

Ontology-Based Data Access

ISWC 2007 Tutorial

Diego Calvanese¹, Domenico Lembo²

¹ Free University of Bozen-Bolzano,



FREIE UNIVERSITÄT BOZEN
LIBERA UNIVERSITÀ DI BOLZANO
FREE UNIVERSITY OF BOZEN - BOLZANO

² Sapienza Università di Roma



SAPIENZA
UNIVERSITÀ DI ROMA

6th Int. Semantic Web Conference
Busan, South Korea – Nov. 12th, 2007



Structure of the tutorial

- 1 Introduction to Ontology-Based Data Access
 - 1 Introduction to ontologies
 - 2 Ontology languages
- 2 Description Logics and the *DL-Lite* family
 - 1 A gentle introduction to DLs
 - 2 DLs as a formal language to specify ontologies
 - 3 Queries in Description Logics
 - 4 The *DL-Lite* family of tractable DLs
- 3 Reasoning in the *DL-Lite* family
 - 1 TBox reasoning
 - 2 TBox & ABox reasoning
 - 3 Complexity of reasoning in Description Logics
- 4 Linking data to ontologies
 - 1 The Description Logic *DL-Lite_A*
 - 2 Connecting ontologies to relational data
- 5 Hands-on session

Part I

Introduction to Ontology-Based Data Access

Outline

- 1 Introduction to ontologies
- 2 Ontology languages

Outline

- 1 Introduction to ontologies
 - Ontologies in information systems
 - Challenges related to ontologies
- 2 Ontology languages

Different meanings of “Semantics”

- ① Part of **linguistics** that studies the meaning of words and phrases.
- ② **Meaning** of a set of symbols in some **representation scheme**.
Provides a means to specify and communicate the intended meaning of a set of “syntactic” objects.
- ③ **Formal semantics of a language** (e.g., an artificial language).
(Meta-mathematical) mechanism to associate to each sentence in a language an element of a symbolic domain that is “outside the language”.

In **information systems** meanings 2 and 3 are the relevant ones:

- In order to talk about semantics we need a representation scheme, i.e., an **ontology**.
- ... but 2 makes no sense without 3.

Ontologies

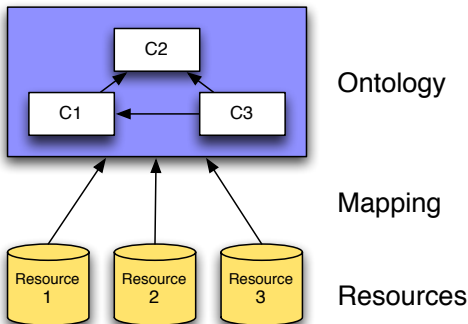
Definition

An **ontology** is a representation scheme that describes a **formal conceptualization** of a domain of interest.

The specification of an ontology comprises several levels:

- **Meta-level**: specifies a set of **modeling categories**.
- **Intensional level**: specifies a set of **conceptual elements** (instances of categories) and of rules to describe the conceptual structures of the domain.
- **Extensional level**: specifies a set of **instances** of the conceptual elements described at the intensional level.

Ontologies at the core of information systems



The usage of all system resources (data and services) is done through the domain conceptualization.

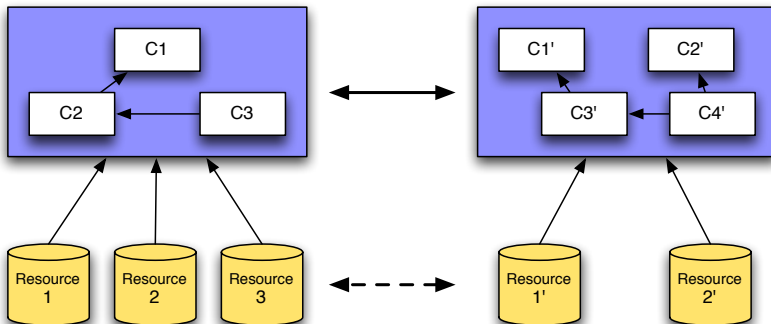
Ontology mediated data access

Desiderata: achieve **logical transparency** in access to data:

- **Hide** to the user where and how data are stored.
- Present to the user a **conceptual view** of the data.
- Use a **semantically rich formalism** for the conceptual view.

Similar to Data Integration, but with a rich conceptual description as the global view.

Ontologies at the core of cooperation



The cooperation between systems is done at the level of the conceptualization.

Three novel challenges

- 1 Languages
- 2 Methodologies
- 3 Tools

... for specifying, building, and managing ontologies to be used in information systems.

The first challenge: ontology languages

- Several proposals for ontology languages have been made.
- **Tradeoff** between **expressive power** of the language and **computational complexity** of dealing with (i.e., performing inference over) ontologies specified in that language.
- Usability needs to be addressed.

In this tutorial:

- We propose variants of ontology languages suited for managing **ontologies in information systems**.
- We discuss in depth the above mentioned **tradeoff** . . .
- . . . paying particular attention to the aspects related to data management.

The second challenge: methodologies

- Building and dealing with ontologies is a complex and challenging task.
- Building **good ontologies** is even more challenging.
- It requires to master the technologies based on semantics, which in turn requires good knowledge about the languages, their semantics, and the implications it has w.r.t. reasoning over the ontology.

In this tutorial:

- We study in depth the **semantics of ontologies**, with an emphasis on their relationship to data in information sources.
- We thus lay the **foundations for the development of methodologies**, though we do not present specific methodologies here.

The third challenge: tools

- According to the principle that “there is no meaning without a language with a formal semantics”, the formal semantics becomes the solid basis for dealing with ontologies.
- Hence every kind of access to an ontology (to extract information, to modify it, etc.), requires to **fully** take into account its semantics.
- We need to resort to tools that provide capabilities to perform **automated reasoning** over the ontology, and the kind of reasoning should be **sound** and **complete** w.r.t. the formal semantics.

In this tutorial:

- We discuss the requirements for such ontology management tools.
- We present a tool that has been specifically designed for optimized access to information sources through ontologies.

A challenge across the three challenges: scalability

When we want to use ontologies to access information sources, we have to address the three challenges of languages, methodologies, and tools by taking into account **scalability** w.r.t.:

- the size of (the intensional level of) the ontology
- the number of ontologies
- the **size of the information sources** that are accessed through the ontology/ontologies.

In this tutorial we pay particular attention to the third aspect, since we work under the realistic assumption that the extensional level (i.e., the data) largely dominates in size the intensional level of an ontology.

Outline

- 1 Introduction to ontologies
- 2 **Ontology languages**
 - Elements of an ontology language
 - Intensional level of an ontology language
 - Extensional level of an ontology language
 - Ontologies and other formalisms
 - Queries

Elements of an ontology language

- **Syntax**
 - Alphabet
 - Languages constructs
 - Sentences to assert knowledge
- **Semantics**
 - Formal meaning
- **Pragmatics**
 - Intended meaning
 - Usage

Static vs. dynamic aspects

The aspects of the domain of interest that can be modeled by an ontology language can be classified into:

- **Static aspects**
 - Are related to the structuring of the domain of interest.
 - Supported by virtually all languages.
- **Dynamic aspects**
 - Are related to how the elements of the domain of interest evolve over time.
 - Supported only by some languages, and only partially (cf. services).

Before delving into the dynamic aspects, we need a good understanding of the static ones.

In this tutorial we concentrate essentially on the static aspects.

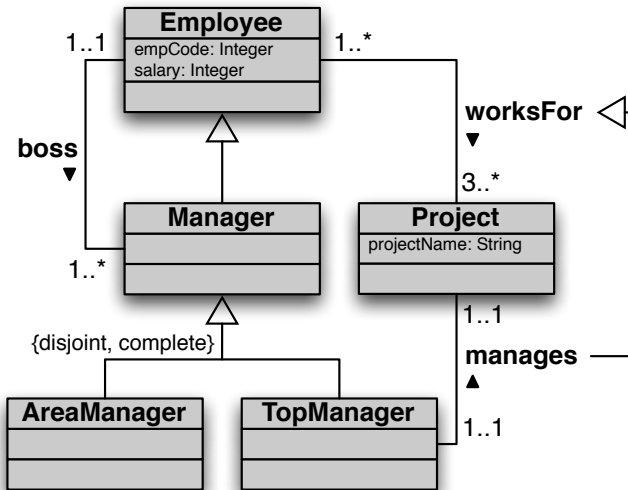
Intensional level of an ontology language

An ontology language for expressing the intensional level usually includes:

- Concepts
- Properties of concepts
- Relationships between concepts, and their properties
- Axioms
- Individuals and facts about individuals
- Queries

Ontologies are typically **rendered as diagrams** (e.g., Semantic Networks, Entity-Relationship schemas, UML Class Diagrams).

Example: ontology rendered as UML Class Diagram



Concepts

Concept

Is an element of the ontology that denotes a collection of instances (e.g., the set of “employees”).

We distinguish between:

- **Intensional definition:**
specification of **name**, **properties**, **relations**, ...
- **Extensional definition:**
specification of the **instances**

Concepts are also called **classes**, **entity types**, **frames**.

Properties

Property

Qualifies an element (e.g., a concept) of an ontology.

Property definition (intensional and extensional):

- Name
- Type:
 - Atomic (integer, real, string, enumerated, ...)
e.g., `eye-color` → { `blu`, `brown`, `green`, `grey` }
 - Structured (date, sets, lists, ...)
e.g., `date` → `day/month/year`
- Default value

Properties are also called **attributes**, **features**, **slots**.

Relationships

Relationship

Expresses an association among concepts.

We distinguish between:

- **Intensional definition:**
specification of involved **concepts**
e.g., **worksFor** is defined on **Employee** and **Project**
- **Extensional definition:**
specification of the instances of the relationship, called **facts**
e.g., **worksFor(domenico, TONES)**

Relationships are also called **associations**, **relationship types**, **roles**.

Axioms

Axiom

Is a logical formula that expresses at the intensional level a condition that must be satisfied by the elements at the extensional level.

Different kinds of axioms/conditions:

- subclass relationships, e.g., $\text{Manager} \sqsubseteq \text{Employee}$
- equivalences, e.g., $\text{Manager} \equiv \text{AreaManager} \sqcup \text{TopManager}$
- disjointness, e.g., $\text{AreaManager} \sqcap \text{TopManager} \equiv \perp$
- (cardinality) restrictions,
e.g., each Employee worksFor at least 3 Project
- ...

Axioms are also called **assertions**.


A special kind of axioms are **definitions**.

Extensional level of an ontology language

At the extensional level we have individuals and facts:

- An **instance** represents an individual (or object) in the extension of a concept.
e.g., `instanceOf(domenico, Employee)`
- A **fact** represents a relationship holding between instances.
e.g., `worksFor(domenico, TONES)`

The three levels of an ontology



Comparison with other formalisms

- Ontology languages vs. knowledge representation languages:
Ontologies **are** knowledge representation schemas.
- Ontology vs. logic:
Logic is a **the** tool for assigning semantics to ontology languages.
- Ontology languages vs. conceptual data models:
Conceptual schema **are** special ontologies, suited for conceptualizing a **single** logical model (database).
- Ontology languages vs. programming languages:
Class definitions **are** special ontologies, suited for conceptualizing a **single** structure for computation.

Classification of ontology languages

- Graph-based
 - Semantic networks
 - Conceptual graphs
 - UML
- Frame based
 - Frame Systems
 - OKBC, XOL
- Logic based
 - **Description Logics** (e.g., *SHOIQ*, *DLR*, *DL-Lite*, *OWL*, ...)
 - Rules (e.g., RuleML, LP/Prolog, F-Logic)
 - First Order Logic (e.g., KIF)
 - Non-classical logics (e.g., Nonmonotonic, probabilistic)

Queries

An ontology language may also include constructs for expressing queries.

Query

In an expression at the intensional level denoting a (possibly structured) collection of individuals satisfying a given condition.

Meta-Query

In an expression at the meta level denoting a collection of ontology elements satisfying a given condition.

Note: One may also conceive queries that span across levels (**object-meta queries**), cf. [RDF, Cali&Kifer VLDB'06]

Ontology languages vs. query languages

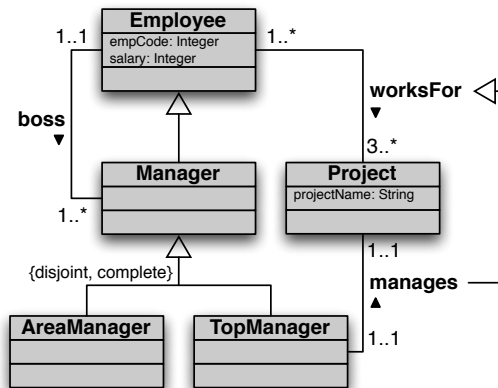
Ontology languages:

- Tailored for capturing intensional relationships.
- Are quite **poor as query languages**:
 - Cannot refer to same object via multiple navigation paths in the ontology,
 - i.e., allow only for a limited form of JOIN, namely chaining.

Instead, **when querying** a data source (either directly, or via the ontology), to retrieve the data of interest, general forms of **joins** are required.

It follows that the constructs for queries may be quite different from the constructs used in the ontology to form concepts and relationships.

Example of query



$q(ce, cm, se, sm) \leftarrow$ worksFor(e, p) \wedge manages(m, p) \wedge boss(m, e) \wedge
 empCode(e, ce) \wedge empCode(m, cm) \wedge
 salary(e, se) \wedge salary(m, sm) $\wedge se \geq sm$

Query answering under different assumptions

There are fundamentally different assumptions when addressing query answering in different settings:


- traditional database assumption
- knowledge representation assumption

Query answering under the database assumption


- Data are completely specified (CWA), and typically large.
- Schema/intensional information used in the design phase.
- At **runtime**, the data is assumed to satisfy the schema, and therefore the **schema is not used**.
- Queries allow for complex navigation paths in the data (cf. SQL).

↪ Query answering amounts to **query evaluation**, which is computationally easy.

Query answering under the database assumption (cont'd)



Query answering under the database assumption – Example



For each concept/relationship we have a (complete) table in the DB.

DB: Employee = { john, mary, nick }
 Manager = { john, nick }
 Project = { prA, prB }
 worksFor = { (john,prA), (mary,prB) }

Query: $q(x) \leftarrow \text{Manager}(x), \text{Project}(p), \text{worksFor}(x, p)$

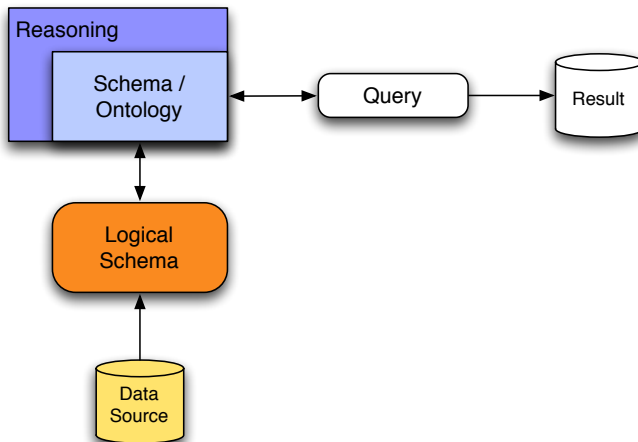
Answer: { john }

Query answering under the KR assumption


- An ontology (or conceptual schema, or knowledge base) imposes constraints on the data.
- Actual data may be incomplete or inconsistent w.r.t. such constraints.
- The system has to take into account intensional information during query answering, and overcome incompleteness or inconsistency.
- Size of the data is not considered critical (comparable to the size of the intensional information).
- Queries are typically simple, i.e., atomic (the name of a concept).

