

Walk-in Introduction to Linked Data

Rinke Hoekstra

Vrije Universiteit Amsterdam
Amsterdam Data Science
rinke.hoekstra@vu.nl



Amsterdam
Data Science



What Happens in an Internet Minute?



And Future Growth is Staggering





What Happens in an Internet Minute?

1,572,877 GB of global IP data transferred¹

10 Million ads displayed²

347,222 Tweets³

3.3 Million pieces of content shared⁴

6.9 Million messages sent⁵

Netflix + Youtube = more than 1/2 of all traffic⁵

438,801 Wiki page views⁷

\$400 Million during Alibaba peak day sales⁶

10 Million WeChat messages at its peak⁹

34.7 Million instant messages (MIM) sent⁸

194,064 app downloads¹⁰

\$133,436 in sales¹¹

31,773 hours of music played¹²

38,194 photos uploaded¹³

57,870 page views¹⁴

100 hours of video uploaded¹⁶

138,889 hours of video watched¹⁶

4.1 Million searches¹⁵

23,148 hours of video watched¹⁷

And Future Growth is Staggering



By 2017, mobile traffic will have grown **13X** in just 5 years¹



In 2017, there will be **3X** more connected devices than people on Earth¹

All digital data created reached **4 zettabytes** in 2013¹⁸

1 "Cisco Visual Networking Index: Forecast and Methodology, 2012-2017" White Paper, Cisco, May 29, 2013. http://www.cisco.com/en/US/solutions/collateral/rs341/rs525/rs537/rs705/rs827/white_paper_c11-481360.pdf

2 "2013 U.S. Digital Future in Focus Whitepaper", February 14, 2013. http://www.comscore.com/Insights/Presentations_and_Whitepapers/2013/2013_US_Digital_Future_in_Focus

3 Twitter IPO - SEC Report <http://www.sec.gov/Archives/edgar/data/1418091/000119312513390321/d564001ds1.htm>

4 "A Focus on Efficiency", a whitepaper from Facebook, Ericsson and Qualcomm, page 6, September, 2013. https://fbcdn-dragon-a.akamaihd.net/hphotos-akash3/851590_229753833859617_1129962605_n.pdf

5 "Global Internet Phenomena Report", Sandvine, 2H 2013. <https://www.sandvine.com/downloads/general/global-internet-phenomena/2013/2h-2013-global-internet-phenomena-report.pdf>

9 "On the eve of Chinese New Year, WeChat processed 10 million messages in one minute", Tech in Asia, February 3, 2014. <http://sg.finance.yahoo.com/news/eve-chinese-wechat-processed-10-041119017.html>

10 "Gartner Says Mobile App Stores Will See Annual Downloads Reach 102 Billion in 2013", Gartner, September, 2013. <http://www.gartner.com/newsroom/id/2592315>

11 Amazon Quarterly Income Statement, Year Ending September 31, 2013. <http://www.google.com/finance?q=NASDAQ:AMZN&fstyle=iiif>

12 Pandora Monthly Audience Metrics Press Releases, January 2013-December 2013. <http://press.pandora.com/phoenix.zhtml?c=251764&p=irol-news&ry=1>

13 Instagram Press Site, February 2013. <http://instagram.com/press/>


14 "The CMO's Guide to Pinterest", iCrossing, Inc., Page 3, November, 2013. <http://www.icrossing.com/>

Data Growth (2007)

THE WORLD'S CAPACITY TO STORE INFORMATION

This chart shows the world's growth in storage capacity for both analog data (books, newspapers, videotapes, etc.) and digital (CDs, DVDs, computer hard drives, smartphone drives, etc.)

In gigabytes or estimated equivalent

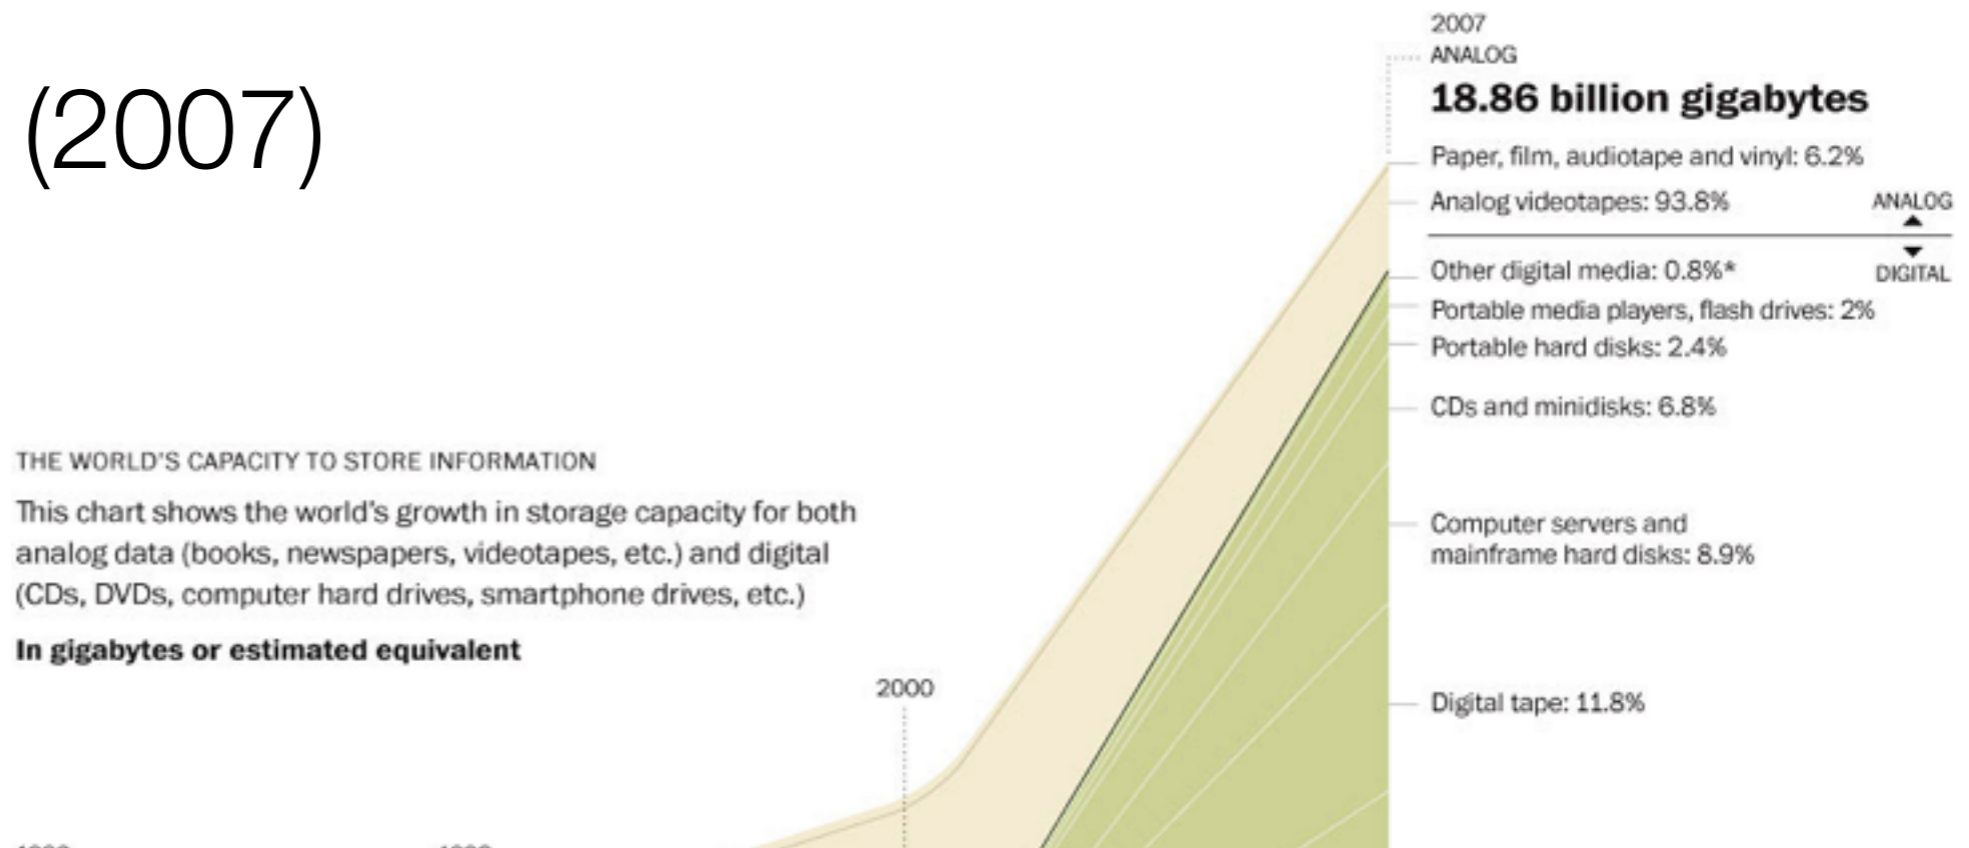


Data Growth (2007)

THE WORLD'S CAPACITY TO STORE INFORMATION

This chart shows the world's growth in storage capacity for both analog data (books, newspapers, videotapes, etc.) and digital (CDs, DVDs, computer hard drives, smartphone drives, etc.)

In gigabytes or estimated equivalent




The **Web** as a huge **Database**?

DIGITAL
0.02 billion

COMPUTING POWER

In 1986, pocket calculators accounted for much of the world's data-processing power.

