ECE_5DA04_TP

Big Graph Databases

Garima Gaur, Ioana Manolescu,

Madhulika Mohanty

INRIA & Ecole Polytechnique

firstname.lastname@inria.fr

DataAI Master

Institut Polytechnique de Paris

Course goal

1. Discuss the main characteristics (dimensions) of Big Database systems, focused around graphs.
2. Present the main concepts underlying such systems.
3. Motivate and present architectural choices made to scale better.

Course organization

- **Instructors**:
  Ioana Manolescu (Inria and Ecole Polytechnique)
  ioana.manolescu@inria.fr
  Madhulika Mohanty (Inria and Ecole Polytechnique)
  madhulika.mohanty@inria.fr
  Garima Gaur (Inria and Ecole Polytechnique)
  garima.gaur@inria.fr

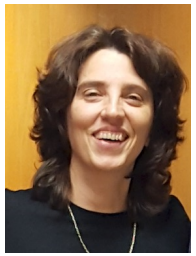

- **Evaluation**: final exam (50%) + lab work (50%)

- All course material in Moodle (Self enrol):
  https://moodle.r2.enst.fr/moodle/course/view.php?id=162

SCAN ME

What characterizes a Big Database system?

- Functionality provided
  - What kind of data can I put in? **Data model**
  - How can I get data out of it? **Query language or API**
  - How does it handle concurrent access? **Concurrency control**
  - How long does a given operation take? **Performance**
- Implementation (internals)
  - How does it cope with scale?
    - for reads? **Data and work distribution**
    - also for writes? **Distributed concurrency control**

Topics to be covered

- Recall/crash course on **relational database management systems (RDBMS)**
  - Data model, query language, performance
- **Graph Data Model**
  - Different data models
  - Query Languages
    - Structured, Semi-structuctured, Unstructured
- **Distributed databases**
- **Reasoning on knowledge**
- **Heterogeneous data integration systems**
  - Local-as-view, global-as-view
  - Mediator systems
  - Dataspaces
  - Data lakes
- Massive parallelism seen in another course (Hadoop, Spark)

What about NoSQL?

1. NoSQL is mostly about distributed concurrency control
   - **Weaker guarantees** than a centralized RDBMS...
   - In exchange for **better performance at scale**
2. NoSQL systems also typically have **other data models and languages** than relational/SQL.
   - Key-values (Redis), JSON store (MongoDB), graphs (Neo4J, Virtuoso)

Today's course plan

1. **Motivation**: Big Data
   - Characteristics
   - Applications
2. **From databases to architectures for Big Graph Data management**
   - Database management system: quick recall (or crash course)
   - What needs to change to handle Big Data?

MOTIVATION: BIG DATA

Defining Big Data: the V's

- Volume
  - Scale
- Velocity
  - Speed of producing and consuming the data
- Variety
  - Very different sources and data types
- Veracity
  - Is the data correct / certain / true?

Where does the data volume come from?

- Human-produced data
  - Web content: Web pages, blogs, social networks, tweets…
  - X: 7 Terabytes (1 tera = 10^{18}) per day
  - Facebook: 10 Terabytes per day

Where does the data volume come from? (1)

- Human-produced data
  - Web content: Web pages, blogs, social networks, tweets…
  - Twitter: 7 Terabytes (1 tera = 10^{18}) per day
  - Facebook: 10 Terabytes per day
- Machine-produced data
  - Log data from all kind of servers
  - Real world devices: banks, telecom, energy, weather, transportation, shipment…
  - Sensors, including on highways
  - and trains

Gazpar (GRDF)

Linky (EDF)

Where does the data volume come from? (2)

- E.g. french railway system: surveillance trains for the normal and high-speed lines (TGV)
- TGV specially equipped for measuring while circulating at 320 km/h:
  - rail geometry
  - train/rail interaction
  - rail signalization and communication devices
  - electric power availability etc.
  - 150 sensors, 20 cameras

Gordon Bell, Microsoft, 2009

"It's like having a multimedia transcript of your life.

By about 2020 [...] our entire life histories will be online and searchable.

Location-aware smartphones and inexpensive digital memory storage in the "cloud" of the Internet make the transition possible and inevitable.

No one will have to fret about storing the details of their lives in their heads anymore. We'll have computers for that.

And this revolution will "*change what it means to be human*"

Huge data volumes lead to distributed storage

The typical architecture for large-scale distributed storage and computing is cloud-based

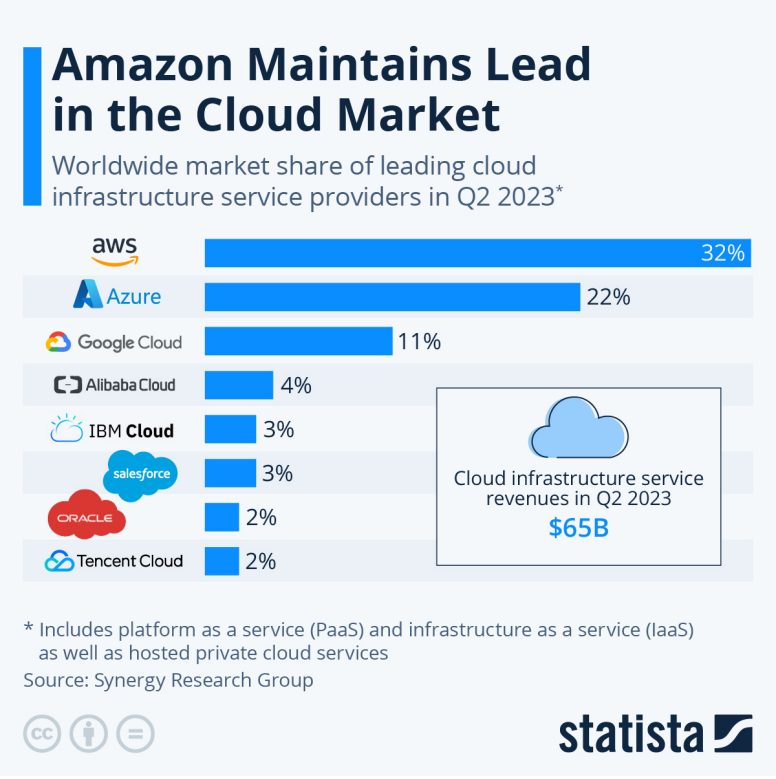

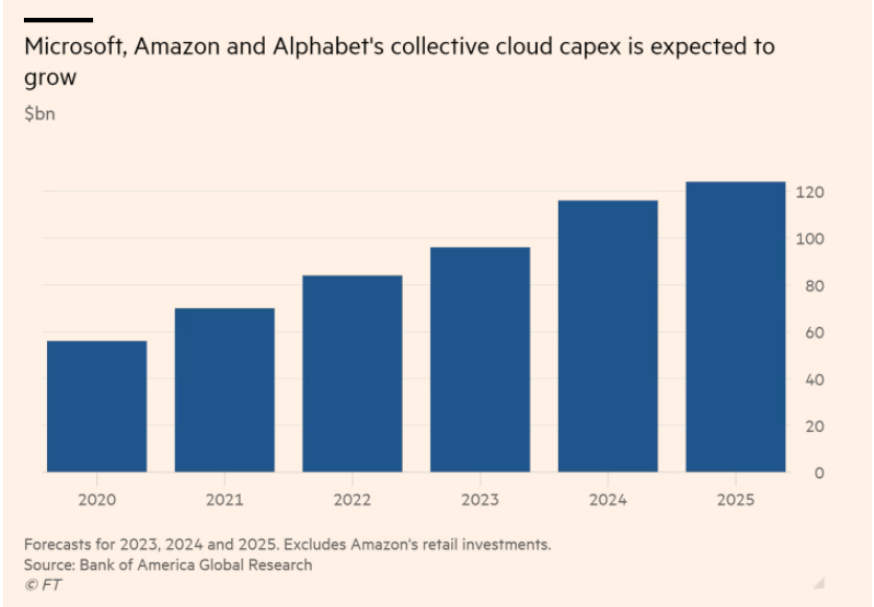

Defining Big Data: the V's

- Volume
  - Scale
- Velocity
  - Speed of producing and consuming the data
- Variety
  - Very different sources and data types
- Veracity
  - Is the data correct / certain / true?

Big Data velocity

- How much data is produced e.g. per second
- Data enters a **pipeline** consisting of storage and/or processing
  - Store-then-process: for off-line data analysis. Storage by itself is a challenge sometimes, e.g. data links to/from clouds are rather slow
  - Process-then-store: for data whose interest is maximized upon arrival (real-time processing)
  - Process-then-discard: sensor/monitoring (if nothing happens)

Sample high-throughput data streams

- French IT company runs a data center of 2000 servers
- 5000 energy efficiency indicators (temperature, electricity consumption etc.) are measured every 20 seconds x 50 Kb per measure result = 170 Terabytes / year

- Unable to store all data → sample (measure more rarely)
- May miss important things when they happen

The importance of current/recent data

- Real-time applications work only / mostly with **the latest data**
  - Embedded control mechanisms based on sensor data, e.g., "*this railway wagon component is breaking*" (now!)
  - Intrusion or malfunctioning detection…
- Keeping **humans engaged**
  - Customer relationship management *while the client is on the phone* with the customer service (see: "Ordering pizza in the future" video)
  - Recommending places where your friends are hanging out *now*

Important aspect of Big Data: ethics

- Current technologies for gathering and processing data raise **risks** wrt personal **freedoms** and **rights**
  - Discriminations, privacy violations, consumer rights, infringement on personal freedom (cf. geo-tagging), political manipulation (cf. focused FB ads)
  - Democracy and liberty may be at risk

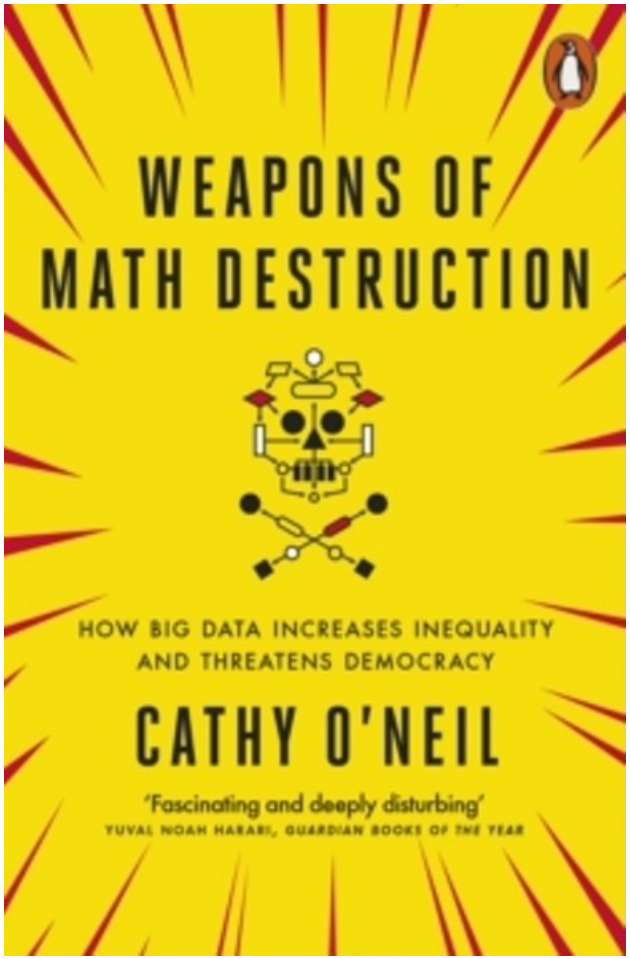

Data ethics problems with real consequences

Scenario 1: Anne carries a connected watch recording her movements in the city

- This makes it easy to know where her husband and children often go, even if they did not agree to sharing this info.

https://t.co/SyLsVTsFTv?amp=1

Scenario 2: Carol shares her DNA information with a DNA analysis company

- Then a health ensurance company buys it to learn that Carol's parents and children share a gene variant associated with an expensive-treatment illness

Defining Big Data: the V's

- Volume
  - Scale
- Velocity
  - Speed of producing and consuming the data
- Variety
  - Different sources, data formats, data types
- Veracity
  - Is the data correct / certain / true?

Big data heterogeneity (variety)

- Each new data type has added up on the old ones
- Enterprise data typically has high per-byte value ($/byte)
  - Hard to explain that "we will not need this database in the future"
- In many areas, legal obligation to keep old data (e.g. railway sensors, telecom, commercial…)

→ Data model & data management system soup

- hierarchical, relational, object-oriented, XML, RDF, JSON, key-value pairs…

Sample relational database

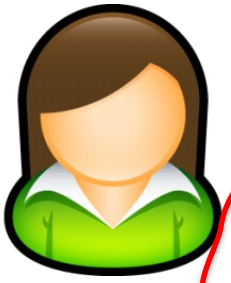

Clients

NumClient	Nom	Adresse	Ville	Age
1	Julie	1 rue Dugommier	Paris	22
…	…	…	…	…

Comptes

NumCompte	Type	Découvert	NumClient
12345	Courant	1000	1
…	…	…	…

Transaction

NumCompte	Montant	Date	Info
12345	-40,00	5/10/11	Retrait
12345	+23,45	6/10/11	Remb. MAIF
12345	-300,00	7/10/11	Chaussures

XML: extensible markup language

W3C, 2008

clients.xml:

```
<clients>
 <client><nom>Julie</nom>
  <address>1, rue Dugommier</address>
  <city>Paris</city><age>22</age>
 </client>
 <client><nom>Marc</nom>…
 </client>
</clients>
```

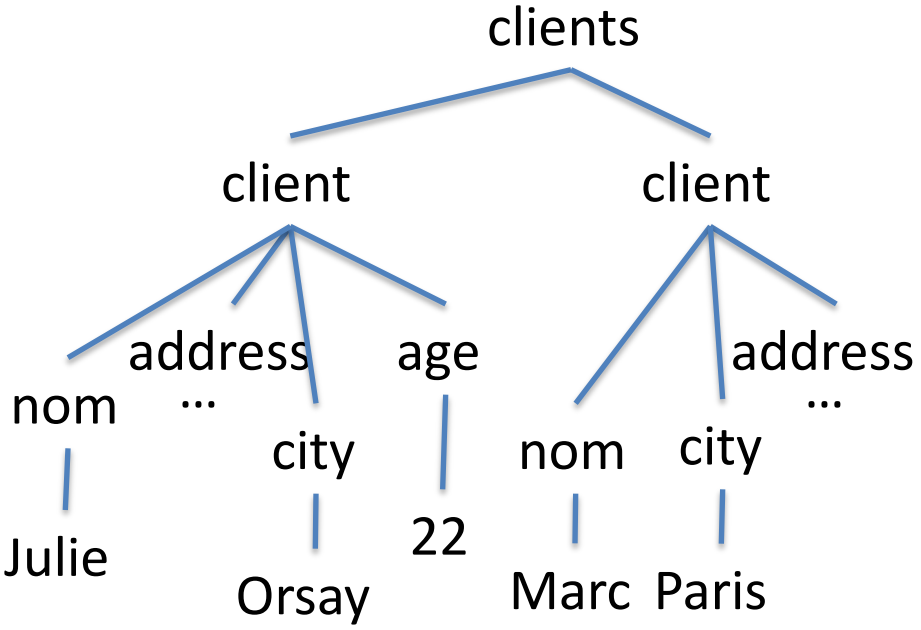

Flexible
Platform-independent
Separate content from presentation
Schema possible (not compulsory)

XML applications

- Main language for the Web: **XHTML, XML Schema, SVG, RSS,** …
- Web Services: **SOAP, WSDL**
- **MathML** (mathematical markup language)
- **CML** (chemical markup language)
- **SMILE** (synchronized multimedia integration language)
- Financial Exchange (**IFX**)
- The Text Encoding Initiative (**TEI**)

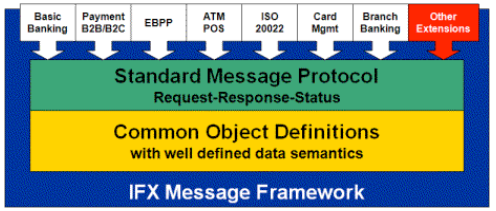

JavaScript Object Notation (JSON)

Human-readable XML

1. `Object` = set of (attribute, value) pairs
2. `Array` = list of values.
3. `Value` = string | number | true | false | null | object | Array

```
{"menu": {
  "header": "SVG Viewer",
  "items": [
    {"id": "Open"},
    {"id": "OpenNew", "label": "Open New"},
    {"id": "ZoomIn", "label": "Zoom In"},
    {"id": "ZoomOut", "label": "Zoom Out"},
    {"id": "OriginalView", "label": "Original View"},
    null,
    {"id": "Quality"},
    {"id": "Pause"},
    {"id": "Mute"},
    {"id": "Help"},
    {"id": "About", "label": "About SVG Viewer..."}
  ]
}}
```

JavaScript Object Notation

- Among the most popular data interchange formats today
- There exist JSON notations for other data formats, e.g., RDF

```
{
  "http://example.org/about" : {
      "http://purl.org/dc/terms/creator" : [ { "value" : "_:anna",
                                               "type" : "bnode" } ] ,
  "_:anna" : {
      "http://xmlns.com/foaf/0.1/name" : [ { "value" : "Anna",
                                             "type" : "literal" } ]
  }
}
```

Critique of tree data models (XML, JSON)

- Each information ends up in only one place
- OK for "classification" applications, structured text
- Fundamentally restrictive for **data = real world**!

