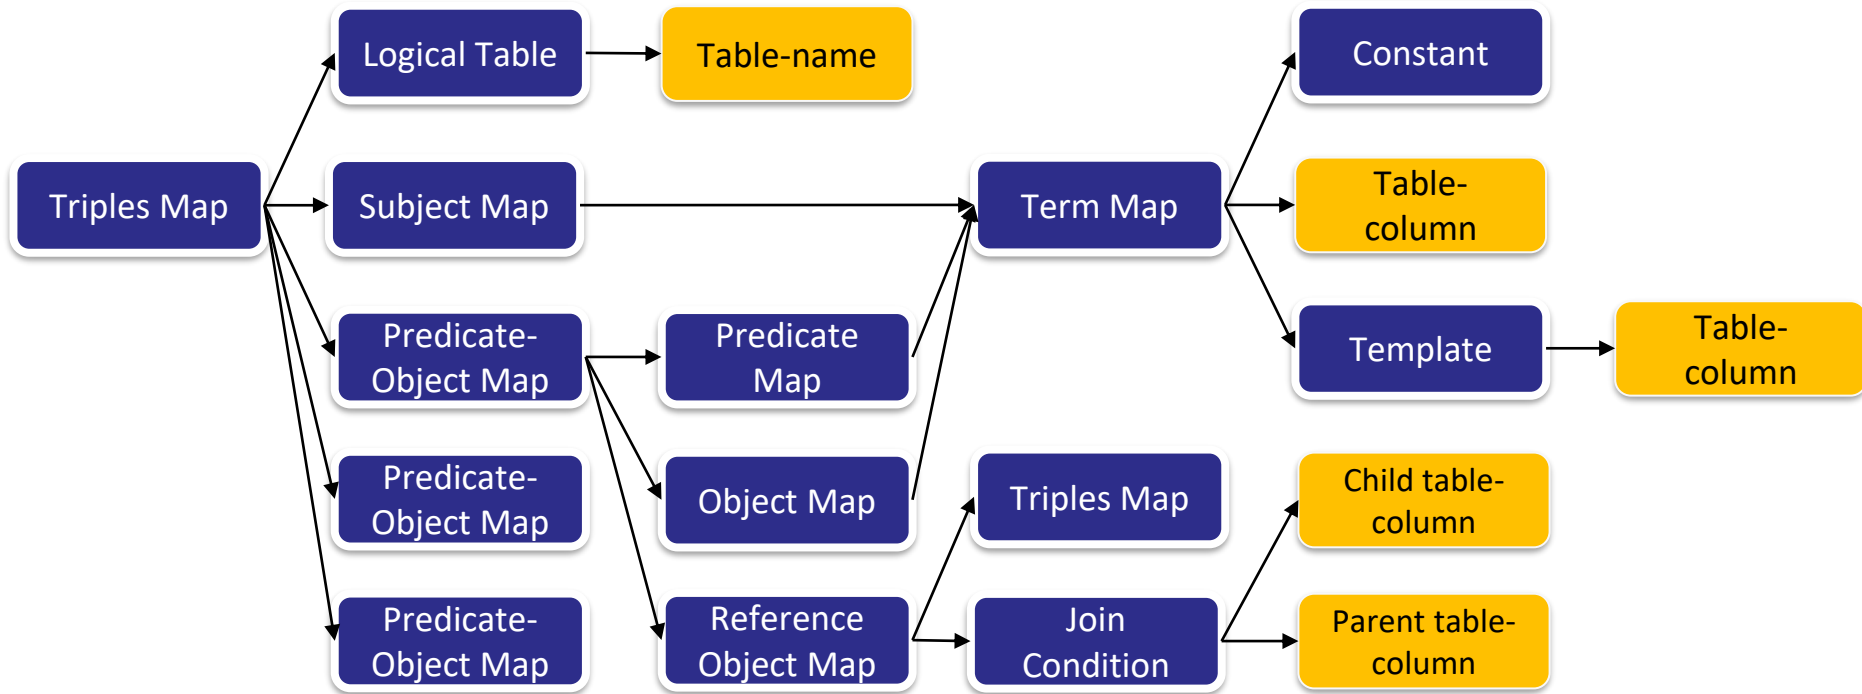
RDF Mapping Language (RML)

- A generic, scalable mapping language defined to express rules that map data in heterogeneous structures and serializations to the RDF data model.
- A superset of R2RML.



Anastasia Dimou, Miel Vander Sande, Pieter Colpaert, Ruben Verborgh, Erik Mannens, and Rik Van de Walle. "**RML: a generic language for integrated RDF mappings of heterogeneous data.**" In *LDOW*. 2014.

From R2RML to RML



From R2RML to RML

R2RML	Generic mapping language
Logical Table	Logical Source (CSV, XML, JSON)
Table Name	Source name / URI
Column	???
per row iteration	???

References to values of heterogeneous resources

```
<PendingOrders>...  
  <Order id="398">  
    <Product>  
      <Id>AE5982</Id>  
      <Name>Samsonite DeLux 45</Name>  
    </Product>  
  </Order>...  
</PendingOrders>
```

```
{ ...,  
  "ProductInStock" :  
    { "ID": "567",  
      "Name": "Samsonite DeLux 45",  
      "type": "suitcase",  
    }, ...  
}
```

XPath for XML

Reference:

"Order@Id"

Iterator:

"/PendingOrders /Order"

JSONPath for JSON

Reference:

"\$. ProductInStock.ID"

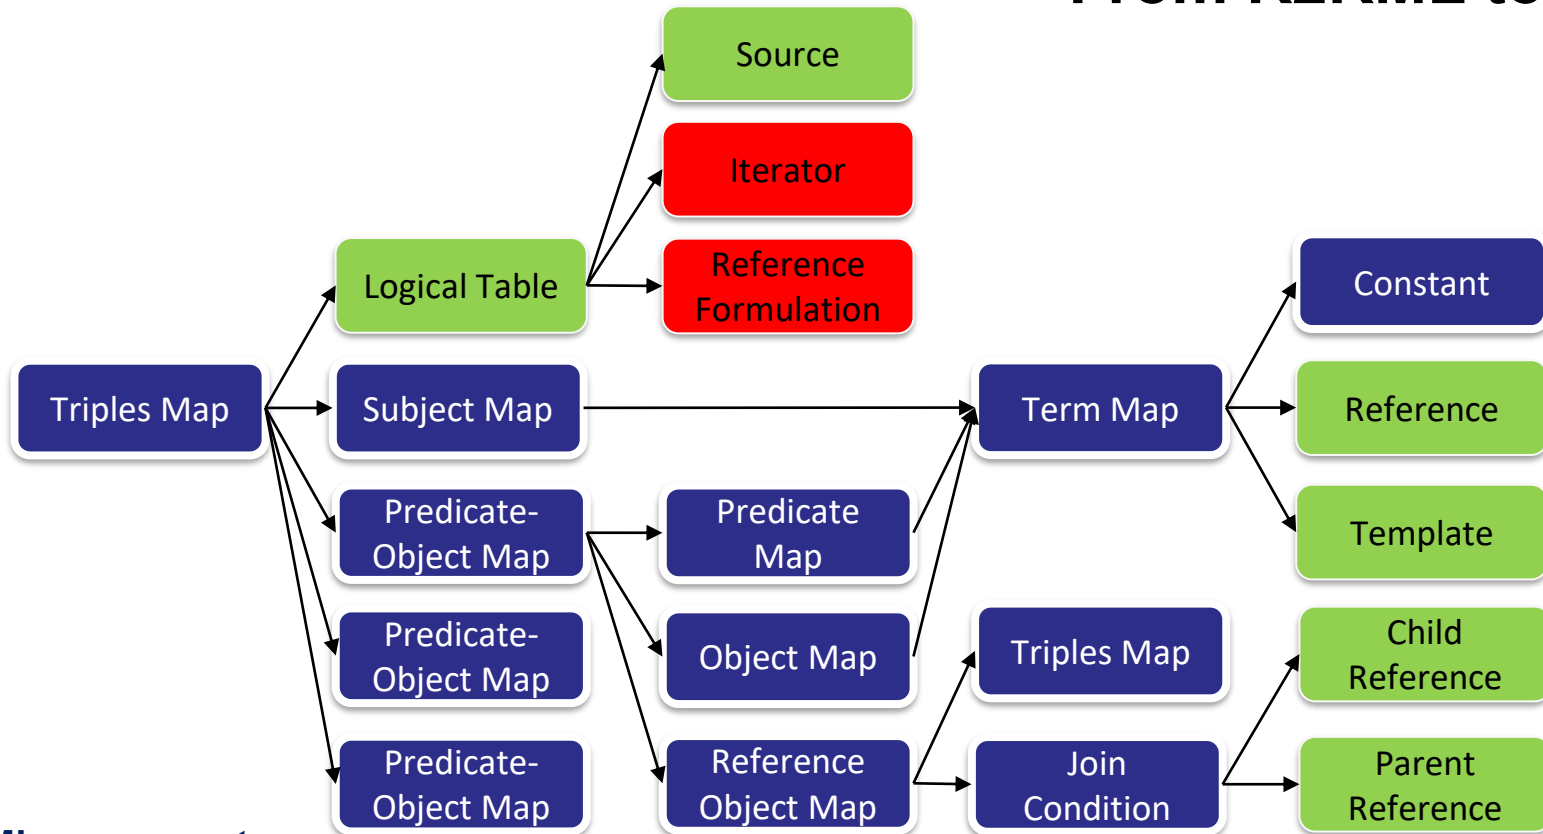
Iterator:

"\$.ProductInStock"

From R2RML to RML

R2RML	R2 RML
Logical Table	Logical Source (CSV, XML, JSON)
Table Name	Source name / URI
Column	Reference (defined Reference Formulation)
per row iteration	defined Iterator

From R2RML to RML



R2RML components

R2RML extended components

RML specific components

RML Basics

Input in JSON format (People.json)

```
[ ... { "Title": "Multimedia Lab",
        "People": [ { "name": "Anastasia", "surname": "Dimou" },
                    { "name": "Ruben", "surname": "Verborgh" } ],
      { "Title": "Media and ICT",
        "People": [ ... ] } , ... ]
```



Output RDF representation

```
ex:Anastasia_Dimou a ex:Person ;
ex:name "Anastasia".

ex:Ruben_Verborgh a ex:Person ;
ex:name "Ruben".
```

RML Logical Source

```
<#PersonMapping>
  rml:logicalSource [
    rml:source "People.json" ;
    rml:referenceFormulation ql:JSONPath ;
    rml:iterator "$.[*].People" ;
  ]
```

Logical Source Mapping contains:

- (i) a reference to input **source(s)**
- (ii) the **referenceFormulation** to specify how to refer to the data (e.g., ql:CSV, ql:CSS3, ql:JSONPath, rr:SQL2008 and ql:Xpath) and
- (iii) the **iterator** that specifies how to iterate over the data.

RML Basics

Input in JSON format (People.json)

```
[ ... { "Title": "Multimedia Lab",  
      "People": [ { "name": "Anastasia", "surname": "Dimou" },  
                  { "name": "Ruben", "surname": "Verborgh" } ],  
  { "Title": "Media and ICT",  
    "People": [ ... ] } , ... ]
```



Output RDF representation

```
ex:Anastasia_Dimou a ex:Person ;  
ex:name "Anastasia".  
  
ex:Ruben_Verborgh a ex:Person ;  
ex:name "Ruben".
```

RML Subject Map

```
<#PersonMapping>  
  rr:subjectMap [  
    rr:template "http://ex.com/Person/{name}_{surname}" ;  
    rr:class ex:Person ;  
  ]
```

.

The Subject Map consists of:

- (i) the URI pattern that defines how each triple's subject is generated and
- (ii) (optionally) its type.

The references to the input data occurs using valid references according to reference formulation specified at the Logical Source

RML Basics

Input in JSON format (People.json)

```
[ ... { "Title": "Multimedia Lab",  
      "People": [ { "name": "Anastasia", "surname": "Dimou" },  
                  { "name": "Ruben", "surname": "Verborgh" } ],  
  { "Title": "Media and ICT",  
    "People": [ ... ] } , ... ]
```



Output RDF representation

```
ex:Anastasia_Dimou a ex:Person ;  
  ex:name "Anastasia".  
  
ex:Ruben_Verborgh a ex:Person ;  
  ex:name "Ruben".
```

Predicate Object Map

```
<#PersonMapping>  
  rr:predicateObjectMap [  
    rr:predicate ex:name ;  
    rr:objectMap [ rml:reference "name" ] ;  
  ]  
.
```

Triples are generated using **Predicate Object Maps** consisting of:

- (i) A **Predicate Map** specifying how the triple's predicate is generated.
- (ii) An **Object Map** specifying how the triple's object(s) are generated.

Looks promising.. But can I operate on values?

- Custom functions can be added into RML engine
- However, the procedure for providing custom functions is implementation specific,
 - standard function specs are currently NOT defined
 - e.g., RML-Mapper: <https://rml.io/docs/rmlmapper/default-functions/>

Default functions in the RMLMapper

This is a list of all functions and their arguments that are available in the [RMLMapper](#) by default. [Additional](#) functions can be added too.

Jump to

[DBpedia Spotlight annotate](#)

[Read local file as string](#)

[True condition](#)

[and](#)

[chomp](#)

[contains](#)

[decide](#)

[endsWith](#)

[equal](#)

[escape](#)

Wait.. are there multiple RML implementations?

- Yes! → <https://rml.io/implementation-report/>

Name	Version	Test date	Contact	Web page
RMLMapper	4.9.0	2020-09-17	Anastasia Dimou	https://github.com/rmlio/rmlmapper-java
CARML	0.3.0	2020-09-15	Pano Maria	https://github.com/carm1/carm1
RocketRML	1.0.6	2019-06-28	Umutcan Simsek	https://github.com/semantifyit/RocketRML
SDM-RDFizer	3.2	2020-06-08	Maria-Esther Vidal	https://github.com/SDM-TIB/SDM-RDFizer
RMLStreamer	2.0.0	2020-06-10	Anastasia Dimou	https://github.com/RMLio/RMLStreamer
Chimera	2.1	2021-02-18	Mario Scrocca	https://github.com/cefriel/chimera

Still.. Writing RML mapping in turtle is cumbersome for some people...

Introducing YARRRML – RML rules in YAML

- YARRRML is a human readable text-based representation for RDF Graph generation rules (i.e., RML/R2RML).
 - A subset of [\[YAML\]](#).
 - Spec: <https://rml.io/yarrml/spec/>
- A convenient way to write (R2)RML mappings
- YARRRML will be translated into (R2)RML mappings before executed by the appropriate engine.

Translation from RML to YARRRML

RML

... (prefixes ommitted)

```
map:map_person_000 a rr:TriplesMap ;
  rml:logicalSource [
    rml:source "data.json" ;
    rml:referenceFormulation ql:JSONPath ;
    rml:iterator "$.persons[*]" ;
  ] ;
rr:subjectMap [
  rr:template "http://example.com/{firstname}" ;
  rr:class foaf:Person ;
];
rr:predicateObjectMap [
  rr:predicate foaf:name ;
  rr:objectMap [
    rml:reference "firstname" ;
    rr:termType rr:Literal ;
  ] ;
];
```

YARRRML

```
prefixes:
  ex: "http://example.com/"

mappings:
  person:
    sources:
      - ['data.json~jsonpath', '$.persons[*]']
    s: http://example.com/${firstname}
    po:
      - [a, foaf:Person]
      - [foaf:name, ${firstname}]
```

Further Pointers

- RML demo using caRML (<https://github.com/carml/carml/>)
 - Example implementation available at: <https://github.com/fekaputra/carmlizer>
 - CLI tool: <https://github.com/netage/carml-cli>
- RML using YARRRML + Matey (on-the-fly)
 - Matey: <https://rml.io/yarrml/matey/>

... and some final notes on RML

- Handling of some **details on RML mappings might be vary** between different engines
- Before implementing (complex) custom functions, **consider pre-processing the data** → might be more efficient due to the limitation of the RML / Query Engines (JSONPath / XPath)
- Before (or during) the RML mapping creations for JSON/XML, **consider using online JSONPath or XPath editors** for checking the correctness of iterator/queries.
- **YARRRML** simplified the creation of RML mappings; however, be wary that you need time to learn it (on top of learning RML basics)
- For you that hate to write in turtle/yaml, **RML editor** (<https://rml.io/tools/rmleditor/>) might be something for you
- Recently, OntoText (the company behind GraphDB) have forked OpenRefine + RDFRefine into one of its (free) product called **Ontotext-Refine**, which is a great alternative to the original OpenRefine (<https://www.ontotext.com/products/ontotext-refine/>)
- For **tabular data**, **consider using OpenRefine + RDF refine** before jumping into RML → OpenRefine is much more agile and powerful for tabular data (pre-)processing.

Knowledge Graph Creation


Text to RDF

Emanuel Sallinger
Katja Hose

Mapping and Transformation Tools



Text



Document Annotation

Document Annotation with explicit semantics, e.g. semantic entities

On July 26, 1928, American film director, screenwriter, producer, cinematographer, and editor **Stanley Kubrick** was born. He is regarded as one of the greatest filmmakers of all time. His

Stanley Kubrick

Stanley Kubrick (/ˈkuːbrɪk/; July 26, 1928 – March 7, 1999) was an American film director, screenwriter, producer, cinematographer, and

birth year 1928

death year 1999

is producer of 2001: A Space Odyssey (film)

ethnicity American Jews

is spouse of Christiane Kubrick

occupation Cinematographer

or short stories, are noted for their unique cinematography, and an inspired use of music scores. Clockwork Orange, The Shining, 'I'm Spartacus' or Dr. Strangelove – they all are movies ever made.

think the thing was probably helped me the most was that to do in the early 50s for someone who actually go and as impossible. It really is terribly easy. All anybody ever, and some imagination.” (Stanley Kubrick in an

<http://scihi.org/breaking-new-grounds-in-cinematography-stanley-kubrick/>

district of New York, as the first of two children into a family with Jewish ancestry. As a child, Stanley was considered intelligent, but he did not achieve particularly high

**Example for
Linked Data
Based Document
Annotation**

Document Annotation allows Advanced Exploration

Ruth Sobotka	Christiane Kubrick		
St Albans	Braunschweig	Manhattan	
Vienna	New York	New York City	
London	United Kingdom	Germany	
2001: A Space Odyssey (film)	Spartacus (film)	Full Metal Jacket	
Dr. Strangelove	American Jews	Cinematographer	
Scottish Gaelic	Hertfordshire	Irish language	
Pound sterling	Painting	Judaism	
English language			
1707	1801	1928	1999

Stanley Kubrick

15 Recommended Articles:

- #1 Breaking New Grounds in Cinematography – Stanley Kubrick
- #2 Benoît Mandelbrot and the Beauty of Mathematics
- #3 Benoît Mandelbrot and the Beauty of Mathematics
- #4 Now I am become Death, the Destroyer of Worlds
- #5 Charles Proteus Steinmetz and the Alternating Current
- #6 George Kistiakowsky – From the Manhattan Project to a Nuclear Weapons

Stanley Kubrick




Stanley Kubrick (/ˈkuːbrɪk/; July 26, 1928 – March 7, 1999) was an American film director, screenwriter, producer, cinematographer, and editor who did much of his work in the United Kingdom. He is regarded as one of the great filmmakers. He was part of the New Hollywood wave of filmmaking. His films, typically adaptations of novels or short stories, are noted for their "dazzling" and unique cinematography, attention to detail in the service of realism, and the evocative use of music. Kubrick's films covered a variety of genres, including war, crime, literary adaptations, romantic and black comedies, horror, epic and science fiction. Kubrick

DBpedia: Stanley Kubrick

Semantic Annotation


- The term “**annotation**” implies to **attach data to some other piece of data**. It establishes a **(typed) relation between the annotated data and the annotating data**
- The term “**annotation**” can denote both the **process of annotating** and the **result of that process**
- **Formal Model:**
 - An **annotation** A is a tuple $(\mathbf{a}_s, \mathbf{a}_p, \mathbf{a}_o, \mathbf{a}_c)$, where
 - \mathbf{a}_s is the **subject** of the annotation (the annotated data)
 - \mathbf{a}_o is the **object** of the annotation (the annotating data)
 - \mathbf{a}_p is the **predicate** (the annotation relation) that defines the type of relationship between \mathbf{a}_s and \mathbf{a}_o , and
 - \mathbf{a}_c is the **context** in which the annotation is made.



- **Formal Model:**

- An **annotation** A is a tuple (a_s, a_p, a_o, a_c) , where
 - a_s is the **subject** of the annotation (the annotated data)
 - a_o is the **object** of the annotation (the annotating data)
 - a_p is the **predicate** (the annotation relation) that defines the type of relationship between a_s and a_o , and
 - a_c is the **context** in which the annotation is made.


```
PREFIX dbr: <http://dbpedia.org/resource/> .  
PREFIX dbo: <http://dbpedia.org/ontology/> .  
  
:9787532717071 dbo:author dbr:Ernest_Hemingway .
```



- **Formal Model:**

- An **annotation** A is a tuple (a_s, a_p, a_o, a_c) , where
 - a_s is the **subject** of the annotation (the annotated data)
 - a_o is the **object** of the annotation (the annotating data)
 - a_p is the **predicate** (the annotation relation) that defines the type of relationship between a_s and a_o , and
 - a_c is the **context** in which the annotation is made.

```
PREFIX dbr: <http://dbpedia.org/resource/> .  
PREFIX dbo: <http://dbpedia.org/ontology/> .  
  
:9787532717071 dbo:author dbr:Ernest_Hemingway .
```




- **Formal Model:**

- An **annotation** A is a tuple (a_s, a_p, a_o, a_c) , where
 - a_s is the **subject** of the annotation (the annotated data)
 - a_o is the **object** of the annotation (the annotating data)
 - a_p is the **predicate** (the annotation relation) that defines the type of relationship between a_s and a_o , and
 - a_c is the **context** in which the annotation is made.

```
PREFIX dbr: <http://dbpedia.org/resource/> .  
PREFIX dbo: <http://dbpedia.org/ontology/> .  
  
:9787532717071 dbo:author dbr:Ernest_Hemingway .
```


Semantic Annotation



- **Formal Model:**


- An **annotation** A is a tuple (a_s, a_p, a_o, a_c) , where
 - a_s is the **subject** of the annotation (the annotated data)
 - a_o is the **object** of the annotation (the annotating data)
 - a_p is the **predicate** (the annotation relation) that defines the type of relationship between a_s and a_o , and
 - a_c is the **context** in which the annotation is made.

```
PREFIX dbr: <http://dbpedia.org/resource/> .  
PREFIX dbo: <http://dbpedia.org/ontology/> .  
  
:9787532717071 dbo:author dbr:Ernest_Hemingway .
```

- created at Dec 11, 2017
- created by Lysander07
- ...

Ontologies for Semantic Annotation

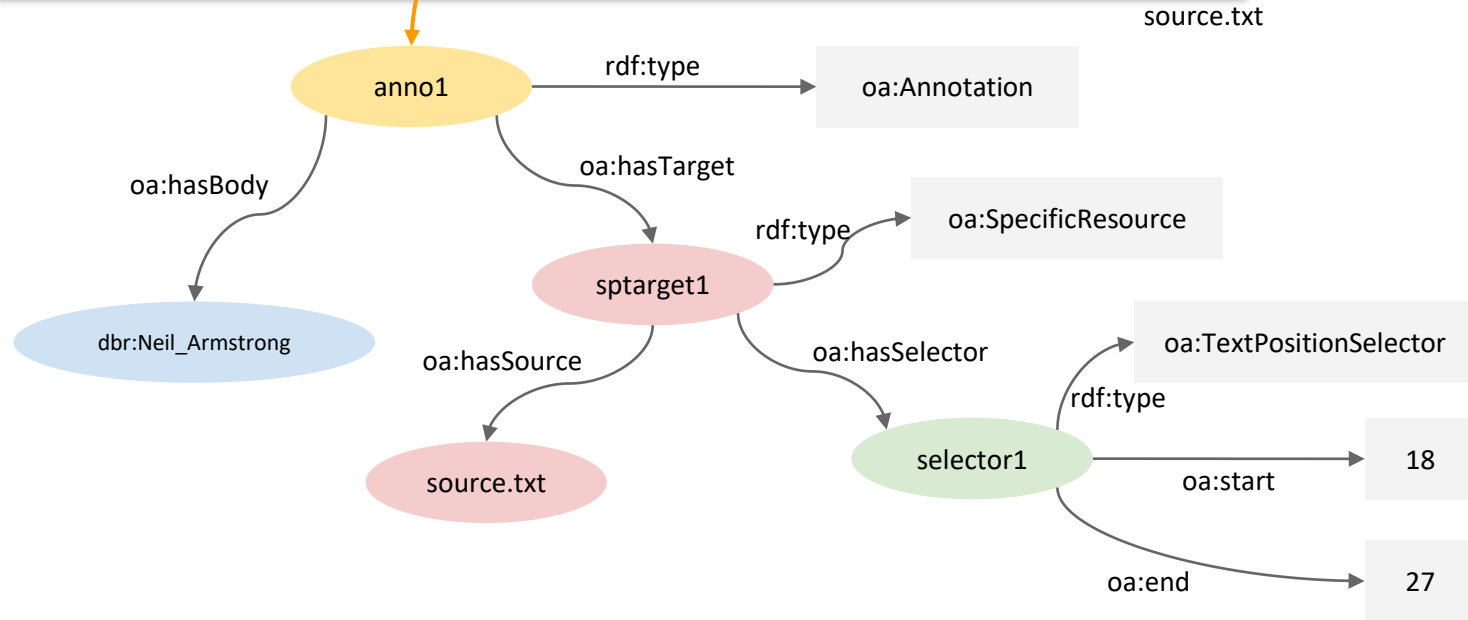
- The **Open Annotation Ontology** is a W3C Standard (2017) [1]
 - “...specifies an interoperable framework for creating associations between related resources, annotations, using a methodology that conforms to the Architecture of the World Wide Web.”



