TECHNISCHE UNIVERSITÄT DRESDEN

Computational Logic ∴ Group

Hannes Strass (based on slides by Bernardo Cuenca Grau, Ian Horrocks, Przemysław Wałęga)

Faculty of Computer Science, Institute of Artificial Intelligence, Computational Logic Group

Introduction and Overview

Lecture 1, 14th Oct 2024 // Foundations of Knowledge Representation, WS 2024/25

Applications often need to represent knowledge

- seats on an aeroplane

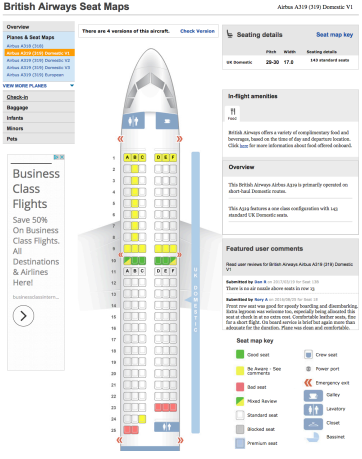

Applications often need to represent knowledge

- seats on an aeroplane
- account transactions

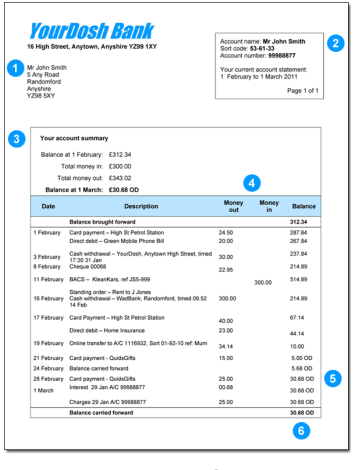

YourDosh Bank
16 High Street, Anytown, Anyshire YZ99 1XY

1 Mr John Smith
5 Any Road
Randomford
Anyshire
YZ98 5AY

2 Account name: **Mr John Smith**
Sort code: **53-61-33**
Account number: **99988877**

Your current account statement:
1 February to 1 March 2011

Page 1 of 1

3 **Your account summary**

Balance at 1 February: £312.34
Total money in: £300.00
Total money out: £343.02
Balance at 1 March: £30.68 OD

4

Date	Description	Money out	Money in	Balance
	Balance brought forward			**312.34**
1 February	Card payment – High St Petrol Station	24.50		287.84
	Direct debit – Green Mobile Phone Bill	20.00		267.84
3 February	Cash withdrawal – YourDosh, Anytown High Street, timed 17:30:31 Jan	30.00		237.84
8 February	Cheque 00068	22.95		214.89
11 February	BACS – KleanKars, ref JSS-999		300.00	514.89
	Standing order – Rent to J Jones			
16 February	Cash withdrawal – WadBank, Randomford, timed 09.52 14 Feb	300.00		214.89
17 February	Card Payment – High St Petrol Station	40.00		67.14
	Direct debit – Home Insurance	23.00		44.14
19 February	Online transfer to A/C 1116932, Sort 01-92-10 ref: Mum	34.14		10.00
21 February	Card payment - QuidsGifts	15.00		5.00 OD
24 February	Balance carried forward			5.68 OD
28 February	Card payment - QuidsGifts	25.00		30.68 OD
1 March	Interest 29 Jan A/C 99988877	00.68		30.68 OD
	Charges 29 Jan A/C 99988877	25.00		30.68 OD
	Balance carried forward			**30.68 OD**

5

6

Applications often need to represent knowledge

- seats on an aeroplane
- account transactions
- tall buildings

British Airways Seat Maps — Airbus A319 (319) Domestic V1

Name	City	State	Height	Feet	Floors	Completed
Federation Tower: East Tower	Moscow	Russia	373.7	1,226	95	2016
OKO: South Tower	Moscow	Russia	354.1	1,162	85	2015
Mercury City Tower	Moscow	Russia	338.8	1,112	75	2013
The Shard[1]	London	United Kingdom	309.7	1,017	87	2012
Eurasia[2]	Moscow	Russia	308.9	1,014	72	2014
CoC: Moscow Tower[3]	Moscow	Russia	301.6	990	76	2010
Skyland Istanbul 1[4][5]	Istanbul	Turkey	293.1	932	65	2017
Skyland Istanbul 2[4][5]	Istanbul	Turkey	293.1	932	65	2017
Metropol Istanbul[6]	Istanbul	Turkey	280	919	66	2017
Emaar Square	Istanbul	Turkey	280	920	62	2018
Naberezhnaya Tower C	Moscow	Russia	268.4	881	59	2007
Triumph Palace[7]	Moscow	Russia	264.1	867	57	2005
Commerzbank Tower[8][9]	Frankfurt	Germany	258.7	848	56	1997
CoC: Saint Petersburg Tower[10]	Moscow	Russia	256.9	843	65	2010
Messeturm	Frankfurt	Germany	256.5	842	55	1990
Nurol Life[11][12]	Istanbul	Turkey	252	827	60	2017
Torre de Cristal	Madrid	Spain	249	817	45	2008
Torre Cepsa	Madrid	Spain	248.3	815	45	2008
Evolution Tower[13]	Moscow	Russia	245.9	807	53	2014
OKO: North Tower[14]	Moscow	Russia	245	804	49	2014
Federation: West Tower	Moscow	Russia	243.2	798	62	2007
Main building of Moscow State University	Moscow	Russia	240	787	36	1953
Imperia Tower	Moscow	Russia	238.7	783	60	2011
Palace of Culture and Science	Warsaw	Poland	237	777	43	1955
Torre PwC	Madrid	Spain	236	774	52	2008
1 Canada Square	London	United Kingdom	235	771	50	1991
Istanbul Sapphire[15][16]	Istanbul	Turkey	234.9	770	54	2010
Tour First	Paris[a]	France	231	758	56	2011
Unicredit Tower	Milan	Italy	231	758	35	2011
Heron Tower	London	United Kingdom	230	755	46	2011

Applications often need to represent knowledge

- seats on an aeroplane
- account transactions
- tall buildings

and to answer questions

- seats available on flight?

