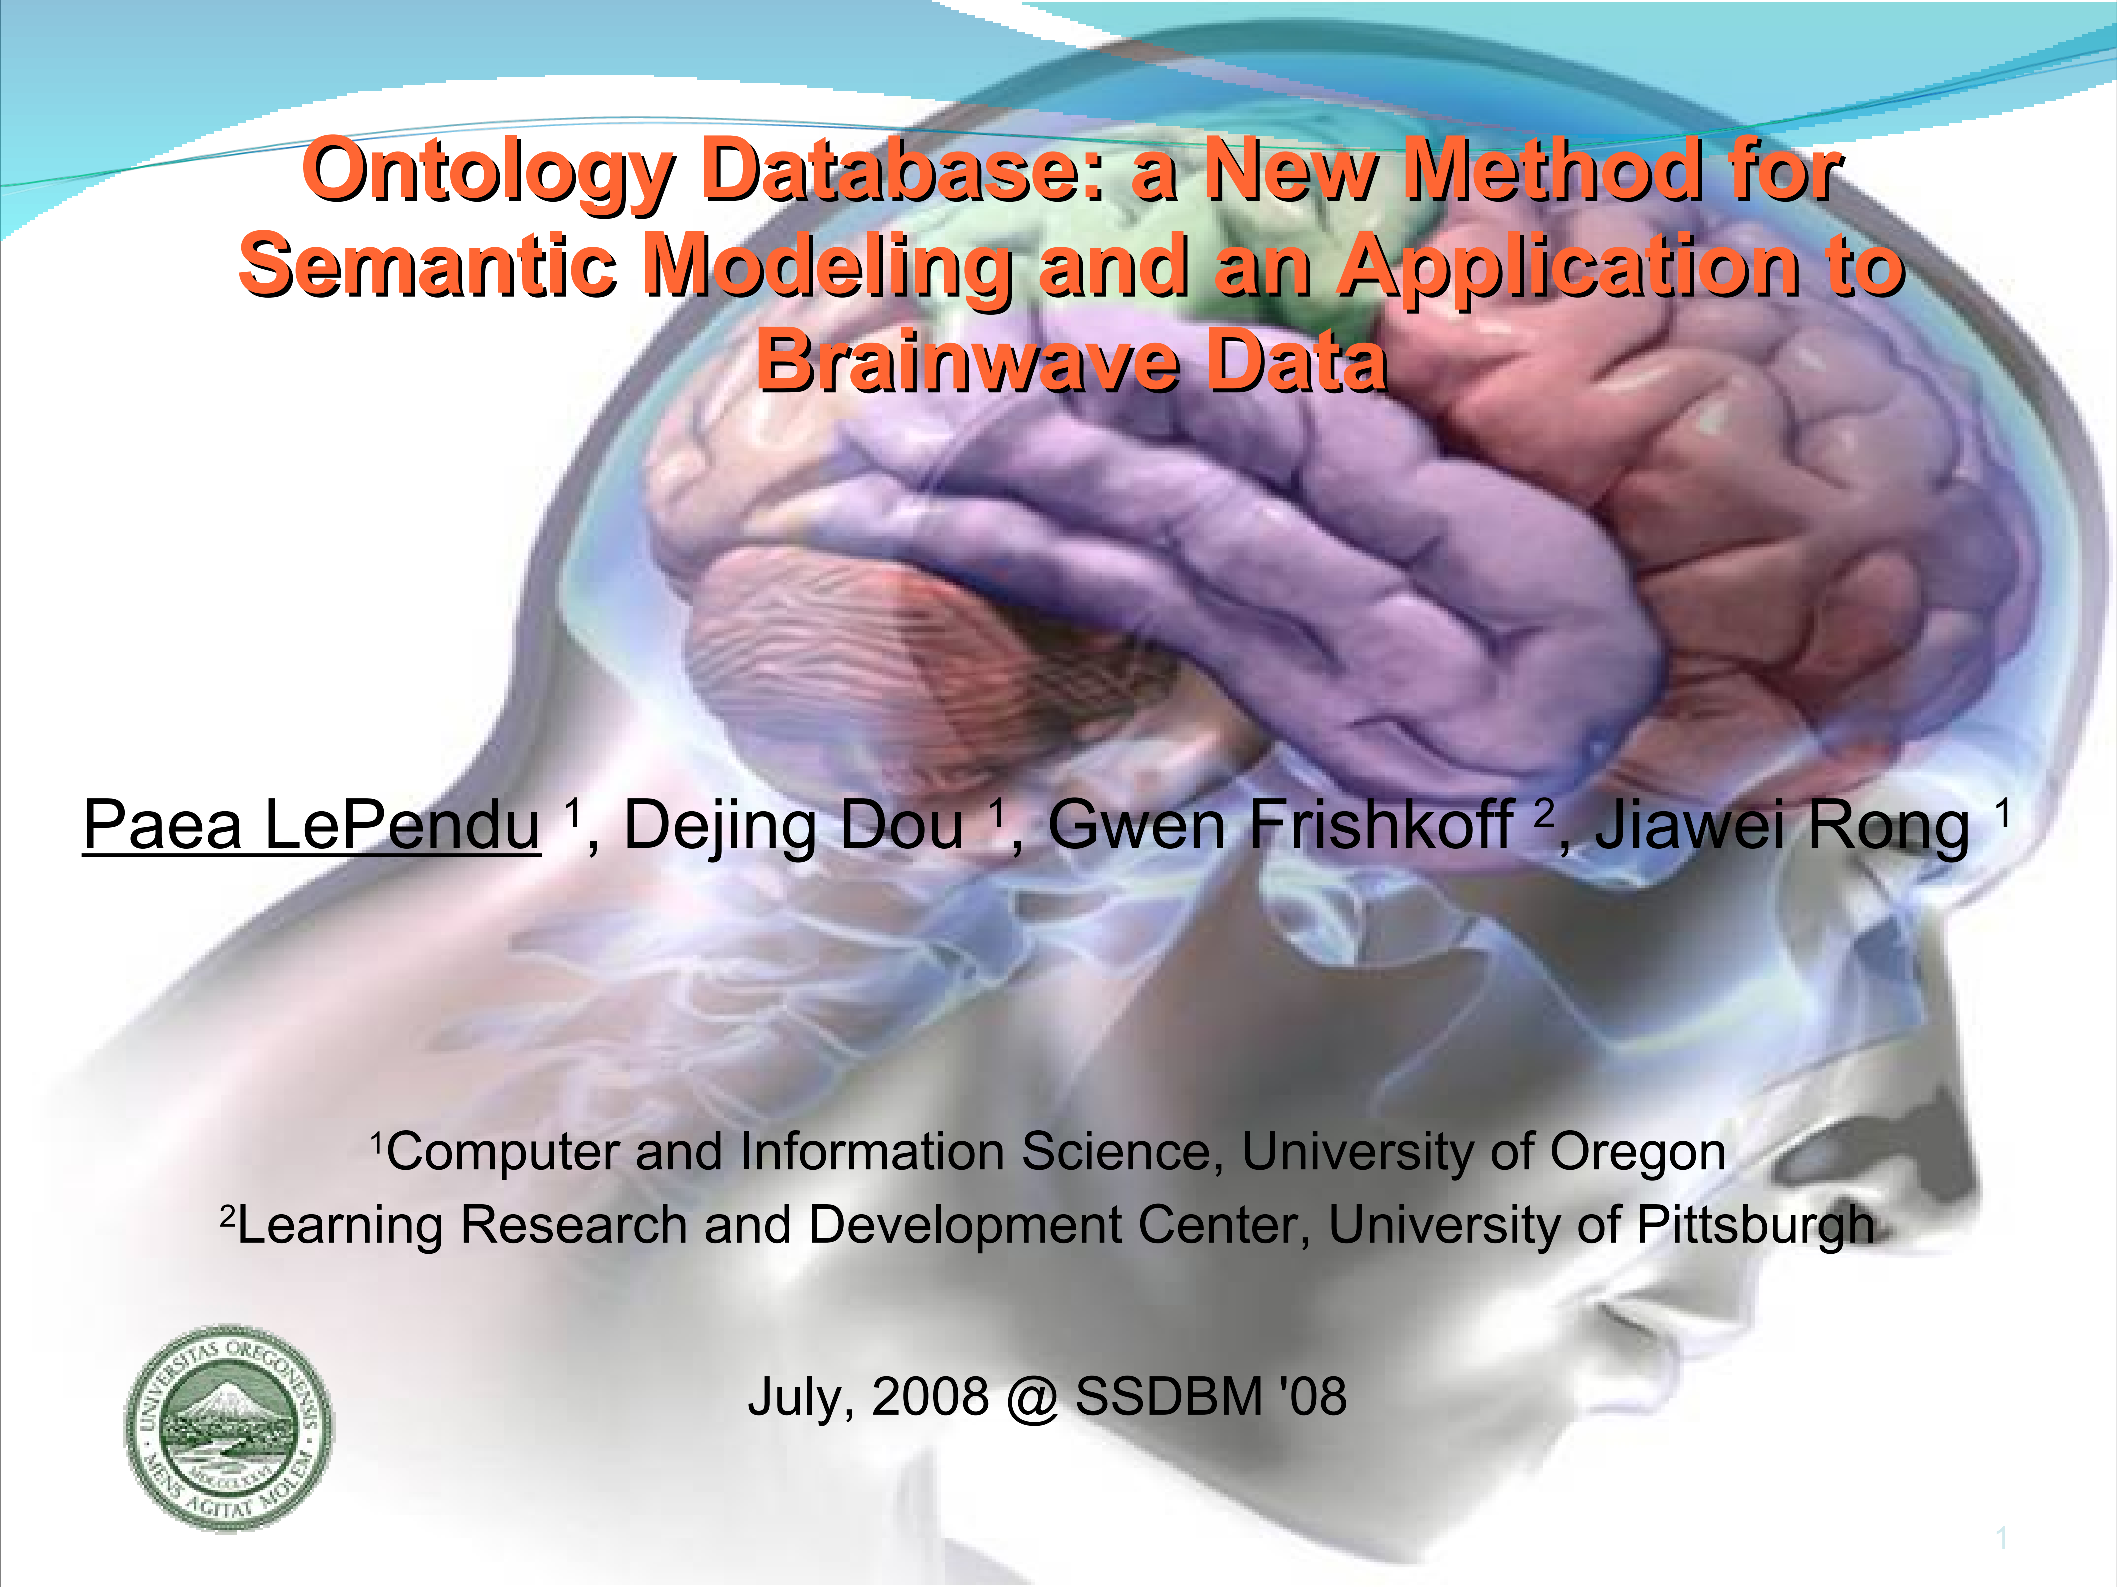

Ontology Database: a New Method for Semantic Modeling and an Application to Brainwave Data

Paea LePendu [1], Dejing Dou [1], Gwen Frishkoff [2], Jiawei Rong [1]

[1]Computer and Information Science, University of Oregon

[2]Learning Research and Development Center, University of Pittsburgh

July, 2008 @ SSDBM '08

Outline

- Background and Related Work
  - Brainwave data and pattern analysis
  - The NEMO project as motivation
  - Domain ontologies
- Ontology Database Methodology
  - Existing, view-based technique
  - New, trigger-based technique
- Benchmarking Analysis
- Discussion and Future Work

Brainwave Data

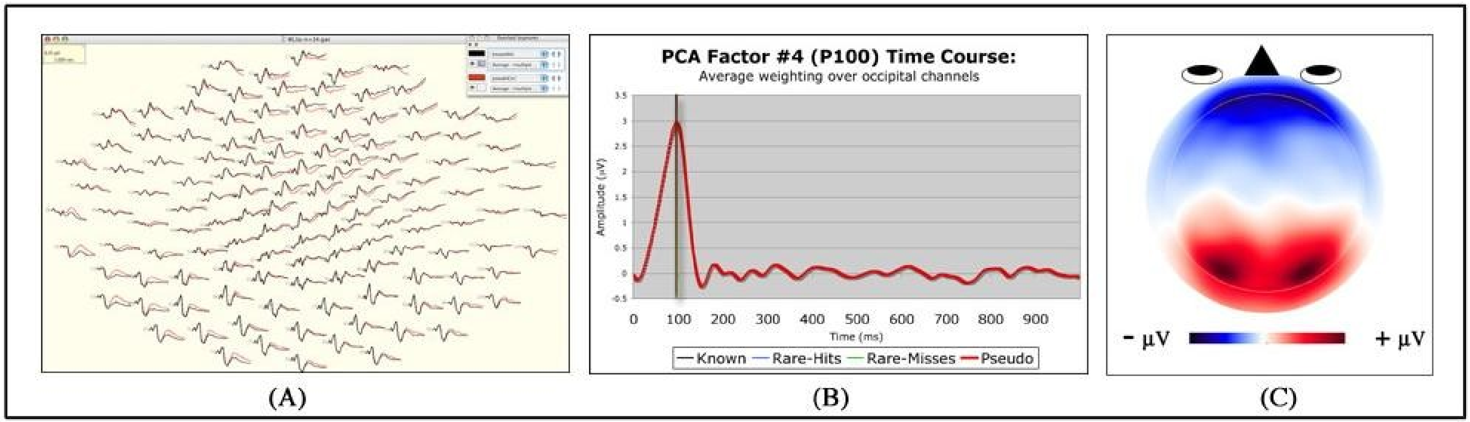

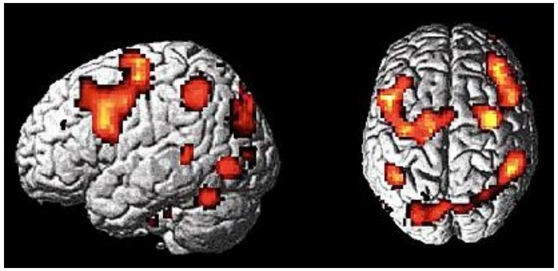

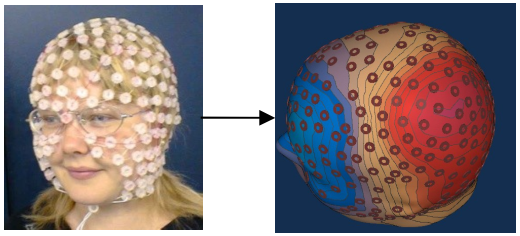

Brainwave Data

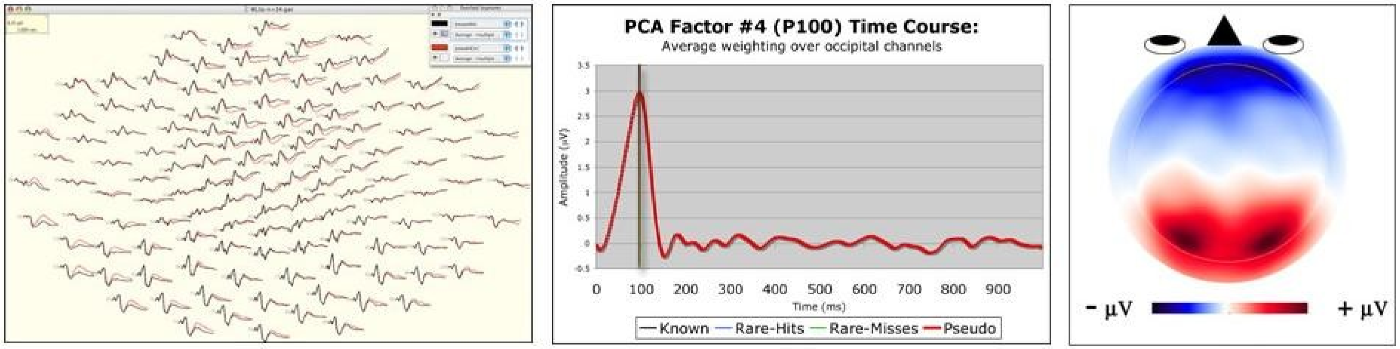

(A) (B) (C)

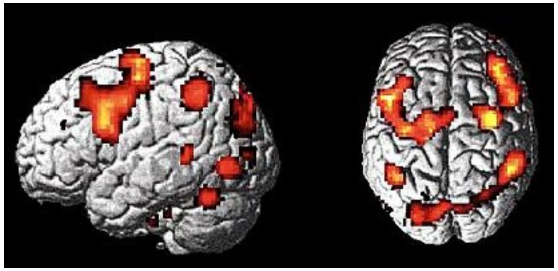

Talk about exponential growth!

