# Fast, just-in-time queries on heterogeneous scientific data

*Anastasia Ailamaki*

EPFL and RAW Labs SA

*work funded by the Swiss NSF and the European Union*

***Most firms estimate that they are only analyzing 12% of the data that they already have.***

*Forrester, 2014*

- Growing data
- Growing heterogeneity
- Data movement restrictions

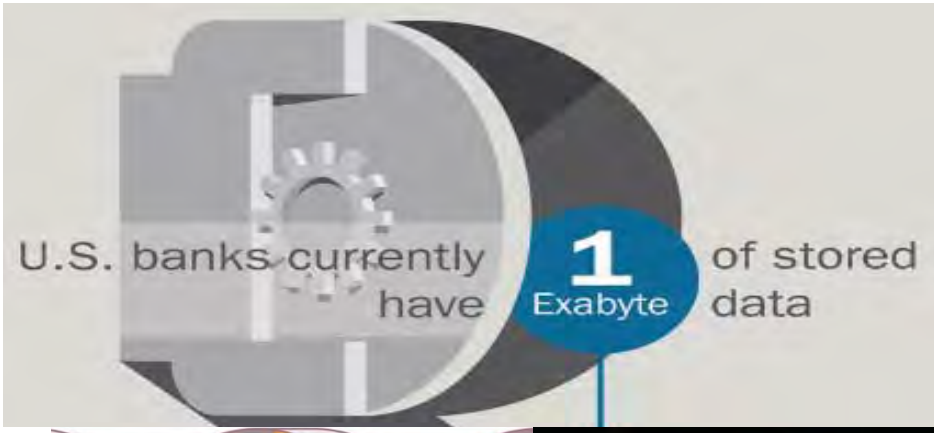


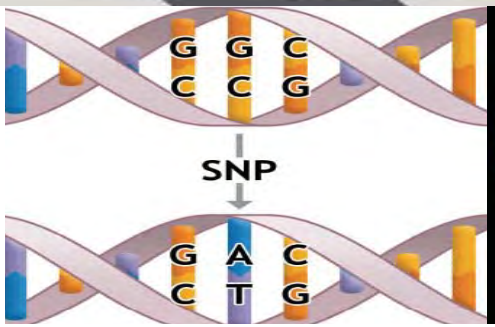


**Available data impedes business & scientific analytics**

# needful data

**[VLDB12a,IISWC12]**

- Cloudera customer jobs use KB-TB data
  (from the PBs gathered daily)
- 80% of jobs access 1-8% of bytes
- 90% of FB jobs read <10% of bytes stored
- 80% of reuse is within 3hrs
- 85% of data used is 1 day old (Yahoo!+UIUC)

**interesting is** 

**data, re-accessed soon**

# The Digital Universe: 50-fold Growth from the Beginning of 2010 to the End of 2020

(The Digital Universe, EMC/IDC 2014)

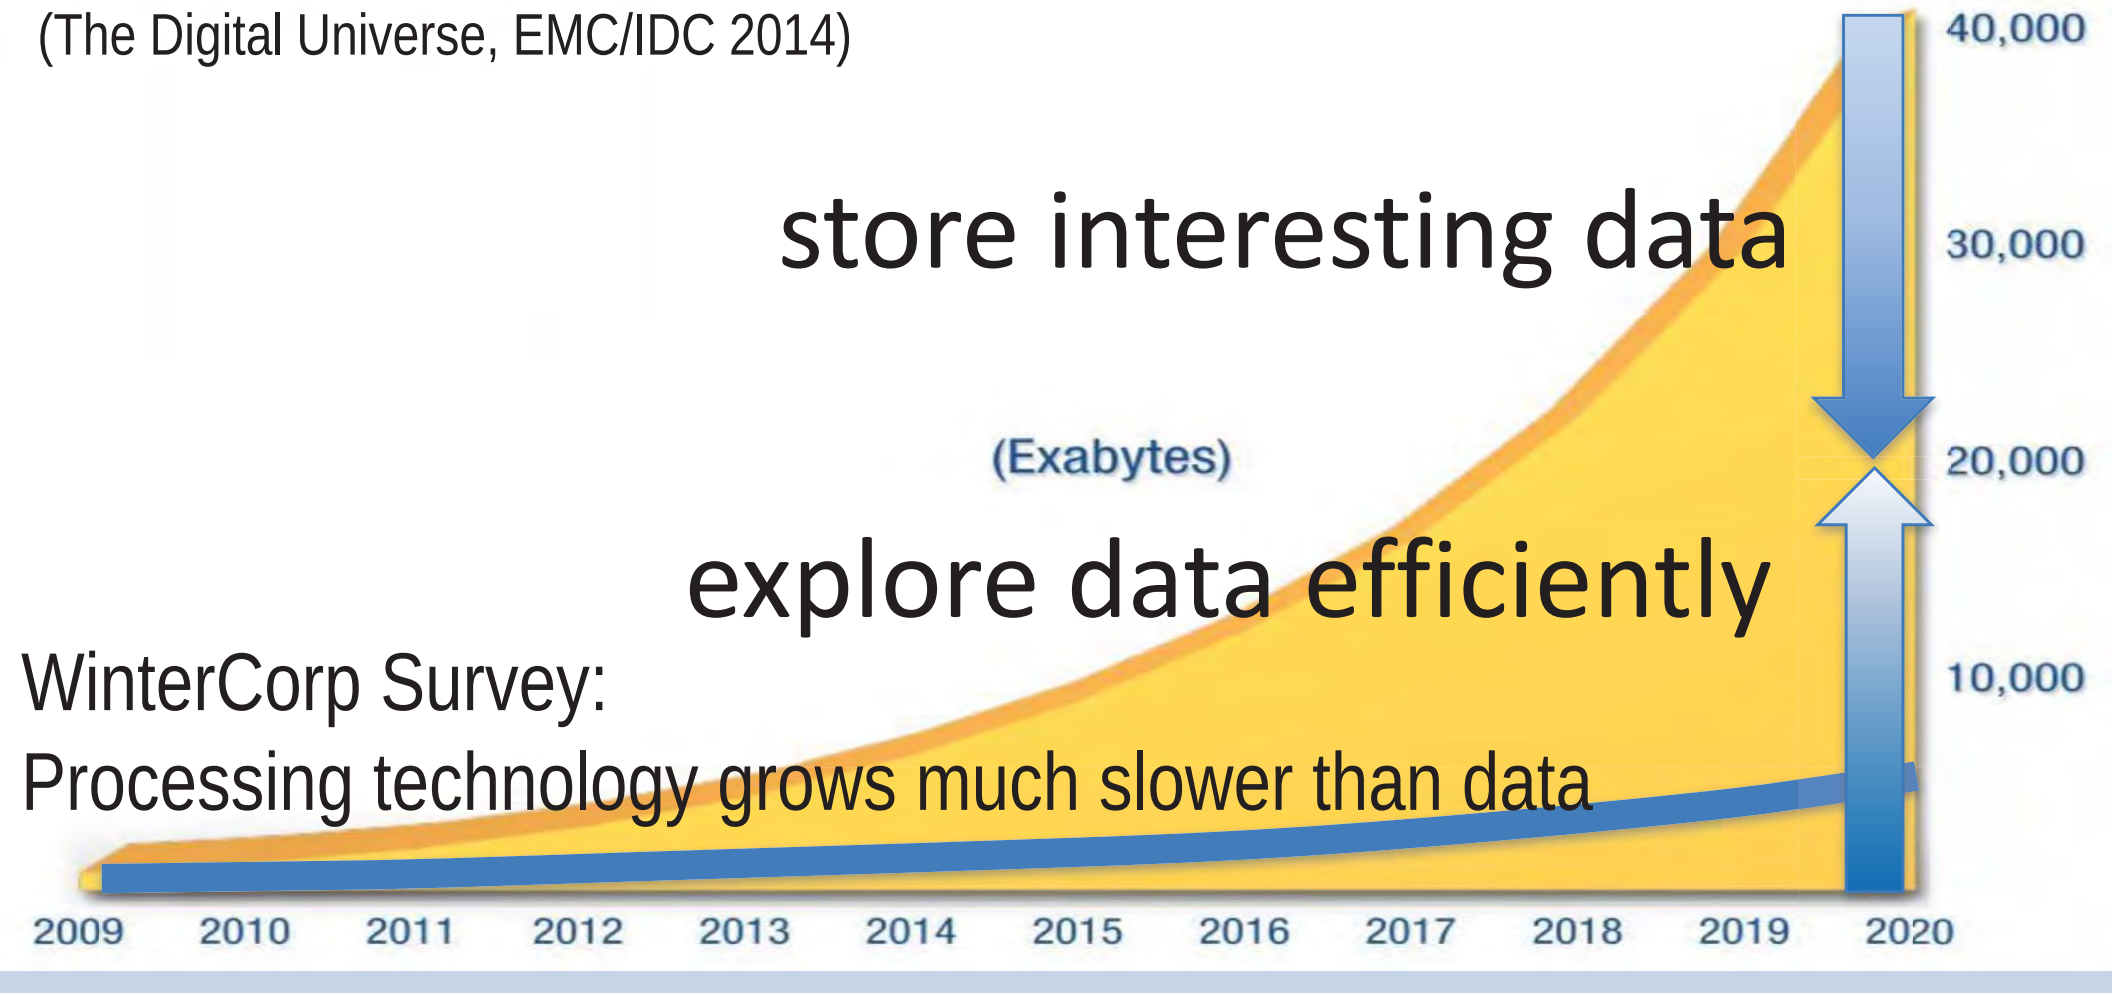


WinterCorp Survey:

Processing technology grows much slower than data

# build database to run queries

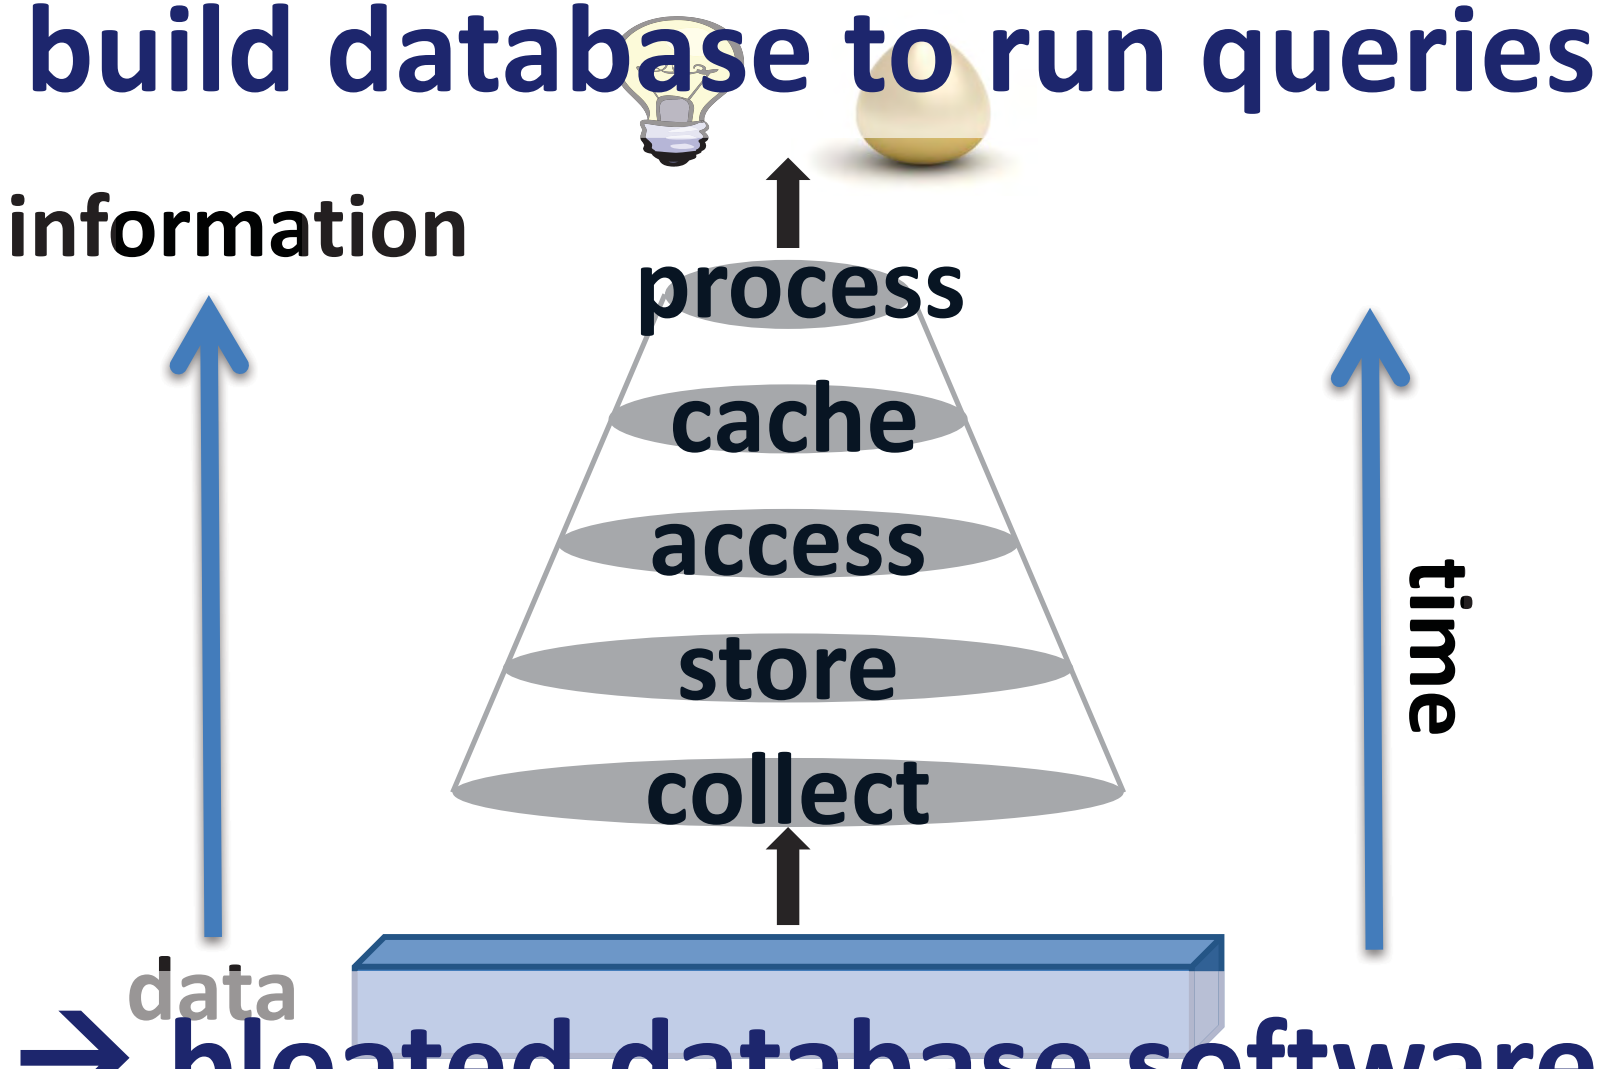


→ bloated database software

# new: one DB per app/data pair

**80% of analysts' time goes to data preparation and configuration**

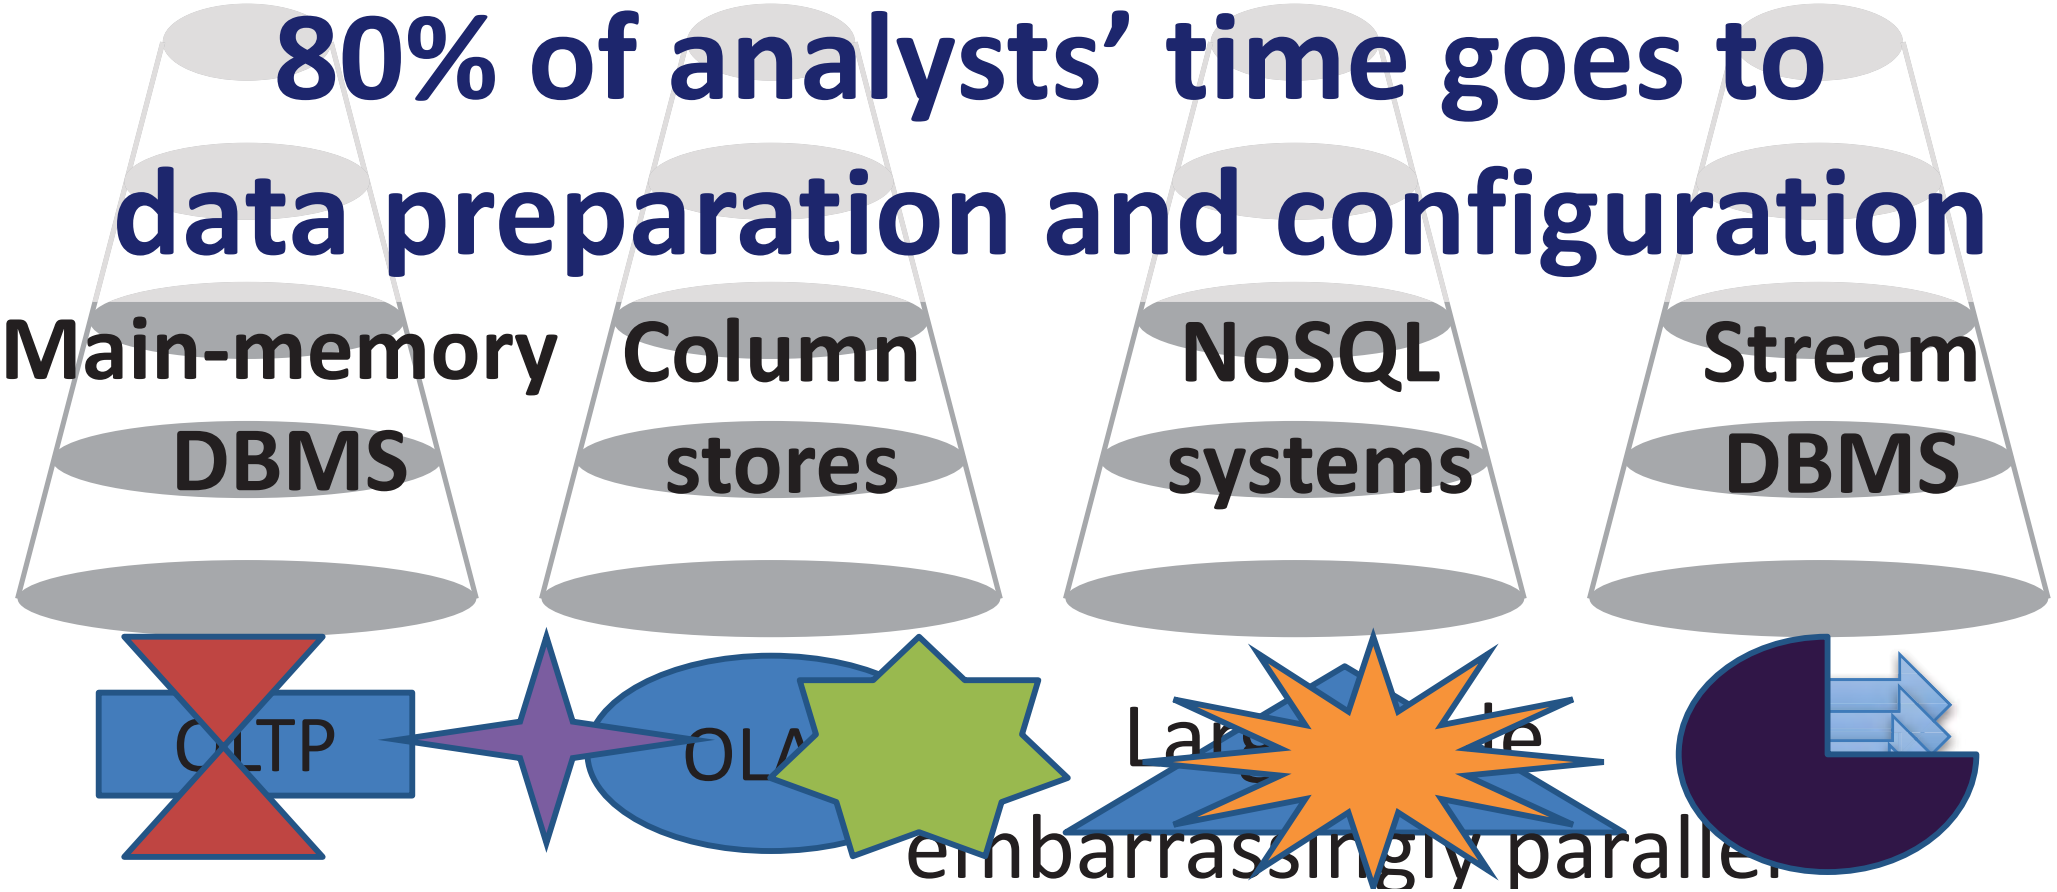


Databases will be extinct

# outline

- inspiring use case: the human brain project
  - brain simulation data
  - querying patient records
- a lean and agile database approach
  - adaptive query processing
  - SQL for all data

# outline

- **inspiring use case: the human brain project**
  - **simulation data sets**
  - **querying patient records**
- a lean and agile database approach
  - adaptive query processing
  - SQL for all data