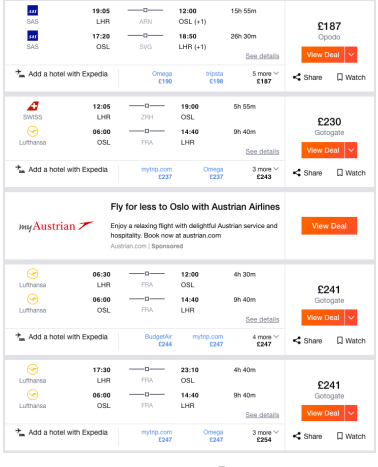

Applications often need to represent knowledge

- seats on an aeroplane
- account transactions
- tall buildings

and to answer questions

- seats available on flight?
- can afford to pay rent?

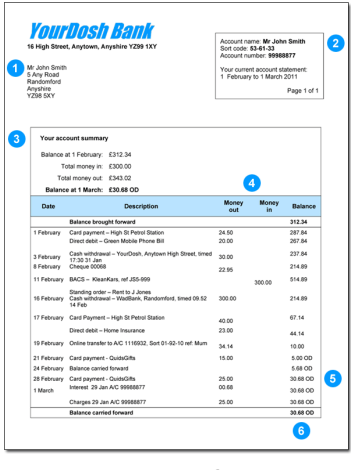

YourDosh Bank
16 High Street, Anytown, Anyshire YZ99 1XY

1 Mr John Smith
5 Any Road
Randomford
Anyshire
YZ98 5AY

2 Account name: **Mr John Smith**
Sort code: **53-61-33**
Account number: **99988877**

Your current account statement:
1 February to 1 March 2011

Page 1 of 1

3 **Your account summary**

Balance at 1 February: £312.34
Total money in: £300.00
Total money out: £343.02
Balance at 1 March: £30.68 OD

4

Date	Description	Money out	Money in	Balance
	Balance brought forward			312.34
1 February	Card payment – High St Petrol Station	24.50		287.84
	Direct debit – Green Mobile Phone Bill	20.00		267.84
3 February	Cash withdrawal – YourDosh, Anytown High Street, timed 17:30:31 Jan	30.00		237.84
8 February	Cheque 00068	22.95		214.89
11 February	BACS – KleanKars, ref JSS-999		300.00	514.89
16 February	Standing order – Rent to J Jones Cash withdrawal – WadBank, Randomford, timed 09.52 14 Feb	300.00		214.89
17 February	Card Payment – High St Petrol Station	40.00		67.14
	Direct debit – Home Insurance	23.00		44.14
19 February	Online transfer to A/C 1116932, Sort 01-92-10 ref: Mum	34.14		10.00
21 February	Card payment - QuidsGifts	15.00		5.00 OD
24 February	Balance carried forward			5.68 OD
28 February	Card payment - QuidsGifts	25.00		30.68 OD
1 March	Interest 29 Jan A/C 99988877	00.68		30.68 OD
	Charges 29 Jan A/C 99988877	25.00		30.68 OD
	Balance carried forward			**30.68 OD**

5

6

Applications often need to represent knowledge

- seats on an aeroplane
- account transactions
- tall buildings

and to answer questions

- seats available on flight?
- can afford to pay rent?
- tallest building in Europe?

What kind of representation?

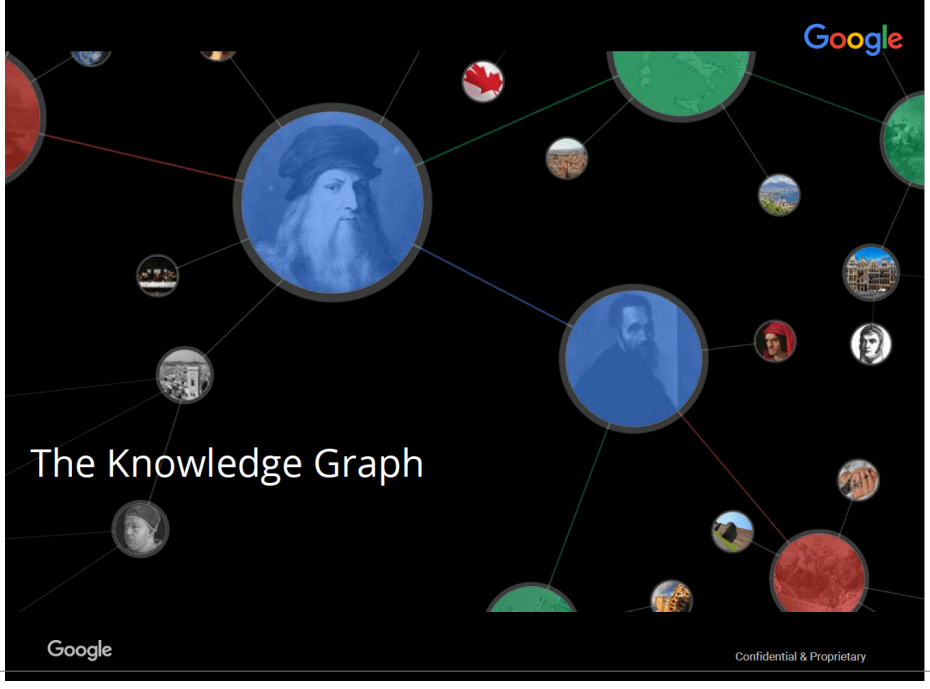
Google
The Knowledge Graph
Google
Confidential & Proprietary

The Knowledge Graph

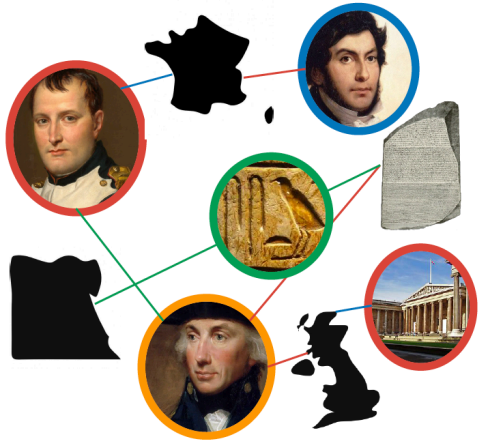

The Knowledge Graph is a comprehensive collection of real-world entities (people, places, things, and concepts) along with relationships and factual attributes that describe them.