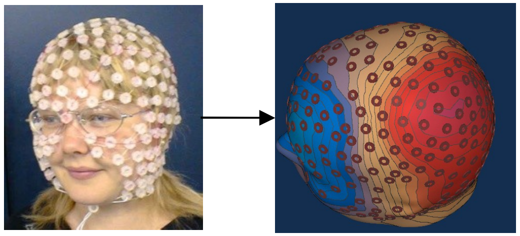

Brainwave Data

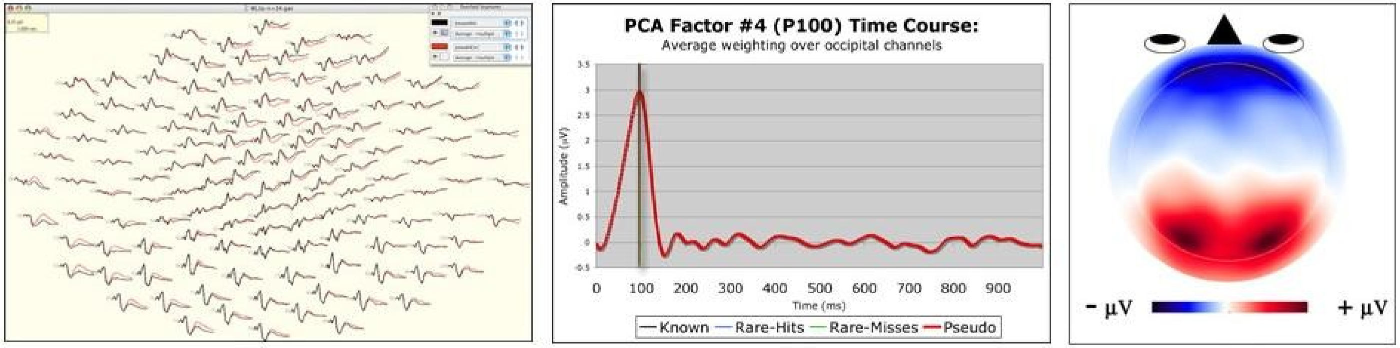

(A) (B) (C)

- Some problems with EEG/ERP data:
  - Complex dimensionality (spatial, temporal, functional)
  - Data sharing
  - Meta-analysis

Brainwave Ontologies

- To address these problems, ontologies are used:
  - Birnlex
  - NEMO (NeuroElectroMagnetic Ontologies)

- Distinct but inter-dependent models

NEMO (NeuroElectroMagnetic Ontology)

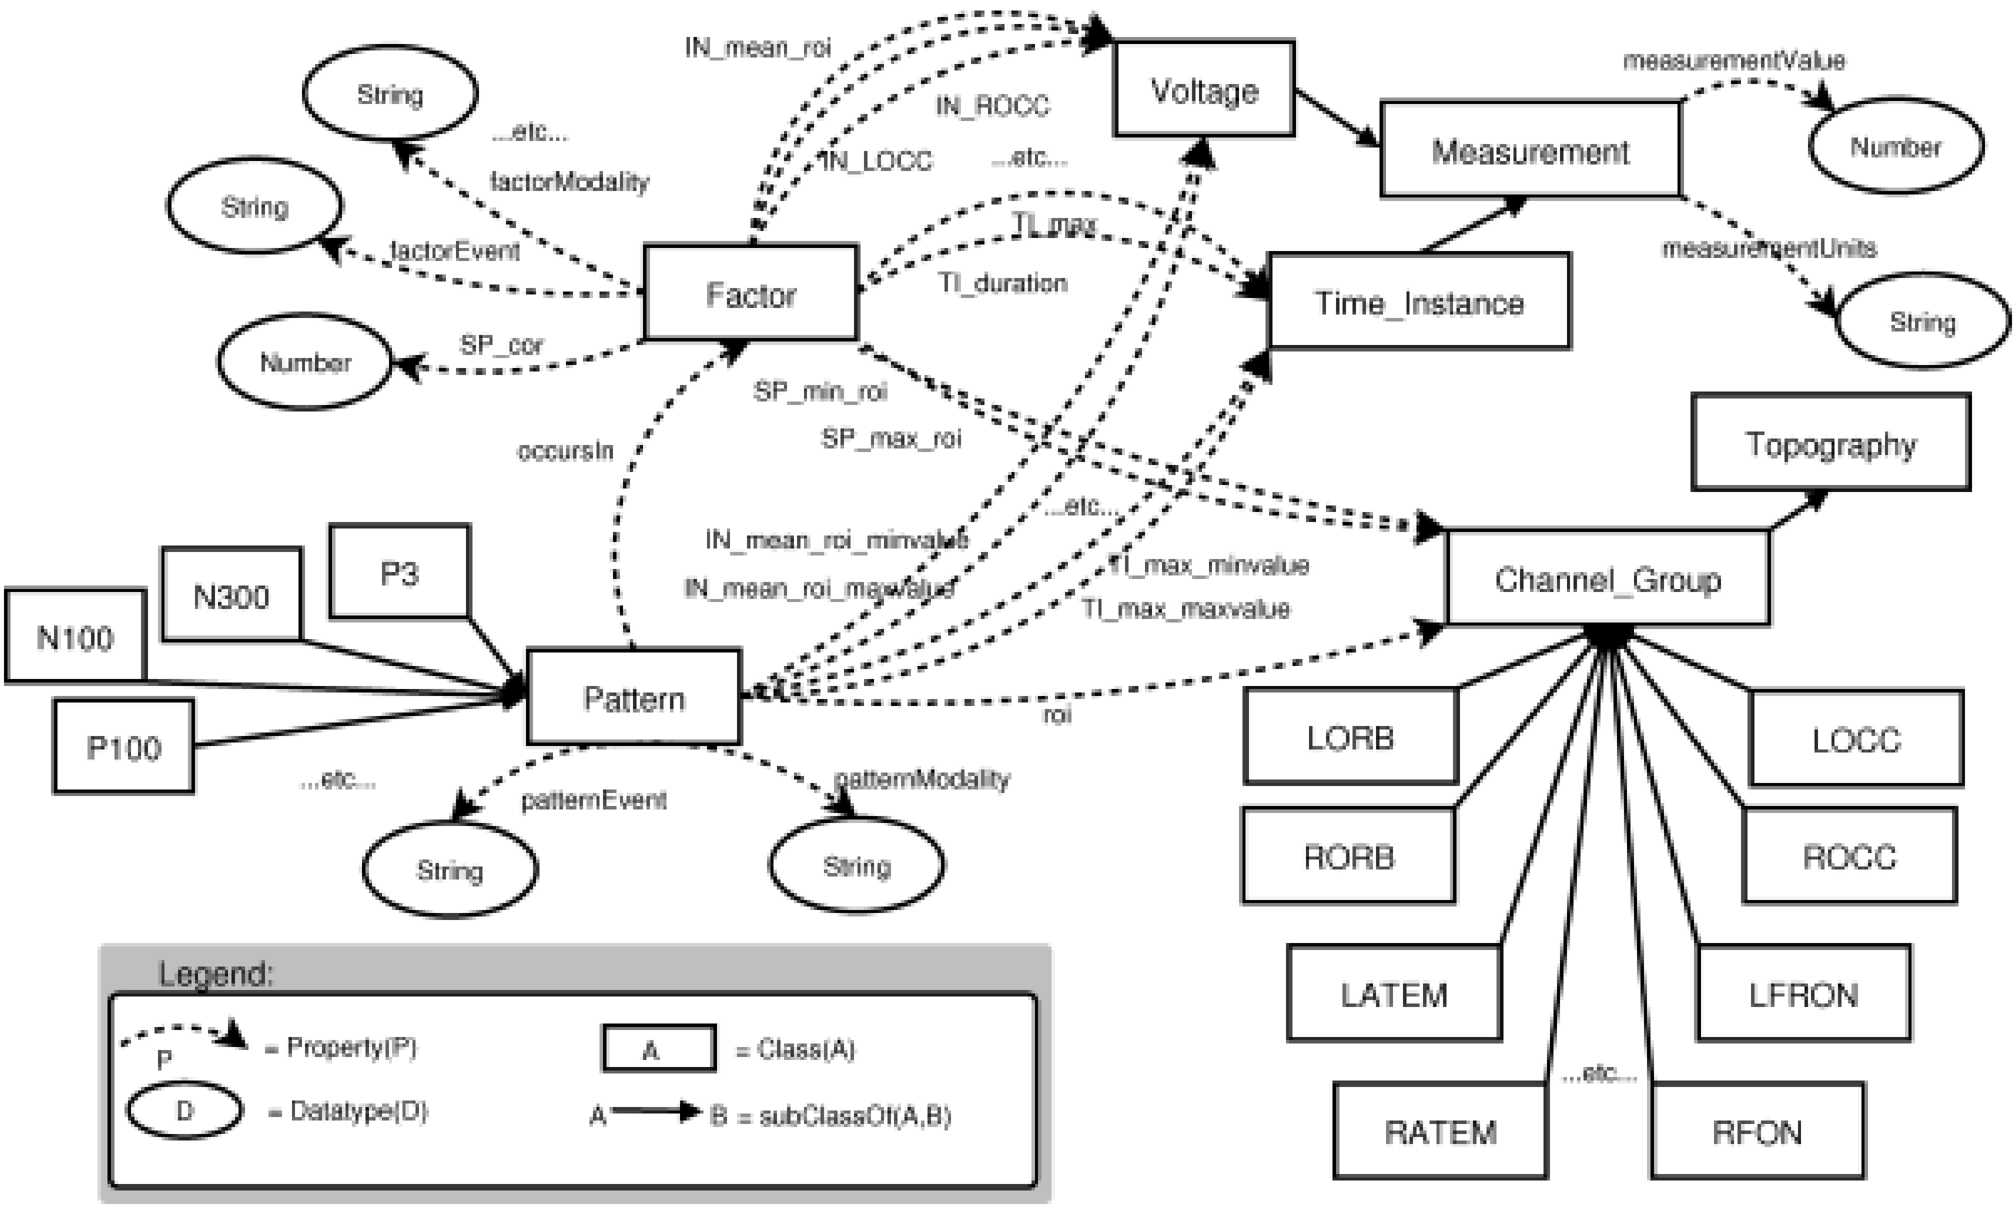

NEMO (NeuroElectroMagnetic Ontology)

Graphical View of the ER-Diagram

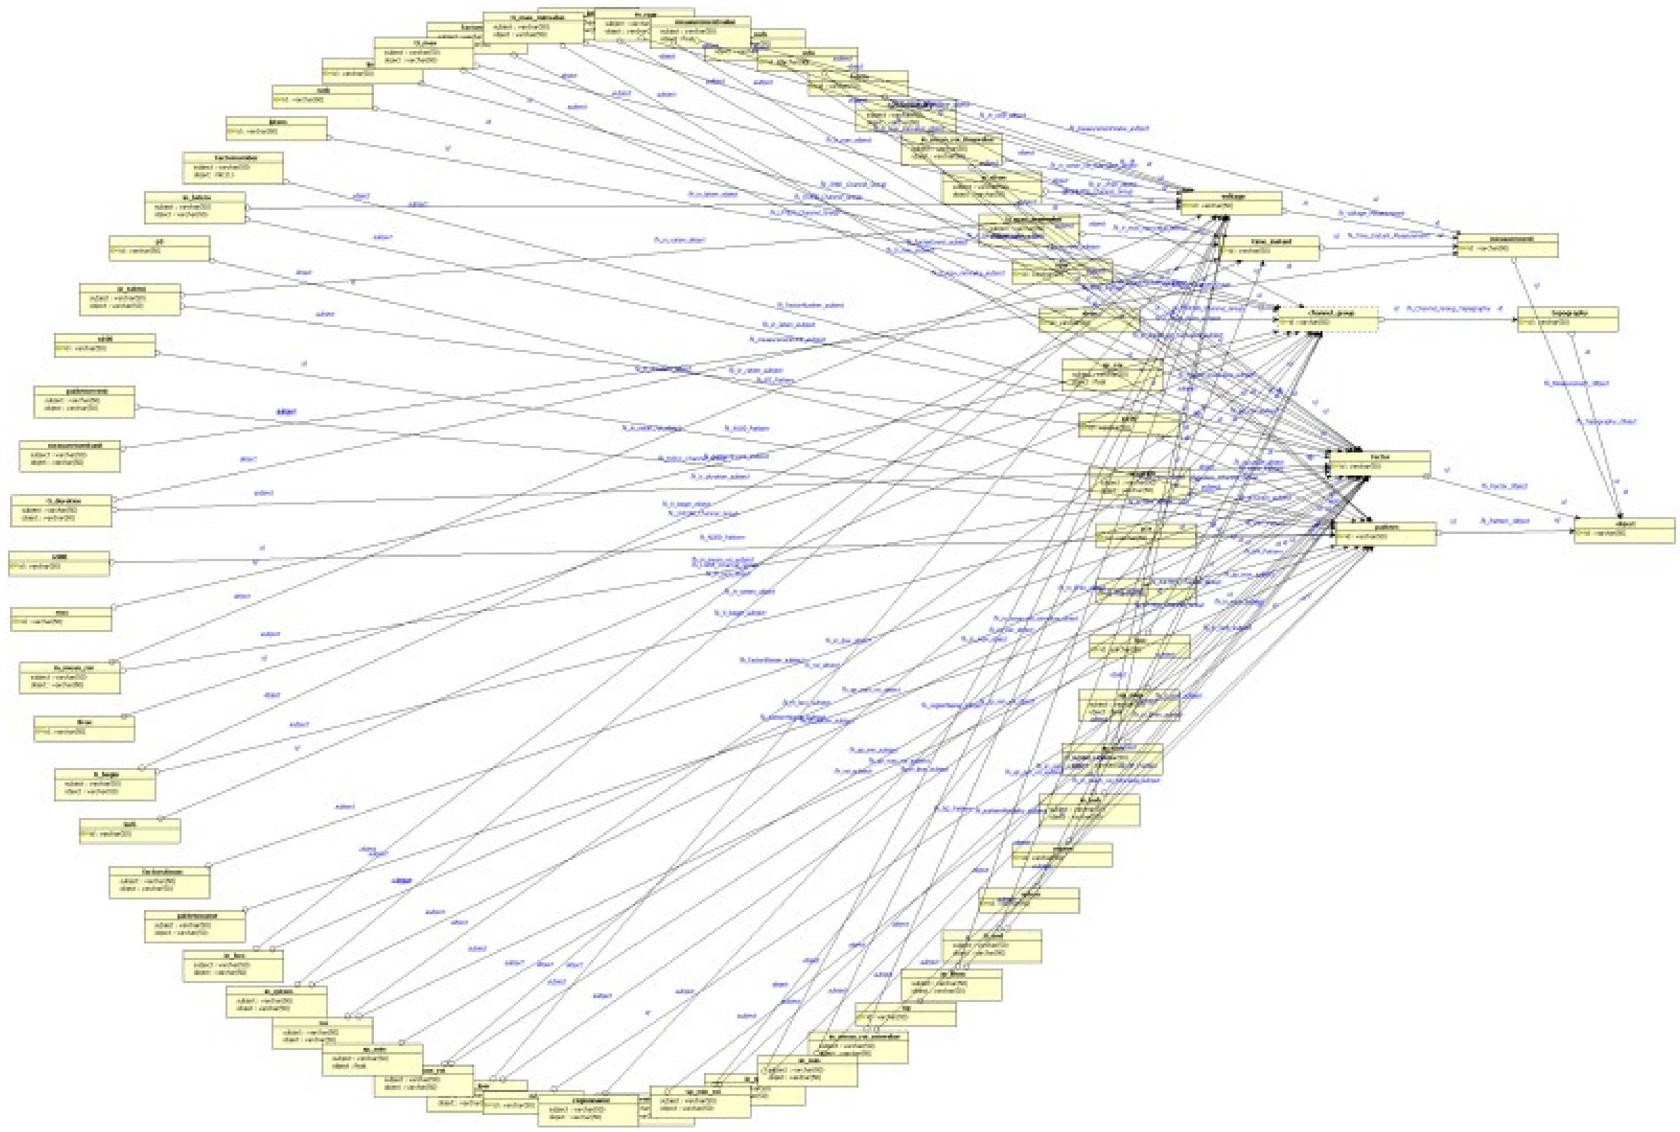

What are Ontologies?

- Machine processible models
- Logic-based formalisms
- Main communities:
  - Knowledge Representation and Reasoning (KRR)
  - Semantic Web

What does it have to do with databases?

- The problem of data scale (vs. model consistency)
  - Billion-triple challenge ISWC '08
- Views (Datalog) are coming back...
  - But databases have since evolved!
  - (e.g., Active Database technology)
- KRDB Group in Bozen-Bolzano, Italy
  - Reuniting Knowledge Representation and DataBases

Ontology Databases

- A Simple Problem Example
  - Some reasoning review
- Bridging the Gap
  - Ontologies and Databases
- Contrast Existing and Proposed Methodology

Example: a Simple Problem

<u>This is what we know :</u>

All sisters are siblings.

Hilary and Lynn are sisters.

<u>This is what we want to know :</u>

Who are siblings? { <x,y> | siblingOf(x,y) }

<u>Obviously, the answer should be :</u>

Hilary and Lynn are siblings. { <Hilary, Lynn> }

A Goal Directed Search

Automated reasoning can solve this easily.

A Goal Directed Search

{ <x,y> | siblingOf(x,y) }

A Goal Directed Search

{ <x,y> | siblingOf(x,y) }

siblingOf(x,y)

A Goal Directed Search

{ <x,y> | siblingOf(x,y) }

$$\frac{\cancel{\emptyset}}{\text{siblingOf}(x,y)}\ \textit{unify?}$$

A Goal Directed Search

$\{ \langle x,y \rangle \mid \text{siblingOf}(x,y) \}$

$$\frac{\emptyset}{\text{siblingOf}(x,y)}\ \emptyset$$

A Goal Directed Search

{ <x,y> | siblingOf(x,y) }

siblingOf(x,y)

A Goal Directed Search

$\{ <x,y> \mid \text{siblingOf}(x,y) \}$

$$\frac{\text{sisterOf}(x,y) \Rightarrow \text{siblingOf}(x,y) \qquad \text{sisterOf}(x, y)}{\text{siblingOf}(x,y)} \quad \textit{modus ponens}$$

A Goal Directed Search

$\{ \langle x,y \rangle \mid \text{siblingOf}(x,y) \}$

$$\dfrac{\dfrac{\dfrac{}{\forall x',y'.\ \text{sisterOf}(x',y') \Rightarrow \text{siblingOf}(x',y')}}{\text{sisterOf}(x,y) \Rightarrow \text{siblingOf}(x,y)}\ \forall_E\ \{x'/x,\ y'/y\} \qquad \text{sisterOf}(x,\ y)}{\text{siblingOf}(x,y)}\ \textit{modus ponens}$$

A Goal Directed Search

$$\{ \langle x,y \rangle \mid \text{siblingOf}(x,y) \}$$

$$\dfrac{\dfrac{\dfrac{}{\forall x',y'.\ \text{sisterOf}(x',y') \Rightarrow \text{siblingOf}(x',y')}}{\text{sisterOf}(x,y) \Rightarrow \text{siblingOf}(x,y)}\ \forall_E\ \{x'/x,\ y'/y\} \qquad \dfrac{}{\text{sisterOf}(x,\ y)}\ \textit{unify?}}{\text{siblingOf}(x,y)}\ \textit{modus ponens}$$

A Goal Directed Search

$$\{ \langle x,y \rangle \mid \text{siblingOf}(x,y) \}$$

$$\dfrac{\dfrac{\dfrac{}{\forall x',y'.\ \text{sisterOf}(x',y') \Rightarrow \text{siblingOf}(x',y')}}{\text{sisterOf}(x,y) \Rightarrow \text{siblingOf}(x,y)}\ \forall_E\ \{x'/x,\ y'/y\} \qquad \dfrac{\text{sisterOf}(\text{Hilary, Lynn})}{\text{sisterOf}(x,\ y)}\ \textit{unify?}}{\text{siblingOf}(x,y)}\ \textit{modus ponens}$$

A Goal Directed Search

$\{ \langle x,y \rangle \mid \text{siblingOf}(x,y) \}$

$$\dfrac{\dfrac{\overline{\forall x',y'.\ \text{sisterOf}(x',y') \Rightarrow \text{siblingOf}(x',y')}}{\text{sisterOf}(x,y) \Rightarrow \text{siblingOf}(x,y)}\ \forall_E\ \{x'/x,\ y'/y\} \qquad \dfrac{\overline{\text{sisterOf}(\text{Hilary, Lynn})}}{\text{sisterOf}(x,\ y)}\ \textit{unify?}}{\text{siblingOf}(x,y)}\ \textit{modus ponens}$$

A Goal Directed Search

$\{ \langle x,y \rangle \mid \text{siblingOf}(x,y) \}$

$$
\dfrac{\dfrac{\overline{\forall x',y'.\ \text{sisterOf}(x',y') \Rightarrow \text{siblingOf}(x',y')}}{\text{sisterOf}(x,y) \Rightarrow \text{siblingOf}(x,y)}\ \forall_E\ \{x'/x,\ y'/y\} \qquad \dfrac{\overline{\text{sisterOf}(\text{Hilary, Lynn})}}{\text{sisterOf}(x,\ y)}\ \textit{unify!}}{\text{siblingOf}(x,y)}\ \textit{modus ponens}
$$

A Goal Directed Search

{ <x,y> | siblingOf(x,y) }

$\forall$x',y'. sisterOf(x',y') $\Rightarrow$ siblingOf(x',y')

sisterOf(Hilary, Lynn)

$\forall_E$ {x'/x, y'/y}

{x/Hilary}

sisterOf(x,y) $\Rightarrow$ siblingOf(x,y)

sisterOf(x, y)

modus ponens

siblingOf(x,y)

A Goal Directed Search

$\{ \langle x,y \rangle \mid \text{siblingOf}(x,y) \}$

$$\dfrac{\dfrac{\dfrac{}{\forall x',y'.\ \text{sisterOf}(x',y') \Rightarrow \text{siblingOf}(x',y')}}{\text{sisterOf}(x,y) \Rightarrow \text{siblingOf}(x,y)}\,\forall_E\ \{x'/x,\ y'/y\} \qquad \dfrac{\dfrac{}{\text{sisterOf}(\text{Hilary, Lynn})}}{\text{sisterOf}(x,\ y)}\,\{x/\text{Hilary},\ y/\text{Lynn}\}}{\text{siblingOf}(x,y)}\ \textit{modus ponens}$$

A Goal Directed Search

$$\{ \langle x,y \rangle \mid \text{siblingOf}(x,y) \}$$

$$\cfrac{\cfrac{\cfrac{}{\forall x',y'.\ \text{sisterOf}(x',y') \Rightarrow \text{siblingOf}(x',y')}}{\text{sisterOf}(\text{Hilary}, \text{Lynn}) \Rightarrow \text{siblingOf}(\text{Hilary}, \text{Lynn})}\,\forall_E \{x'/\text{Hilary}, y'/\text{Lynn}\} \qquad \cfrac{}{\text{sisterOf}(\text{Hilary}, \text{Lynn})}}{\text{siblingOf}(\text{Hilary}, \text{Lynn})}\ \textit{modus ponens}$$

A Goal Directed Search

$$\langle Hilary, Lynn\rangle \in \quad \{ \langle x,y\rangle \mid siblingOf(x,y) \}$$

$$\dfrac{\dfrac{\dfrac{}{\forall x',y'.\ sisterOf(x',y') \Rightarrow siblingOf(x',y')}}{sisterOf(Hilary, Lynn) \Rightarrow siblingOf(Hilary, Lynn)}\ \forall_{E}\ \{x'/Hilary,\ y'/Lynn\} \qquad \dfrac{}{sisterOf(Hilary, Lynn)}}{siblingOf(Hilary, Lynn)}\ \textit{modus ponens}$$

Key Question #1

If data storage and querying is our main goal...