Percentage of available processing power by device:



PC hard disks: 44.5%
123 billion gigabytes

*Other includes chip cards, memory cards, floppy disks, mobile phones/PDAs, cameras/camcorders, video games

2007 DIGITAL
276.12 billion gigabytes

Web Data \neq Web **of** Data



```
{
  "id": "524200751",
  "name": "Rinke Hoekstra",
  "birthday": "09/10/1977",
  "education": [
    {
      "school": {
        "id": "110440178984531",
        "name": "Bertrand Russell College"
      },
      "type": "High School"
    },
    {
      "school": {
        "id": "113928981951563",
        "name": "University of Amsterdam"
      },
      "degree": {
        "id": "102167843158588",
        "name": "Master's Degree"
      },
      "year": {
        "id": "137409666290034",
```

```
untitled
{
  "_total": 1,
  "values": [{
    "_key": "~",
    "firstName": "Rinke",
    "headline": "Researcher at Vrije Universiteit Amsterdam and Universiteit van Amsterdam",
    "id": "78GGmSNgZq",
    "lastName": "Hoekstra",
    "pictureUrl":
    • "http://m3.lidn.com/mpr/mprx/0_XGm9LH71dwUnvW4ZXFSdLE7fWWDZvoWZQ8sWLwSfnuJeEmY4eXYVwIJ3bFSb9DeNL3uHo2P5bVvC"
  }]
}
```

Line: 11 Column: 2 JavaScript Tab Size: 2

DIGITAL
276.12 billion gigabytes

Web 2.0 - problems



JSON only prescribes **data structure**

Data is **accessible** but still locked in **silos**

Data ownership \approx market share

Data **integration** is a major issue

Web 2.0 - problems

```
{
  "id": "524200751",
  "name": "Rinke Hoekstra",
  "birthday": "09/10/1977",
  "education": [
    {
      "school": {
        "id": "110440178984531",
        "name": "Bertrand Russell College"
      },
      "type": "High School"
    },
    {
      "school": {
        "id": "113928981951563",
        "name": "University of Amsterdam"
      },
      "degree": {
        "id": "102167843158588",
        "name": "Master's Degree"
      },
      "year": {
        "id": "137409666290034",
        "name": "1995"
      },
      "concentration": [
        {
          "id": "110719132288970",
          "name": "AI"
        }
      ]
    },
    {
      "type": "Graduate School"
    }
  ]
}
```

JSON only prescribes **data structure**

Data is **accessible** but still locked in **silos**

Data ownership \approx market share

Data **integration** is a major issue

Web 2.0 - problems





JSON only prescribes **data structure**

Data is **accessible** but still locked in **silos**

Data ownership \approx market share

Data **integration** is a major issue

```
{
  "id": "524200751",
  "name": "Rinke Hoekstra",
  "birthday": "09/10/1977",
  "education": [
    {
      "school": {
        "id": "110440178984531",
        "name": "Bertrand Russell College"
      },
      "type": "High School"
    },
    {
      "school": {
        "id": "113928981951563",
        "name": "University of Amsterdam"
      },
      "degree": {
        "id": "102167843158588",
        "name": "Master's Degree"
      },
      "year": {
        "id": "137409666290034",
        "name": "1995"
      }
    },
    {
      "concentration": [
        {
          "id": "110719132288970",
          "name": "AI"
        }
      ]
    }
  ],
  "type": "Graduate School"
}
```

	524200751
	78GGmSNgZq

```
{
  "_total": 1,
  "values": [
    {
      "_key": "-",
      "firstName": "Rinke",
      "headline": "Researcher at Vrije Universiteit Amsterdam and Universiteit van Amsterdam",
      "id": "78GGmSNgZq",
      "lastName": "Hoekstra",
      "pictureUrl":
      "http://m3.llicdn.com/mpr/mprx/8_XGm9LH71dwUnvM4ZXF5dLE7fWWDZvoRZQ8sRLwSfnuJeEmY4eXYVwI13bF5b9DeNL3uHo2PSbWvC"
    }
  ]
}
```

Gov 2.0 - problems

Databases only prescribe **data structure**

Data is **accessible** but still locked in **silos**

Data ownership \approx power

Data **integration** is a major issue

Gov 2.0 - problems

```
{
  "id": "524200751",
  "name": "Rinke Hoekstra",
  "birthday": "09/10/1977",
  "education": [
    {
      "school": {
        "id": "110440178984531",
        "name": "Bertrand Russell College"
      },
      "type": "High School"
    },
    {
      "school": {
        "id": "113928981951563",
        "name": "University of Amsterdam"
      },
      "degree": {
        "id": "102167843158588",
        "name": "Master's Degree"
      },
      "year": {
        "id": "137409666290034",
        "name": "1995"
      },
      "concentration": [
        {
          "id": "110719132288970",
          "name": "AI"
        }
      ]
    },
    {
      "type": "Graduate School"
    }
  ]
}
```

Databases only prescribe **data structure**



Data is **accessible** but still locked in **silos**

Data ownership \approx power

Data **integration** is a major issue

Gov 2.0 - problems

```
{
  "id": "524200751",
  "name": "Rinke Hoekstra",
  "birthday": "09/10/1977",
  "education": [
    {
      "school": {
        "id": "110440178984531",
        "name": "Bertrand Russell College"
      },
      "type": "High School"
    },
    {
      "school": {
        "id": "113928981951563",
        "name": "University of Amsterdam"
      },
      "degree": {
        "id": "102167843158588",
        "name": "Master's Degree"
      },
      "year": {
        "id": "137409666290034",
        "name": "1995"
      }
    },
    {
      "concentration": [
        {
          "id": "110719132288970",
          "name": "AI"
        }
      ]
    }
  ],
  "type": "Graduate School"
}
```

	524200751
	78GGmSNgZq

```
{
  "_total": 1,
  "values": [
    {
      "_key": "-",
      "firstName": "Rinke",
      "headline": "Researcher at Vrije Universiteit Amsterdam and Universiteit van Amsterdam",
      "id": "78GGmSNgZq",
      "lastName": "Hoekstra",
      "pictureUrl": "http://m3.l1cdn.com/mpr/mprx/8_XGm9LH71dwUnvM4ZXF5dLE7fWWDZvovR2Q8sRLwSfnuJeEmY4eXYVwI33bFSb9DeNL3uHo2PSbWvC"
    }
  ]
}
```



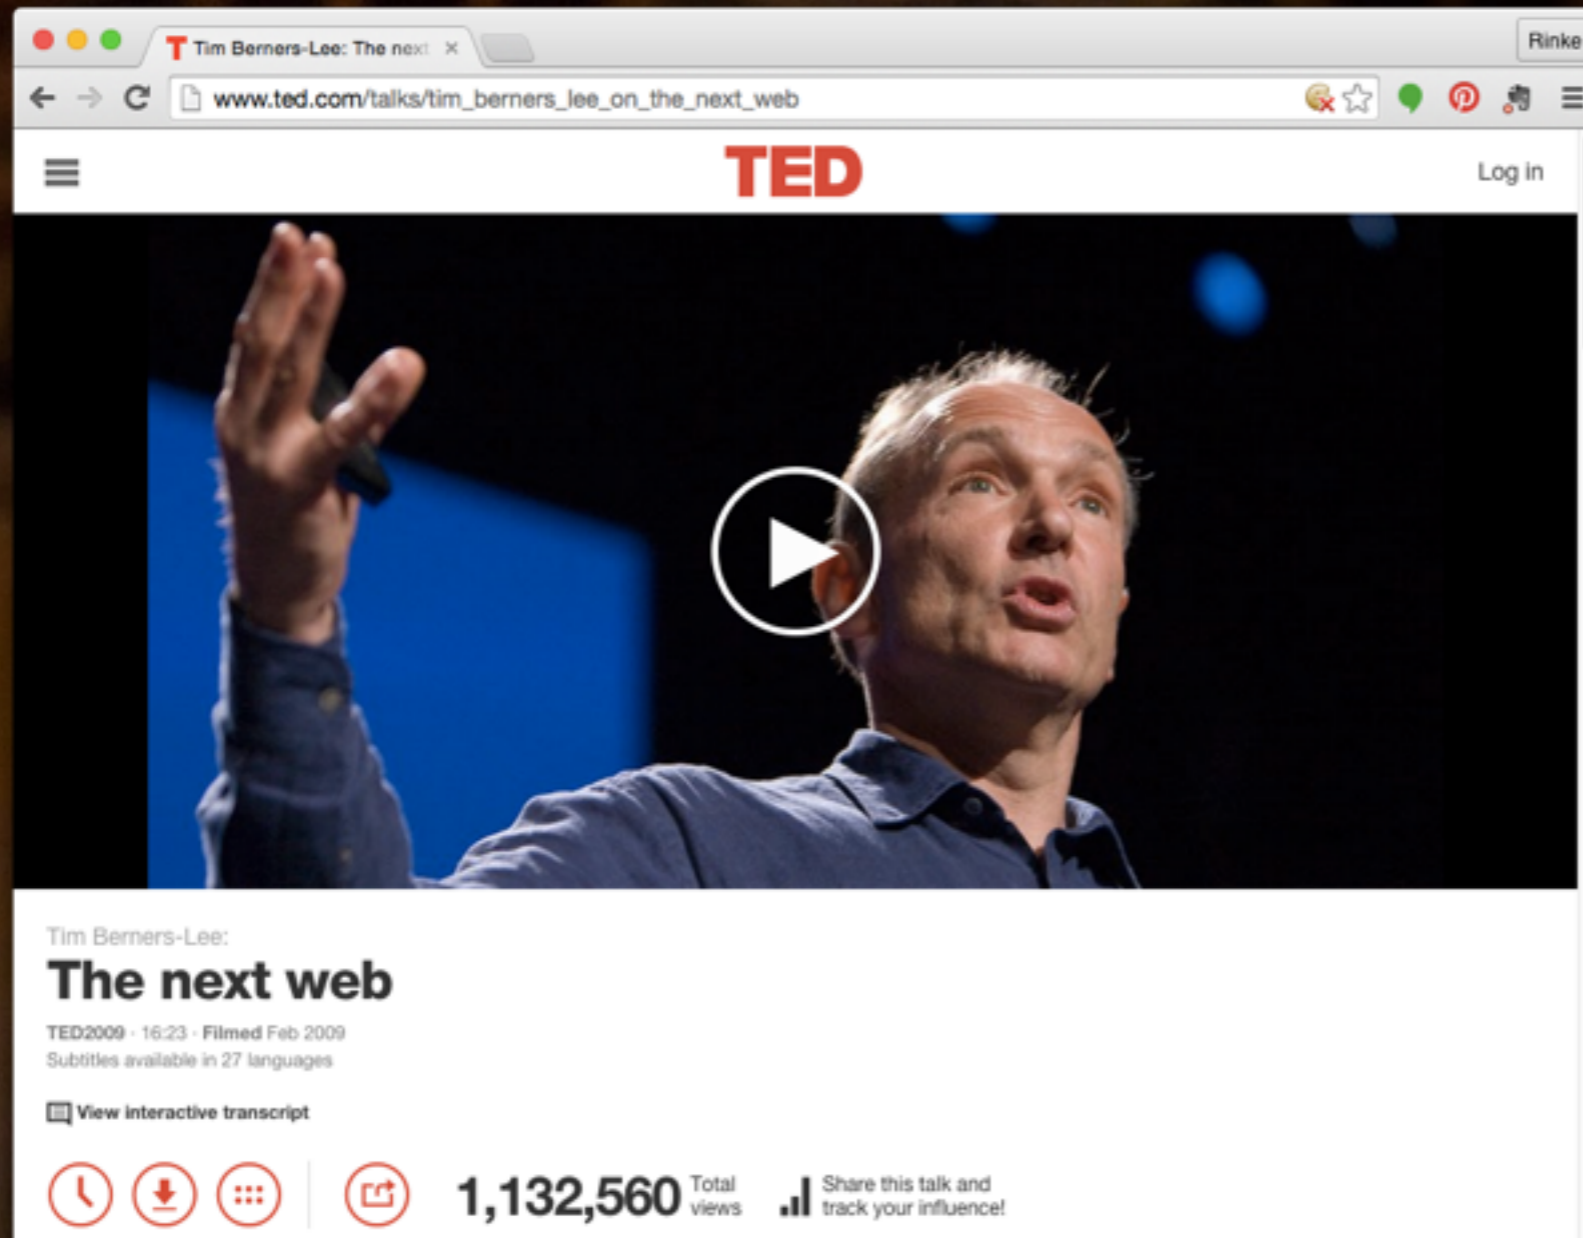
Databases only prescribe **data structure**