# human brain project

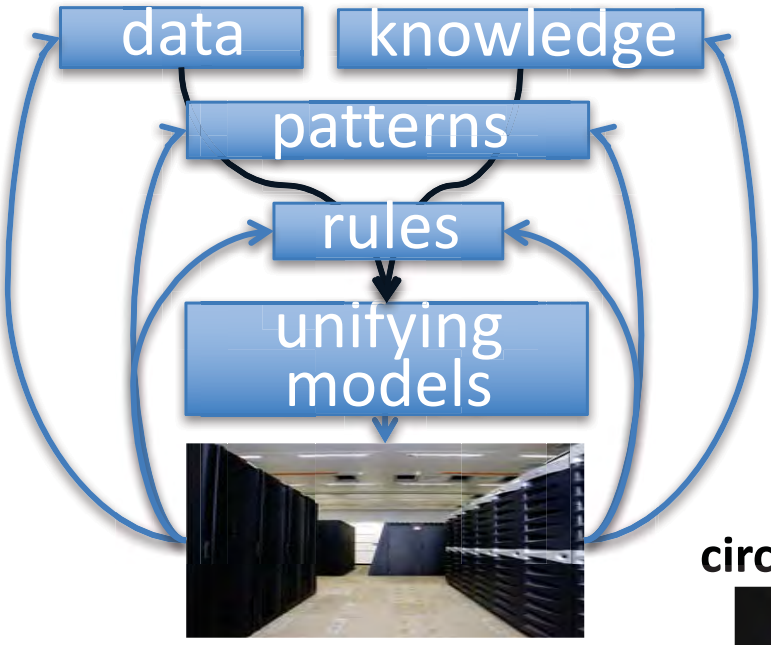


cognition

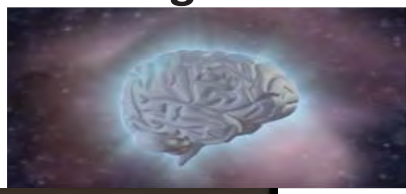

whole brain

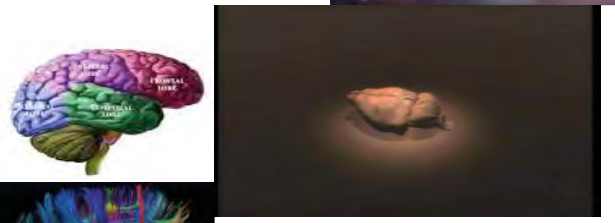

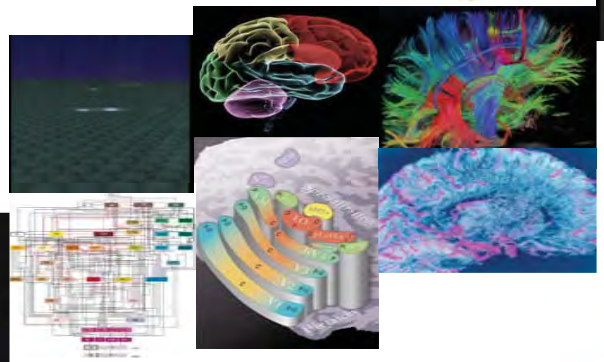

circuits

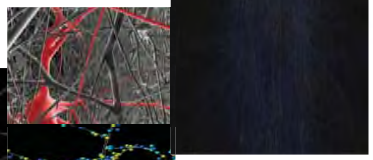

synapses

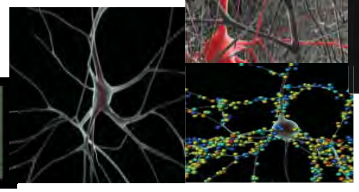

neurons


integrate clinical and simulation data

# brain simulation

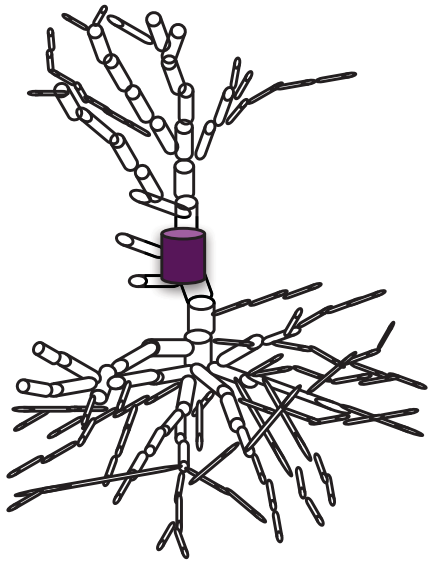

**single neuron,**
**modeled with 3D cylinders**

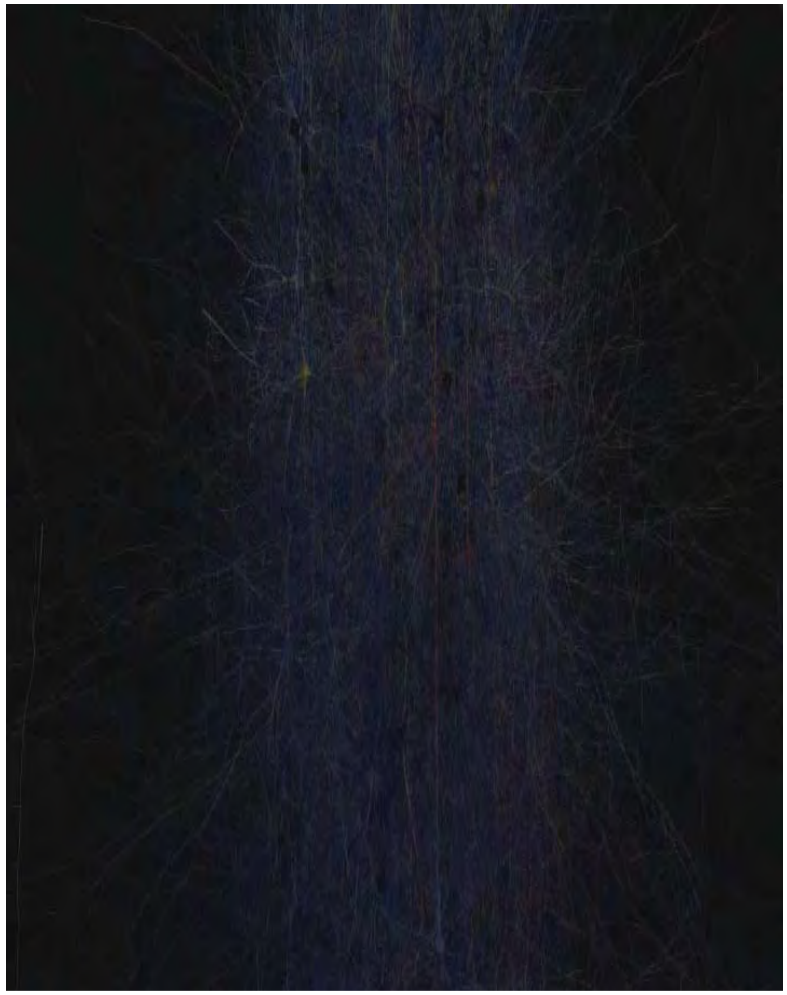

**database the brain**

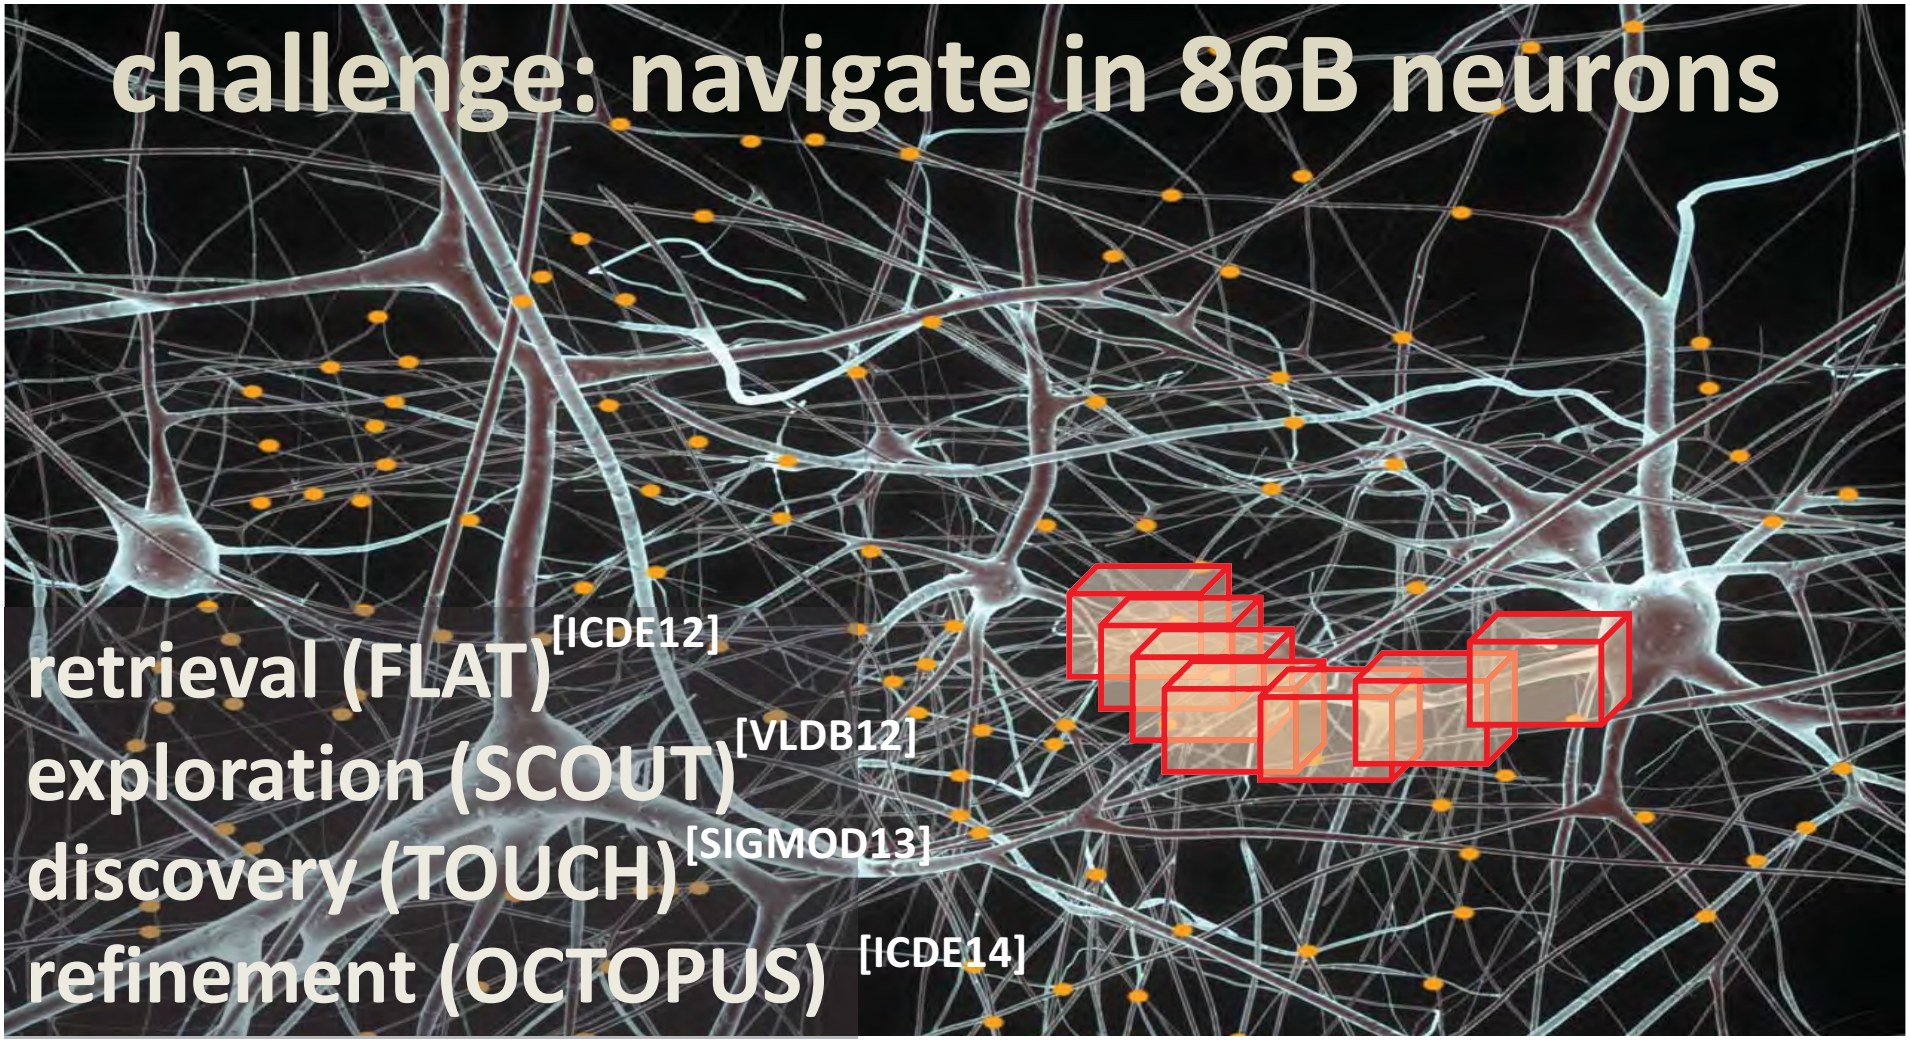


# 1000x simulations

# patient data: medical informatics


**health services**
**epidemiology data visualization**
**hypothesis testing**
**clinical trials**

# biological disease signatures

the coupling of
***clinical measurements*** with
***validated biomarkers***

**Example**: Alzheimer's disease

| Clinical - Phenotype | Proteomic Biomarkers | Genomic Biomarkers | [illegible]kers |
|---|---|---|---|
| **Cognition:** memory | **CSF protein**: beta amyloid [illegible] | [illegible]: APP, [illegible]SEN1, PSEN2 | **Volumetric change**: hippocampus, inferior temporal cortex… |
| **Functional capacity** | | | |
| **General physical heal[illegible]** | | **Common genetic variant**: APOE e4e4 … | **Beta amyloid imaging** |