Key Question #1

...do we really need all this reasoning?

Ontology Databases

Bringing ontologies and databases together.

Ontology Databases

Class

Property

Datatype

Axioms

Objects

Facts

How do we bridge these?

Relation

Attribute

Datatype

keys

constraints

views

triggers

tuples

Ontology Databases

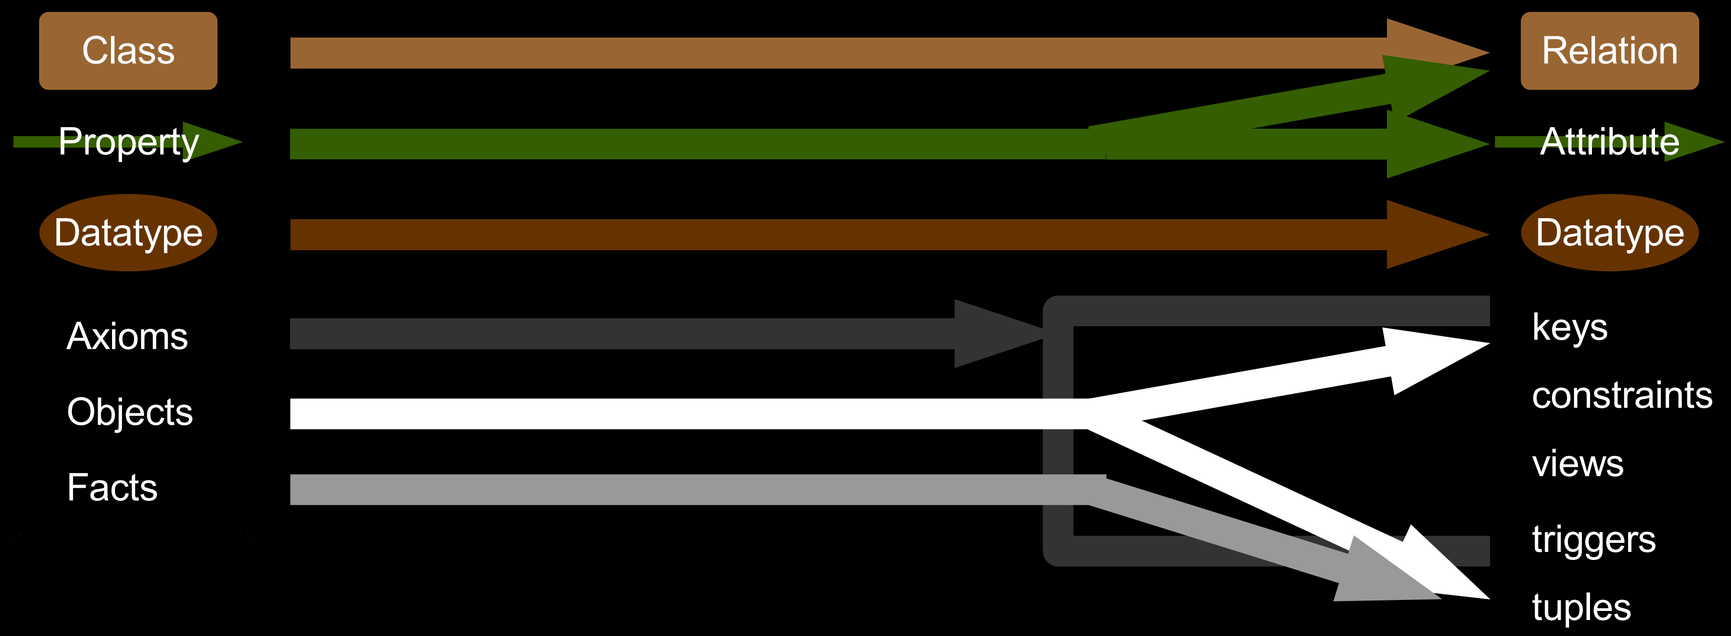

Ontology Databases

Class

Property

Datatype

Axioms

Objects

Facts

Here's an example.

Relation

Attribute

Datatype

keys

constraints

views

triggers

tuples

Ontology Databases

datatype-properties

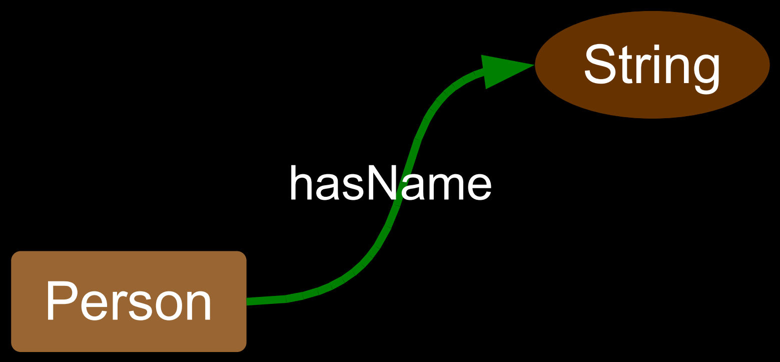

Ontology Databases

HasName

Subject	Object (varchar)
AishaSun	Aisha Sun
HilaryMeade	Hilary Meade
LynnMeade	Lynn Meade

foreign-key constraint
referencing

Person

Id
AishaSun
HilaryMeade
LynnMeade

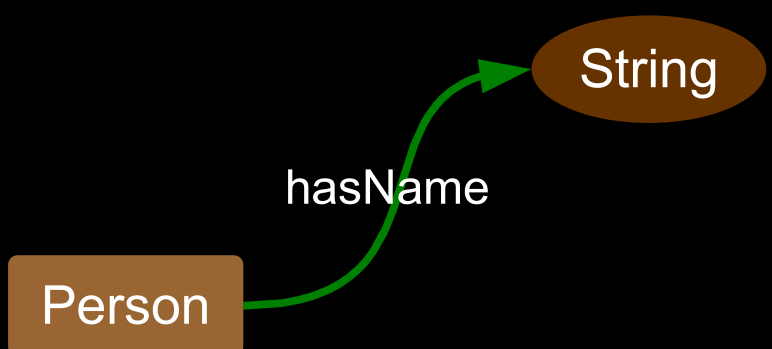

Ontology Databases

object-properties

Ontology Databases

Female

Id
AishaSun
HilaryMeade
LynnMeade

HusbandOf

Subject	Object
MahmudReece	LynnMeade

Male

Id
MahmudReece

Ontology Databases

Female

Id
AishaSun
HilaryMeade
LynnMeade

HusbandOf

Subject	Object
MahmudReece	LynnMeade

foreign-key constraint referencing

Male

Id
MahmudReece

::domain

Ontology Databases

Female

Id
AishaSun
HilaryMeade
LynnMeade

foreign-key constraint referencing

HusbandOf

Subject	Object
MahmudReece	LynnMeade

foreign-key constraint referencing

::range

Male

Id
MahmudReece

Ontology Databases

subClass axioms

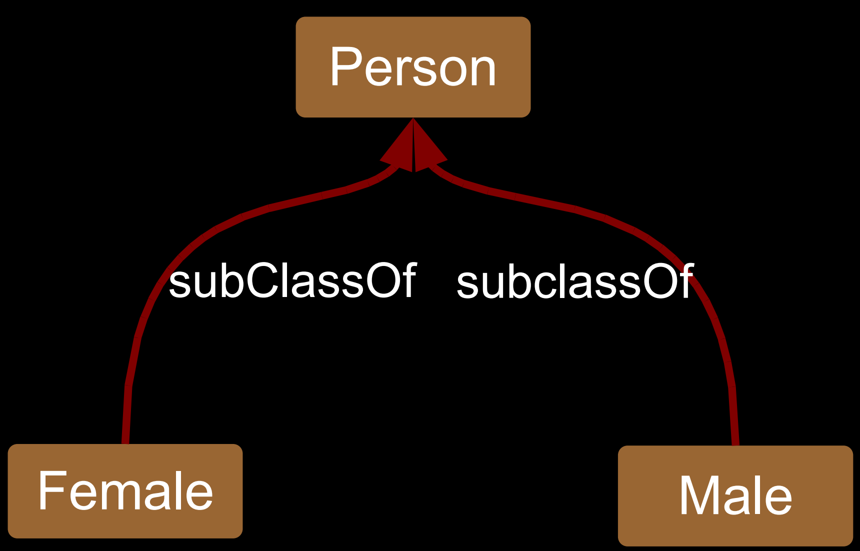

Ontology Databases

subClass axioms

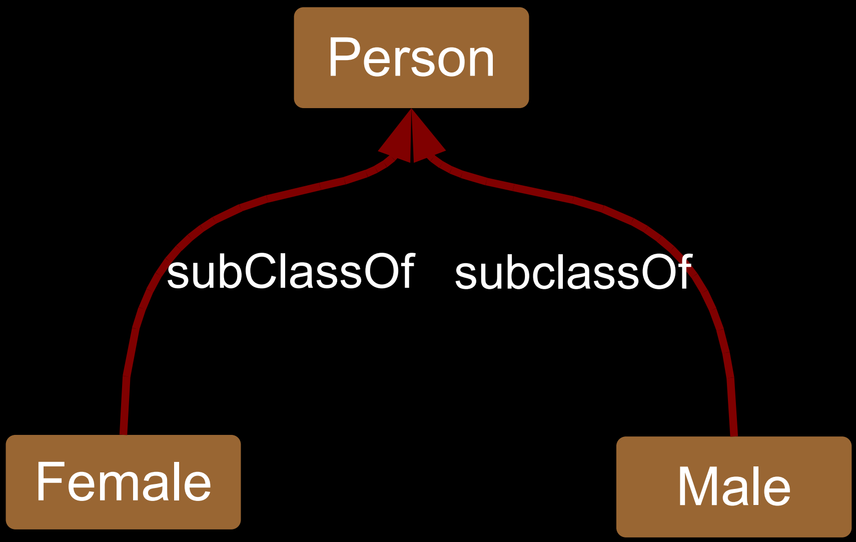

Two approaches.

Ontology Databases

subClass axioms

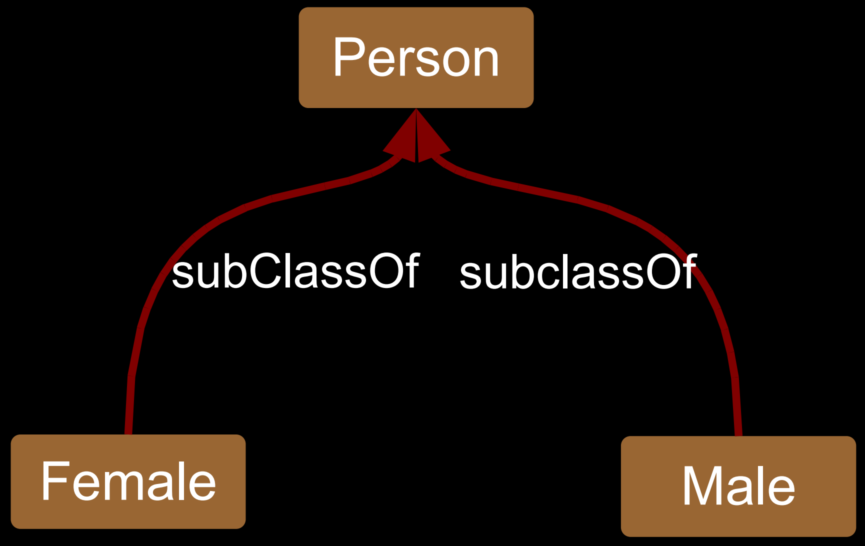

1. View-based approach.

Ontology Databases

```
CREATE VIEW v_Person(id) AS
    SELECT id FROM Person
    UNION
    SELECT id FROM Male
```

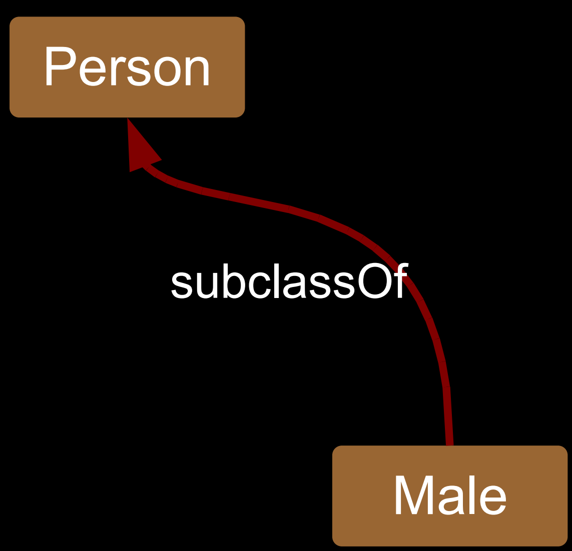

Person

Id

Male

Id
MahmudReece

1. View-based approach.

Ontology Databases

```
CREATE VIEW v_Person(id) AS
    SELECT id FROM Person
    UNION
    SELECT id FROM Male
```

v_Person

Id
MahmudReece

Person

Id

Male

Id
MahmudReece

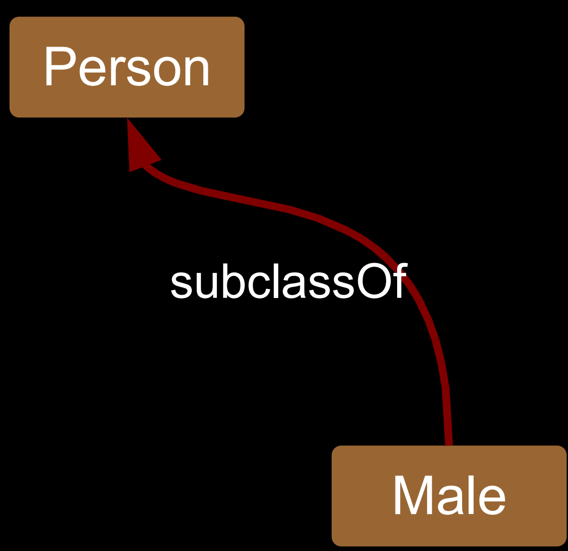

1. View-based approach.

Ontology Databases

```
CREATE VIEW v_Person(id) AS
    SELECT id FROM Person
    UNION
    SELECT id FROM Male
```

v_Person

Id
MahmudReece

Person

Id

Female

Id
AishaSun
HilaryMeade
LynnMeade

Male

Id
MahmudReece

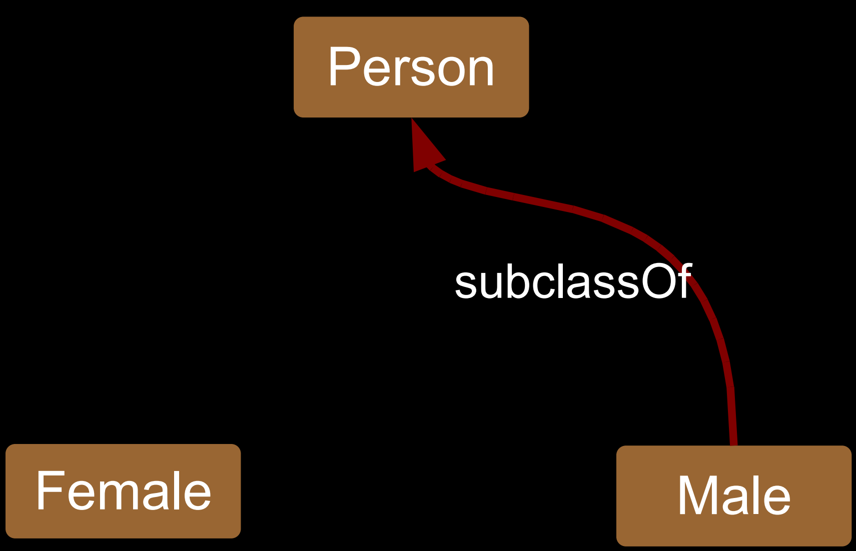

1. View-based approach.

Ontology Databases

```
CREATE VIEW v_Person(id) AS
    SELECT id FROM Person
    UNION
    SELECT id FROM Male
    UNION
    SELECT id FROM Female
```

v_Person

Id
MahmudReece
AishaSun
HilaryMeade
LynnMeade

Person

Id

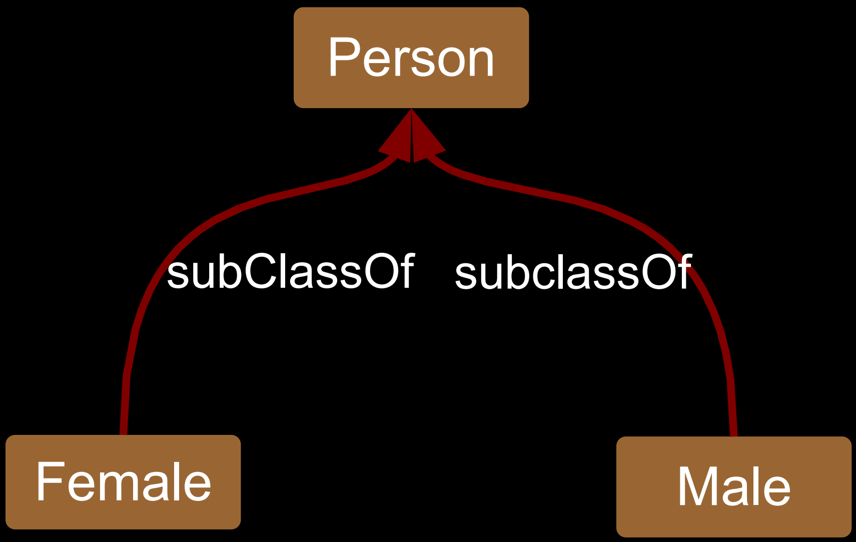

Female

Id
AishaSun
HilaryMeade
LynnMeade

Male

Id
MahmudReece

1. View-based approach.

Ontology Databases

DLDB [Pan & Heflin, 2003] implements the view-based approach to store and retrieve voluminous Semantic Web data.