Tim Berners-Lee, WWW proposal, CERN, 1998:
"*Many systems are organised hierarchically. A tree has the practical advantage of giving every node a unique name. However,* ***it does not allow the system to model the real world***."
(On newsgroups)*: "Typically, a discussion under one newsgroup will develop into a different topic, at which point* ***it ought to be in a different part of the tree***."

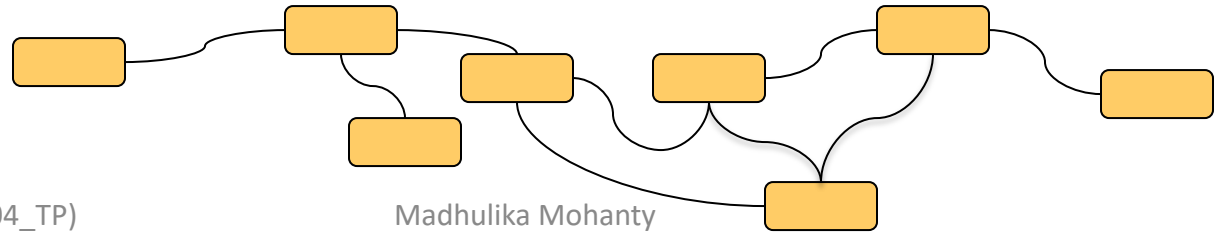

Graph data format for the Web: RDF

- Resource Description Format, W3C, 2003
- Resources have properties with values.
- URIs (Universal Resource Identifiers) identify resources
- Resources, properties, or values may be specified by an URI.
- Values may be constants

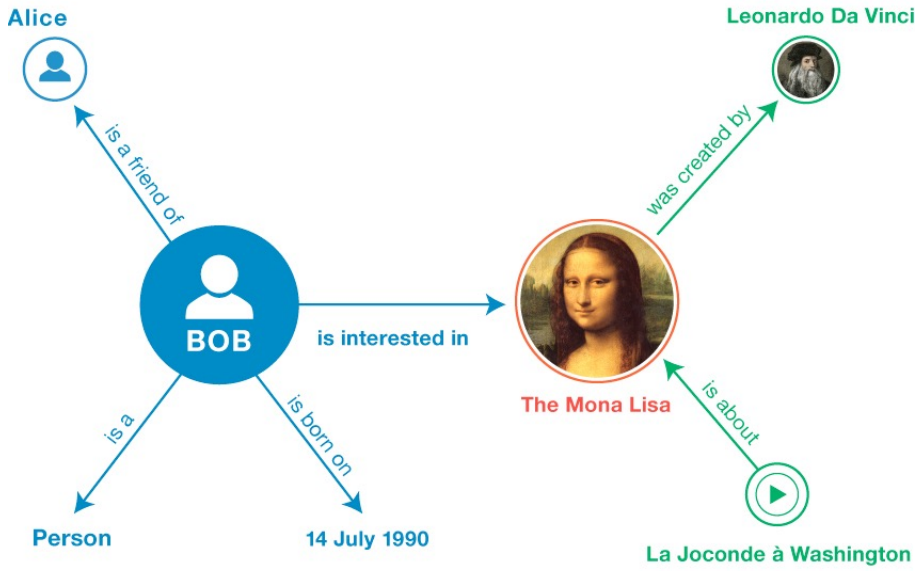

RDF feature: blank nodes

- Unnamed resource

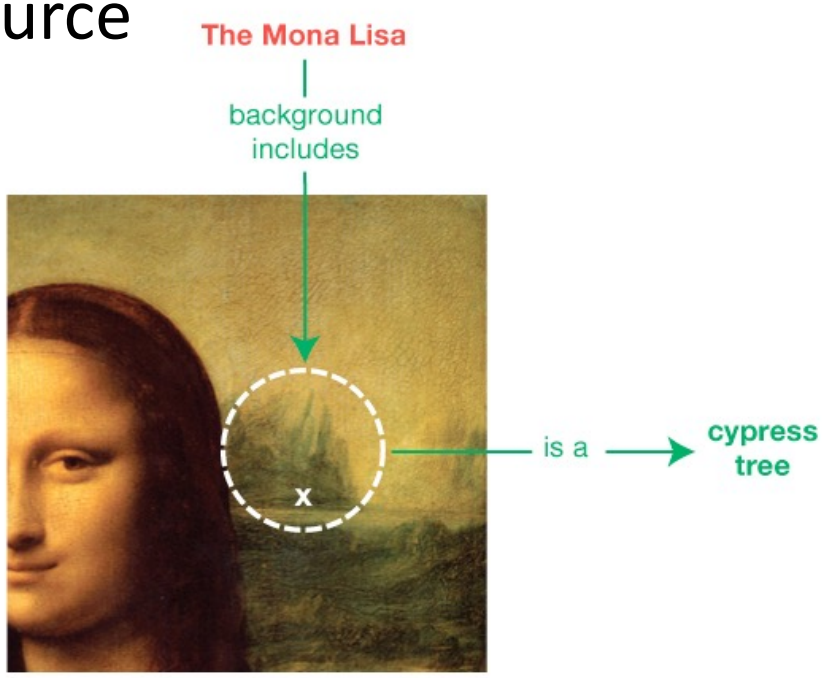

- « Labeled null »

RDF graphs

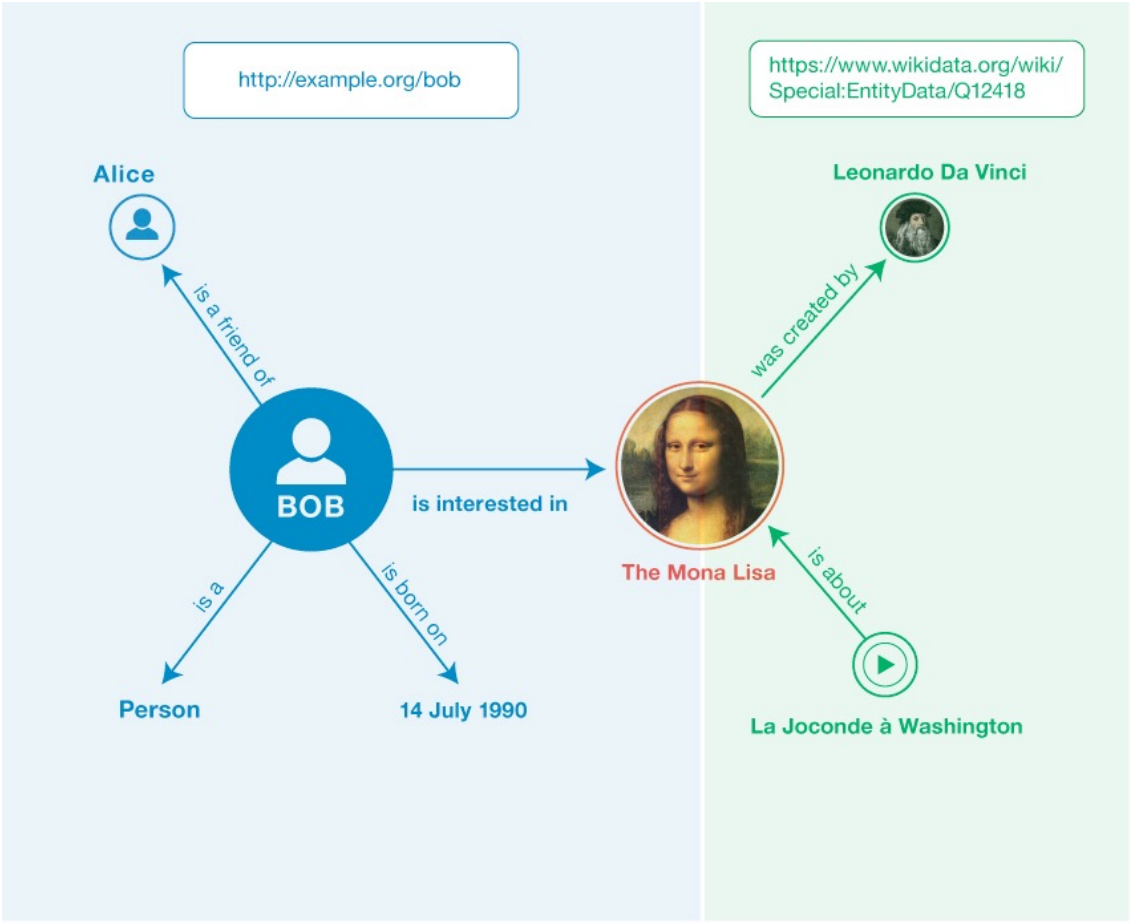
http://example.org/bob
https://www.wikidata.org/wiki/Special:EntityData/Q12418
Alice
Leonardo Da Vinci
is a friend of
was created by
BOB
is interested in
The Mona Lisa
is about
is a
is born on
Person
14 July 1990
La Joconde à Washington

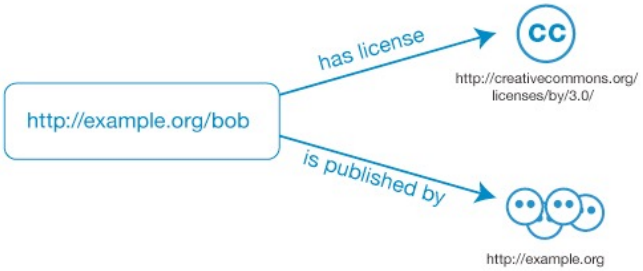
has license
CC
http://creativecommons.org/licenses/by/3.0/
http://example.org/bob
is published by
http://example.org

RDF Schema constructs

Construct	Syntactic form	Description
Class (a class)	**C** rdf:type rdfs:Class	**C** (a resource) is an RDF class
Property (a class)	**P** rdf:type rdf:Property	**P** (a resource) is an RDF property
type (a property)	**I** rdf:type **C**	**I** (a resource) is an instance of **C** (a class)
subClassOf (a property)	**C1** rdfs:subClassOf **C2**	**C1** (a class) is a subclass of **C2** (a class)
subPropertyOf (a property)	**P1** rdfs:subPropertyOf **P2**	**P1** (a property) is a sub-property of **P2** (a property)
domain (a property)	**P** rdfs:domain **C**	domain of **P** (a property) is **C** (a class)
range (a property)	**P** rdfs:range **C**	range of **P** (a property) is **C** (a class)

Typed RDF graph

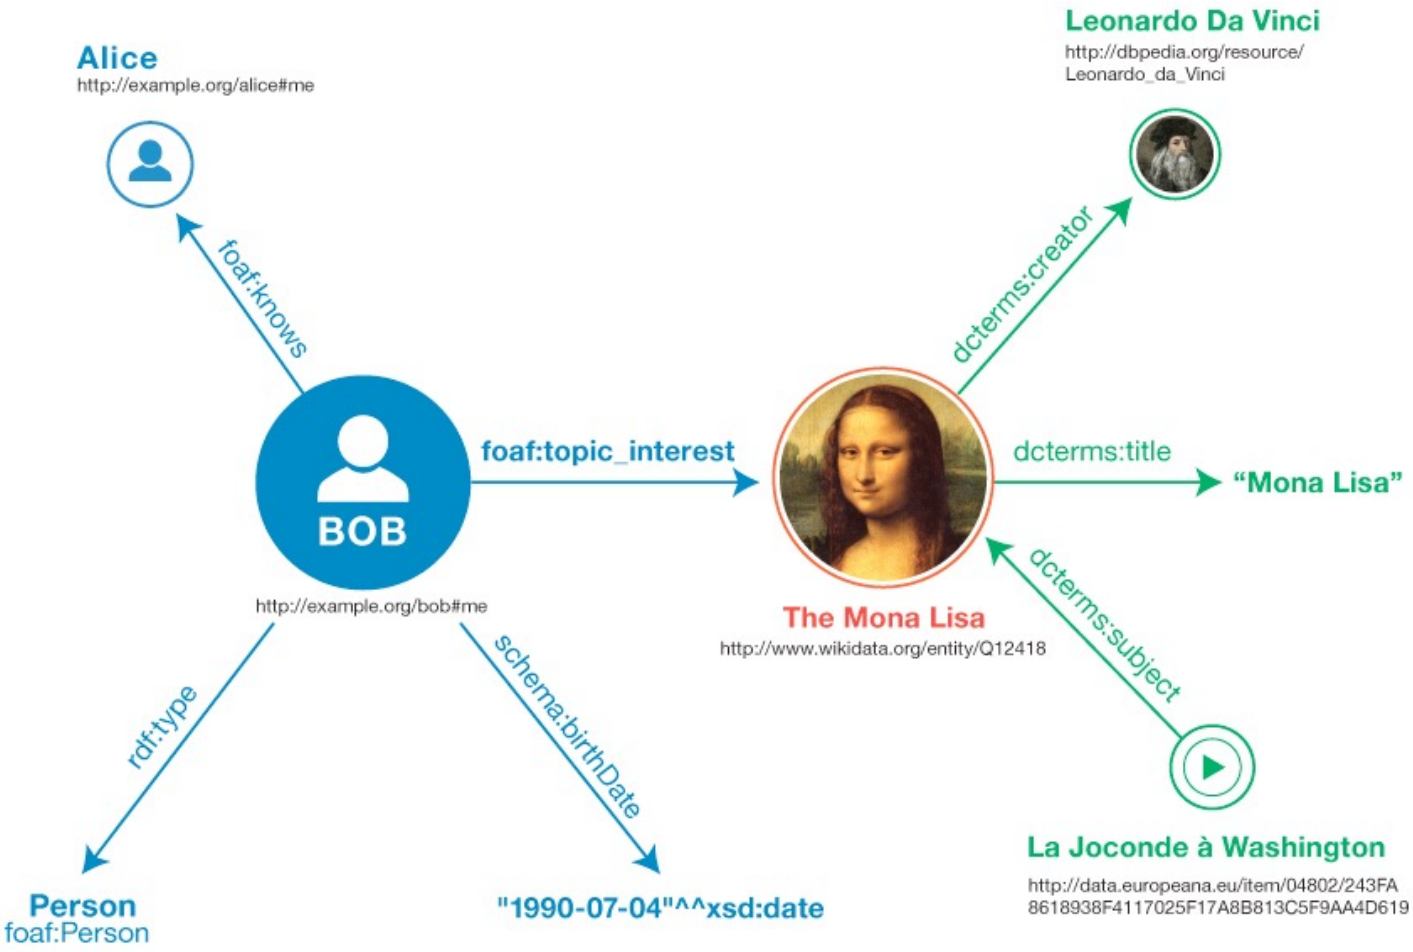

RDF reasoning

- RDF allow expressing **data *and* knowledge**
- Example:
  - if X *teaches a class*
  - then X *is a person*, *is an instructor*, *belongs to the school giving the class*, and *works for the university which includes the school*
- Reasoning exploits knowledge to infer *implicit data*

Linked Open Data (LOD) cloud

Linked Open Government Data project (logd.rw.rpi.edu): 10^{10} triples.