Google

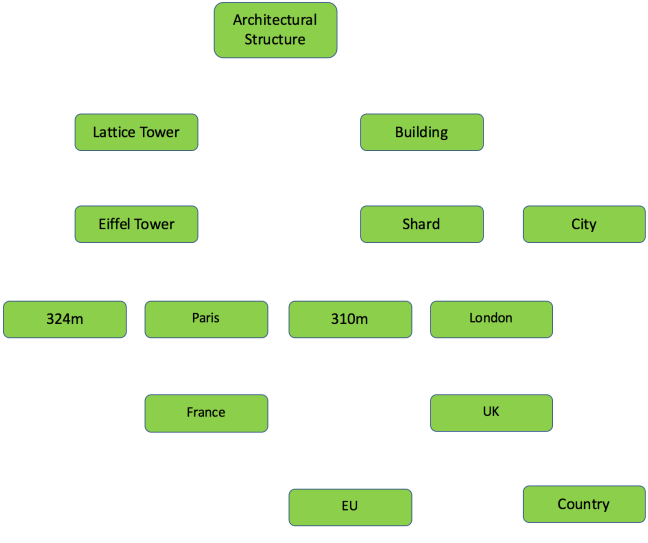
Architectural Structure
Lattice Tower
Building
Eiffel Tower
Shard
City
324m
Paris
310m
London
France
UK
EU
Country

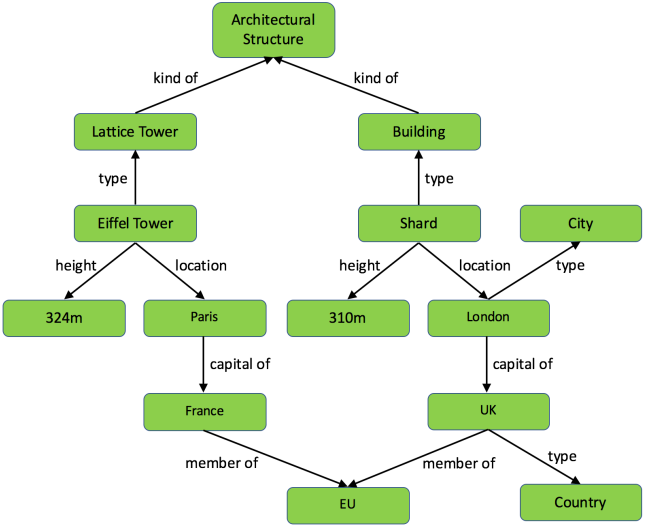
Architectural Structure
kind of
kind of
Lattice Tower
Building
type
type
Eiffel Tower
Shard
City
height
location
height
location
type
324m
Paris
310m
London
capital of
capital of
France
UK
member of
member of
type
EU
Country

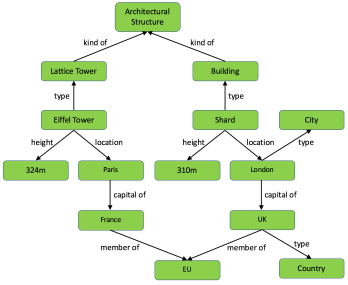

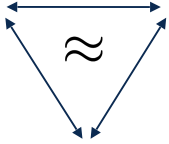

Architectural Structure			
name	**location**	**height**	**kind**
Shard	London	310m	Building
Eiffel Tower	Paris	324m	Lattice Tower
…	…	…	…

City	
name	**capital of**
London	UK
Paris	France
…	…

member of	
country	**organisation**
France	EU
UK	EU
…	…

Building(Shard)
City(London)
location(Shard,London)
height(Shard,310m)
capitalOf(London,UK)
…

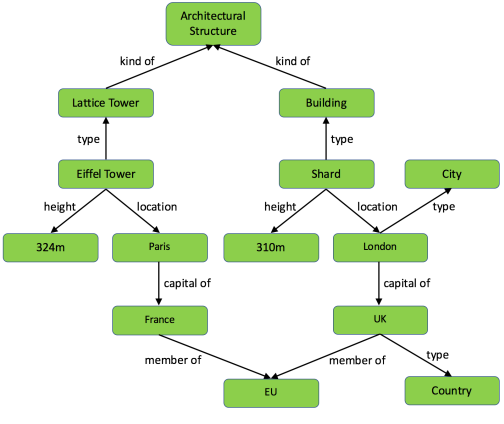

Reasoning is the process of **answering queries** w.r.t. the represented knowledge

What is the height of the Eiffel Tower?

```
SELECT ?x
WHERE { EiffelTower height ?
x. }
```

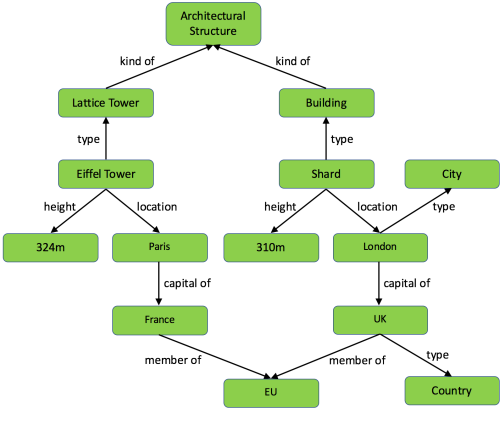

Reasoning is the process of **answering queries** w.r.t. the represented knowledge

What is the height of the Eiffel Tower?

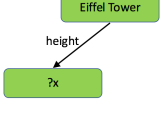

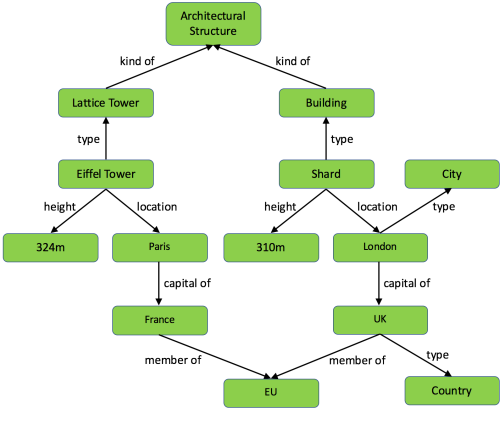

- Reasoning is the process of **answering queries** w.r.t. the represented knowledge

- What is the height of the Eiffel Tower?