- The **Body** provides the information which is annotating the **Target**

Open Annotation Ontology Example

On July 16, 1969, **Armstrong**, along with Edwin E. Aldrin, Jr., and Michael Collins, blasted off in the Apollo 11 vehicle toward the Moon.



Confidence: 0.75

Language:

n-best candidates

On July 26, 1928, American film director, screenwriter, producer, [cinematographer](#), and editor [Stanley Kubrick](#) was born. He is regarded as one of the greatest filmmakers of all time. His films, typically adaptations of novels or short stories, are noted for their unique cinematography, attention to details to achieve realism and an inspired use of music scores. Clockwork [Orange](#), 2001, [Full Metal Jacket](#), [Lolita](#), [The Shining](#), 'I'm [Spartacus](#)' or Dr. Strangelove – they all are considered as being some of the best movies ever made.

[Stanley Kubrick](#) was born in the [Bronx](#) district of [New York](#), as the first of two children into a family with Jewish ancestry. As a child, Stanley was considered intelligent, but he did not achieve particularly high grades at school. Concerned to find a remedy for his son's poor academic performance, Jacques (Jacob) Leonard [Kubrick](#) sent him to stay with his uncle in [Pasadena, California](#) in 1940, hoping that a change of scenery would also improve Stanley's grades. Returning back home in 1941 for his last year of [grammar school](#), there seemed to be little change in his attitude or his results. http://dbpedia.org/resource/Stanley_Kubrick But then, Stanley was introduced by his father to the game of [chess](#), with the desired result. [Kubrick](#) took to the game passionately, and quickly became a skilled player. Actually, he later also played [chess](#) for money in [Greenwich Village](#) to supplement his income, before making his first movie.

- Linking unstructured information sources to the LOD cloud through DBpedia
- Web Service, Java/Scala API, jQuery plugin, Java and PHP clients
- Currently supporting more than 15 languages!

Source: <https://demo.dbpedia-spotlight.org/>

PermIDs (formerly: OpenCalais)

Intelligent Tagging Demo

The screenshot displays the Permid.org viewer interface. On the left, a sidebar titled "FOUND IN DOCUMENT" shows a list of entities and relations. The "ENTITIES" section includes "City", "Person", "Position", and "Province Or State". The "RELATION" section includes "Person Attributes", "Person Career", and "Person Location". The main area, titled "DOCUMENT VIEW", shows a document snippet with highlighted text. Three pop-up windows are overlaid on the document, providing detailed information for a specific mention of Stanley Kubrick. Each pop-up shows a relevance score of 80% and a list of attributes such as "persontype", "nationality", "confidencelevel", and "firstname".

FOUND IN DOCUMENT

ENTITIES

- City
- Person
- Position
- Province Or State

RELATION

- Person Attributes
- Person Career
- Person Location

DOCUMENT VIEW

On July 26, 1928, American film director, screenwriter, producer, cinematographer, and editor Stanley Kubrick was born. He is regarded as one of the greatest filmmakers of all time. His films, many of which are adaptations of novels or short stories, are known for their attention to details to achieve realism and an inspiration to filmmakers. Some of his most famous works include Full Metal Jacket, Lolita, The Shining, 'I'm Still a处女, and The Godfather. Kubrick was born in the Bronx district of New York City. He was of Jewish ancestry. As a child, Stanley was considered a prodigy and earned high grades at school. Concerned to improve his performance, Jacques (Jacob) Leonard Kubrick sent Stanley to a boarding school in California in 1940, hoping that a change of scenery would also improve Stanley's grades. Returning home in 1941 for his last year of grammar school, there seemed to be little change in his

PERSON Stanley Kubrick (Mention 1 of 18)

Relevance 80%

persontype	N/A
nationality	N/A
confidencelevel	0.999
firstname	Stanley

PERSON Stanley Kubrick (Mention 3 of 18)

Relevance 80%

persontype	N/A
nationality	N/A
confidencelevel	0.999
firstname	Stanley

RELATION PERSON CAREER

person	Stanley Kubrick
position	editor
careertype	professional
status	current
forenduserdisplay	true

[View RDF](#)

- NLP and machine learning to extract entities, facts, and events
- Commercial product (<https://www.refinitiv.com/en>)

Source: <https://permid.org/onecalaisViewer>

- Common problem with general purpose, open-domain semantic annotation tools
- Best results require bespoke customisation



General Architecture for Text Engineering

- Open-source framework and development environment
- Started 1996, large developer community
- Used worldwide by many organisations to build bespoke solutions; e.g. Press Association and the National Archive
- Information Extraction in many languages

GATE Cloud NLP Services



Home

Services

Dashboard

Show only items tagged: **Basque (1)** **Biomed (1)** **Bulgarian (1)** **Catalan (1)** **Chunker (1)** **Croatian (1)** **Custom (2)** **Czech (1)** **Danish (1)** **Dutch (2)** **English (19)**
Environment (2) **Estonian (1)** **Finnish (1)** **French (4)** **German (5)** **Greek (1)** **Indonesian (1)** **Latvian (1)** **Measurements (2)** **Morphology (1)** **Named Entity (23)**
OpenNLP (3) **Opinion Mining (4)** **Part-of-Speech (23)** **Polish (1)** **Politics (2)** **Portugese (1)** **Romanian (2)** **Russian (3)** **Server (2)** **Slovak (1)** **Slovenian (1)**
SoBigData (3) **Spanish (3)** **Summarization (2)** **Swedish (1)** **Term Recognition (2)** **Twitter (14)** **Veracity (1)** **Welsh (1)**

English Named Entity Recognizer



Identify names of *persons, locations, organizations*, as well as *money amounts, time and date expressions* in English texts automatically.

1,200 free requests / day
Larger batches GBP0.80 / CPU hour

English Named Entity Recognizer for Tweets



Analyse tweets for names of *persons, locations, organizations* and other entities. Also performs normalization of abbreviations and common shorthands ("brb", "gr8", "2day", etc.).

1,200 free requests / day
Larger batches GBP0.80 / CPU hour

Twitter Collector



Collect tweets, view tweet statistics, and store results in your dashboard for further analysis.

GBP0.05 / CPU hour

BioYODIE Named Entity Disambiguation



Identifies biomedical named entities of various types, and disambiguates them against UMLS.

800 free requests / day
Batch processing not available

YODIE Named Entity Disambiguation (English)



Identifies named entities of various types in English documents, and disambiguates them against DBpedia.

800 free requests / day
Batch processing not available

YODIE Named Entity Disambiguation (German)



Identifies named entities of various types in German documents, and disambiguates them against DBpedia.

800 free requests / day
Batch processing not available

YODIE Named Entity Disambiguation

Annotation types: 📍 Mention

On July 26, 1928, 📍 American film director, screenwriter, producer, cinematographer, and editor 📍 Stanley Kubrick was born. He is regarded as one of the greatest filmmakers of all time. His films, typically adaptations of novels or 📍 short stories, are noted for their unique cinematography, attention to details to achieve realism and an inspired use of music scores. Clockwork Orange, 2001, 📍 Full Metal Jacket, 📍 Lolita, 📍 The Shining, 'I'm Spartacus' or 📍 Dr. Strangelove – they all are considered as being some of the best movies ever made. 📍 Stanley Kubrick was born in the Bronx 📍 district of New York, as the first of two children into a family with 📍 Jewish ancestry. As a child, Stanley was considered intelligent, but he did not achieve particularly high grades at school. Concerned to find a remedy for his son's poor academic performance, Jacques (📍 Jacob) Leonard Kubrick sent him to stay with his uncle in 📍 Pasadena, California in 1940, hoping that a change of scenery would also improve Stanley's grades. Returning back home in 1941 for his last year of grammar school, there seemed to be little change in his attitude or his results. But then, Stanley was introduced by his father to the game of chess, with the desired result. 📍 Kubrick took to the game passionately, and quickly became a skilled player. Actually, he later also played chess for money in 📍 Greenwich Village to supplement his income, before making his first movie.

Annotations at this location

Mention

confidence	0.965
dbpInterestingClasses	dbpedia-owl:Person
dbpSpecificClasses	dbpedia-owl:Actor
inst	http://dbpedia.org/resource/Stanley_Kubrick

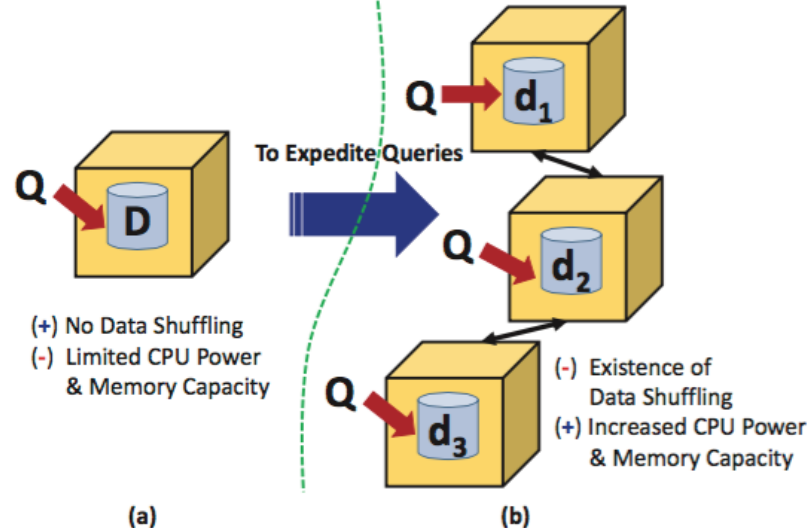
Knowledge Graph Creation

RDF Storage

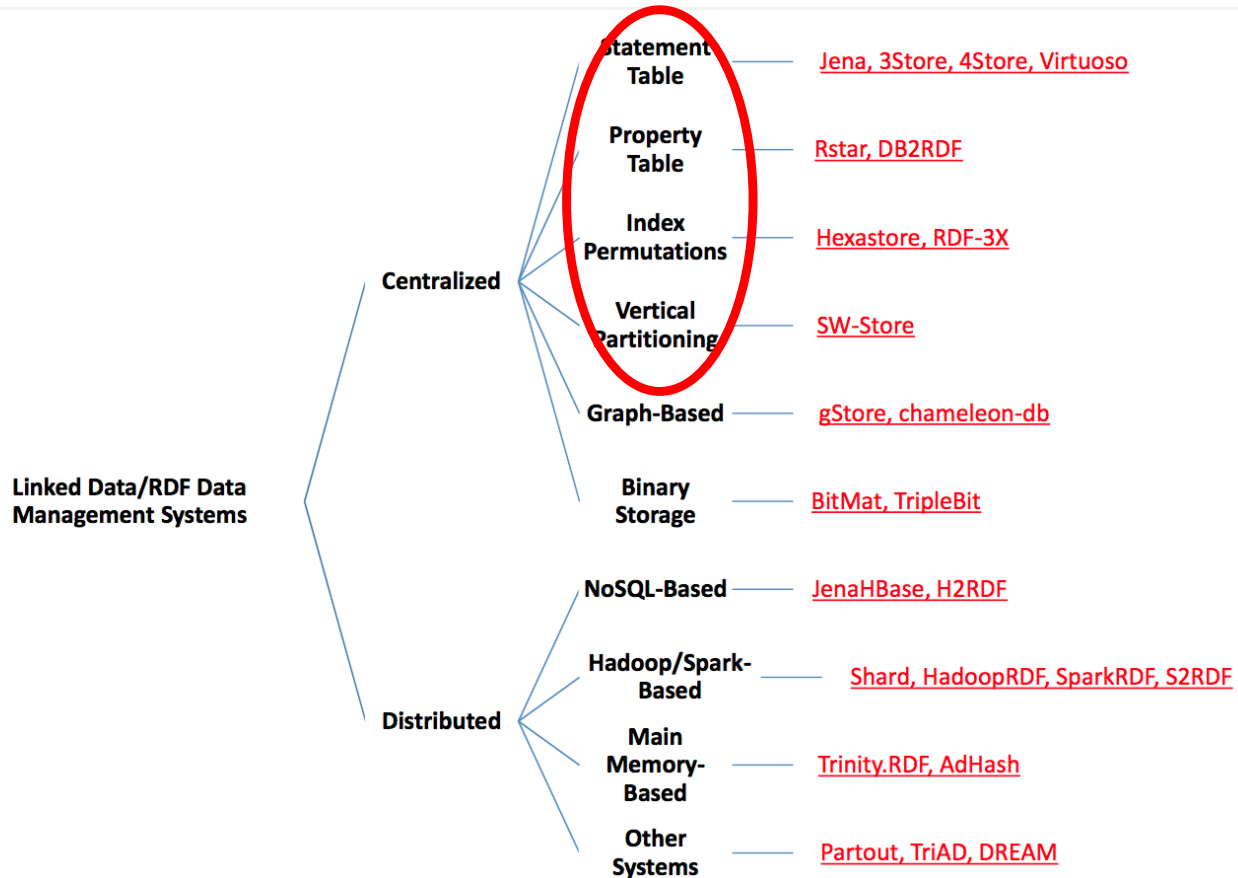
Emanuel Sallinger
Katja Hose

RDF Storage Systems

- Centralized systems (a):
 - storage and processing of RDF data on a single node
- Distributed systems (b):
 - storage and processing of RDF data on a multiple nodes



A Taxonomy of RDF Storage Systems



Statement Table

Triple statements are directly stored in a table-like structure as a linearized list of triples.

Subject	Predicate	Object
Product12345	rdf:type	bsbm:Product
Product12345	rdfs:label	Canon Ixus 2010
Product12345	bsbm:producer	bsbm- inst:Producer1234
...
Producer1234	rdf:label	Canon
Producer1234	foaf:homepage	http://www.canon.com
...

Example systems: Jena SDB, 3Store, 4Store, Virtuoso.

Property Table

Cluster commonly accessed nodes in the graph together in a single table to avoid the expensive cost of many self-join operations on the large statement table encoding the RDF data

Product Property Table

Subject	Type	Label	NumericProperty1	aaa
Product12345	bsbm:Product	Canon Ixus 2010	NULL	...
...

Left-Over Triples


Subject	Predicate	Object
Producer1234	foaf:homepage	http://www.canon.com
...

- Example systems: DB2RDF, Jena2 SDB.

Index Permutations

All possible combinations the three components of triples is indexed and materialized

Rationale: any query can be answered using the available indices so that it allows fast access to all parts of the triples by sorted lists and fast merge-joins



- Example systems: Hexastore, RDF-3x.

Vertical Partitioning

Creates a table for each unique property in the dataset

- Containing subjects and objects connected by the property
- For multi-valued attributes, each distinct value is listed in a successive row
- Tables are indexed by subject for quick retrieval

<rdf:type>

Subject	Object
Product12345	bsbm:Product

<rdfs:label>

Subject	Object
Product12345	Canon Ixus 2010
Producer1234	Canon

<aaa>

Subject	Object
uuu	xxx
...	...


- Example systems: SW-Store.

**188.399 Introduction to
Semantic Systems**
(2.0 VU / 3.0 ECTS)
2023W

Katja Hose
Emanuel Sallinger

iss@dbai.tuwien.ac.at

Short History of Knowledge Graphs



- Expressing meaning
- Knowledge representation
- Ontologies
- Agents
- Evolution of knowledge

Killer app: Agents

- Retrieve information about prescribed treatment/therapy
- Look for providers/specialists with good ratings in proximity
- Schedule times according to multiple busy schedules
- Share information with other agents and “negotiate”

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](#)

The entertainment system was belting out the Beatles' "We Can Work It Out" when the phone rang. When Pete answered, his phone turned the sound down by sending a message to all the other *local* devices that had a *volume control*. His sister, Lucy, was on the line from the doctor's office: "Mom needs to see a specialist and then has to have a series of physical therapy sessions. Biweekly or something. I'm going to have my agent set up the appointments." Pete immediately agreed to share the chauffeuring. At the



doctor's office, Lucy instructed her Semantic Web agent through her handheld Web browser. The agent promptly retrieved information about Mom's *prescribed treatment* from the doctor's agent, looked up several lists of *providers*, and checked for the ones *in-plan* for Mom's insurance within a *20-mile radius* of her home and with a *rating of excellent or very good* on trusted rating services. It then began trying to find a match between available *appointment times* (supplied by the agents of individual providers through their Web sites) and Pete's and Lucy's busy schedules. (The emphasized keywords indicate terms whose semantics, or meaning, were defined for the agent through the Semantic Web.)

In a few minutes the agent presented them with a plan. Pete didn't like it—University Hospital was all the way across town from Mom's place, and he'd be driving back in the middle of rush hour. He set his own agent to redo the search with stricter preferences about *location* and *time*. Lucy's agent, having *complete trust* in Pete's agent in the context of the present task, automatically assisted by supplying access certificates and shortcuts to the data it had already sorted through.

“A really important thing about data is:
the more things you have to connect
together, the more powerful it is.”



Google ✕ 🔍 Sign in

[All](#) [Images](#) [Videos](#) [News](#) [Maps](#) [More](#)

[Settings](#) [Tools](#)

About 463,000,000 results (0,73 seconds)

Top stories

Stay, go, delay: Joe Biden is trapped and has no good choices in Afghanistan

The Guardian · 14 mins ago

Joe Biden won the presidency by making the most of his lucky breaks

The Washington Post · 6 hours ago

[View all](#)

en.wikipedia.org > wiki > Joe_Biden

Joe Biden - Wikipedia

Joseph Robinette Biden Jr. (born November 20, 1942) is an / 46th president of the United States. A member of the Democr

Vice President: [Kamala Harris](#) Parents: Jo
President: [Barack Obama](#) Occupation

[Family of Joe Biden](#) · [Presidency of Joe Biden](#) · [Timeline of th](#)

www.whitehouse.gov > administration > president-biden

Joe Biden: The President | The White Hous

Joe Biden · [Joe Biden's Family](#) · [A Leader in the Senate](#) · [TI United States](#) · [A New Chapter.](#)

Microsoft Bing 🗨️ 🔍 🔄 🔍

ALL NEWS IMAGES VIDEOS

51,700,000 Results [Open links in new tab](#)



Joe Biden

[Quotes](#)

[Speeches](#)

[Career](#)

[Early life](#)

News about Joe Biden

bing.com/news



Donald Trump at CPAC: Ex-president expected to tear into ...

Trump is not expected to declare a 2024 presidential candidacy - he is likely to discuss ...

MSN · 21m



Joe Biden Is Covertly Defunding the Police by Matt Vespa

Townhall · 22h



Joe Biden's COVID relief bill with stimulus checks passes ...

Des Moines Register · 2d



Trump's unconditional retreat from Afghanistan leaves ...

USA Today · 2d

[See more news about joe biden](#)



Joe Biden

Joseph Robinette Biden Jr. is an American politician serving as the 46th president of the United States. A member of the Democratic Party, he served as the 47th vice president from 2009 to 2017 under Barack Obama and represented Delaware in the United States Senate from 1973 to 2009.



[Wikipedia](#)



[Facebook](#)



[Instagram](#)



[Official site](#)



[IMDb](#)

Born: 11/20/1942 (age 78) · [Scranton, PA](#)

Height: 1.83 m

Inauguration date: [January 20, 2021](#)

Transition plan: [Build Back Better](#)

“

Knowledge graphs provide the structured data and factual knowledge that drive many products and make them more intelligent and "magical".

”

N. Noy, Y. Gao, A. Jain, A. Narayanan, A. Patterson, J. Taylor:
 Industry-scale knowledge graphs: lessons and challenges.
 Commun. ACM 62(8): 36-43 (2019)

Common characteristics of the knowledge graphs.

	Data model	Size of the graph	Development stage
Microsoft	The types of entities, relations, and attributes in the graph are defined in an ontology.	~2 billion primary entities, ~55 billion facts	Actively used in products
Google	Strongly typed entities, relations with domain and range inference	1 billion entities, 70 billion assertions	Actively used in products
Facebook	All of the attributes and relations are structured and strongly typed, and optionally indexed to enable efficient retrieval, search, and traversal.	~50 million primary entities, ~500 million assertions	Actively used in products
eBay	Entities and relation, well-structured and strongly typed	Expect around 100 million products, >1 billion triples	Early stages of development and deployment
IBM	Entities and relations with evidence information associated with them.	Various sizes. Proven on scales documents >100 million, relationships >5 billion, entities >100 million	Actively used in products and by clients





WIKIDATA

1.9B Facts



210M Facts



52M Facts



<http://linkedlifedata.com/sources.html>

6.7B Facts



<https://pubchem.ncbi.nlm.nih.gov/docs/rdf>


132B Facts

2022

- Constantly growing
- Ca. 1255 datasets
- Ca. 400 billion triples

Linked Data principles

1. Use URIs as names for things
2. Use HTTP URIs so that people can look up those names
3. When someone looks up a URI, provide useful information
4. Include links to other URIs so that they can discover more things



Take back control of your data


The New York Times

SUBSCRIBE NOW

Digital Privacy | How to Read Privacy Labels | How to Protect Your Digital Life | 10 Simple Tips | How Apps Collect Your Data

He Created the Web. Now He's Out to Remake the Digital World.

Tim Berners-Lee wants to put people in control of their personal data. He has technology and a start-up pursuing that goal. Can he succeed?



Tim Berners Lee's startup Inrupt releases Solid privacy platform for enterprises

Ron Miller @ron_miller / 6:01 AM GMT+1 • November 9, 2020

Comment

The screenshot shows a Financial Times article page. At the top, the 'FINANCIAL TIMES' logo is visible. Below it, a navigation bar includes 'HOME', 'WORLD', 'US', 'COMPANIES', 'TECH', 'MARKETS', 'CLIMATE', 'OPINION', 'WORK & CAREERS', 'LIFE & ARTS', and 'HOW TO SPEND IT'. The article is categorized under 'Technology' with a '+ Add to myFT' button. The main headline reads 'NHS signs up for Tim Berners-Lee pilot to reinvent web'. Below the headline, a sub-headline states 'BBC, NatWest and government of Flanders also test new Solid servers from Inrupt'. A large photograph of Tim Berners-Lee speaking at a microphone is featured. Below the photo, a caption reads: 'Tim Berners-Lee: "There is a feeling, a zeitgeist, that change is really overdue." © AFP/Getty Images'. At the bottom of the article preview, it says 'John Thornhill in London NOVEMBER 9 2020' and shows a comment icon with the number '118' and a share icon.






— — **In a data fabric, knowledge graphs connect** siloed data and **simplify** complicated data.

“A knowledge graph in a data fabric solves the data integration task and bridges the gap between disparate corporate datasets and the actual business requirements.

It allows business users to model, explore and discover how datasets are connected in a human readable model where there is little difference between the conceptual and the physical machine readable models.”

Source: Demystifying the Data Fabric | 9.17.20 | Gartner Group



A knowledge graph built on semantics and graph data models fits into the data orchestration and discovery layers in a data fabric.

The Data Fabric is an architecture for modern data management that anticipates the need to connect data across the enterprise at speed and scale.

Knowledge graph is central to the success of any Data Fabric as an agile integration overlay that blends complex and diverse data into easy-to-consume data products for end-users, data scientists, and applications.