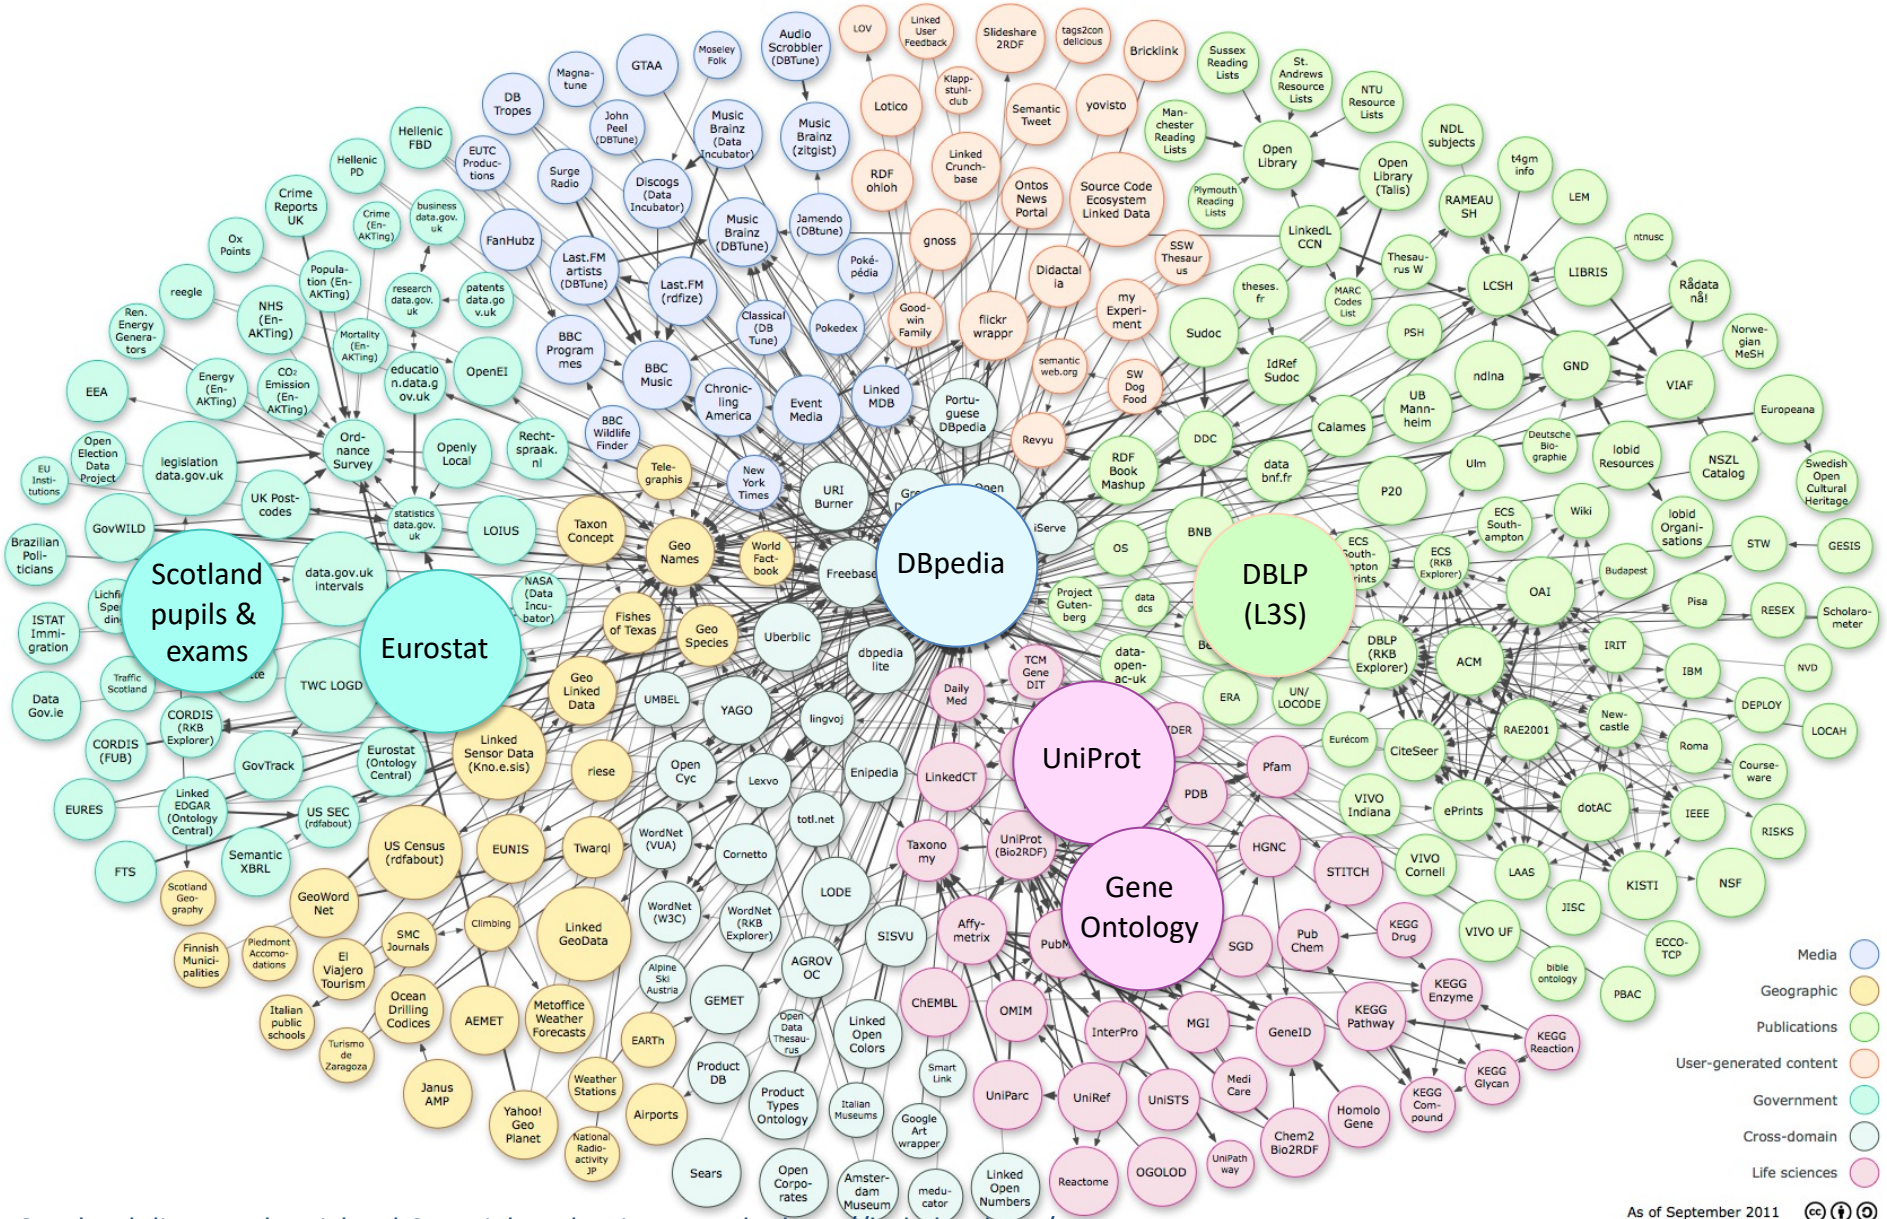

LOD cloud 2009

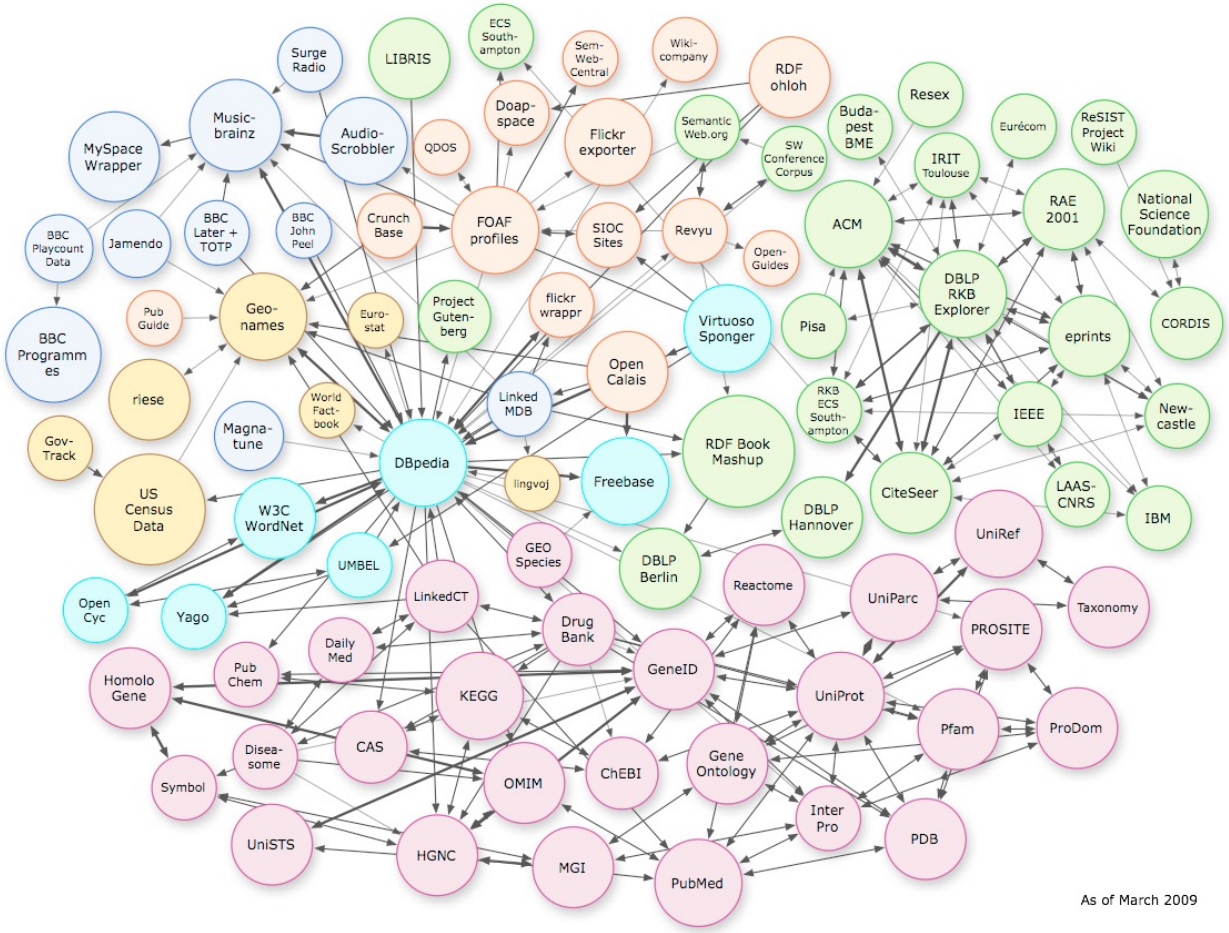

LOD cloud 2014

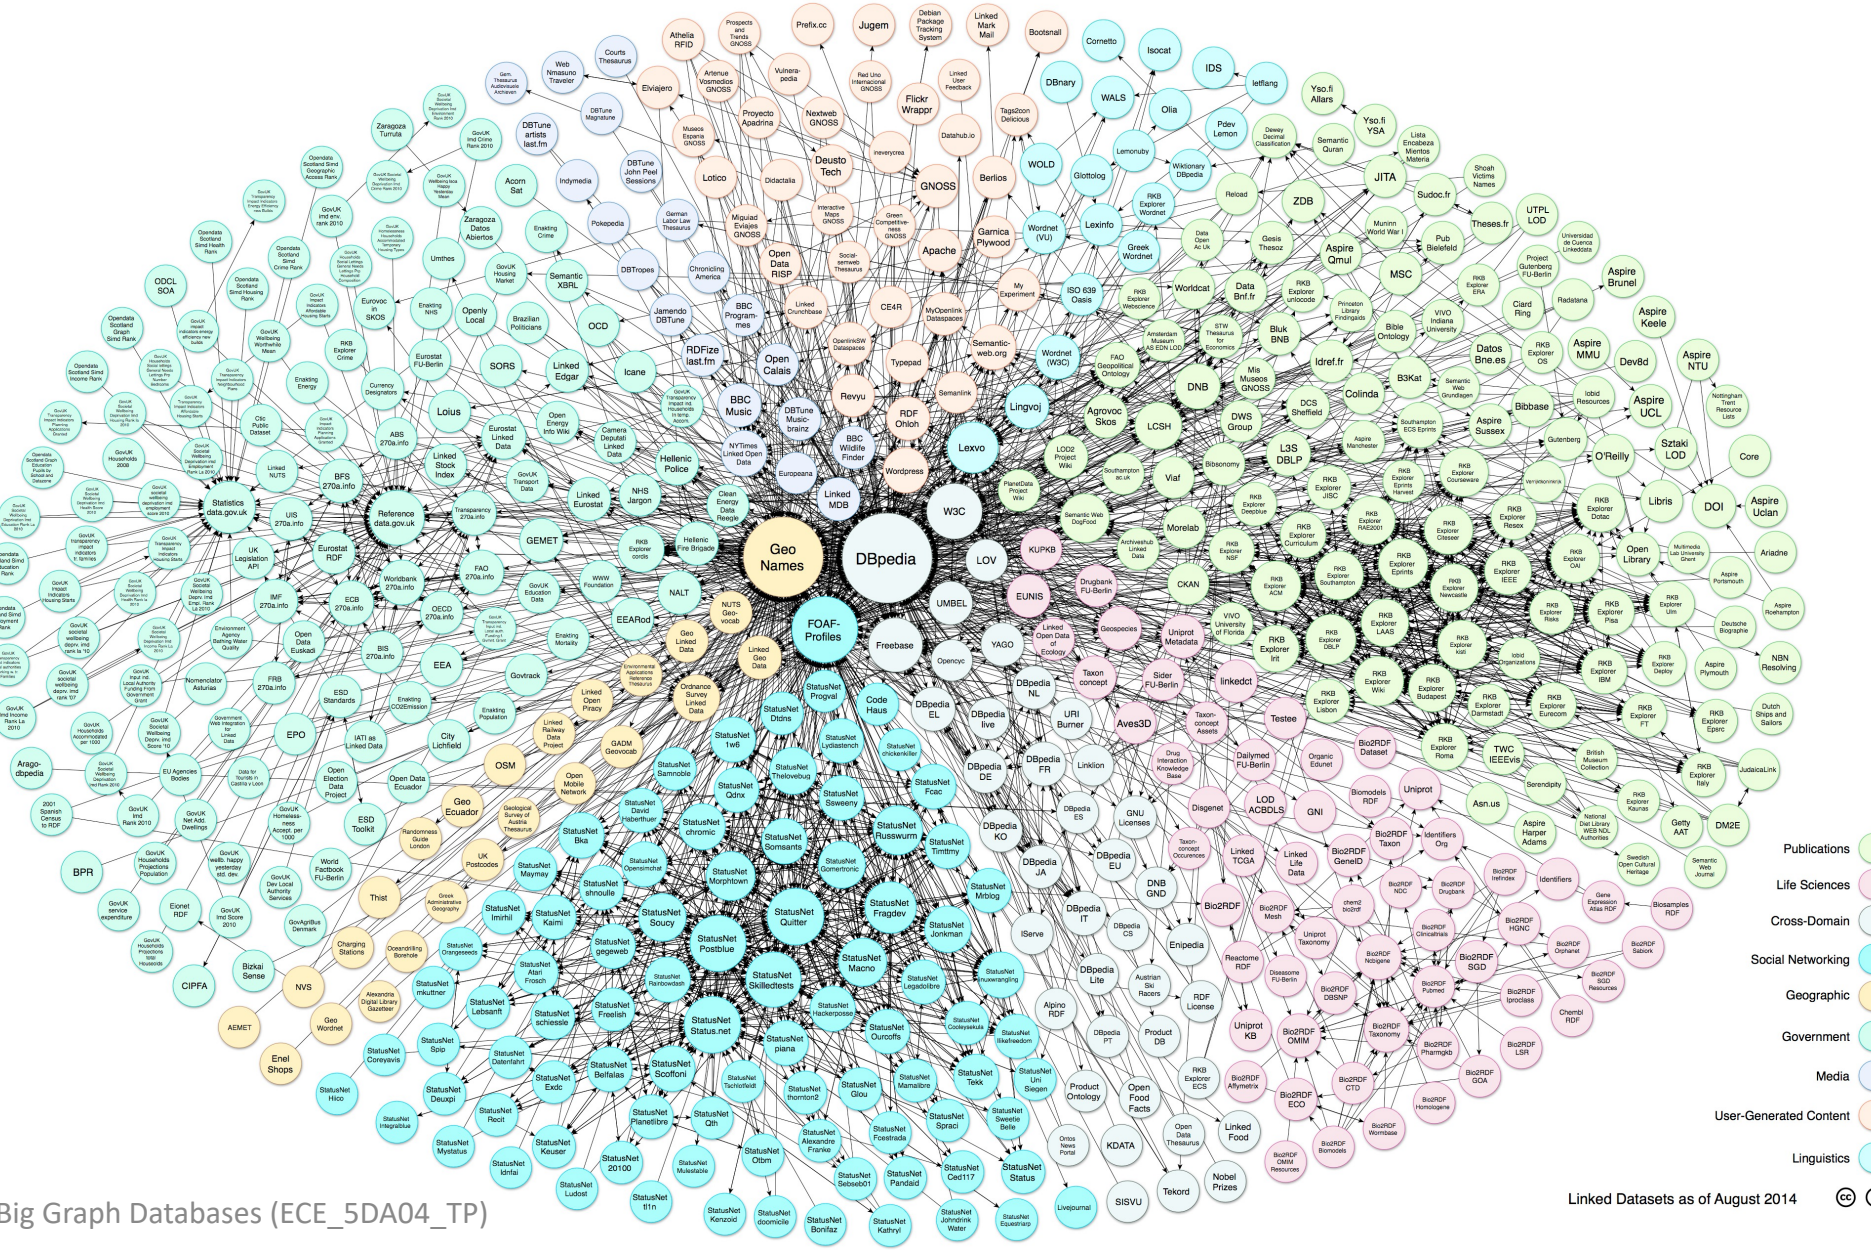

LOD cloud 2017

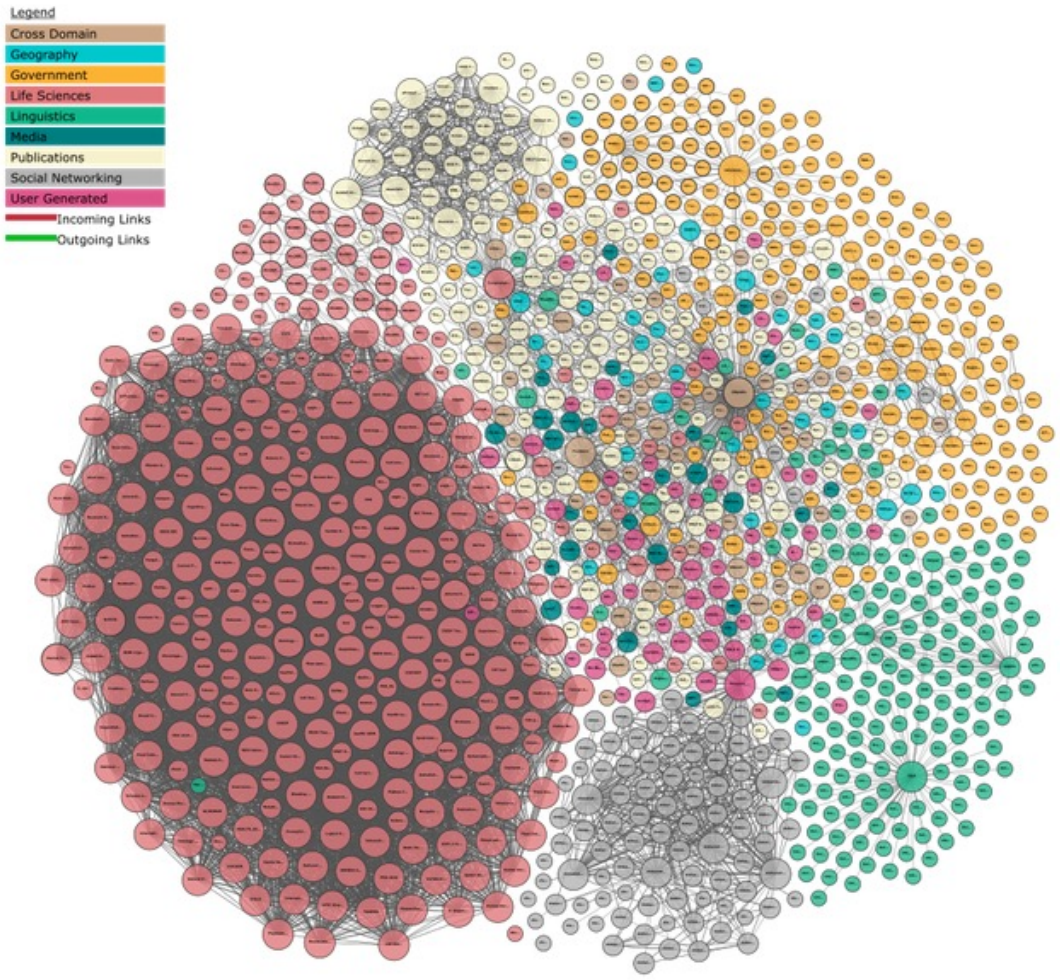

Linking Open Data cloud diagram 2017, by Andrejs Abele, John P. McCrae, Paul Buitelaar, Anja Jentzsch and Richard Cyganiak.
http://lod-cloud.net/

LOD cloud 2021

http://lod-cloud.net/

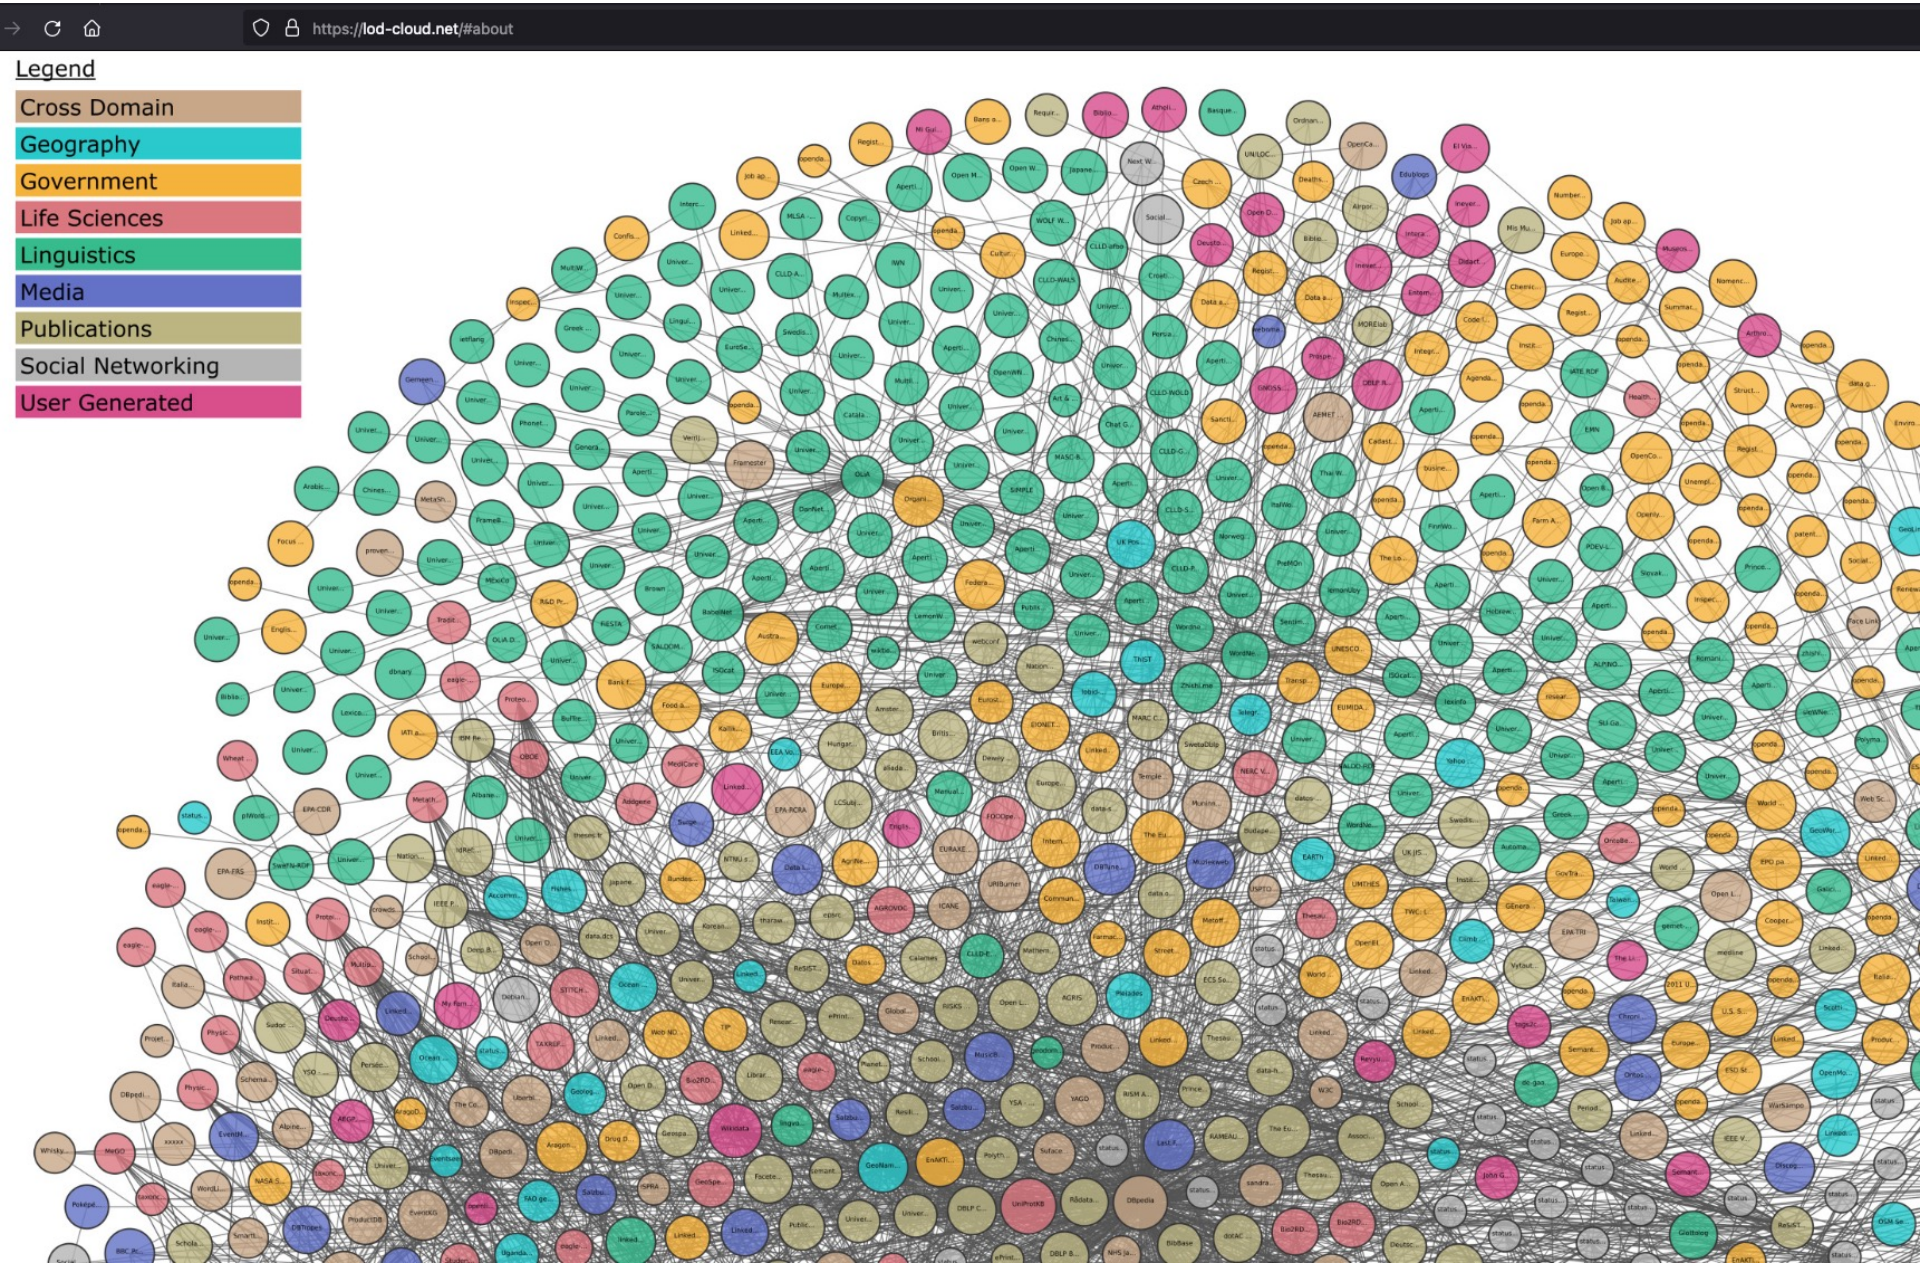

Open vs. linked data

1. **Linked Data**:
   "recommended best practice for exposing, sharing, and connecting pieces of data, information, and knowledge on the Semantic Web using **URIs** and **RDF**"
   - (Tim Berners-Lee) vision for the Web
2. **Open Data**:
   "idea that certain data should be freely available to everyone to use and republish as they wish, without restrictions from copyright, patents or other mechanisms of control"
   - In principle, orthogonal to the Linked aspect
   - In practice, Linked is a technical means toward Open

Open Data: data.gov (US)

DATA.GOV
EMPOWERING PEOPLE

Search our catalogs.. SEARCH

HOME | DATA | TOOLS | COMMUNITY | METRICS | OPEN DATA SITES | GALLERY | WHAT'S NEW UPDATED

PREVIOUSLY HIGHLIGHTED DATASETS AND TOOLS

Below is a gallery of datasets and tools that have been highlighted on the Data.gov home page. Click on "View More" to learn more about the data and link to the data itself.