↪ Query answering amounts to **logical inference**, which is computationally more costly.

Query answering under the KR assumption (cont'd)



Query answering under the KR assumption – Example



Partial DB assumption: we have a (complete) table in the database only for some concepts/relationships.

DB: $\text{Manager} = \{ \text{john}, \text{nick} \}$

$\text{Project} = \{ \text{prA}, \text{prB} \}$

$\text{worksFor} = \{ (\text{john}, \text{prA}), (\text{mary}, \text{prB}) \}$


Query: $q(x) \leftarrow \text{Employee}(x)$

Answer: $\{ \text{john}, \text{nick}, \text{mary} \}$

Rewritten query: $q(x) \leftarrow \text{Employee}(x) \vee \text{Manager}(x) \vee \text{worksFor}(\text{prA}, x)$

Query answering under the KR assumption – Example 2

◀ hasFather



Each person has a father, who is a person

Tables in the DB may be **incompletely** specified.

DB: Person = { john, nick, toni }
 hasFather \supseteq { (john,nick), (nick,toni) }

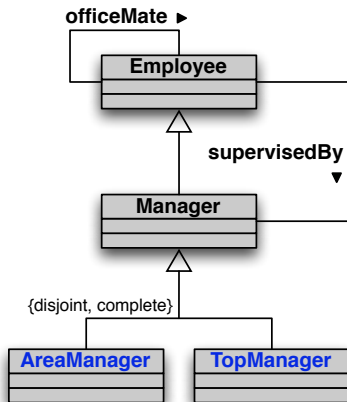
Queries: $q_1(x, y) \leftarrow \text{hasFather}(x, y)$
 $q_2(x) \leftarrow \text{hasFather}(x, y)$
 $q_3(x) \leftarrow \text{hasFather}(x, y_1), \text{hasFather}(y_1, y_2), \text{hasFather}(y_2, y_3)$
 $q_4(x, y_3) \leftarrow \text{hasFather}(x, y_1), \text{hasFather}(y_1, y_2), \text{hasFather}(y_2, y_3)$

Answers: to q_1 : { (john,nick), (nick,toni) }
 to q_2 : { john, nick, toni }
 to q_3 : { john, nick, toni }
 to q_4 : { }


Rewritten queries: see later

QA under the KR assumption – Andrea's Example


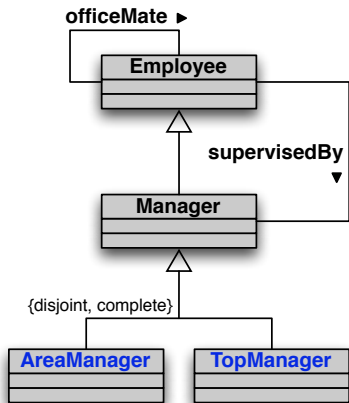
Tables may be **incompletely** specified.



Employee = { andrea, nick, mary, john }
 Manager = { andrea, nick, mary }
 AreaManager \supseteq { nick }
 TopManager \supseteq { mary }
 supervisedBy = { (john, andrea), (john, mary) }
 officeMate = { (mary, andrea), (andrea, nick) }



QA under the KR assumption – Andrea's Example (cont'd)



$q(x) \leftarrow \text{supervisedBy}(x, y), \text{TopManager}(y), \text{officeMate}(y, z), \text{AreaManager}(z)$

Answer: { john }

To determine this answer, we need to resort to reasoning by cases.

Rewritten query? ???There is none (at least not in SQL).

Query answering in Ontology-Based Data Access

In OBDA, we have to face the difficulties of both assumptions:

- The actual **data** is stored in external information sources (i.e., databases), and thus its size is typically **very large**.
- The ontology introduces **incompleteness** of information, and we have to do logical inference, rather than query evaluation.
- We want to take into account at **runtime** the **constraints** expressed in the ontology.
- We want to answer **complex database-like queries**.
- We may have to deal with multiple information sources, and thus face also the problems that are typical of data integration.

Researchers are starting only now to tackle this difficult and challenging problem. In the rest of this tutorial we provide an insight in state-of-the-art technology in this area.

Part II

Description Logics and the *DL-Lite* family

Outline

- 3 A gentle introduction to Description Logics
- 4 DLs as a formal language to specify ontologies
- 5 Queries in Description Logics
- 6 The *DL-Lite* family of tractable Description Logics

Outline

- 3 A gentle introduction to Description Logics
 - Ingredients of Description Logics
 - Description language
 - Description Logics ontologies
 - Reasoning in Description Logics
- 4 DLs as a formal language to specify ontologies
- 5 Queries in Description Logics
- 6 The *DL-Lite* family of tractable Description Logics

What are Description Logics?

Description Logics [BCM⁺03] are **logics** specifically designed to represent and reason on structured knowledge:

The domain is composed of **objects** and is structured into:

- **concepts**, which correspond to classes, and denote sets of objects
- **roles**, which correspond to (binary) relationships, and denote binary relations on objects

The knowledge is asserted through so-called **assertions**, i.e., logical axioms.

Origins of Description Logics

Description Logics stem from early days Knowledge Representation formalisms (late '70s, early '80s):

- Semantic Networks: graph-based formalism, used to represent the meaning of sentences
- Frame Systems: frames used to represent prototypical situations, antecedents of object-oriented formalisms

Problems: **no clear semantics**, reasoning not well understood

Description Logics (a.k.a. Concept Languages, Terminological Languages) developed starting in the mid '80s, with the aim of providing semantics and inference techniques to knowledge representation systems.

Current applications of Description Logics

DLs have evolved from being used “just” in KR.

Novel applications of DLs:

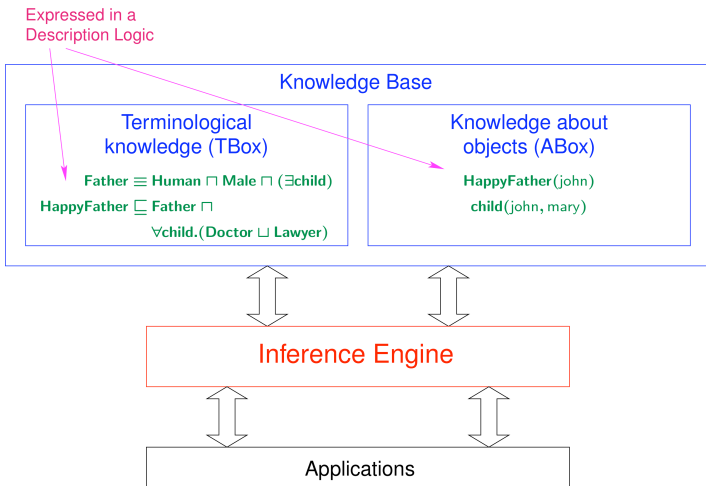
- Databases:
 - schema design, schema evolution
 - query optimization
 - integration of heterogeneous data sources, data warehousing
- Conceptual modeling
- Foundation for the Semantic Web (variants of OWL correspond to specific DLs)
- ...

Ingredients of a Description Logic

A **Description Logic** is characterized by:

- 1 A **description language**: how to form concepts and roles
 $\text{Human} \sqcap \text{Male} \sqcap \exists \text{hasChild} \sqcap \forall \text{hasChild}.(\text{Doctor} \sqcup \text{Lawyer})$
- 2 A mechanism to **specify knowledge** about concepts and roles (i.e., a **TBox**)
 $\mathcal{T} = \{ \text{Father} \equiv \text{Human} \sqcap \text{Male} \sqcap \exists \text{hasChild}, \text{HappyFather} \sqsubseteq \text{Father} \sqcap \forall \text{hasChild}.(\text{Doctor} \sqcup \text{Lawyer}) \}$
- 3 A mechanism to specify **properties of objects** (i.e., an **ABox**)
 $\mathcal{A} = \{ \text{HappyFather}(\text{john}), \text{hasChild}(\text{john}, \text{mary}) \}$
- 4 A set of **inference services**: how to reason on a given KB
 $\mathcal{T} \models \text{HappyFather} \sqsubseteq \exists \text{hasChild}.(\text{Doctor} \sqcup \text{Lawyer})$
 $\mathcal{T} \cup \mathcal{A} \models (\text{Doctor} \sqcup \text{Lawyer})(\text{mary})$

Architecture of a Description Logic system



Description language

A description language is characterized by a set of **constructs** for building **complex concepts** and **roles** starting from atomic ones:

- **concepts** correspond to classes: interpreted as sets of objects
- **roles** corr. to relationships: interpreted as binary relations on objects

Formal semantics is given in terms of interpretations.

An **interpretation** $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$ consists of:

- a nonempty set $\Delta^{\mathcal{I}}$, the domain of \mathcal{I}
- an interpretation function $\cdot^{\mathcal{I}}$, which maps
 - each individual a to an element $a^{\mathcal{I}}$ of $\Delta^{\mathcal{I}}$
 - each atomic concept A to a subset $A^{\mathcal{I}}$ of $\Delta^{\mathcal{I}}$
 - each atomic role P to a subset $P^{\mathcal{I}}$ of $\Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$

The interpretation function is extended to complex concepts and roles according to their syntactic structure.

Concept constructors

Construct	Syntax	Example	Semantics
atomic concept	A	Doctor	$A^I \subseteq \Delta^I$
atomic role	P	hasChild	$P^I \subseteq \Delta^I \times \Delta^I$
atomic negation	$\neg A$	\neg Doctor	$\Delta^I \setminus A^I$
conjunction	$C \sqcap D$	Hum \sqcap Male	$C^I \cap D^I$
(unqual.) exist. res.	$\exists R$	\exists hasChild	$\{ a \mid \exists b. (a, b) \in R^I \}$
value restriction	$\forall R.C$	\forall hasChild.Male	$\{ a \mid \forall b. (a, b) \in R^I \rightarrow b \in C^I \}$
bottom	\perp		\emptyset

(C , D denote arbitrary concepts and R an arbitrary role)

The above constructs form the basic language \mathcal{AL} of the family of \mathcal{AL} languages.

Additional concept and role constructors

Construct	\mathcal{AL}	Syntax	Semantics
disjunction	\mathcal{U}	$C \sqcup D$	$C^{\mathcal{I}} \cup D^{\mathcal{I}}$
top		\top	$\Delta^{\mathcal{I}}$
qual. exist. res.	\mathcal{E}	$\exists R.C$	$\{ a \mid \exists b. (a, b) \in R^{\mathcal{I}} \wedge b \in C^{\mathcal{I}} \}$
(full) negation	\mathcal{C}	$\neg C$	$\Delta^{\mathcal{I}} \setminus C^{\mathcal{I}}$
number restrictions	\mathcal{N}	$(\geq k R)$	$\{ a \mid \#\{b \mid (a, b) \in R^{\mathcal{I}}\} \geq k \}$
		$(\leq k R)$	$\{ a \mid \#\{b \mid (a, b) \in R^{\mathcal{I}}\} \leq k \}$
qual. number restrictions	\mathcal{Q}	$(\geq k R.C)$	$\{ a \mid \#\{b \mid (a, b) \in R^{\mathcal{I}} \wedge b \in C^{\mathcal{I}}\} \geq k \}$
		$(\leq k R.C)$	$\{ a \mid \#\{b \mid (a, b) \in R^{\mathcal{I}} \wedge b \in C^{\mathcal{I}}\} \leq k \}$
inverse role	\mathcal{I}	R^{-}	$\{ (a, b) \mid (b, a) \in R^{\mathcal{I}} \}$
role closure	reg	\mathcal{R}^*	$(R^{\mathcal{I}})^*$

Many different DL constructs and their combinations have been investigated.

Further examples of DL constructs

- Disjunction: $\forall \text{hasChild} . (\text{Doctor} \sqcup \text{Lawyer})$
- Qualified existential restriction: $\exists \text{hasChild} . \text{Doctor}$
- Full negation: $\neg(\text{Doctor} \sqcup \text{Lawyer})$
- Number restrictions: $(\geq 2 \text{ hasChild}) \sqcap (\leq 1 \text{ sibling})$
- Qualified number restrictions: $(\geq 2 \text{ hasChild} . \text{Doctor})$
- Inverse role: $\forall \text{hasChild}^{-} . \text{Doctor}$
- Reflexive-transitive role closure: $\exists \text{hasChild}^* . \text{Doctor}$


Reasoning on concept expressions

An interpretation \mathcal{I} is a **model** of a concept C if $C^{\mathcal{I}} \neq \emptyset$.

Basic reasoning tasks:

- 1 **Concept satisfiability**: does C admit a model?
- 2 **Concept subsumption** $C \sqsubseteq D$: does $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$ hold for all interpretations \mathcal{I} ?

Subsumption used to build the concept hierarchy:



Note: (1) and (2) are mutually reducible if DL is propositionally closed.

Complexity of reasoning on concept expressions

Complexity of concept satisfiability: [DLNN97]

$\mathcal{AL}, \mathcal{ALN}$	PTIME
$\mathcal{ALU}, \mathcal{ALUN}$	NP-complete
\mathcal{ALE}	coNP-complete
$\mathcal{ALC}, \mathcal{ALCN}, \mathcal{ALCI}, \mathcal{ALCQI}$	PSPACE-complete

Observations:

- Two sources of complexity:
 - union (\mathcal{U}) of type NP,
 - existential quantification (\mathcal{E}) of type coNP.

When they are combined, the complexity jumps to PSPACE.

- Number restrictions (\mathcal{N}) do not add to the complexity.

Structural properties vs. asserted properties

We have seen how to build complex **concept and roles expressions**, which allow one to denote classes with a complex structure.

However, in order to represent real world domains, one needs the ability to **assert properties** of classes and relationships between them (e.g., as done in UML class diagrams).

The assertion of properties is done in DLs by means of an **ontology** (or knowledge base).

Description Logics ontology (or knowledge base)

Is a pair $\mathcal{O} = \langle \mathcal{T}, \mathcal{A} \rangle$, where \mathcal{T} is a **TBox** and \mathcal{A} is an **ABox**:

Description Logics **TBox**

Consists of a set of **assertions** on concepts and roles:

- Inclusion assertions on concepts: $C_1 \sqsubseteq C_2$
- Inclusion assertions on roles: $R_1 \sqsubseteq R_2$
- Property assertions on (atomic) roles:

(transitive P)	(symmetric P)	(domain $P C$)
(functional P)	(reflexive P)	(range $P C$) ...

Description Logics **ABox**

Consists of a set of **membership assertions** on individuals:

- for concepts: $A(c)$
- for roles: $P(c_1, c_2)$ (we use c_i to denote individuals)

Description Logics knowledge base – Example

Note: We use $C_1 \equiv C_2$ as an abbreviation for $C_1 \sqsubseteq C_2$, $C_2 \sqsubseteq C_1$.