1. View-based approach.

Ontology Databases

subClass axioms

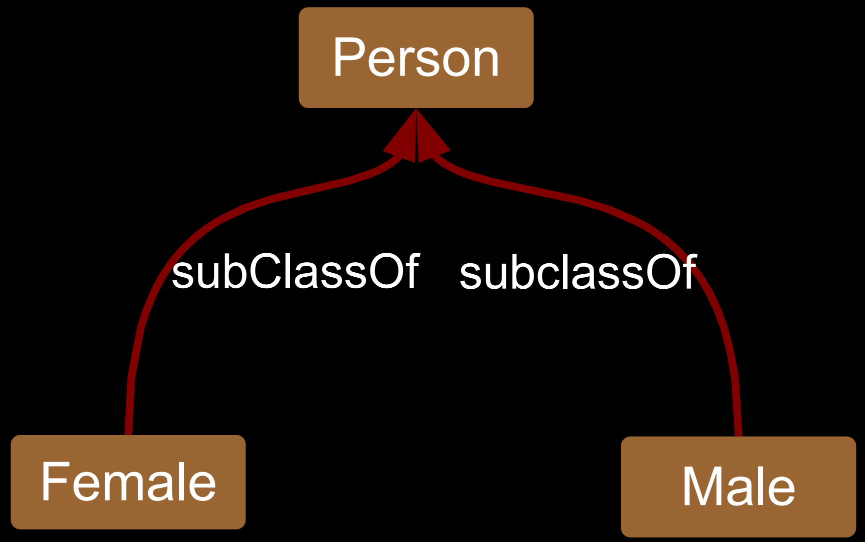

Ontology Databases

subClass axioms

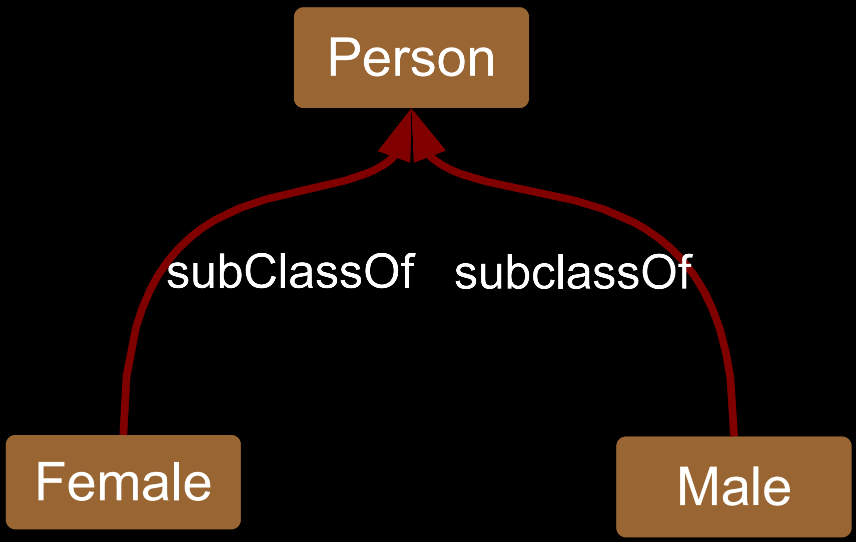

2. Trigger-based approach.

Ontology Databases

Person

Id

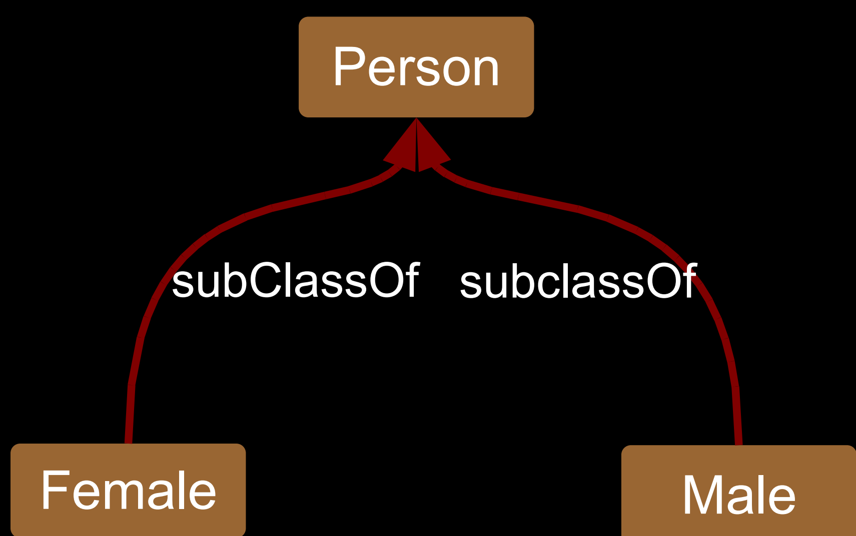

Female

Id

Male

Id

2. Trigger-based approach.

Ontology Databases

Person

Id

Female

Id

Male

Id
MahmudReece

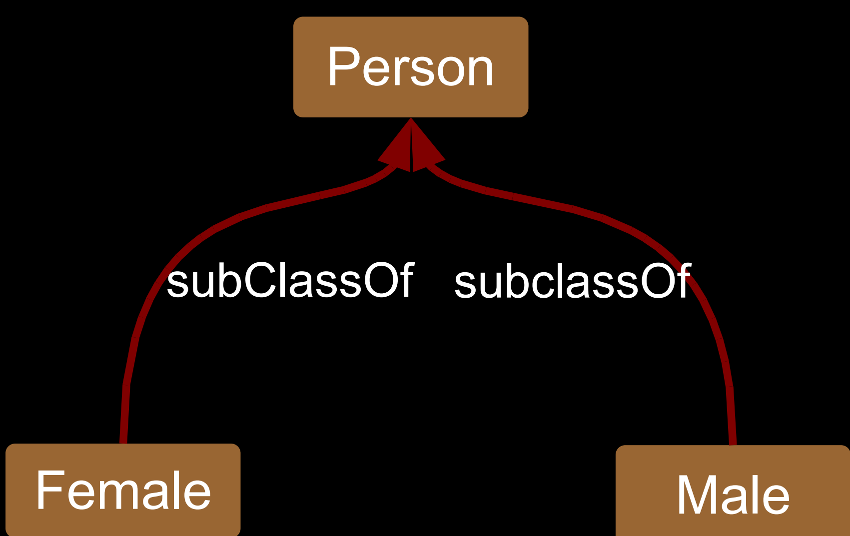

2. Trigger-based approach.

Ontology Databases

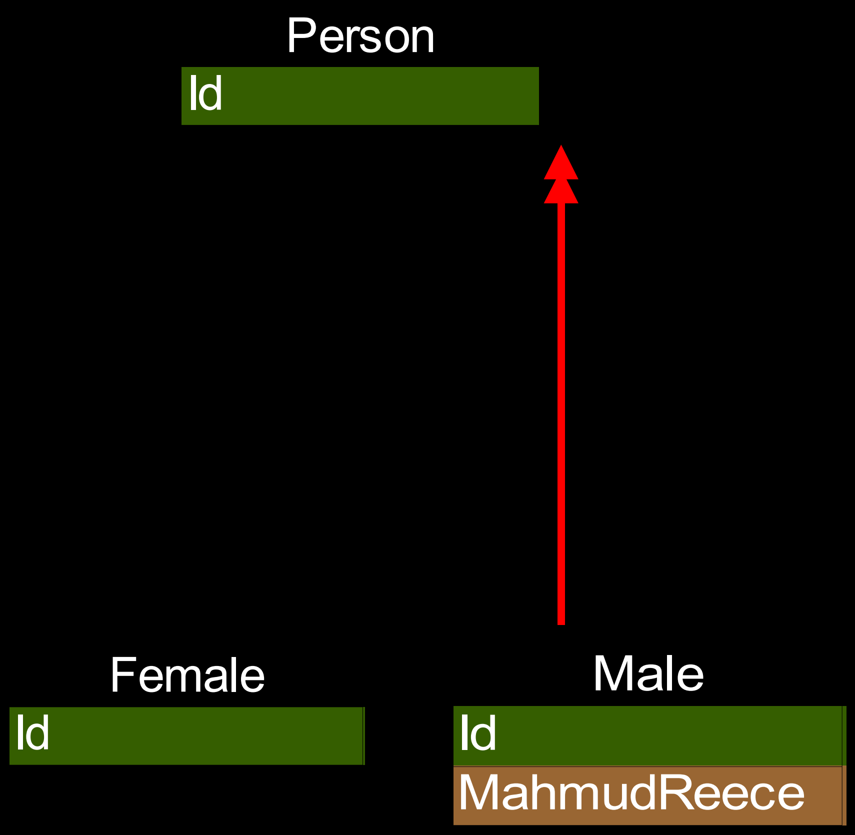

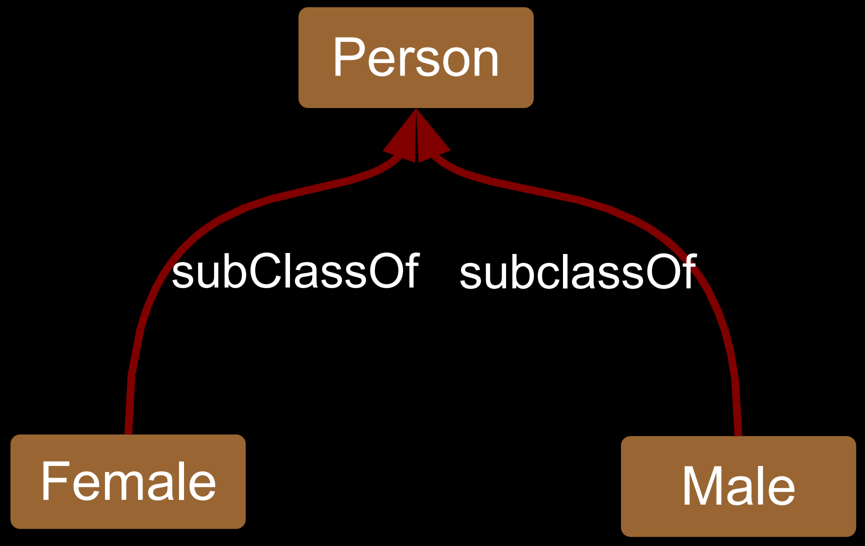

2. Trigger-based approach.

Ontology Databases

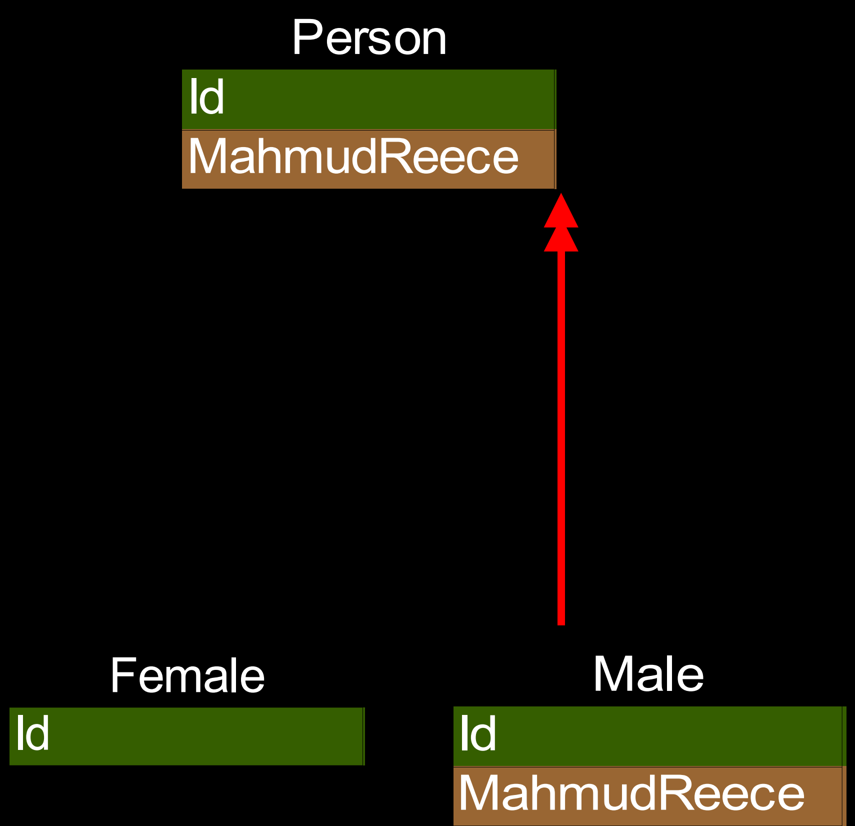

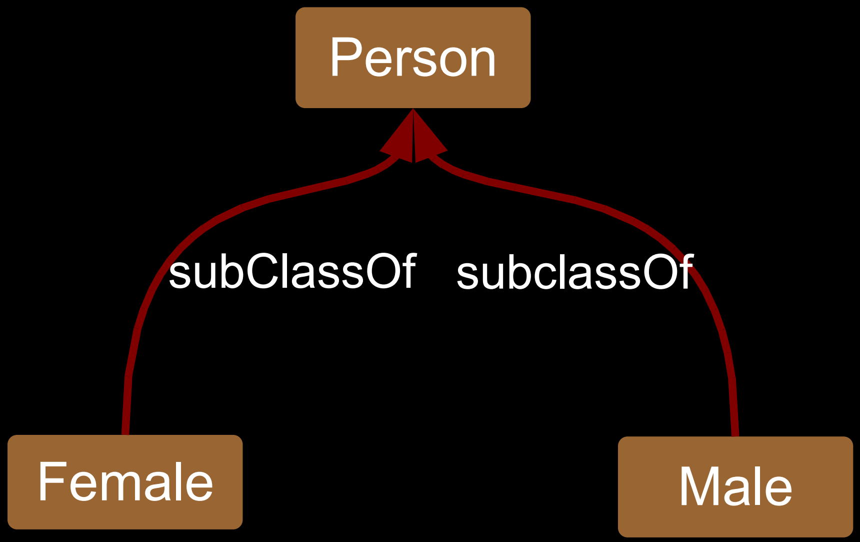

2. Trigger-based approach.

Ontology Databases

Person

Id
MahmudReece

Female

Id

Male

Id
MahmudReece

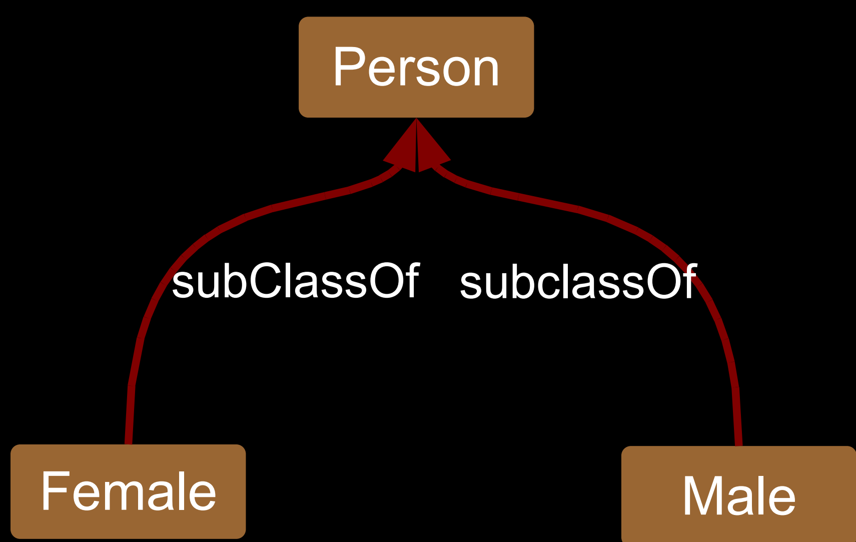

2. Trigger-based approach.

Ontology Databases

Person

Id
MahmudReece

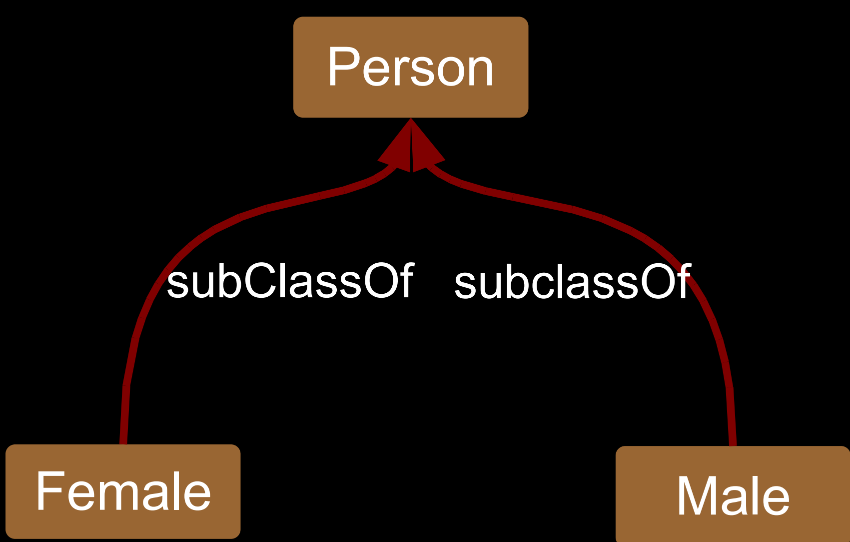

Female

Id
AishaSun

Male

Id
MahmudReece

2. Trigger-based approach.

Ontology Databases

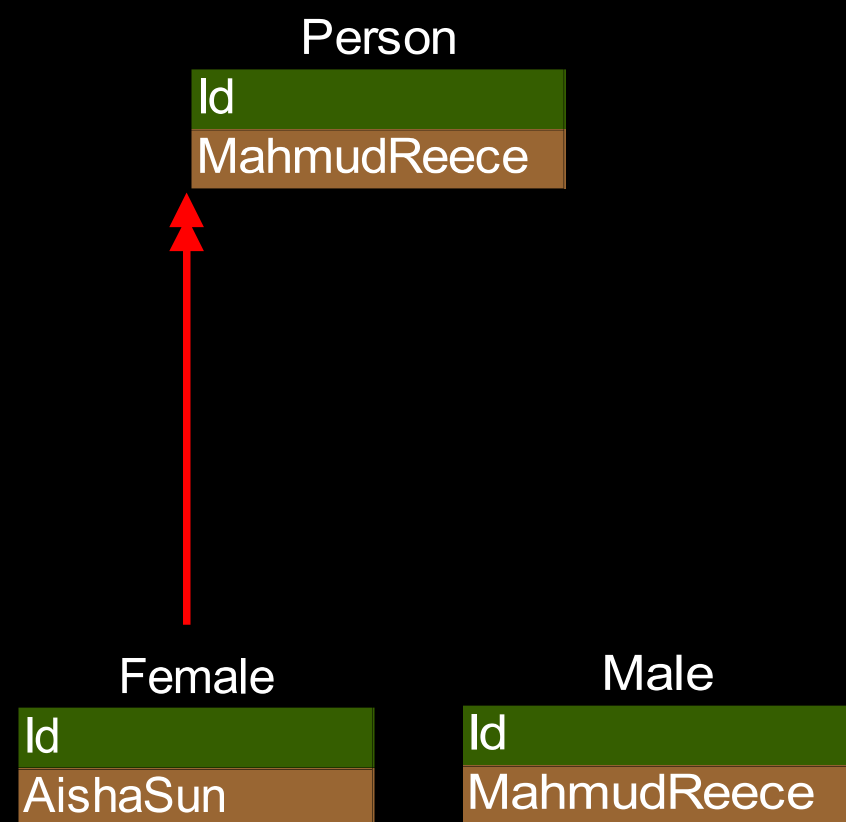

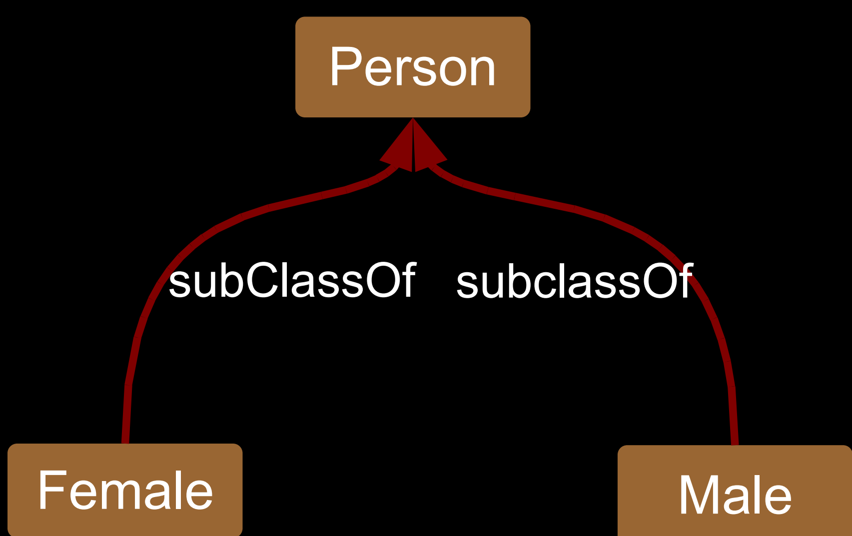

2. Trigger-based approach.

Ontology Databases

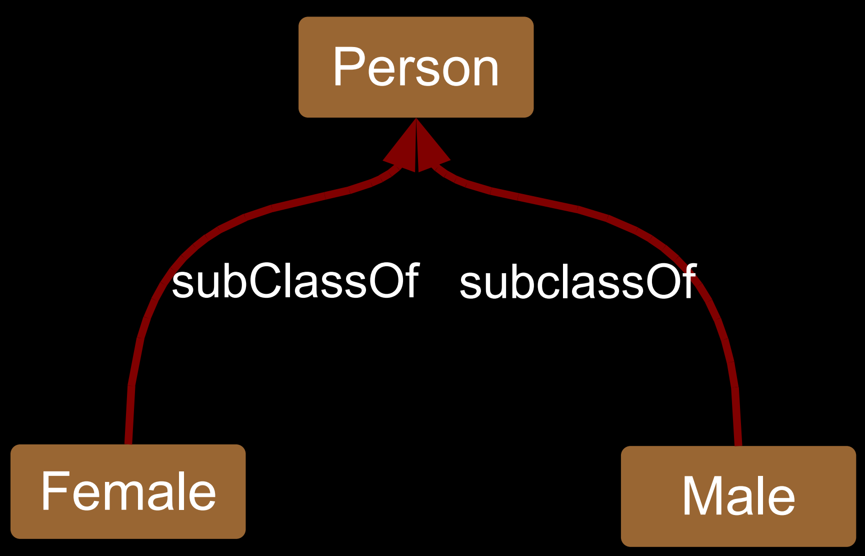

Person

Id
MahmudReece
AishaSun

Female

Id
AishaSun

Male

Id
MahmudReece

2. Trigger-based approach.

Ontology Databases

Person

Id
MahmudReece
AishaSun

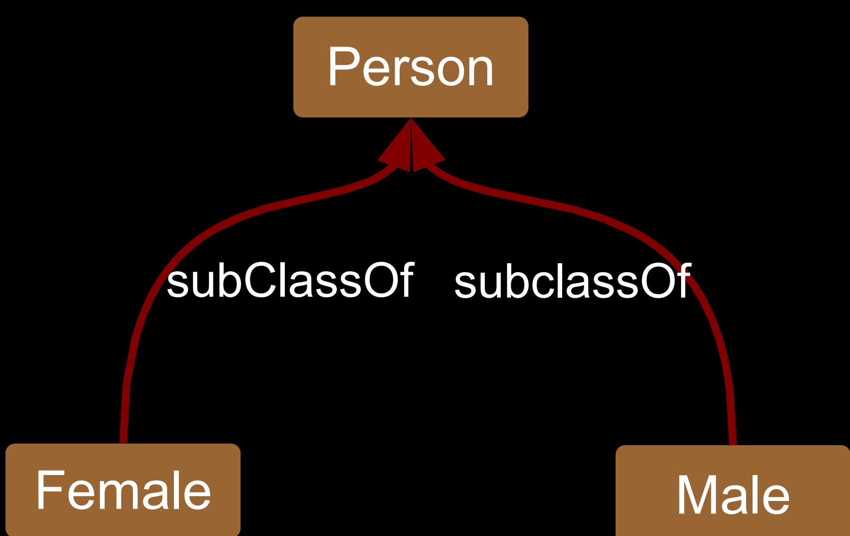

Female

Id
AishaSun

Male

Id
MahmudReece

2. Trigger-based approach.

Ontology Databases

Person

Id
MahmudReece
AishaSun

Female

Id
AishaSun
HilaryMeade

Male

Id
MahmudReece

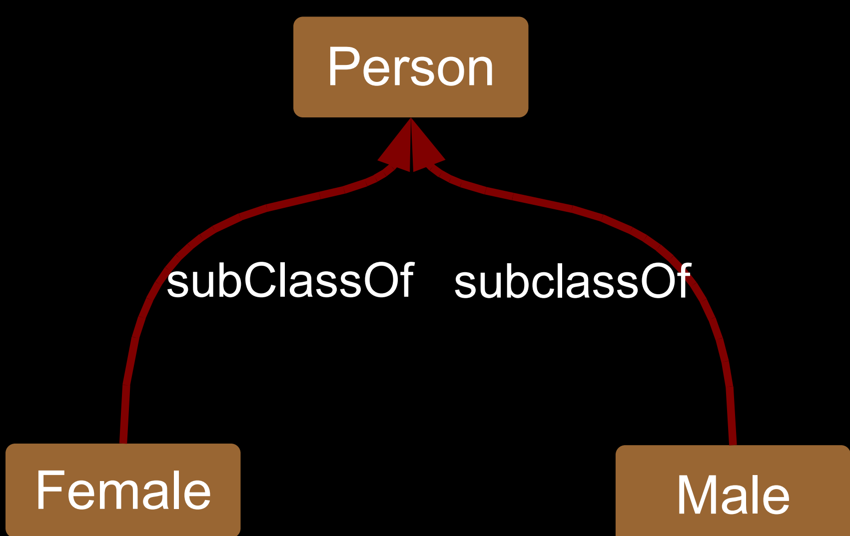

2. Trigger-based approach.

Ontology Databases

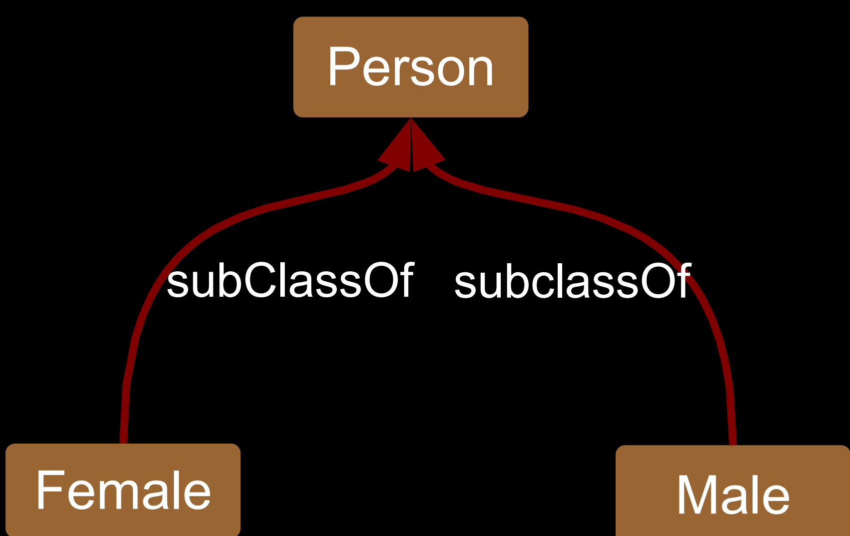

Person

Id
MahmudReece
AishaSun