Pure Logic Reasoning

Pure Learning



Pure Logic Reasoning

Pure Learning


$$(Katja, \text{rdf:type}, Woman) \wedge (Woman, \text{subClassOf}, Person) \rightarrow (Katja, \text{rdf:type}, Person)$$



Pure Logic Reasoning

Pure Learning




Person
Male

Person
Female


PC chairs opening EDBT 2023

role

conference



Who is who in the picture?
And how certain are you?



PC chairs opening EDBT 2023
role conference

EDBT 2023 has two PC chairs*
These are Katja and Sourav
Katja is a Woman, Sourav is a Man
Men and Women are Persons

The PC chairs of EDBT 2023 are (likely) on the picture
One is (likely) a male person, the other one (likely) female

* Partial Completeness Assumption (PCA)
Galárraga, Teflioudi, Hose, Suchanek. Fast Rule Mining in Ontological Knowledge Bases with AMIE+, VLDBJ 2015

Data and Knowledge Engineering




Data Management



Data Analytics



Machine Learning



Extraction and Integration

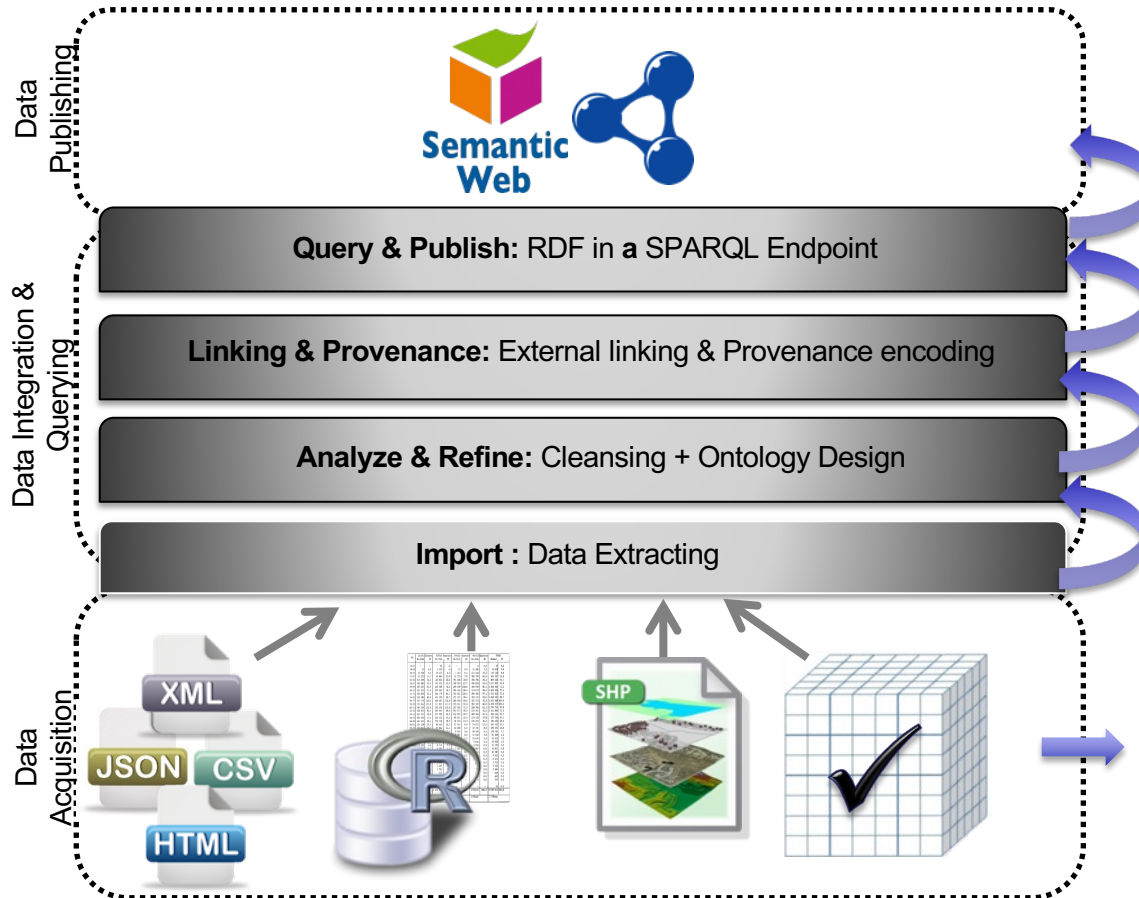


Query Processing and Optimization



Provenance, Evolution, Metadata, and Privacy





Task: Integrate data and updates into a knowledge graph & enable efficient querying and analytics

Challenges: **Statistical knowledge graphs** require novel data modeling, management, and query optimization techniques

Existing standards: RDF Cube (QB), QB4OLAP, Provenance (PROV-O), RDF-star

Multidimensional Data Cubes

Nath, Romero, Pedersen, Hose
High-Level ETL for Semantic Data
Warehouses. In: SWJ 2021

Gür, Pedersen, Hose, Midtgaard
Multidimensional Enrichment of Spatial RDF
Data for SOLAP. In: SWJ 2021

Ok Google, what's the capital of Denmark?

The capital of Denmark is Aalborg.

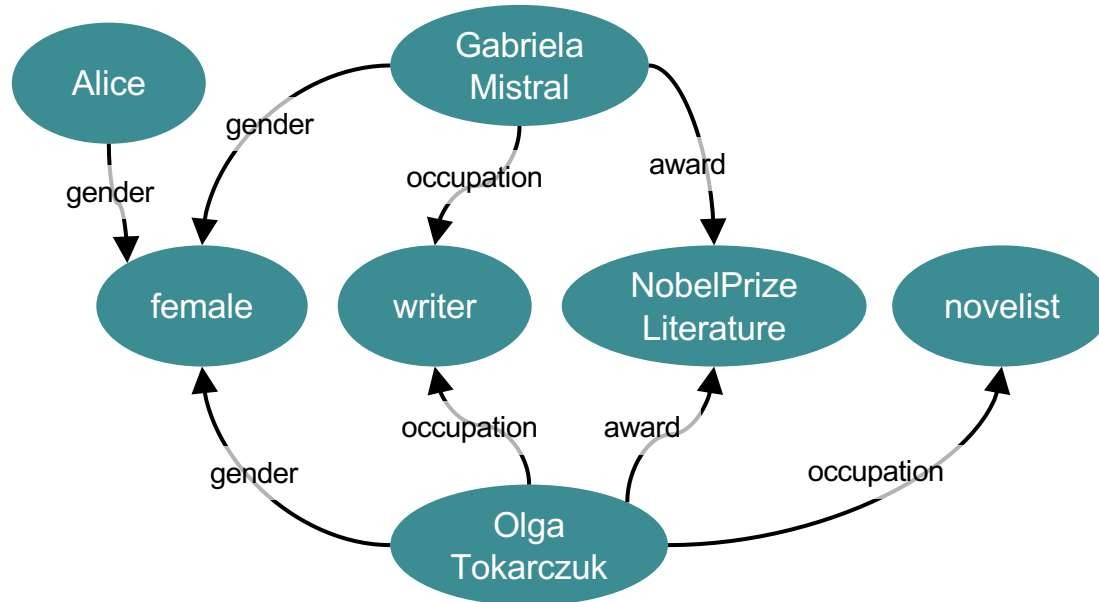
Is there a war in Ukraine?

No, there's a special military operation though.

???



Example Knowledge Graph




How-provenance: How are query answers generated?

Example Query

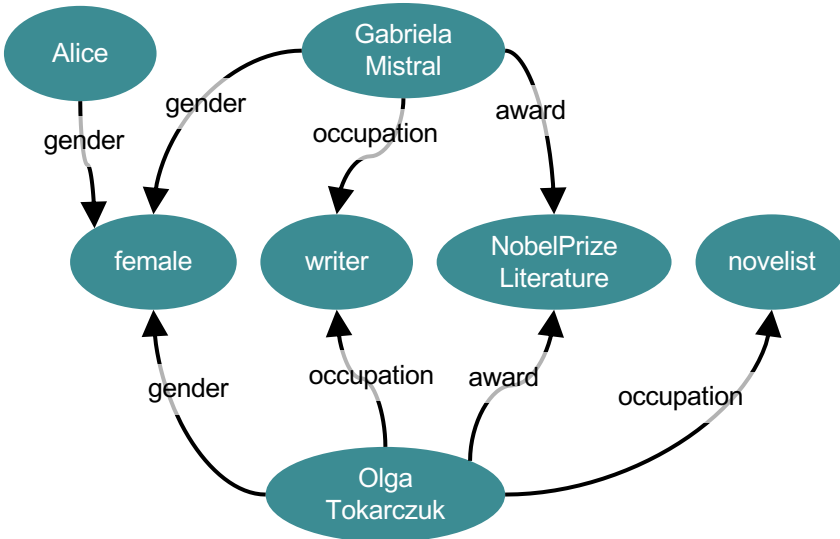
Which women have occupations?

```
SELECT ?x
WHERE {
  ?x gender 'female' .
  ?x occupation ?y
}
```



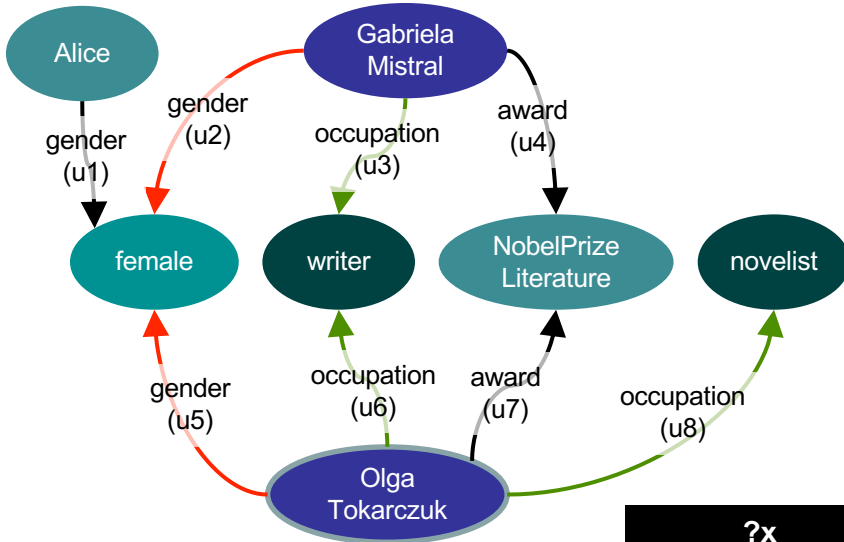
Answers

?x
GabrielaMistral
OlgaTokarczuk



How-provenance: How are query answers generated?


Adding a unique identifier per edge (triple) enables us to track provenance.



Which women have occupations?

```


SELECT ?x
WHERE {
  ?x gender 'female' .
  ?x occupation ?y
}
  
```




Answers

?x	?x	ance
GabrielaMistral	GabrielaMistral	i3
OlgaTokarczuk	OlgaTokarczuk	u5 ⊗ u8

?s	?p	?o
<uri1>	<p1>	<o1>
<uri1>	<p2>	<o2>
<uri2>	<p1>	<o3>
<uri2>	<p3>	<o4>
<uri3>	<p1>	<o1>
<uri4>	<p2>	<o5>
...
<uri22>	<p1>	<o14>
<uri22>	<p3>	<o15>



?s	<p1>	<p2>	<p3>	...
<uri1>	<o1>	<o2>	-	...
<uri2>	<o3>	-	<o4>	...
<uri3>	<o1>	-	-	...
<uri4>	-	<o5>	-	...
...
<uri22>	<o14>	-	<o15>	...



Knowledge graph management and querying

KG Management

dump centralized client/server distributed/parallel federated P2P streams

Tentris [Bigerl et al. 2020]
 Tunable-LSH [Aluç et al. 2019]
MARVEL [Ibragimov et al. 2016]
 RDFox [Nenov et al. 2015]
 chameleon-db [Aluç...2013]
 TripleBit [Yuan et al. 2013]
 gStore [Zou et al. 2011]
 dipLODocuRDF [Wylot...2011]
 BitMat [Atre et al. 2009]

WiseKG [Azzam et al. 2021]
SPF [Aebeloe et al. 2020]
 SmartKG [Azzam et al. 2020]
 SaGe [Minier et al 2019]
 brTPF [Hartig et al. 2016]
 TPF [Verborgh et al. 2016]
 ...

DP2RPQ [Wang et al. 2020]
 gStore-D2 [Peng et al. 2019]
 WORQ [Madkour et al. 2018]
 S2RDF [Wu et al. 2017]
 DiploCloud [Wylot et al. 2016]
 CliqueSquare [Goasdoué...2015]
Partout [Galárraga et al. 2014]
WARP [Hose et al. 2013]
 TriAD [Gurajada et al. 2014]
 H2RDF [Papailiou et al. 2013]
 SHAPE [Zhang et al. 2013]
 EAGRE [Zhang et al. 2013]
 Trinity.RDF [Zeng et al. 2013]
 H-RDF-3X [Huang et al. 2011]
 HadoopRDF [Husain et al. 2011]
 SHARD [Rohloff et al. 2011]

MultiQuery [Peng et al. 2021]
Odyssey [Montoya et al 2017]
 Lusail [Mansour et al. 2017]
 SemaGrow [Charalambidis...2015]
CoDA [Ibragimov et al. 2015]
 HiBiSCuS [Saleem et al.2014]
FedX [Schwarte et al. 2011]
 SPLENDID [Görlitz et al. 2011]
 Anapsid [Acosta et al.2011]
 DARQ [Quillitz et al. 2008]

Lothbrok (Aebeloe et al. 2023)
ColChain [Aebeloe et al. 2021]
PIQNIC [Aebeloe et al. 2019]
 Cyclades [Folz et al. 2016]
 RDFpeers [Cai et al. 2014]
 ATLAS [Kaoudi et al.2010]
 Unistore [Karnstedt et al.2007]
 Bibster [Haase et al. 2004]
 InfoQuilt [Arumugam...2002]





runtime

throughput

runtime

runtime


availability



Federated architectures

SPARQL Query

```
SELECT DISTINCT * WHERE {  
  ?album dbo:artist ?artist .  
  ?album rdfs:type dbo:album .  
  ?album dbo:releaseDate ?date .  
  ?artist dbo:genre ?genre .  
  ?artist foaf:name ?name .  
  ?song dbo:writer ?artist  
}
```



Federation Engine

Sources

Local Optimization & Execution



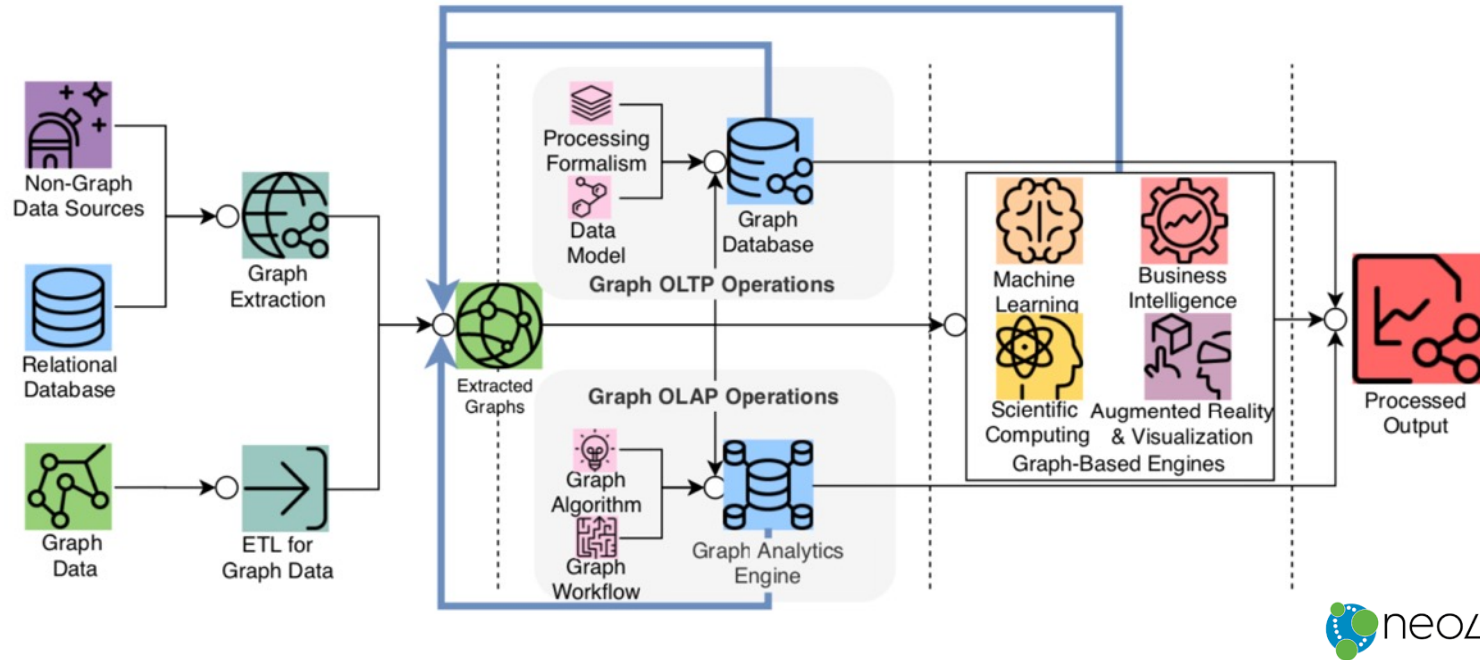
Query Result

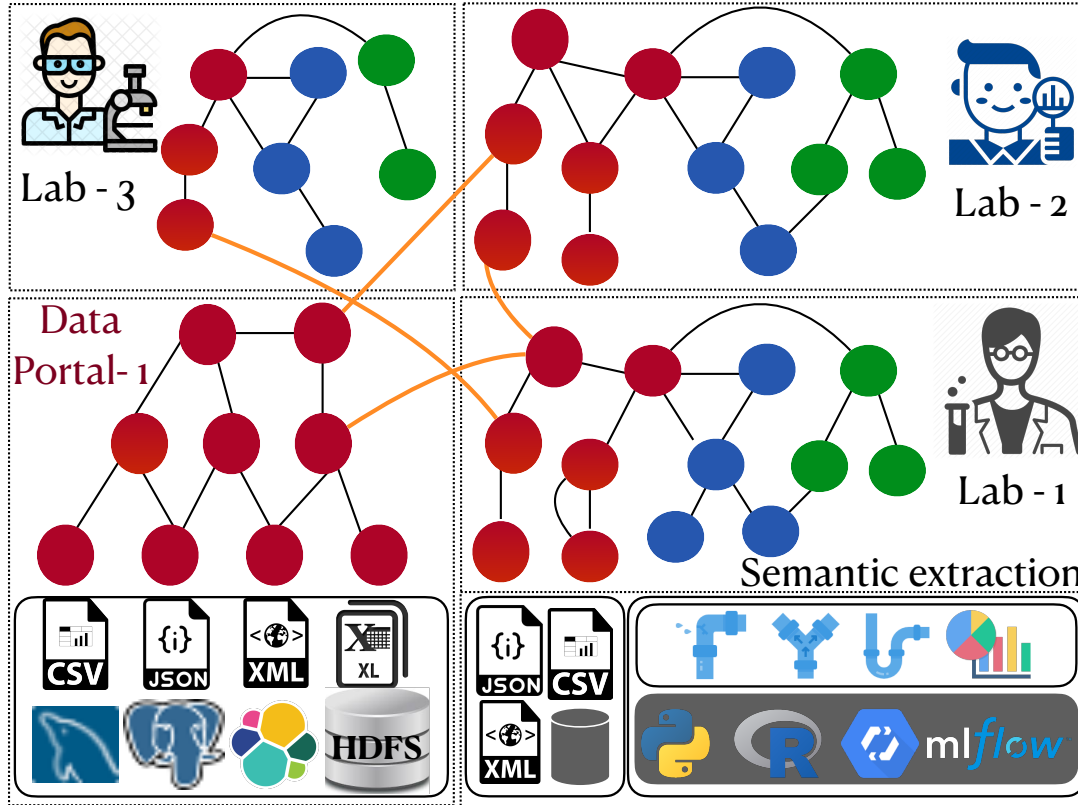
```
{?album=A_Night_at_the_Opera_(Queen_song),  
 ?artist=dbr:Queen_(band),  
 ?date="1975-11-21",  
 ?genre=dbr:Rock_music,  
 ?name="Queen"@en,  
 ?song=dbr:Bohemian_Rhapsody},  
 ...
```

<https://relweb.cs.aau.dk/>

Hernández, Galárraga, Hose: Computing How-Provenance for SPARQL Queries via Query Rewriting. PVLDB 14(13): 3389-3401, 2021

Azzam, Aebeloe, Montoya, Keles, Polleres, Hose: WiseKG: Balanced Access to Web Knowledge Graphs. TheWebConf, 1422-1434, 2021





Use Cases



Manufacturing and Production



Machine Learning



Sustainability and
Circular Economy



Medicine and Bio Science



Media and Entertainment



Finance and Economics



Transportation



Manufacturing and Production




Machine Learning



Sustainability and
Circular Economy



Medicine and Bio Science



Media and Entertainment




Finance and Economics




Transportation








Patient data and electronic health records



Manufacturing and Production




Machine Learning



Sustainability and
Circular Economy



Medicine and Bio Science



Media and Entertainment



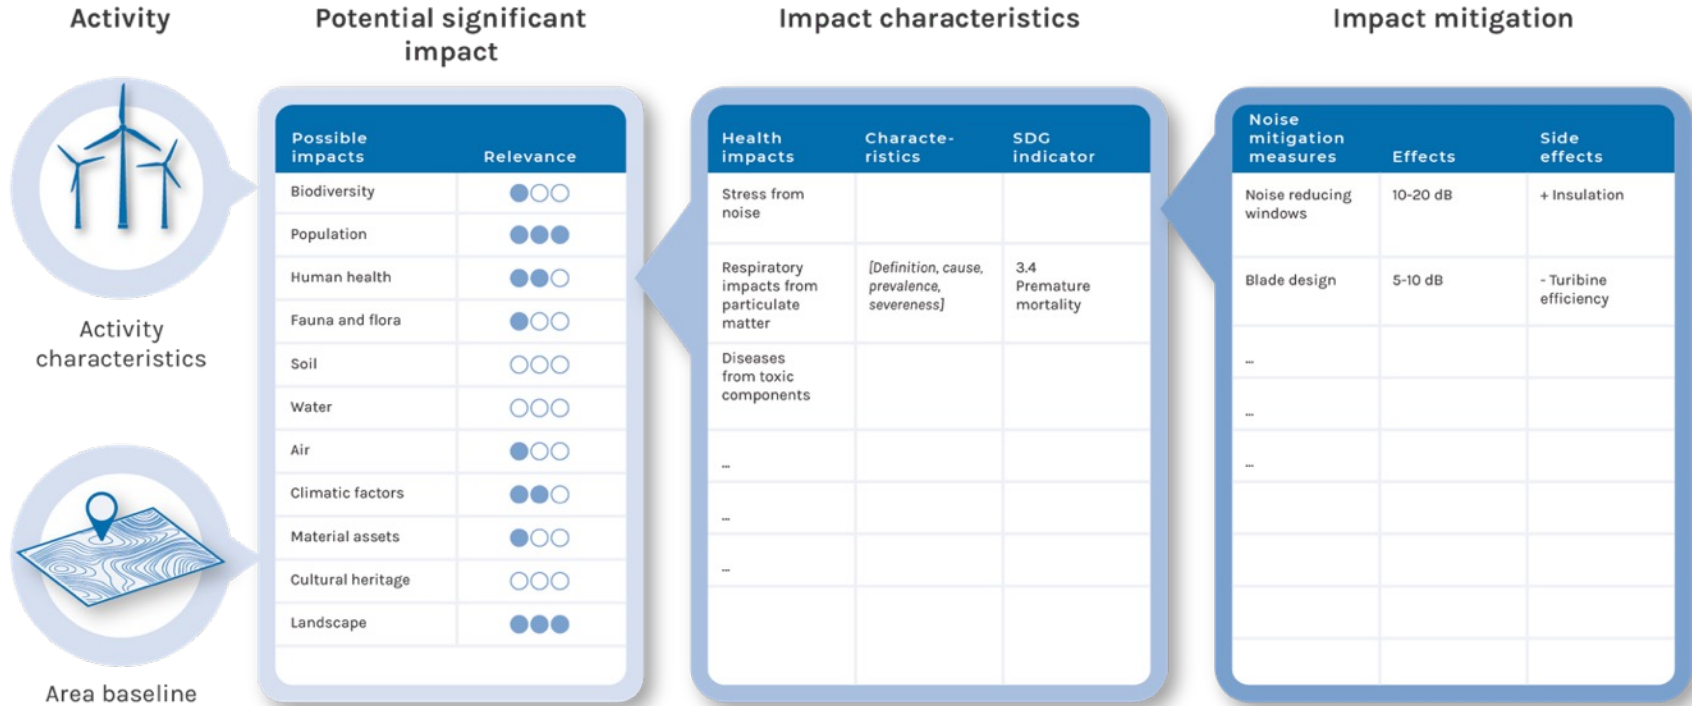
Finance and Economics



Transportation




Digitally supported environmental assessments



<https://dreamsproject.dk>

K. Hose, J. Bjerva: Fra korrelation til kausalitet: ML, NLP og relation extraction skal forstå miljørapporter. In: ING/DATATECH Synspunkt, 2021

Manufacturing and Production



Machine Learning



Sustainability and
Circular Economy



Medicine and Bio Science



Media and Entertainment




Finance and Economics




Transportation






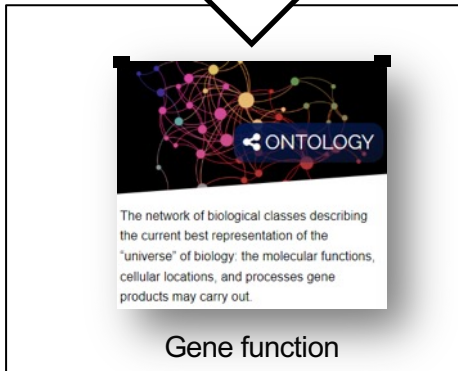

Metadata

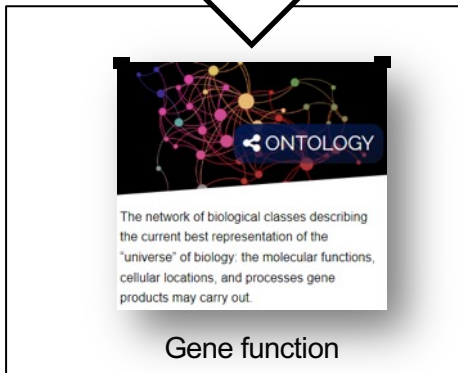



Satellite & drone imaging




Soil properties





Satellite & drone imaging



Metadata




External ontologies, legacy databases, Open Data, FAIR data, etc.