TBox assertions:

- Inclusion assertions on concepts:

$Father \equiv Human \sqcap Male \sqcap \exists hasChild$
 $HappyFather \sqsubseteq Father \sqcap \forall hasChild.(Doctor \sqcup Lawyer \sqcup HappyPerson)$
 $HappyAnc \sqsubseteq \forall descendant.HappyFather$
 $Teacher \sqsubseteq \neg Doctor \sqcap \neg Lawyer$

- Inclusion assertions on roles:

$hasChild \sqsubseteq descendant$ $hasFather \sqsubseteq hasChild^{-}$

- Property assertions on roles:

(**transitive** descendant), (**reflexive** descendant),
 (**functional** hasFather)

ABox membership assertions:

- $Teacher(mary)$, $hasFather(mary, john)$, $HappyAnc(john)$

Semantics of a Description Logics knowledge base

The semantics is given by specifying when an interpretation \mathcal{I} satisfies an assertion:

- $C_1 \sqsubseteq C_2$ is satisfied by \mathcal{I} if $C_1^{\mathcal{I}} \subseteq C_2^{\mathcal{I}}$.
- $R_1 \sqsubseteq R_2$ is satisfied by \mathcal{I} if $R_1^{\mathcal{I}} \subseteq R_2^{\mathcal{I}}$.
- A property assertion (**prop** P) is satisfied by \mathcal{I} if $P^{\mathcal{I}}$ is a relation that has the property **prop**.
(Note: domain and range assertions can be expressed by means of concept inclusion assertions.)
- $A(c)$ is satisfied by \mathcal{I} if $c^{\mathcal{I}} \in A^{\mathcal{I}}$.
- $P(c_1, c_2)$ is satisfied by \mathcal{I} if $(c_1^{\mathcal{I}}, c_2^{\mathcal{I}}) \in P^{\mathcal{I}}$.

We adopt **the unique name assumption**, i.e., $c_1^{\mathcal{I}} \neq c_2^{\mathcal{I}}$, for $c_1 \neq c_2$.

Models of a Description Logics ontology

Model of a DL knowledge base

An interpretation \mathcal{I} is a **model** of $\mathcal{O} = \langle \mathcal{T}, \mathcal{A} \rangle$ if it satisfies all assertions in \mathcal{T} and all assertions in \mathcal{A} .

\mathcal{O} is said to be **satisfiable** if it admits a model.

The fundamental reasoning service from which all other ones can be easily derived is ...

Logical implication

\mathcal{O} **logically implies** and assertion α , written $\mathcal{O} \models \alpha$, if α is satisfied by all models of \mathcal{O} .

TBox reasoning

- **Concept Satisfiability:** C is satisfiable wrt \mathcal{T} , if there is a model \mathcal{I} of \mathcal{T} such that $C^{\mathcal{I}}$ is not empty, i.e., $\mathcal{T} \not\models C \equiv \perp$.
- **Subsumption:** C_1 is subsumed by C_2 wrt \mathcal{T} , if for every model \mathcal{I} of \mathcal{T} we have $C_1^{\mathcal{I}} \subseteq C_2^{\mathcal{I}}$, i.e., $\mathcal{T} \models C_1 \sqsubseteq C_2$.
- **Equivalence:** C_1 and C_2 are equivalent wrt \mathcal{T} if for every model \mathcal{I} of \mathcal{T} we have $C_1^{\mathcal{I}} = C_2^{\mathcal{I}}$, i.e., $\mathcal{T} \models C_1 \equiv C_2$.
- **Disjointness:** C_1 and C_2 are disjoint wrt \mathcal{T} if for every model \mathcal{I} of \mathcal{T} we have $C_1^{\mathcal{I}} \cap C_2^{\mathcal{I}} = \emptyset$, i.e., $\mathcal{T} \models C_1 \sqcap C_2 \equiv \perp$.
- **Functionality implication:** A functionality assertion (**funct** R) is logically implied by \mathcal{T} if for every model \mathcal{I} of \mathcal{T} , we have that $(o, o_1) \in R^{\mathcal{I}}$ and $(o, o_2) \in R^{\mathcal{I}}$ implies $o_1 = o_2$, i.e., $\mathcal{T} \models (\text{funct } R)$.

Analogous definitions hold for role satisfiability, subsumption, equivalence, and disjointness.

Reasoning over an ontology

- **Ontology Satisfiability:** Verify whether an ontology \mathcal{O} is satisfiable, i.e., whether \mathcal{O} admits at least one model.
- **Concept Instance Checking:** Verify whether an individual c is an instance of a concept C in \mathcal{O} , i.e., whether $\mathcal{O} \models C(c)$.
- **Role Instance Checking:** Verify whether a pair (c_1, c_2) of individuals is an instance of a role R in \mathcal{O} , i.e., whether $\mathcal{O} \models R(c_1, c_2)$.
- **Query Answering:** see later ...

Reasoning in Description Logics – Example

- Inclusion assertions on concepts:

$$\begin{aligned} \text{Father} &\equiv \text{Human} \sqcap \text{Male} \sqcap \exists \text{hasChild} \\ \text{HappyFather} &\sqsubseteq \text{Father} \sqcap \forall \text{hasChild}. (\text{Doctor} \sqcup \text{Lawyer} \sqcup \text{HappyPerson}) \\ \text{HappyAnc} &\sqsubseteq \forall \text{descendant}. \text{HappyFather} \\ \text{Teacher} &\sqsubseteq \neg \text{Doctor} \sqcap \neg \text{Lawyer} \end{aligned}$$

- Inclusion assertions on roles:

$$\text{hasChild} \sqsubseteq \text{descendant} \qquad \text{hasFather} \sqsubseteq \text{hasChild}^-$$

- Property assertions on roles:

(**transitive** descendant), (**reflexive** descendant), (**functional** hasFather)

The above TBox logically implies: **HappyAncestor** \sqsubseteq **Father**.

- Membership assertions:

Teacher(mary), **hasFather**(mary, john), **HappyAnc**(john)

The above TBox and ABox logically imply: **HappyPerson**(mary)

Complexity of reasoning over DL ontologies

Reasoning over DL ontologies is much more complex than reasoning over concept expressions:

- **Bad news:**
 - without restrictions on the form of TBox assertions, reasoning over DL ontologies is already **EXPTIME-hard**, even for very simple DLs (see, e.g., [Don03]).
- **Good news:**
 - We can add a lot of expressivity (i.e., essentially all DL constructs seen so far), while still staying within the EXPTIME upper bound.
 - There are DL reasoners that perform reasonably well in practice for such DLs (e.g, Racer, Pellet, Fact++, ...) [MH03].

Outline

- 3 A gentle introduction to Description Logics
- 4 DLs as a formal language to specify ontologies
 - DLs to specify ontologies
 - DLs vs. OWL
 - DLs vs. UML Class Diagrams
- 5 Queries in Description Logics
- 6 The *DL-Lite* family of tractable Description Logics

Relationship between DLs and ontology formalisms

- Description Logics are nowadays advocated to provide the foundations for ontology languages.
- Different versions of the **Ontology Web Language (OWL)** have been defined as syntactic variants of certain Description Logics.
- DLs are also ideally suited to capture the fundamental features of conceptual modeling formalisms used in information systems design:
 - **Entity-Relationship diagrams**, used in database conceptual modeling
 - **UML Class Diagrams**, used in the design phase of software applications

We briefly overview these correspondences, highlighting essential DL constructs, also in light of the tradeoff between expressive power and computational complexity of reasoning.

DLs vs. OWL

The Ontology Web Language (OWL) comes in different variants:

- **OWL-Lite** is a variant of the DL $SHIN(D)$, where:
 - \mathcal{S} stands for \mathcal{ALC} extended with **transitive roles**
 - \mathcal{H} stands for **role hierarchies** (i.e., role inclusion assertions)
 - \mathcal{I} stands for **inverse roles**
 - \mathcal{N} stands for (unqualified) **number restrictions**
 - (D) stand for **data types**, which are necessary in any practical knowledge representation language
- **OWL-DL** is a variant of $SHOIQ(D)$, where:
 - \mathcal{O} stands for **nominals**, which means the possibility of using individuals in the TBox (i.e., the intensional part of the ontology)
 - \mathcal{Q} stands for qualified number restrictions

DL constructs vs. OWL constructs

OWL constructor	DL constructor	Example
intersectionOf	$C_1 \sqcap \dots \sqcap C_n$	Human \sqcap Male
unionOf	$C_1 \sqcup \dots \sqcup C_n$	Doctor \sqcup Lawyer
complementOf	$\neg C$	\neg Male
oneOf	$\{a_1\} \sqcup \dots \sqcup \{a_n\}$	{john} \sqcup {mary}
allValuesFrom	$\forall P.C$	\forall hasChild.Doctor
someValuesFrom	$\exists P.C$	\exists hasChild.Lawyer
maxCardinality	$(\leq n P)$	$(\leq 1$ hasChild)
minCardinality	$(\geq n P)$	$(\geq 2$ hasChild)

DL axioms vs. OWL axioms

OWL axiom	DL syntax	Example
subClassOf	$C_1 \sqsubseteq C_2$	Human \sqsubseteq Animal \sqcap Biped
equivalentClass	$C_1 \equiv C_2$	Man \equiv Human \sqcap Male
disjointWith	$C_1 \sqsubseteq \neg C_2$	Man $\sqsubseteq \neg$ Female
sameIndividualAs	$\{a_1\} \equiv \{a_2\}$	{presBush} \equiv {G.W.Bush}
differentFrom	$\{a_1\} \sqsubseteq \neg\{a_2\}$	{john} $\sqsubseteq \neg\{peter\}$
subPropertyOf	$P_1 \sqsubseteq P_2$	hasDaughter \sqsubseteq hasChild
equivalentProperty	$P_1 \equiv P_2$	hasCost \equiv hasPrice
inverseOf	$P_1 \equiv P_2^-$	hasChild \equiv hasParent ⁻
transitiveProperty	$P^+ \sqsubseteq P$	ancestor ⁺ \sqsubseteq ancestor
functionalProperty	$\top \sqsubseteq (\leq 1 P)$	$\top \sqsubseteq (\leq 1$ hasFather)
inverseFunctionalProperty	$\top \sqsubseteq (\leq 1 P^-)$	$\top \sqsubseteq (\leq 1$ hasSSN ⁻)

DLs vs. UML Class Diagrams

There is a tight correspondence between variants of DLs and UML Class Diagrams [BCDG05].

- We can devise two transformations:
 - one that associates to each UML Class Diagram \mathcal{D} a DL TBox $\mathcal{T}_{\mathcal{D}}$.
 - one that associates to each DL TBox \mathcal{T} a UML Class Diagram $\mathcal{D}_{\mathcal{T}}$.
- The transformations are not model-preserving, but are based on a correspondence between instantiations of the Class Diagram and models of the associated ontology.
- The transformations are **satisfiability-preserving**, i.e., a class C is consistent in \mathcal{D} iff the corresponding concept is satisfiable in \mathcal{T} .

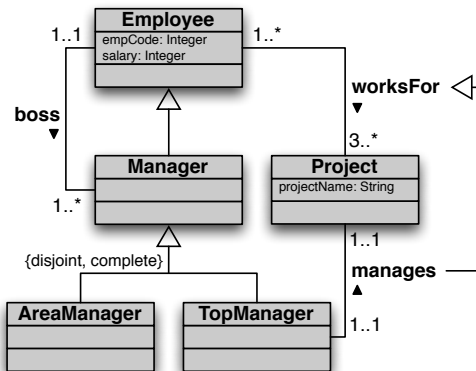
Encoding UML Class Diagrams in DLs

The ideas behind the encoding of a UML Class Diagram \mathcal{D} in terms of a DL TBox $\mathcal{T}_{\mathcal{D}}$ are quite natural:

- Each class is represented by an atomic concept.
- Each attribute is represented by a role.
- Each binary association is represented by a role.
- Each non-binary association is reified, i.e., represented as a concept connected to its components by roles.
- Each part of the diagram is encoded by suitable assertions.

We illustrate the encoding by means of an example.

Encoding UML Class Diagrams in DLs – Example



- Manager ⊆ Employee
- AreaManager ⊆ Manager
- TopManager ⊆ Manager
- Manager ⊆ AreaManager ⊔ TopManager
- AreaManager ⊆ ¬TopManager
- Employee ⊆ ∃salary
- ∃salary⁻ ⊆ Integer
- ∃worksFor ⊆ Employee
- ∃worksFor⁻ ⊆ Project
- Employee ⊆ ∃worksFor
- Project ⊆ (≥ 3 worksFor⁻)

(**func**t manages)
 (**func**t manages⁻)

- manages ⊆ worksFor
- ...

Note: Domain and range of associations are expressed by means of concept inclusions.



Reasoning on UML Class Diagrams using DLs

- The two encodings show that DL TBoxes and UML Class Diagrams essentially have the **same expressive power**.
- Hence, reasoning over UML Class Diagrams has the same complexity as reasoning over ontologies in expressive DLs, i.e., EXPTIME -complete.
- The high complexity is caused by:
 - ① the possibility to use disjunction (covering constraints)
 - ② the interaction between role inclusions and functionality constraints (maximum 1 cardinality)

Without (1) and restricting (2), reasoning becomes simpler [ACK⁺07]:

- NLOGSPACE -complete in combined complexity
- in LOGSPACE in data complexity (see later)

Efficient reasoning on UML Class Diagrams

We are interested in using UML Class Diagrams to specify ontologies in the context of Ontology-Based Data Access.

Questions

- Which is the right combination of constructs to allow in UML Class Diagrams to be used for OBDA?
- Are there techniques for query answering in this case that can be derived from Description Logics?
- Can query answering be done efficiently in the size of the data?
- If yes, can we leverage relational database technology for query answering?

Outline

- 3 A gentle introduction to Description Logics
- 4 DLs as a formal language to specify ontologies
- 5 Queries in Description Logics**
 - Queries over Description Logics ontologies
 - Certain answers
 - Complexity of query answering
- 6 The *DL-Lite* family of tractable Description Logics

Queries over Description Logics ontologies

We need more complex queries than simple concept (or role) expressions.

A **conjunctive query** $q(\vec{x})$ over an ontology $\mathcal{O} = \langle \mathcal{T}, \mathcal{A} \rangle$ has the form:

$$q(\vec{x}) \leftarrow \exists \vec{y}. \text{conj}(\vec{x}, \vec{y})$$

where:

- \vec{x} is a tuple of so-called **distinguished** variables.
The number of variables in \vec{x} is called the **arity** of q .
- \vec{y} is a tuple of so-called **non-distinguished** variables,
- $q(\vec{x})$ is called the **head** of q .
- $\text{conj}(\vec{x}, \vec{y})$, called the **body** of q , is a conjunction of atoms, where each atom:
 - has as predicate symbol an atomic concept or role of \mathcal{T} ,
 - may use variables in \vec{x} and \vec{y} ,
 - may use constants that are individuals of \mathcal{A} .

Queries over Description Logics ontologies (cont'd)

Note: we may also use for CQs a simplified notation

$$q(\vec{x}) \leftarrow \text{body}(\vec{x}, \vec{y})$$

where $\text{body}(\vec{x}, \vec{y})$ is a sequence constituted by the atoms in $\text{conj}(\vec{x}, \vec{y})$.

Example of conjunctive query

$$q(x, y) \leftarrow \exists p. \text{Employee}(x) \wedge \text{Employee}(y) \wedge \text{Project}(p) \wedge \\ \text{boss}(x, y) \wedge \text{worksFor}(x, p) \wedge \text{worksFor}(y, p)$$

In simplified notation:

$$q(x, y) \leftarrow \text{Employee}(x), \text{Employee}(y), \text{Project}(p), \\ \text{boss}(x, y), \text{worksFor}(x, p), \text{worksFor}(y, p)$$

Note: a CQ corresponds to a select-project-join SQL query.

Certain answers to a query

Let $\mathcal{O} = \langle \mathcal{T}, \mathcal{A} \rangle$ be an ontology, \mathcal{I} an interpretation for \mathcal{O} , and $q(\vec{x}) \leftarrow \exists \vec{y}. conj(\vec{x}, \vec{y})$ a CQ.

The **answer** to $q(\vec{x})$ over \mathcal{I} , denoted $q^{\mathcal{I}}, \dots$

is the set of **tuples \vec{c} of constants of \mathcal{A}** such that the formula $\exists \vec{y}. conj(\vec{c}, \vec{y})$ evaluates to true in \mathcal{I} .