HilaryMeade

Female

Id
AishaSun
HilaryMeade

Male

Id
MahmudReece

2. Trigger-based approach.

Ontology Databases

Person

Id
MahmudReece
AishaSun
HilaryMeade

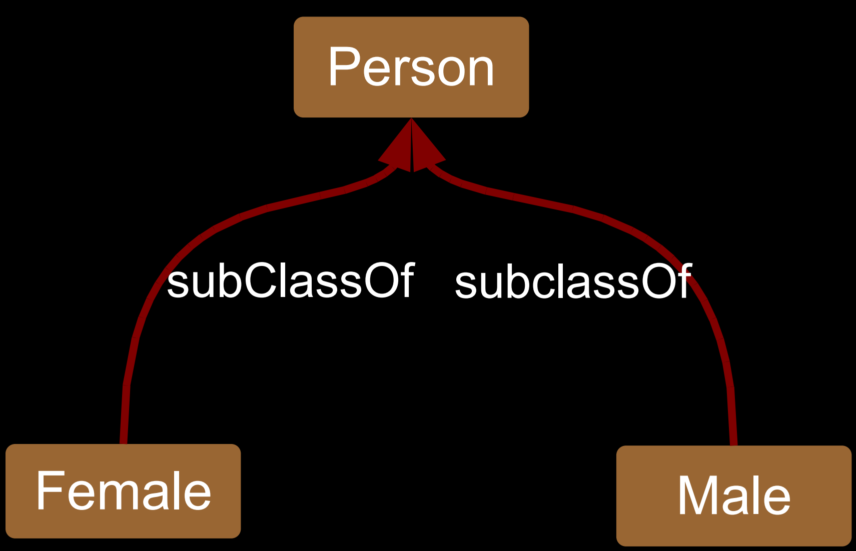

Female

Id
AishaSun
HilaryMeade

Male

Id
MahmudReece

2. Trigger-based approach.

Ontology Databases

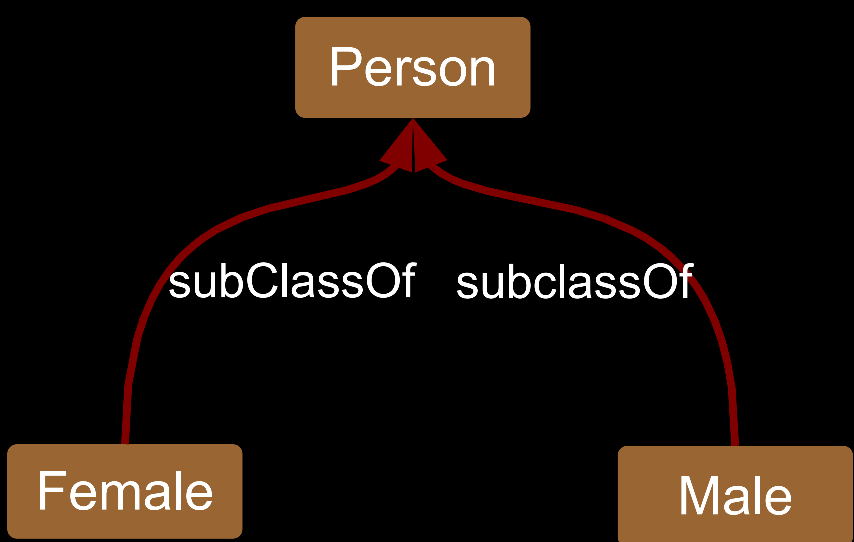

Person

Id
MahmudReece
AishaSun
HilaryMeade

Female

Id
AishaSun
HilaryMeade
LynnMeade

Male

Id
MahmudReece

2. Trigger-based approach.

Ontology Databases

Person

Id
MahmudReece
AishaSun
HilaryMeade
LynnMeade

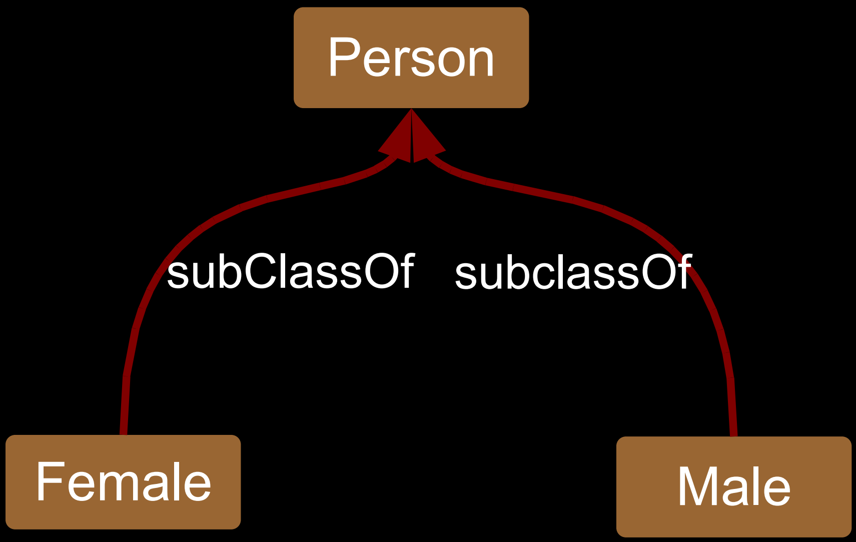

Female

Id
AishaSun
HilaryMeade
LynnMeade

Male

Id
MahmudReece

2. Trigger-based approach.

Ontology Databases

Person

Id
MahmudReece
AishaSun
HilaryMeade
LynnMeade

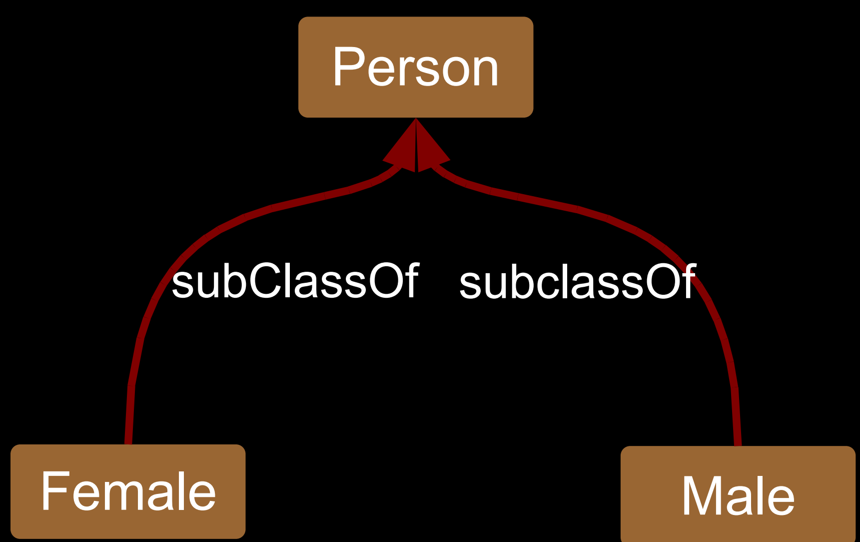

Female

Id
AishaSun
HilaryMeade
LynnMeade

Male

Id
MahmudReece

2. Trigger-based approach.

Ontology Databases

subProperty axioms

2. Trigger-based approach.

Ontology Databases

subProperty axioms

(basically the same idea)

2. Trigger-based approach.

Ontology Databases

OntoDB [SSDBM '08] implements the trigger-based approach.

2. Trigger-based approach.

Ontology Databases

Class

Property

Datatype

Axioms

Objects

Facts

Relation

Attribute

Datatype

keys

constraints

triggers

tuples

Ontology Databases

Class

Property

Datatype

Axioms

Objects

Facts

Now we have bridged these.

Relation

Attribute

Datatype

keys

constraints

triggers

tuples

Ontology Databases

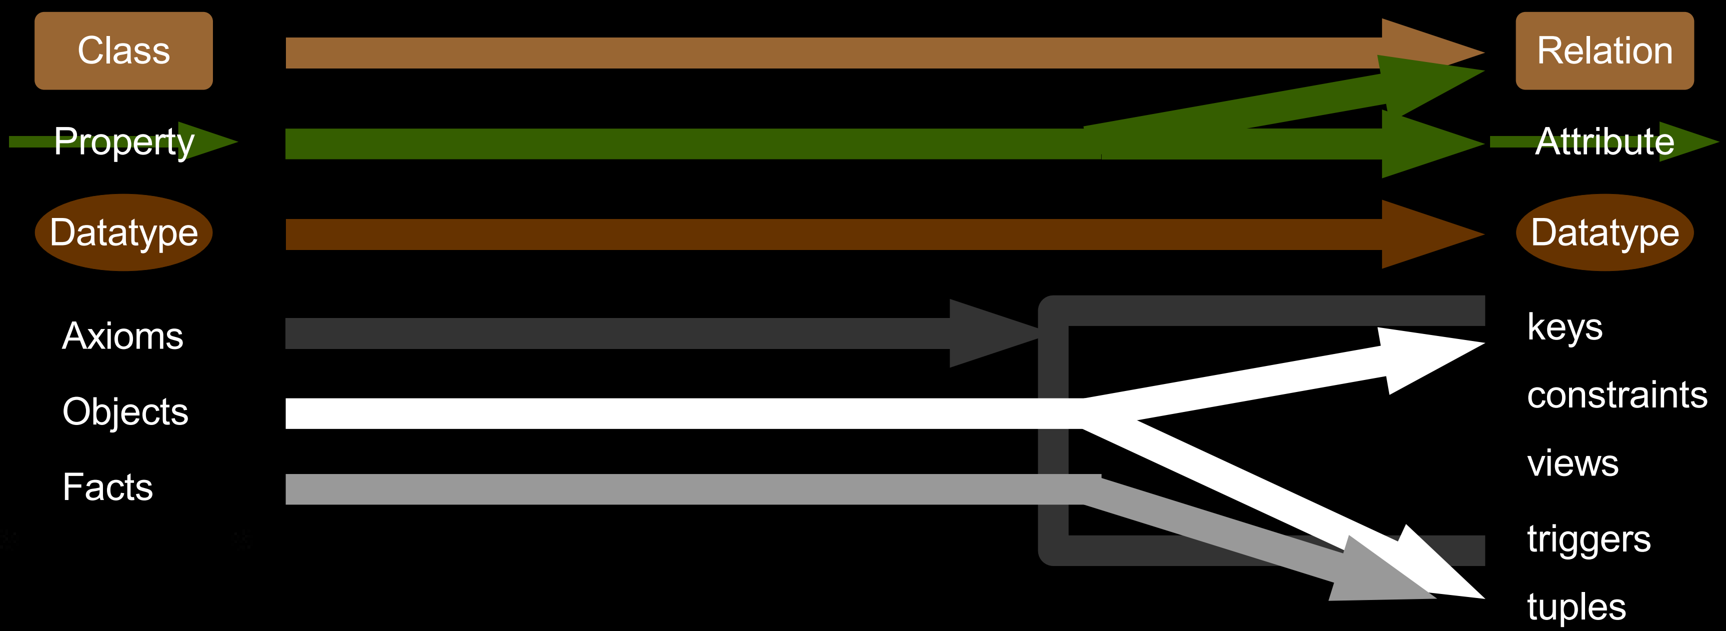

Ontology Databases

Class		Relation
Property		Attribute
Datatype		Datatype
Axioms		keys
Objects	So what?	constraints
Facts		
		triggers
		tuples

A Simple Problem (revisited)

This is what we know :

All sisters are siblings.