GEO Gene Expression Omnibus

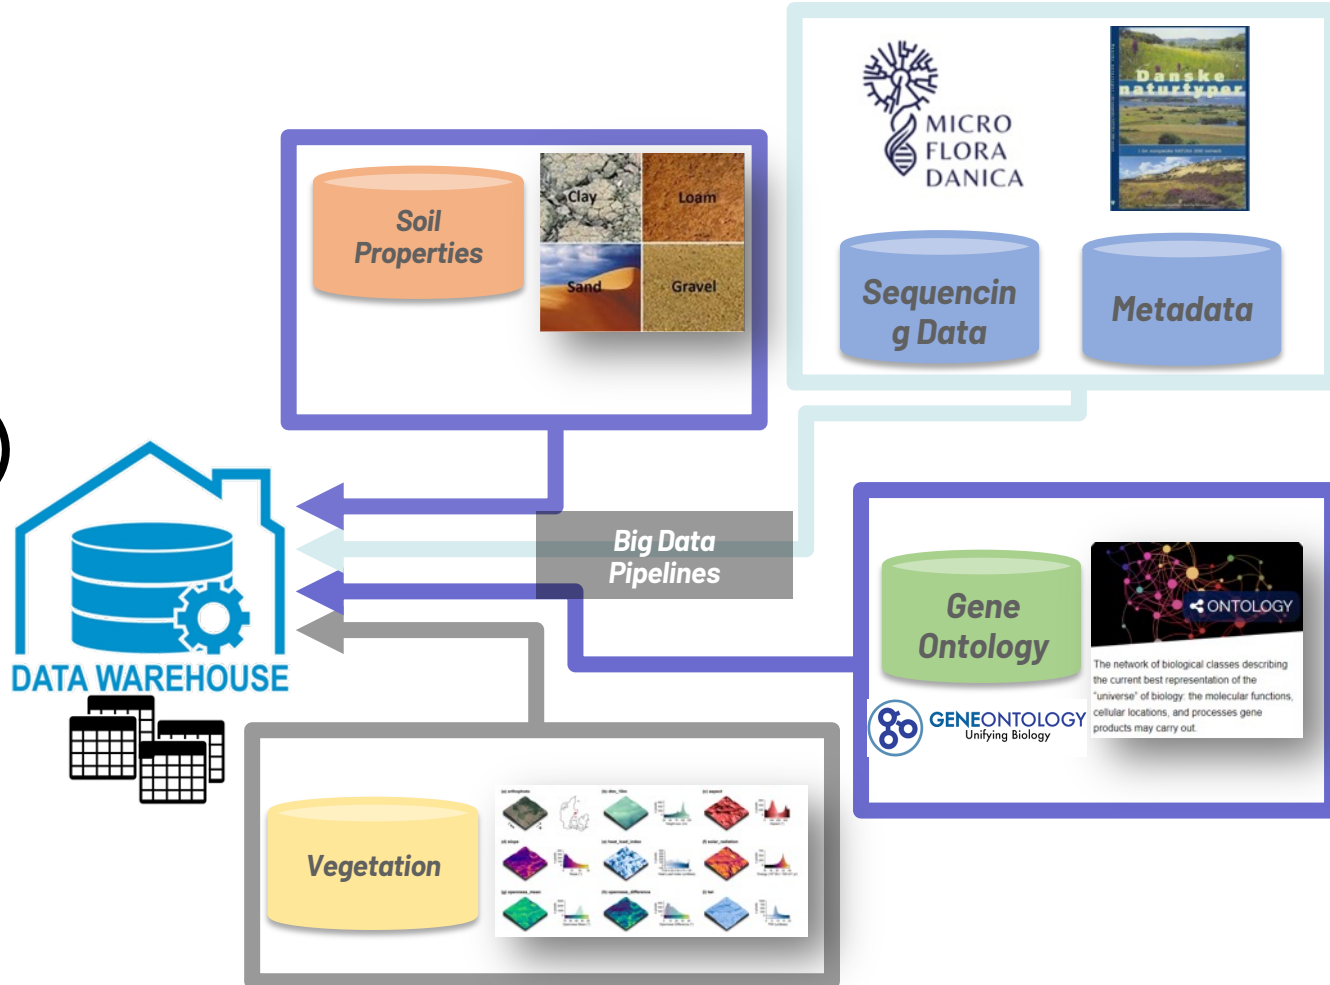
F A I R
Findable Accessible Interoperable Reusable

Soil properties



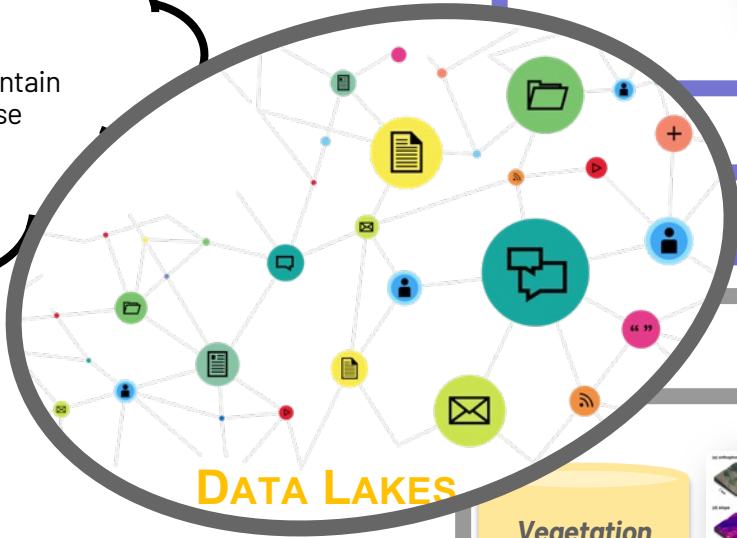
Traditional approach: data warehouses

Which habitats contain the most diverse microbes?




New approach: data lakes and data fabrics integrate data in a knowledge graph

Which habitats contain the most diverse microbes?



Soil Properties



Sequencing Data

Metadata


Big Data Pipelines

Gene Ontology



The network of biological classes describing the current best representation of the "universe" of biology: the molecular functions, cellular locations, and processes gene products may carry out.

Vegetation



New approach: data lakes and data fabrics integrate data in a knowledge graph



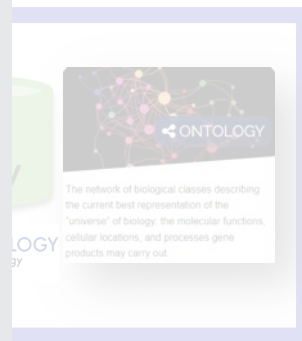
Updates Available!



Which habitats contain the most diverse microbial communities?



GOVERNANCE





Knowledge Graphs

Real-World Application Teaser

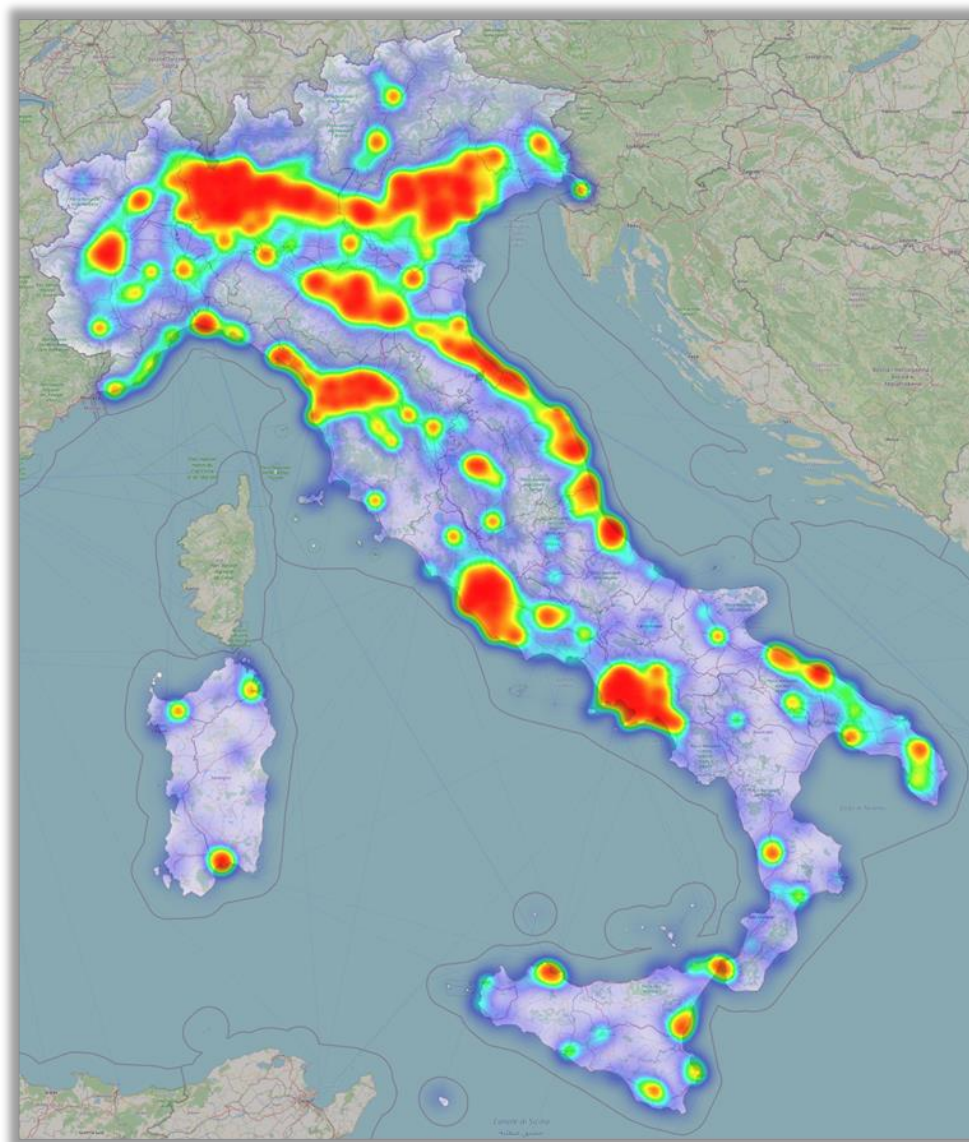
Teaser

Emanuel Sallinger

Confidential:
registered participants only, not for resharing



Economic Impact of Crises





Hostile Takeovers



REUTERS Business Markets World Politics TV More

BUSINESS NEWS MARCH 26, 2020 / 11:36 AM / 2 MONTHS AGO

EU leaders to shield strategic firms from hostile interest amid crisis

Francesco Guarascio, Gabriela Baczynska 4 MIN READ

BRUSSELS (Reuters) - European Union leaders will on Thursday shield strategic firms in sectors such as healthcare, infrastructure and other firms seen as having strategic value from hostile foreign takeovers, draft EU summit conclusions show.





Anti-Money Laundering

FINANCIAL TIMES

Danske Bank AS [+ Add to myFT](#)

Danske Bank chairman ousted by main shareholder after scandal

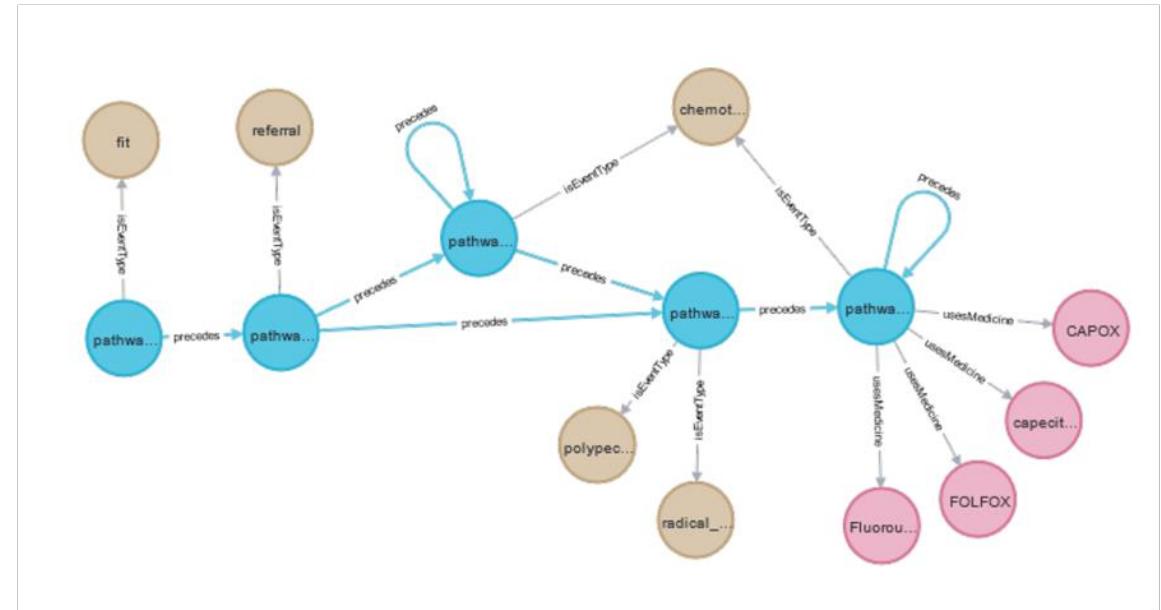

Maersk family brings in new blood to stabilise lender in wake of €200bn money laundering



€200bn money laundering




Cancer Pathways






Supply Chains





Media Intelligence






Open Knowledge Network Innovation Sprint





Baby Formula Supply Chain Issue



Individual

Supply Chain
+ Finance

PPP / US CARES Effectiveness



Small Business

COVID Impact
+ Finance

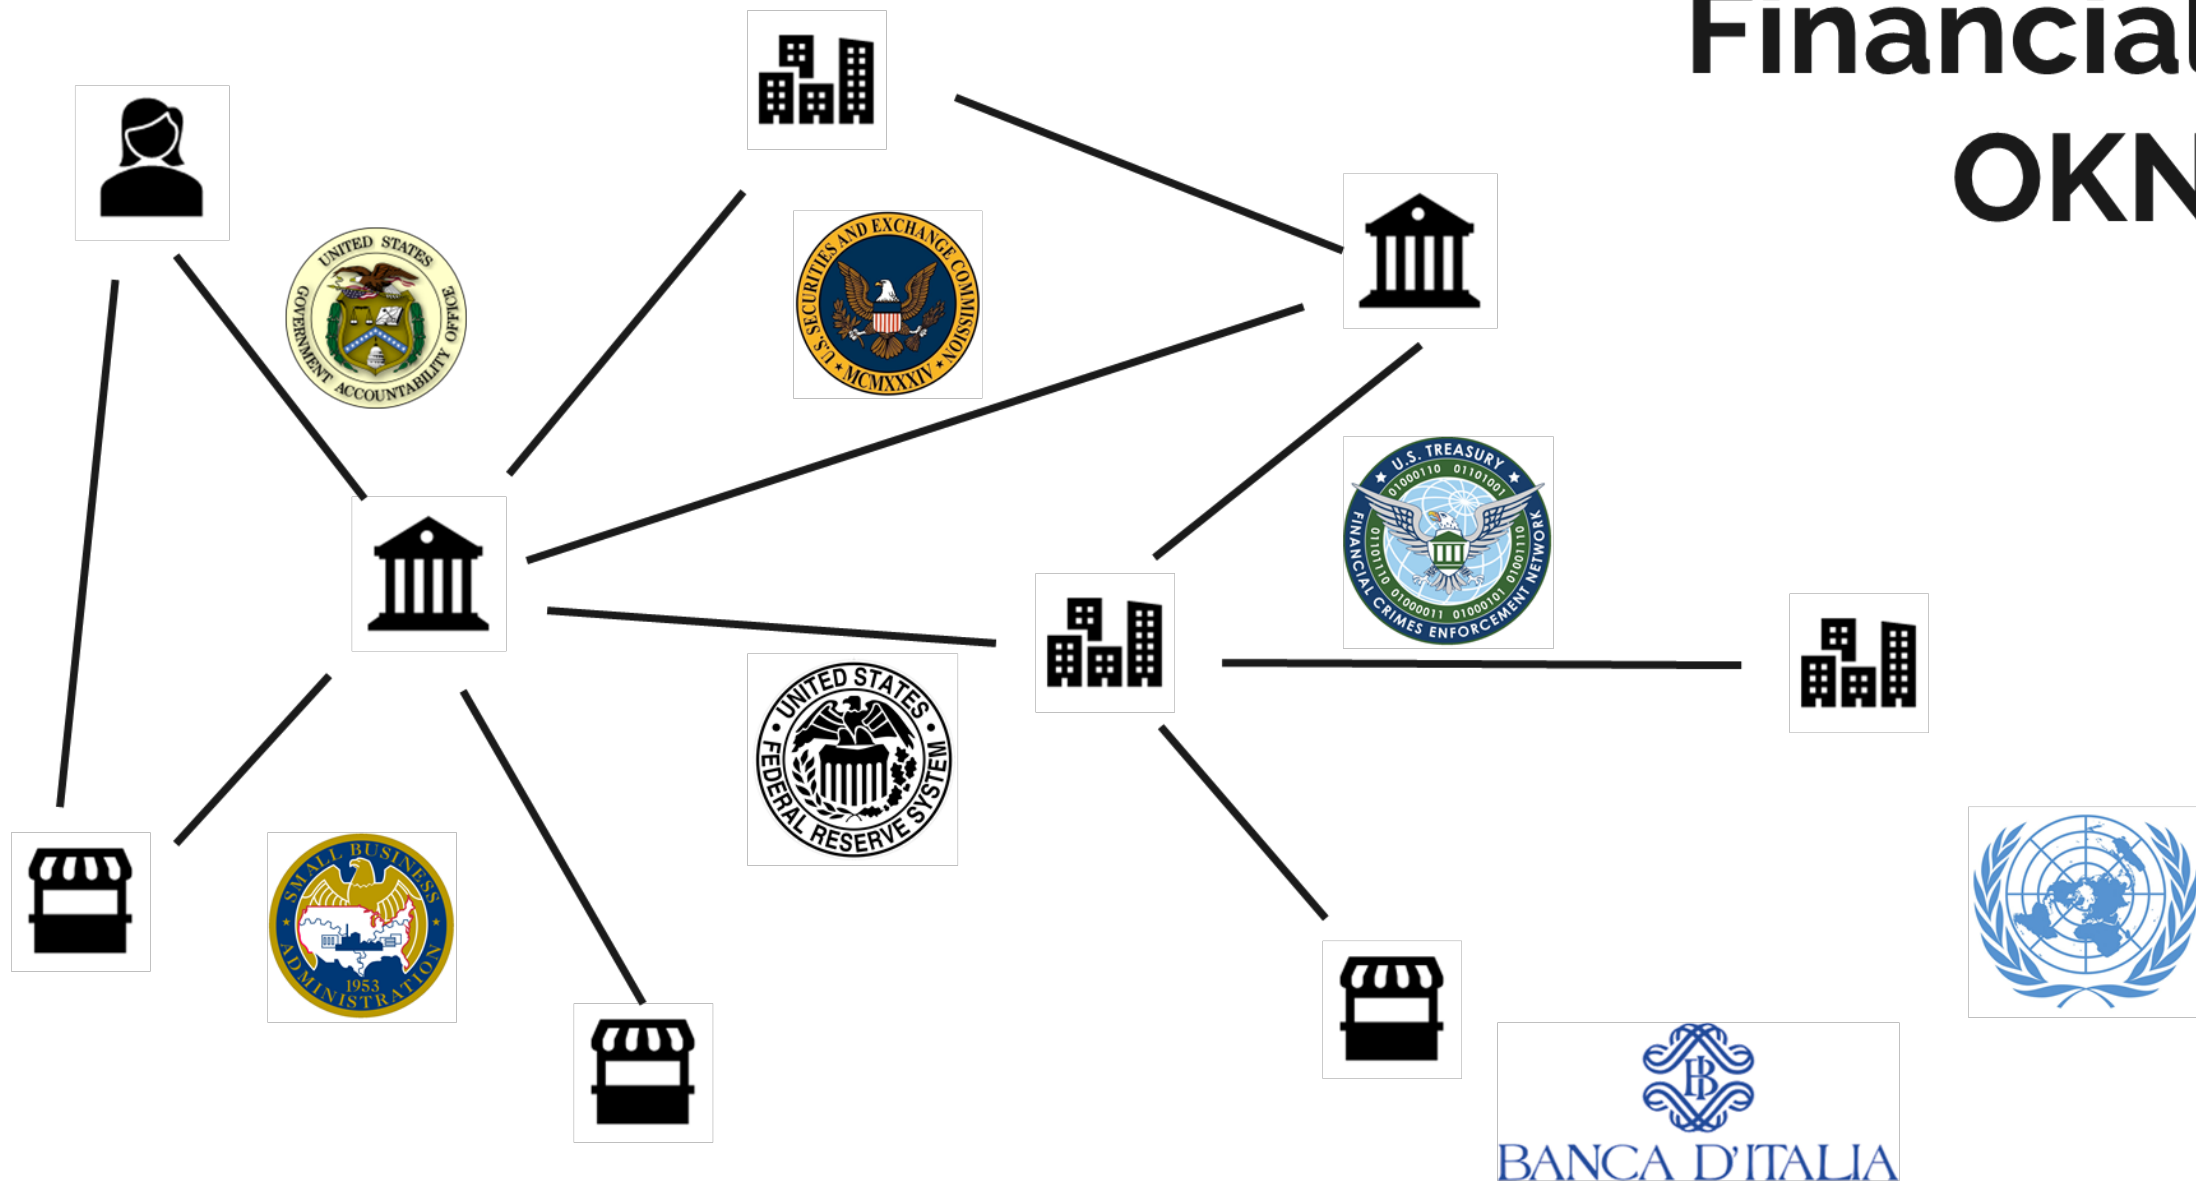
Climate Finance



Environmental Models
+ Finance



Financial OKN





Course Teaser

"192.116 Knowledge Graphs"

Teaser

Emanuel Sallinger

Features slides with confidential material:
registered participants only, not for resharing



“What is the faculty of informatics at TU Wien”?