Data is **accessible** but still locked in **silos**

Data ownership \approx power

Data **integration** is a major issue

The Vision for a **Web of Data**



The screenshot shows a web browser window displaying a TED talk. The browser's address bar shows the URL www.ted.com/talks/tim_berniers_lee_on_the_next_web. The TED logo is visible in the top right of the page. The video player shows a play button over a still image of Tim Berners-Lee speaking. Below the video, the title "The next web" is displayed, along with the speaker's name "Tim Berners-Lee:". Additional information includes "TED2009 · 16:23 · Filmed Feb 2009" and "Subtitles available in 27 languages". There is a link to "View interactive transcript". At the bottom, there are icons for clock, download, and share, followed by the text "1,132,560 Total views" and "Share this talk and track your influence!".



The four **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** using RDF
4. Include **links to other URIs**, so that they can **discover** more things



The four **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** using RDF
4. Include **links to other URIs**, so that they can **discover** more things

Um... what's RDF?



Anybody can say Anything about Anything




The screenshot shows a YouTube video player interface. The main video is titled "Gov 2.0 Expo 2010: Tim Berners-Lee, 'Open, Linked Data for a Global Community'". The video is from the channel "O'Reilly" and has 22,732 views. The video player shows a man (Tim Berners-Lee) speaking on a stage with a "GOV 2.0 EXPO" backdrop. To the right of the video player is a "Up next" list of related videos, including "Gov 2.0 Expo 2010: Kathy 'Creating Passionate Citiz" by O'Reilly (4,394 views), "Gov 2.0 Expo 2010: Gary Vaynerchuk, 'The" by O'Reilly (4,581 views), "Jack Dangermond Intervie Gov 2.0" by O'Reilly (2,961 views), "Gov 2.0 Expo 2010: Tim O 'Government as a Platfor" by O'Reilly (10,570 views), "Tim Berners-Lee: The nex of open, linked data" by TED (156,713 views), "What is Linked Data?" by Manu Sporny (37,952 views), and "The Semantic Web of Data Berners-Lee" by dilshod2006 (30,190 views).



The **Resource Description Framework** (RDF)

- RDF is a standard **data model** for data interchange on the Web
- It facilitates **data merging** even if the underlying schemas/models differ.
- It extends the **linking structure** of the Web to use URIs to name the **relationship** between things, as well as the two ends of the link (a **triple**)
- It allows data to be **mixed**, **exposed** and **shared** across different applications
- The links form a directed, labeled graph: a **web of data**



RDF - Triples

- All information in RDF is expressed as **triples**; two-placed predicates; relations
- A triple consists of a **subject**, a **predicate** and an **object**:

subject	predicate	object
The Netherlands	has capital	Amsterdam
Amsterdam	has mayor	Eberhard van der Laan
Eberhard van der Laan	birth year	1955

- Another word for a triple is a **statement** or a **fact**
- The elements of an RDF triple are either **URI references**, **blank nodes**, or, **literals**.

RDF - Uniform Resource Identifiers (URIs)

- The Resource Description Framework talks about **resources**
(almost anything is a resource)
- Resources are **identified by** URIs, or URIs **denote** resources
(URIs can only refer to a resource, they **are not** the resource, and **multiple** URIs can denote the **same** resource)

The Netherlands http://dbpedia.org/resource/The_Netherlands

has capital <http://dbpedia.org/ontology/capital>

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

has mayor <http://dbpedia.org/ontology/leaderName>

Eberhard van der Laan http://dbpedia.org/resource/Eberhard_van_der_Laan

- Internationalised Resource Identifiers are URIs that allow unicode characters

RDF - URIs and CURIES (or QNames)

- URIs are often long and hard to read and write
- RDF **syntaxes** often use an abbreviation mechanism: **namespaces** and **prefixes**

```
@prefix dbpedia:    <http://dbpedia.org/resource/>  
@prefix dbo:       <http://dbpedia.org/ontology/>
```

- We can then map **compact URIs** (CURIs) to **full URIs**

```
dbpedia:The_Netherlands http://dbpedia.org/resource/The\_Netherlands  
dbo:capital              http://dbpedia.org/ontology/capital
```

RDF - URIs and Data

- We can now state that the capital of The Netherlands is Amsterdam

http://dbpedia.org/resource/The_Netherlands <http://dbpedia.org/ontology/capital> <http://dbpedia.org/resource/Amsterdam> .

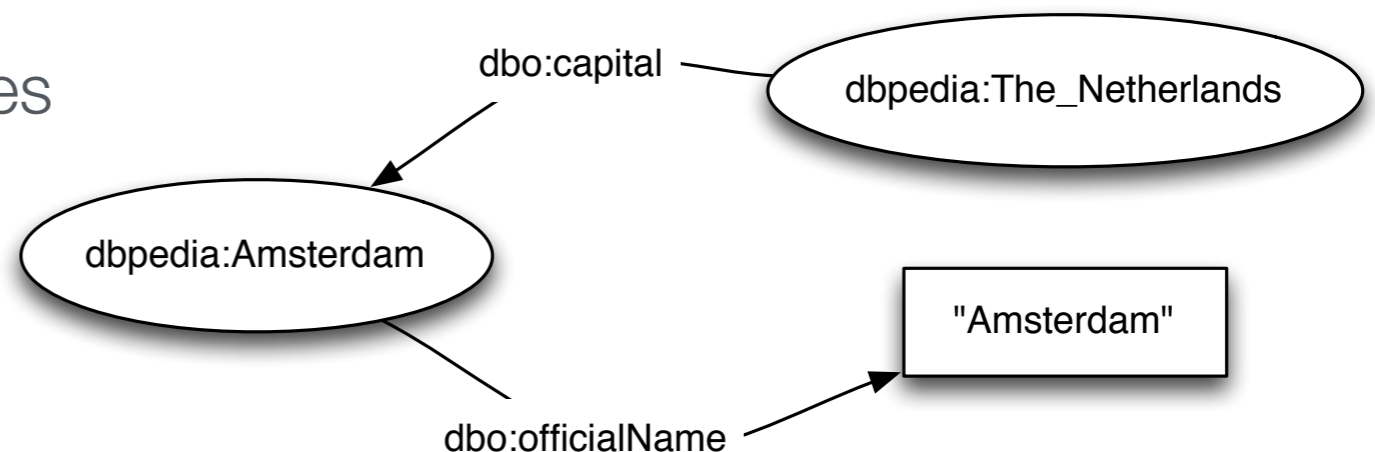
- Or using prefixes

dbpedia:The_Netherlands dbo:capital dbpedia:Amsterdam .

- An **RDF graph** is a **set of triples**, e.g.:

dbpedia:The_Netherlands dbo:capital dbpedia:Amsterdam .
dbpedia:Amsterdam dbo:officialName "Amsterdam" .

... is a graph that contains two triples



RDF - Literals

- **Literals** are used to represent "literal" data values

- All literals have a **datatype**

- Datatypes are also **resources**, referenced via URIs, and written as:

```
dbpedia:Amsterdam dbo:areaTotal "21932000"^^xsd:double.
```

- Default: if no datatype is specified, the datatype is assumed to be `xsd:string`

```
dbpedia:Amsterdam dbo:officialName "Amsterdam".
```

- One can specify the **language** of a string using a **language tag**:

```
dbpedia:The_Hague rdfs:label "Den Haag"@nl.
```

```
dbpedia:The_Hague rdfs:label "The Hague"@en.
```

RDF - Why HTTP URIs?