This gallery displays just a tiny fraction of the datasets available to you on Data.gov. As we continue to add datasets, tools and highlights, we encourage you to explore all the valuable resources in our raw data, tools, and geodata catalogs.

* Displaying 62 datasets and tools.

Open Data: data.gov (US)

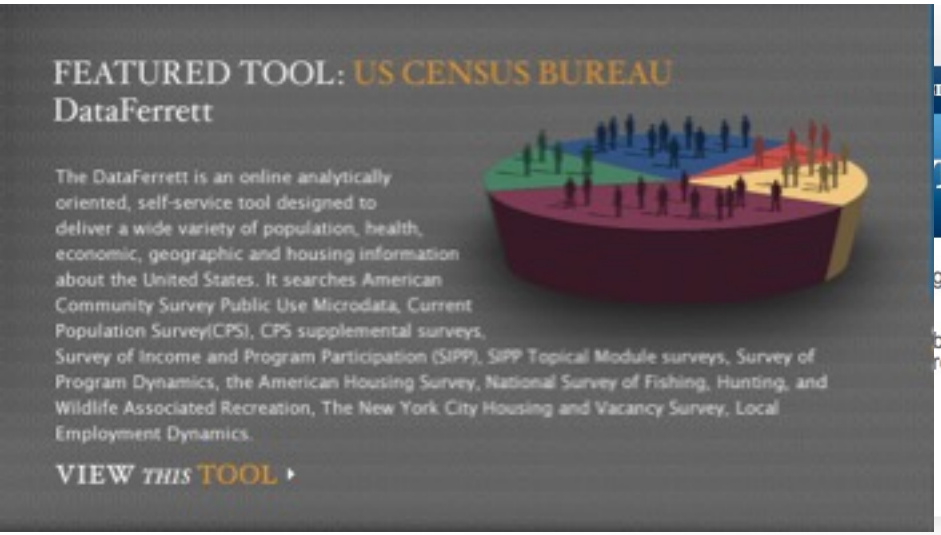

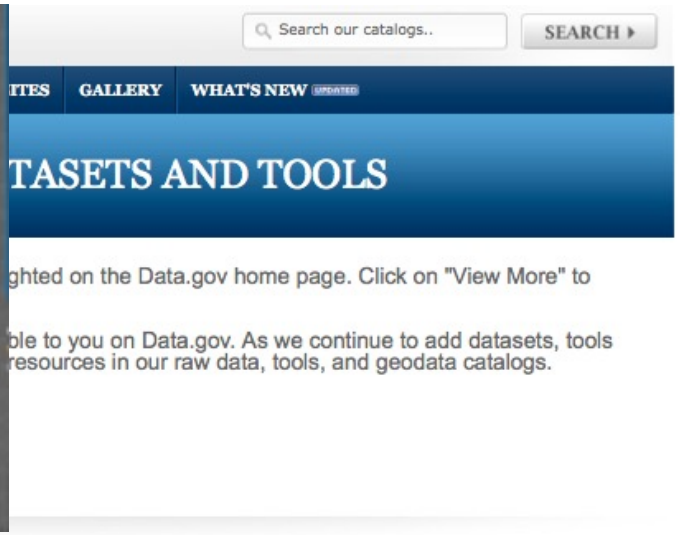

Open Data: data.gov (US)

Open Data: data.gov (US)

Open Data: data.gov (US)

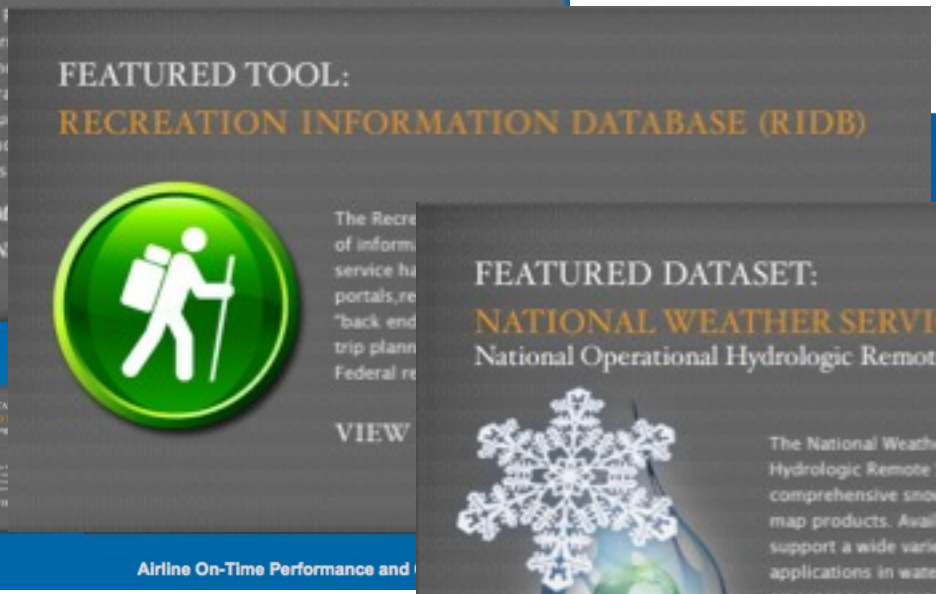

Open Data from Etalab (FR)

GDP per French region (Le Journal Du Net)

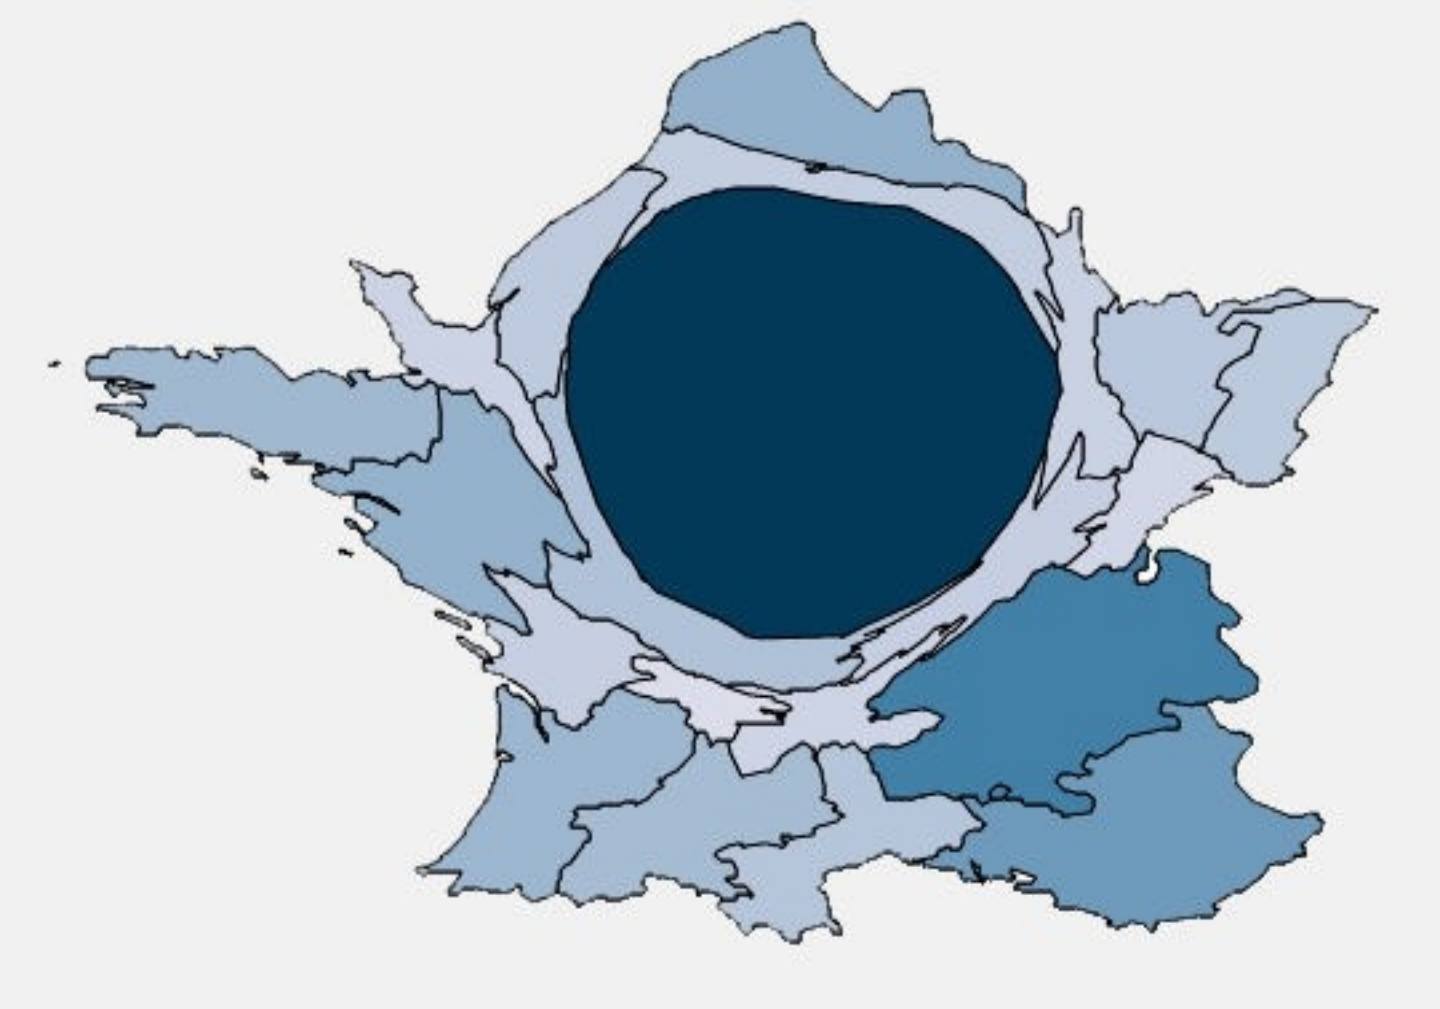

Open Data from Etalab (FR)

Organic agriculture per French region

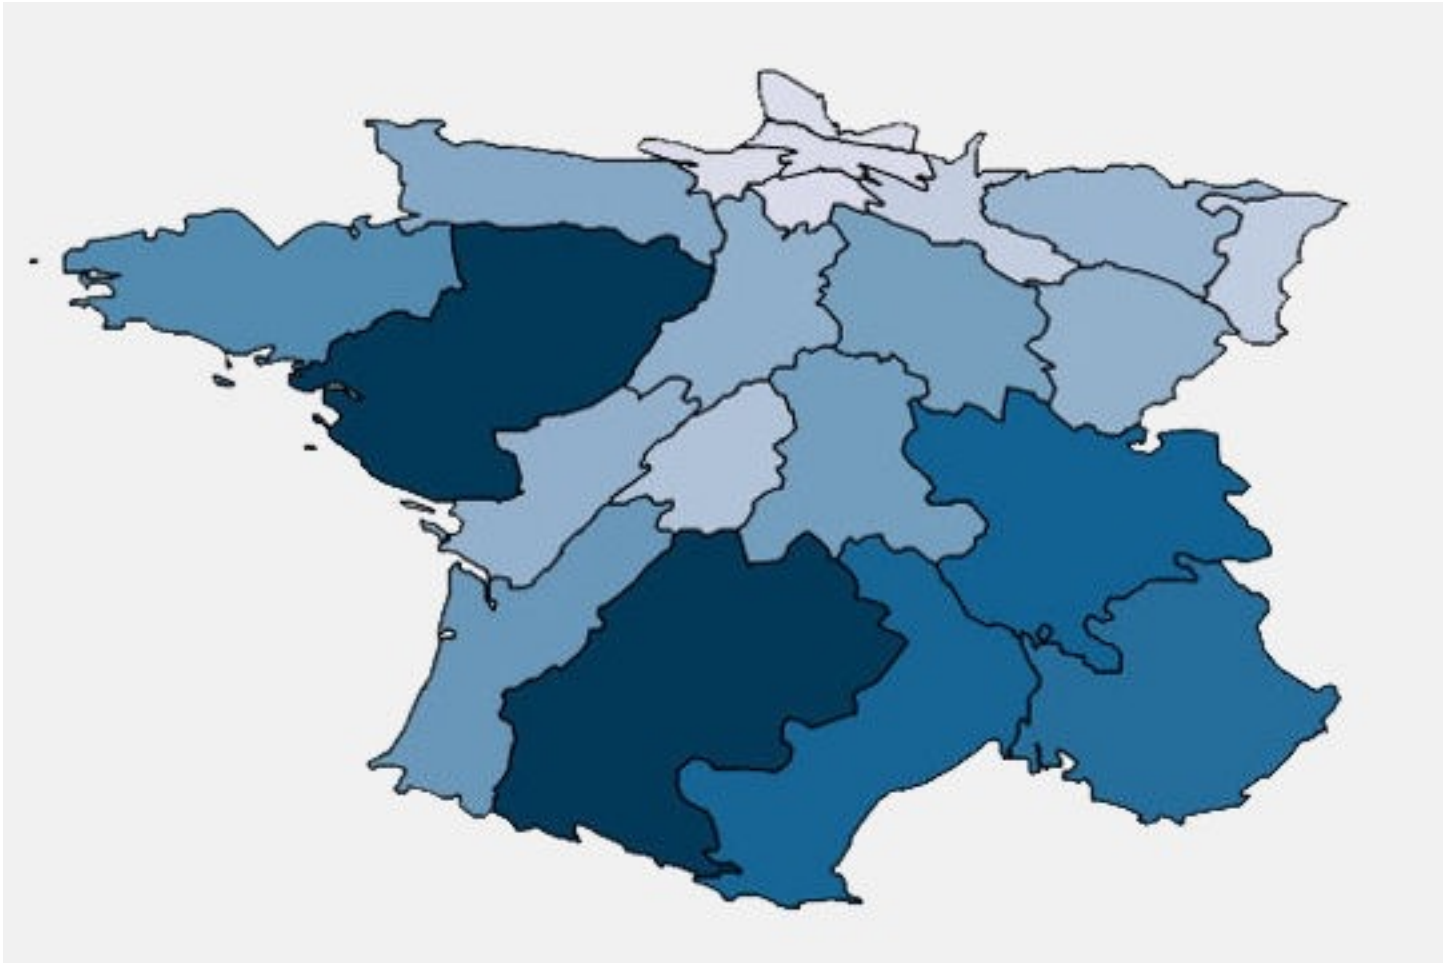

Open Data important for democracy

- Journalists and/or NGO workers play increasingly important role explaining society functioning
  - E.g., ICIJ Panama Papers investigation