We are interested in finding those answers that hold in all models of an ontology.

The **certain answers** to $q(\vec{x})$ over $\mathcal{O} = \langle \mathcal{T}, \mathcal{A} \rangle$, denoted $cert(q, \mathcal{O}), \dots$

are the **tuples \vec{c} of constants of \mathcal{A}** such that $\vec{c} \in q^{\mathcal{I}}$, for **every model \mathcal{I}** of \mathcal{O} .

Query answering over ontologies

Query answering over an ontology \mathcal{O}


Is the problem of **computing the certain answers** to a query over \mathcal{O} .

Computing certain answers is a form of **logical implication**:

$$\vec{c} \in \text{cert}(q, \mathcal{O}) \quad \text{iff} \quad \mathcal{O} \models q(\vec{c})$$

Note: **instance checking** is a special case of query answering: it amounts to answering the boolean query $q() \leftarrow A(c)$ (resp., $q() \leftarrow P(c_1, c_2)$) over \mathcal{O} (in this case \vec{c} is the empty tuple).

Query answering over ontologies – Example



TBox \mathcal{T} :

$\exists \text{hasFather}$	\sqsubseteq	Person
$\exists \text{hasFather}^-$	\sqsubseteq	Person
Person	\sqsubseteq	$\exists \text{hasFather}$

ABox \mathcal{A} : Person(john), Person(nick), Person(toni)
hasFather(john,nick), hasFather(nick,toni)

Queries:

$q_1(x, y) \leftarrow \text{hasFather}(x, y)$

$q_2(x) \leftarrow \exists y. \text{hasFather}(x, y)$

$q_3(x) \leftarrow \exists y_1, y_2, y_3. \text{hasFather}(x, y_1) \wedge \text{hasFather}(y_1, y_2) \wedge \text{hasFather}(y_2, y_3)$

$q_4(x, y_3) \leftarrow \exists y_1, y_2. \text{hasFather}(x, y_1) \wedge \text{hasFather}(y_1, y_2) \wedge \text{hasFather}(y_2, y_3)$

Certain answers:

$\text{cert}(q_1, \langle \mathcal{T}, \mathcal{A} \rangle)$	$= \{ (\text{john}, \text{nick}), (\text{nick}, \text{toni}) \}$
$\text{cert}(q_2, \langle \mathcal{T}, \mathcal{A} \rangle)$	$= \{ \text{john}, \text{nick}, \text{toni} \}$
$\text{cert}(q_3, \langle \mathcal{T}, \mathcal{A} \rangle)$	$= \{ \text{john}, \text{nick}, \text{toni} \}$
$\text{cert}(q_4, \langle \mathcal{T}, \mathcal{A} \rangle)$	$= \{ \}$

Unions of conjunctive queries

We consider also unions of CQs.

A **union of conjunctive queries** (UCQ) has the form:

$$q(\vec{x}) \leftarrow \exists \vec{y}_1. \text{conj}(\vec{x}, \vec{y}_1) \vee \dots \vee \exists \vec{y}_k. \text{conj}(\vec{x}, \vec{y}_k)$$

where each $\vec{y}_i. \text{conj}(\vec{x}, \vec{y}_i)$ is the body of a CQ.

Example

$$q(x) \leftarrow (\text{Manager}(x) \wedge \text{worksFor}(x, \text{tones})) \vee \\ (\exists y. \text{boss}(x, y) \wedge \text{worksFor}(y, \text{tones}))$$

The (certain) answers to a UCQ are defined analogously to those for CQs.

Data and combined complexity

When measuring the complexity of answering a query $q(\vec{x})$ over an ontology $\mathcal{O} = \langle \mathcal{T}, \mathcal{A} \rangle$, various parameters are of importance.

Depending on which parameters we consider, we get different complexity measures:

- **Data complexity**: TBox and query are considered fixed, and only the size of the ABox (i.e., the data) matters.
- **Query complexity**: TBox and ABox are considered fixed, and only the size of the query matters.
- **Schema complexity**: ABox and query are considered fixed, and only the size of the TBox (i.e., the schema) matters.
- **Combined complexity**: no parameter is considered fixed.

In the OBDA setting, **the size of the data largely dominates** the size of the conceptual layer (and of the query).

↪ **Data complexity** is the relevant complexity measure.

Complexity of query answering in DLs

Answering (U)CQs over DL ontologies has been studied extensively:

- **Combined complexity:**
 - NP-complete for plain databases (i.e., with an empty TBox)
 - EXPTIME-complete for \mathcal{ALC} [CDGL98, Lut07]
 - 2EXPTIME-complete for very expressive DLs (with inverse roles) [CDGL98, Lut07]
- **Data complexity:**
 - in LOGSPACE for plain databases
 - coNP-hard with disjunction in the TBox [DLNS94, CDGL⁺06b]
 - coNP-complete for very expressive DLs [LR98, OCE06, GHLS07]

Questions

- Can we find interesting families of DLs for which the query answering problem can be solved efficiently?
- If yes, can we leverage relational database technology for query answering?

Outline

- 3 A gentle introduction to Description Logics
- 4 DLs as a formal language to specify ontologies
- 5 Queries in Description Logics
- 6 The *DL-Lite* family of tractable Description Logics
 - The *DL-Lite* family
 - Syntax of *DL-Lite_F* and *DL-Lite_R*
 - Semantics of *DL-Lite*
 - Properties of *DL-Lite*

The *DL-Lite* family

- Is a family of DLs optimized according to the tradeoff between expressive power and data complexity of query answering.
- We present now two incomparable languages of this family, *DL-Lite_F*, *DL-Lite_R* (we use *DL-Lite* to refer to both).
- We will see that *DL-Lite* has nice computational properties:
 - PTIME in the size of the TBox (schema complexity)
 - LOGSPACE in the size of the ABox (data complexity)
 - enjoys FOL-rewritability
- We will see that *DL-Lite_F* and *DL-Lite_R* are in some sense the maximal DLs with these nice computational properties, which are lost with minimal additions of constructs.

Hence, *DL-Lite* provides a positive answer to our basic questions, and sets the foundations for Ontology-Based Data Access.

DL-Lite_F ontologies

TBox assertions:

- Concept inclusion assertions: $Cl \sqsubseteq Cr$, with:

$$\begin{array}{l}
 Cl \longrightarrow A \mid \exists Q \\
 Cr \longrightarrow A \mid \exists Q \mid \neg A \mid \neg \exists Q \\
 Q \longrightarrow P \mid P^-
 \end{array}$$

- Functionality assertions: **(funct Q)**

ABox assertions: $A(c)$, $P(c_1, c_2)$, with c_1, c_2 constants

Observations:

- Captures all the basic constructs of UML Class Diagrams and ER
- Notable exception: covering constraints in generalizations.

DL-Lite_R ontologies

TBox assertions:

- Concept inclusion assertions: $Cl \sqsubseteq Cr$, with:

$$\begin{array}{l} Cl \longrightarrow A \mid \exists Q \\ Cr \longrightarrow A \mid \exists Q \mid \neg A \mid \neg \exists Q \\ Q \longrightarrow P \mid P^- \end{array}$$

- Role inclusion assertions: $Q \sqsubseteq R$, with:

$$R \longrightarrow Q \mid \neg Q$$

ABox assertions: $A(c)$, $P(c_1, c_2)$, with c_1, c_2 constants

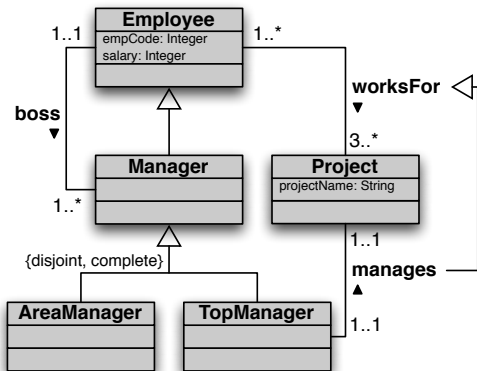
Observations:

- Drops functional restrictions in favor of ISA between roles.
- Extends (the DL fragment of) the ontology language RDFS.

Semantics of *DL-Lite*

Construct	Syntax	Example	Semantics
atomic conc.	A	Doctor	$A^I \subseteq \Delta^I$
exist. restr.	$\exists Q$	$\exists \text{child}^-$	$\{d \mid \exists e. (d, e) \in Q^I\}$
at. conc. neg.	$\neg A$	$\neg \text{Doctor}$	$\Delta^I \setminus A^I$
conc. neg.	$\neg \exists Q$	$\neg \exists \text{child}$	$\Delta^I \setminus (\exists Q)^I$
atomic role	P	child	$P^I \subseteq \Delta^I \times \Delta^I$
inverse role	P^-	child^-	$\{(o, o') \mid (o', o) \in P^I\}$
role negation	$\neg Q$	$\neg \text{manages}$	$(\Delta_O^I \times \Delta_O^I) \setminus Q^I$
conc. incl.	$C_l \sqsubseteq C_r$	$\text{Father} \sqsubseteq \exists \text{child}$	$C_l^I \subseteq C_r^I$
role incl.	$Q \sqsubseteq R$	$\text{hasFather} \sqsubseteq \text{child}^-$	$Q^I \subseteq R^I$
funct. asser.	(funct Q)	(funct succ)	$\forall d, e, e'. (d, e) \in Q^I \wedge (d, e') \in Q^I \rightarrow e = e'$
mem. asser.	$A(c)$	$\text{Father}(\text{bob})$	$c^I \in A^I$
mem. asser.	$P(c_1, c_2)$	$\text{child}(\text{bob}, \text{ann})$	$(c_1^I, c_2^I) \in P^I$

DL-Lite – Example



- Manager ⊆ Employee
- AreaManager ⊆ Manager
- TopManager ⊆ Manager
- AreaManager ⊆ ¬TopManager
- Employee ⊆ ∃salary
- ∃salary ⊆ Integer
- ∃worksFor ⊆ Employee
- ∃worksFor ⊆ Project
- Employee ⊆ ∃worksFor
- Project ⊆ ∃worksFor
- ⋮

Additionally, in $DL-Lite_{\mathcal{F}}$: (**func** manages), (**func** manages⁻), ...
 in $DL-Lite_{\mathcal{R}}$: manages ⊆ worksFor

Note: in $DL-Lite$ we cannot capture: – completeness of the hierarchy,
 – number restrictions



Properties of *DL-Lite*

- The TBox may contain **cyclic dependencies** (which typically increase the computational complexity of reasoning).

Example: $A \sqsubseteq \exists P$, $\exists P^- \sqsubseteq A$

- We have not included in the syntax \sqcap on the right hand-side of inclusion assertions, but it can trivially be added, since

$$Cl \sqsubseteq Cr_1 \sqcap Cr_2 \quad \text{is equivalent to} \quad \begin{array}{l} Cl \sqsubseteq Cr_1 \\ Cl \sqsubseteq Cr_2 \end{array}$$

- A domain assertion on role P has the form: $\exists P \sqsubseteq A_1$
A range assertion on role P has the form: $\exists P^- \sqsubseteq A_2$

Properties of $DL\text{-Lite}_{\mathcal{F}}$

$DL\text{-Lite}_{\mathcal{F}}$ does **not** enjoy the **finite model property**.

Example

TBox \mathcal{T} : $\text{Nat} \sqsubseteq \exists \text{succ}$ $\exists \text{succ}^- \sqsubseteq \text{Nat}$
 $\text{Zero} \sqsubseteq \text{Nat} \sqcap \neg \exists \text{succ}^-$ (**funct succ**⁻)

ABox \mathcal{A} : **Zero(0)**

$\mathcal{O} = \langle \mathcal{T}, \mathcal{A} \rangle$ admits only infinite models.

Hence, it is satisfiable, but **not finitely satisfiable**.

Hence, reasoning w.r.t. arbitrary models is different from reasoning w.r.t. finite models only.

Properties of $DL-Lite_{\mathcal{R}}$

- The TBox may contain **cyclic dependencies**.
- $DL-Lite_{\mathcal{R}}$ **does enjoy the finite model property**. Hence, reasoning w.r.t. finite models is the same as reasoning w.r.t. arbitrary models.
- With role inclusion assertions, we can simulate **qualified existential quantification** in the rhs of an inclusion assertion $A_1 \sqsubseteq \exists Q.A_2$.

To do so, we introduce a new role Q_{A_2} and:

- the role inclusion assertion $Q_{A_2} \sqsubseteq Q$
- the concept inclusion assertions:

$$\begin{array}{l} A_1 \sqsubseteq \exists Q_{A_2} \\ \exists Q_{A_2}^- \sqsubseteq A_2 \end{array}$$

In this way, we can consider $\exists Q.A$ in the right-hand side of an inclusion assertion as an abbreviation.

Complexity results for *DL-Lite*

- 1 We have seen that $DL-Lite_{\mathcal{F}}$ and $DL-Lite_{\mathcal{R}}$ can capture the essential features of prominent conceptual modeling formalisms.
- 2 In the next part, we will analyze reasoning in *DL-Lite*, and establish the following characterization of its computational properties:
 - Ontology satisfiability is **polynomial** in the size of **TBox** and **ABox**.
 - Query answering is:
 - **P**TIME in the size of the **TBox**.
 - **LOGSPACE** in the size of the **ABox**, and **FOL-rewritable**, which means that we can leverage for it relational database technology.
- 3 We will also see that *DL-Lite* is essentially the maximal DL enjoying these nice computational properties.

From (1), (2), and (3) we get the following claim:

DL-Lite is the representation formalism that is best suited to underly Ontology-Based Data Management systems.

Query answering in $DL-Lite_{\mathcal{R}}$: Query reformulation (cont'd)

Conversely, for the query

$$q(x) \leftarrow \text{teaches}(x, \text{databases})$$

$\text{Professor} \sqsubseteq \exists \text{teaches}$

as a logic rule: $\text{teaches}(z_1, z_2) \leftarrow \text{Professor}(z_1)$

$\text{teaches}(x, \text{databases})$ does not unify with $\text{teaches}(z_1, z_2)$, since the **existentially quantified variable z_2** in the head of the rule **does not unify** with the constant databases .

In this case the PI **does not apply** to the atom $\text{teaches}(x, \text{databases})$.

The same holds for the following query, where y is **distinguished**

$$q(x, y) \leftarrow \text{teaches}(x, y)$$



Query answering in $DL-Lite_{\mathcal{R}}$: Query reformulation (cont'd)

An analogous behavior with join variables

$$q(x) \leftarrow \text{teaches}(x, y), \text{Course}(y)$$

$$\text{Professor} \sqsubseteq \exists \text{teaches}$$

as a logic rule: $\text{teaches}(z_1, z_2) \leftarrow \text{Professor}(z_1)$

The PI above does not apply to the atom $\text{teaches}(x, y)$.

Conversely, the PI

$$\exists \text{teaches}^- \sqsubseteq \text{Course}$$

as a logic rule: $\text{Course}(z_2) \leftarrow \text{teaches}(z_1, z_2)$

applies to the atom $\text{Course}(y)$.

We add to the perfect reformulation the query

$$q(x) \leftarrow \text{teaches}(x, y), \text{teaches}(z, y)$$

Query answering in $DL-Lite_{\mathcal{R}}$: Query reformulation (cont'd)

We now have the query

$$q(x) \leftarrow \text{teaches}(x, y), \text{teaches}(z, y)$$

The PI

Professor $\sqsubseteq \exists \text{teaches}$

as a logic rule: $\text{teaches}(z_1, z_2) \leftarrow \text{Professor}(z_1)$

does not apply to $\text{teaches}(x, y)$ nor $\text{teaches}(z, y)$, since y is in join.