- HTTP URIs have a **global scope**, unique throughout the Web
(c.f. e.g. keys in relational databases which are only unique within a table)
 - Helps to avoid **name clashes** <http://abc-co.com/category/item/123>
<http://xyz-co.com/product/123>
 - Grounded in **society** (DNS registration)
- HTTP URIs are also **addresses** that indicate **authority** and enable **trust**
 - Exploit the well-functioning machinery of Web **browsing**
 - Track data by **following** the resource identifiers found in triples

URIs are Addresses



The screenshot shows a web browser window with the address bar containing `dbpedia.org`. The page title is "About: Amsterdam". Below the title, it states: "An Entity of Type : List of municipalities of the Netherlands, from Named Graph : <http://dbpedia.org>, within Data Space : dbpedia.org".

The main text describes Amsterdam: "Amsterdam (English /'æmstərdæm/ or /,æmstər'dæm/; Dutch: [ˌʔɑmstər'dɑm]) is the capital city and most populous city of the Kingdom of the Netherlands. Its status as the Dutch capital is mandated by the Constitution of the Netherlands though it is not the seat of the Dutch government, which is The Hague. Amsterdam has a population of 825,080 within the city proper, 1,317,663 in the urban area and 1,590,520 in the metropolitan area. The city region has an approximate population of 2,431,000."

Property	Value
dbo:PopulatedPlace/areaTotal	• 219.32

Amsterdam (English /'æmstərdæm/ or /,æmstər'dæm/; Dutch: [ˌʔɑmstər'dɑm]) is the capital city and most populous city of the Kingdom of the Netherlands. Its status as the Dutch capital is mandated by the Constitution of the Netherlands though it is not the seat of the Dutch government, which is The Hague. Amsterdam has a population of 825,080 within the city proper, 1,317,663 in the urban area and 1,590,520 in the metropolitan area. The city region has an approximate population of 2,431,000. The city is located in the province of North Holland in the west of the country, and is also North Holland's largest city. It comprises much of the northern part of the Randstad, one of the larger conurbations in Europe, with a population of approximately 7 million. Amsterdam's name derives from Amstelredamme, indicative of the city's origin as a dam of the river Amstel. Originating as a small fishing village

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

RDF - Why triples?

- Any information format can be **transformed** to triples

Tables:	row	column	cell
Trees:	parent	path	child

- Relationships are made **explicit**; they are elements in their own right
 - Unlike **database columns** and **binary predicates**
 - The predicate is an element in the triple, and **can be described** in RDF
 - "Self-documenting"

From **Tables** to **Graphs** ...

surname	age	occupation	sex
Fumes	45	cigar maker	female
Bridges	32	civil engineer	female
Moves	17	dancer	male

From **Tables** to **Graphs** ...

- Tables and graphs are both “just” **representation formats**

surname	age	occupation	sex
Fumes	45	cigar maker	female
Bridges	32	civil engineer	female
Moves	17	dancer	male

From **Tables** to **Graphs** ...

- Tables and graphs are both “just” **representation formats**

surname	age	occupation	sex
Fumes	45	cigar maker	female
Bridges	32	civil engineer	female
Moves	17	dancer	male



From **Tables** to **Graphs** ...

- Tables and graphs are both “just” **representation formats**

surname	age	occupation	sex
Fumes	45	cigar maker	female
Bridges	32	civil engineer	female
Moves	17	dancer	male

- Each **row** is a **record** *about* something



From **Tables** to **Graphs** ...

- Tables and graphs are both “just” **representation formats**


surname	age	occupation	sex
Fumes	45	cigar maker	female
Bridges	32	civil engineer	female
Moves	17	dancer	male

- Each **row** is a **record** *about* something
- Each **column** is an **attribute** of that thing



From **Tables** to **Graphs** ...

- Tables and graphs are both “just” **representation formats**



ID	surname	age	occupation	sex
1	Fumes	45	cigar maker	female
2	Bridges	32	civil engineer	female
3	Moves	17	dancer	male

- Each **row** is a **record** *about* something
- Each **column** is an **attribute** of that thing
- A **primary key** identifies a record (local)



From **Tables** to **Graphs** ...

- Tables and graphs are both “just” **representation formats**

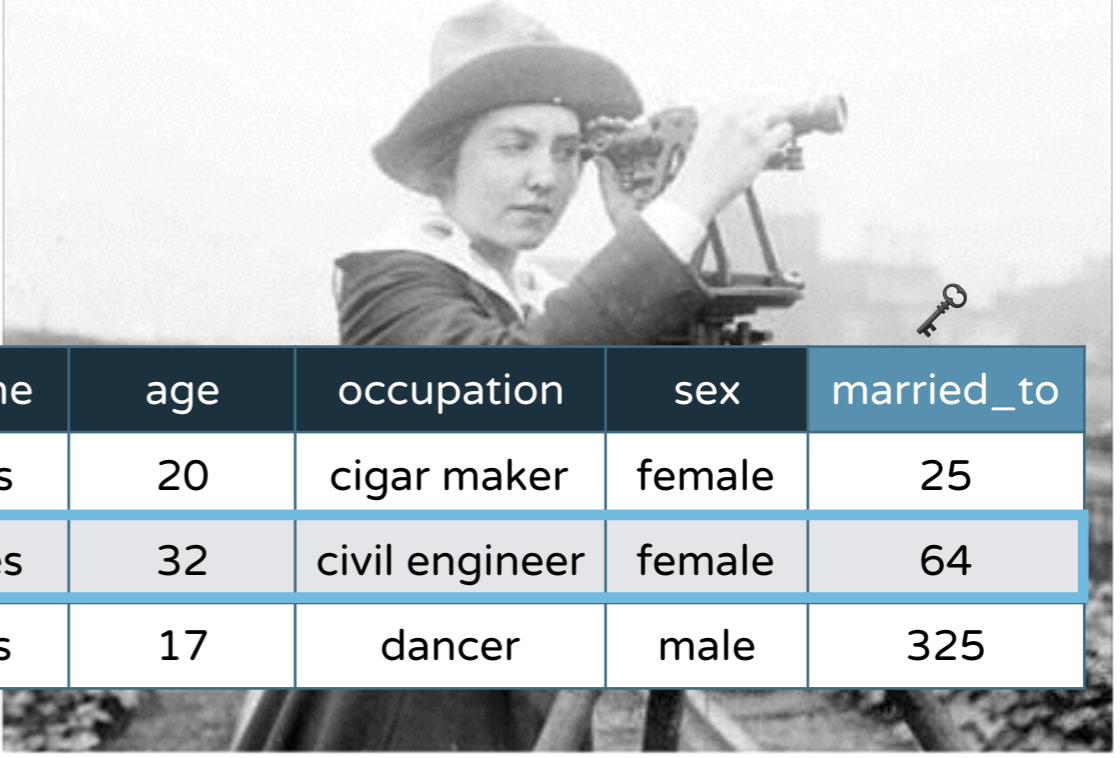
ID	surname	age	occupation	sex	married_to
1	Fumes	20	cigar maker	female	25
2	Bridges	32	civil engineer	female	64
3	Moves	17	dancer	male	325

- Each **row** is a **record** *about* something
- Each **column** is an **attribute** of that thing
- A **primary key** identifies a record (local)
- A **secondary key** can identify an (external) record



From **Tables** to **Graphs**... (the simple version)

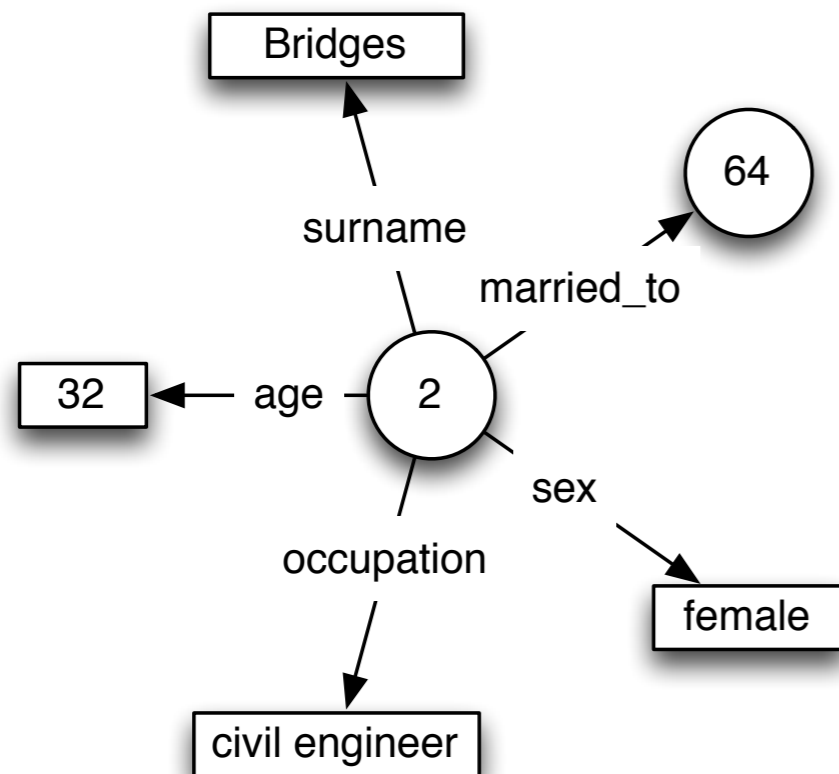
- Each **node** represents a **thing** (*resource*) or **value** (*literal*)
- Each **edge** represents an **attribute** (*property*)



ID	surname	age	occupation	sex	married_to
1	Fumes	20	cigar maker	female	25
2	Bridges	32	civil engineer	female	64
3	Moves	17	dancer	male	325

From **Tables** to **Graphs**... (the simple version)