memory loss/tau level/APOE e4e4/hippocampus atrophy/amyloid

# medical informatics platform

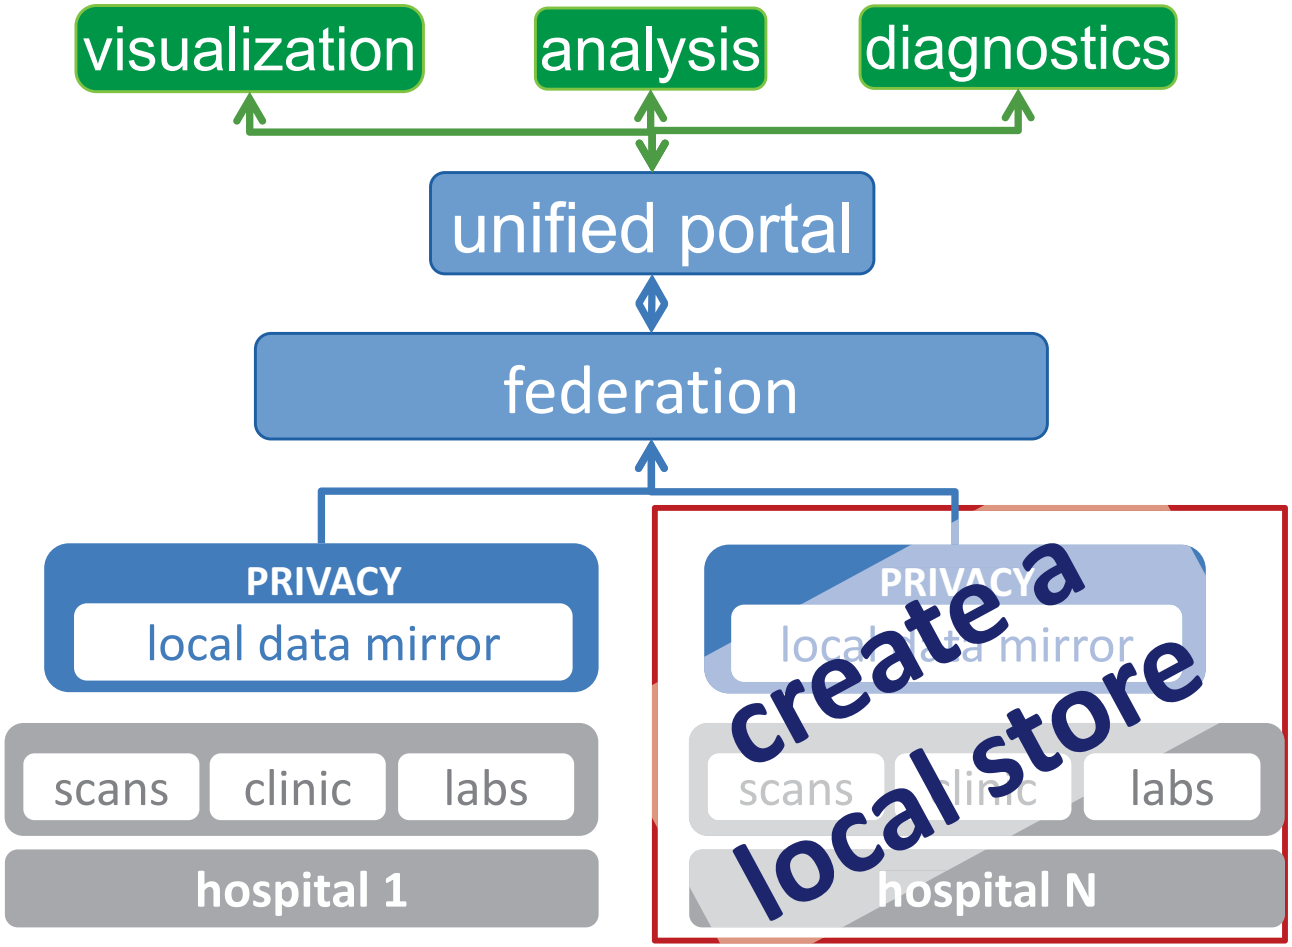

# image preprocessing (Lausanne)

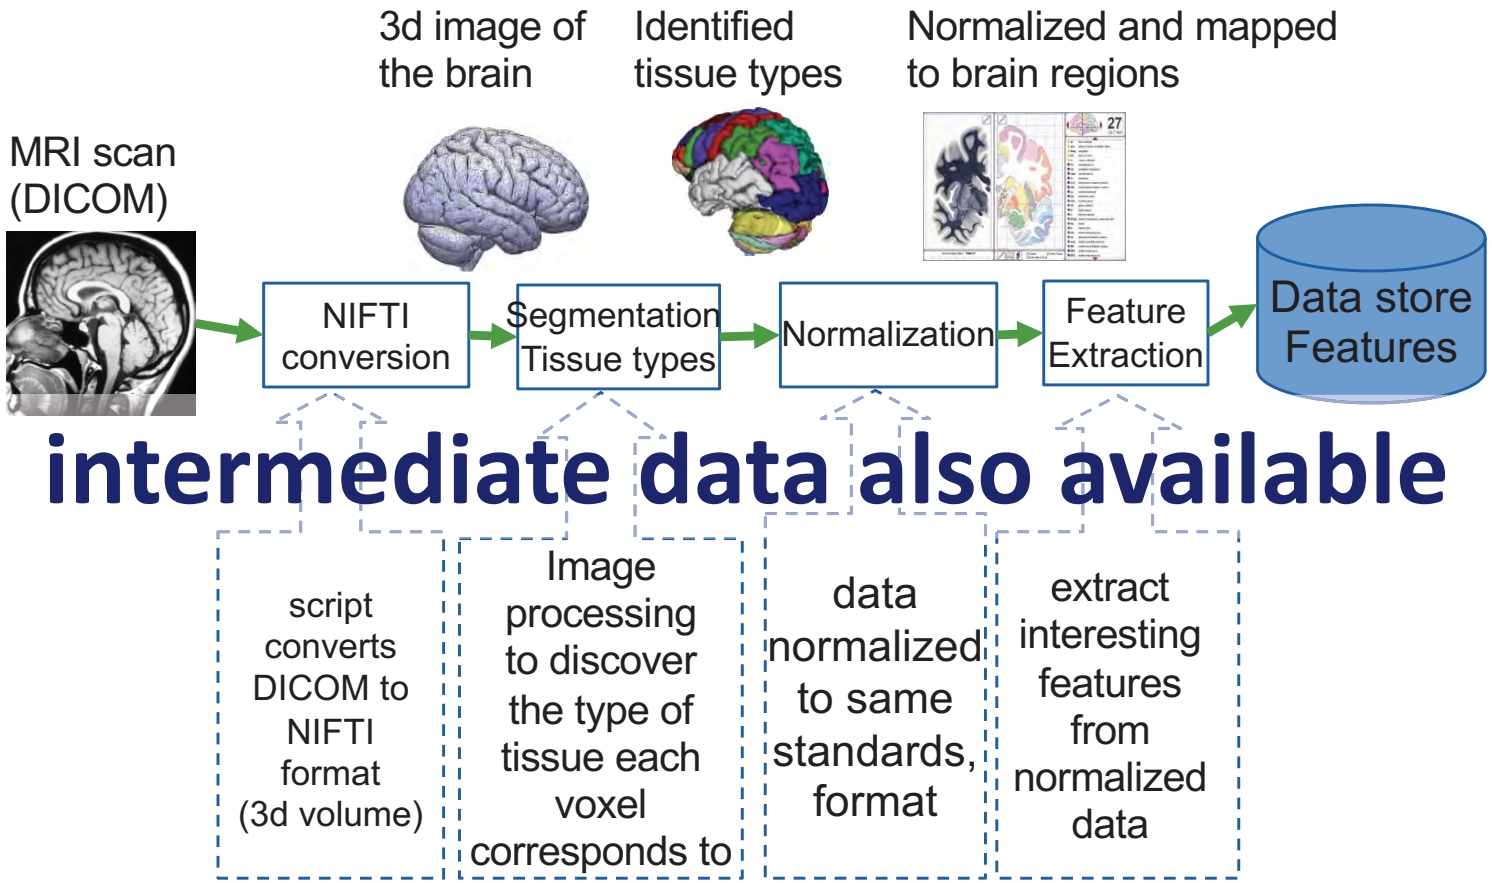

# two-level anonymization

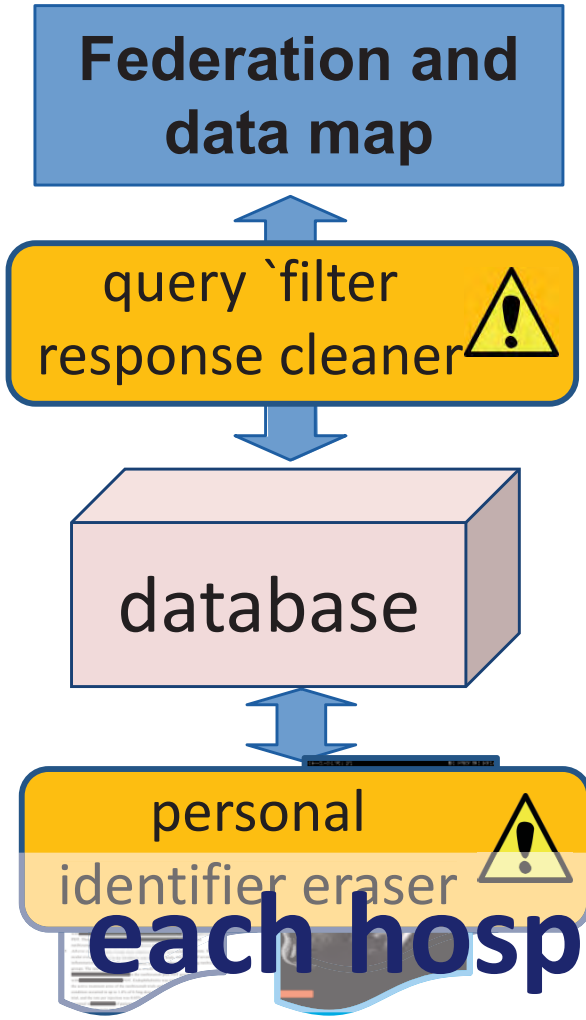


exclude fields
aggregate values
remove keywords

anon headers
match words
mask images

**each hospital's responsibility**

# access to private hospital data

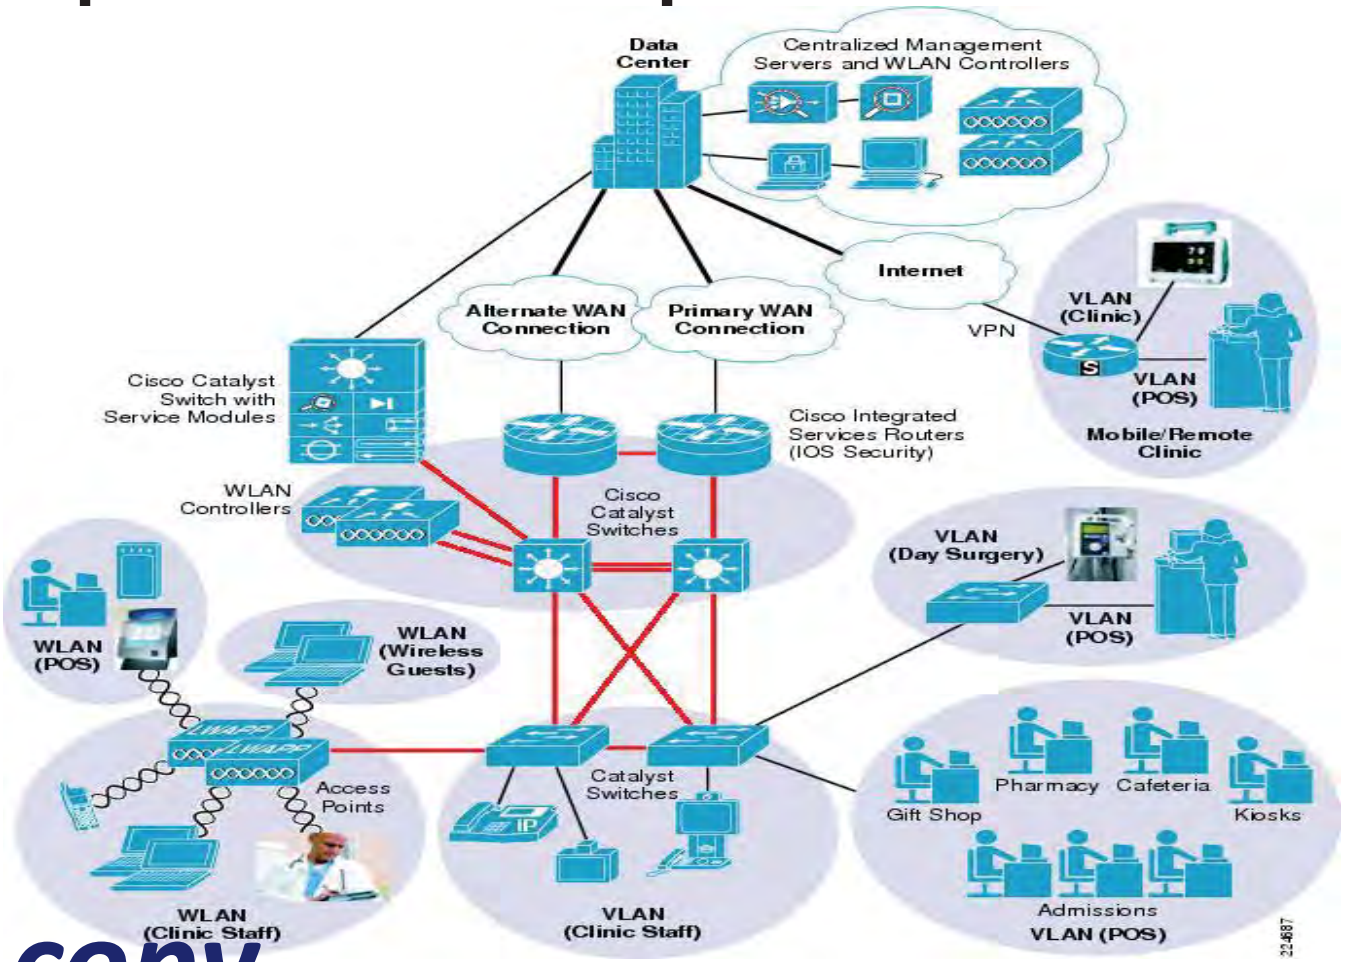


***no move, no copy***

# outline

- inspiring use case: the human brain project
  - simulation data sets
  - querying patient records
- **a lean and agile database approach**
  - **Just-in-time query processing**
  - **SQL for all data**

# clinical+genetic+imaging data → signature

## Patients (CSV)

| id | Protein: AACT | Age | Phenotype | … |
|---|---|---|---|---|
| 1 | 1.4 | 45 | Trauma | … |
| 2 | 2 | 55 | Chronic Symptoms | … |
| 3 | 0.2 | 56 | … | … |

## Brain_GrayMatter (Binary)

| | 0 | 1 | … | n |
|---|---|---|---|---|
| 0 | 0.45 | 0.75 | … | 0.1 |
| 1 | 0.33 | 0.3 | … | 0.38 |
| … | … | … | … | … |