The image shows a Google search interface. The search bar contains the text "What is the faculty of informatics at TU Wien?". Below the search bar, there are navigation options: All, Images, Maps, News, Shopping, and More. The search results show "About 489,000 results (0,95 seconds)".

The first search result is for "informatics.tuwien.ac.at" with the title "TU Wien Informatics". The description states: "TU Wien Informatics is one of Europe's leading research, teaching, and innovation institutions in computer science and Austria's largest faculty of its kind." It lists links for Professors, Bachelor Programs, Master Programs, and Organizational Structure.

The second search result is for "informatics.tuwien.ac.at > people > professors" with the title "Professors — TU Wien Informatics". The description says: "Our dedicated faculty fuel the dynamic environment in both research and teaching. They are a creative, diverse and passionate group of excellent researchers ..."

The third search result is for "informatics.tuwien.ac.at > orgs" with the title "Organizational Structure — TU Wien Informatics". The description says: "TU Wien Informatics is organized in 4 institutes, 22 research units, the Office of the Dean, and a Communications and IT ... Services of the Faculty of Informatics"

The fourth search result is for "www.tuwien.at > studies > studies > bachelor-programmes" with the title "Computer Science and Business Informatics | TU Wien". The description says: "Study content. The bachelor's program Media Informatics and Visual Computing combines the teaching of technologies and methods in the fields of computer vision, computer graphics, visualization and augmented/mixed/virtual reality with education in the design of innovative interfaces."

The fifth search result is for "www.tuwien.at > studies > studies > master-programmes" with the title "Computer Science | TU Wien". The description is partially visible: "The EPIC ... as it is known, is a first-class in studies conducted at the Faculty of ..."

On the right side of the search results, there is a knowledge panel for "TU Wien Informatics". It features a map of Wieden, ORF, and Junkhaus Wien. Below the map are buttons for "Website", "Directions", and "Save". The panel shows a rating of 4.9 stars from 11 Google reviews and identifies it as a "University department in Vienna". It provides the following information:

- Located in: Vienna University of Technology
- Address: Favoritenstraße 9/11, 1040 Wien
- Hours: Closed · Opens 8AM Mon
- Phone: 01 5880119501

There is a warning icon and text: "Hours or services may differ". Below this are links for "Suggest an edit" and "Own this business?". The panel also has a "Questions & answers" section with the text "Be the first to ask a question" and an "Ask a question" button. At the bottom, there are links for "Reviews", "Write a review", and "Add a photo".



"How likely will the largest bank in Italy be taken over in 2023?"

The image shows a Google search interface. The search bar contains the text "how likely will intesa sanpaolo be taken over in 2023?". Below the search bar, there are navigation tabs for "All", "News", "Books", "Videos", "Shopping", and "More". The search results are displayed below, showing about 14,800,000 results in 0.54 seconds. The first result is from Reuters, titled "Intesa sets prudent 2023 goal as its cuts assets to boost capital", dated 3 days ago. The second result is from the Intesa Sanpaolo website, titled "Full year 2022 results – Highlights - Gruppo Intesa Sanpaolo", dated 3 days ago. The third result is also from the Intesa Sanpaolo website, titled "2023 Financial Calendar - Gruppo Intesa Sanpaolo", dated Dec 21, 2022. The fourth result is from the Intesa Sanpaolo website, titled "Italy can beat the crisis". Carlo Messina looks ahead to 2023, dated Oct 14, 2022. A large yellow rectangular box highlights the right side of the search results area.



"How likely will the largest bank in Italy be taken over in 2023?"

how likely will intesa sanpaolo be taken over in 2023?



It is difficult to predict with certainty whether Intesa Sanpaolo will be taken over in 2023. Takeovers are typically the result of a complex set of factors, including the strategic goals of the potential acquirer, the financial performance of the target company, and regulatory approval.

That being said, Intesa Sanpaolo is a large and well-established banking group, and it would likely take a significant offer from a well-capitalized acquirer to successfully complete a takeover. Additionally, the bank's strategic importance to the Italian economy would likely be a major consideration for regulators and government officials.

Therefore, it is not likely that Intesa Sanpaolo will be taken over in 2023 but it can't be completely ruled out. As the year progresses, it is important to keep an eye on any developments in the merger and acquisition market that might affect the bank's future.





Knowledge Graphs

1. The **technology** used by Google and others
2. A meeting place of **databases, data science** and **Artificial Intelligence** research
3. A **skillset** to solve fascinating problems



Anti-Money Laundering

FINANCIAL TIMES

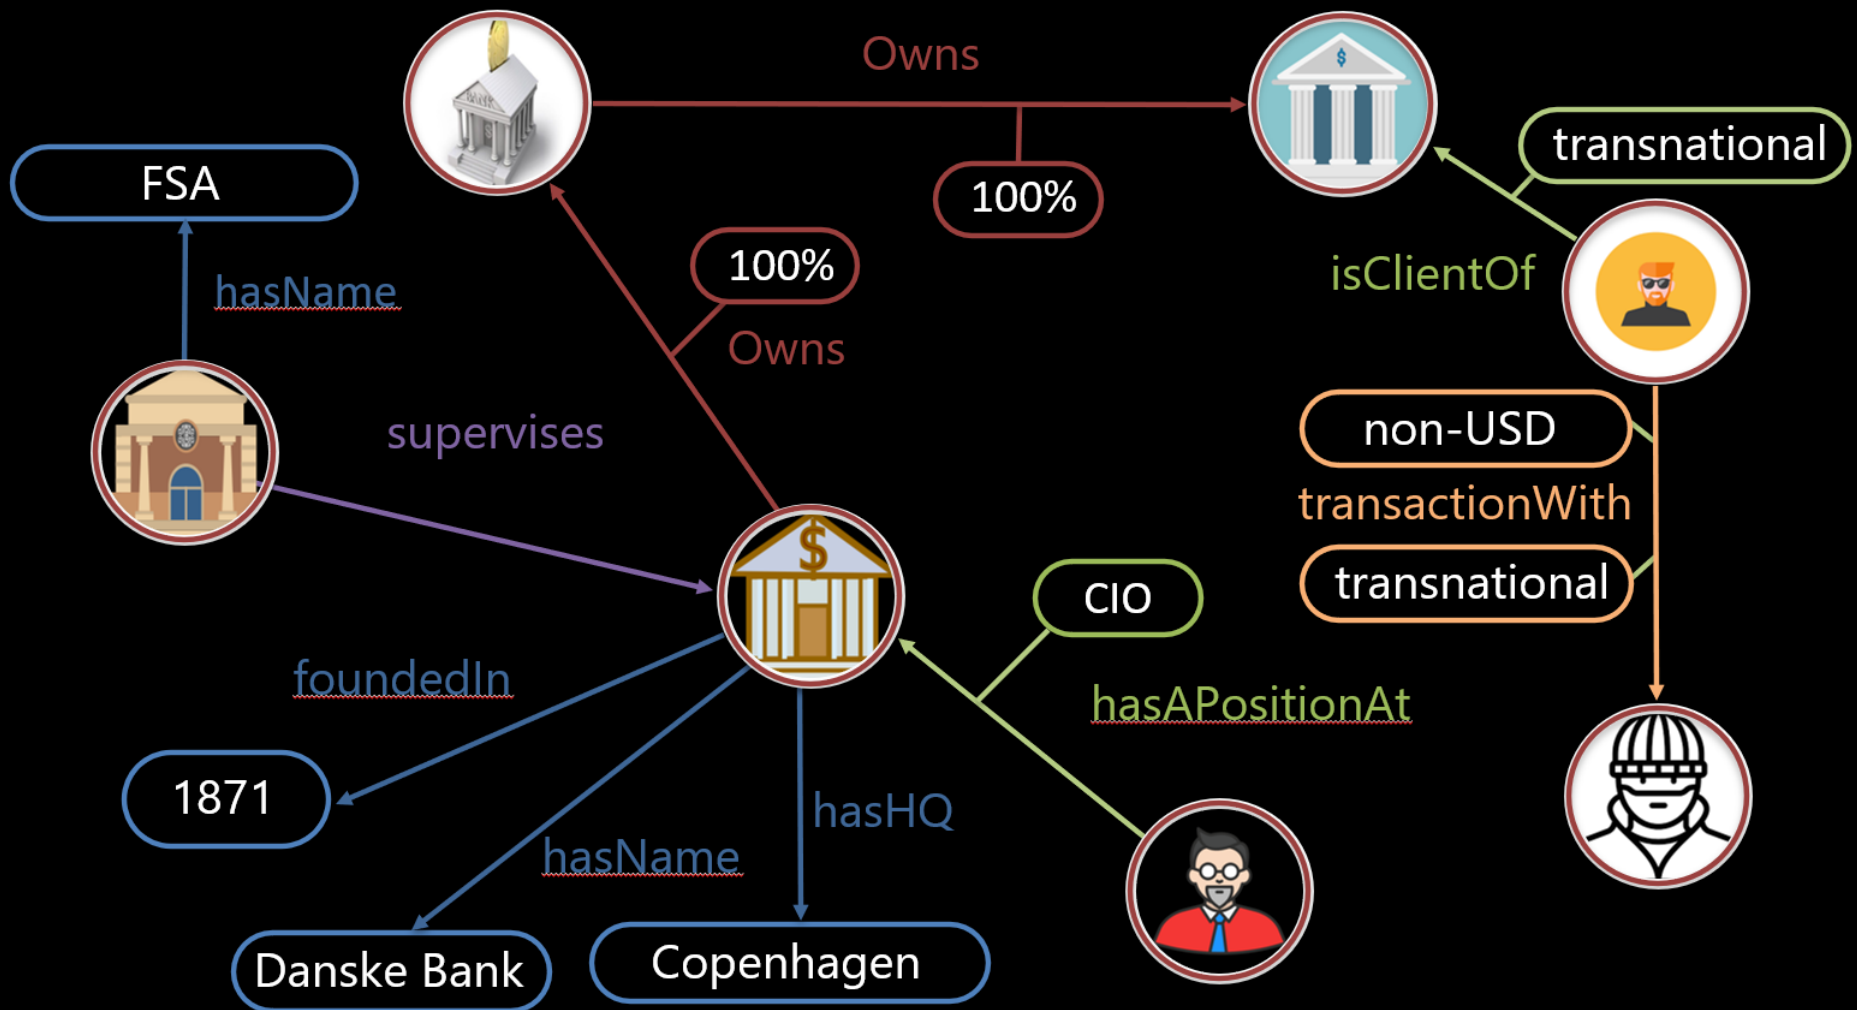
Danske Bank AS [+ Add to myFT](#)

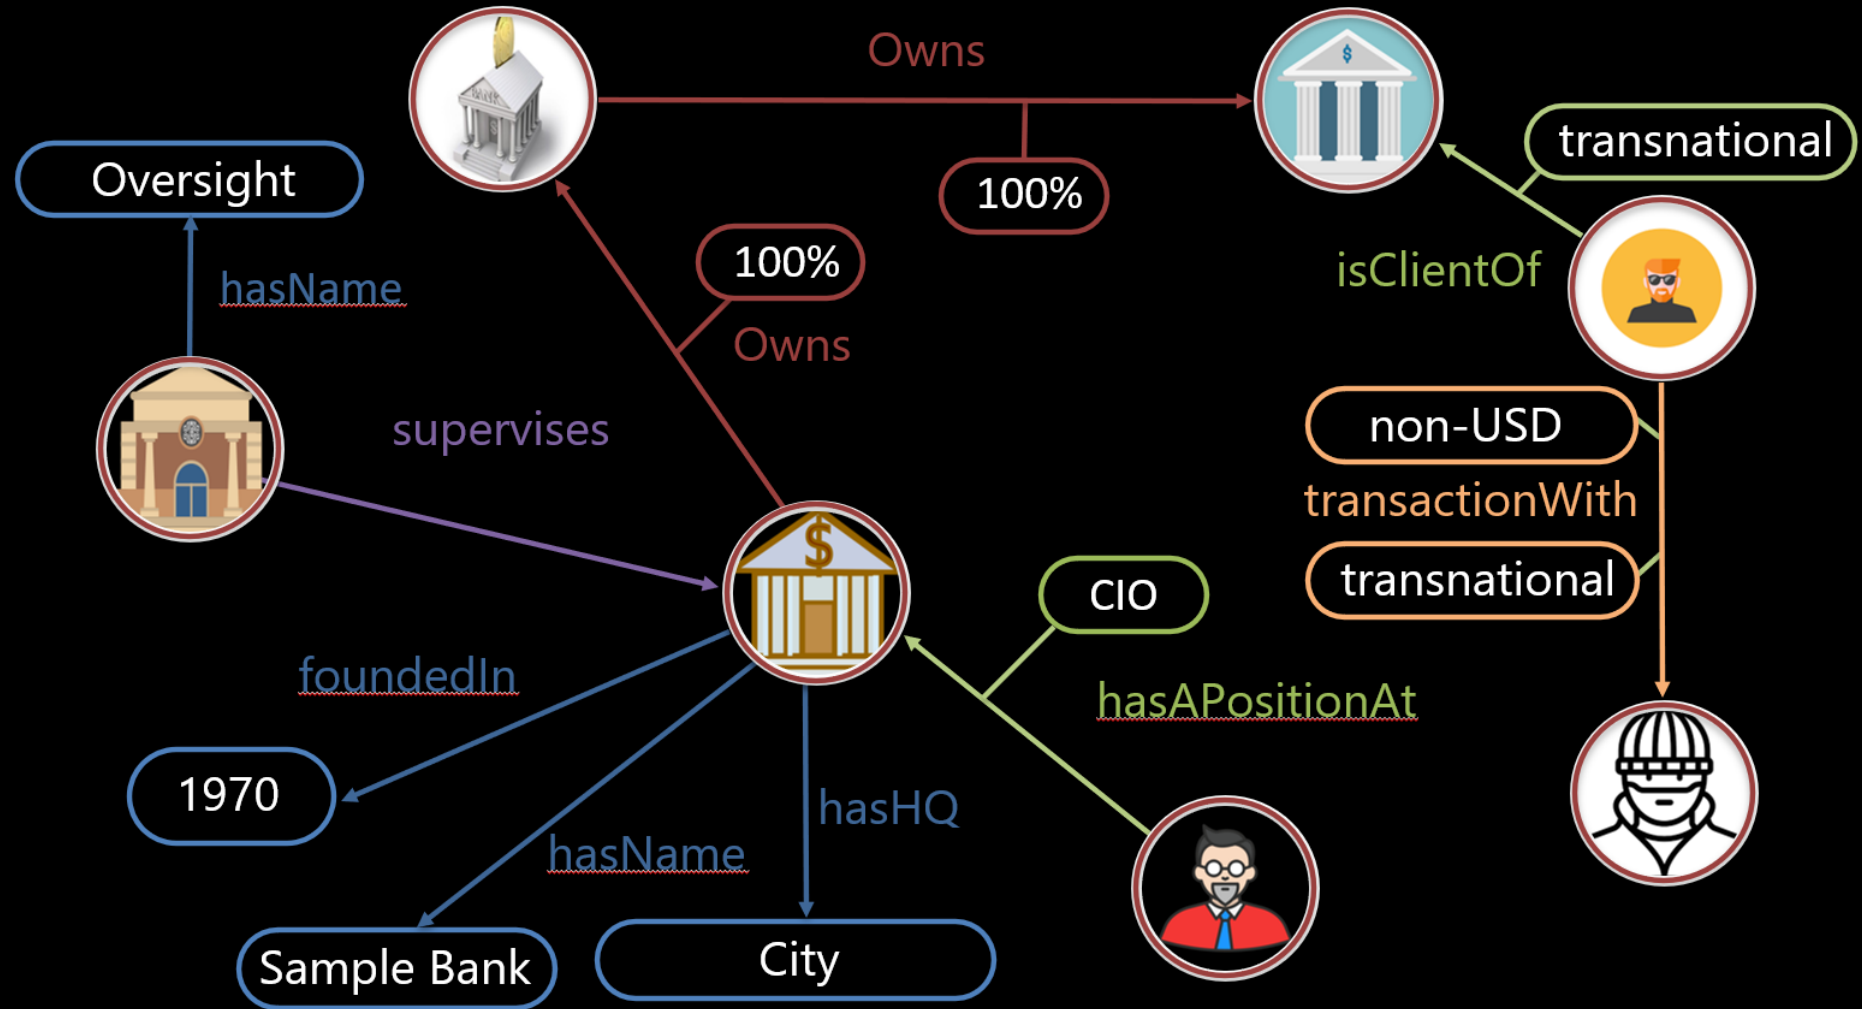
Danske Bank chairman ousted by main shareholder after scandal


Maersk family brings in new blood to stabilise lender in wake of €200bn money laundering


€200bn money laundering
















Data Models


Overview of data models
in different communities






Data Models

Overview of data models in different communities





KG Embeddings

Widely-applied, large family of ML models



Data Models

Overview of data models in different communities




Logical Knowledge in KGs

Highly expressive, diverse family of logical models.



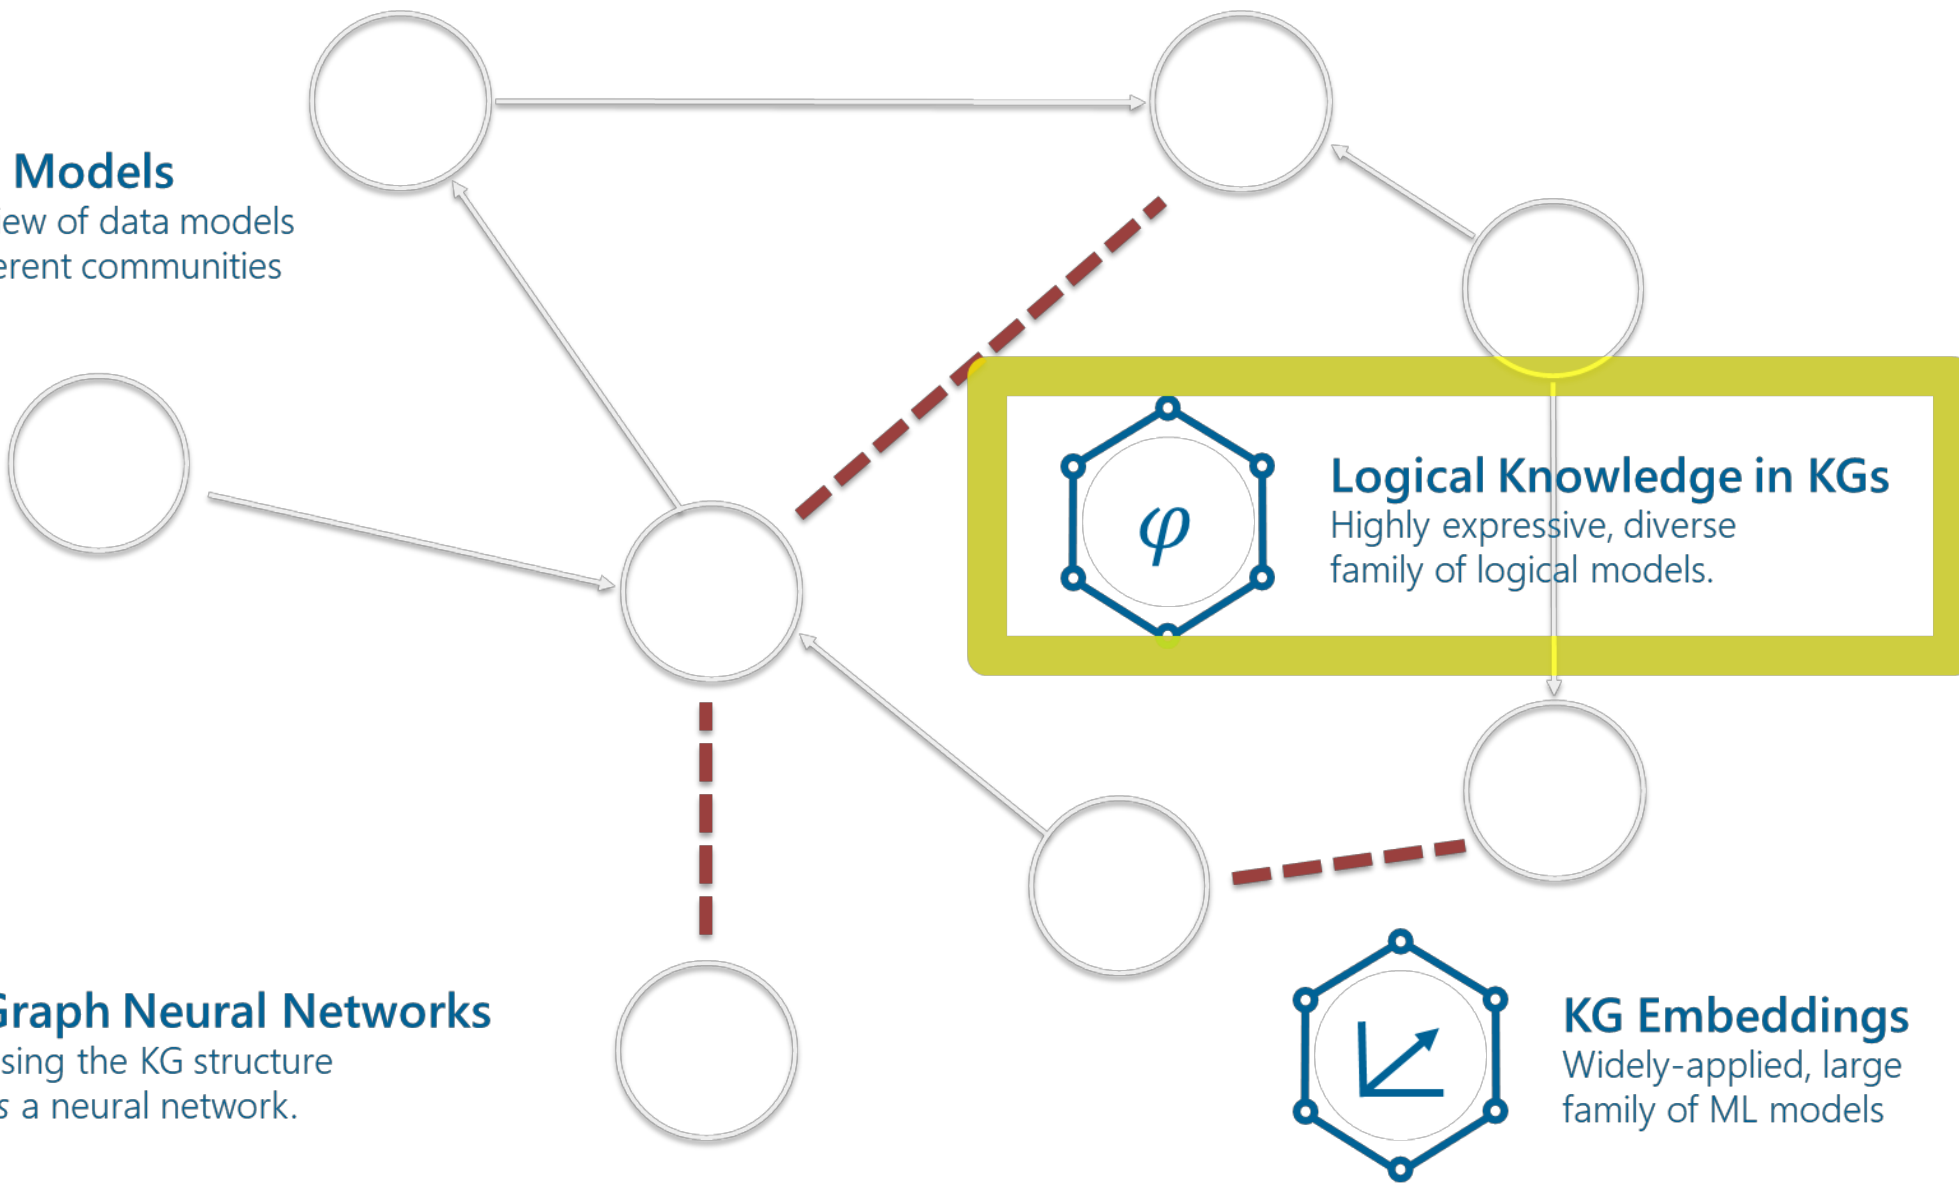
KG Embeddings

Widely-applied, large family of ML models



Data Models

Overview of data models in different communities




Graph Neural Networks

Using the KG structure as a neural network.