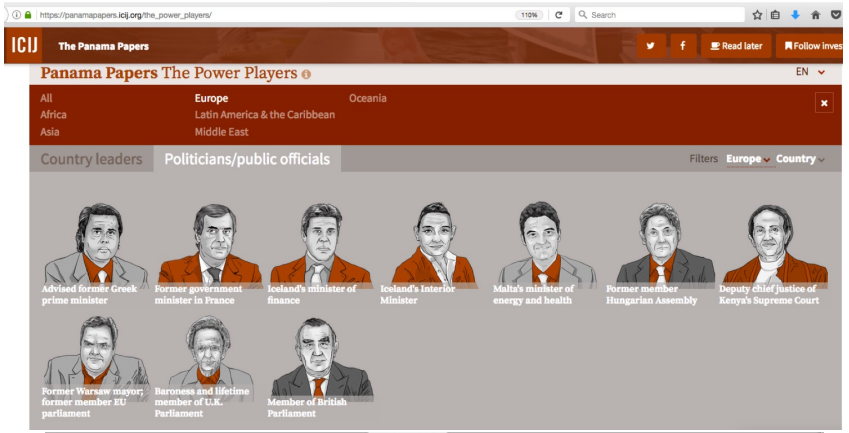

Open Data important for democracy

- Journalists and/or NGO workers play increasingly important role explaining society functioning
  - E.g., Météo France data on average yearly temperature, Les Décodeurs (Le Monde) https://www.lemonde.fr/les-decodeurs/article/2021/01/06/visualisez-le-rechauffement-climatique-en-france-et-dans-votre-ville-avec-nos-barres-de-rechauffement_6065388_4355770.html

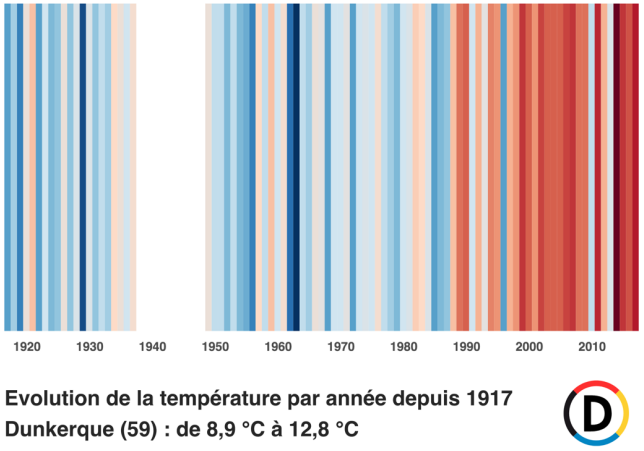

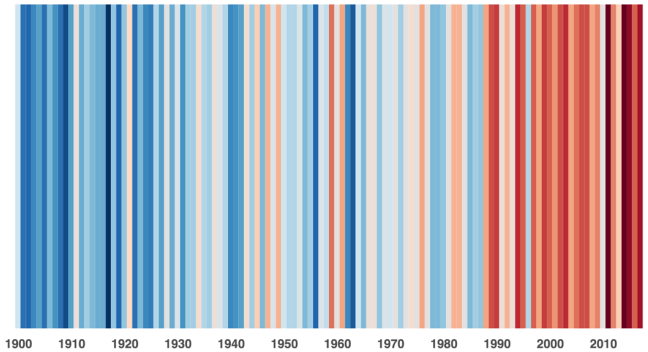

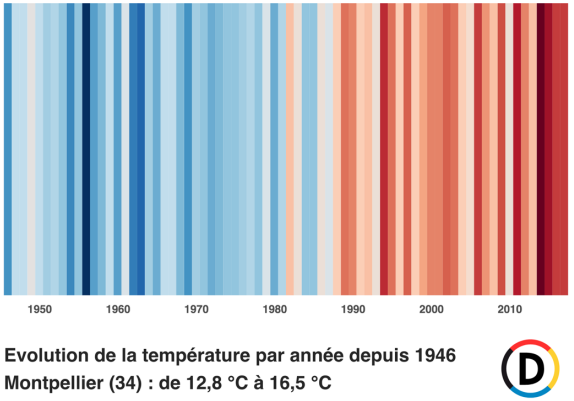

US Budget allocation, 2016

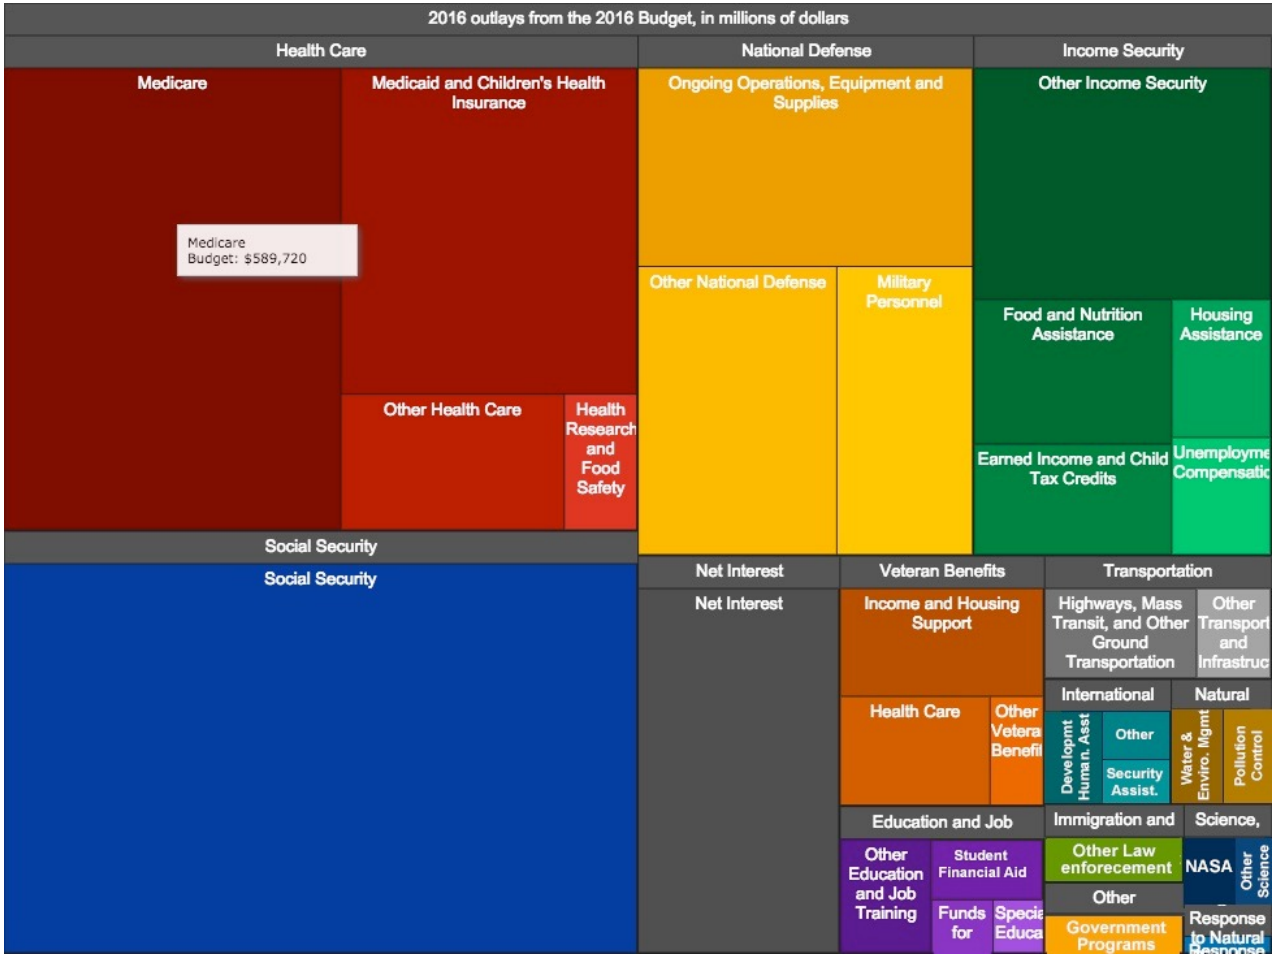

https://cdns.tblsft.com/sites/default/files/pages/5_2016_budget_visualization_safe_img.jpg

Property graphs: the other "standard" for graph data

- Initially introduced by the Neo4J system

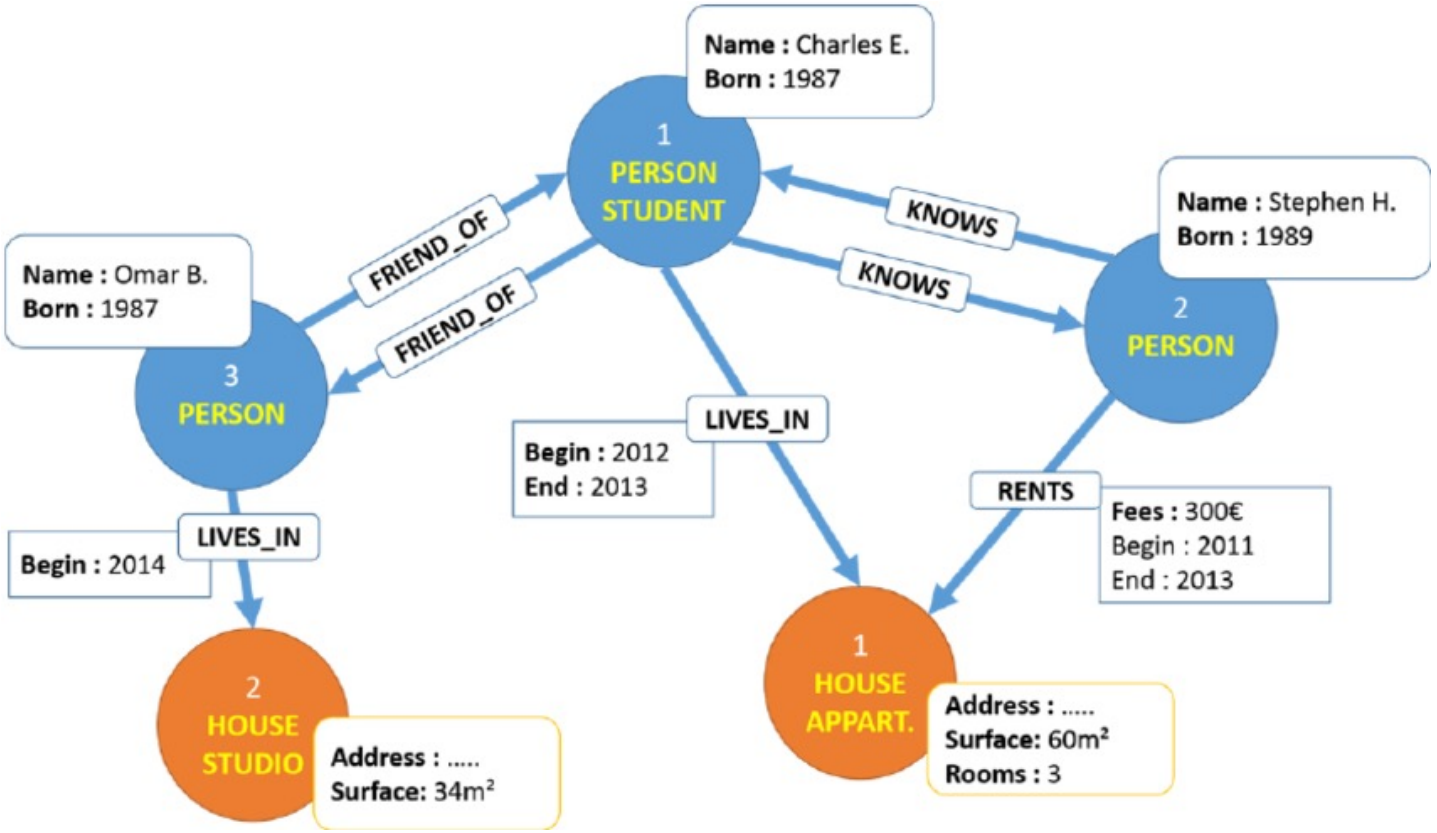

Property graphs: the other "standard" for graph data

- A property graph can be loaded from or dumped into a set of CSV files.
  - One CSV for each node type, e.g., one for PERSON, one for STUDENT, one for HOUSE, etc.
  - One CSV for each edge label, e.g., FRIENDS_OF, LIVES_IN, KNOWS etc.
- A graph can be queried using the Cypher or GQL languages (SQL for graphs)

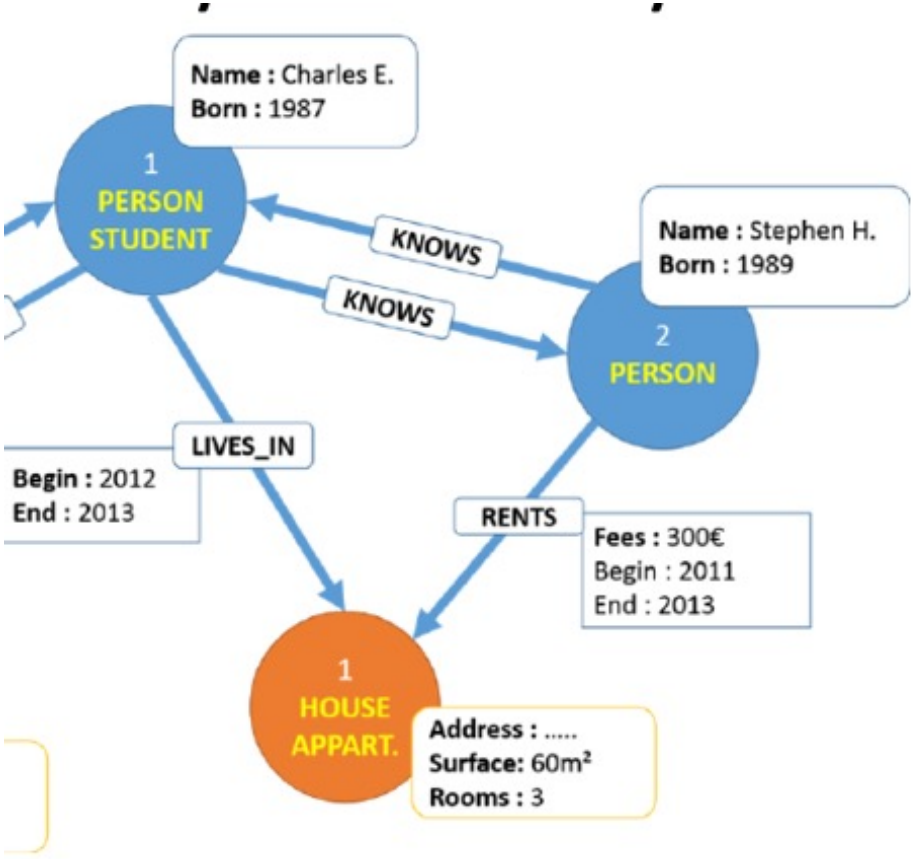

Property graphs: the other “standard” for graph data

- Embraced by all major database vendors (Oracle, DB2, Microsoft)
- Linked Data Benchmark Council (LDBC)
- Many exciting developments

