

Knowledge Bases and Databases

Part 2: Ontology-Based Access to Information

Diego Calvanese

Faculty of Computer Science
Master of Science in Computer Science

A.Y. 2007/2008



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

Overview of Part 2: Ontology-based access to information

- 1 Introduction to ontology-based access to information
 - 1 Introduction to ontologies
 - 2 Ontology languages
- 2 Description Logics and the *DL-Lite* family
 - 1 An 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 Linking ontologies to relational data
 - 1 The impedance mismatch problem
 - 2 OBDA systems
 - 3 Query answering in OBDA systems
- 4 Reasoning in the *DL-Lite* family
 - 1 TBox reasoning
 - 2 TBox & ABox reasoning
 - 3 Complexity of reasoning in Description Logics
 - 4 The Description Logic *DL-Lite_A*



Chapter I

Introduction to ontology-based access to information



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

Def.: **Ontology**

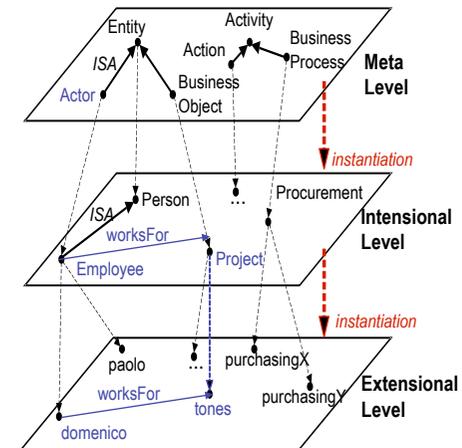
is a representation scheme that describes a **formal conceptualization** of a domain of interest.

The specification of an ontology comprises several levels:

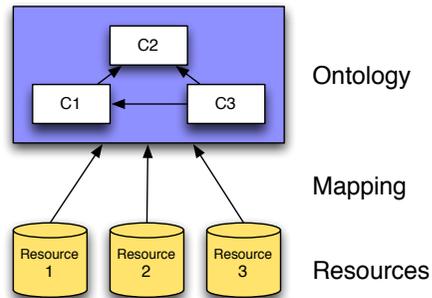
- **Meta-level**: specifies a set of **modeling categories**.
- **Intensional level**: specifies a set of **conceptual elements** (instances of categories) and of rules to describe the conceptual structures of the domain.
- **Extensional level**: specifies a set of **instances** of the conceptual elements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



The three levels of an ontology



Ontologies at the core of information systems



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



Ontology mediated access to data

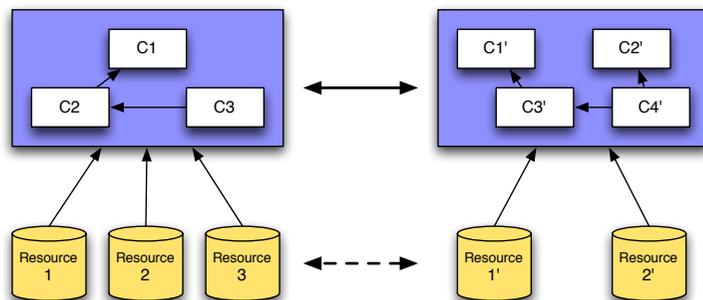
Desiderata: achieve **logical transparency** in access to data:

- **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.

We will see that this setting is similar to the one of Data Integration. The difference is that here the ontology provides a rich conceptual description as the information managed by the system.



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.



Challenge 1: 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 course:

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



Challenge 2: Methodologies

- Developing and dealing with ontologies is a complex and challenging task.
- Developing **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 course:

- 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 discuss specific ontology-development methodologies here.



Challenge 3: 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 course:

- We discuss the requirements for such ontology management tools.
- We will work with 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 course 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 course we concentrate essentially on the static aspects.