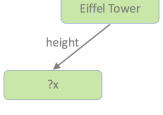

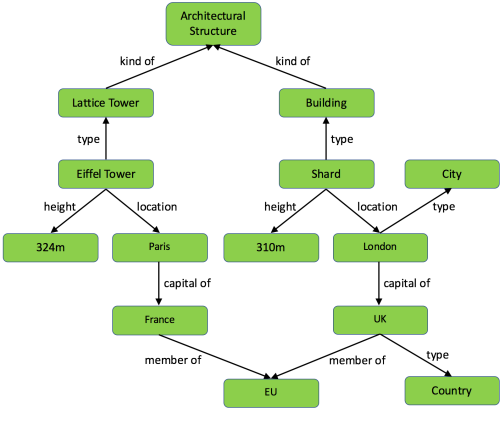

- Reasoning is the process of **answering queries** w.r.t. the represented knowledge
- What is the height of the Eiffel Tower?
- **324m**

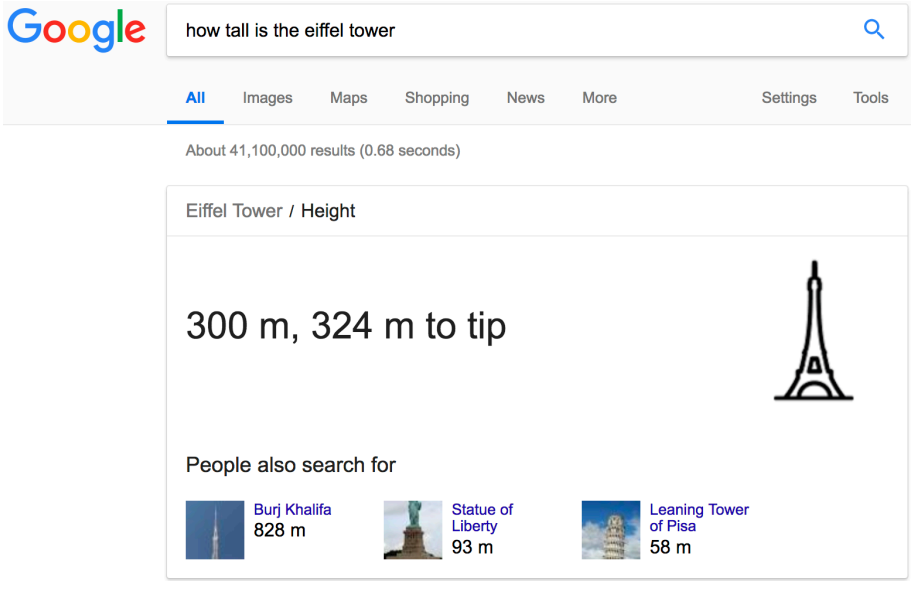
Google
how tall is the eiffel tower
All
Images
Maps
Shopping
News
More
Settings
Tools
About 41,100,000 results (0.68 seconds)
Eiffel Tower / Height
300 m, 324 m to tip
People also search for
Burj Khalifa
828 m
Statue of Liberty
93 m
Leaning Tower of Pisa
58 m

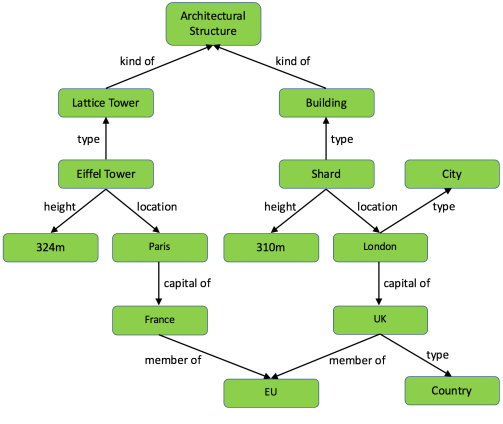

- What is the difference between a Building, a Lattice Tower and an Architectural Structure?
- Is the Eiffel Tower a Building; is it an Architectural Structure?
- Special meaning of, e.g., type and kind of edges?

- **Semantics**: the study of meaning

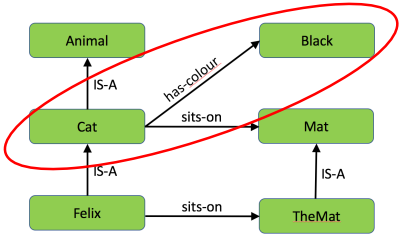

[Quillian, 1967]

(Precise) semantics needed in order to define what (correct) query answers should be

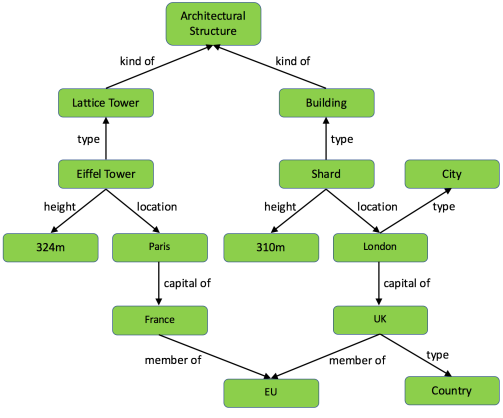

- What country is the Shard located in?

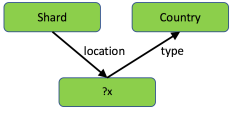

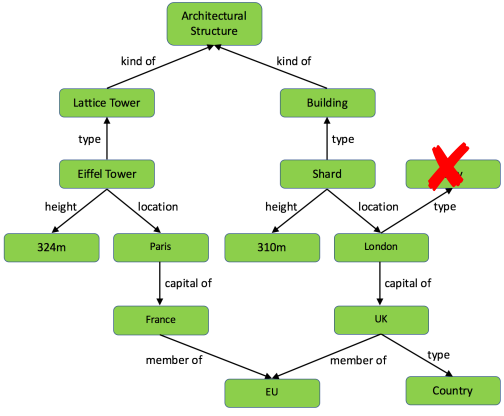

- What country is the Shard located in?
- ???