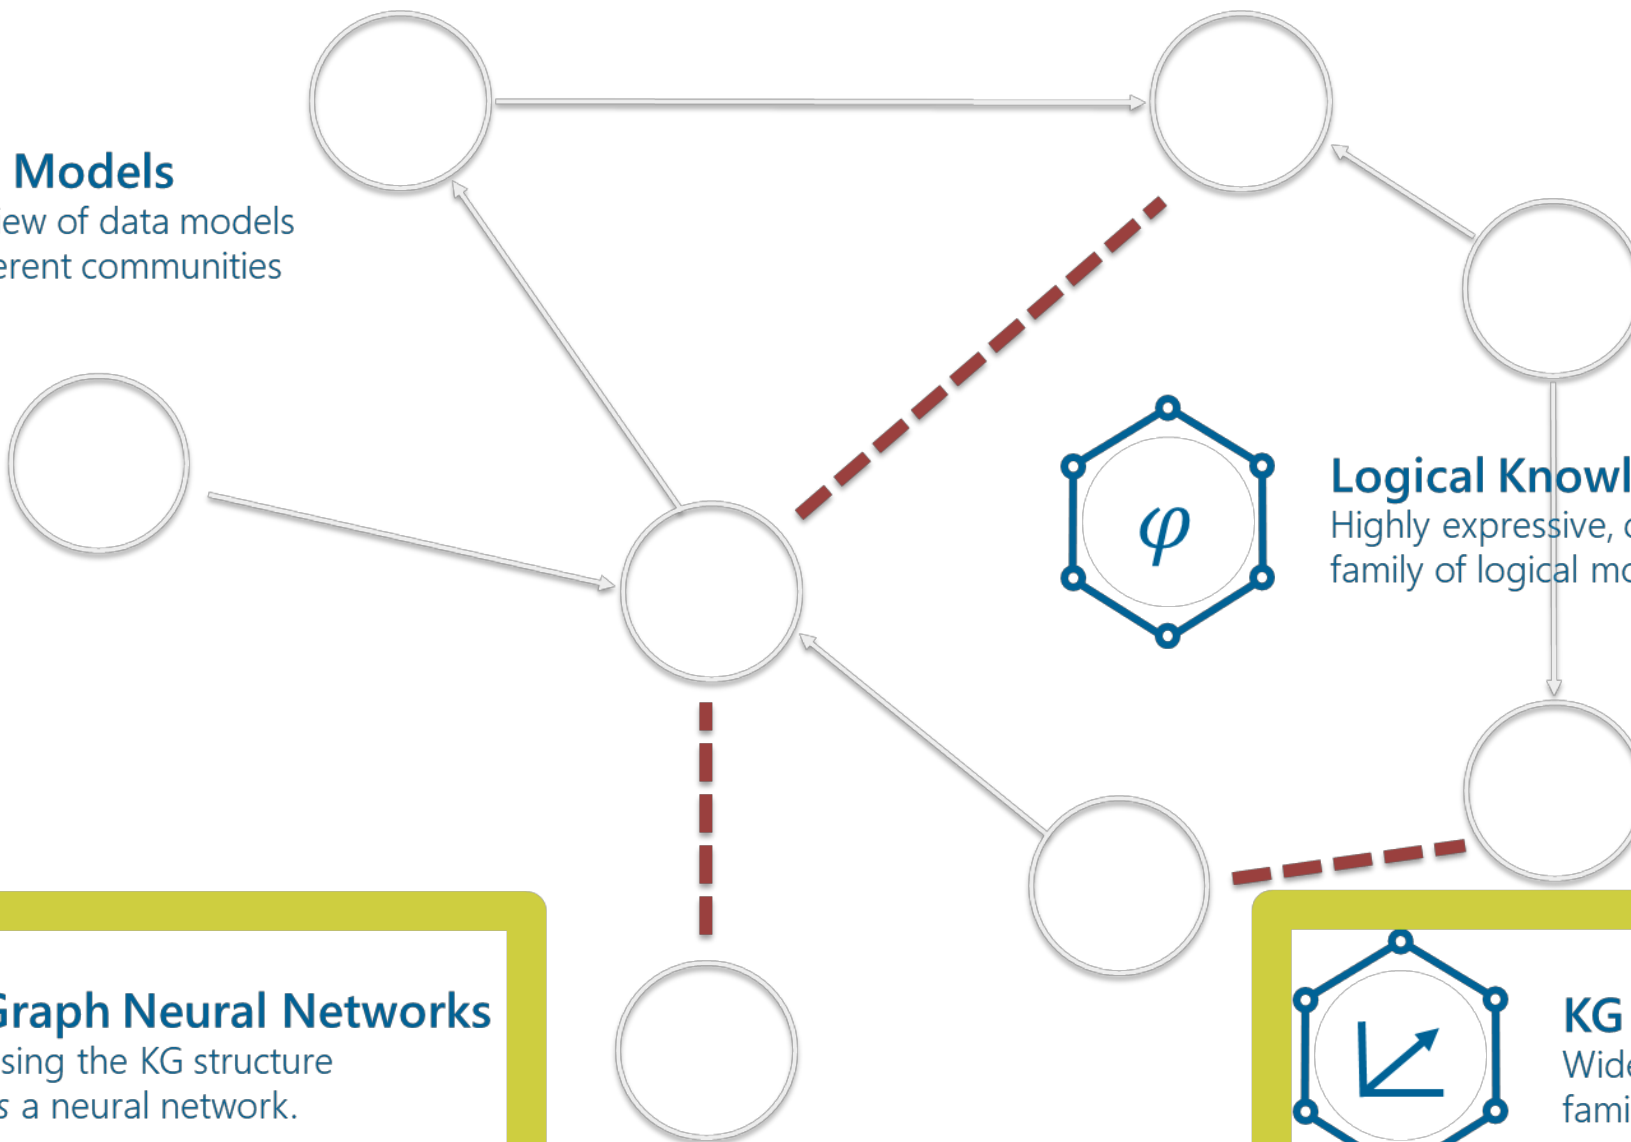
KG Embeddings

Widely-applied, large family of ML models




Data Models

Overview of data models in different communities




Logical Knowledge in KGs

Highly expressive, diverse family of logical models.




Graph Neural Networks

Using the KG structure as a neural network.



KG Embeddings

Widely-applied, large family of ML models



Data Models

Overview of data models in different communities


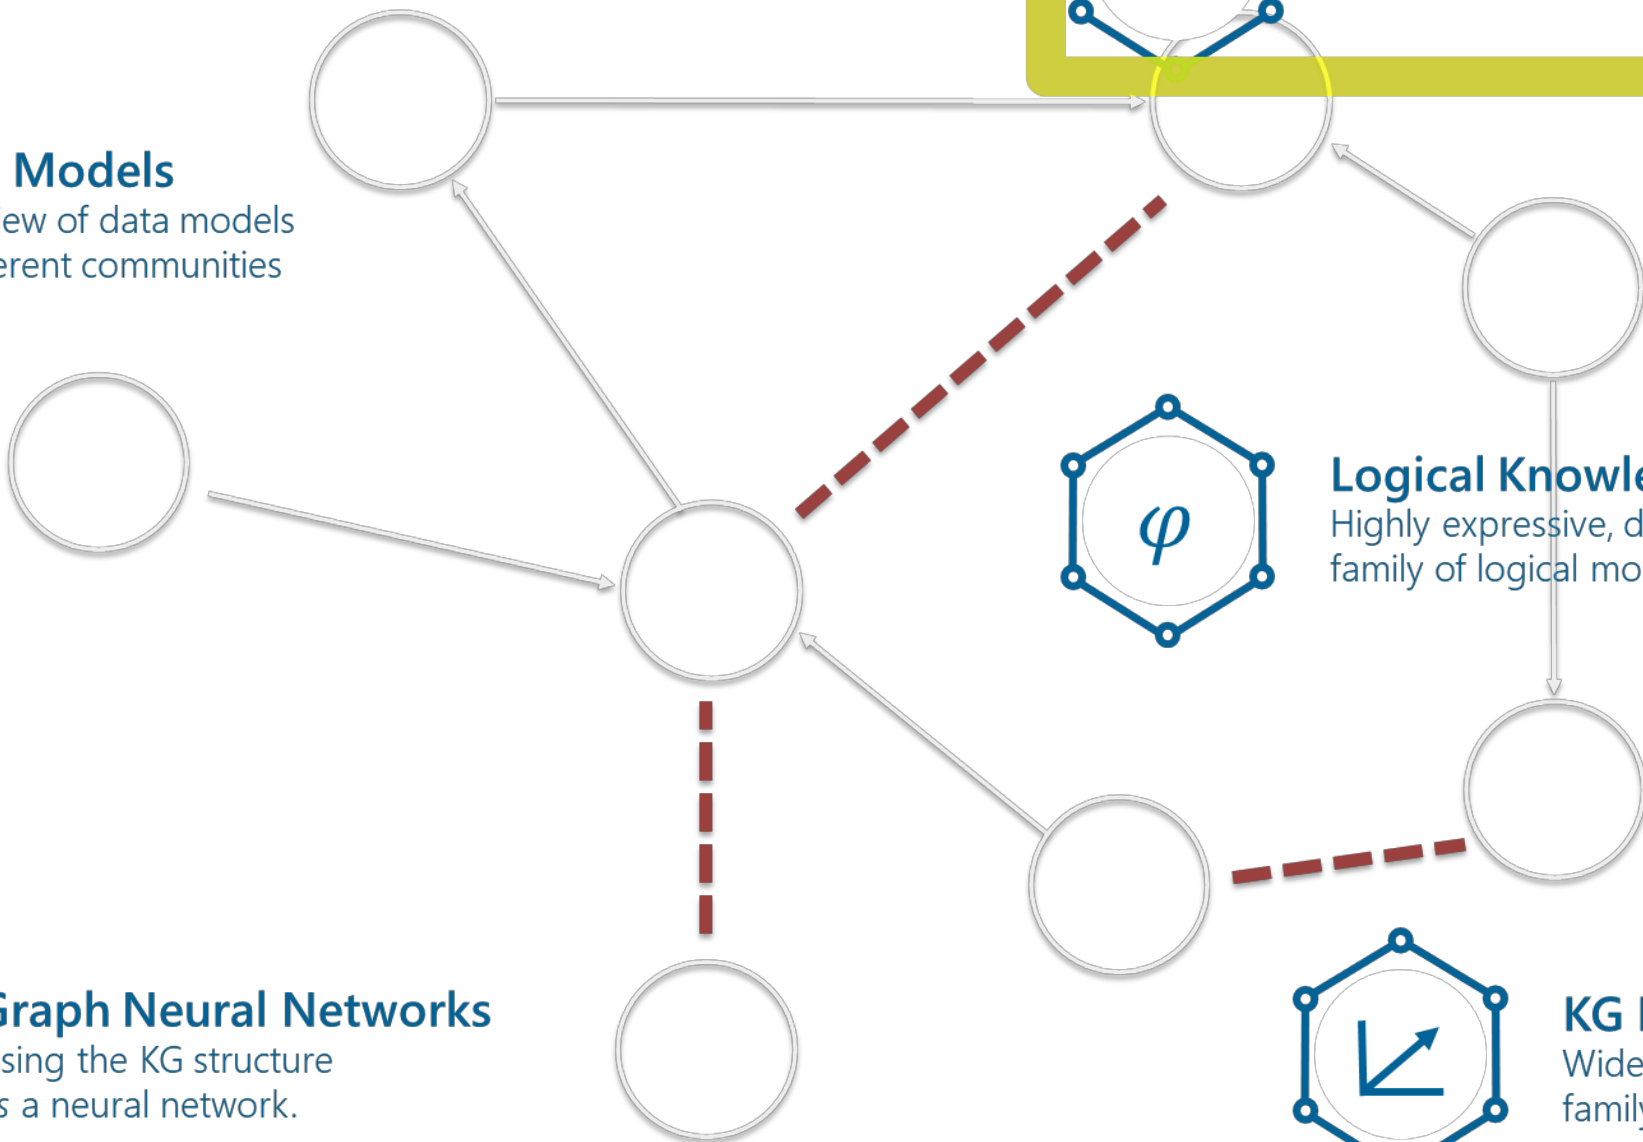


Graph Neural Networks

Using the KG structure as a neural network.


Large Language Models

Human-like text capabilities.



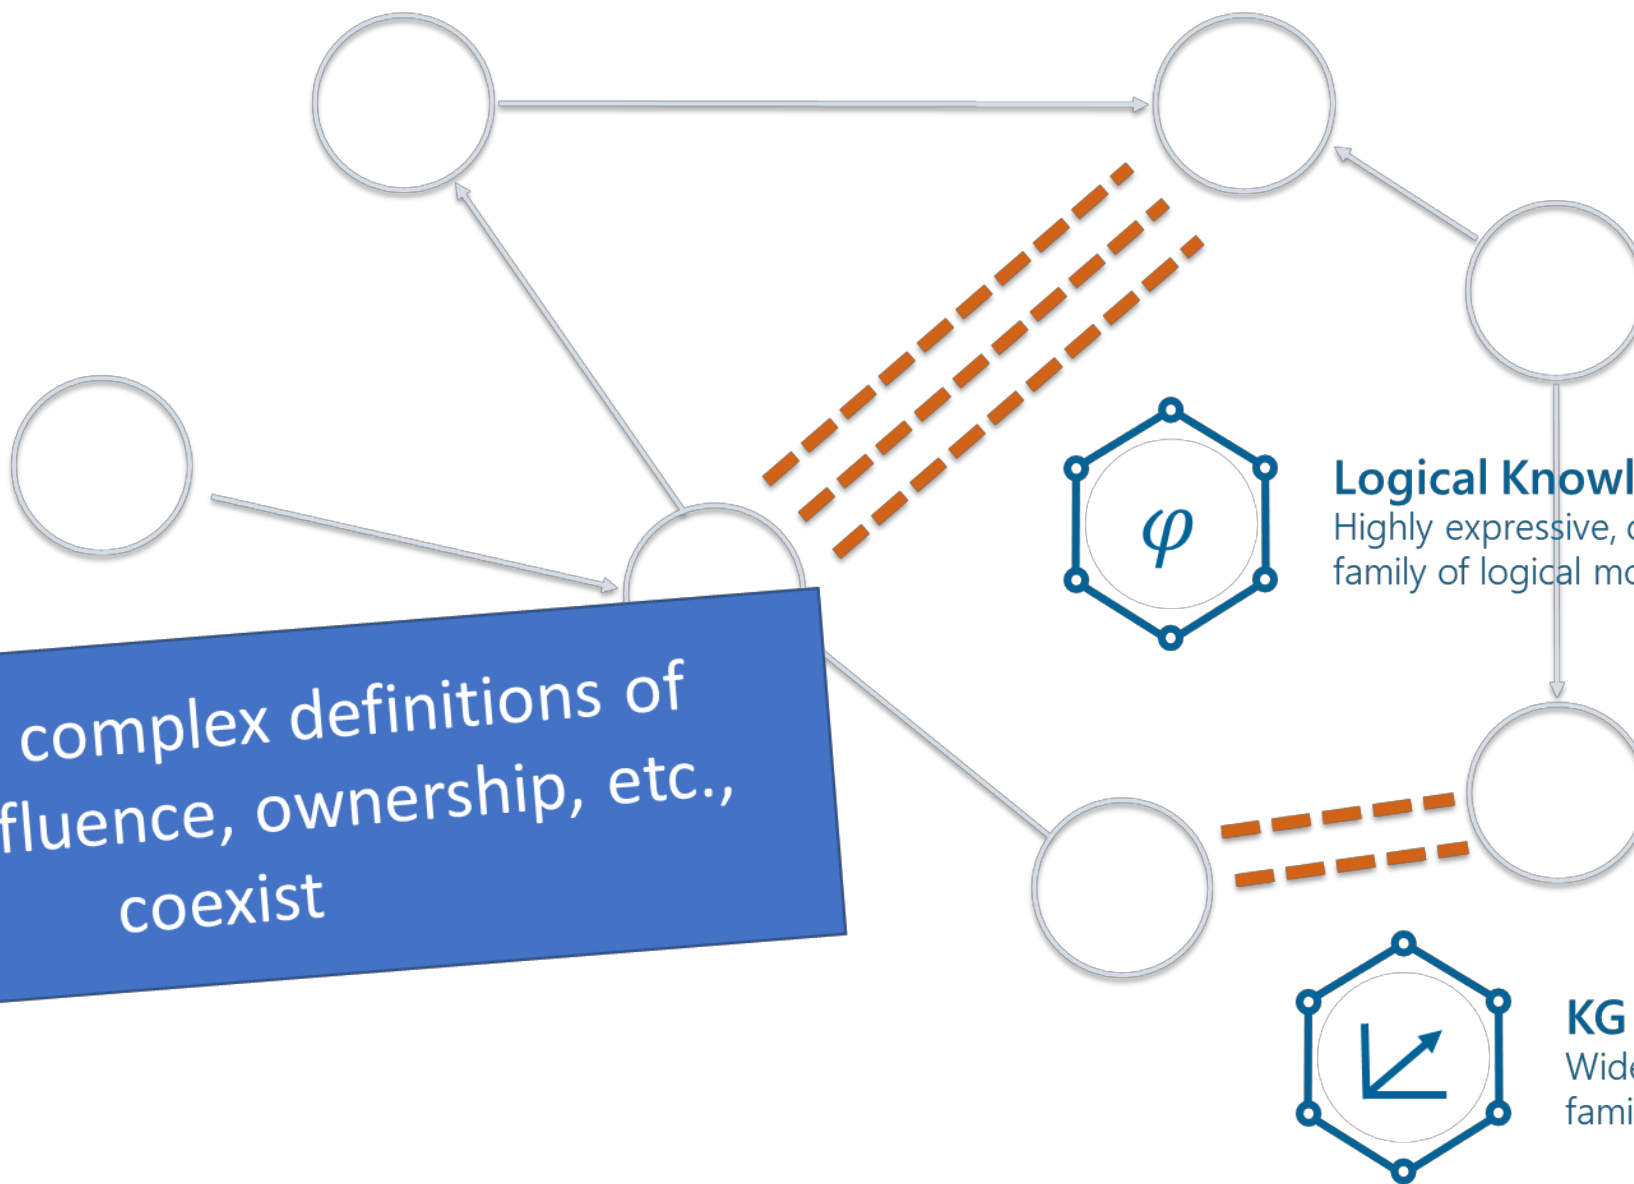
Logical Knowledge in KGs

Highly expressive, diverse family of logical models.



KG Embeddings

Widely-applied, large family of ML models



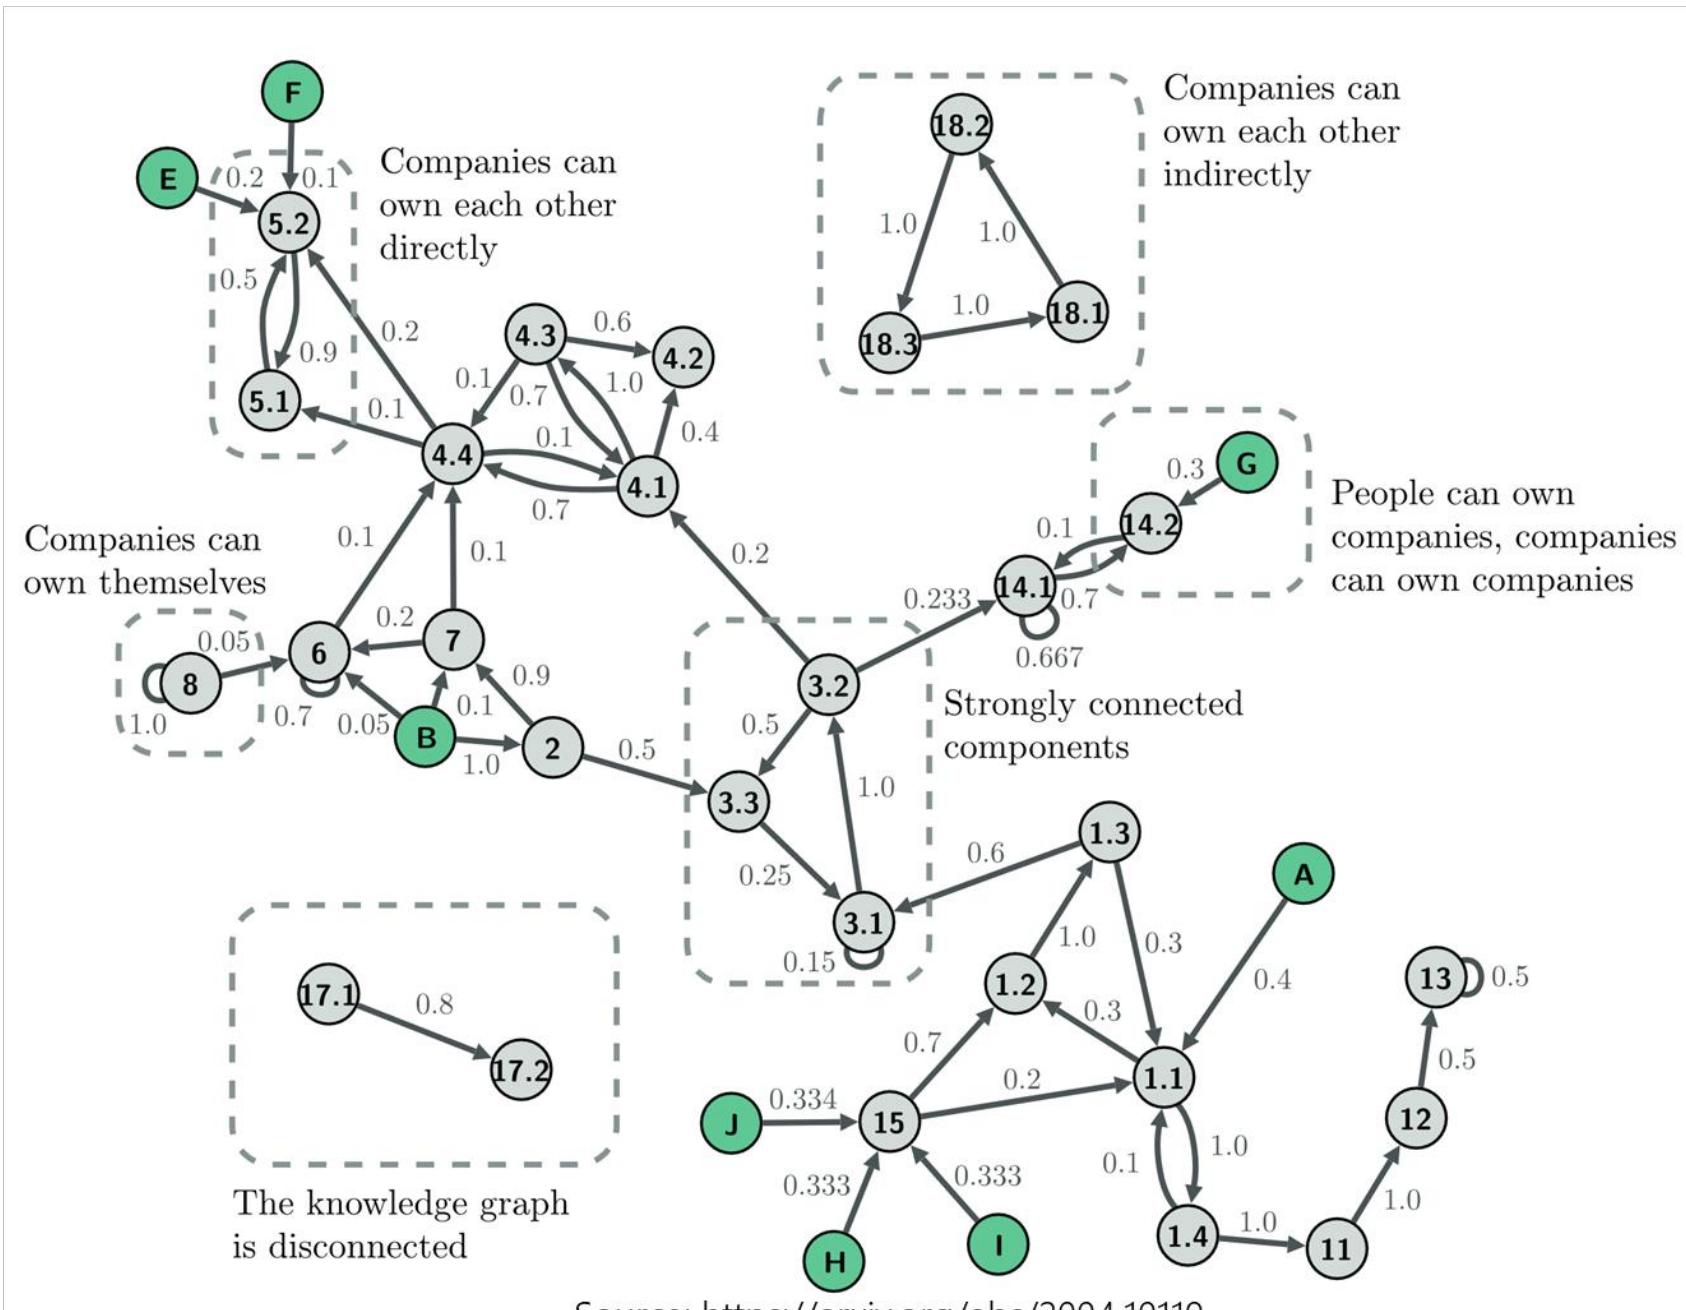
Multiple, complex definitions of control, influence, ownership, etc., coexist