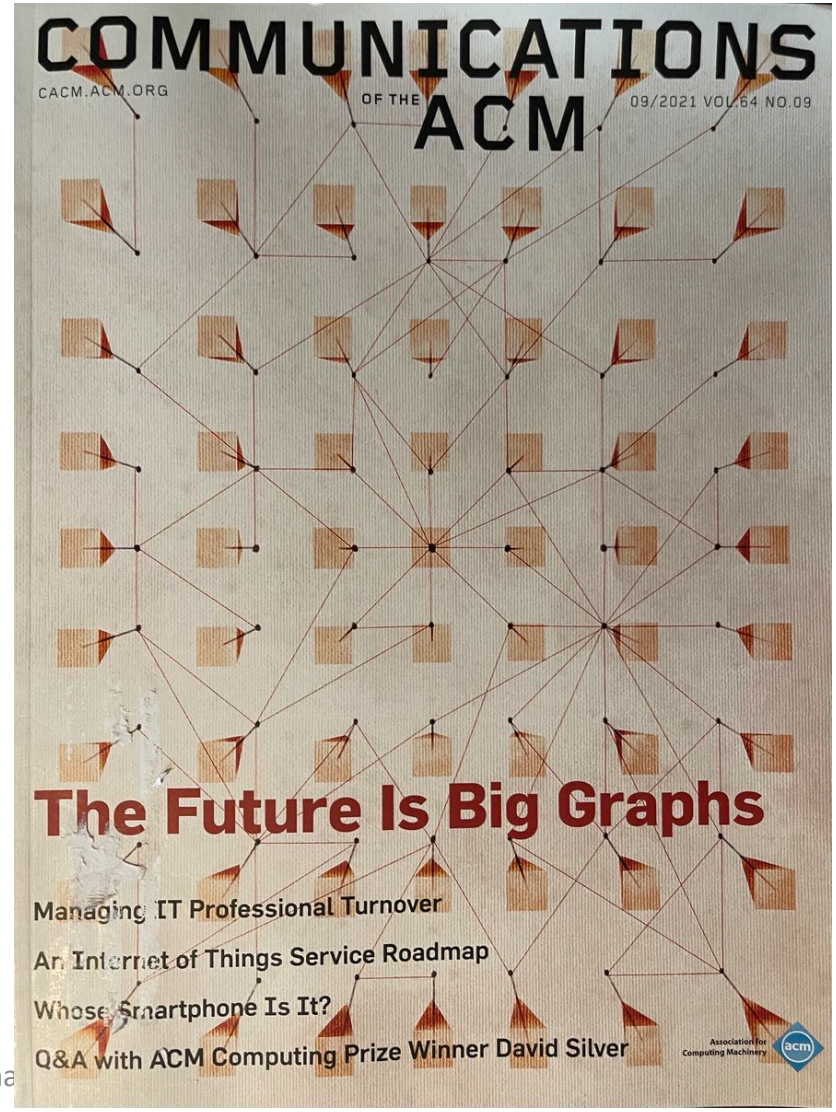

Object-oriented databases

- Studied in 1980's-2000
- Goal: adopt benefits of object-oriented programming (rich structure, encapsulation, inheritance, strong typing, etc.) to databases
- C++ , then Java classes with persistence

```
class myClass1{ ...}
        persistent class myClass2{...}
        public static void main {... commit() }
```

- Native OO databases mostly disappeared.
- Major relational databases have *OO extensions*
  - Declare a table as containing objects of a certain type
  - Some method support
  - Little-known, not central to the SQL standard.

Key-value databases

- Simplest possible data model
  - V1: 1 key → 1 value
    - e.g., "1" → "Julie, Paris 12e"
  - V2: 1 key → 1 record → 1 property → 1 value
    - e.g., "1" → ● → "name" → "Julie"
      → "address" → "Paris 12e"
- Records are grouped in named collections
- In a collection, records may have different property sets ➔ Heterogeneity!
- A property may be defined with unique values or multiple values, with or without duplicates
- No query language, only put and get.

Big data heterogeneity (variety)

Data model & data management system soup

- hierarchical, relational, object-oriented, XML, RDF, JSON, key-value pairs…

Traditionally this has been solved (time and $ permitting) with **data migration** / **ETL** (extract-transform-load)

- Heterogeneous data and high throughput may make ETL impractical for Big Data →
  Need to **exploit big, heterogeneous data as is**

Varied Big Data has huge value potential

- Real estate ad from Zillow (US):

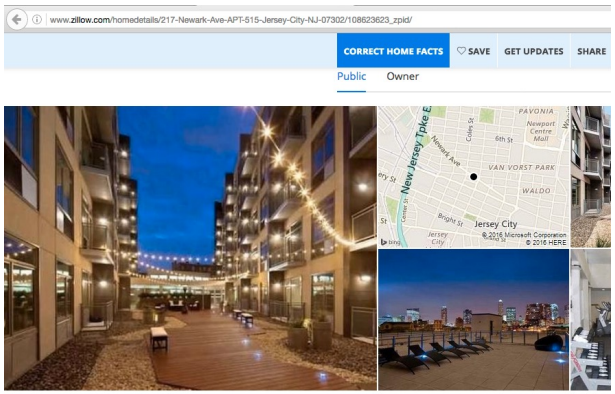

217 Newark Ave APT 515
Jersey City, NJ 07302
-- beds · -- baths · 871 sqft Edit
Edit home facts for a more accurate Zestimate.

Thinking About Selling?
Find a local agent who can give you a professional estimate of your home value.
Find an Agent

2 Bedroom 2 Bath w/ Garage Parking at The Saffron. Slee modern design and details that include stainless steel appliances, Caesarstone counters, Bamboo floors and soaring ceiling heights. Master bedroom with walk in clos and master bathroom. Great storage space, W/D and dishwasher. Manhattan views from spacious roof deck lounge and sun deck or relax in common courtyard. On si fitness center. Low maint. & taxes. 2.5 blocks from the Gr St. PATH.

Home price
$ 505,000
Down payment
$ 101,000 20 %
Loan program
30-year fixed
Interest rate See current rates
3.304 %
Include taxes/ins.

Your payment
$1,770
P&I
$1,770
Get pre-qualified

Price / Tax History

Price History | Tax History

DATE	EVENT	PRICE	AGENTS
05/19/14	Sold	$505,000 +1.2%	Stacey Bonavitacola
03/09/14	Listed for sale	$499,000 +23.5%	
		$404,000	

Downtown

me values will increase 6.3% next year, compared to a 4% hole. Among Downtown homes, this home is valued 8.4% less ome, and is valued 5.2% less per square foot.

aradise)

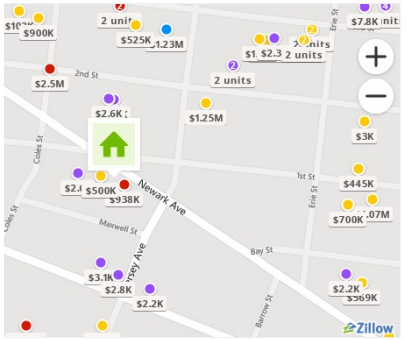

Home Expenses

INTERNET, TV & PHONE

Bundle Services and Save — $50-$100 / month
Bundle your monthly services and consider ways to "cut the cord" with low-cost online alternatives to cable such as streaming and on-demand video.
Learn More

SECURITY

Make Your Home a Fortress — $14.99 / month
- 24/7 protection for just $14.99/mo
- No long-term contracts, so you're free to cancel anytime

Try It Today

Powered by SimpliSafe

Nearby Schools in Jersey City

SCHOOL RATING		GRADES	DISTANCE
3 out of 10	Number 4 Middle (assigned)	6-8 & ungraded	0.3 mi
10 out of 10	Dr. Ronald Mc Nair Academy High (assigned)	9-12 & ungraded	0.3 mi
5 out of 10	Number 5 Elementary	PK-8 & ungraded	0.4 mi

More schools in Jersey City
Data by GreatSchools.org

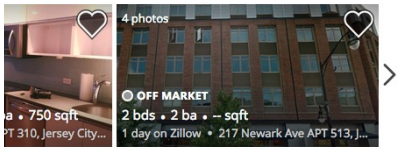

Improve Your Home Value

PROJECT	ADDED VALUE	PROJECT COST
Entry Door Replacement	+$2,337	$3,156
Vinyl Window Replacement	+$10,353	$15,030
Bathroom Remodel	+$7,297	$18,841
Major Kitchen Remodel	+$44,675	$61,399

See More Home Improvement Inspiration

Nearby Similar Sales

- SOLD: $551,700
  Sold on 4/4/2016
  2 beds, 1.0 baths, 800 sqft
  158 Wayne St APT 401A, Jersey City, NJ 07302
- SOLD: $589,000
  Sold on 10/15/2015
  2 beds, 1.0 baths, 987 sqft
  280 Monmouth St APT B, Jersey City, NJ 07302
- SOLD: $594,000
  Sold on 6/1/2016
  2 beds, 1.0 baths, 1013 sqft
  227 Christopher Columbus Dr APT 217B, Jer...
- SOLD: $599,000
  Sold on 5/27/2016
  2 beds, 1.0 baths, 978 sqft
  341 Monmouth St APT 311D, Jersey City, NJ ...
- SOLD: $600,000
  Sold on 10/5/2015
  2 beds, 2.0 baths, 1016 sqft
  287 8th St APT 4B, Jersey City, NJ 07302

See sales similar to 217 Newark Ave APT 515

A brief history of data models

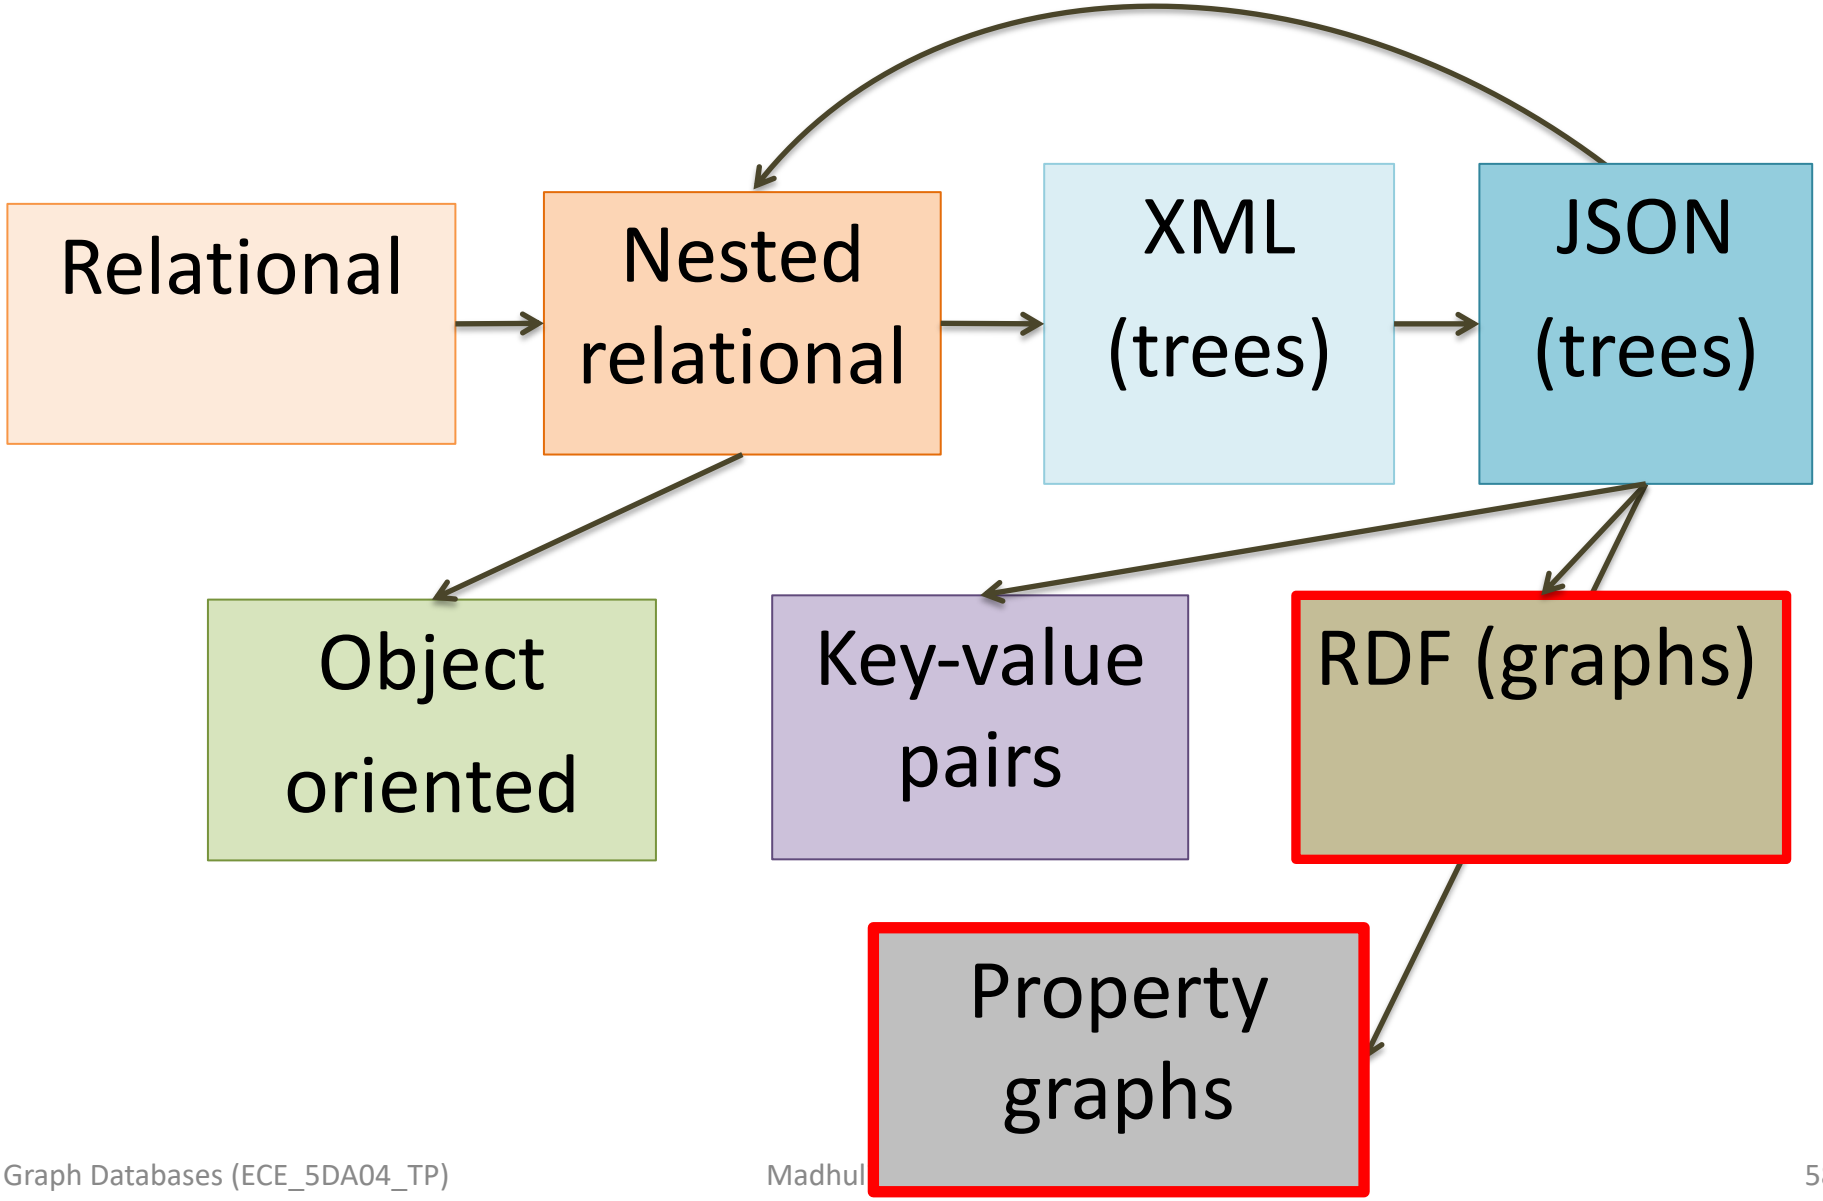

Defining Big Data: the V's

- Volume
  - Scale
- Velocity
  - Speed of producing and consuming the data
- Variety
  - Very different sources and data types
- Veracity
  - Is the data correct / certain / true?