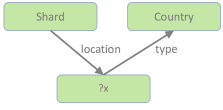

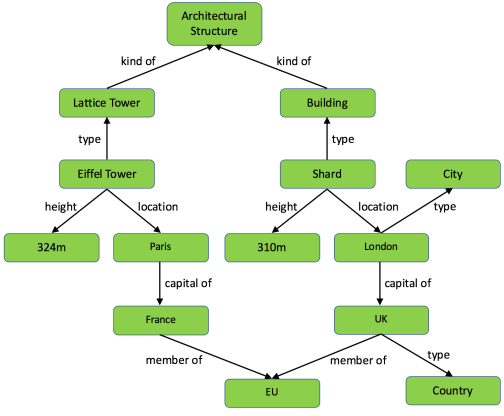

- What Architectural Structures are located in London?

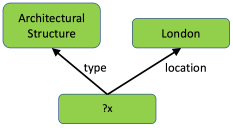

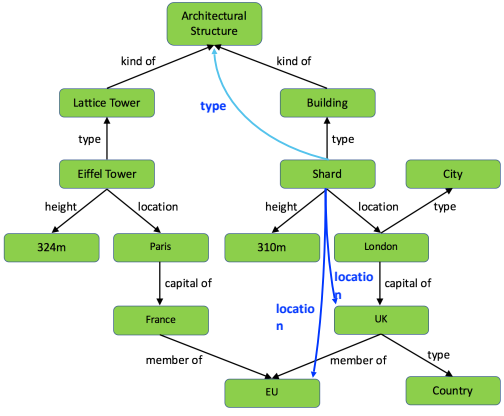
Architectural Structure
kind of
kind of
Lattice Tower
type
Building
type
type
Eiffel Tower
Shard
City
height
location
height
location
type
324m
Paris
310m
London
capital of
location
capital of
France
location
UK
member of
member of
type
EU
Country

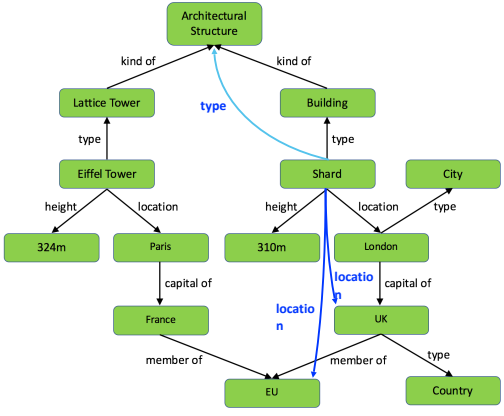

- Every Building located in London is also located in
  - UK
  - EU
  - England
  - Northern Hemisphere
  - …
- Need to add a **very** large number of edges

A city that is the capital of a country is a (geographical) part of that country†

A thing that is located in a city that is a (geographical) part of a country is also located in that country

A thing that is located in a country that is a member of a supranational union is also located in that supranational union

† Part-whole relationships are complicated! They are the subject of a whole field of study in logic and philosophy: mereology

- **Semantics**: the study of meaning

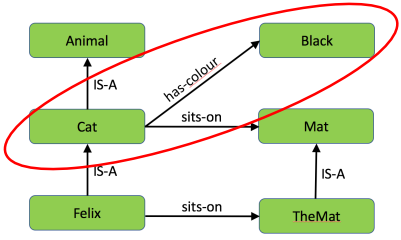

[Quillian, 1967]

(Precise) semantics needed in order to define what (correct) query answers should be

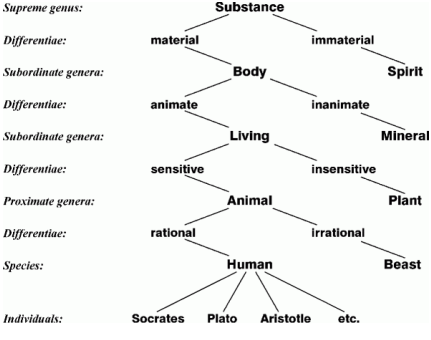

"All men are mortal, all Greeks are men, therefore all Greeks are mortal"
(syllogism)

Modern KR languages are often based on logic

Typically (subsets of) First Order Predicate Calculus

Gottlob Frege

Charles Sanders Peirce

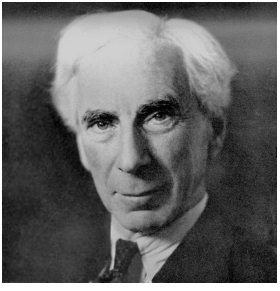

Bertrand Russell

29

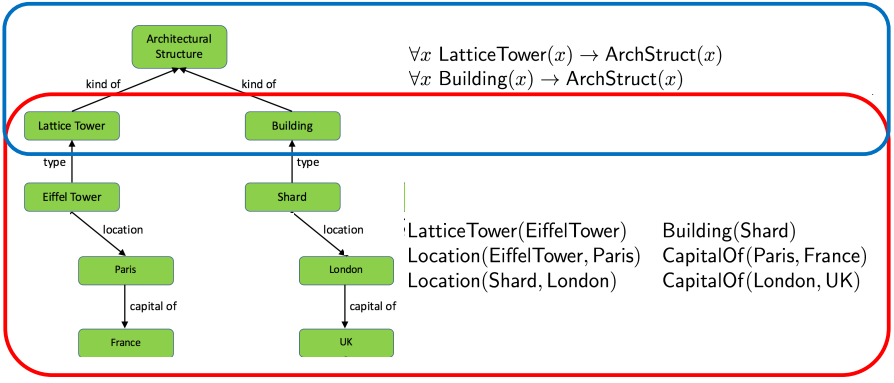

facts / data / (RDF) graph

ontology / conceptual schema

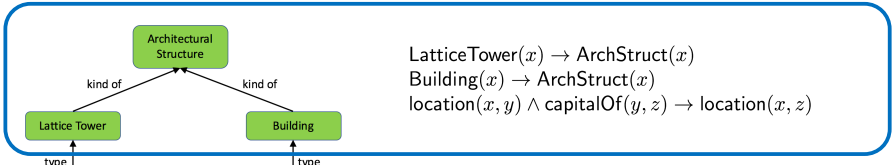

$\mathsf{LatticeTower}(x) \rightarrow \mathsf{ArchStruct}(x)$
$\mathsf{Building}(x) \rightarrow \mathsf{ArchStruct}(x)$
$\mathsf{location}(x, y) \wedge \mathsf{capitalOf}(y, z) \rightarrow \mathsf{location}(x, z)$