Intensional level of an ontology language

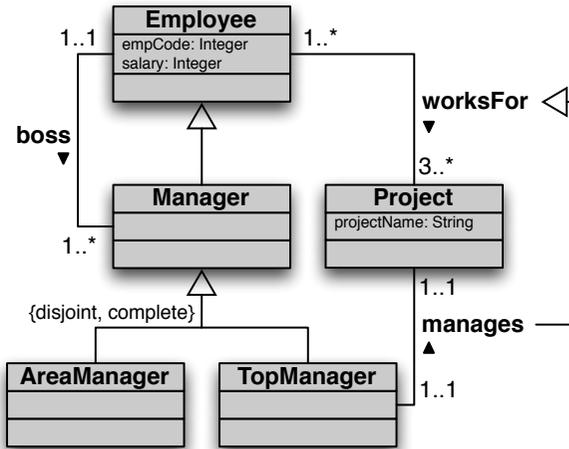
An ontology language for expressing the intensional level usually includes:

- Concepts
- Properties of concepts
- Relationships between concepts, and their properties
- Axioms
- 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

Def.: Concept

Is an element of an 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

Def.: Property

Is an element of an ontology that qualifies another element (e.g., a concept or a relationship).

Property definition (intensional and extensional):

- Name
- Type: may be either
 - atomic (integer, real, string, enumerated, ...), or
e.g., **eye-color** → { **blu, brown, green, grey** }
 - structured (date, set, list, ...)
e.g., **date** → **day/month/year**
- The definition may also specify a default value.

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

Relationships

Def.: Relationship

Is an element of an ontology that 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

Def.: 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., **Manager** \sqsubseteq **Employee**
- equivalences, e.g., **Manager** \equiv **AreaManager** \sqcup **TopManager**
- disjointness, e.g., **AreaManager** \sqcap **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**.



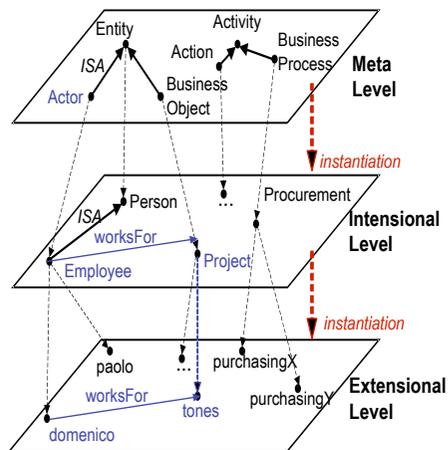
Extensional level of an ontology language

At the extensional level we have individuals and facts:

- An **instance** represents an individual (or object) in the extension of a concept.
e.g., **domenico** is an instance of **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 **the** tool for assigning semantics to ontology languages.
- Ontology languages vs. conceptual data models:
Conceptual schemas **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 class diagrams, Entity-Relationship schemas
- 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., non-monotonic, probabilistic)



Queries

An ontology language may also include constructs for expressing queries.

Def.: Query

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

Def.: Meta-Query

Is 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], [CK06]



Ontology languages vs. query languages

Ontology languages:

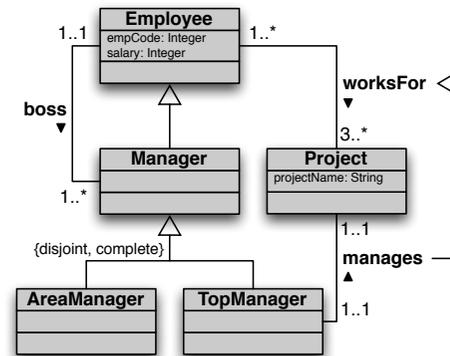
- Tailored for capturing intensional relationships.
- Are quite **poor as query languages**:
 - Cannot refer to same object via multiple navigation paths in the ontology,
 - i.e., allow only for a limited form of JOIN, namely chaining.

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



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


Query answering under different assumptions

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

- traditional database assumption
- knowledge representation assumption



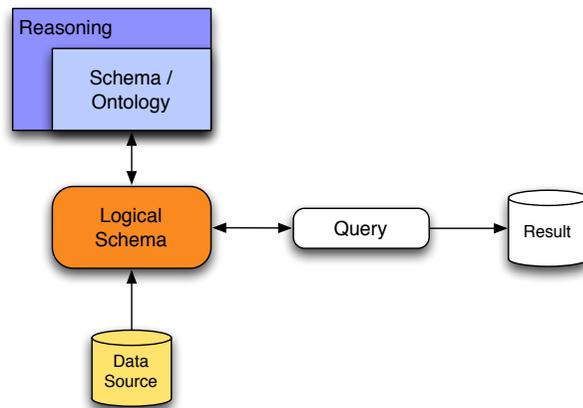
Query answering under the database assumption

- 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).