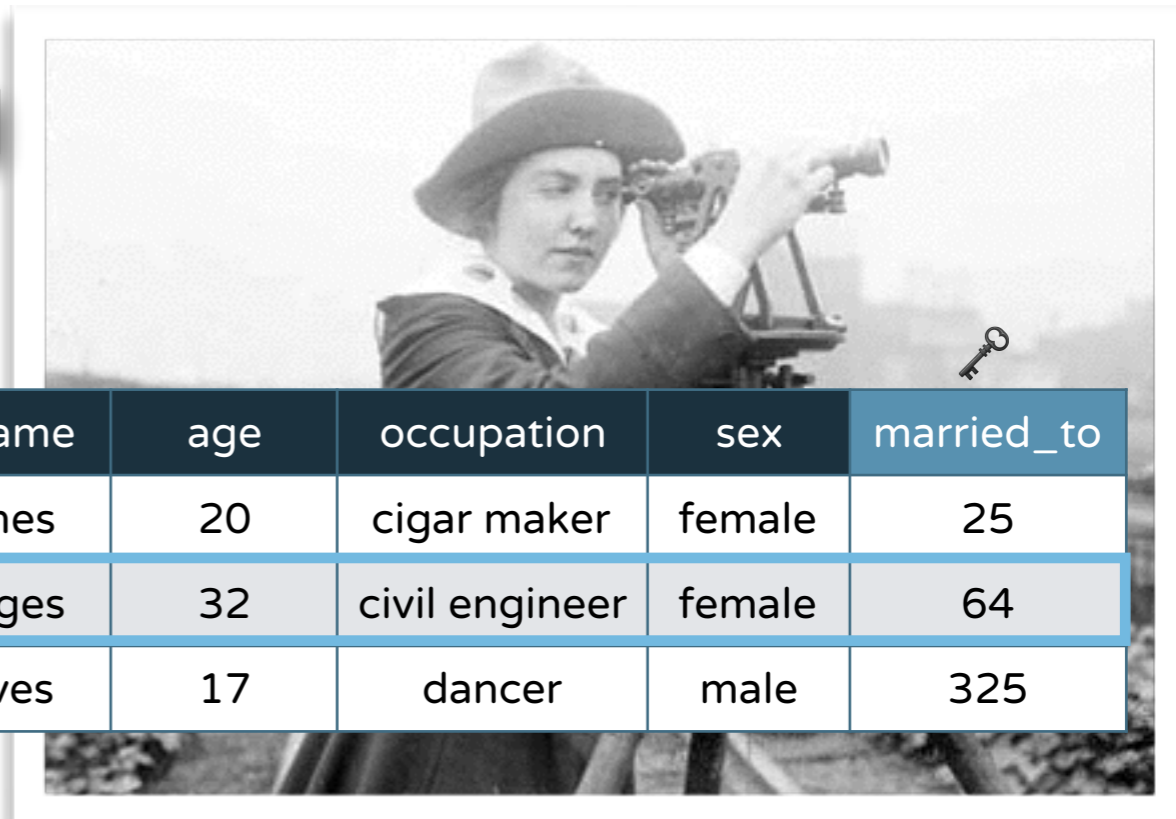
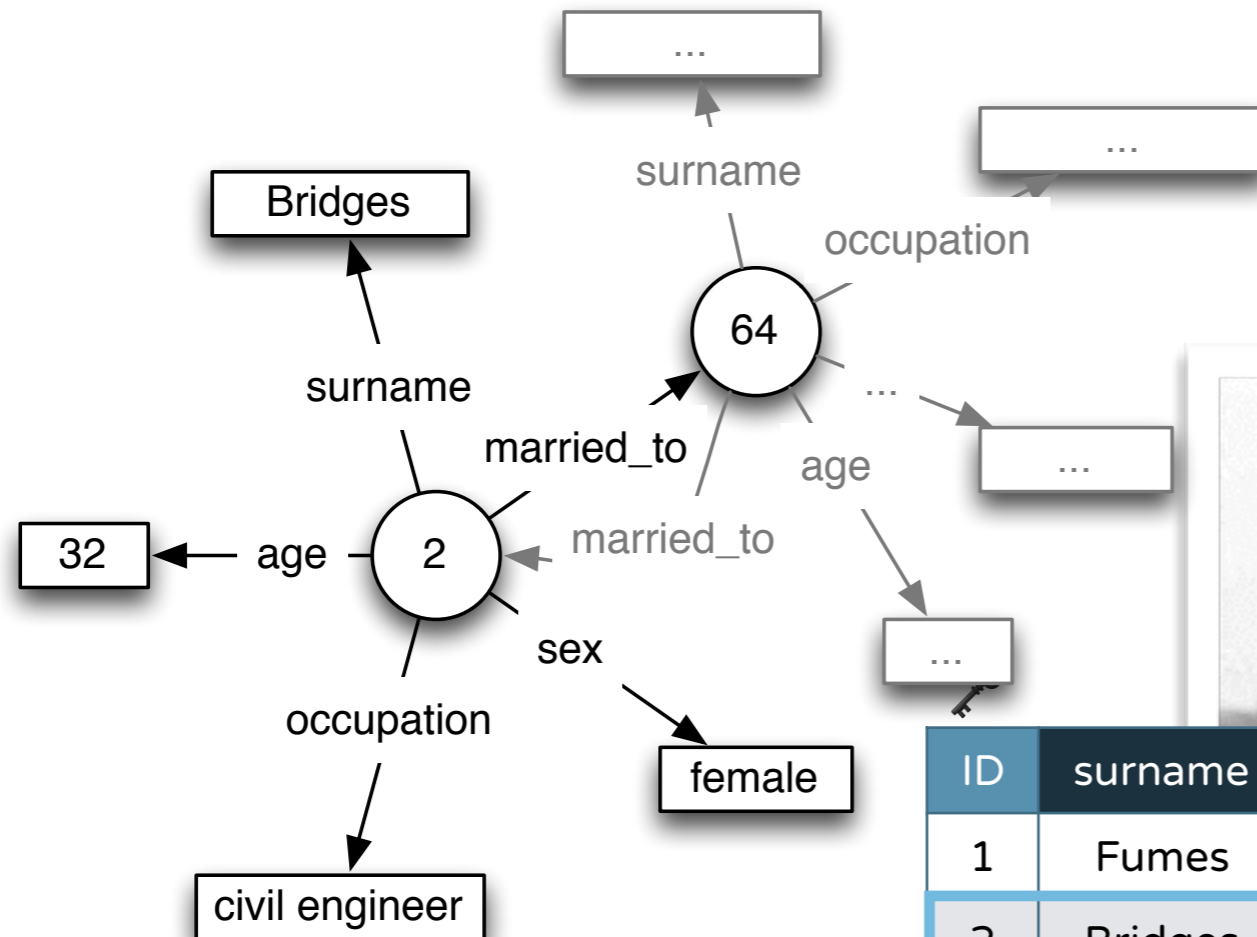
- Each **node** represents a **thing** (*resource*) or **value** (*literal*)
- Each **edge** represents an **attribute** (*property*)



ID	surname	age	occupation	sex	married_to
1	Fumes	20	cigar maker	female	25
2	Bridges	32	civil engineer	female	64
3	Moves	17	dancer	male	325

From **Tables** to **Graphs**... (the simple version)


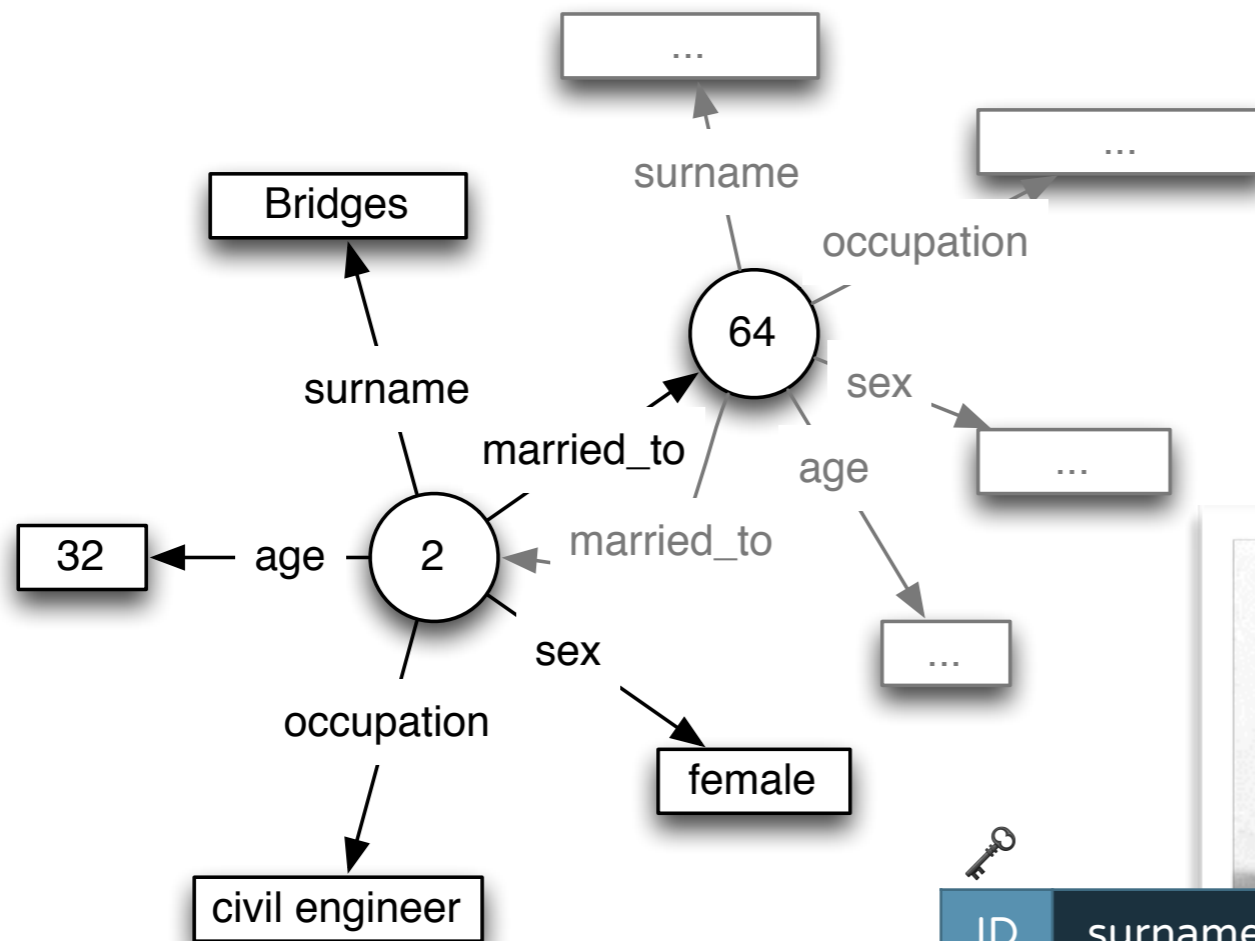
- Each **node** represents a **thing** (*resource*) or **value** (*literal*)
- Each **edge** represents an **attribute** (*property*)