$\mathsf{LatticeTower}(\mathsf{EiffelTower})$
$\mathsf{Location}(\mathsf{EiffelTower}, \mathsf{Paris})$
$\mathsf{Location}(\mathsf{Shard}, \mathsf{London})$
$\mathsf{Building}(\mathsf{Shard})$
$\mathsf{CapitalOf}(\mathsf{Paris}, \mathsf{France})$
$\mathsf{CapitalOf}(\mathsf{London}, \mathsf{UK})$

facts / data / (RDF) graph

ontology / conceptual schema

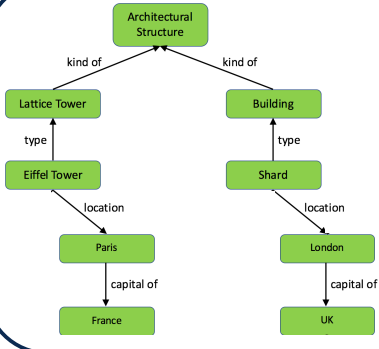

$\mathsf{LatticeTower}(x) \rightarrow \mathsf{ArchStruct}(x)$
$\mathsf{Building}(x) \rightarrow \mathsf{ArchStruct}(x)$
$\mathsf{location}(x, y) \wedge \mathsf{capitalOf}(y, z) \rightarrow \mathsf{location}(x, z)$

$\mathsf{LatticeTower}(\mathsf{EiffelTower})$
$\mathsf{Location}(\mathsf{EiffelTower}, \mathsf{Paris})$
$\mathsf{Location}(\mathsf{Shard}, \mathsf{London})$
$\mathsf{Building}(\mathsf{Shard})$
$\mathsf{CapitalOf}(\mathsf{Paris}, \mathsf{France})$
$\mathsf{CapitalOf}(\mathsf{London}, \mathsf{UK})$

facts / data / (RDF) graph
ontology / conceptual schema
} knowledge base

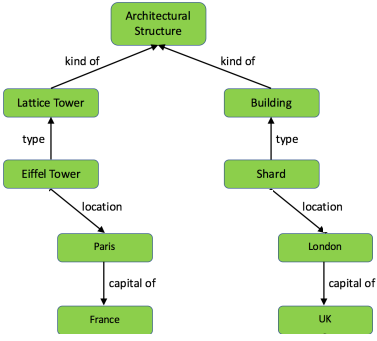

$\mathsf{LatticeTower}(x) \rightarrow \mathsf{ArchStruct}(x)$
$\mathsf{Building}(x) \rightarrow \mathsf{ArchStruct}(x)$
$\mathsf{location}(x, y) \wedge \mathsf{capitalOf}(y, z) \rightarrow \mathsf{location}(x, z)$

$\mathsf{LatticeTower}(\mathsf{EiffelTower})$ $\quad$ $\mathsf{Building}(\mathsf{Shard})$
$\mathsf{Location}(\mathsf{EiffelTower}, \mathsf{Paris})$ $\quad$ $\mathsf{CapitalOf}(\mathsf{Paris}, \mathsf{France})$
$\mathsf{Location}(\mathsf{Shard}, \mathsf{London})$ $\quad$ $\mathsf{CapitalOf}(\mathsf{London}, \mathsf{UK})$

$\mathcal{K} \models \mathsf{ArchStruct}(\mathsf{EiffelTower})$?

True
True
True
True
False

Devise algorithms that compute query answers
E.g., using natural deduction rules:

$$\frac{\forall x P(x) \to R(x)}{\forall x Q(x) \to R(x) \quad \forall x P(x) \to Q(x)} \qquad \frac{\forall x \mathsf{Greek}(x) \to \mathsf{Mortal}(x)}{\forall x \mathsf{Man}(x) \to \mathsf{Mortal}(x) \quad \forall x \mathsf{Greek}(x) \to \mathsf{Man}(x)}$$

$$\frac{Q(a)}{\forall x P(x) \to Q(x) \quad P(a)} \qquad \frac{\mathsf{ArchStruct}(\mathsf{EiffelTower})}{\forall x\, \mathsf{Building}(x) \to \mathsf{ArchStruct}(x) \quad \mathsf{Building}(\mathsf{EiffelTower})}$$

Can check/prove algorithms are sound and complete w.r.t. semantics

Turing showed that some problems cannot be completely solved using standard computational model

- halting problem
- **FOL entailment problem**

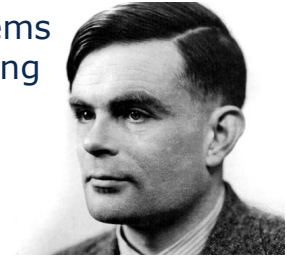

Even if decidable, reasoning might be of inherently high complexity and so take an infeasibly long time

“Scruffy” approach:
Ad-hoc representation
Efficient but (at least) incomplete algorithms

- ✓ Can use arbitrarily powerful representation
- ✓ Favourable scalability properties
- ✗ Incomplete answers
  - ✗ Degree of incompleteness unknown
  - ✗ Incompleteness can easily become unsoundness

“Neat” approach:

Study KR languages to find appropriate balance of expressive power and computability

Design algorithms that work well in typical cases

Develop highly optimised implementations

✓ Precisely defined semantics
✓ Formal properties well understood
✓ Sound and complete reasoning
✗ Limited representation power
✗ Optimisations may not offer robust scalability

Family of logic-based KR languages

Most are decidable subsets of FO logic

Provide a range of different constructors

- Booleans (and, or, not)
- Restricted forms of quantification (exists, forall)
- Counting (atmost, atleast)
- …