| m | 0.12 | 0 | … | 0.47 |

**signature:**

***age > 50***

***AND***

***amygdala.Vol > 0.3***

***AND***

***AACT < 1***

## BrainRegions (JSON)

```
[{"id": 1,
 "amygdala": {"X":15,"Y":20, “Vol”: 0.5},
 "hippocampus": {"X":17, "Y":10, “Vol”:0.2}},
{"id": 2, ...},
{"id": 3, ...}]
```

# queries on heterogeneous data

**cannot load into a Database System!**

- diverse formats
- legacy software
- privacy limitations
- data “owned” by one database

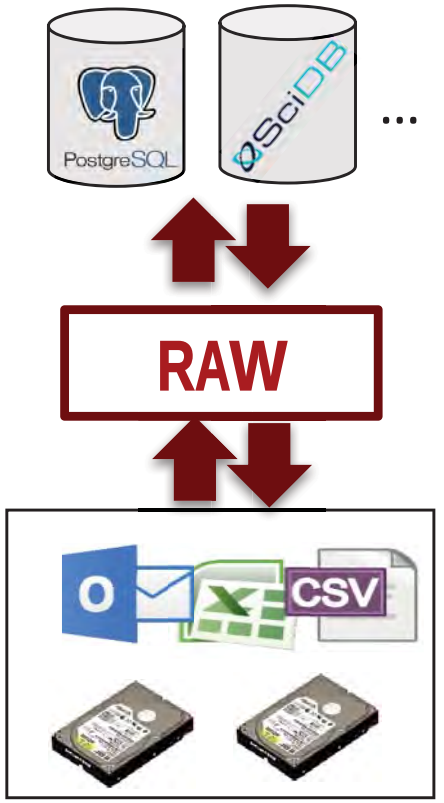


**RAW: interface to raw data**

SQL, SCALA, notebooks

code-generated engine

**key: data virtualization**

# Adapting a query engine to data

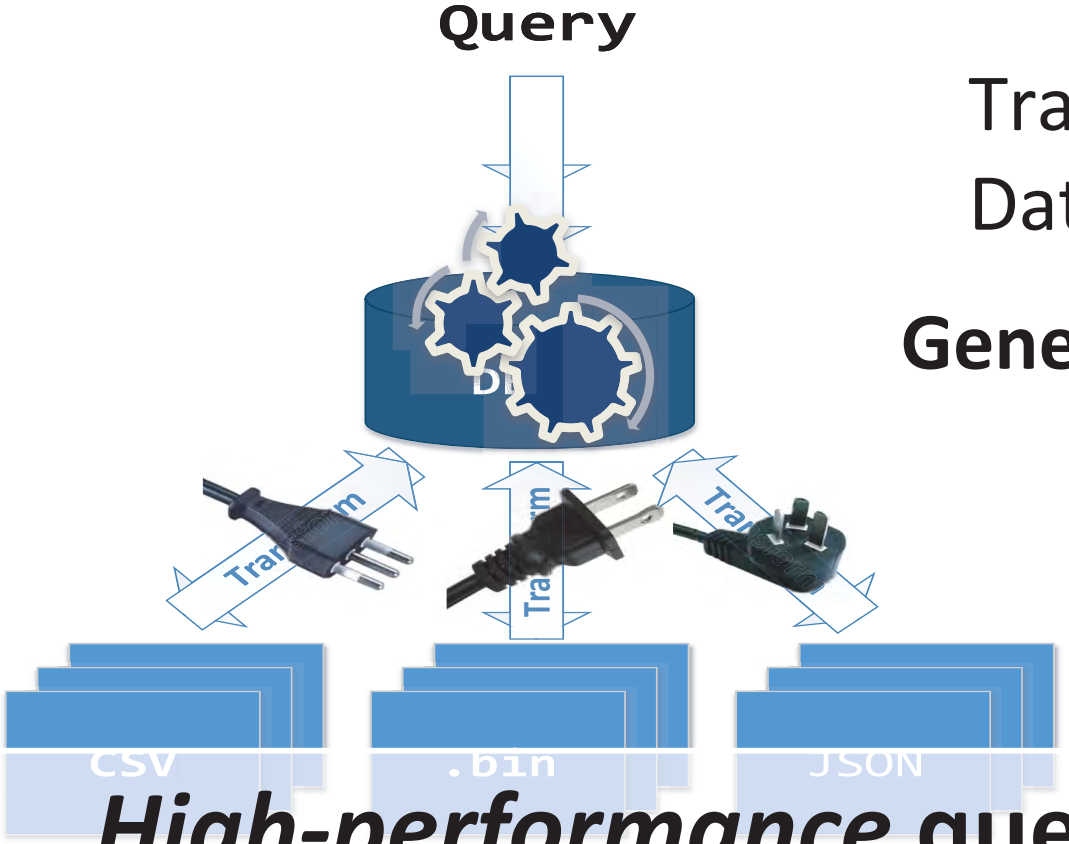


Traditional:
Data adapts to engine

**Generate plug-in** per data source

Treat each source as
*native storage format*

***High-performance* querying… while using original data formats, files, and scripts**

# How it works

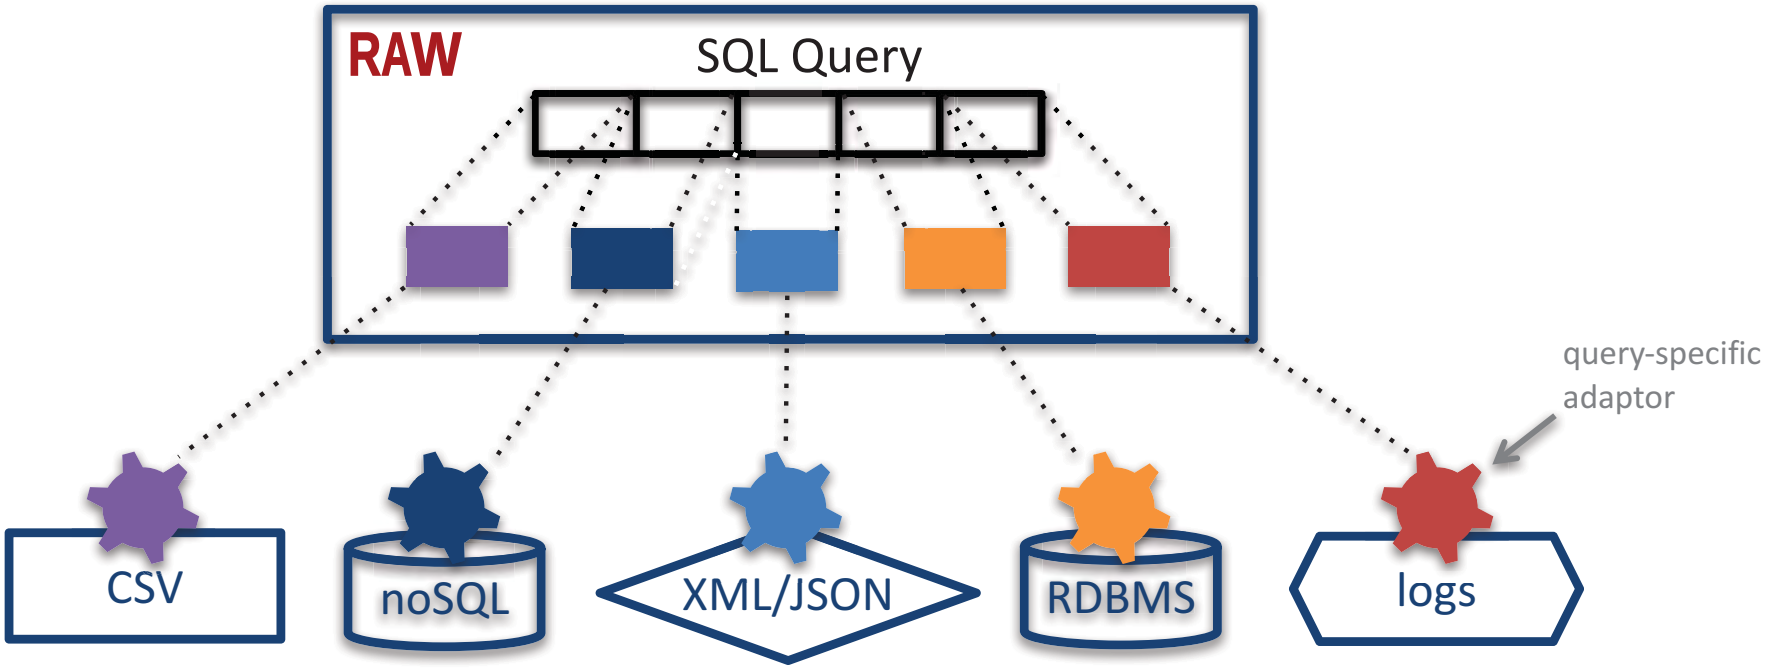


**query-specific adaptors auto-generated**

**Tools and data cached for efficiency**

As queries run, RAW gathers information on usage of data and generates software tools.