However, we can transform the above query by **unifying** the atoms $\text{teaches}(x, y)$, $\text{teaches}(z_1, y)$. This rewriting step is called **reduce**, and produces the following query

$$q(x) \leftarrow \text{teaches}(x, y)$$

We can now apply the PI above, and add to the reformulation the query

$$q(x) \leftarrow \text{Professor}(x)$$

Query answering in $DL-Lite_{\mathcal{R}}$: Query reformulation (cont'd)

Reformulate the CQ q into a set of queries: apply to q in all possible ways the Pls in the TBox \mathcal{T} :

$$\begin{array}{lll}
 A_1 \sqsubseteq A_2 & \dots, A_2(x), \dots & \rightsquigarrow \dots, A_1(x), \dots \\
 \exists P \sqsubseteq A & \dots, A(x), \dots & \rightsquigarrow \dots, P(x, -), \dots \\
 \exists P^- \sqsubseteq A & \dots, A(x), \dots & \rightsquigarrow \dots, P(-, x), \dots \\
 A \sqsubseteq \exists P & \dots, P(x, -), \dots & \rightsquigarrow \dots, A(x), \dots \\
 A \sqsubseteq \exists P^- & \dots, P(-, x), \dots & \rightsquigarrow \dots, A(x), \dots \\
 \exists P_1 \sqsubseteq \exists P_2 & \dots, P_2(x, -), \dots & \rightsquigarrow \dots, P_1(x, -), \dots
 \end{array}$$

($-$ denotes an **unbound** variable, i.e., a variable that appears only once)

This corresponds to exploiting ISAs, role typing, and mandatory participation to obtain new queries that could contribute to the answer.

Unifying atoms can make applicable rules that could not be applied otherwise.

Query answering in $DL\text{-Lite}_{\mathcal{R}}$: ABox storage

ABox \mathcal{A} stored as a **relational database** in a standard RDBMS as follows:

- For each **atomic concept** A used in the ABox:
 - define a **unary relational table** tab_A
 - populate tab_A with each $\langle d \rangle$ such that $A(c) \in \mathcal{A}$
- For each **atomic role** P used in the ABox,
 - define a **binary relational table** tab_P
 - populate tab_P with each $\langle a, b \rangle$ such that $P(c_1, c_2) \in \mathcal{A}$

We denote with **DB**(\mathcal{A}) the database obtained as above.

Query answering in $DL\text{-Lite}_{\mathcal{R}}$: Query evaluation

Let $r_{q,\mathcal{T}}$ be the UCQ returned by the algorithm PerfectRef (q, \mathcal{T})

- We denote by $\mathbf{SQL}(r_{q,\mathcal{T}})$ the encoding of $r_{q,\mathcal{T}}$ into an SQL query over $DB(\mathcal{A})$.
- We indicate with $\mathbf{Eval}(\mathbf{SQL}(r_{q,\mathcal{T}}), DB(\mathcal{A}))$ the evaluation of $\mathbf{SQL}(r_{q,\mathcal{T}})$ over $DB(\mathcal{A})$.

Query answering in $DL-Lite_{\mathcal{R}}$

Theorem

Let \mathcal{T} be a $DL-Lite_{\mathcal{R}}$ TBox, \mathcal{T}_P the set of Pls in \mathcal{T} , q a CQ over \mathcal{T} , and let $r_{q,\mathcal{T}} = \text{PerfectRef}(q, \mathcal{T}_P)$. Then, for each ABox \mathcal{A} such that $\langle \mathcal{T}, \mathcal{A} \rangle$ is satisfiable, we have that $\text{cert}(q, \langle \mathcal{T}, \mathcal{A} \rangle) = \text{Eval}(\text{SQL}(r_{q,\mathcal{T}}), \text{DB}(\mathcal{A}))$.

In other words, query answering over a satisfiable $DL-Lite_{\mathcal{R}}$ ontology is FOL-rewritable.

Notice that we did not mention NIs of \mathcal{T} in the theorem above. Indeed, when the ontology is satisfiable, we can ignore NIs and answer queries as NIs were not specified in \mathcal{T} .

Query answering in $DL\text{-Lite}_{\mathcal{R}}$: Example

TBox: Professor $\sqsubseteq \exists \text{teaches}$
 $\exists \text{teaches}^{-} \sqsubseteq \text{Course}$

Query: $q(x) \leftarrow \text{teaches}(x, y), \text{Course}(y)$

Perfect Reformulation: $q(x) \leftarrow \text{teaches}(x, y), \text{Course}(y)$
 $q(x) \leftarrow \text{teaches}(x, y), \text{teaches}(-, y)$
 $q(x) \leftarrow \text{teaches}(x, -)$
 $q(x) \leftarrow \text{Professor}(x)$

ABox: $\text{teaches}(\text{John}, \text{databases})$
 $\text{Professor}(\text{Mary})$

It is easy to see that $\text{Eval}(\text{SQL}(r_{q, \mathcal{T}}), \text{DB}(\mathcal{A}))$ in this case produces the set $\{\text{John}, \text{Mary}\}$.

Query answering in $DL\text{-Lite}_{\mathcal{R}}$: An interesting case

TBox: $\text{Person} \sqsubseteq \exists \text{hasFather}$ ABox: $\text{Person}(\text{Mary})$
 $\exists \text{hasFather}^- \sqsubseteq \text{Person}$

Query: $q(x) \leftarrow \text{Person}(x), \text{hasFather}(x, y_1), \text{hasFather}(y_1, y_2), \text{hasFather}(y_2, y_3)$

$q(x) \leftarrow \text{Person}(x), \text{hasFather}(x, y_1), \text{hasFather}(y_1, y_2), \text{hasFather}(y_2, -)$

\Downarrow Apply $\text{Person} \sqsubseteq \exists \text{hasFather}$ to the atom $\text{hasFather}(y_2, -)$

$q(x) \leftarrow \text{Person}(x), \text{hasFather}(x, y_1), \text{hasFather}(y_1, y_2), \text{Person}(y_2)$

\Downarrow Apply $\exists \text{hasFather}^- \sqsubseteq \text{Person}$ to the atom $\text{Person}(y_2)$

$q(x) \leftarrow \text{Person}(x), \text{hasFather}(x, y_1), \text{hasFather}(y_1, y_2), \text{hasFather}(-, y_2)$

\Downarrow Unify atoms $\text{hasFather}(y_1, y_2)$ and $\text{hasFather}(-, y_2)$

$q(x) \leftarrow \text{Person}(x), \text{hasFather}(x, y_1), \text{hasFather}(y_1, y_2)$

\Downarrow

\dots

$q(x) \leftarrow \text{Person}(x), \text{hasFather}(x, -)$

\Downarrow Apply $\text{Person} \sqsubseteq \exists \text{hasFather}$ to the atom $\text{hasFather}(x, -)$

$q(x) \leftarrow \text{Person}(x)$

Query answering in $DL\text{-Lite}_{\mathcal{F}}$

If we limit our attention to PIs, we can say that $DL\text{-Lite}_{\mathcal{F}}$ ontologies are $DL\text{-Lite}_{\mathcal{R}}$ ontologies of a special kind (i.e., with no PIs between roles).

As for **NIs and functionality assertions**, it is possible to show that they can be disregarded in query answering over satisfiable $DL\text{-Lite}_{\mathcal{F}}$ ontologies.

The following result is therefore straightforward.

Theorem

Let \mathcal{T} be a $DL\text{-Lite}_{\mathcal{F}}$ TBox, \mathcal{T}_P the set of PIs in \mathcal{T} , q a CQ over \mathcal{T} , and let $r_{q,\mathcal{T}} = \text{PerfectRef}(q, \mathcal{T}_P)$. Then, for each ABox \mathcal{A} such that $\langle \mathcal{T}, \mathcal{A} \rangle$ is satisfiable, we have that $\text{cert}(q, \langle \mathcal{T}, \mathcal{A} \rangle) = \text{Eval}(\text{SQL}(r_{q,\mathcal{T}}), \text{DB}(\mathcal{A}))$.

In other words, query answering over a satisfiable $DL\text{-Lite}_{\mathcal{F}}$ ontology is FOL-rewritable.

DL-Lite_R ontologies

Unsatisfiability in *DL-Lite_R* ontologies can be however caused by NIs

Example: **TBox** \mathcal{T} : Professor $\sqsubseteq \neg$ Student
 \exists teaches \sqsubseteq Professor

ABox \mathcal{A} : teaches(John, databases)
 Student(John)

In what follows we provide a mechanism to establish, in an efficient way, whether a *DL-Lite_R* ontology is satisfiable.

Checking satisfiability of $DL\text{-Lite}_{\mathcal{R}}$ ontologies

Satisfiability of a $DL\text{-Lite}_{\mathcal{R}}$ ontology $\mathcal{O} = \langle \mathcal{T}, \mathcal{A} \rangle$ is reduced to evaluating a FOL-query (in fact a UCQ) over $DB(\mathcal{A})$

We proceed as follows:

- ① Let \mathcal{T}_P the set of PIs in \mathcal{T}
- ② For each NI between concepts (resp. roles) in \mathcal{T} , we ask $\langle \mathcal{T}_P, \mathcal{A} \rangle$ if there exists some individual (resp. pair of individuals) that contradicts N , i.e., we pose over $\langle \mathcal{T}_P, \mathcal{A} \rangle$ a boolean CQ q_N such that $\langle \mathcal{T}_P, \mathcal{A} \rangle \models q_N$ iff $\langle \mathcal{T}_P \cup \{N\}, \mathcal{A} \rangle$ is unsatisfiable
- ③ We exploit PerfectRef to verify if $\langle \mathcal{T}_P, \mathcal{A} \rangle \models q_N$, i.e., we compute $\text{PerfectRef}(q_N, \mathcal{T}_P)$, and evaluate it (in fact its SQL encoding) over $DB(\mathcal{A})$.

Example

PIs \mathcal{T}_P : $\exists \text{teaches} \sqsubseteq \text{Professor}$

NI N : $\text{Professor} \sqsubseteq \neg \text{Student}$

Query q_N : $q() \leftarrow \text{Student}(x), \text{Professor}(x)$

Perfect Reformulation: $q() \leftarrow \text{Student}(x), \text{Professor}(x)$
 $q() \leftarrow \text{Student}(x), \text{teaches}(x, -)$

ABox \mathcal{A} : $\text{teaches}(\text{John}, \text{databases})$
 $\text{Student}(\text{John})$

It is easy to see that $\langle \mathcal{T}_P, \mathcal{A} \rangle \models q_N$, and that
 $\langle \mathcal{T}_P \cup \{\text{Professor} \sqsubseteq \neg \text{Student}\}, \mathcal{A} \rangle$ is unsatisfiable.

Queries for NIs

For each NI in \mathcal{T} we compute a boolean CQ according to the following rules:

$A_1 \sqsubseteq \neg A_2$	\rightsquigarrow	$q() \leftarrow A_1(x), A_2(x)$
$\exists P \sqsubseteq \neg A$ or $A \sqsubseteq \neg \exists P$	\rightsquigarrow	$q() \leftarrow P(x, y), A(x)$
$\exists P^- \sqsubseteq \neg A$ or $A \sqsubseteq \neg \exists P^-$	\rightsquigarrow	$q() \leftarrow P(y, x), A(x)$
$\exists P_1 \sqsubseteq \neg \exists P_2$	\rightsquigarrow	$q() \leftarrow P_1(x, y), P_2(x, z)$
$\exists P_1 \sqsubseteq \neg \exists P_2^-$	\rightsquigarrow	$q() \leftarrow P_1(x, y), P_2(z, x)$
$\exists P_1^- \sqsubseteq \neg \exists P_2$	\rightsquigarrow	$q() \leftarrow P_1(x, y), P_2(y, z)$
$\exists P_1^- \sqsubseteq \neg \exists P_2^-$	\rightsquigarrow	$q() \leftarrow P_1(x, y), P_2(z, y)$
$P_1 \sqsubseteq \neg P_2$ or $P_1^- \sqsubseteq \neg P_2^-$	\rightsquigarrow	$q() \leftarrow P_1(x, y), P_2(x, y)$
$P_1^- \sqsubseteq \neg P_2$ or $P_1 \sqsubseteq \neg P_2^-$	\rightsquigarrow	$q() \leftarrow P_1(x, y), P_2(y, x)$

Given a NI $N \in \mathcal{T}$, we denote with q_N the corresponding CQ.

$DL\text{-Lite}_{\mathcal{R}}$: From satisfiability to query answering

Lemma [separation]

Let $\mathcal{O} = \langle \mathcal{T}, \mathcal{A} \rangle$ be a $DL\text{-Lite}_{\mathcal{R}}$ ontology, and \mathcal{T}_P the set of PIs in \mathcal{T} . Then, \mathcal{O} is unsatisfiable iff there exists a NI $N \in \mathcal{T}$ such that $\langle \mathcal{T}_P, \mathcal{A} \rangle \models q_N$.

The lemma relies on the properties that NIs do not interact with each other, and interaction between NIs and PIs can be managed through PerfectRef. Notably, each NI can be processed individually.

$DL-Lite_{\mathcal{R}}$: FOL-rewritability of satisfiability

From the lemma above and the theorem on query answering for satisfiable $DL-Lite_{\mathcal{R}}$, we get the following result

Theorem

Let $\mathcal{O} = \langle \mathcal{T}, \mathcal{A} \rangle$ be a $DL-Lite_{\mathcal{R}}$ ontology, and \mathcal{T}_P the set of PIs in \mathcal{T} . Then, \mathcal{O} is unsatisfiable iff there exists a NI $N \in \mathcal{T}$ such that $\text{Eval}(\text{SQL}(\text{PerfectRef}(q_N, \mathcal{T}_P)), \text{DB}(\mathcal{A}))$ returns *true*.

In other words, satisfiability of $DL-Lite_{\mathcal{R}}$ ontology can be reduced to FOL-query evaluation.

$DL-Lite_{\mathcal{F}}$ ontologies

Unsatisfiability in $DL-Lite_{\mathcal{F}}$ ontologies can be caused by **NIs and functionality assertions**.

Example: **TBox** \mathcal{T} : Professor $\sqsubseteq \neg$ Student
 \exists teaches \sqsubseteq Professor
 (**funct** teaches $^{-}$)

ABox \mathcal{A} : teaches(John, databases)
 teaches(Michael, databases)

In what follows we extend to $DL-Lite_{\mathcal{F}}$ ontologies the technique for $DL-Lite_{\mathcal{R}}$ ontology satisfiability given before.

Checking satisfiability of $DL\text{-Lite}_{\mathcal{F}}$ ontologies

Satisfiability of a $DL\text{-Lite}_{\mathcal{F}}$ ontology $\mathcal{O} = \langle \mathcal{T}, \mathcal{A} \rangle$ is reduced to evaluating a FOL-query over $DB(\mathcal{A})$.

We deal with NIs exactly as done in $DL\text{-Lite}_{\mathcal{R}}$ ontologies (indeed, limited to NIs, $DL\text{-Lite}_{\mathcal{F}}$ ontologies are $DL\text{-Lite}_{\mathcal{R}}$ ontologies of a special kind).

As for **functionality assertions** we proceed as follows:

- 1 For each functionality assertion $F \in \mathcal{T}$ we ask if there exists two pairs of individuals in \mathcal{A} that contradict F , i.e., we pose over \mathcal{A} a boolean FOL query q_F such that $\mathcal{A} \models q_F$ iff $\langle \{F\}, \mathcal{A} \rangle$ is **unsatisfiable**
- 2 To verify if $\mathcal{A} \models q_F$, we evaluate $SQL(q_F)$ over $DB(\mathcal{A})$.