ID	surname	age	occupation	sex	married_to
1	Fumes	20	cigar maker	female	25
2	Bridges	32	civil engineer	female	64
3	Moves	17	dancer	male	325

From **Tables** to **Graphs**... (the “better” version)

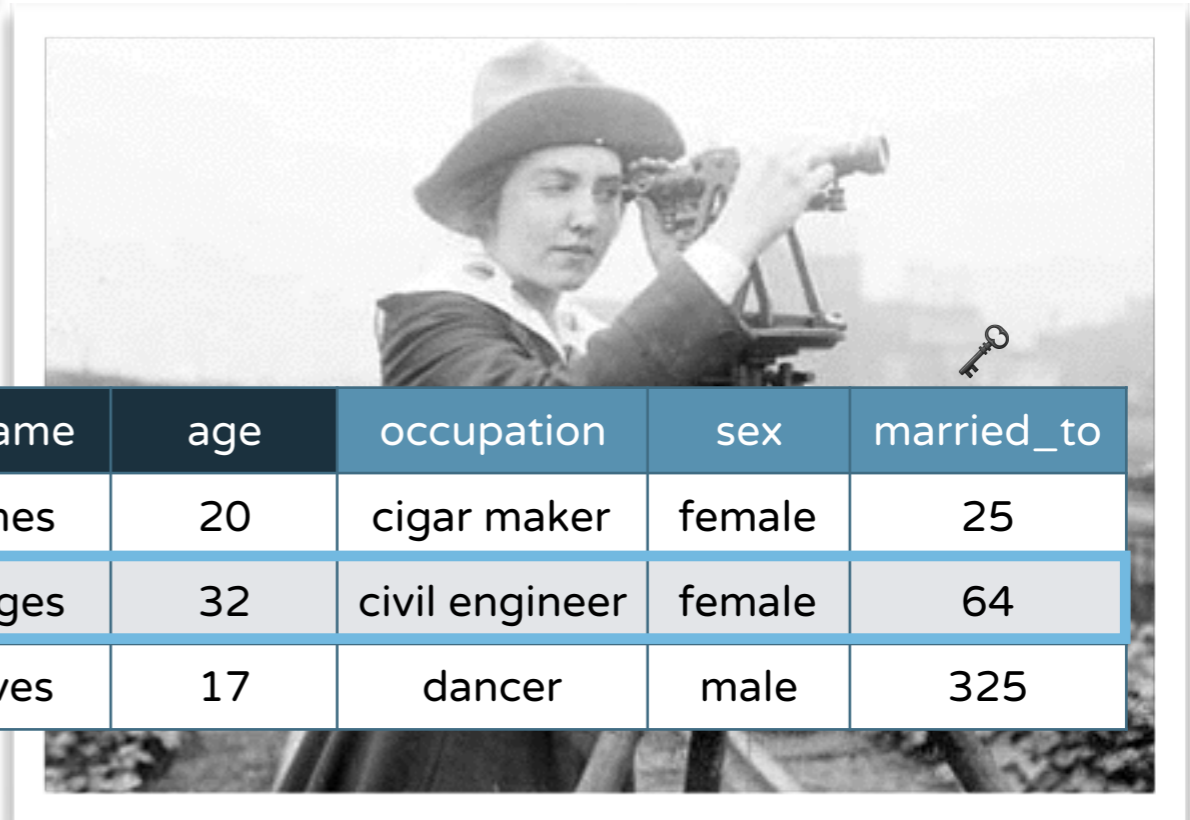
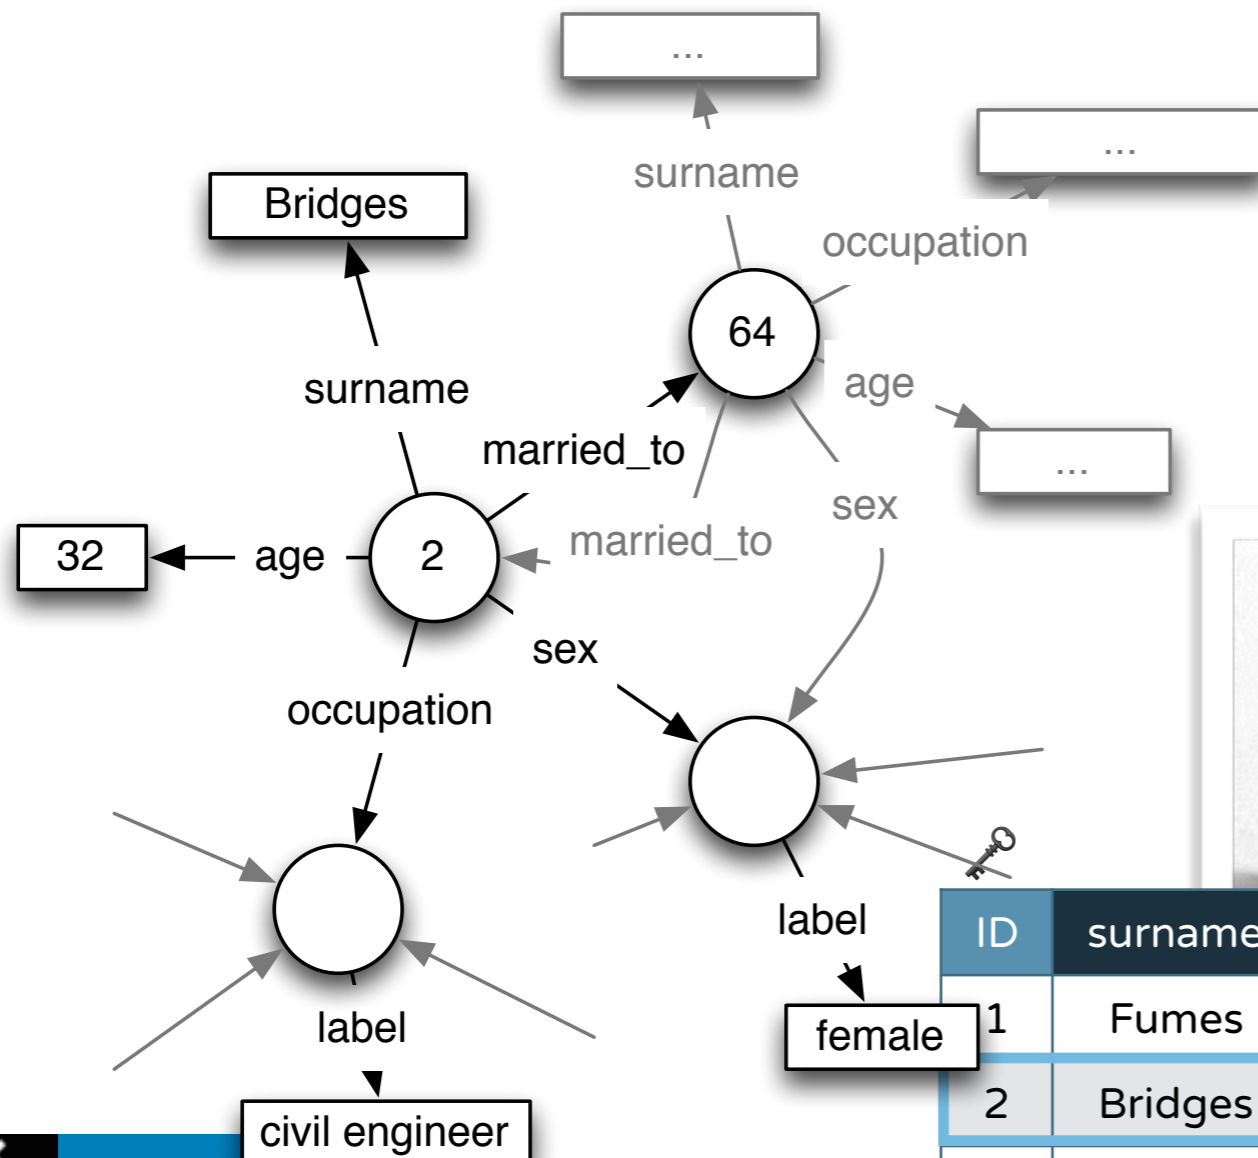
- Some **values** are actually **keys**



ID	surname	age	occupation	sex	married_to
1	Fumes	20	cigar maker	female	25
2	Bridges	32	civil engineer	female	64
3	Moves	17	dancer	male	325

From **Tables** to **Graphs**... (the “better” version)

- Some **values** are actually **keys**



ID	surname	age	occupation	sex	married_to
1	Fumes	20	cigar maker	female	25
2	Bridges	32	civil engineer	female	64
3	Moves	17	dancer	male	325