Hilary and Lynn are sisters.

This is what we want to know :

Who are siblings?

Obviously, the answer should be :

Hilary and Lynn are siblings.

A Simple Problem (revisited)

<u>This is what we know :</u>

All sisters are siblings.

Hilary and Lynn are sisters.

SiblingOf

Subject	Object

SisterOf

Subject	Object

<u>This is what we want to know :</u>

Who are siblings?

<u>Obviously, the answer should be :</u>

Hilary and Lynn are siblings.

A Simple Problem (revisited)

<u>This is what we know :</u>

All sisters are siblings.

Hilary and Lynn are sisters.

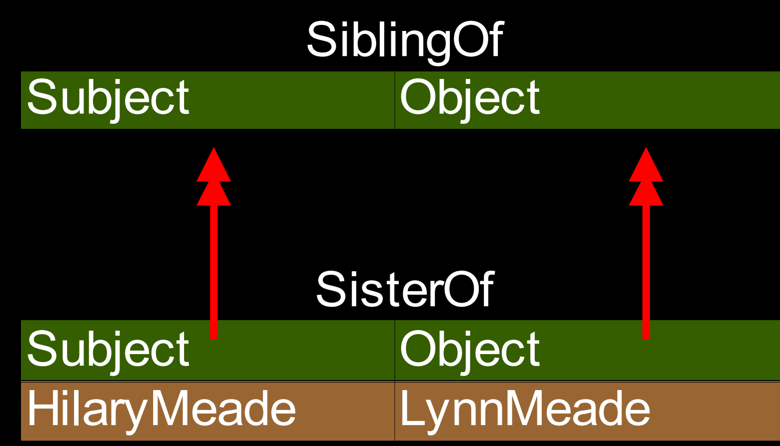

<u>This is what we want to know :</u>

Who are siblings?

<u>Obviously, the answer should be :</u>

Hilary and Lynn are siblings.

A Simple Problem (revisited)

This is what we know :

All sisters are siblings.

Hilary and Lynn are sisters.

SiblingOf

Subject	Object
HilaryMeade	LynnMeade

SisterOf

Subject	Object
HilaryMeade	LynnMeade

This is what we want to know :

Who are siblings?

Obviously, the answer should be :

Hilary and Lynn are siblings.

A Simple Problem (revisited)

This is what we know :

All sisters are siblings.

Hilary and Lynn are sisters.

SiblingOf

Subject	Object
HilaryMeade	LynnMeade

SisterOf

Subject	Object
HilaryMeade	LynnMeade

This is what we want to know :

Who are siblings?

Obviously, the answer should be :

Hilary and Lynn are siblings.

A Simple Problem (revisited)

<u>This is what we know :</u>

All sisters are siblings.

Hilary and Lynn are sisters.

SiblingOf

Subject	Object
HilaryMeade	LynnMeade

SisterOf

Subject	Object
HilaryMeade	LynnMeade

<u>This is what we want to know :</u>

Who are siblings?

$\{ \langle x,y \rangle \mid \text{siblingOf}(x,y) \}$

<u>Obviously, the answer should be :</u>

Hilary and Lynn are siblings.

A Simple Problem (revisited)

<u>This is what we know :</u>

All sisters are siblings.

Hilary and Lynn are sisters.

SiblingOf

Subject	Object
HilaryMeade	LynnMeade

SisterOf

Subject	Object
HilaryMeade	LynnMeade

<u>This is what we want to know :</u>

Who are siblings?

$\{ \langle x,y \rangle \mid \text{siblingOf}(x,y) \}$

<u>Obviously, the answer should be :</u>

Hilary and Lynn are siblings.

Just look it up!

A Simple Problem (revisited)

<u>This is what we know :</u>

All sisters are siblings.

Hilary and Lynn are sisters.

SiblingOf

Subject	Object
HilaryMeade	LynnMeade

SisterOf

Subject	Object
HilaryMeade	LynnMeade

<u>This is what we want to know :</u>

Who are siblings?

$\{ \langle x,y \rangle \mid \text{siblingOf}(x,y) \}$

<u>Obviously, the answer should be :</u>

Hilary and Lynn are siblings.

Just look it up!

A Simple Problem (revisited)

<u>This is what we know :</u>

All sisters are siblings.

Hilary and Lynn are sisters.

SiblingOf

Subject	Object
HilaryMeade	LynnMeade

SisterOf

Subject	Object
HilaryMeade	LynnMeade

<u>This is what we want to know :</u>

Who are siblings?

$\{ \langle x,y \rangle \mid \text{siblingOf}(x,y) \}$

<u>Obviously, the answer should be :</u>

Hilary and Lynn are siblings.

{ <Hilary, Lynn> }

A Data-Driven Search

This process is data-driven,

loosely based on forward chaining.

A Data-Driven Search

This process is data-driven,

loosely based on forward chaining.

Clearly, we are trading space for query time.

A Data-Driven Search

This process is data-driven,

loosely based on forward chaining.

Clearly, we are trading space for query time.

(We eagerly propagate data.)

Key Question #2

In eagerly propagating data, do we incur a significant load-time cost?

Key Question #2

In eagerly propagating data, do we incur a significant load-time cost?

Probably?

Key Question #3

Do we actually improve query time?

Key Question #3

Do we actually improve query time?

Most likely.

Lehigh University Benchmark

A standard benchmarking suite, which includes:

- the university ontology (department, faculty, student...)

- standard dataset generator

- a set of 14 queries testing various features:
    - subsumption depth
    - instance checking
    - meta features (subProperty, inverse)
    - completeness
    - stars and chains (kinds of joins)

[Lehigh University, SWAT lab, under Jeff Heflin's direction]

Lehigh University Benchmark

ParentClass	Class
AdministrativeStaff	SystemsStaff
Course	GraduateCourse
Employee	Faculty
Faculty	Lecturer
Faculty	PostDoc
Faculty	Professor
Object	Director
Object	Employee
Object	Organization
Object	Person
Object	Publication
Object	Schedule
Object	Student
Object	TeachingAssistant
Object	Work
Organization	Department
Organization	ResearchGroup
Organization	University
Person	GraduateStudent
Professor	AssistantProfessor
Professor	AssociateProfessor
Professor	FullProfessor
Publication	Software
Publication	Specification
Student	ResearchAssistant
Student	UndergraduateStudent
Work	Course
Work	Research
...etc...	

Property
advisor
affiliatedOrganizationOf
affiliateOf
degreeFrom
doctoralDegreeFrom
emailAddress
hasAlumnus
headOf
listedCourse
mastersDegreeFrom
member
memberOf
name
officeNumber
publicationAuthor
publicationDate
publicationResearch
researchInterest
researchProject
softwareDocumentation
softwareVersion
subOrganizationOf
takesCourse
teacherOf
teachingAssistantOf
title
undergraduateDegreeFrom
worksFor
...etc...

Lehigh University Benchmark

Radial Tree View

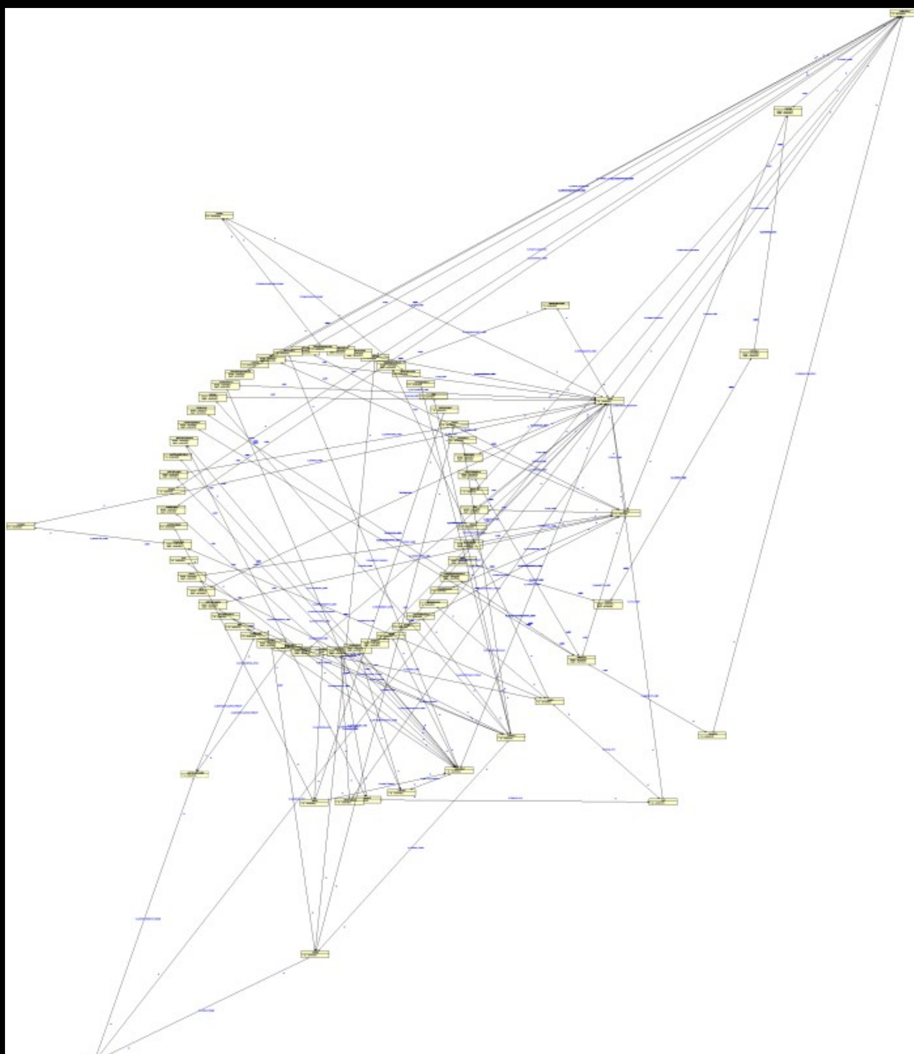

Radial Isometric View

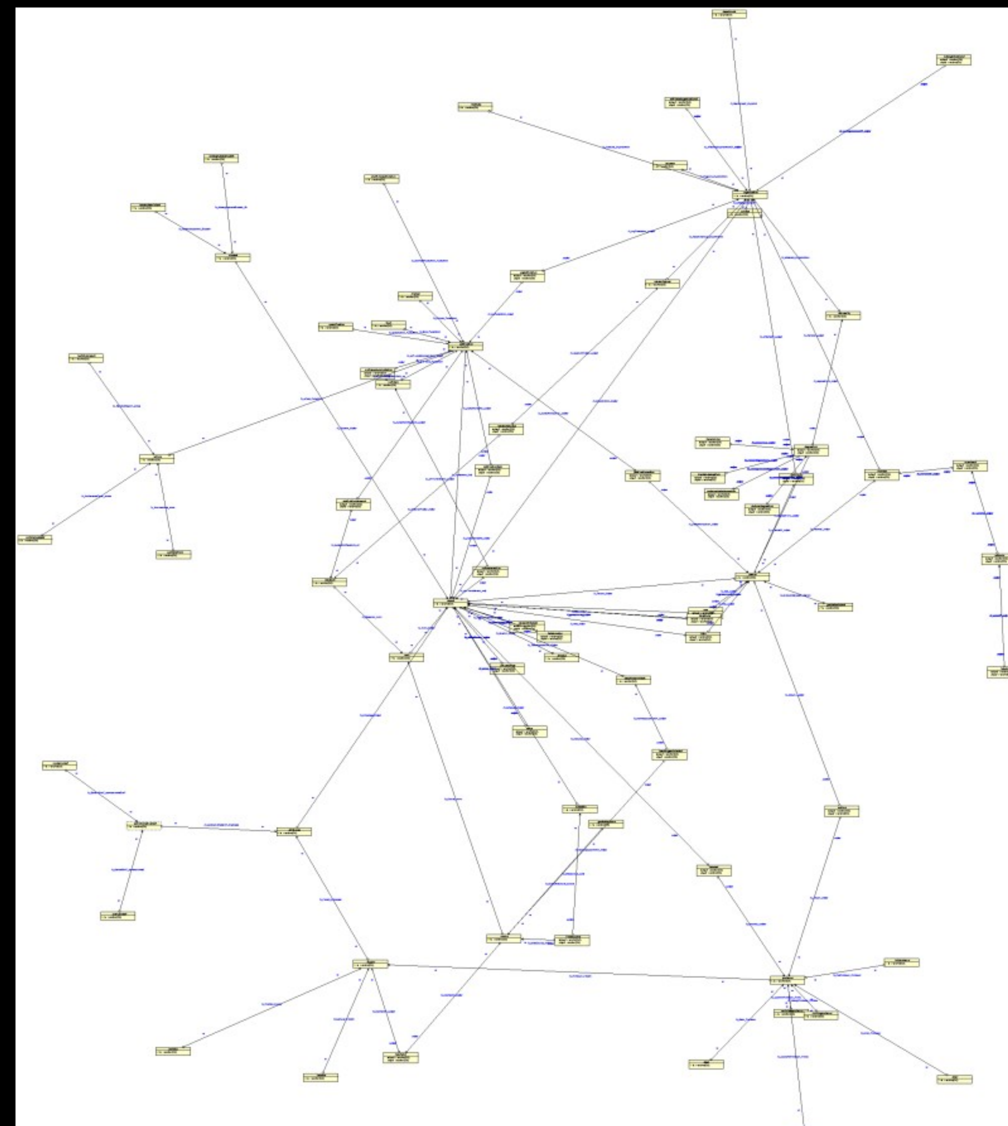

Lehigh University Benchmark

Load Time (1.5 million facts)
(10 Universities, 20 Departments)

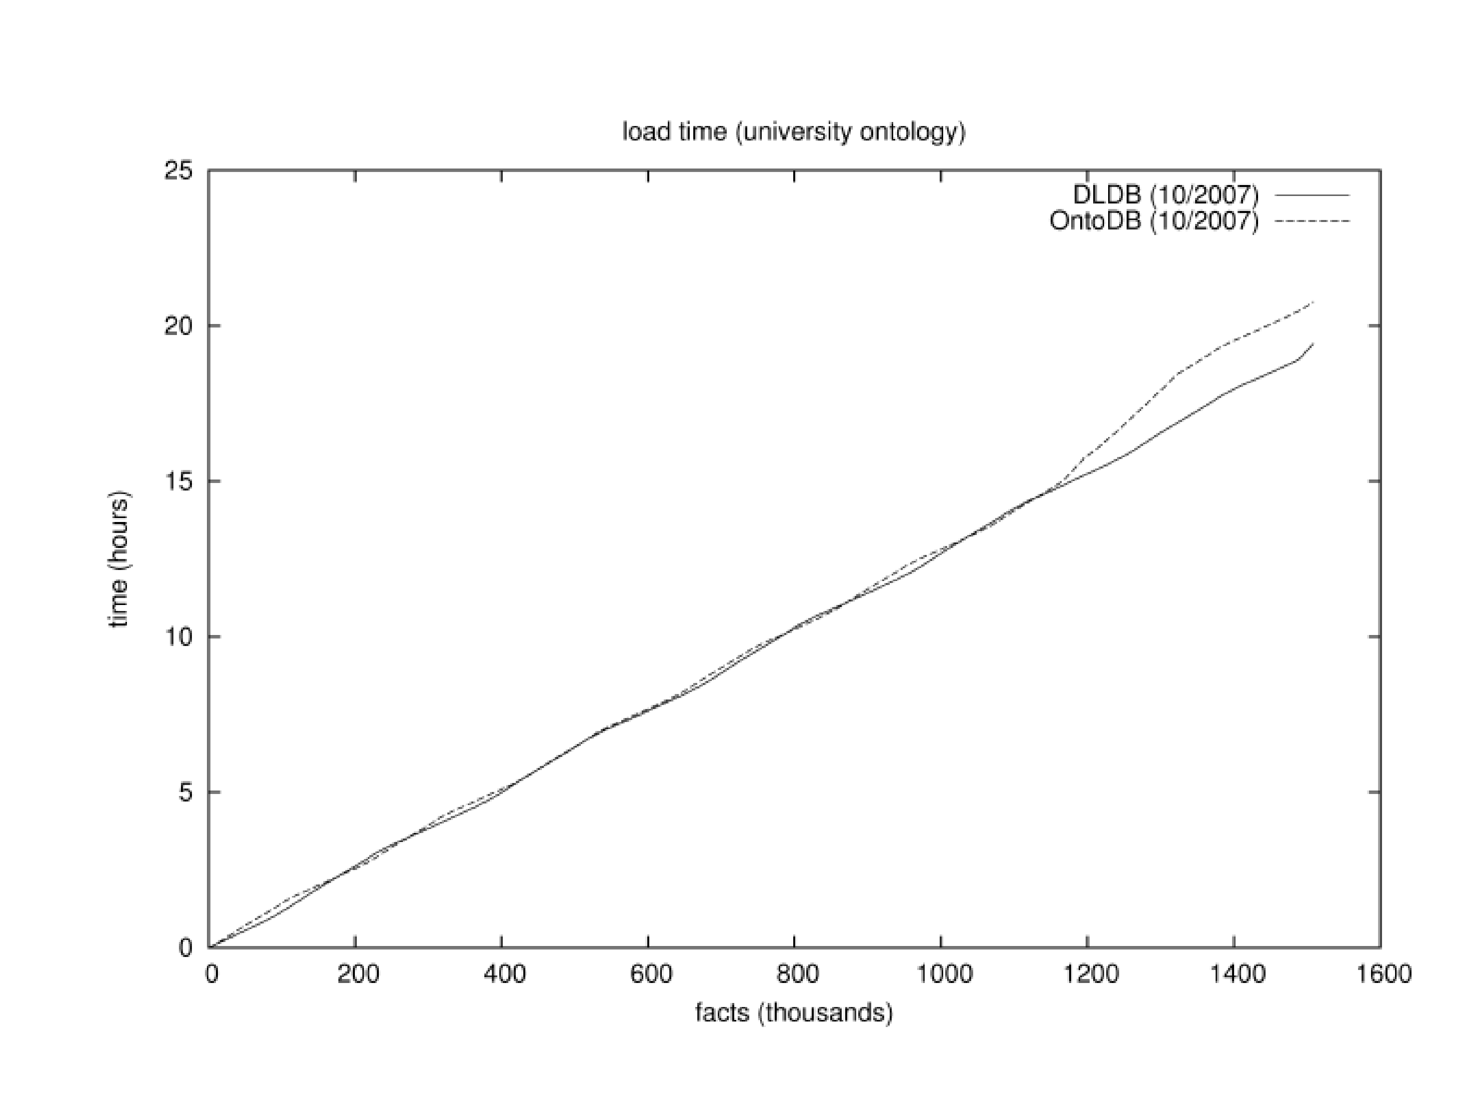

Lehigh University Benchmark

In trading space, do we incur a significant load-time cost?