Example

Functionality F : (**funct** $teaches^-$)

Query q_F : $q() \leftarrow TeachesTo(x, y), TeachesTo(z, y), x \neq z$

ABox \mathcal{A} : $teaches(John, databases)$
 $teaches(Michael, databases)$

It is easy to see that $\mathcal{A} \models q_F$, and that $\langle \{(\mathbf{funct} \mathit{teaches}^-)\}, \mathcal{A} \rangle$, is **unsatisfiable**.

Queries for functionality assertions

For each functionality assertion in \mathcal{T} we compute a FOL query according to the following rules:

$$\begin{aligned} (\mathbf{funct} P) &\rightsquigarrow q() \leftarrow P(x, y), P(x, z), y \neq z \\ (\mathbf{funct} P^-) &\rightsquigarrow q() \leftarrow P(x, y), P(z, y), x \neq z \end{aligned}$$

Given a functionality assertion $F \in \mathcal{T}$, we denote with q_F the corresponding FOL query.

$DL-Lite_{\mathcal{R}}$: From satisfiability to query answering

Lemma

Let $\mathcal{O} = \langle \mathcal{T}, \mathcal{A} \rangle$ be a $DL-Lite_{\mathcal{F}}$ ontology, and \mathcal{T}_P the set of PIs in \mathcal{T} . Then, \mathcal{O} is unsatisfiable iff one of the following condition holds

- (a) there exists a NI $N \in \mathcal{T}$ such that $\langle \mathcal{T}_P, \mathcal{A} \rangle \models q_N$
- (b) there exists a functionality assertion $F \in \mathcal{T}$ such that $\mathcal{A} \models q_F$.

The lemma relies on the properties that NIs do not interact with each other, and interaction between NIs and PIs can be managed through PerfectRef.

It also exploits the properties that **NIs and PIs do not interact with functionalities**: indeed, **no functionality assertions are contradicted in a $DL-Lite_{\mathcal{F}}$ ontology \mathcal{O} , beyond those explicitly contradicted by the ABox.**

Notably, the lemma asserts that to check ontology satisfiability, each NI and each functionality can be processed individually.

DL-Lite_R: FOL-rewritability of satisfiability

By the lemma above and the theorem on query answering for satisfiable *DL-Lite_F*, the following result follows

Theorem

Let $\mathcal{O} = \langle \mathcal{T}, \mathcal{A} \rangle$ be a *DL-Lite_F* ontology, and \mathcal{T}_P the set of Pls in \mathcal{O} . Then, \mathcal{O} is unsatisfiable iff one of the following condition holds.

- (a) there exists a NI $N \in \mathcal{T}$ such that
 $\text{Eval}(\text{SQL}(\text{PerfectRef}(q_N, \mathcal{T}_P)), \text{DB}(\mathcal{A}))$ returns *true*
- (b) **there exists a functionality assertion $F \in \mathcal{T}$ such that**
 $\text{Eval}(\text{SQL}(q_F), \text{DB}(\mathcal{A}))$ **returns *true*.**

In other words, satisfiability of a *DL-Lite_F* ontology can be reduced to FOL-query evaluation.

Complexity of query answering over satisfiable ontologies

Theorem

Query answering over both *DL-Lite_R* and *DL-Lite_F* ontologies is

- ① **NP-complete** in the size of **query and ontology** (combined comp.).
- ② **P TIME** in the size of the **ontology**.
- ③ **LOGSPACE** in the size of the **ABox** (data complexity).

Proof (sketch)

- ① We **guess** the derivation of one of the CQs of the perfect rewriting, and an assignment to its existential variables. Checking the derivation and evaluating the guessed CQ over the ABox is then polynomial in combined complexity. NP-hardness follows from combined complexity of evaluating CQs over a database.
- ② The number of CQs in the perfect reformulation is polynomial in the size of the TBox, and we can get them in **P TIME**.
- ③ Is the data complexity of evaluating FOL queries over a database.

Complexity of ontology satisfiability

Theorem

Checking satisfiability of both $DL-Lite_{\mathcal{R}}$ and $DL-Lite_{\mathcal{F}}$ ontologies is

- ① **P**TIME in the size of the **ontology** (combined complexity).
- ② **L**OGSPACE in the size of the **ABox** (data complexity).

Proof (sketch)

Follows directly from the algorithm for ontology satisfiability and the complexity of query answering over satisfiable ontologies.

Summary of results on data complexity

	Cl	Cr	\mathcal{F}	\mathcal{R}	Data complexity of query answering
1	$DL-Lite_{\mathcal{F}}$		✓	—	in LOGSPACE
2	$DL-Lite_{\mathcal{R}}$		—	✓	in LOGSPACE
3	$A \mid \exists P.A$	A	—	—	NLOGSPACE-hard
4	A	$A \mid \forall P.A$	—	—	NLOGSPACE-hard
5	A	$A \mid \exists P.A$	✓	—	NLOGSPACE-hard
6	$A \mid \exists P.A \mid A_1 \sqcap A_2$	A	—	—	PTime-hard
7	$A \mid A_1 \sqcap A_2$	$A \mid \forall P.A$	—	—	PTime-hard
8	$A \mid A_1 \sqcap A_2$	$A \mid \exists P.A$	✓	—	PTime-hard
9	$A \mid \exists P.A \mid \exists P^-.A$	$A \mid \exists P$	—	—	PTime-hard
10	A	$A \mid \exists P.A \mid \exists P^-.A$	✓	—	PTime-hard
11	$A \mid \exists P.A$	$A \mid \exists P.A$	✓	—	PTime-hard
12	$A \mid \neg A$	A	—	—	coNP-hard
13	A	$A \mid A_1 \sqcup A_2$	—	—	coNP-hard
14	$A \mid \forall P.A$	A	—	—	coNP-hard

All NLOGSPACE and PTime hardness results hold already for atomic queries.

Observations

- *DL-Lite-family* is FOL-rewritable, hence **LOGSPACE** – holds also with n -ary relations \rightsquigarrow *DLR-Lite_F* and *DLR-Lite_R*.
- **RDFS** is a subset of *DL-Lite_R* \rightsquigarrow is FOL-rewritable, hence **LOGSPACE**.
- *Horn-SHIQ* [HMS05] is **P_{TIME}-hard** even for instance checking (line 11).
- **DLP** [GHVD03] is **P_{TIME}-hard** (line 6)
- **EL** [BBL05] is **P_{TIME}-hard** (line 6).

NLOGSPACE-hard cases

Instance checking (and hence query answering) is NLOGSPACE-hard in data complexity for:

	Cl	Cr	\mathcal{F}	\mathcal{R}	Data complexity
3	$A \mid \exists P.A$	A	–	–	NLOGSPACE-hard

By reduction from reachability in directed graphs

4	A	$A \mid \forall P.A$	–	–	NLOGSPACE-hard
---	-----	----------------------	---	---	----------------

Follows from 3 by replacing $\exists P.A_1 \sqsubseteq A_2$ with $A_1 \sqsubseteq \forall P^-.A_2$

5	A	$A \mid \exists P.A$	✓	–	NLOGSPACE-hard
---	-----	----------------------	---	---	----------------

Proved by simulating in the reduction $\exists P.A_1 \sqsubseteq A_2$
via $A_1 \sqsubseteq \exists P^-.A_2$ and (funct P^-)

Path System Accessibility

Instance of Path System Accessibility: $PS = (N, E, S, t)$ with

- N a set of nodes
- $E \subseteq N \times N \times N$ an accessibility relation
- $S \subseteq N$ a set of source nodes
- $t \in N$ a terminal node

Accessibility of nodes is defined inductively:

- each $n \in S$ is accessible
- if $(n, n_1, n_2) \in E$ and n_1, n_2 are accessible, then also n is accessible

Given PS , checking whether t is accessible, is PTIME-complete.

Reduction from Path System Accessibility

Given an instance $PS = (N, E, S, t)$, we construct

- TBox \mathcal{T} consisting of the inclusion assertions

$$\exists P_1.A \sqsubseteq B_1$$

$$B_1 \sqcap B_2 \sqsubseteq A$$

$$\exists P_2.A \sqsubseteq B_2$$


$$\exists P_3.A \sqsubseteq A$$

- ABox \mathcal{A} encoding the accessibility relation using P_1 , P_2 , and P_3 , and asserting $A(s)$ for each source node $s \in S$

$$e_1 = (n, \cdot, \cdot)$$

$$e_2 = (n, s_1, s_2)$$

$$e_3 = (n, \cdot, \cdot)$$



Result:

$$\langle \mathcal{T}, \mathcal{A} \rangle \models A(t) \text{ iff } t \text{ is accessible in } PS.$$

Part IV

Linking data to ontologies

Outline

- 10 The Description Logic $DL-Lite_{\mathcal{A}}$
- 11 Connecting ontologies to relational data

Outline

- 10 The Description Logic $DL\text{-Lite}_{\mathcal{A}}$
 - Missing features in $DL\text{-Lite}$
 - Combining functionality and role inclusions
 - Syntax and semantics of $DL\text{-Lite}_{\mathcal{A}}$
 - Reasoning in $DL\text{-Lite}_{\mathcal{A}}$

- 11 Connecting ontologies to relational data

What is missing in *DL-Lite* wrt popular data models?

Let us consider UML class diagrams that have the following features:

- functionality of associations (i.e., roles)
- inclusion (i.e., ISA) between associations
- attributes of concepts and associations, possibly functional
- covering constraints in hierarchies

What can we capture of these while maintaining FOL-rewritability?

- ① We can **forget about covering constraints**, since they make query answering coNP-hard in data complexity (see Part 3).
- ② **Attributes of concepts** are “syntactic sugar” (they could be modeled by means of roles), but their functionality is an issue.
- ③ We could also add **attributes of roles** (we won't discuss this here).
- ④ **Functionality and role inclusions** are present separately (in *DL-Lite_F* and *DL-Lite_R*), but *were not allowed to be used together*.

Let us first analyze this last point.

Combining functionalities and role inclusions

We have seen till now that:

- By including in *DL-Lite* both functionality of roles and qualified existential quantification (i.e., $\exists P.A$), query answering becomes NLOGSPACE-hard (and PTIME-hard with also inverse roles) in data complexity (see Part 3).
- Qualified existential quantification can be simulated by using role inclusion assertions (see Part 2).
- When the data complexity of query answering is NLOGSPACE or above, the DL does not enjoy FOL-rewritability.

As a consequence of these results, we get:

To preserve FOL-rewritability, we need to restrict the interaction of functionality and role inclusions.

Let us analyze on an example the effect of an unrestricted interaction

Combining functionalities and role inclusions – Example

$$\text{TBox } \mathcal{T}: \quad A \sqsubseteq \exists P \qquad P \sqsubseteq S$$

$$\qquad \exists P^- \sqsubseteq A \qquad (\text{funct } S)$$

$$\text{ABox } \mathcal{A}: \quad A(c_1), S(c_1, c_2), S(c_2, c_3), \dots, S(c_{n-1}, c_n)$$

$$\begin{array}{rcl}
 A(c_1), & A \sqsubseteq \exists P & \models P(c_1, x), \text{ for some } x \\
 P(c_1, x), & P \sqsubseteq S & \models S(c_1, x) \\
 S(c_1, x), & S(c_1, c_2), (\text{funct } S) & \models x = c_2 \\
 P(c_1, c_2), & \exists P^- \sqsubseteq A & \models A(c_2) \\
 A(c_2), & A \sqsubseteq \exists P & \dots \\
 & & \models A(c_n)
 \end{array}$$

Hence, we get:

- If we add $B(c_n)$ and $B \sqsubseteq \neg A$, the ontology becomes inconsistent.
- Similarly, the answer to the following query over $\langle \mathcal{T}, \mathcal{A} \rangle$ is *true*:

$$q() \leftarrow A(z_1), S(z_1, z_2), S(z_2, z_3), \dots, S(z_{n-1}, z_n), A(z_n)$$

Restrictions on combining functionalities and role inclusions

Note: The number of unification steps above **depends on the data**. Hence this kind of deduction cannot be mimicked by a FOL (or SQL) query, since it requires a form of **recursion**. As a consequence, we get:

Combining functionality and role inclusions is problematic.

It breaks **separability**, i.e., functionality assertions may force existentially quantified objects to be unified with existing objects.

Note: the problems are caused by the **interaction** among:

- an inclusion $P \sqsubseteq S$ between roles,
- a functionality assertion (**funct** S) on the super-role, and
- a cycle of concept inclusion assertions $A \sqsubseteq \exists P$ and $\exists P^- \sqsubseteq A$.

Since we do not want to limit cycles of ISA, we pose suitable restrictions on the combination of functionality and role inclusions

Features of $DL\text{-Lite}_A$

$DL\text{-Lite}_A$ is a Description Logic designed to capture as much features as possible of conceptual data models, while preserving nice computational properties for query answering.

- Enjoys **FOL-rewritability**, and hence is LOGSPACE in data complexity.
- Allows for **both functionality assertions and role inclusion assertions**, but restricts in a suitable way their interaction.
- Takes into account the distinction between **objects** and **values**:
 - Objects are elements of an abstract interpretation domain.
 - Values are elements of concrete data types, such as integers, strings, ecc.
- Values are connected to objects through **attributes**, rather than roles (we consider here only concept attributes and not role attributes [CDGL⁺06a]).

Syntax of the $DL\text{-Lite}_{\mathcal{A}}$ description language

- Concept expressions:

$$\begin{aligned} B &\longrightarrow A \mid \exists Q \mid \delta(U) \\ C &\longrightarrow \top_C \mid B \mid \neg B \mid \exists Q.C \end{aligned}$$

- Role expressions:

$$\begin{aligned} Q &\longrightarrow P \mid P^- \\ R &\longrightarrow Q \mid \neg Q \end{aligned}$$

- Value-domain expressions: (each T_i is one of the RDF datatypes)

$$\begin{aligned} E &\longrightarrow \rho(U) \\ F &\longrightarrow \top_D \mid T_1 \mid \cdots \mid T_n \end{aligned}$$

- Attribute expressions:

$$V \longrightarrow U \mid \neg U$$

Semantics of $DL-Lite_{\mathcal{A}}$ – Objects vs. values

We make use of an alphabet Γ of constants, partitioned into:

- an alphabet Γ_O of object constants.
- an alphabet Γ_V of value constants, in turn partitioned into alphabets Γ_{V_i} , one for each RDF datatype T_i .

The interpretation domain $\Delta^{\mathcal{I}}$ is partitioned into:

- a domain of objects $\Delta_O^{\mathcal{I}}$
- a domain of values $\Delta_V^{\mathcal{I}}$

The semantics of $DL-Lite_{\mathcal{A}}$ descriptions is determined as usual, considering the following:

- The interpretation $C^{\mathcal{I}}$ of a concept C is a subset of $\Delta_O^{\mathcal{I}}$.
- The interpretation $R^{\mathcal{I}}$ of a role R is a subset of $\Delta_O^{\mathcal{I}} \times \Delta_O^{\mathcal{I}}$.
- The interpretation $val(v)$ of each value constant v in Γ_V and RDF datatype T_i is given a priori (e.g., all strings for `xsd:string`).
- The interpretation $V^{\mathcal{I}}$ of an attribute V is a subset of $\Delta_O^{\mathcal{I}} \times \Delta_V^{\mathcal{I}}$.