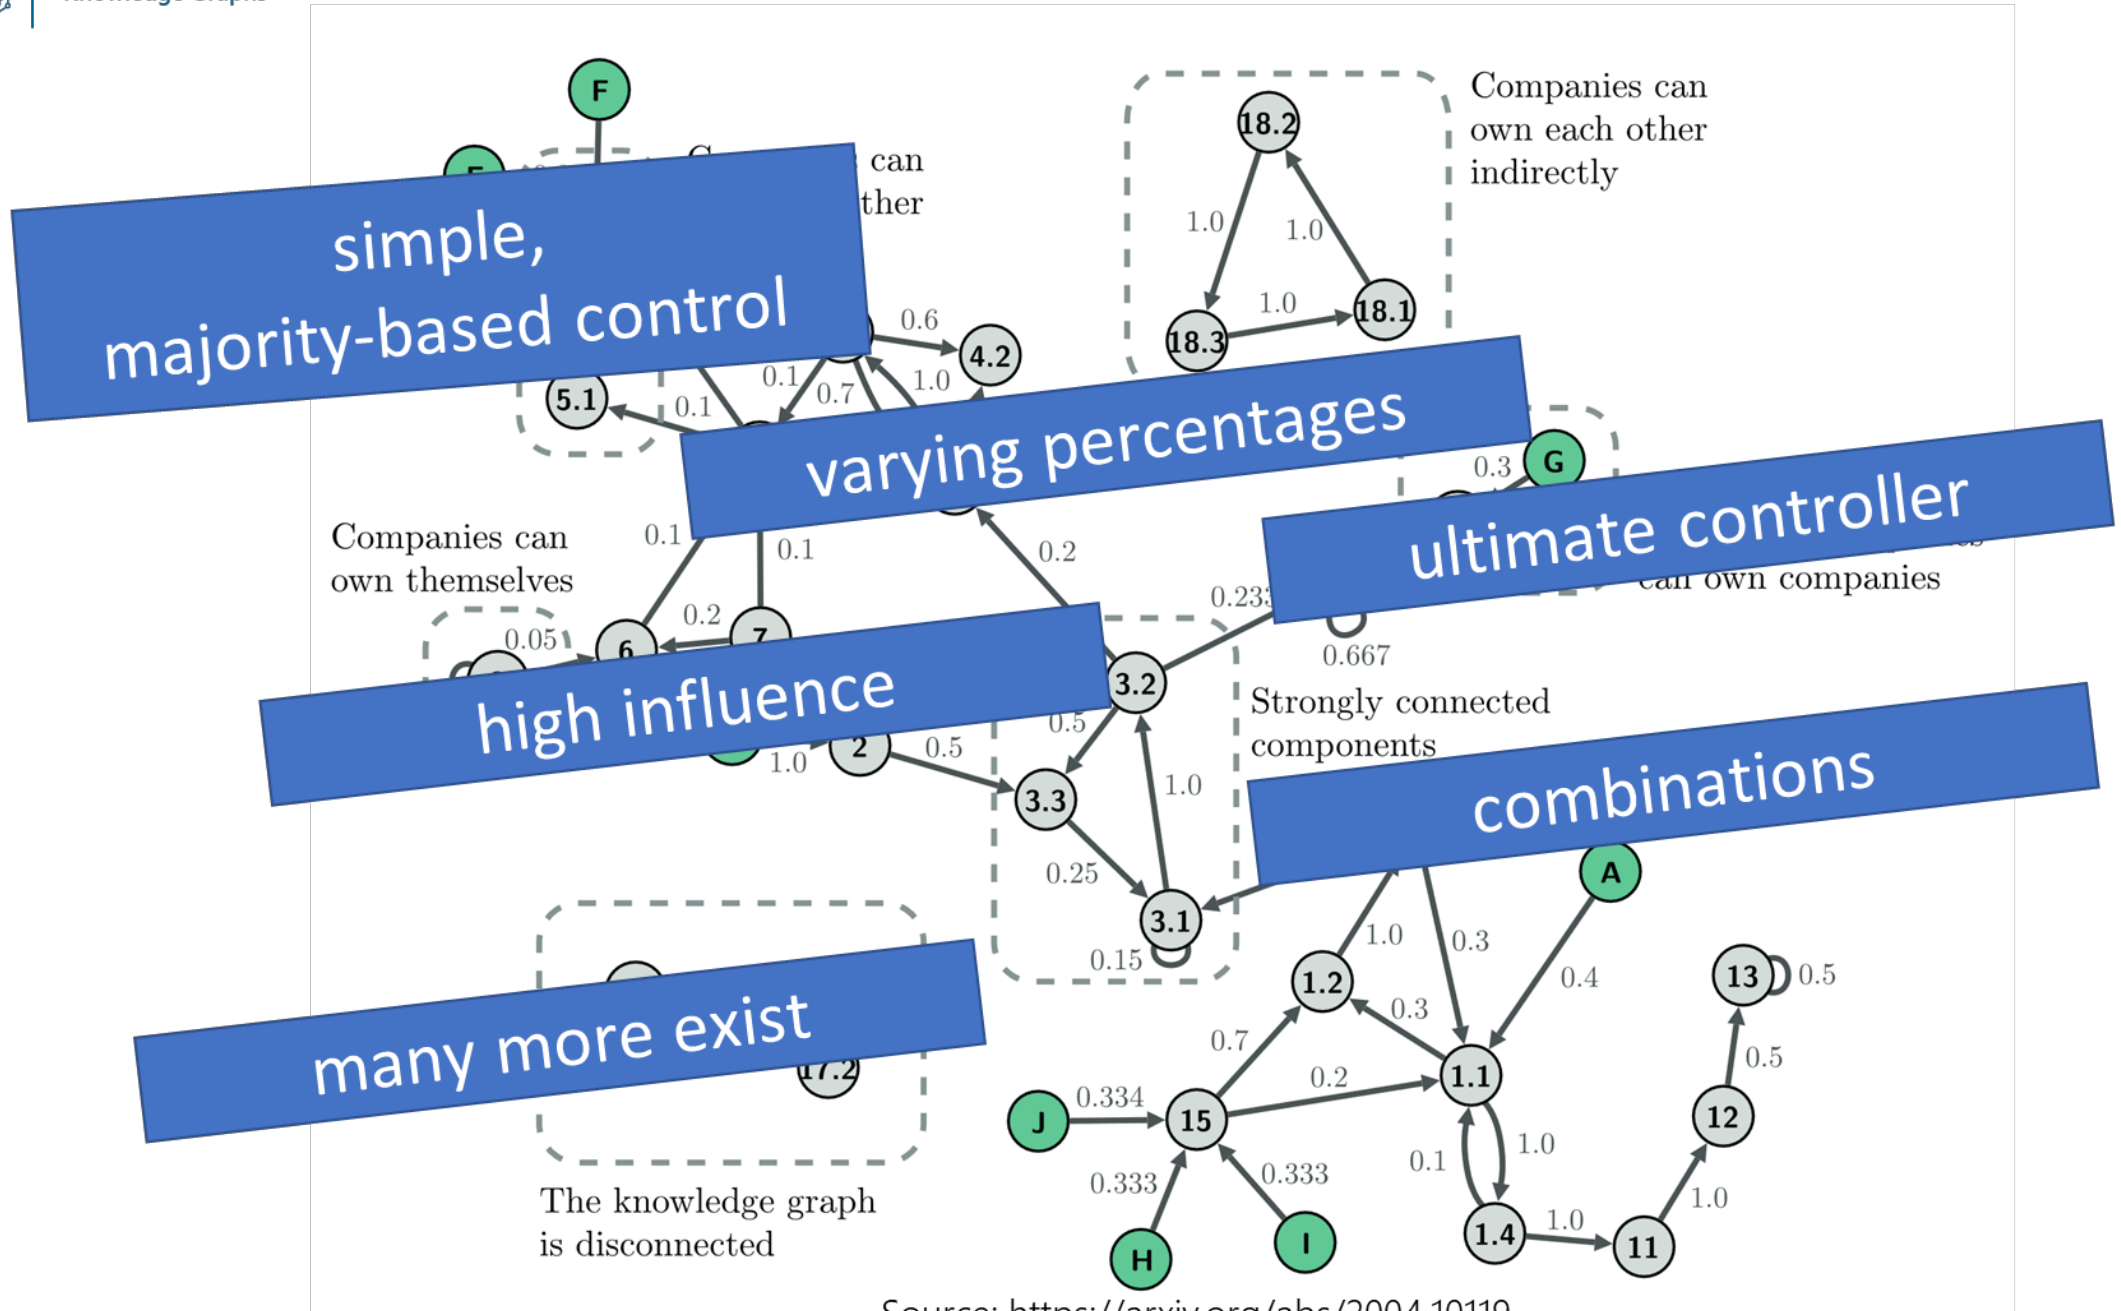


Logical Knowledge in KGs
Highly expressive, diverse family of logical models.



KG Embeddings
Widely-applied, large family of ML models







Knowledge Graphs

Research and Pushing the Boundaries

Teaser

Emanuel Sallinger

Confidential:
registered participants only, not for resharing



2016-2019 Selected Highlights



LANGUAGE

*Data Wrangling for Big Data:
Towards a Lingua Franca for Data
Wrangling*



LOGIC

*Swift Logic for Big Data and
Knowledge Graphs*



PROJECT

*The VADA Architecture for Cost-
Effective Data Wrangling*



DATA SCIENCE

*Data Science with Vadalog:
Bridging Machine Learning and
Reasoning*



SYSTEM

*The Vadalog System: Datalog-
based Reasoning for Knowledge
Graphs*



PROJECT ARCHITECTURE

*VADA: an architecture for end
user informed data preparation*



SPACE EFFICIENCY

*The Space-Efficient
Core of Vadalog*



ENTERPRISE AI

*Knowledge Graphs and Enterprise
AI: The Promise of an Enabling
Technology*



RECOMMENDER SYSTEMS

*Feature Engineering and
Explainability with Vadalog: A
Recommender Systems Application*






2020 Selected Highlights

EMERGENCY RESPONSE
(not peer-reviewed)

COVID-19: LOCKDOWN
COVID-19 and Company Knowledge Graphs: Assessing Golden Powers and Economic Impact of Selective Lockdown via AI Reasoning



ENTERPRISE AI IN PRACTICE
Weaving Enterprise Knowledge Graphs: The Case of Company Ownership Graphs



Declarative AI 2020


COVID-19: TAKEOVERS
Reasoning on Company Takeovers during the COVID-19 Crisis with Knowledge Graphs

Declarative AI 2020

PROBABILISTIC
Reasoning Under Uncertainty in Knowledge Graphs

Declarative AI 2020

MONEY LAUNDERING
Rule-based Anti-Money Laundering in Financial Intelligence Units: Experience and Vision



KNOWLEDGE GRAPHS
Knowledge Graphs: The Layered Perspective



BOOK
Knowledge Graphs and Big Data Processing



KG EMBEDDINGS
Reasoning in Knowledge Graphs: An Embeddings Spotlight





2021 Selected Highlights




CONFIDENTIALITY

Financial Data Exchange with Statistical Confidentiality: A Reasoning-based Approach



COMPANY CONTROL

Distributed Company Control in Company Shareholding Graphs



TEMPORAL

Monotonic Aggregation for Temporal Datalog



HARMFUL JOINS

Eliminating Harmful Joins in Warded Datalog+/-.



INDUSTRIAL BLOCKCHAIN

Rule-based Blockchain Knowledge Graphs: Declarative AI for Solving Industrial Blockchain Challenges



JOINS

Traversing Knowledge Graphs with Good Old (and New) Joins



HYBRID AI

Augmenting Logic-based Knowledge Graphs: The Case of Company Graphs



BLOCKCHAIN VISION

Blockchains as Knowledge Graphs - Blockchains for Knowledge Graphs





2022 Selected Highlights



DATA SCIENCE

Data science with Vadalog: Knowledge Graphs with machine learning and reasoning in practice



AI ARCHITECTURE

Vadalog: A modern architecture for automated reasoning with large knowledge graphs



SHY AND WARDED

On the Relationship between Shy and Warded Datalog+/-



CS AND ECONOMY

Reasoning on company takeovers: From tactic to strategy



KG DESIGN

Model-Independent Design of Knowledge Graphs - Lessons Learnt From Complex Financial Graphs



RULE LEARNING

Rule Learning over Knowledge Graphs with Genetic Logic Programming



38th IEEE International Conference on Data Engineering

TEMPORAL

iTemporal: An Extensible Generator of Temporal Benchmarks



EGDs

Exploiting the Power of Equality-generating Dependencies in Ontological Reasoning





DATA SCIENCE

*Data Science with Vadalog:
Bridging Machine Learning and
Reasoning*



HYBRID AI

*Augmenting Logic-based
Knowledge Graphs: The Case of
Company Graphs*



38th IEEE
International
Conference on
Data Engineering


RULE LEARNING

*Rule Learning over Knowledge
Graphs with
Genetic Logic Programming*




PATTERN AWARENESS

*Pattern-Aware and Noise-
Resilient Embedding Models*



Logical
Knowledge



Knowledge Graph
Embeddings



Graph Neural
Networks



Representations

Logical and ML-based representations for KGs.



KG Embeddings
Widely-applied, large family of ML models.



Logical Knowledge in KGs
Highly expressive, diverse family of logical models.



Graph Neural Networks
Using the KG structure as a neural network.




Data Models
Overview of data models in different communities.



Architectures
The big picture of building IT architectures for KGs.



Scalable Reasoning
Making use of the knowledge in the KG.



KG Creation
How to create a KG from heterogeneous data?



KG Evolution
How to update, correct and complete a KG?

Systems

Systems to bring KGs into practice.

Applications

Real-world applications of KGs.



Real-World Applications
Overview of diverse applications.



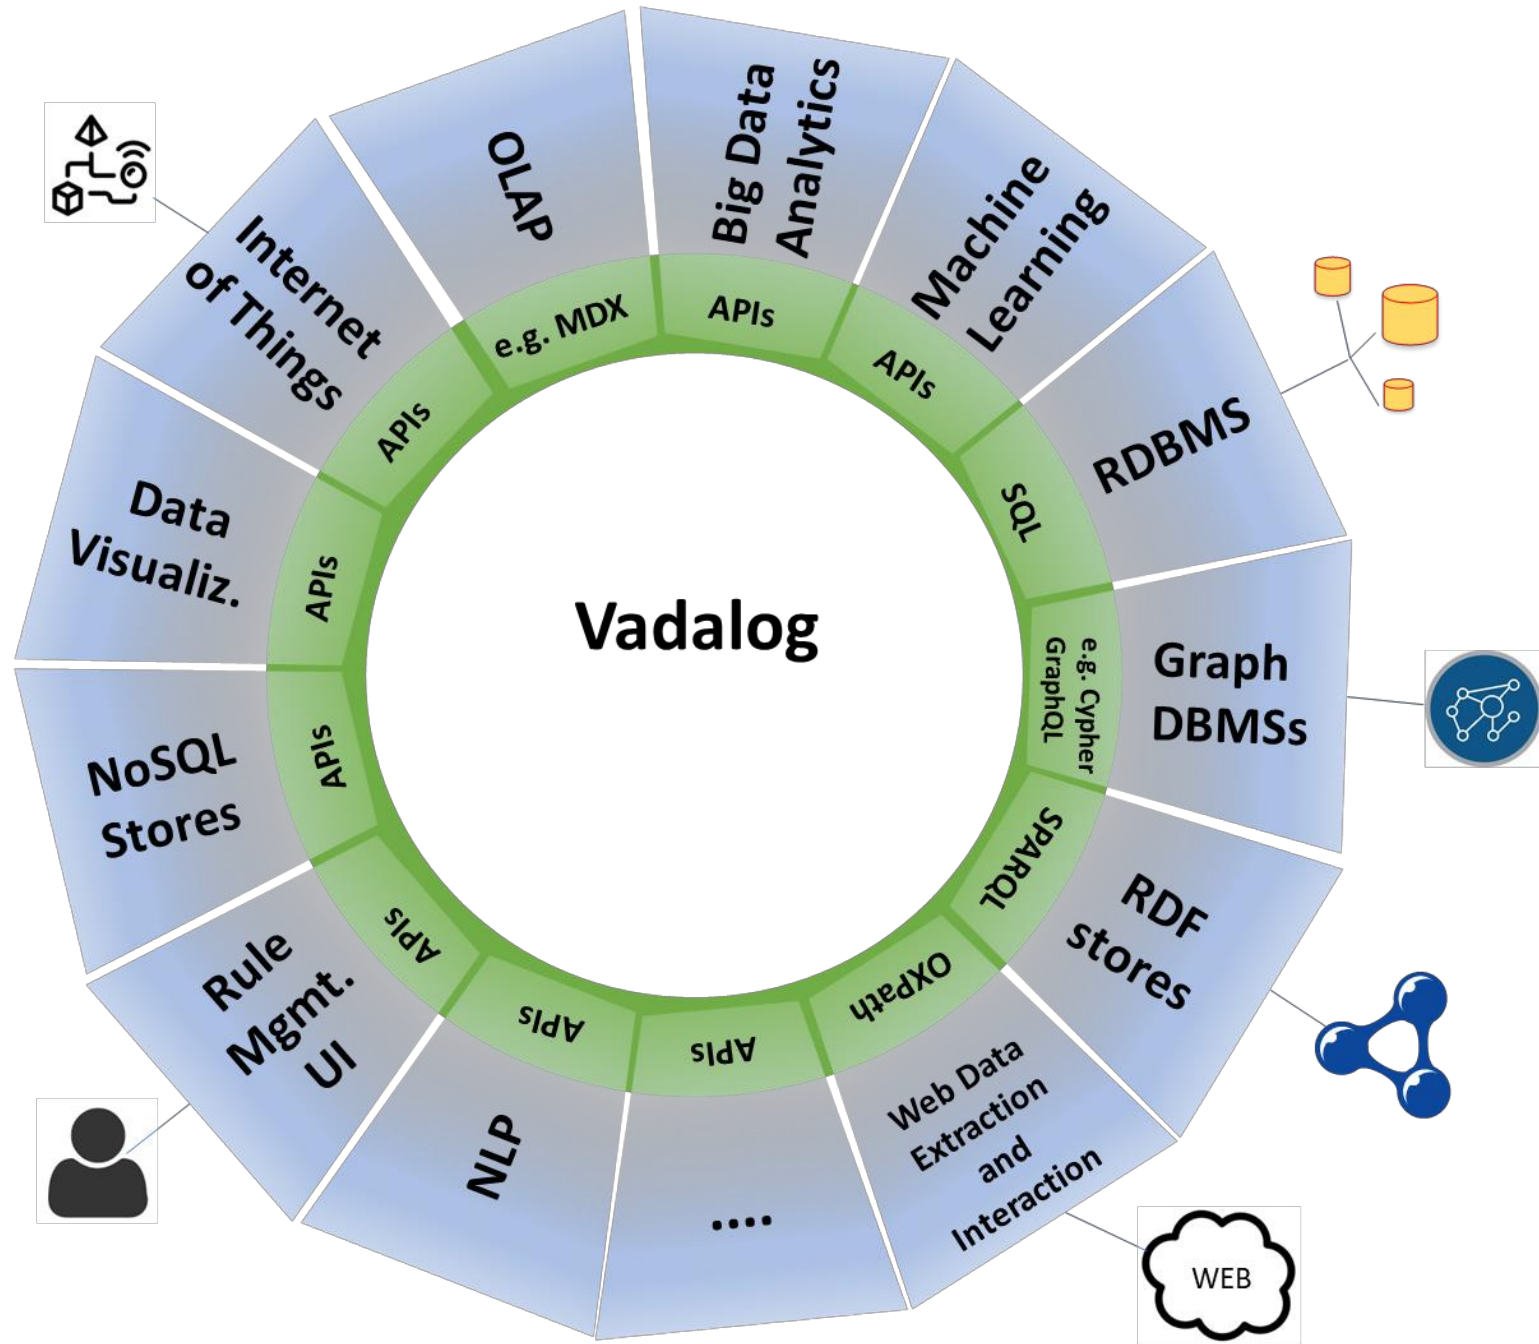
Financial KGs
Concrete applications in finance and economics.



Services
Which service to provide based on KGs?



Connections
... between KGs, AI, ML and Data Science.