Decidability/complexity and (efficient) algorithms known for many combinations of constructors

Highly optimised implementations for various "sweet spot" languages

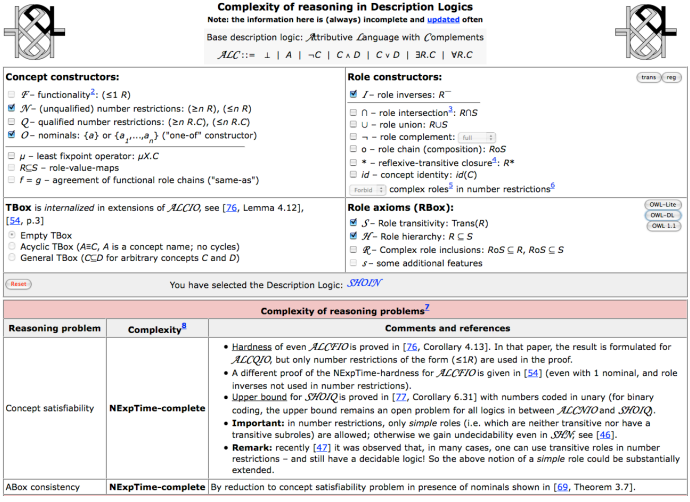

Complexity of reasoning problems[7]

Reasoning problem	Complexity[8]	Comments and references
Concept satisfiability	**NExpTime-complete**	• Hardness of even $\mathcal{ALCFIO}$ is proved in [76, Corollary 4.13]. In that paper, the result is formulated for $\mathcal{ALCQIO}$, but only number restrictions of the form (≤1R) are used in the proof. • A different proof of the NExpTime-hardness for $\mathcal{ALCFIO}$ is given in [54] (even with 1 nominal, and role inverses not used in number restrictions). • Upper bound for $\mathcal{SHOIQ}$ is proved in [77, Corollary 6.31] with numbers coded in unary (for binary coding, the upper bound remains an open problem for all logics in between $\mathcal{ALCNIO}$ and $\mathcal{SHOIQ}$. • **Important:** in number restrictions, only *simple* roles (i.e. which are neither transitive nor have a transitive subroles) are allowed; otherwise we gain undecidability even in $\mathcal{SHN}$, see [46]. • **Remark:** recently [47] it was observed that, in many cases, one can use transitive roles in number restrictions – and still have a decidable logic! So the above notion of a *simple* role could be substantially extended.
ABox consistency	**NExpTime-complete**	By reduction to concept satisfiability problem in presence of nominals shown in [69, Theorem 3.7].

List of reasoners

Reasoner	Instution	Download	Publication
BaseVISor	VIStology, Inc.	Download	Core publication
BUNDLE	University of Ferrara	Download	Core publication
CEL	Technische Universität Dresden	Download	Core publication
Chainsaw	The University of Manchester	Download	Core publication
Clipper	Vienna University of Technology	Download	Core publication
DBOWL	University of Malaga	Download	Core publication
DeLorean	Not given	Download	Core publication
DistEL	Wright State University	Download	Core publication
DRAOn	University of Paris 8, IUT of Montreuil	Download	Core publication
DReW	Vienna University of Technology	Download	Core publication
ELepHant	Not given	Download	Core publication
ELK	University of Ulm, Germany	Download	Core publication
ELOG	Not given	Download	Core publication
FaCT++	The University of Manchester	Download	Core publication
fuzzyDL	ISTI – CNR	Download	Core publication
HermiT	University of Oxford	Download	Core publication
jcel	Technische Universität Dresden	Download	Core publication
JFact	The University of Manchester	Download	Core publication
Konclude	University of Ulm, derivo GmbH	Download	Core publication
LiFR	Centre for Research and Technology Hellas (CERTH)	Download	Core publication
Mastro	Sapienza University of Rome	Download	Core publication
MORe	University of Oxford	Download	Core publication
ontop	Free University of Bozen-Bolzano	Download	Core publication

Standardised KR language

- RDF provides a graphical data model
- OWL provides a DL-based ontology language

OWL ontology:

$$\mathsf{LatticeTower}(x) \rightarrow \mathsf{ArchStruct}(x)$$
$$\mathsf{Building}(x) \rightarrow \mathsf{ArchStruct}(x)$$
$$\mathsf{location}(x, y) \wedge \mathsf{capitalOf}(y, z) \rightarrow \mathsf{location}(x, z)$$

RDF data:

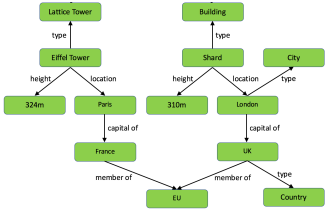

Standardised KR language

- RDF provides a graphical data model
- OWL provides a DL-based ontology language

Developed as part of **W3C's Semantic Web** project

> "A new form of Web content that is meaningful to computers will unleash a revolution of new possibilities" **(!)**

Standardised KR language

- RDF provides a graphical data model
- OWL provides a DL-based ontology language

Developed as part of **W3C's Semantic Web** project

Now **widely used** in science, healthcare and Industry

Often referred to as "**semantic technology**"

Based on powerful but still decidable DL (**SROIQ**)

Three "profiles" based on tractable subsets

- **QL**: based on the DL-Lite description logic
- **EL**: based on the EL description logic
- **RL**: based on the DL fragment of Datalog (aka DLP)

Different algorithmic techniques

- (**Hyper-**) **Tableau** for full language
- **Query rewriting** for QL
- **Consequence-based** for EL
- **Materialisation** for RL

Highly optimised implementations

- Several of which have been developed here at Oxford


R D F
WOL

Tools:

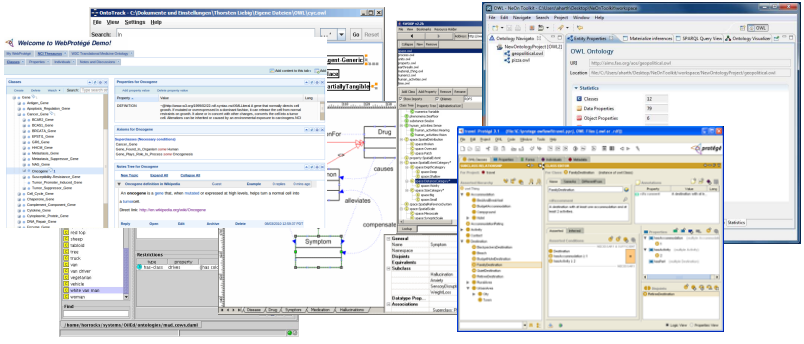
Welcome to WebProtégé Demo!
OWL Ontology
Drug
causes
alleviates
compensates
Symptom

Reasoners:

Hermi FaCT++ ORACLE

Racer pellet uOnto Querying ONTOlogies

KAON 2 CEL semantic web framework Jena TrOWL www.trowl.eu

Applications: Question Answering

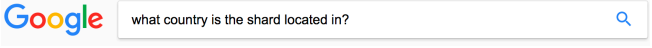


Google
what country is the shard located in?
All
Maps
News
Images
Shopping
More
Settings
Tools
About 6,400,000 results (0.81 seconds)
The Shard / Country
Norway
United Kingdom
Poland
Germany
Ukraine
France
Spain
Italy
Turkey
Map data ©2018
United Kingdom
People also search for
View 10+ more
England
London
Great Britain
Scotland
United States of America
Wales
Ireland

SNOMED CT
The global language of healthcare

SNOMED is a **huge** medical ontology
- More than 500,000 terms!

Why SNOMED? Let's ask Healthcare experts!
- "We need a clinical ontology that is universal, so any term I use is the same as every other colleague around the country"
- "SNOMED is the glue that binds the clinical community together and is the platform for all clinically relevant information"

Used to annotate patient records in more than 20 countries, including UK, USA, New Zealand, …

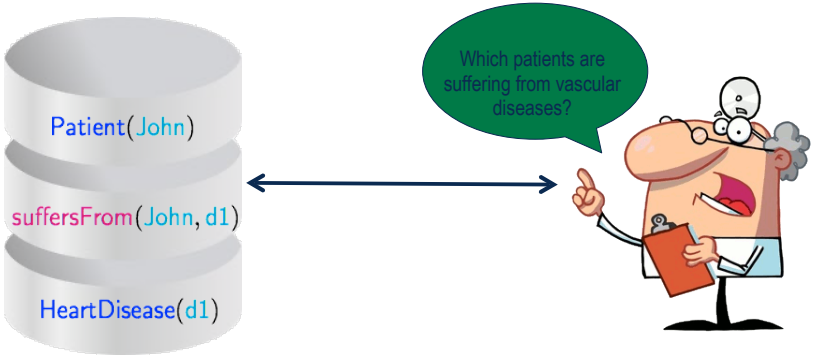
Patient(John)
suffersFrom(John, d1)
HeartDisease(d1)
Which patients are suffering from vascular diseases?

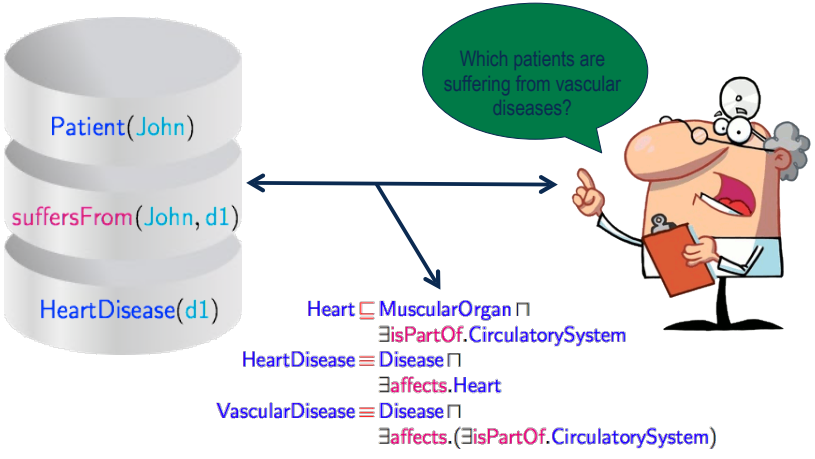
Patient(John)
suffersFrom(John, d1)
HeartDisease(d1)
Which patients are suffering from vascular diseases?
Heart ⊑ MuscularOrgan ⊓ ∃isPartOf.CirculatorySystem
HeartDisease ≡ Disease ⊓ ∃affects.Heart
VascularDisease ≡ Disease ⊓ ∃affects.(∃isPartOf.CirculatorySystem)

The end?

Extensions

- Arithmetic functions and aggregation
- Reasoning about time
- Data streams

Algorithms

- Consequence-based reasoning
- Hybrid rewriting/materialisation

Optimisation and implementation

- Incremental reasoning
- Query planning
- HPC, including large-scale and distributed architectures

Tools and applications

Course Structure

Logics for KRR
- Propositional and First Order Logic
- Ontological modelling
- The role of reasoning

Horn logics and Datalog
- Expressivity and formal properties
- Reasoning

Description Logic
- Motivation and foundations
- Model theory
- Tableau reasoning
- Reasoning with data
- Lightweight description logics
- Ontology based data access

Course Structure

Description Logic -v- Datalog
- Expressiveness and decidability
- Combining DL and Datalog
- Other decidable fragments of FOL

Ontology Languages and Semantic Technologies
- RDF & OWL
- SROIQ
- Non-DL features
- Profiles
- Tools and Reasoners

Nonmonotonic Reasoning
- Limitations of FOL
- Closed World Assumption and negation as failure
- Stable model semantics

Reading List

Primary Text

- An Introduction to Description Logic. Franz Baader, Ian Horrocks, Carsten Lutz, Uli Sattler

Supplementary Texts

- Handbook of Knowledge Representation. Frank van Harmelen, Vladimir Lifschitz and Bruce Porter (Eds). Foundations of Artificial Intelligence, 2008.
- Foundations of Semantic Web Technologies. Chapman & Hall/ CRC Textbooks in Computing. Pascal Hitzler, Markus Kroetsch, and Sebastian Rudolph, 2009.