Semantics of the $DL\text{-Lite}_{\mathcal{A}}$ constructs

Construct	Syntax	Example	Semantics
top concept	\top_C		$\top_C^I = \Delta_O^I$
atomic concept	A	Doctor	$A^I \subseteq \Delta_O^I$
existential restriction	$\exists Q$	$\exists \text{child}^-$	$\{o \mid \exists o'. (o, o') \in Q^I\}$
qualified exist. restriction	$\exists Q.C$	$\exists \text{child.Male}$	$\{o \mid \exists o'. (o, o') \in Q^I \wedge o' \in C^I\}$
concept negation	$\neg B$	$\neg \exists \text{child}$	$\Delta^I \setminus B^I$
attribute domain	$\delta(U)$	$\delta(\text{salary})$	$\{o \mid \exists v. (o, v) \in U^I\}$
atomic role	P	child	$P^I \subseteq \Delta_O^I \times \Delta_O^I$
inverse role	P^-	child^-	$\{(o, o') \mid (o', o) \in P^I\}$
role negation	$\neg Q$	$\neg \text{manages}$	$(\Delta_O^I \times \Delta_O^I) \setminus Q^I$
top domain	\top_D		$\top_D^I = \Delta_V^I$
datatype	T_i	<code>xsd:int</code>	$\text{val}(T_i) \subseteq \Delta_V^I$
attribute range	$\rho(U)$	$\rho(\text{salary})$	$\{v \mid \exists o. (o, v) \in U^I\}$
atomic attribute	U	salary	$U^I \subseteq \Delta_O^I \times \Delta_V^I$
attribute negation	$\neg U$	$\neg \text{salary}$	$(\Delta_O^I \times \Delta_V^I) \setminus U^I$
object constant	c	john	$c^I \in \Delta_O^I$
value constant	v	'john'	$\text{val}(v) \in \Delta_V^I$

$DL\text{-Lite}_{\mathcal{A}}$ assertions

TBox assertions can have the following forms:

$B \sqsubseteq C$	concept inclusion assertion
$Q \sqsubseteq R$	role inclusion assertion
$E \sqsubseteq F$	value-domain inclusion assertion
$U \sqsubseteq V$	attribute inclusion assertion
(funct Q)	role functionality assertion
(funct U)	attribute functionality assertion

ABox assertions: $A(c)$, $P(c, c')$, $U(c, d)$,
 where c, c' are object constants
 d is a value constant

Semantics of the $DL\text{-Lite}_{\mathcal{A}}$ assertions

Assertion	Syntax	Example	Semantics
conc. incl.	$B \sqsubseteq C$	Father $\sqsubseteq \exists\text{child}$	$B^{\mathcal{I}} \subseteq C^{\mathcal{I}}$
role incl.	$Q \sqsubseteq R$	father $\sqsubseteq \text{anc}$	$Q^{\mathcal{I}} \subseteq R^{\mathcal{I}}$
v.dom. incl.	$E \sqsubseteq F$	$\rho(\text{age}) \sqsubseteq \text{xsd:int}$	$E^{\mathcal{I}} \subseteq F^{\mathcal{I}}$
attr. incl.	$U \sqsubseteq V$	offPhone $\sqsubseteq \text{phone}$	$U^{\mathcal{I}} \subseteq V^{\mathcal{I}}$
role funct.	(funct Q)	(funct father)	$\forall o, o, o''. (o, o') \in Q^{\mathcal{I}} \wedge (o, o'') \in Q^{\mathcal{I}} \rightarrow o' = o''$
att. funct.	(funct U)	(funct ssn)	$\forall o, v, v'. (o, v) \in U^{\mathcal{I}} \wedge (o, v') \in U^{\mathcal{I}} \rightarrow v = v'$
mem. asser.	$A(c)$	Father(bob)	$c^{\mathcal{I}} \in A^{\mathcal{I}}$
mem. asser.	$P(c_1, c_2)$	child(bob, ann)	$(c_1^{\mathcal{I}}, c_2^{\mathcal{I}}) \in P^{\mathcal{I}}$
mem. asser.	$U(c, d)$	phone(bob, '2345')	$(c^{\mathcal{I}}, \text{val}(d)) \in U^{\mathcal{I}}$

Restriction on TBox assertions in $DL\text{-Lite}_{\mathcal{A}}$ ontologies

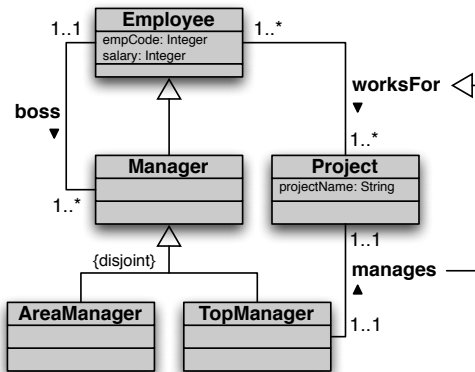
As shown, to ensure FOL-rewritability, we have to impose a **restriction** on the use of functionality and role/attribute inclusions.

Restriction on $DL\text{-Lite}$ TBoxes

No functional role or attribute can be specialized by using it in the right-hand side of a role or attribute inclusion assertions.

Formally:

- If $\exists P.C$ or $\exists P^-.C$ appears in \mathcal{T} , then **(*funct* P)** and **(*funct* P^-)** are **not in \mathcal{T}** .
- If $Q \sqsubseteq P$ or $Q \sqsubseteq P^-$ is in \mathcal{T} , then **(*funct* P)** and **(*funct* P^-)** are **not in \mathcal{T}** .
- If $U_1 \sqsubseteq U_2$ is in \mathcal{T} , then **(*funct* U_2)** is **not in \mathcal{T}** .

$DL\text{-Lite}_{\mathcal{A}}$ – Example

Manager \sqsubseteq Employee
 AreaManager \sqsubseteq Manager
 TopManager \sqsubseteq Manager
 AreaManager \sqsubseteq \neg TopManager

Employee \sqsubseteq δ (salary)
 δ (salary) \sqsubseteq Employee
 ρ (salary) \sqsubseteq xsd:int
 (funcnt salary)

\exists worksFor \sqsubseteq Employee
 \exists worksFor $^-$ \sqsubseteq Project
 Employee \sqsubseteq \exists worksFor
 Project \sqsubseteq \exists worksFor $^-$

(funcnt manages)

(funcnt manages $^-$)

manages \sqsubseteq worksFor

⋮

Note: in $DL\text{-Lite}_{\mathcal{A}}$ we still cannot capture:

- completeness of the hierarchy
- number restrictions

Reasoning in $DL-Lite_{\mathcal{A}}$ – Separation

It is possible to show that, by virtue of the restriction on the use of role inclusion and functionality assertions, all nice properties of $DL-Lite_{\mathcal{F}}$ and $DL-Lite_{\mathcal{R}}$ continue to hold also for $DL-Lite_{\mathcal{A}}$.

In particular, w.r.t. **satisfiability of a $DL-Lite_{\mathcal{A}}$ ontology \mathcal{O}** , we have:

- NIs do not interact with each other.
- NIs and PIs do not interact with functionality assertions.

We obtain that for $DL-Lite_{\mathcal{A}}$ a **separation** result holds:

- Each NI and each functionality can be checked independently from the others.
- A functionality assertion is contradicted in an ontology $\mathcal{O} = \langle \mathcal{T}, \mathcal{A} \rangle$ only if it is explicitly contradicted by its ABox \mathcal{A} .

Ontology satisfiability in $DL\text{-Lite}_{\mathcal{A}}$

Due to the separation property, we can associate

- to each NI N a boolean CQ q_N , and
- to each functionality assertion F a boolean CQ q_F .

and check satisfiability of \mathcal{O} by suitably evaluating q_N and q_F .

Theorem

Let $\mathcal{O} = \langle \mathcal{T}, \mathcal{A} \rangle$ be a $DL\text{-Lite}_{\mathcal{A}}$ ontology, and \mathcal{T}_P the set of PIs in \mathcal{O} . Then, \mathcal{O} is unsatisfiable iff one of the following condition holds:

- There exists a NI $N \in \mathcal{T}$ such that $\text{Eval}(\text{SQL}(\text{PerfectRef}(q_N, \mathcal{T}_P)), \text{DB}(\mathcal{A}))$ returns *true*.
- There exists a functionality assertion $F \in \mathcal{T}$ such that $\text{Eval}(\text{SQL}(q_F), \text{DB}(\mathcal{A}))$ returns *true*.

Query answering in $DL-Lite_{\mathcal{A}}$

- Queries over $DL-Lite_{\mathcal{A}}$ ontologies are analogous to those over $DL-Lite_{\mathcal{R}}$ and $DL-Lite_{\mathcal{F}}$ ontologies, except that they can also make use of attribute and domain atoms.
- Exploiting the previous result, the query answering algorithm of $DL-Lite_{\mathcal{R}}$ can be easily extended to deal with $DL-Lite_{\mathcal{A}}$ ontologies:
 - Assertions involving attribute domain and range can be dealt with as for role domain and range assertions.
 - $\exists Q.C$ in the right hand-side of concept inclusion assertions can be eliminated by making use of role inclusion assertions.
 - Disjointness of roles and attributes can be checked similarly as for disjointness of concepts, and does not interact further with the other assertions.

Complexity of reasoning in $DL-Lite_{\mathcal{A}}$

As for ontology satisfiability, $DL-Lite_{\mathcal{A}}$ maintains the nice computational properties of $DL-Lite_{\mathcal{R}}$ and $DL-Lite_{\mathcal{F}}$ also w.r.t. query answering. Hence, we get the same characterization of computational complexity.

Theorem

For $DL-Lite_{\mathcal{A}}$ ontologies:

- Checking **satisfiability of the ontology** is
 - **P**TIME in the size of the **ontology** (combined complexity).
 - **LOGSPACE** in the size of the **ABox** (data complexity).
- **TBox reasoning** is **P**TIME in the size of the **TBox**.
- **Query answering** is
 - **NP-complete** in the size of the query and the ontology (comb. com.).
 - **P**TIME in the size of the **ontology**.
 - **LOGSPACE** in the size of the **ABox** (data complexity).

Outline

- 10 The Description Logic $DL-Lite_A$
- 11 Connecting ontologies to relational data
 - The impedance mismatch problem
 - Ontology-Based Data Access System
 - Query answering in Ontology-Based Data Access Systems

Managing ABoxes

In all the previous discussion, we have assumed that the data is maintained in the **ABox** of the ontology:

- The ABox is perfectly compatible with the TBox:
 - the vocabulary of concepts, roles, and attributes is the one used in the TBox.
 - An ABox “stores” abstract objects, and these objects and their properties are those returned by queries over the ontology.
- There may be different ways to manage the ABox from a physical point of view:
 - Description Logics reasoners maintain the ABox in main-memory data structures.
 - When ABoxes become large, managing them in secondary storage may be required, but this is again handled directly by the reasoner.

Data in external sources

There are several situations where the assumptions of having the data in an ABox managed directly by the ontology system (e.g., a Description Logics reasoner) is not feasible or realistic:

- When the ABox is very large, so that it requires relational database technology.
- When have no direct control over the data since it belongs to some external organization, which controls the access to it.
- When multiple data sources need to be accessed, such as in Information Integration.

We would like to deal with such situation by keeping the data in the external (relational) storage, and performing **query answering** by leveraging the capabilities of the **relational engine**.

The impedance mismatch problem


We have to deal with the **impedance mismatch problem**:

- Sources store data, which is constituted by values taken from concrete domains, such as strings, integers, codes, . . .
- Instead, instances of concepts and relations in an ontology are (abstract) objects.

The solution is to define a **mapping language** that allows for specifying how to transform data into objects:

- Basic idea: use Skolem functions in the head of the mapping to “generate” the objects.
- Semantics: objects are denoted by terms (of exactly one level of nesting), and different terms denote different objects (unique name assumption on terms).

Impedance mismatch – Example



Actual data is stored in a DB:
 – An Employee is identified by her *SSN*.
 – A Project is identified by its *name*.

$D_1[SSN: String, PrName: String]$
 Employees and Projects they work for

$D_2[Code: String, Salary: Int]$
 Employee's Code with salary

$D_3[Code: String, SSN: String]$
 Employee's Code with SSN

...

Intuitively:

- An employee should be created from her *SSN*: **pers(SSN)**
- A project should be created from its *Name*: **proj(PrName)**

Creating object identifiers

We need to associate to the data in the tables objects in the ontology.

- We introduce an alphabet Λ of **function symbols**, each with an associated arity.
- To denote values, we use value constants in Γ_V as before.
- To denote objects, we use **object terms** instead of object constants. An object term has the form $\mathbf{f}(d_1, \dots, d_n)$, with $\mathbf{f} \in \Lambda$, and each d_i a value constant in Γ_V .

Example

- If a person is identified by its *SSN*, we can introduce a function symbol **pers/1**. If *VRD56B25* is a *SSN*, then **pers(VRD56B25)** denotes a person.
- If a person is identified by its *name* and *dateOfBirth*, we can introduce a function symbol **pers/2**. Then **pers(Vardi, 25/2/56)** denotes a person.

Mapping assertions

Mapping assertions are used to extract the data from the DB to populate the ontology.

We make use of **variable terms**, which are as object terms, but with variables instead of values as arguments of the functions.


A **mapping assertion** between a database \mathcal{D} and a TBox \mathcal{T} has the form

$$\Phi \rightsquigarrow \Psi$$

where

- Φ is an arbitrary SQL query of arity $n > 0$ over \mathcal{D} .
- Ψ is a conjunctive query over \mathcal{T} of arity $n' > 0$ **without non-distinguished variables**, possibly involving variable terms.

Mapping assertions – Example



D_1 [*SSN: String, PrName: String*]

Employees and Projects they work for

D_2 [*Code: String, Salary: Int*]

Employee's Code with salary

D_3 [*Code: String, SSN: String*]

Employee's Code with SSN

...

M_1 : SELECT SSN, PrName
 FROM D_1

\rightsquigarrow Employee(**pers**(SSN)),
 Project(**proj**(PrName)),
 projectName(**proj**(PrName), PrName),
 workFor(**pers**(SSN), **proj**(PrName))

M_2 : SELECT SSN, Salary
 FROM D_2, D_3
 WHERE $D_2.CODE = D_3.CODE$

\rightsquigarrow Employee(**pers**(SSN)),
 salary(**pers**(SSN), Salary)

Ontology-Based Data Access System

The mapping assertions are a crucial part of an Ontology-Based Data Access System.

Ontology-Based Data Access System

is a triple $\mathcal{O} = \langle \mathcal{T}, \mathcal{M}, \mathcal{D} \rangle$, where

- \mathcal{T} is a TBox.
- \mathcal{D} is a relational database.
- \mathcal{M} is a set of mapping assertions between \mathcal{T} and \mathcal{D} .

Note: we could consider also mapping assertions between the datatypes of the database and those of the ontology.

Semantics of mappings

We first need to define the semantics of mappings.

Definition

An interpretation \mathcal{I} **satisfies** a mapping assertion $\Phi(\vec{x}) \rightsquigarrow \Psi(\vec{t}, \vec{y})$ in \mathcal{M} **with respect to a database** \mathcal{D} , if for each tuple of values $\vec{v} \in \text{Eval}(\Phi, \mathcal{D})$, and for each ground atom in $\Psi[\vec{x}/\vec{v}]$, we have that:

- if the ground atom is $A(s)$, then $s^{\mathcal{I}} \in A^{\mathcal{I}}$.
- if the ground atom is $T(s)$, then $s^{\mathcal{I}} \in T^{\mathcal{I}}$.
- if the ground atom is $P(s_1, s_2)$, then $(s_1^{\mathcal{I}}, s_2^{\mathcal{I}}) \in P^{\mathcal{I}}$.
- if the ground atom is $U(s_1, s_2)$, then $(s_1^{\mathcal{I}}, s_2^{\mathcal{I}}) \in U^{\mathcal{I}}$.