Neurosymbolic Reasoning

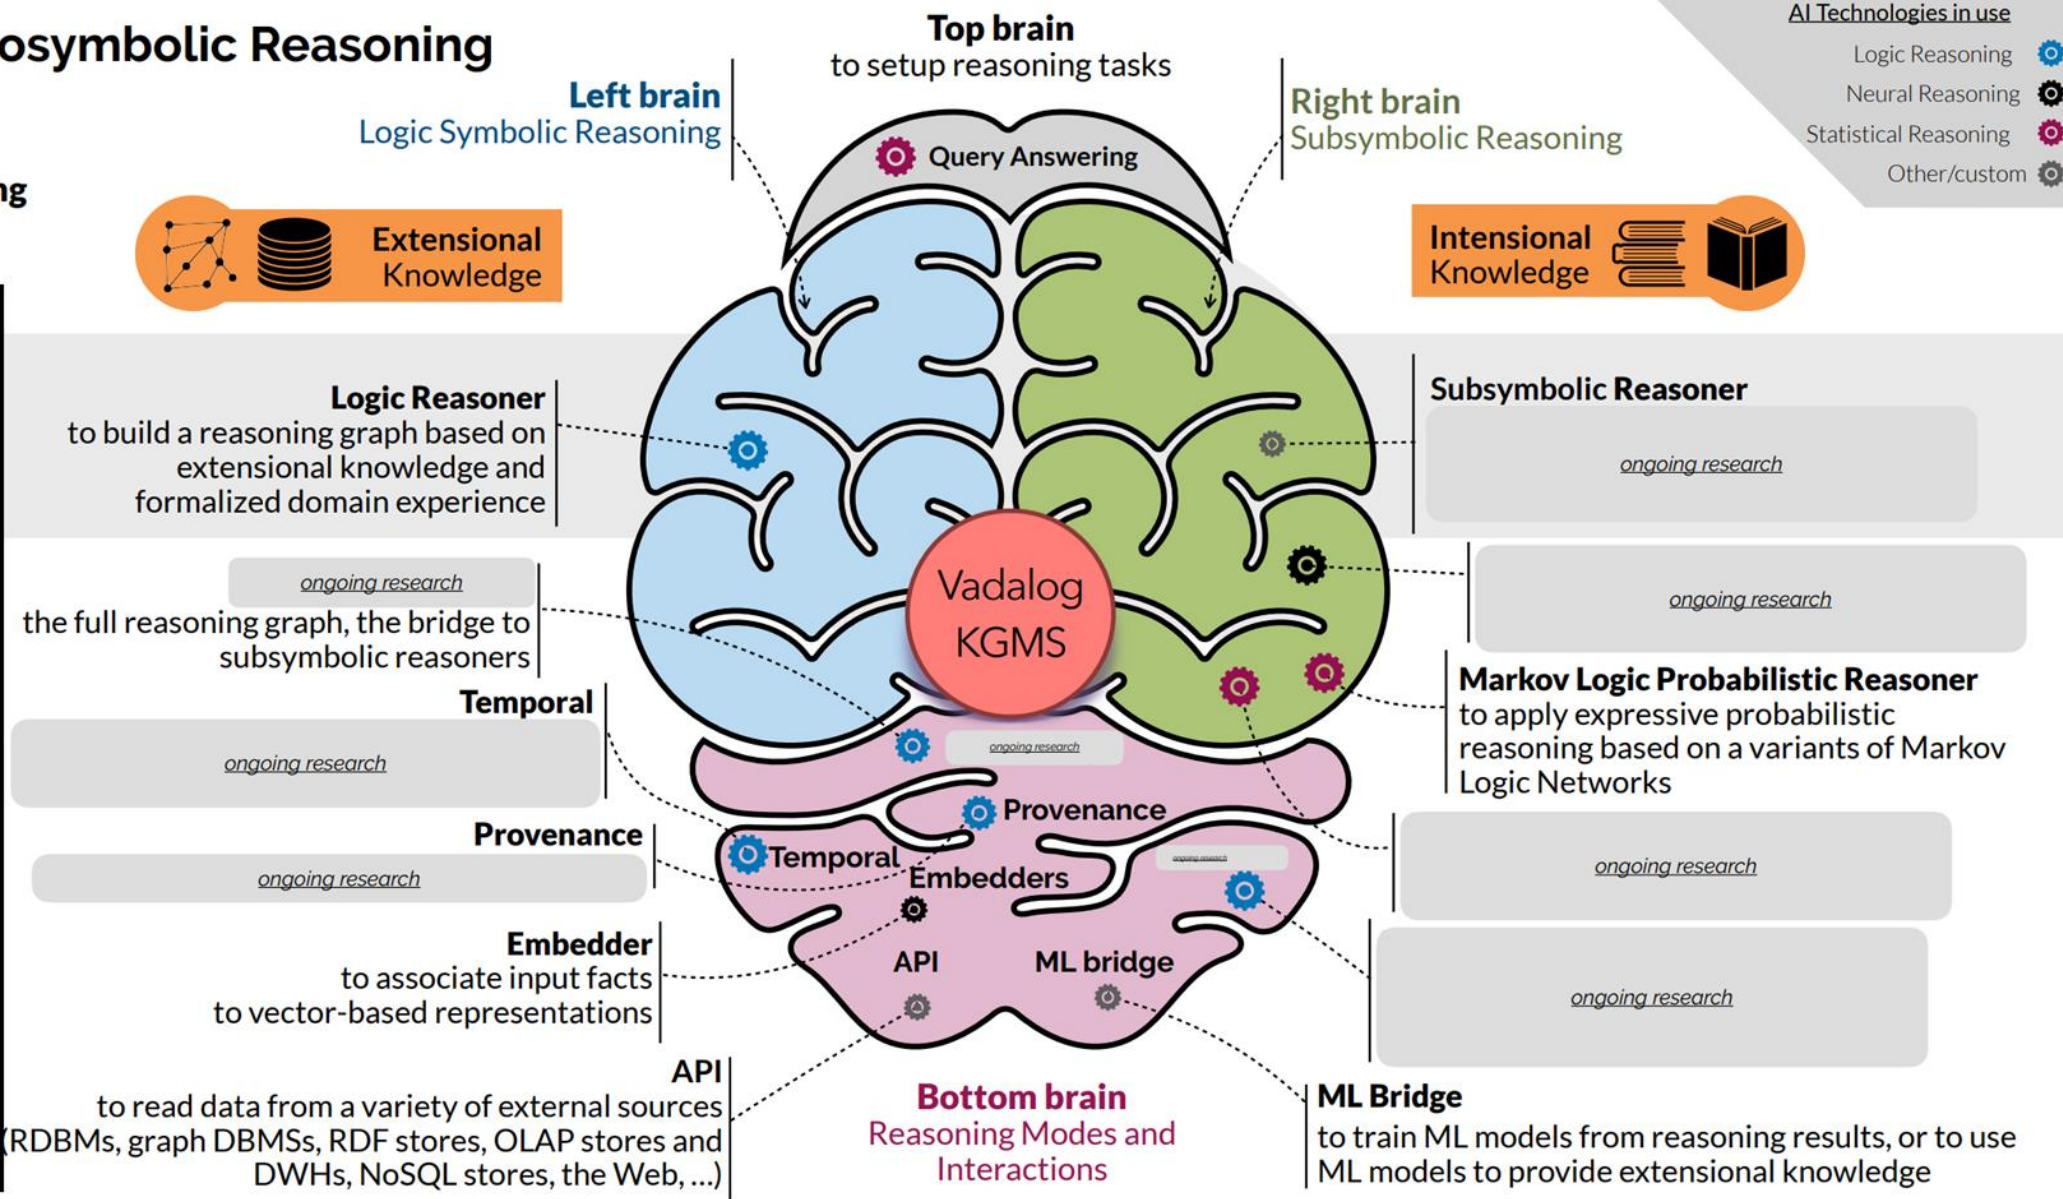
Reasoning Tasks

Decision Support
AI-driven Applications
ML-able tasks

Extensional Knowledge

Intensional Knowledge

- AI Technologies in use
- Logic Reasoning
 - Neural Reasoning
 - Statistical Reasoning
 - Other/custom





Knowledge Graphs
Real-World Applications – Diving Deeper

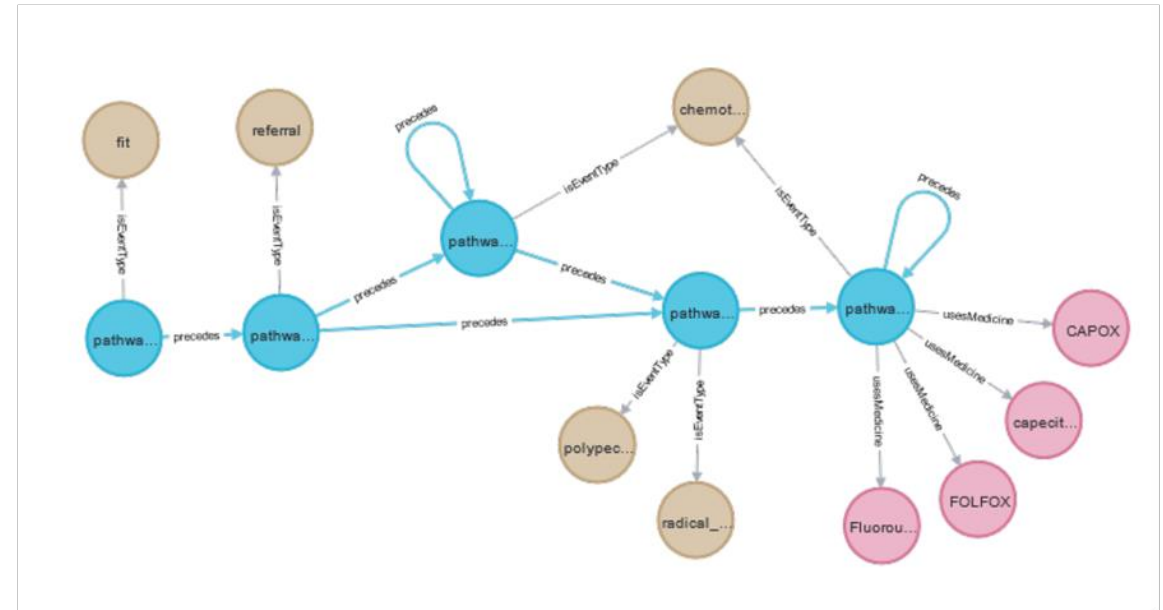
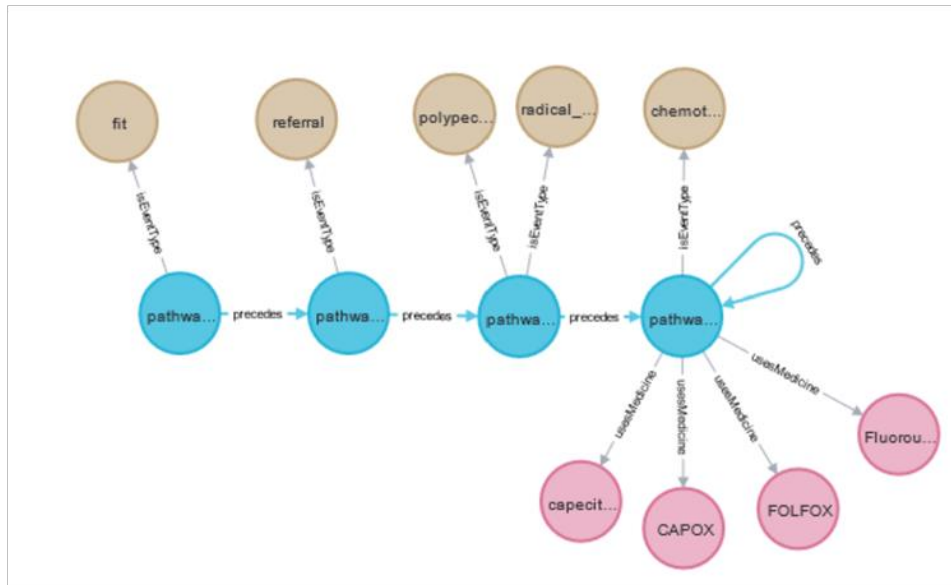
Teaser

Emanuel Sallinger

Confidential:
registered participants only, not for resharing

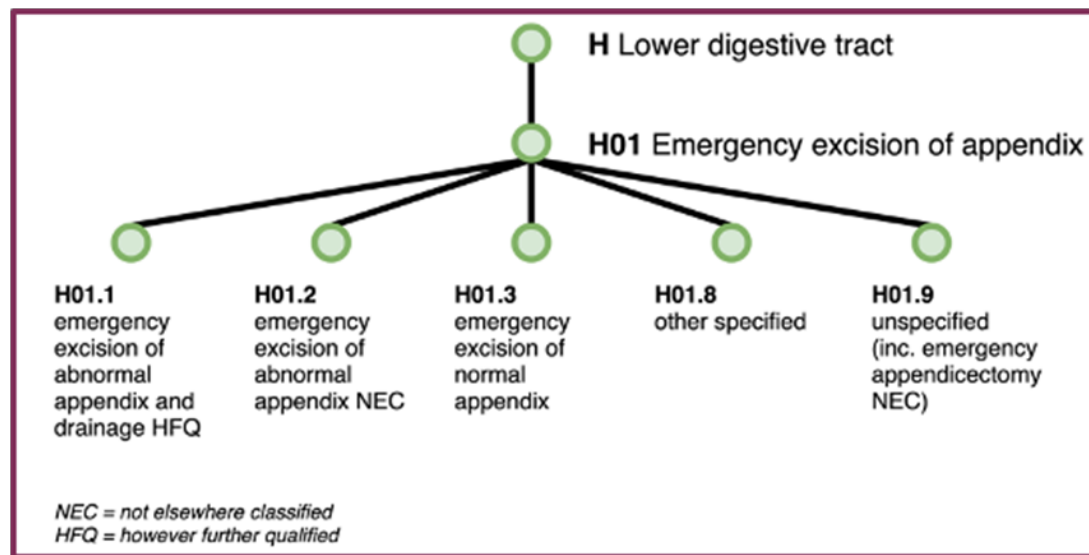



Cancer Pathways



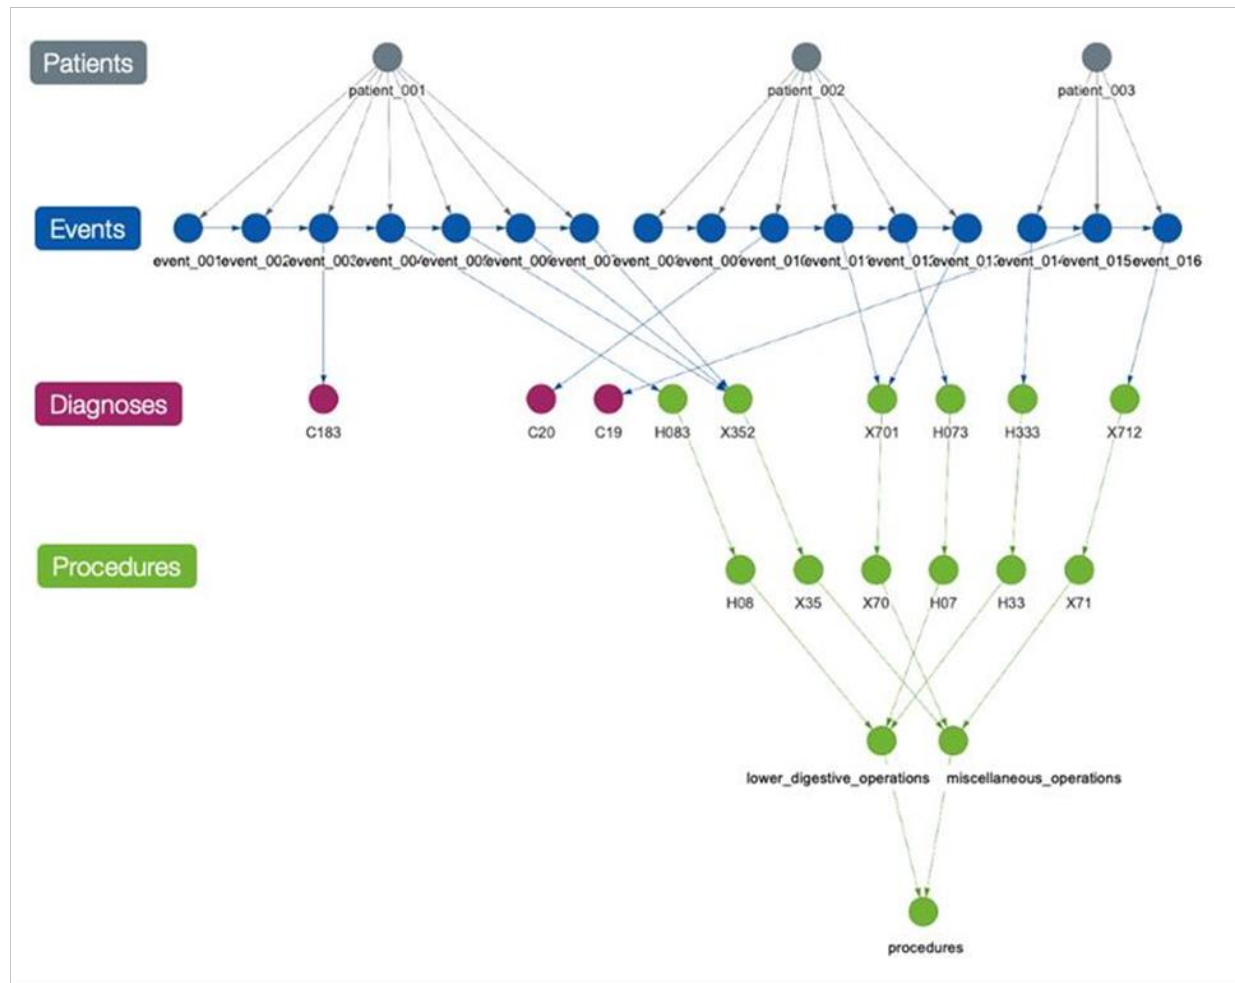
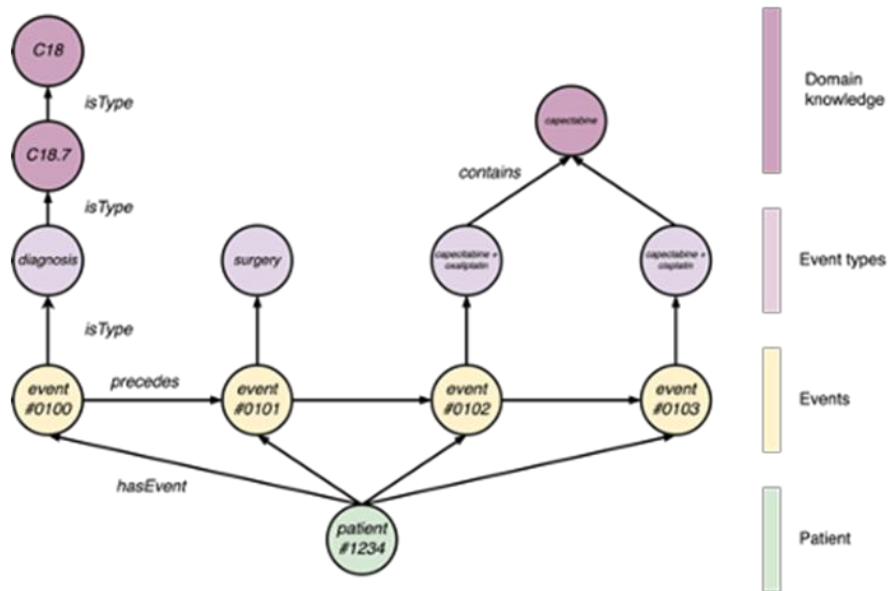


Cancer Pathways






Cancer Pathways







Supply Chains







Supply Chains



Historic Data,
Machine Learning




Supply Chains



→ **Rules:** Express what you want, not how you want it done.

```
path(X, Y, Time) :-
  start(X, T0),
  section(X, Y),
  Time = T0 + stat:mean(X, Y)
```

```
path(X, Z, Time) :-
  path(X, Y, T1),
  section(Y, Z),
  Time = T1 + stat:mean(Y, Z).
```






Temporal: DatalogMTL

Datalog extended with operators from the Metric Temporal Logic

Temporal Operators



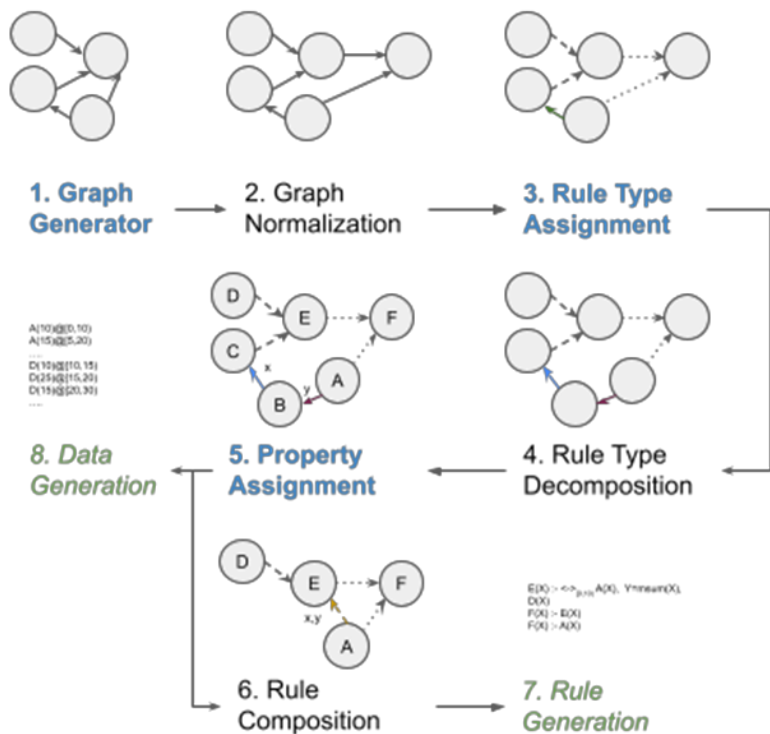


iTemporal Temporal Benchmark Suite



TEMPORAL

*iTemporal: An Extensible
Generator of Temporal
Benchmarks*



iTemporal - Temporal Benchmark Generator

Settings: 9 Store Load

Settings: 1 3 4 5 6
G R A P D

Rule Assignment
Core (S)/Temporal (S)/Aggr: 0 1
Core (M)/Temporal (M): 0 1
Int./Union: 0 1
+>[-]!<>[+]: 0 1
Since/Until: 0 1
ITA/MMWTA/STA: 0 1

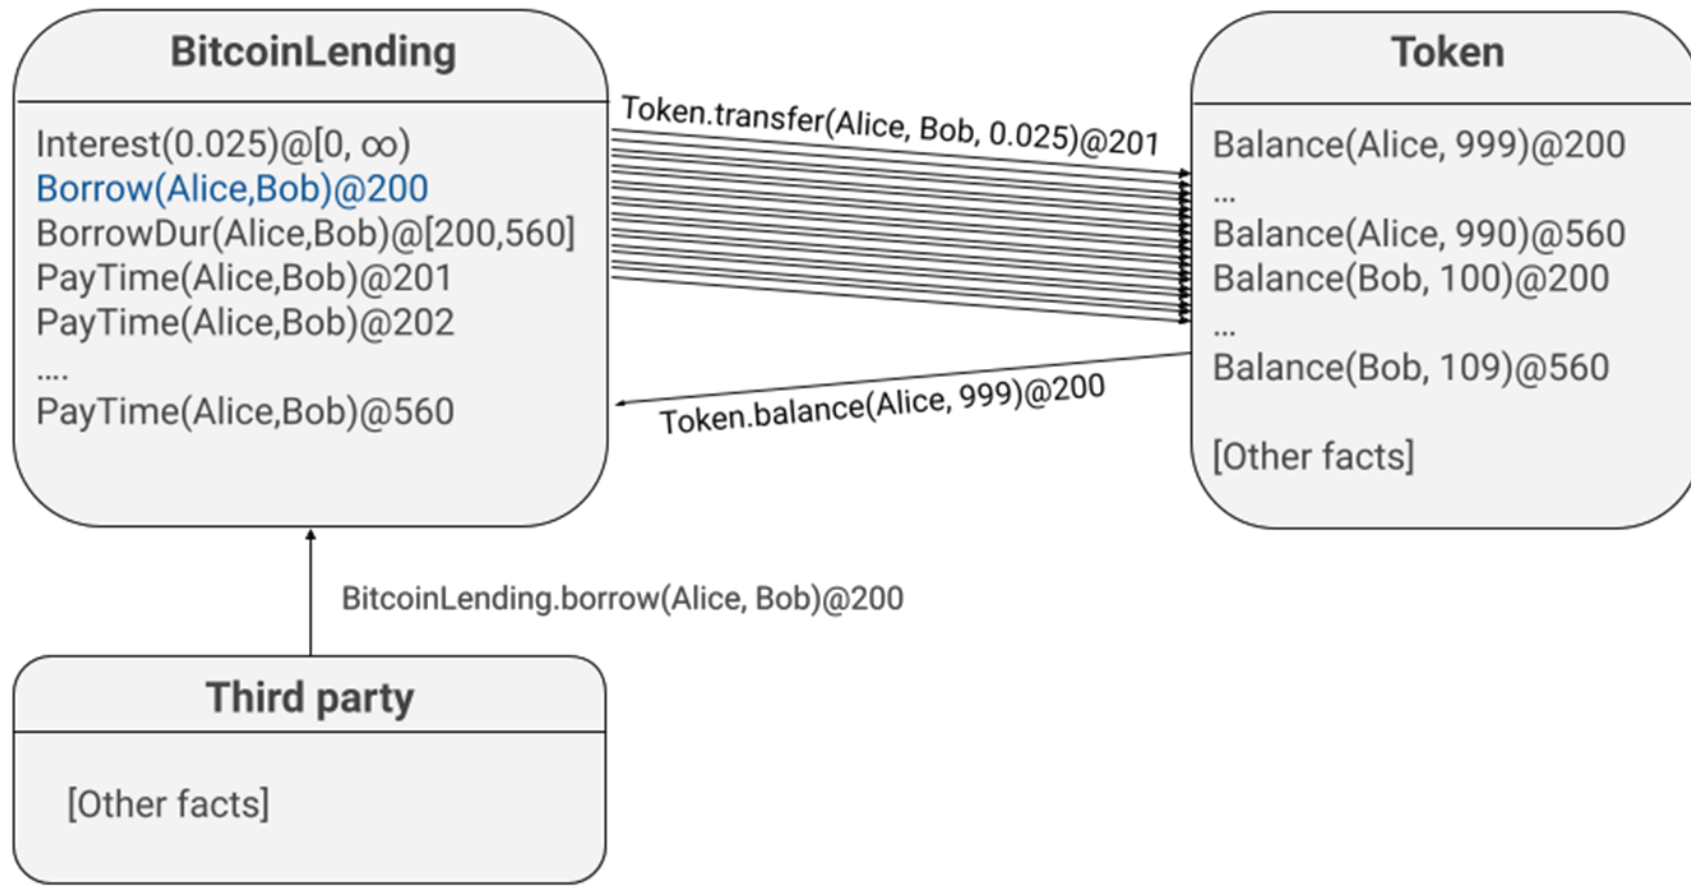
Generate Rule Types

7 8
Edge Type: DiamondMinusEdge
From Node: g262
To Node: g311
Term Order: 1,2,3,4
Left Interval Endpoint: -1
Right Interval Endpoint: -1

(Temporal-S) This edge match the logic of the diamond-minus operator of DatalogMTL




Modelling Smart Contracts with Temporal Datalog









“ExpressivE”: Expressive Knowledge Graph Embeddings



Logical Knowledge



Knowledge Graph Embeddings





Media Intelligence