Big Data veracity

- Is this true? (What is the probability?)
- Contradictory sources (1 vs. 2 clocks)
- Errors in the data
  - Humans introduce many errors
  - Sensors may have failures or erroneous readings
    - Light or heating sensors in a building
    - Wear and tear
- Tackled by **data curation / cleaning / quality** tools for regular (or at least homogeneous) data
- Recent ML methods for this

Big Data veracity

- Data reconciliation / entity extraction
- Large-scale Knowledge Bases such as YAGO, DBPedia, Google Knowledge Base (> Freebase)
  - Reference database of core facts

    - Places (city/country etc.), people (public figures, scientists, artists etc.), events (born, died, emigrated, was created…), time
  - Ontologies automatically extracted from Web and other specific sources → they may have *errors* and are *incomplete*
  - Manual curation and improvement, e.g., YAGO 4.5

Big picture on Big Data

1. **Volume, velocity, variety, veracity**
2. Probably not all the data has the same **value** $/B
   - This is why enterprise databases are preserved
   - Very large, unstructured, uncertain-value data may instead be stored in larger-scale, lower-performance systems, mostly based on Hadoop or Spark
     - Massively parallel processing
     - Iterative, machine learning

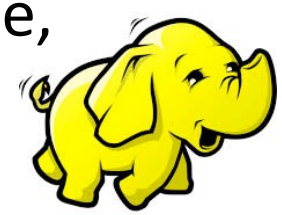

3. Deployment model changes (→ **cloud**)

From databases to Big Data

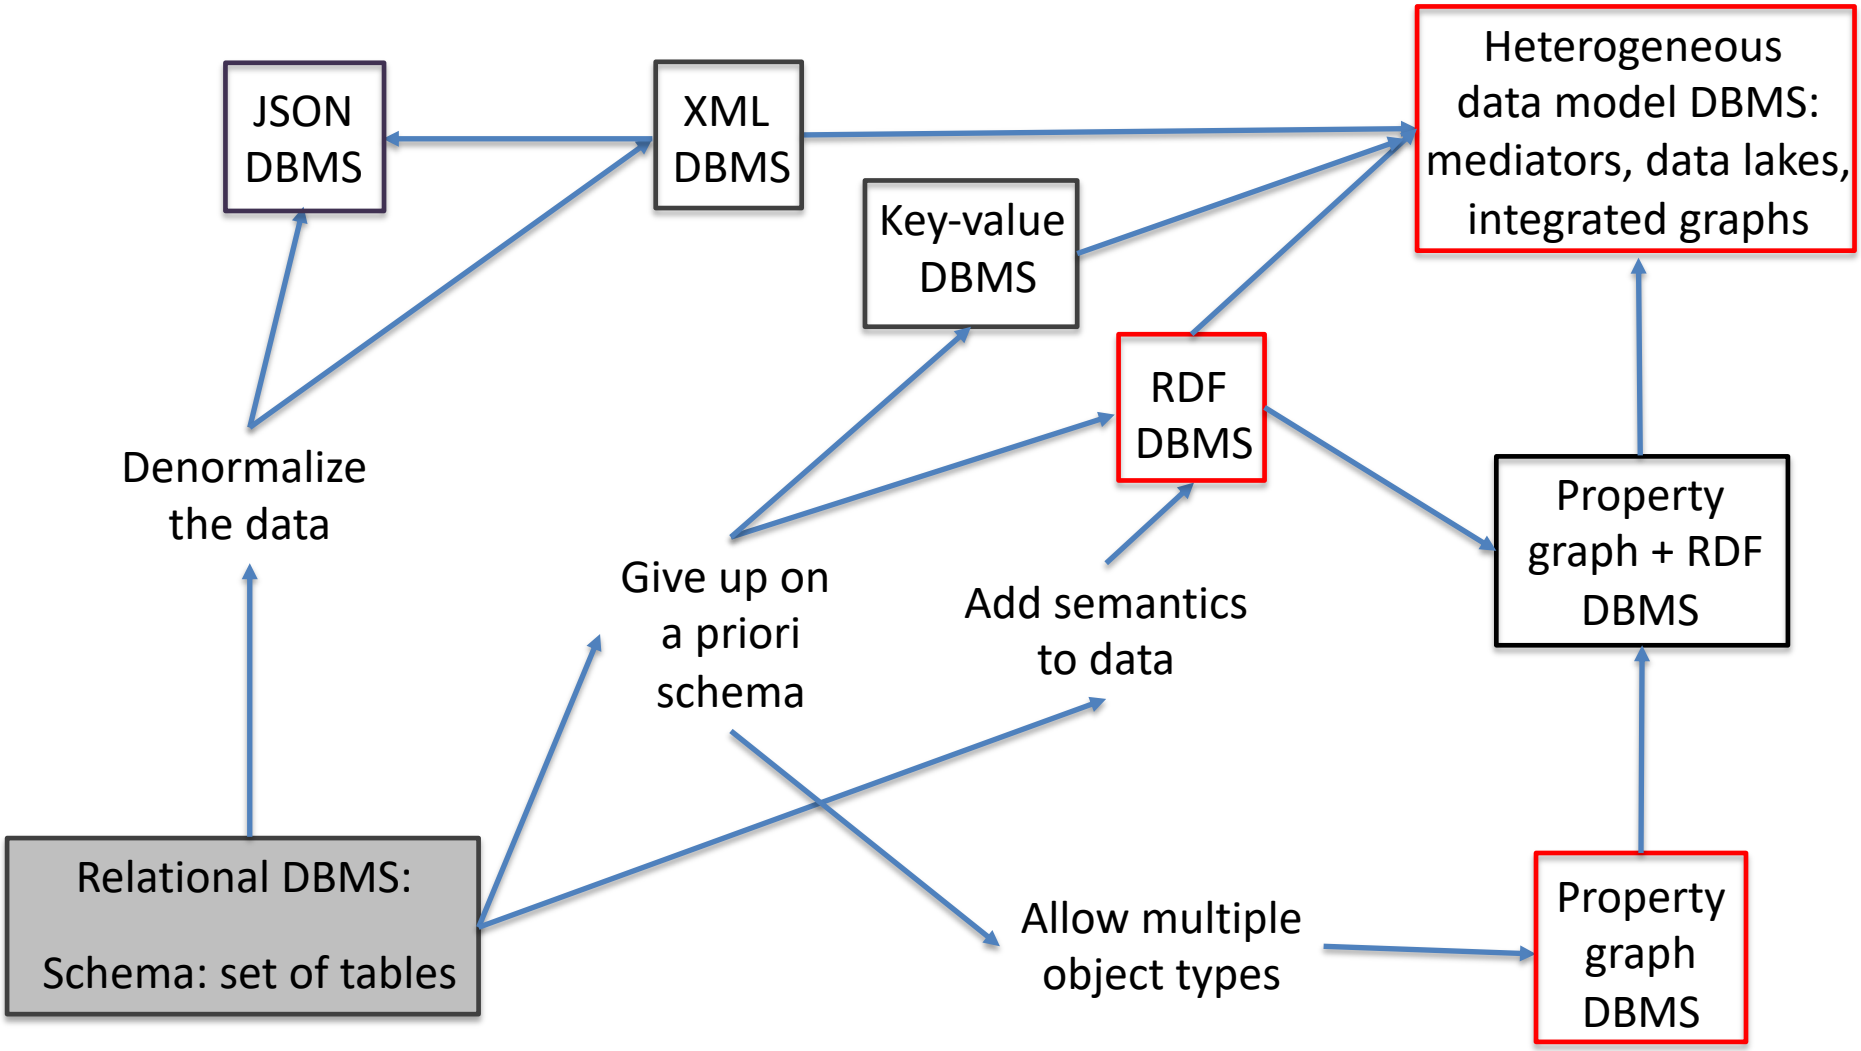

From databases to Big Data

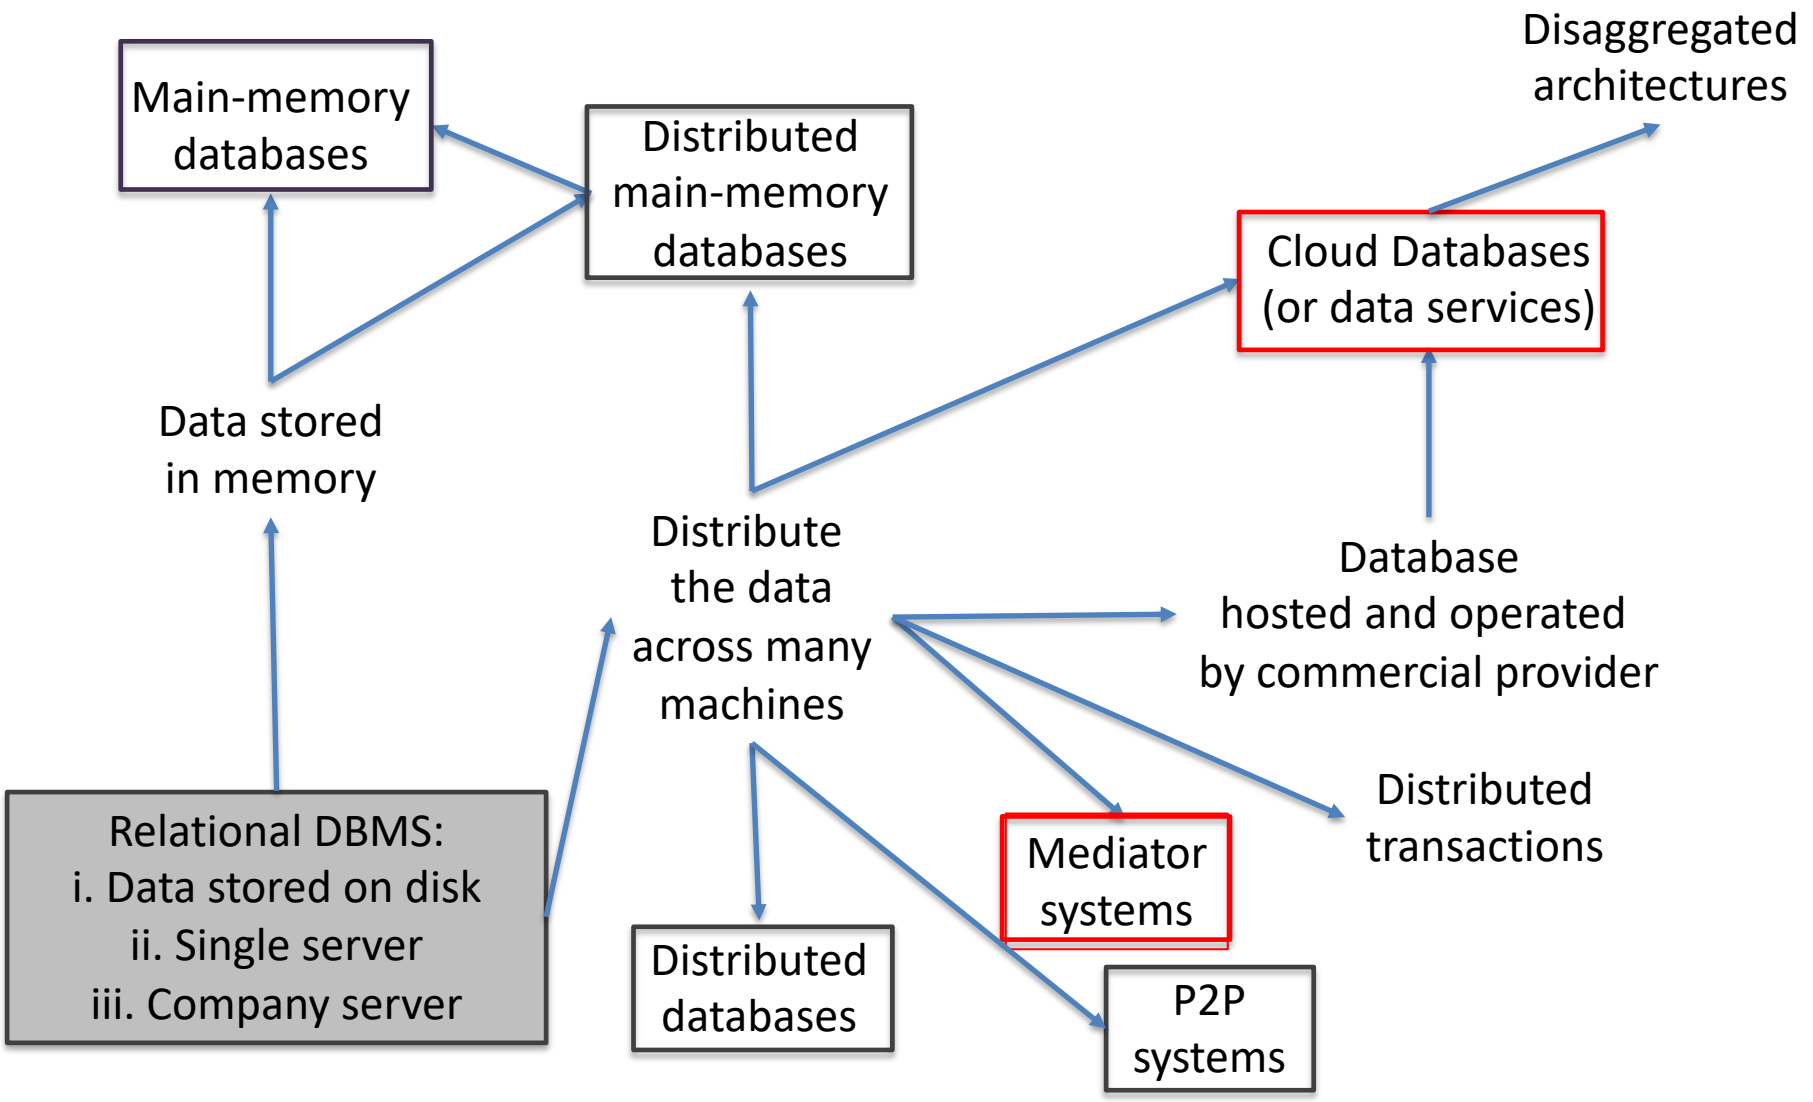

Questions?