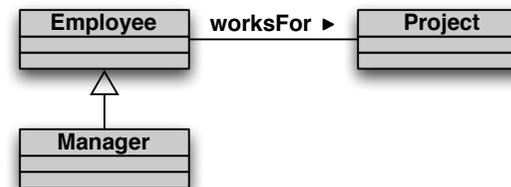
→ Query answering amounts to **query evaluation**, which is computationally easy.



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

Answer: { john }



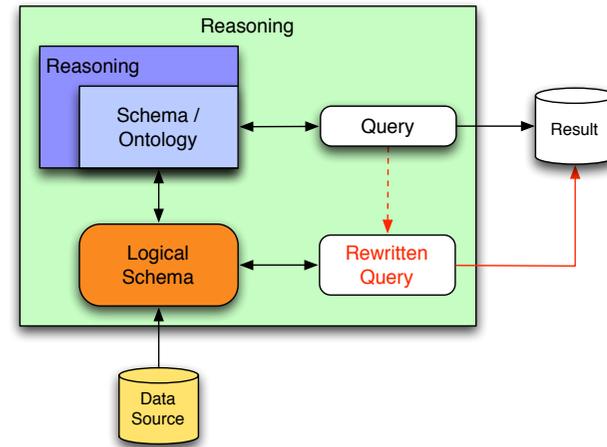
Query answering under the KR assumption

- 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).

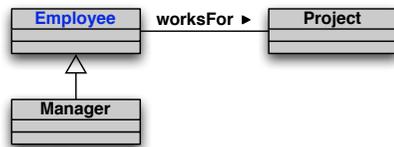
↪ Query answering amounts to **logical inference**, which is computationally more costly.



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

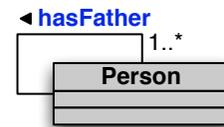
Query: $q(x) \leftarrow \text{Employee}(x)$

Answer: { john, nick, mary }

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



Query answering under the KR assumption – Example 2



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) }

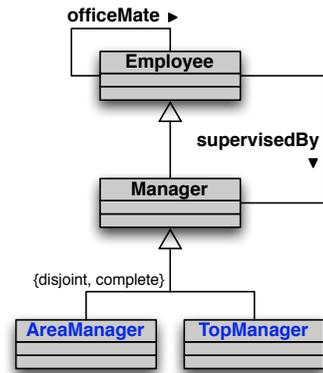
Queries: $q_1(x, y) \leftarrow \text{hasFather}(x, y)$
 $q_2(x) \leftarrow \exists y. \text{hasFather}(x, y)$
 $q_3(x) \leftarrow \exists y_1, y_2, y_3. \text{hasFather}(x, y_1), \text{hasFather}(y_1, y_2), \text{hasFather}(y_2, y_3)$
 $q_4(x, y_3) \leftarrow \exists y_1, y_2. \text{hasFather}(x, y_1),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

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



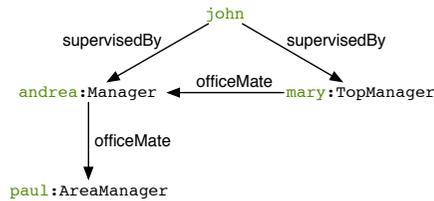
QA under the KR assumption – Andrea's Example