# RAW

just-built useful tools

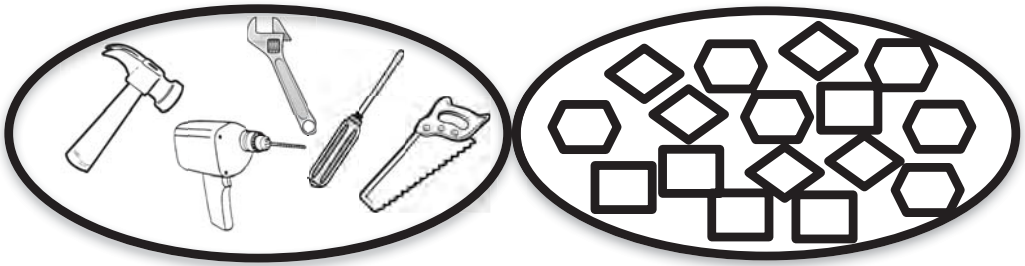

just-gathered query-specific data

query-specific adaptor

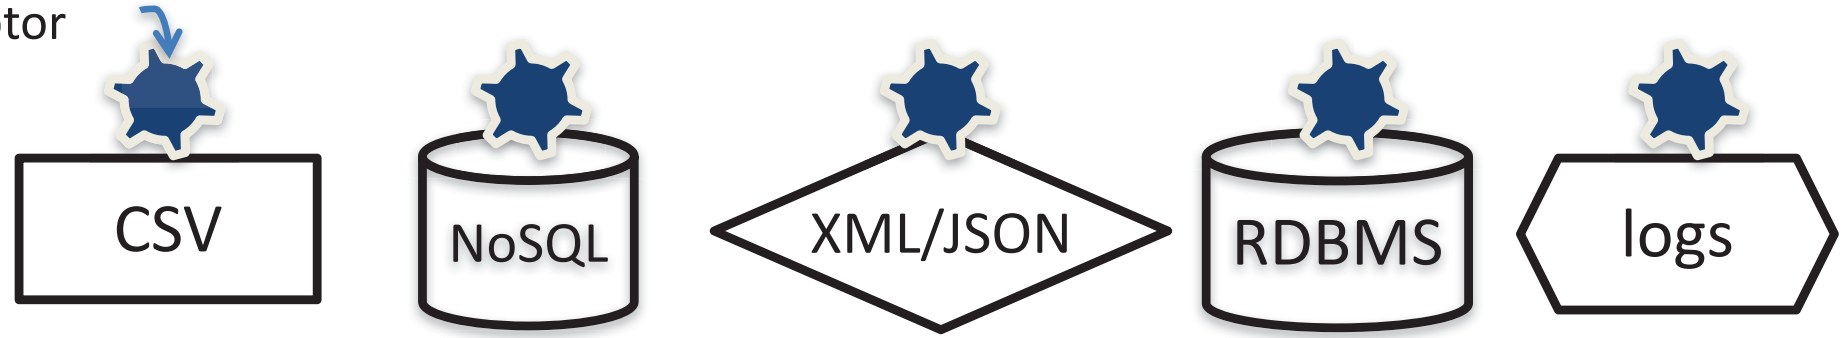

# What you can do with RAW 1/3

- Query raw files (CSV, JSON, XML, …) directly *without ETL:*
  - without data ingestion
  - without data cleaning
  - without “flattening hierarchies”
  - … i.e. without “losing information” during ETL

- How?
  - With an extended SQL language
  - No building scripts or using separate ETL tools on the side

# What you can do with RAW 2/3

- Query datasets regardless of size or location:
  - Small CSV or Excel file on your Dropbox
  - PBs of Parquet data on Hadoop
  - Tables on a RDBMS

- How?
  - RAW code-generates to your execution environment: local node, Hadoop/Spark cluster, RDBMS backend, …

# What you can do with RAW 3/3

- Build and preview your queries/reporting/analytics, live, step-by-step
- Build dashboards with a few clicks
- How?
  - The user interface is built into RAW's engine.
  - RAW comes with an IDE.

RAW
Just ask.

Queries Sources

Settings

**Your Queries** 6 New +

Search query...

- patients_overview
- patients_gender_city
- patients_map
- authors_with_ratings
- authors

**Your Sources** 7 Add +

Search schema...

- authors
- patients
- publications
- author_ratings
- log_example

**Editor**

```
1 select d.description, d.sub_category, d.category
2   from patients p,
3   p.patientBasedData.diagnosis d
4
```

Plot sunburst Full results

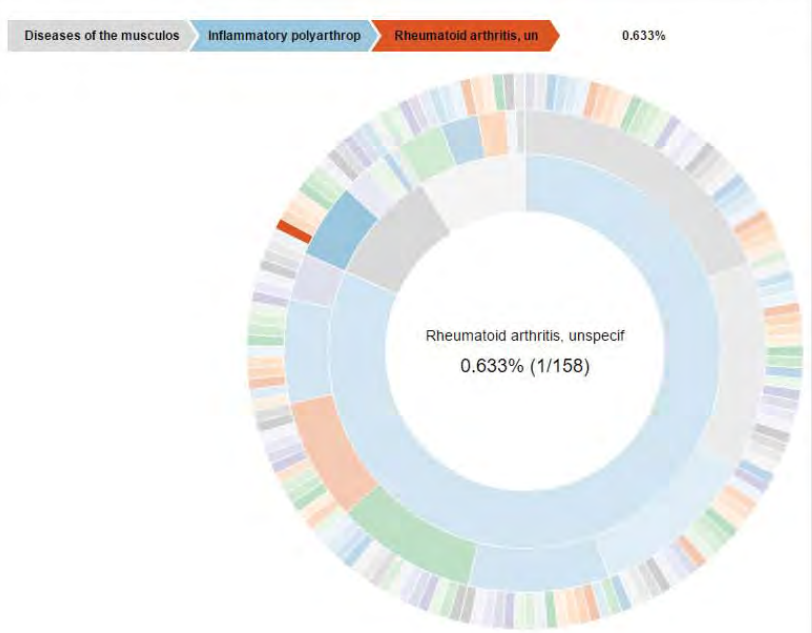


Editor
Plot
No plots available
Query failed during execution
select * from log_example parse as r"""(\d+)"""
incompatible input: 199.72.81.55 - - [01/Jul/2015:00:00:01 -0400] "GET /history/apollo/ HTTP/1.0" 200 6245

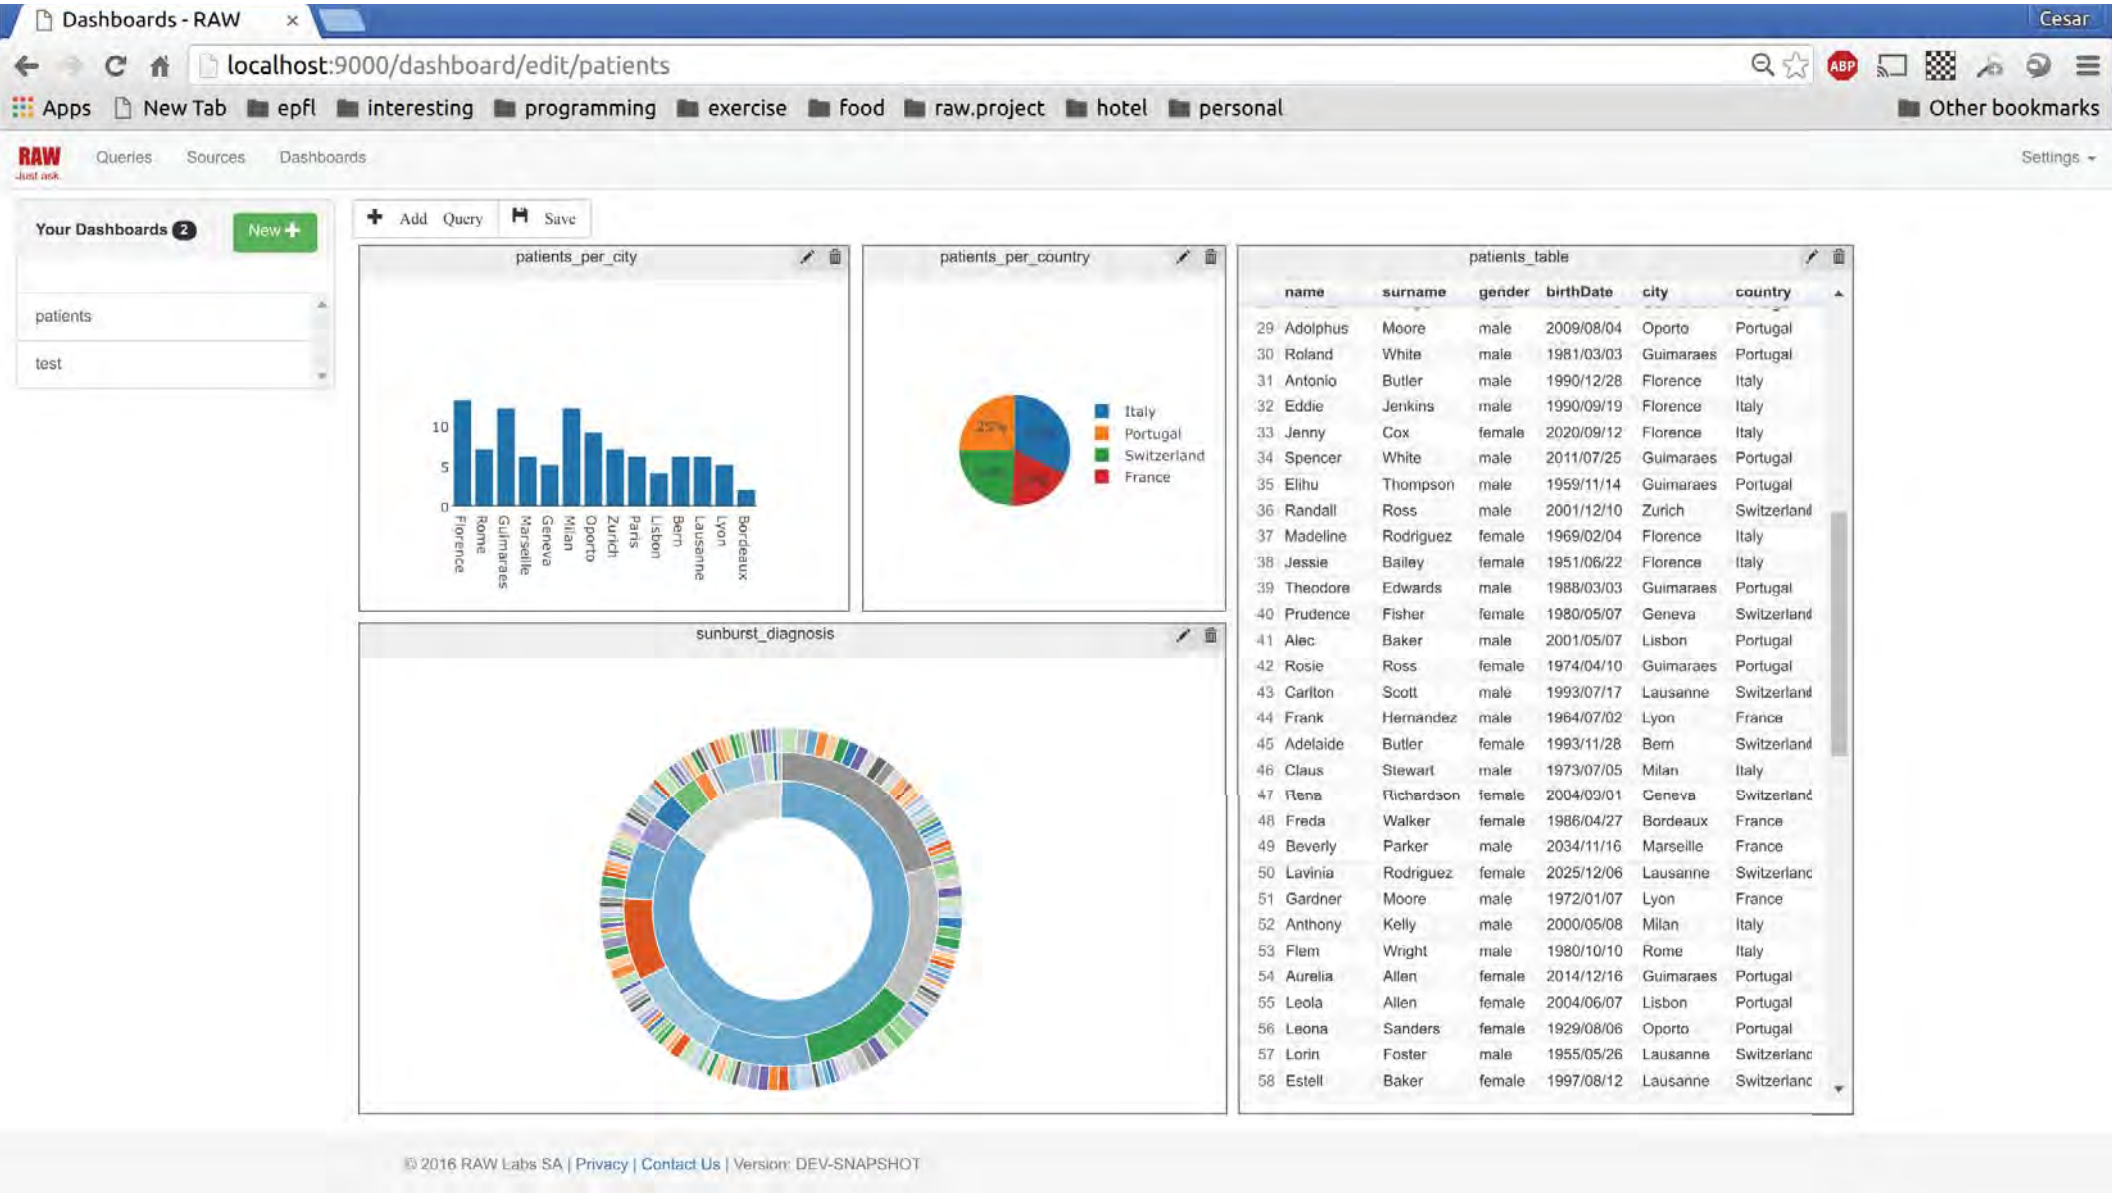
Dashboards - RAW
Cesar
localhost:9000/dashboard/edit/patients
Apps
New Tab
epfl
interesting
programming
exercise
food
raw.project
hotel
personal
Other bookmarks
RAW
Just ask.
Queries
Sources
Dashboards
Settings
Your Dashboards 2
New
patients
test
Add Query
Save
patients_per_city
Florence
Rome
Guimaraes
Marseille
Geneva
Milan
Oporto
Zurich
Paris
Lisbon
Bern
Lausanne
Lyon
Bordeaux
patients_per_country
Italy
Portugal
Switzerland
France
sunburst_diagnosis
patients_table
name surname gender birthDate city country
29 Adolphus Moore male 2009/08/04 Oporto Portugal
30 Roland White male 1981/03/03 Guimaraes Portugal
31 Antonio Butler male 1990/12/28 Florence Italy
32 Eddie Jenkins male 1990/09/19 Florence Italy
33 Jenny Cox female 2020/09/12 Florence Italy
34 Spencer White male 2011/07/25 Guimaraes Portugal
35 Elihu Thompson male 1959/11/14 Guimaraes Portugal
36 Randall Ross male 2001/12/10 Zurich Switzerland
37 Madeline Rodriguez female 1969/02/04 Florence Italy
38 Jessie Bailey female 1951/06/22 Florence Italy
39 Theodore Edwards male 1988/03/03 Guimaraes Portugal
40 Prudence Fisher female 1980/05/07 Geneva Switzerland
41 Alec Baker male 2001/05/07 Lisbon Portugal
42 Rosie Ross female 1974/04/10 Guimaraes Portugal
43 Carlton Scott male 1993/07/17 Lausanne Switzerland
44 Frank Hernandez male 1964/07/02 Lyon France
45 Adelaide Butler female 1993/11/28 Bern Switzerland
46 Claus Stewart male 1973/07/05 Milan Italy
47 Rena Richardson female 2004/09/01 Geneva Switzerland
48 Freda Walker female 1986/04/27 Bordeaux France
49 Beverly Parker male 2034/11/16 Marseille France
50 Lavinia Rodriguez female 2025/12/06 Lausanne Switzerland
51 Gardner Moore male 1972/01/07 Lyon France
52 Anthony Kelly male 2000/05/08 Milan Italy
53 Flem Wright male 1980/10/10 Rome Italy
54 Aurelia Allen female 2014/12/16 Guimaraes Portugal
55 Leola Allen female 2004/06/07 Lisbon Portugal
56 Leona Sanders female 1929/08/06 Oporto Portugal
57 Lorin Foster male 1955/05/26 Lausanne Switzerland
58 Estell Baker female 1997/08/12 Lausanne Switzerland
© 2016 RAW Labs SA | Privacy | Contact Us | Version: DEV-SNAPSHOT