SPARQL - Query Syntax

PREFIX: the namespace prefixes used in the SPARQL query

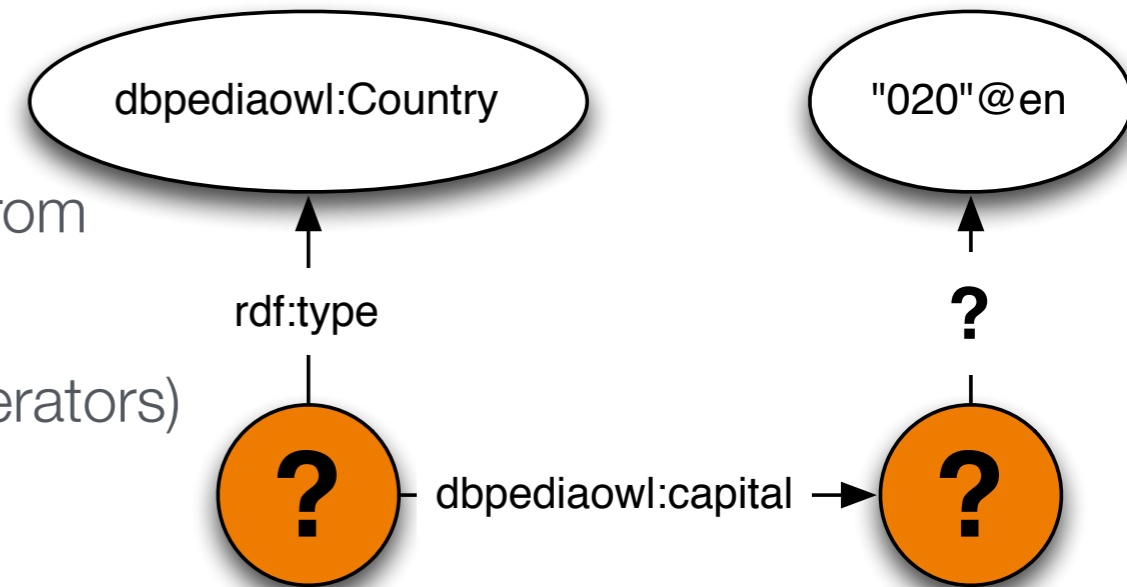
```
PREFIX dbo: <http://dbpedia.org/ontology/>

SELECT ?city WHERE {
  ?city dbo:areaCode "020" .
} LIMIT 10
```

SELECT: the entities (variables) you want to return

WHERE: the (sub)graph you want to get information from

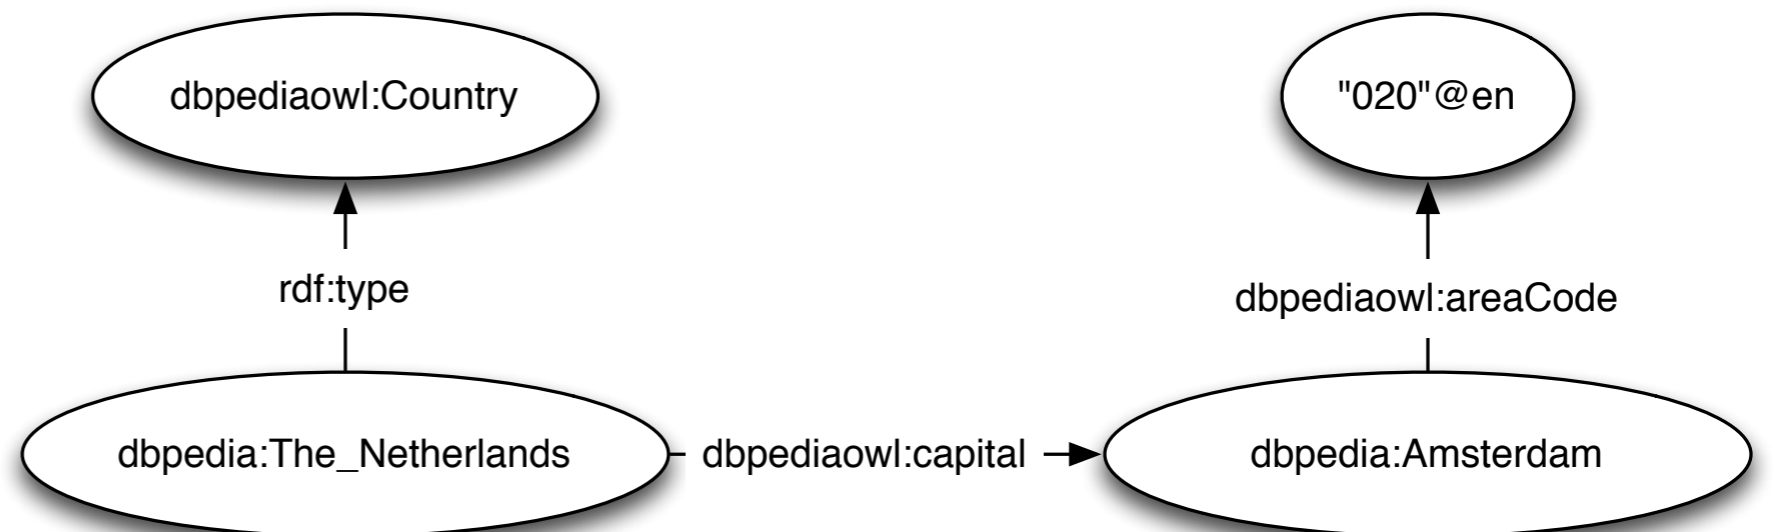
... including additional constraints on results (using operators)



SPARQL - Triple Patterns

- A **graph pattern** consists of multiple **triple patterns**
- A **triple pattern** is a triple with **zero** or **more** variables

```
?x dbo:capital dbpedia:Amsterdam .  
?x dbo:capital ?y .  
?x dbo:areaCode "020" .  
?x ?p ?y
```



SPARQL - Triple Patterns

- A **graph pattern** consists of multiple **triple patterns**
- A **triple pattern** is a triple with **zero** or **more** variables

```
?x dbo:capital dbpedia:Amsterdam .  
?x dbo:capital ?y .  
?x dbo:areaCode "020" .  
?x ?p ?y
```

```
dbpedia:Netherlands rdf:type dbo:Country ;  
dbpedia:Amsterdam dbo:capital dbpedia:Amsterdam .  
dbpedia:Amsterdam dbo:areaCode "020" .
```