Intuitively, \mathcal{I} **satisfies** $\Phi \rightsquigarrow \Psi$ w.r.t. \mathcal{D} if all facts obtained by evaluating Φ over \mathcal{D} and then propagating the answers to Ψ , hold in \mathcal{I} .

Note: $\Psi[\vec{x}/\vec{v}]$ denotes Ψ where each x_i has been substituted with v_i .

Semantics of an OBDA system

Model of an OBDA system

An interpretation \mathcal{I} is a **model** of $\mathcal{O} = \langle \mathcal{T}, \mathcal{M}, \mathcal{D} \rangle$ if:

- \mathcal{I} is a model of \mathcal{T} ;
- \mathcal{I} satisfies \mathcal{M} w.r.t. \mathcal{D} , i.e., satisfies every assertion in \mathcal{M} w.r.t. \mathcal{D} .

An OBDA system \mathcal{O} is **satisfiable** if it admits at least one model.

Answering queries over an OBDA system

In an OBDA system $\mathcal{O} = \langle \mathcal{T}, \mathcal{M}, \mathcal{D} \rangle$

- Queries are posed over the TBox \mathcal{T} .
- The data needed to answer queries is stored in the database \mathcal{D} .
- The mapping \mathcal{M} is used to bridge the gap between \mathcal{T} and \mathcal{D} .

Two approaches to exploit the mapping:

- bottom-up approach: simpler, but less efficient
- top-down approach: more sophisticated, but also more efficient

Note: Both approaches require to first **split** the TBox queries in the mapping assertions into their constituent atoms. This is possible, since all variables in such queries are distinguished.

Bottom-up approach to query answering

Consists in a straightforward application of the mappings:

- ① Propagate the data from \mathcal{D} through \mathcal{M} , materializing an ABox $\mathcal{A}_{\mathcal{M},\mathcal{D}}$ (the constants in such an ABox are values and object terms).
- ② Apply to $\mathcal{A}_{\mathcal{M},\mathcal{D}}$ and to the TBox \mathcal{T} , the satisfiability and query answering algorithms developed for $DL-Lite_A$.

This approach has several drawbacks (hence is only theoretical):

- The technique is no more LOGSPACE in the data, since the ABox $\mathcal{A}_{\mathcal{M},\mathcal{D}}$ to materialize is in general polynomial in the size of the data.
- $\mathcal{A}_{\mathcal{M},\mathcal{D}}$ may be very large, and thus it may be infeasible to actually materialize it.
- Freshness of $\mathcal{A}_{\mathcal{M},\mathcal{D}}$ with respect to the underlying data source(s) may be an issue, and one would need to propagate updates (cf. Data Warehousing).

Top-down approach to query answering

Consists of three steps:

- ① **Reformulation:** Compute the perfect reformulation $q' = \text{PerfectRef}(q, \mathcal{T}_P)$ of the original query q , using the PIs \mathcal{T}_P of the TBox \mathcal{T} .
- ② **Unfolding:** Compute from q' a new query q'' by unfolding q' using (the split version of) the mappings \mathcal{M} .
 - Essentially, each atom in q' that unifies with an atom in Ψ is substituted with the corresponding query Φ over the database.
 - The unfolded query q'' is such that $\text{Eval}(q'', \mathcal{D}) = \text{Eval}(q', \mathcal{A}_{\mathcal{M}, \mathcal{D}})$.
- ③ **Evaluation:** Delegate the evaluation of q'' to the relational DBMS managing \mathcal{D} .

For details, see [PLC⁺07].

Computational complexity of query answering

Theorem

Query answering in a $DL-Lite_{\mathcal{A}}$ OBDM system $\mathcal{O} = \langle \mathcal{T}, \mathcal{M}, \mathcal{D} \rangle$ is

- 1 **NP-complete** in the size of the query.
- 2 **P TIME** in the size of the **TBox** \mathcal{T} and the **mappings** \mathcal{M} .
- 3 **LOGSPACE** in the size of the **database** \mathcal{D} .

Moreover, the **LOGSPACE** result is actually a consequence of the fact that query answering in such a setting can be reduced to evaluating an SQL query over the relational database.

Part V

Hands-on session

Outline

- 12 The MASTRO system
- 13 XML MASTRO input files

Outline

- 12 The MASTRO system
- 13 XML MASTRO input files

The system MASTRO

- MASTRO is a Java-based tool for **Ontology-based Data Access**.
- It allows for the specification of Ontologies in the DL $DL-Lite_{\mathcal{A}}$.
- In MASTRO, $DL-Lite_{\mathcal{A}}$ TBoxes are connected to an **external RDBMS** through suitable mappings.
- In each mapping a generic SQL query over the external RDBMS is put in correspondence with a CQ without existential variables expressed over the $DL-Lite_{\mathcal{A}}$ TBox.

MASTRO vs QUONTO

- At its core, MASTRO uses QUONTO (<http://www.dis.uniroma1.it/~quonto/>) a reasoner for $DL-Lite_A$, which provides query reformulation services (QUONTO implements the algorithm PerfectRef).
- Notice that QUONTO is not designed to support ontology-based data access, and therefore it is not able to deal with mappings to external RDBMs \leadsto MASTRO provides this support.

Reasoning in MASTRO

- The basic services provided by MASTRO are
 - **Specification of a $DL\text{-}Lite_{\mathcal{A}}$ OBDA System**
 - **Query answering**, for computing certain answer for (unions of) conjunctive queries over $DL\text{-}Lite_{\mathcal{A}}$ OBDA Systems
 - **Consistency check**, for verifying satisfiability of $DL\text{-}Lite_{\mathcal{A}}$ OBDA Systems

These are the only services supported by the version of MASTRO demonstrated at the tutorial.

- However, the full version of MASTRO also allows for **TBox reasoning**, **meta-level reasoning**, and **ontology updates**.
- We are currently working also on **query answering of complex (i.e., FOL) queries**, introduction on **new $DL\text{-}Lite$ constructs** (e.g., identification assertions).

Input formats

- MASTRO has its own Java-based interface, and accepts inputs in a **proprietary XML format**.
- That is, to give as input a TBox, an ABox, i.e., a set of mapping assertions to an external RDBMS, and a query, we must specify them into XML, according to a specific DTD.
- Nonetheless, the XML syntax to be used is very simple.

Outline

- 12 The MASTRO system
- 13 XML MASTRO input files

XML TBox: alphabet

```
<alphabet>
  <atomicC>professor</atomicC>
  <atomicC>assistantProf</atomicC>
    . . . . .
  <atomicCA>name</atomicCA>
    . . . .
  <atomicR>WORKS_FOR</atomicR>
    . . . .
  <atomicRA>date</atomicRA>
</alphabet>
```

XML TBox: inclusion assertions

$$\text{assistantProfessor} \sqsubseteq \neg \text{fullProf}$$

```

<inclusionAssertion>
  <basicC>
    <atomicC>assistantProf</atomicC>
  </basicC>
  <generalC>
    <signedC sign="negative">
      <basicC>
        <atomicC>fullProf</atomicC>
      </basicC>
    </signedC>
  </generalC>
</inclusionAssertion>

```

$$\exists \text{TAKES_COURSE}^- \sqsubseteq \text{course}$$

```

<inclusionAssertion>
  <basicC>
    <exists>
      <basicR dir="inverse">
        <atomicC>TAKES_COURSE</atomicC>
      </basicR>
    </exists>
  </basicC>
  <generalC>
    <signedC sign="positive">
      <basicC>
        <atomicC>course</atomicC>
      </basicC>
    </signedC>
  </generalC>
</inclusionAssertion>

```

XML TBox: functionality assertions

(**funct** *ENROLLED*)

```
<funct>  
  <basicR dir="direct">  
    <atomicR>ENROLLED</atomicR>  
  </basicR>  
</funct>
```

(**funct** term)

```
<funct>  
  <atomicCA>term</atomicCA>  
</funct>
```

XML ABox: mapping specification

```

SELECT C_Name, Term,   ~→ course(course(C_Name)),
      Prof_Id          name(course(C_Name), C_Name),
FROM Course_Tab       term(course(C_Name), Term),
                      professor(prof(Prof_Id)),
                      TEACHES(prof(Prof_Id), course(C_Name))

```

It is probably better to see the mapping as follows

```

SELECT C_Name, Term,   ~→ course(X), X = course(C_Name),
      Prof_Id          name(X, Y), Y = C_Name,
FROM Course_Tab       term(X, Z), Z = Term,
                      professor(W), W = prof(Prof_Id),
                      TEACHES(W, X)

```

XML ABox: mapping specification

```

SELECT C_Name, Term,   ~> course(X), X = course(C_Name),
      Prof_Id         name(X, Y), Y = C_Name,
FROM Course_Tab      term(X, Z), Z = Term,
                    professor(W), W = prof(Prof_Id), TEACHES(W, X)

```

```

<mapping>
  <head>
    <CQBody>
      .....
      <atom>
        <AtomicConceptAttribute name="term">
          <term>
            <var name="X"/>
          </term>
          <term>
            <var name="Z"/>
          </term>
        </AtomicConceptAttribute>
      </atom>
      .....
    </CQBody>
  </head>
  .....

```


XML ABox: mapping specification

```

SELECT C_Name, Term,   ~→ course(X), X = course(C_Name),
      Prof_Id         name(X, Y), Y = C_Name,
FROM Course_Tab       term(X, Z), Z = Term,
                      professor(W), W = prof(Prof_Id), TEACHES(W, X)

```

```

<mapping>
  .....
  <map>
    <objMap>
      <dtVar>X</dtVar>
      <sqlObjVar funct="course">C_Name</sqlObjVar>
    </objMap>
  </map>
  .....
  <map>
    <valueMap>
      <dtVar>Z</dtVar>
      <sqlValueVar type="xs:integer">Term</sqlValueVar>
    </valueMap>
  </map>
  <body>SELECT C_Name, Term, Prof_Id FROM Course_Tab</body>
  .....
</mapping>

```

Acknowledgements

- Giuseppe De Giacomo
- Enrico Franconi
- Maurizio Lenzerini
- Marco Ruzzi
- Antonella Poggi
- Riccardo Rosati
- Sergio Tessaris

References I

- [ACK⁺07] A. Artale, D. Calvanese, R. Kontchakov, V. Ryzhikov, and M. Zakharyashev.
Reasoning over extended ER models.
In *Proc. of the 26th Int. Conf. on Conceptual Modeling (ER 2007)*, volume 4801 of *Lecture Notes in Computer Science*, pages 277–292. Springer, 2007.
- [BBL05] F. Baader, S. Brandt, and C. Lutz.
Pushing the \mathcal{EL} envelope.
In *Proc. of the 19th Int. Joint Conf. on Artificial Intelligence (IJCAI 2005)*, pages 364–369, 2005.
- [BCDG05] D. Berardi, D. Calvanese, and G. De Giacomo.
Reasoning on UML class diagrams.
Artificial Intelligence, 168(1–2):70–118, 2005.

References II

- [BCM⁺03] F. Baader, D. Calvanese, D. McGuinness, D. Nardi, and P. F. Patel-Schneider, editors.
The Description Logic Handbook: Theory, Implementation and Applications.
Cambridge University Press, 2003.
- [CDGL98] D. Calvanese, G. De Giacomo, and M. Lenzerini.
On the decidability of query containment under constraints.
In Proc. of the 17th ACM SIGACT SIGMOD SIGART Symp. on Principles of Database Systems (PODS'98), pages 149–158, 1998.
- [CDGL⁺05a] D. Calvanese, G. De Giacomo, D. Lembo, M. Lenzerini, and R. Rosati.
DL-Lite: Tractable description logics for ontologies.
In Proc. of the 20th Nat. Conf. on Artificial Intelligence (AAAI 2005), pages 602–607, 2005.

References III

- [CDGL⁺05b] D. Calvanese, G. De Giacomo, D. Lembo, M. Lenzerini, and R. Rosati.
Tailoring OWL for data intensive ontologies.
In Proc. of the Workshop on OWL: Experiences and Directions (OWLED 2005), 2005.
- [CDGL⁺06a] D. Calvanese, G. De Giacomo, D. Lembo, M. Lenzerini, A. Poggi, and R. Rosati.
Linking data to ontologies: The description logic DL-Lite_A.
In Proc. of the 2nd Workshop on OWL: Experiences and Directions (OWLED 2006), 2006.
- [CDGL⁺06b] D. Calvanese, G. De Giacomo, D. Lembo, M. Lenzerini, and R. Rosati.
Data complexity of query answering in description logics.
In Proc. of the 10th Int. Conf. on the Principles of Knowledge Representation and Reasoning (KR 2006), pages 260–270, 2006.

References IV

- [CDGL⁺07] D. Calvanese, G. De Giacomo, D. Lembo, M. Lenzerini, and R. Rosati.
Tractable reasoning and efficient query answering in description logics:
The DL-Lite family.
J. of Automated Reasoning, 39(3):385–429, 2007.
- [DLNN97] F. M. Donini, M. Lenzerini, D. Nardi, and W. Nutt.
The complexity of concept languages.
Information and Computation, 134:1–58, 1997.
- [DLNS94] F. M. Donini, M. Lenzerini, D. Nardi, and A. Schaerf.
Deduction in concept languages: From subsumption to instance
checking.
J. of Logic and Computation, 4(4):423–452, 1994.
- [Don03] F. M. Donini.
Complexity of reasoning.
In Baader et al. [BCM⁺03], chapter 3, pages 96–136.

References V

- [GHLS07] B. Glimm, I. Horrocks, C. Lutz, and U. Sattler.
Conjunctive query answering for the description logic *SHIQ*.
In *Proc. of the 20th Int. Joint Conf. on Artificial Intelligence (IJCAI 2007)*, pages 399–404, 2007.
- [GHVD03] B. N. Groszof, I. Horrocks, R. Volz, and S. Decker.
Description logic programs: Combining logic programs with description logic.
In *Proc. of the 12th Int. World Wide Web Conf. (WWW 2003)*, pages 48–57, 2003.
- [HMS05] U. Hustadt, B. Motik, and U. Sattler.
Data complexity of reasoning in very expressive description logics.
In *Proc. of the 19th Int. Joint Conf. on Artificial Intelligence (IJCAI 2005)*, pages 466–471, 2005.

References VI

- [LR98] A. Y. Levy and M.-C. Rousset.
Combining Horn rules and description logics in CARIN.
Artificial Intelligence, 104(1–2):165–209, 1998.
- [Lut07] C. Lutz.
Inverse roles make conjunctive queries hard.
In Proc. of the 2007 Description Logic Workshop (DL 2007), 2007.
- [MH03] R. Möller and V. Haarslev.
Description logic systems.
In Baader et al. [BCM⁺03], chapter 8, pages 282–305.
- [OCE06] M. M. Ortiz, D. Calvanese, and T. Eiter.
Characterizing data complexity for conjunctive query answering in expressive description logics.
In Proc. of the 21st Nat. Conf. on Artificial Intelligence (AAAI 2006), 2006.

References VII

- [PLC⁺07] A. Poggi, D. Lembo, D. Calvanese, G. De Giacomo, M. Lenzerini, and R. Rosati.
Linking data to ontologies.
J. on Data Semantics, 2007.
To appear.