No!

Lehigh University Benchmark

In trading space, do we incur a significant load-time cost?

No!

(This was surprising.)

Lehigh University Benchmark

Do we actually improve query time?

Lehigh University Benchmark

Query Performance

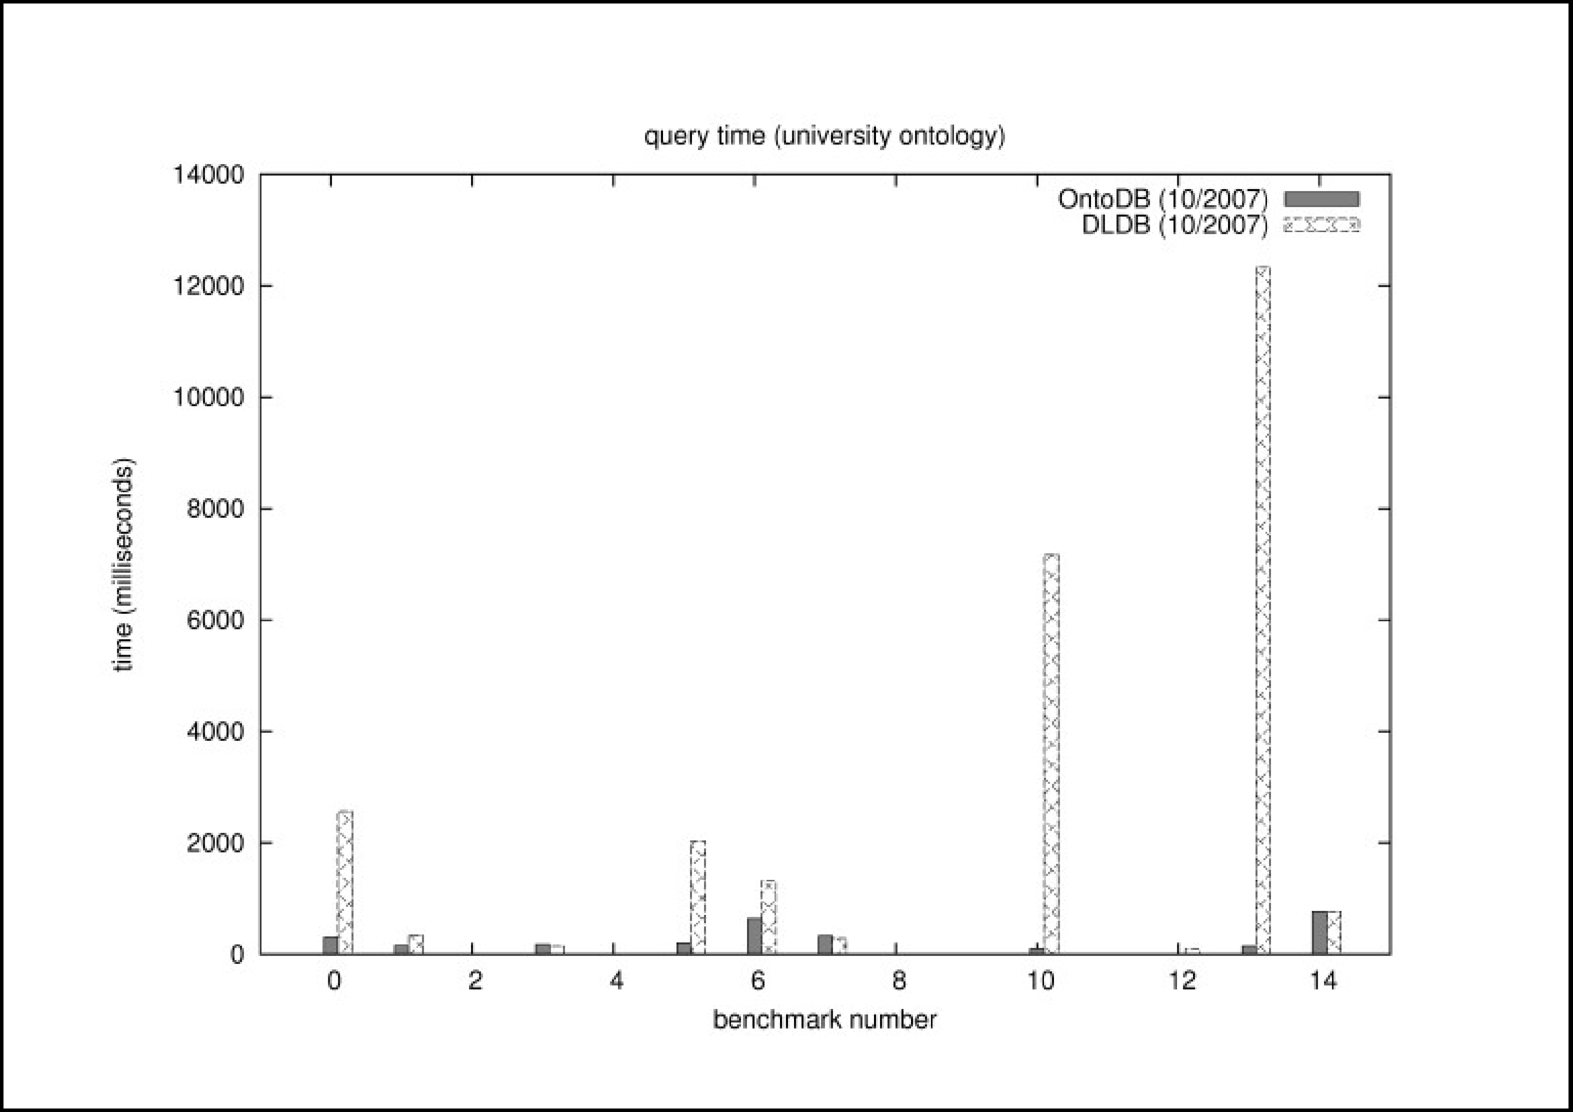

Lehigh University Benchmark

Query Performance
(logarithmic time)

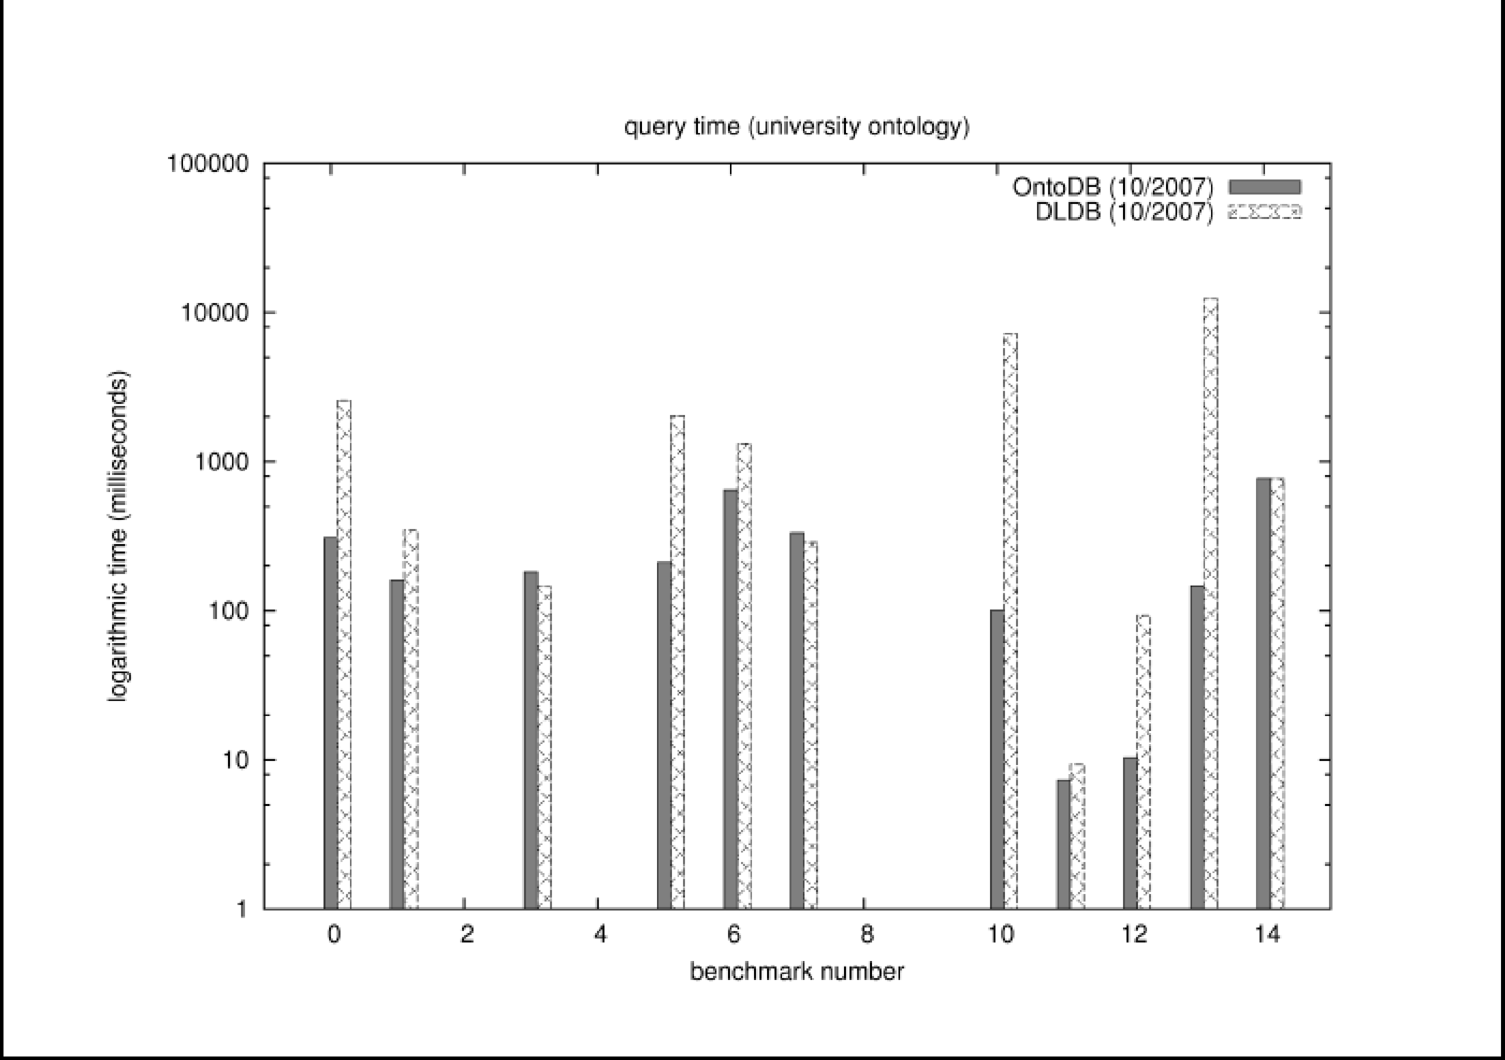

Lehigh University Benchmark

Do we actually improve query time?

Yes!

Lehigh University Benchmark

Do we actually improve query time?

Yes!

As we expected.

NEMO

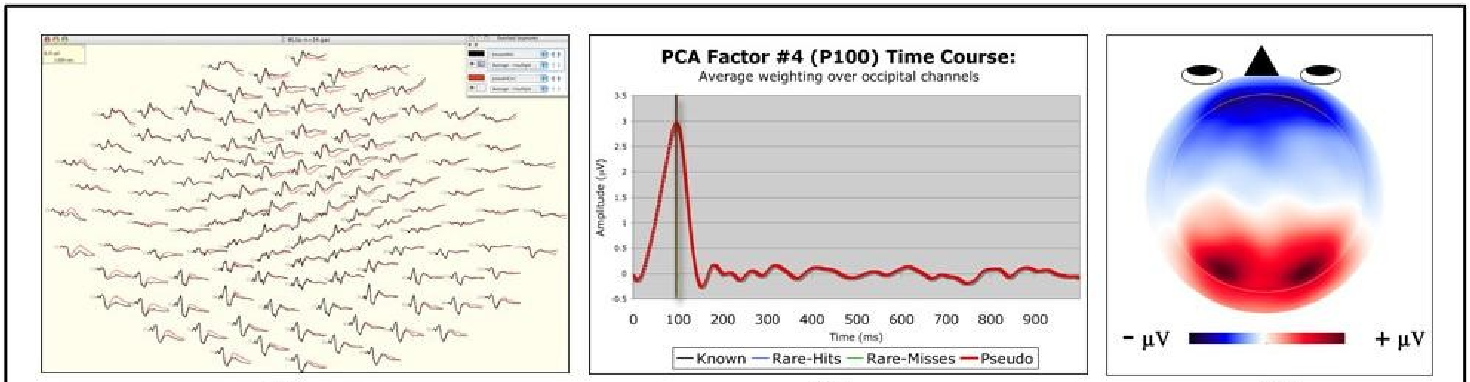

PCA Factor #4 (P100) Time Course:
Average weighting over occipital channels
Amplitude (µV)
Time (ms)
Known
Rare-Hits
Rare-Misses
Pseudo
- µV
+ µV

(A) (B) (C)

NEMO

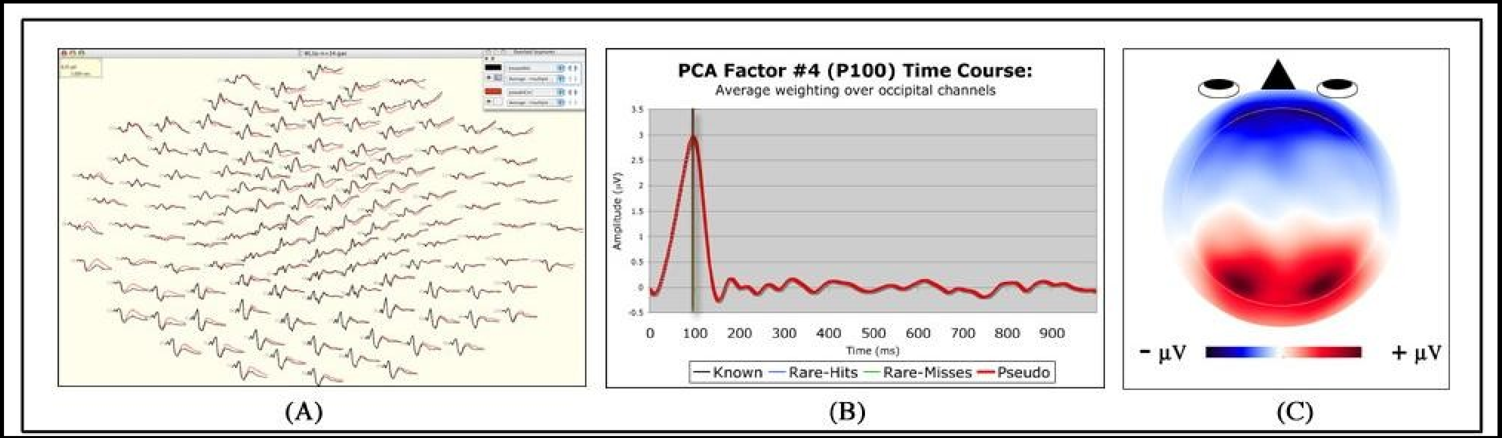

Expert queries answered 100% correctly.

Less than 10 millisecond average response time, regardless of query complexity.