dbpediaowl:Country

"020"@en

dbpedia:The_Netherlands

dbpediaowl:capital

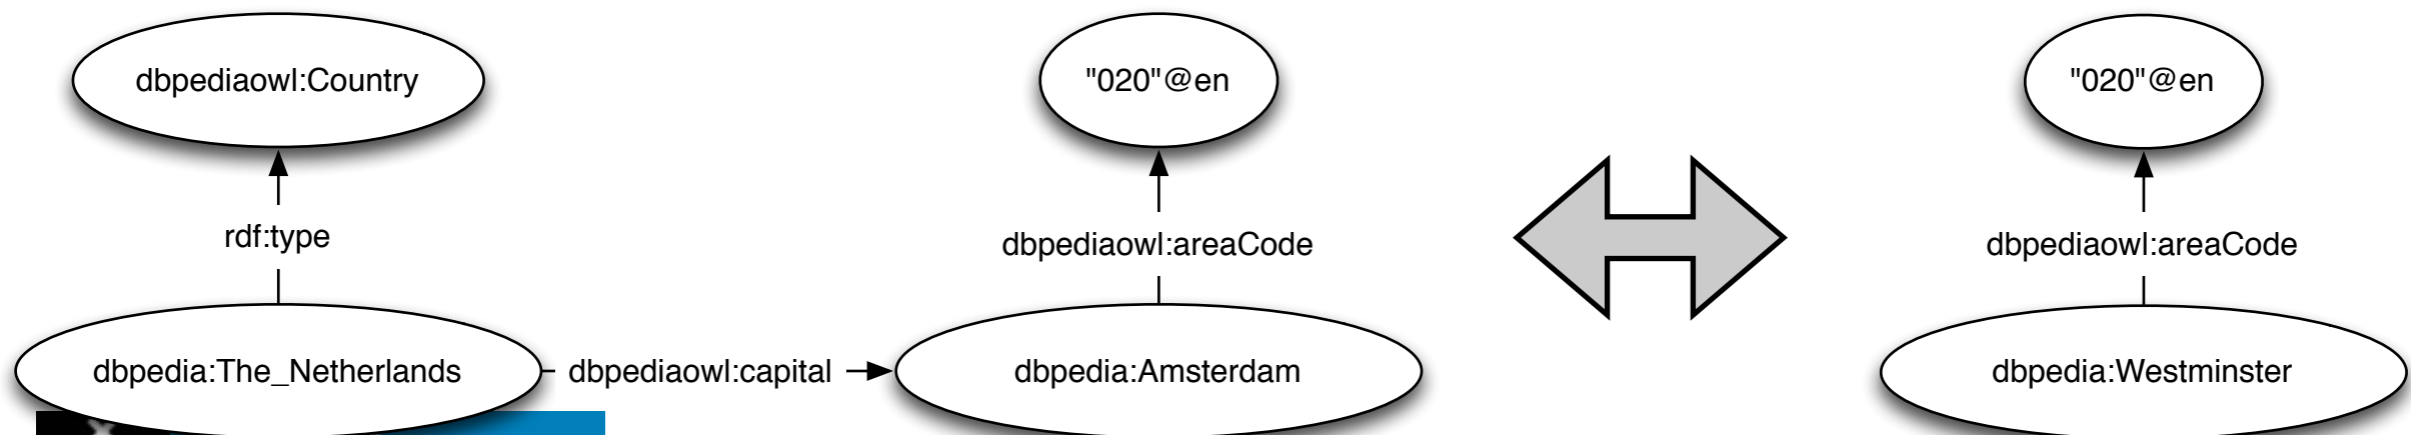
dbpedia:Amsterdam

SPARQL - Triple patterns form a **conjunction**

Every triple pattern in the graph pattern should **match**

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

SELECT ?x WHERE {
  ?x      dbo:capital      ?y .
  ?y      dbo:areaCode    "020" .
} LIMIT 10
```

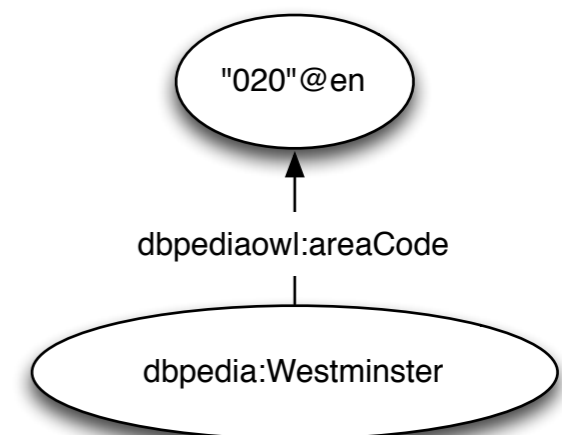

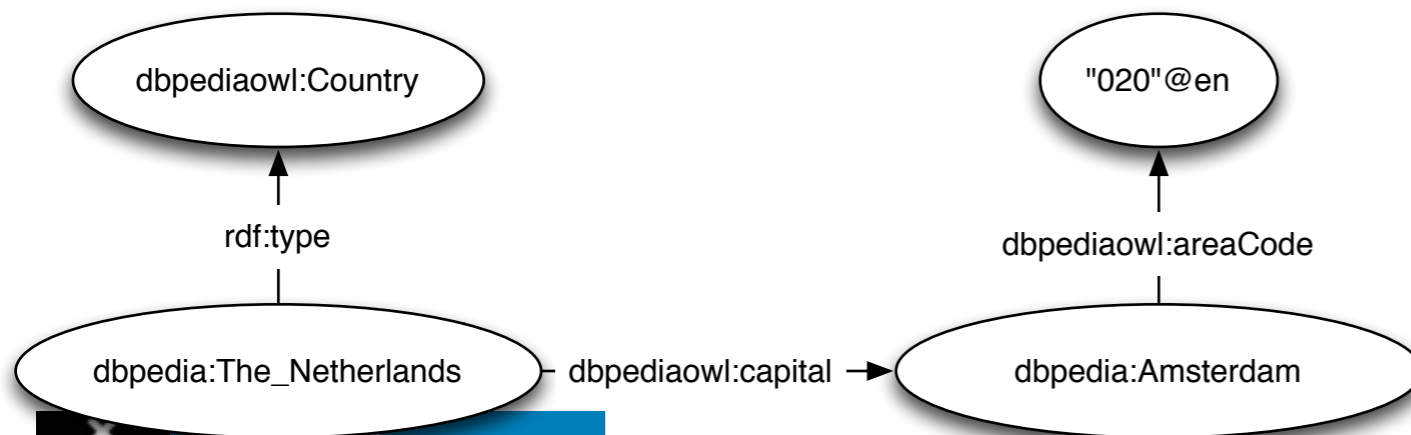
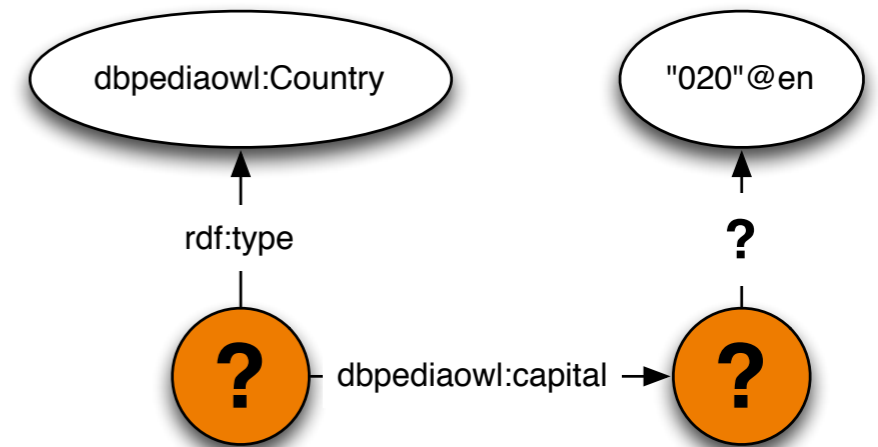


SPARQL - Triple patterns form a **conjunction**

Every triple pattern in the graph pattern should **match**

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

SELECT ?x WHERE {
  ?x      dbo:capital      ?y .
  ?y      dbo:areaCode    "020" .
} LIMIT 10
```

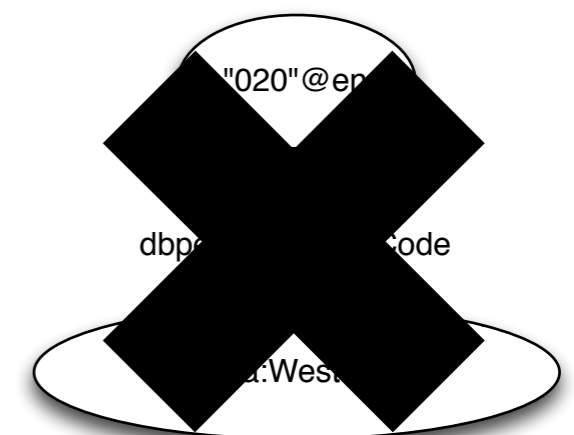

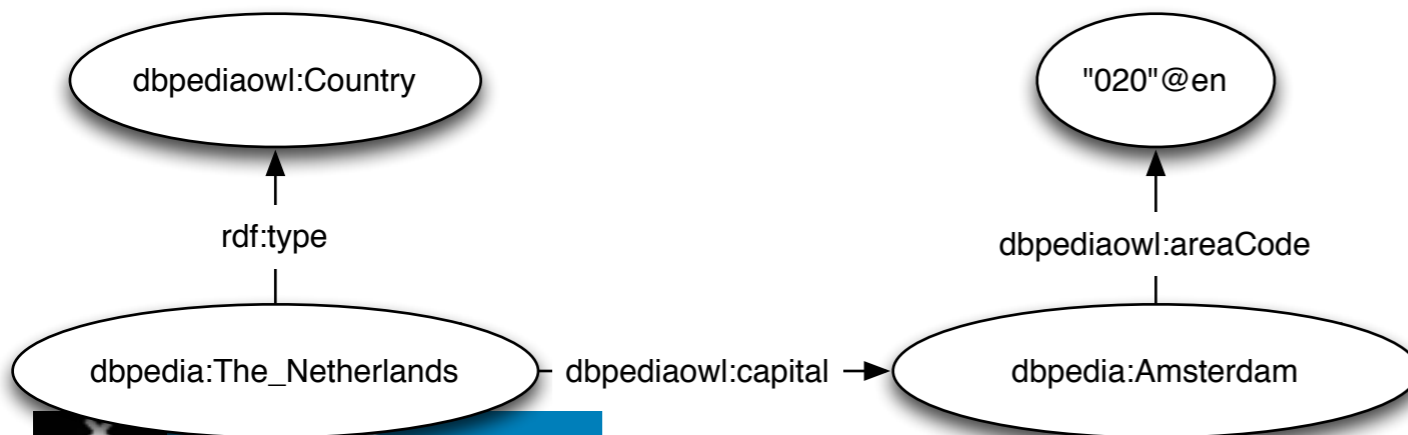
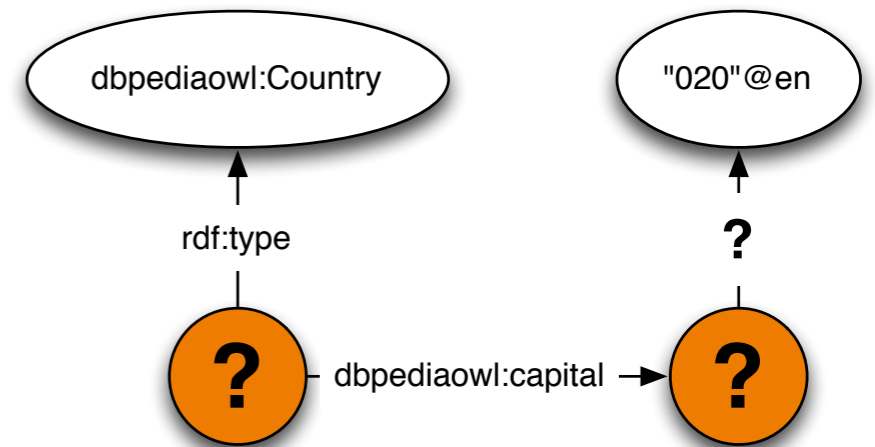


SPARQL - Triple patterns form a **conjunction**

Every triple pattern in the graph pattern should **match**

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

SELECT ?x WHERE {
  ?x      dbo:capital      ?y .
  ?y      dbo:areaCode    "020" .
} LIMIT 10
```



<http://www.w3.org/TR/sparql11-query>

YASGUI x Rinke

yasgui.org

Query 1 x +

http://dbpedia.org/sparql

```
1 PREFIX dbo: <http://dbpedia.org/ontology/>
2 SELECT ?x WHERE {
3   ?x dbo:capital ?y .
4   ?y dbo:areaCode "020" .
5 }
6 LIMIT 10
```

Table Raw Response Pivot Table Google Chart

Showing 1 to 10 of 10 entries (in 0.028 seconds) Search: Show 50 entries

x
1 http://dbpedia.org/resource/Dutch_government-in-exile
2 http://dbpedia.org/resource/Netherlands
3 http://dbpedia.org/resource/Dutch_Republic
4 http://dbpedia.org/resource/United_Kingdom_of_the_Netherlands
5 http://dbpedia.org/resource/Reichskommissariat_Niederlande

38.606.408.854 triples and counting!



LOD Laundromat

The LOD Laundromat provides access to all Linked Open Data (LOD) in the world. It does this by crawling the LOD cloud, and converting all its contents in a standards-compliant way (gzipped N-Triples), removing all data stains such as syntax errors, duplicates, and blank nodes.

Laundry Basket

The LOD Laundry Basket contains the URLs of dirty datasets that are waiting to be cleaned by the [LOD Washing Machine](#). This is