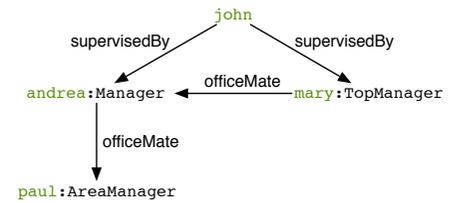
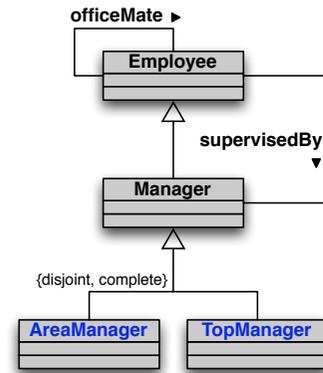
Tables may be **incompletely** specified.

```

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



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



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

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 this course we will study state-of-the-art technology in this area.



Chapter II

Description Logics and the *DL-Lite* family



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)
- ...



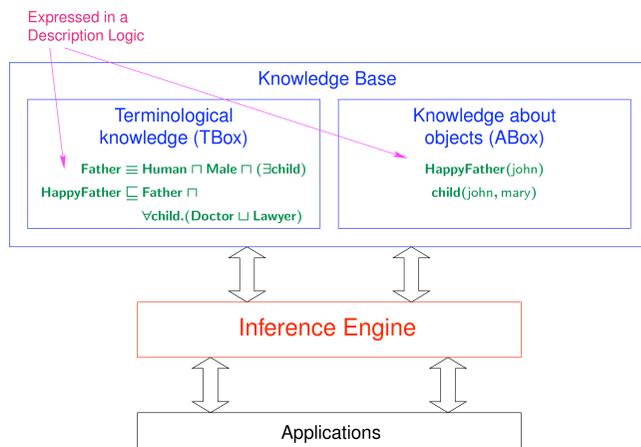
Ingredients of a Description Logic

A DL 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)
 $\mathcal{T} = \{ \text{Father} \sqsubseteq \text{Human} \sqcap \text{Male} \sqcap \exists \text{hasChild}, \text{HappyFather} \sqsubseteq \text{Father} \sqcap \forall \text{hasChild} . (\text{Doctor} \sqcup \text{Lawyer}) \}$
- 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 provides the means for defining:

- **concepts**, corresponding to classes: interpreted as sets of objects;
- **roles**, corresponding to relationships: interpreted as binary relations on objects.

To define concepts and roles:

- We start from a (finite) alphabet of **atomic concepts** and **atomic roles**, i.e., simply names for concept and roles.
- Then, by applying specific **constructors**, we can build **complex concepts** and **roles**, starting from the atomic ones.

A **description language** is characterized by the set of constructs that are available for that.



Semantics of a description language

The **formal semantics** of DLs is given in terms of interpretations.

Def.: 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 $\Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$

Note: A DL interpretation is analogous to a FOL interpretation, except that, by tradition, it is specified in terms of a function $\cdot^{\mathcal{I}}$ rather than a set of (unary and binary) relations.

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	$A^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}}$
atomic role	P	hasChild	$P^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$
atomic negation	$\neg A$	\neg Doctor	$\Delta^{\mathcal{I}} \setminus A^{\mathcal{I}}$
conjunction	$C \sqcap D$	Hum \sqcap Male	$C^{\mathcal{I}} \cap D^{\mathcal{I}}$
(unqual.) exist. res.	$\exists R$	\exists hasChild	$\{ a \mid \exists b. (a, b) \in R^{\mathcal{I}} \}$
value restriction	$\forall R.C$	\forall hasChild.Male	$\{ a \mid \forall b. (a, b) \in R^{\mathcal{I}} \rightarrow b \in C^{\mathcal{I}} \}$
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}	$(\geq k R)$ $(\leq k R)$	$\{ a \mid \#\{b \mid (a, b) \in R^{\mathcal{I}}\} \geq k \}$ $\{ a \mid \#\{b \mid (a, b) \in R^{\mathcal{I}}\} \leq k \}$
qual. number restrictions	\mathcal{Q}	$(\geq k R.C)$ $(\leq k R.C)$	$\{ a \mid \#\{b \mid (a, b) \in R^{\mathcal{I}} \wedge b \in C^{\mathcal{I}}\} \geq k \}$ $\{ a \mid \#\{b \mid (a, b) \in R^{\mathcal{I}} \wedge b \in C^{\mathcal{I}}\} \leq k \}$
inverse role	\mathcal{I}	R^-	$\{ (a, b) \mid (b, a) \in R^{\mathcal{I}} \}$
role closure	reg	R^*	$(R^{\mathcal{I}})^*$

Many different DL constructs and their combinations have been investigated.



Further examples of DL constructs