# deployment: hospital data mirror

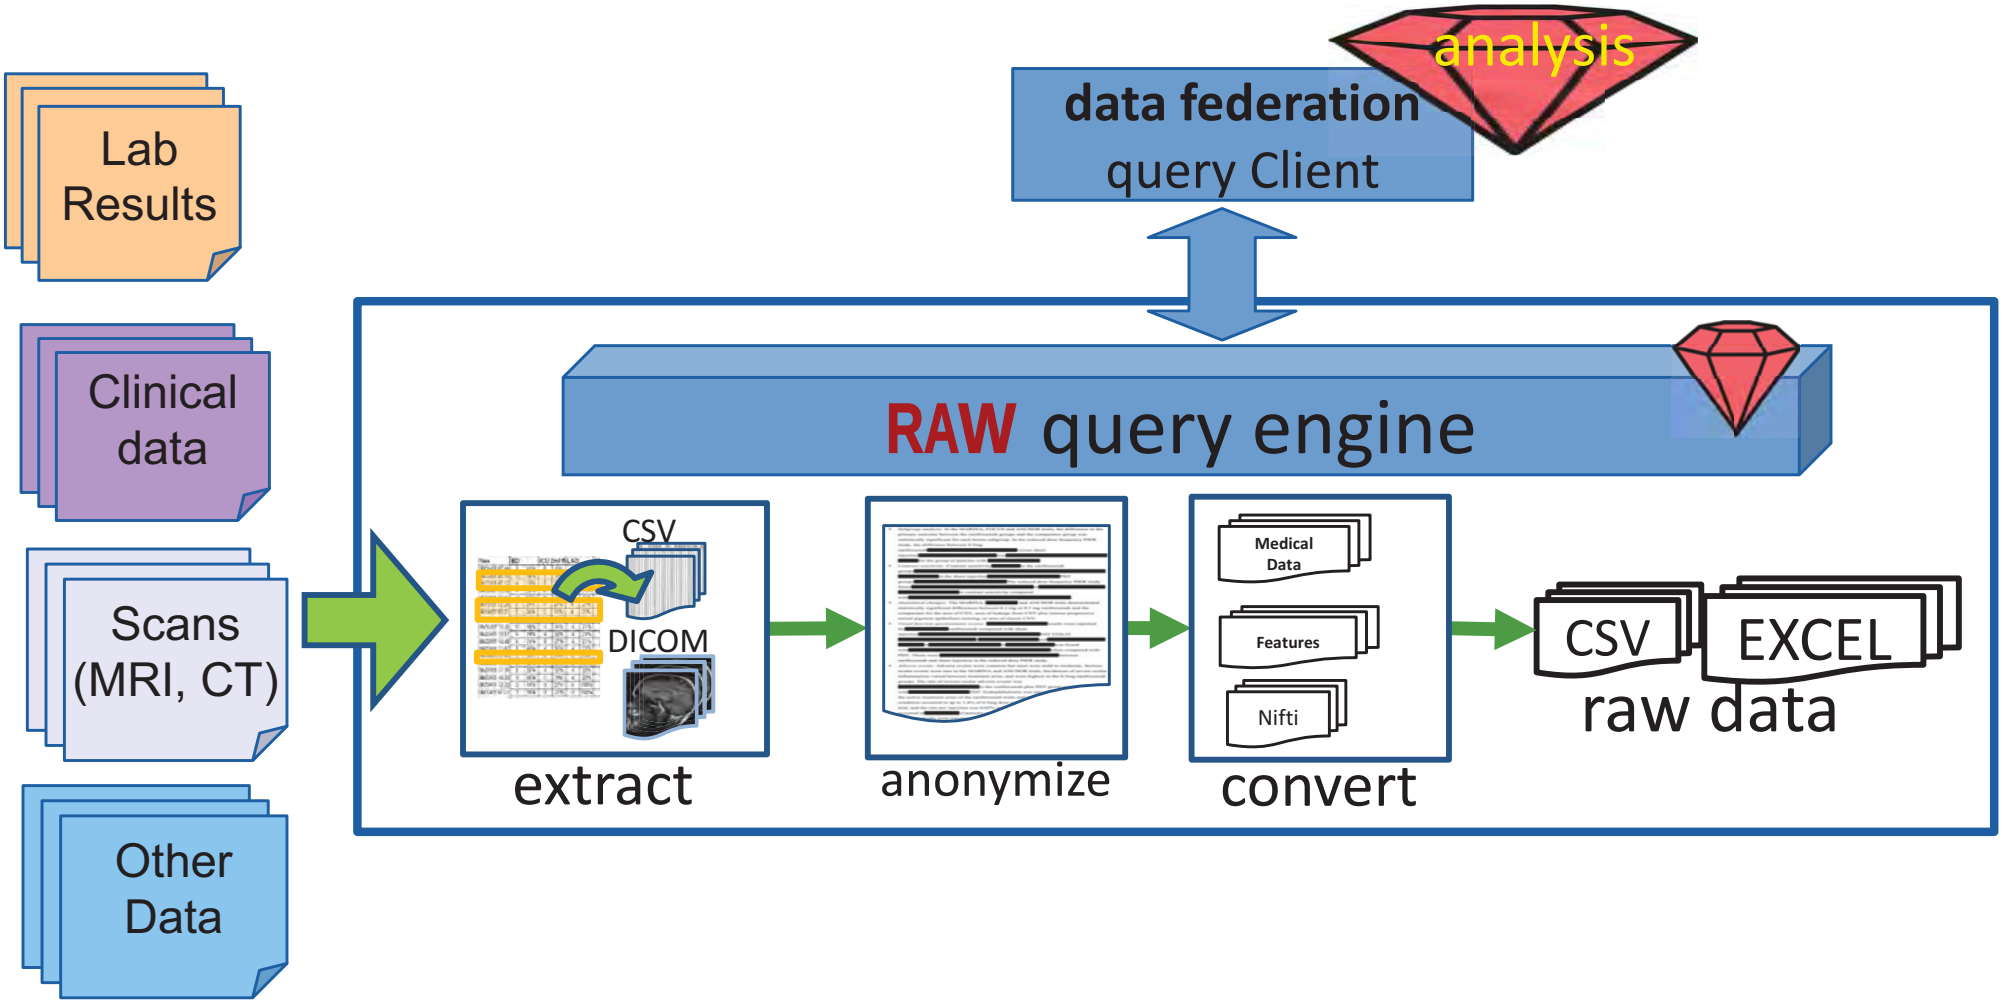

# Medical Informatics Platform

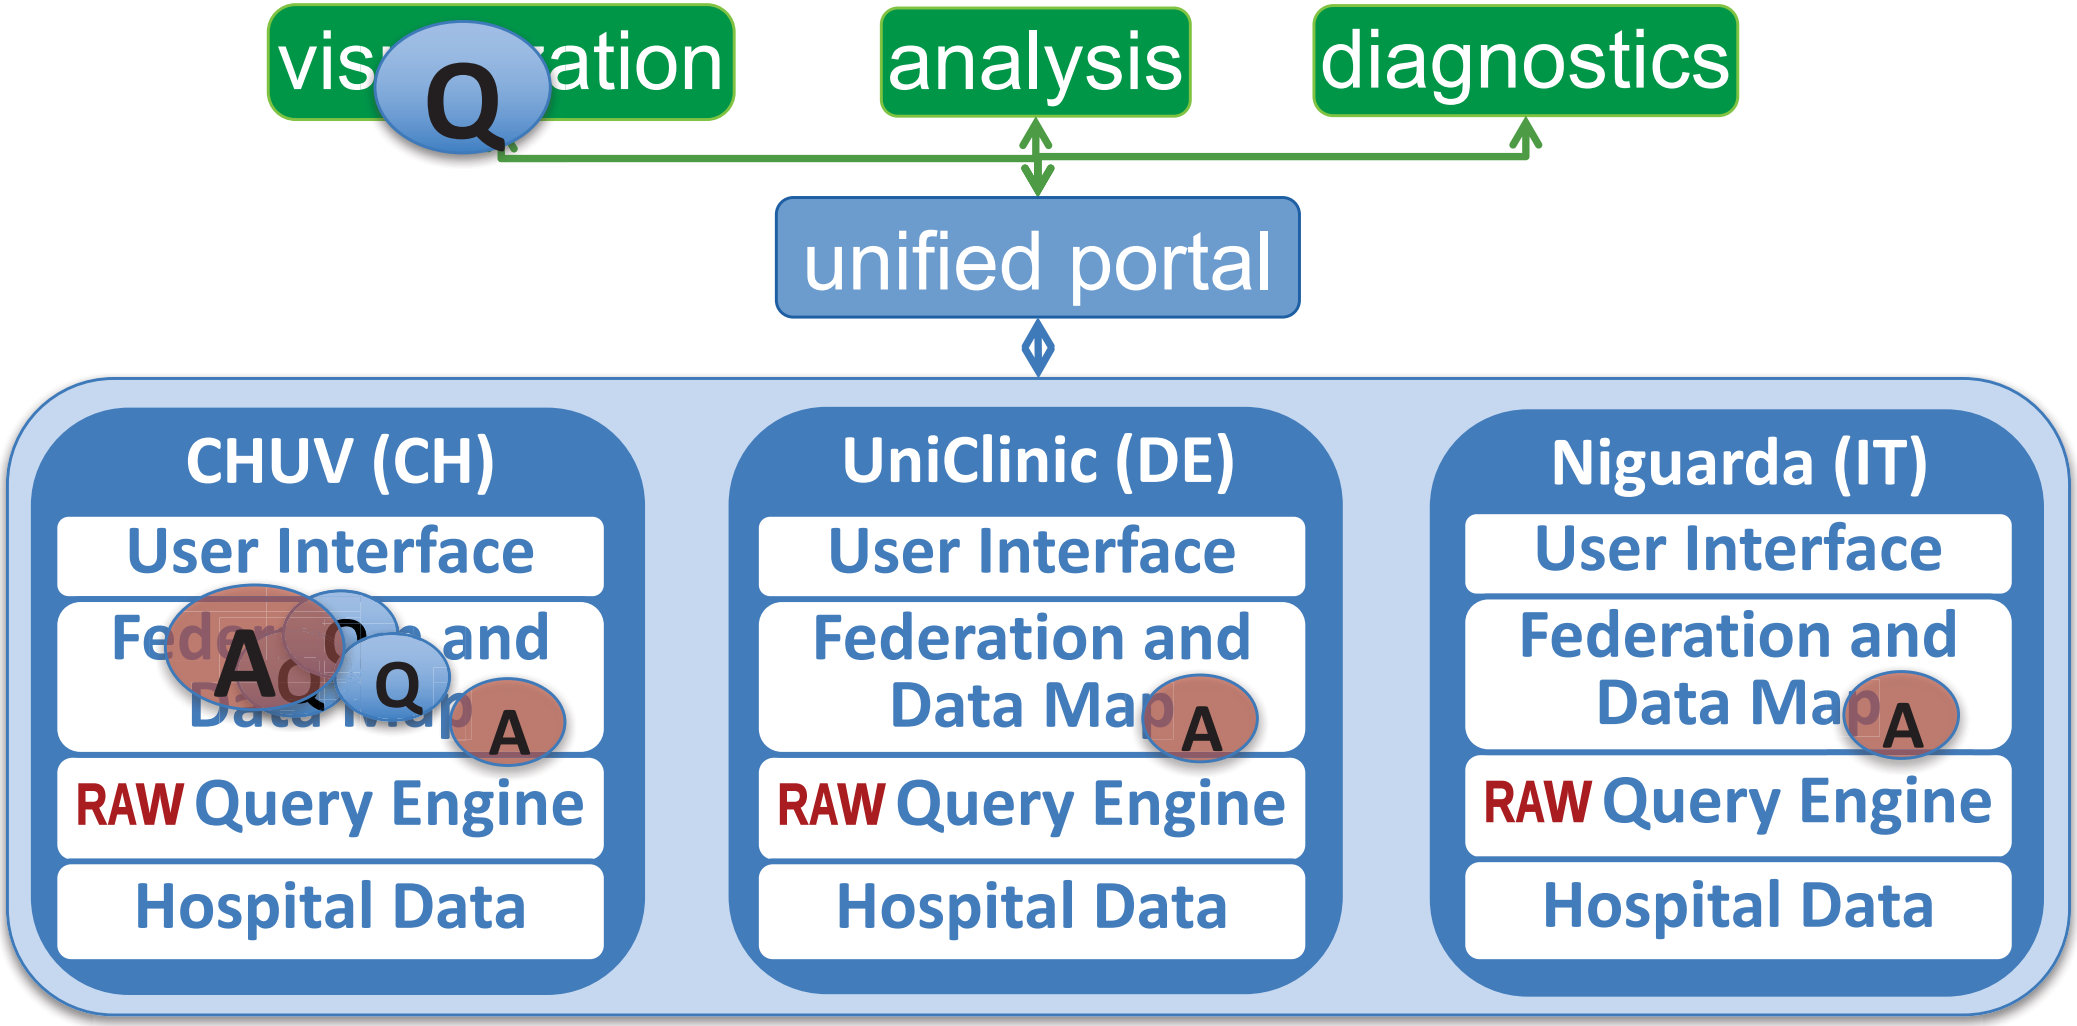

# the user's perspective


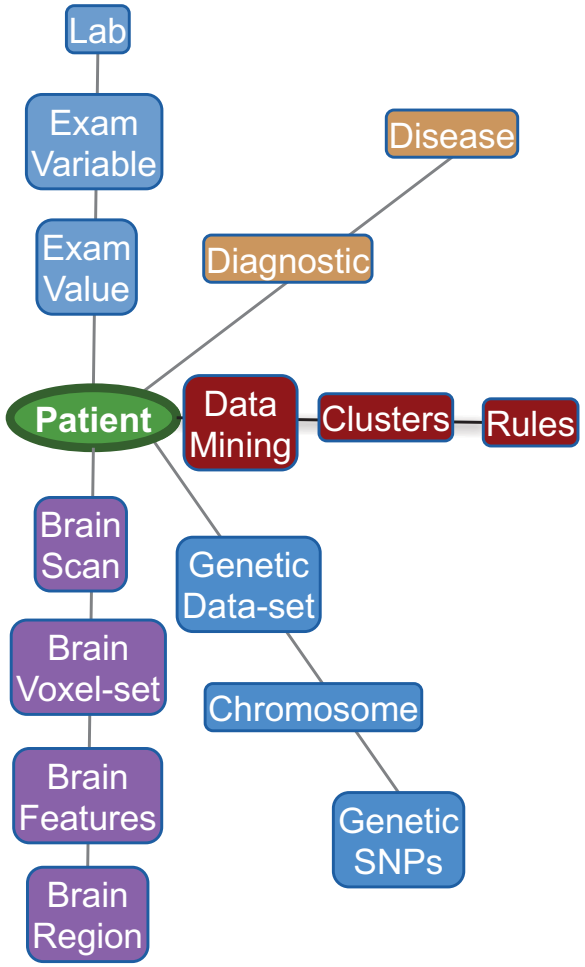

# Summary

- currently data management cost grows with data owned
- impossible to pre-cook a database system suitable for all data
- from manual ingestion to automatic adaptation: just-in-time query processing

***Thank you!***

# References

**[ICDE05]** M. Stonebraker, U. Çetintemel: "One Size Fits All": An Idea Whose Time Has Come and Gone. ICDE 2005

**[CIDR07]** S. Idreos, M. Kersten, S. Manegold: Database Cracking. CIDR 2007

**[CIDR11]** S. Idreos, I. Alagiannis, R. Johnson, A. Ailamaki. Here are my Data Files. Here are my Queries. Where are my Results? CIDR 2011

**[VLDB11]** T. Neumann: Efficiently Compiling Efficient Query Plans for Modern Hardware. VLDB 2011

**[SIGMOD12]** I. Alagiannis, R. Borovica, M. Branco, S. Idreos, A. Ailamaki: NoDB: efficient query execution on raw data files. SIGMOD 2012

**[ICDE12]** F. Tauheed, L. Biveinis, T. Heinis, F. Schürmann, H. Markram, A. Ailamaki: Accelerating Range Queries for Brain Simulations. ICDE 2012

# References

**[IISWC12]** C. Abad, N. Roberts, Y. Lu, and R. Campbell: A Storage-Centric Analysis of MapReduce Workloads: File Popularity, Temporal Locality and Arrival Patterns, IISWC 2012

**[VLDB12]** F. Tauheed, T. Heinis, F. Schürmann, H. Markram, A. Ailamaki: SCOUT: Prefetching for Latent Feature Following Queries. VLDB 2012

**[VLDB12a]** Y. Chen, S. Alspaugh, and R. Katz: Interactive Query Processing in Big Data Systems: A Cross-Industry Study of MapReduce Workloads, VLDB 2012

**[VLDB12b]** Y. Ahmad, O. Kennedy, C. Koch, M. Nikolic: DBToaster: Higher-order Delta Processing for Dynamic, Frequently Fresh Views. VLDB 2012

**[SSDBM12]** M. Ivanova, M. Kersten, S. Manegold: Data Vaults: A Symbiosis between Database Technology and Scientific File Repositories. SSDBM 2012

# References

**[SIGMOD13]** Sadegh Nobari, Farhan Tauheed, Thomas Heinis, Panagiotis Karras, Stéphane Bressan, Anastasia Ailamaki: TOUCH: in-memory spatial join by hierarchical data-oriented partitioning. SIGMOD 2013

**[SIGMOD13a]** D. DeWitt, A. Halverson, R. Nehme, S. Shankar, J. Aguilar-Saborit, A. Avanes, M. Flasza, J. Gramling: Split query processing in polybase. SIGMOD 2013

**[EDBT13]** Pinar Tözün, Ippokratis Pandis, Cansu Kaynak, Djordje Jevdjic, Anastasia Ailamaki: From A to E: analyzing TPC's OLTP benchmarks: the obsolete, the ubiquitous, the unexplored. EDBT 2013

**[EDBT13a]** A. Abouzied, D. Abadi, A. Silberschatz: Invisible loading: access-driven data transfer from raw files into database systems. EDBT 2013

**[DAMON13]** Pinar Tözün, Brian Gold, Anastasia Ailamaki: OLTP in wonderland: where do cache misses come from in major OLTP components? DaMoN 2013

# References

**[SIGMOD14]** I. Alagiannis, S. Idreos, A. Ailamaki: H2O: A Hands-Free Adaptive Store. SIGMOD 2014

**[ICDE14]** Farhan Tauheed, Thomas Heinis, Felix Shurmann, Henry Markram, Anastasia Ailamaki: OCTOPUS: Efficient Query Execution on Dynamic Mesh Datasets. ICDE 2014

**[VLDB14]** Manos Karpathiotakis, Miguel Branco, Ioannis Alagiannis, Anastasia Ailamaki: Adaptive Query Processing on RAW Data. PVLDB 7(12) (2014)

**[VLDB16]** Manos Karpathiotakis, Ioannis Alagiannis, Anastasia Ailamaki: Fast Queries Over Heterogeneous Data Through Engine Customization. VLDB 2016