NEMO

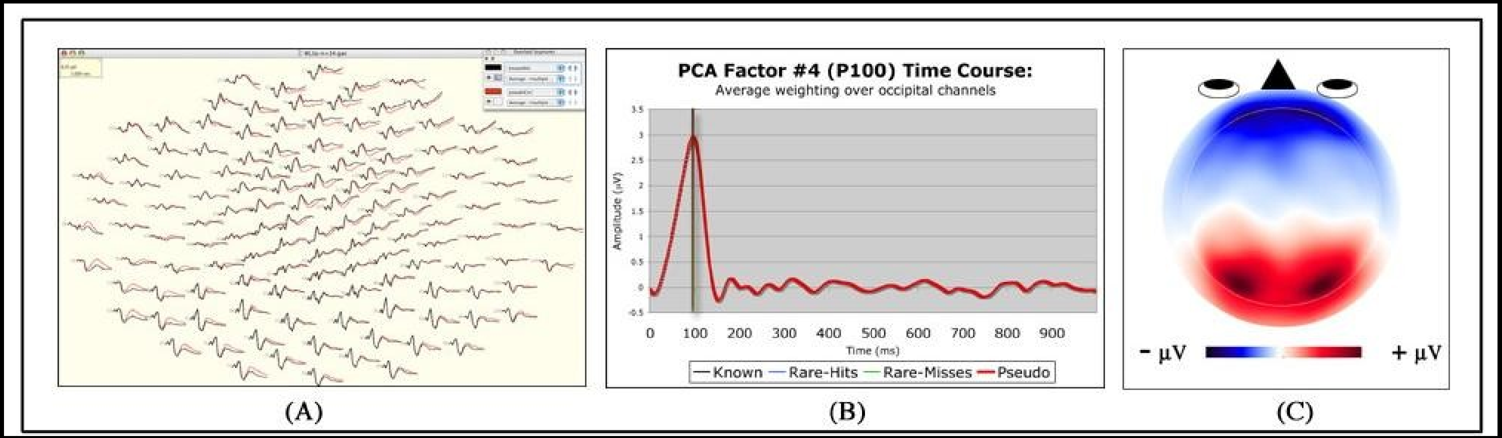

(A) (B) (C)

- *Show the region of interest for all ERP patterns that occur between 0 and 300ms.*

- *Which PCA factor do P100 patterns most often appear in?*

- *What is the range of intensity mean for the region of interest for N100 patterns?*

- *Show the patterns whose region of interest is left occipital and occurs between 220 and 300ms.*

NEMO

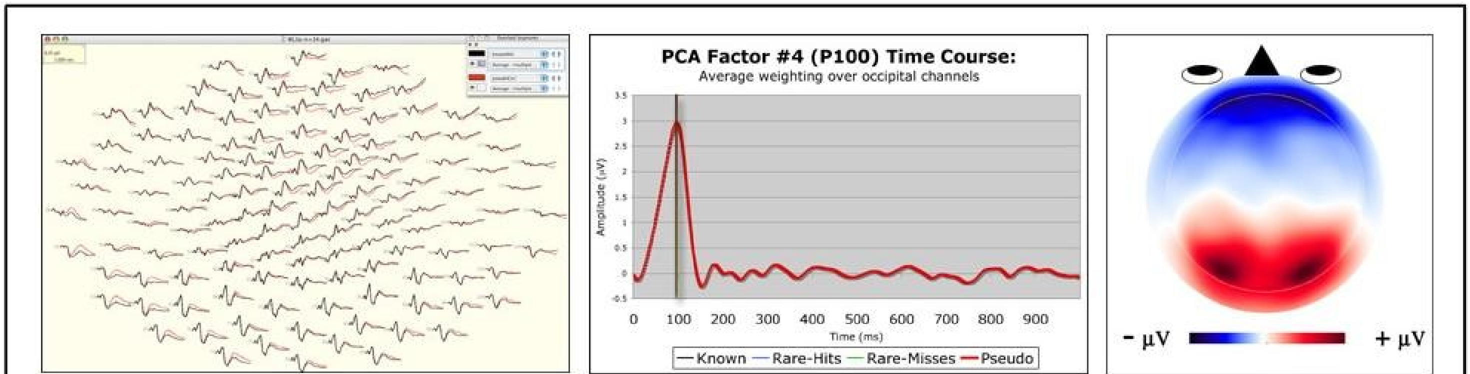

(A) (B) (C)

Main points:

NEMO

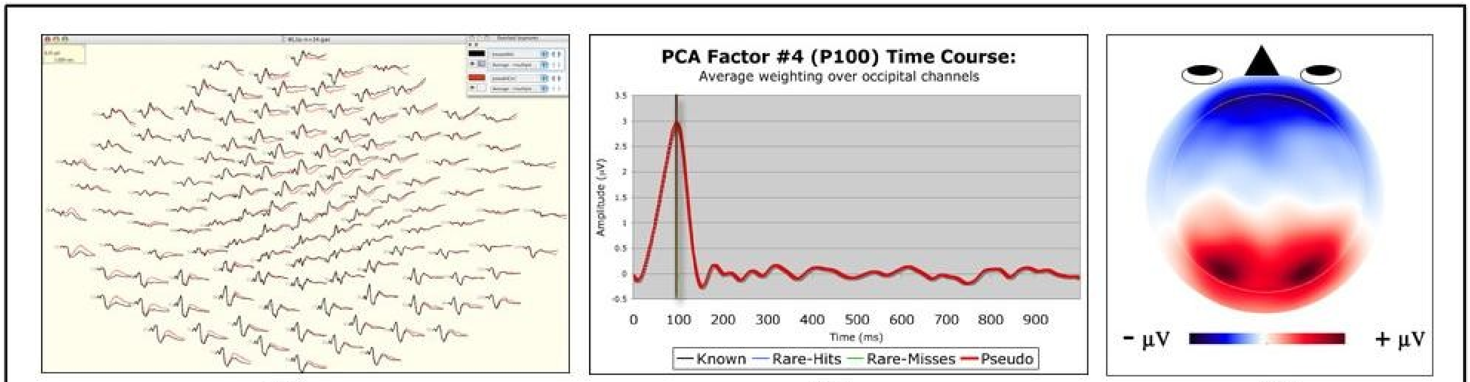

(A) (B) (C)

Main points:

Ontology-based Modeling

NEMO

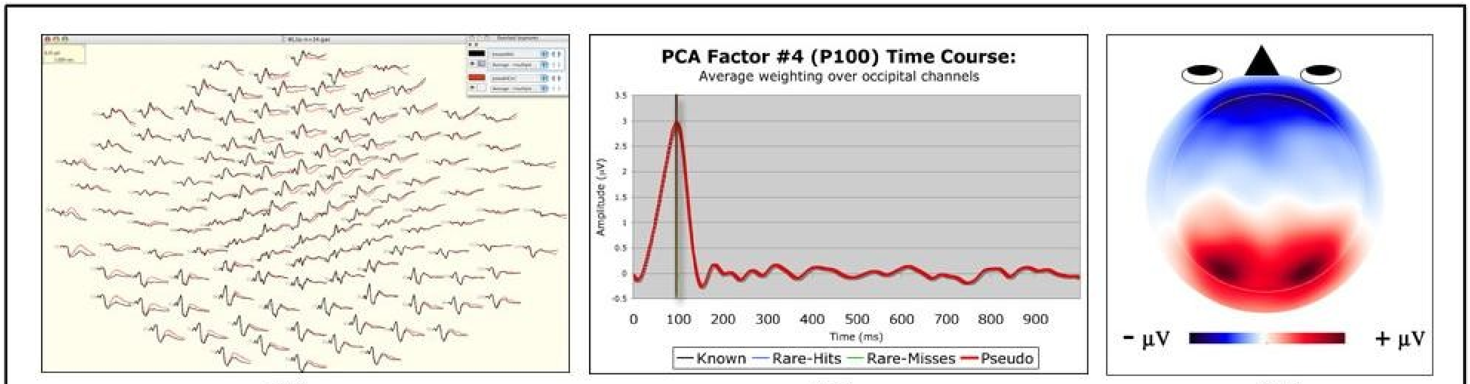

(A) (B) (C)

Main points:

- Ontology-based Modeling
- Ontology-based Query Answering Process

NEMO

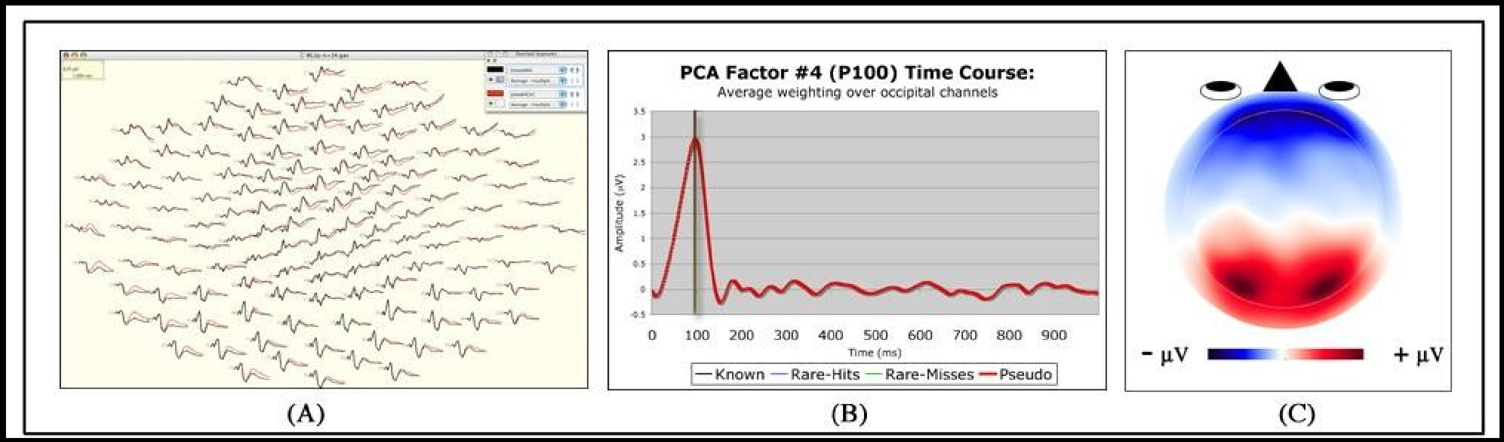

Main points:

- Ontology-based Modeling
- Ontology-based Query Answering Process
- Cross-lab information modeling, storage and analysis

Ontology Databases

Ongoing Work

Disjunctive Logical Models
Scalable T-Box Reasoning (model-based)
Meta-analyses (cross-lab integration)

Thank you!

paea@cs.uoregon.edu

Ontology Databases

Questions?

paea@cs.uoregon.edu

Ontology Databases

Class

Property

Datatype

Axioms

Objects

Facts

Relation

Attribute

Datatype

keys

constraints

views

triggers

tuples

Ontology Databases

Class

Property

Datatype

Axioms

Objects

Facts

datatype-property

Relation

Attribute

Datatype

keys

constraints

views

triggers

tuples

Ontology Databases

Ontology Databases

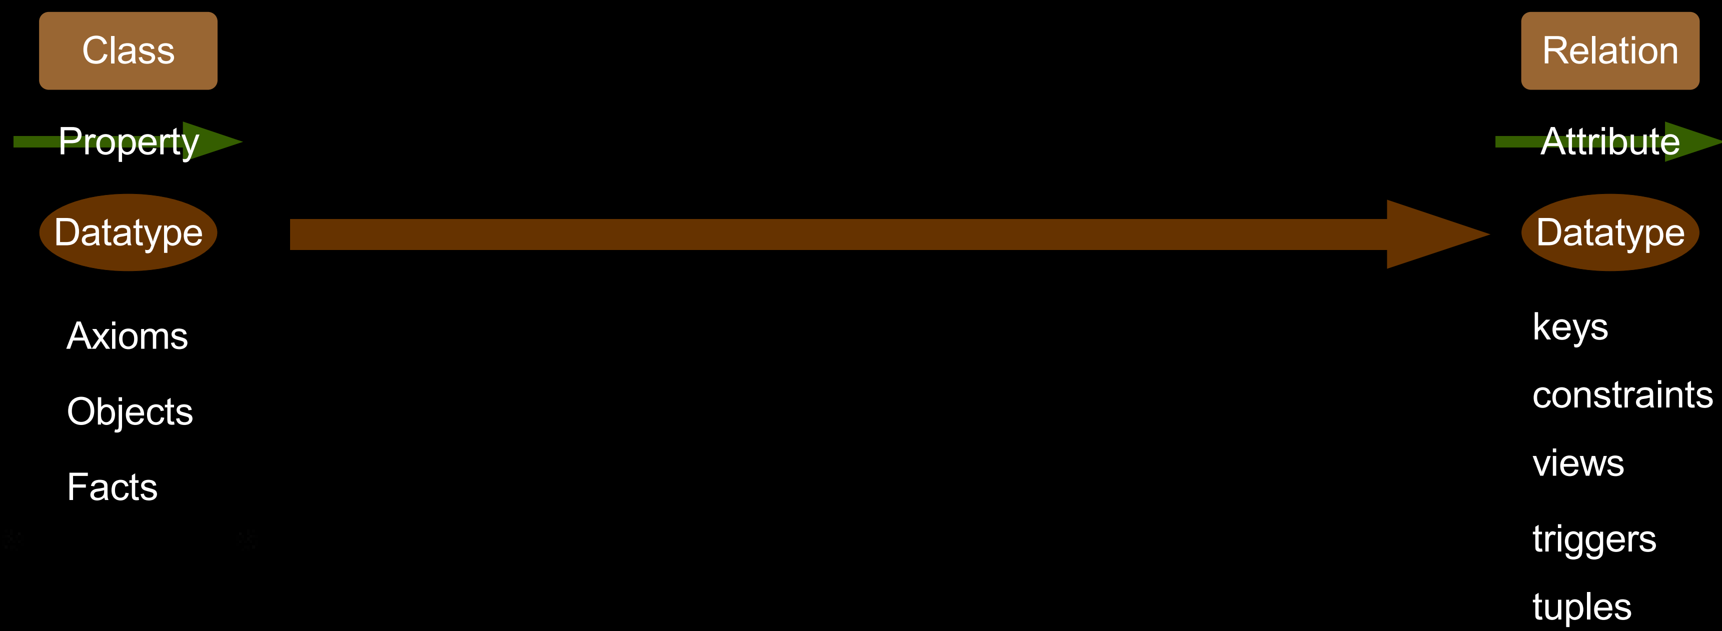

Ontology Databases

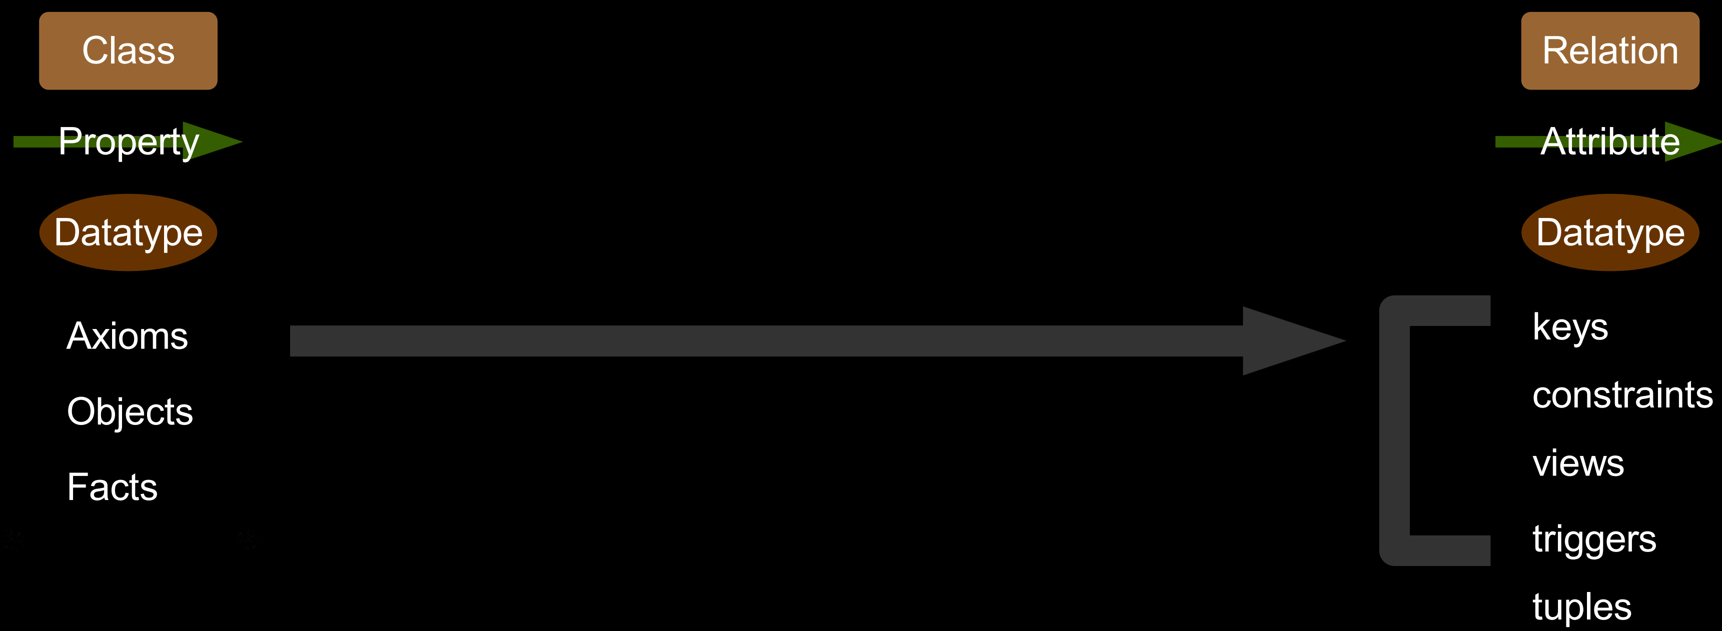

Ontology Databases

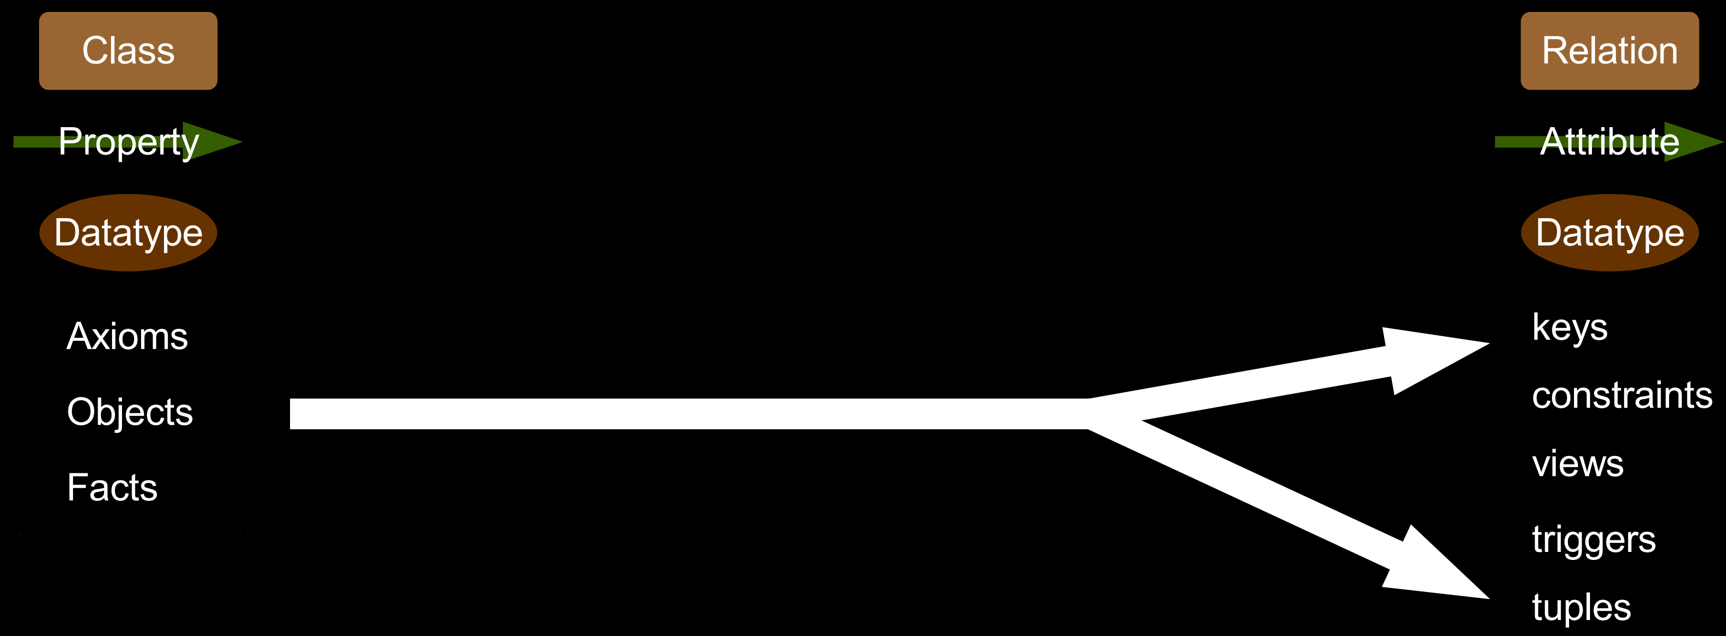

Ontology Databases

Class

Property

Datatype

Axioms

Objects

Facts

Relation

Attribute

Datatype

keys

constraints

views

triggers

tuples