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



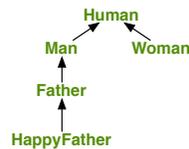
Reasoning on concept expressions

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**:

Def.: 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 $P C$) ...

Def.: 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

Def.: 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 ...

Def.: Logical implication

\mathcal{O} **logically implies** an 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:
 - Father \sqsubseteq 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)

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

- DLs are nowadays advocated to provide the foundations for ontology languages.
- Different versions of the W3C standard **Ontology Web Language (OWL)** have been defined as syntactic variants of certain DLs.
- 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:
 - 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 Chap. 2: Description Logics and the DL-Lite family

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} .



DLs vs. UML Class Diagrams Chap. 2: Description Logics and the DL-Lite family

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:

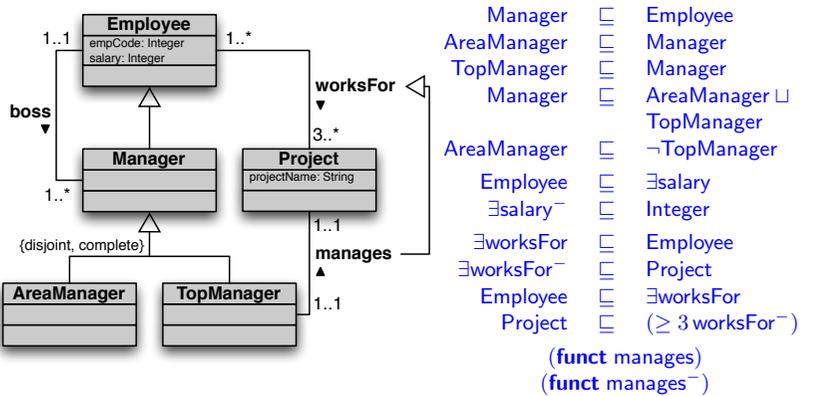
- Each class is represented by an atomic concept.
- Each attribute is represented by a role.
- Each binary association is represented by a role.
- Each non-binary association is reified, i.e., represented as a concept connected to its components by roles.
- Each part of the diagram is encoded by suitable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

We illustrate the encoding by means of an example.



DLs vs. UML Class Diagrams Chap. 2: Description Logics and the DL-Lite family

Encoding UML Class Diagrams in DLs – Example



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

- Manager ⊆ Employee
- AreaManager ⊆ Manager
- TopManager ⊆ Manager
- Manager ⊆ AreaManager ⊔ TopManager
- AreaManager ⊆ ∃salary
- Employee ⊆ ∃salary
- ∃worksFor ⊆ Employee
- ∃worksFor ⊆ Project
- Employee ⊆ ∃worksFor
- Project ⊆ (≥ 3 worksFor)
- (funct manages)
- (funct manages)
- manages ⊆ worksFor
- ...



DLs vs. UML Class Diagrams Chap. 2: Description Logics and the DL-Lite family

Encoding DL TBoxes in UML Class Diagrams

The encoding of an \mathcal{ALC} TBox \mathcal{T} in terms of a UML Class Diagram \mathcal{D} is based on the following observations:

- We can restrict the attention to \mathcal{ALC} TBoxes, that are constituted by concept inclusion assertions of a simplified form (single atomic concept on the left, and a single concept constructor on the right).
- For each such inclusion assertion, the encoding introduces a portion of UML Class Diagram, that may refer to some common classes.

Reasoning in the encoded \mathcal{ALC} -fragment is already **EXPTIME-hard**. From this, we obtain:

Theorem
Reasoning over UML Class Diagrams is EXPTIME-hard.



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., EXP-TIME-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

Traditionally, simple concept (or role) expressions have been considered as queries over DL ontologies.

We need more complex forms of queries, as those used in databases.

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

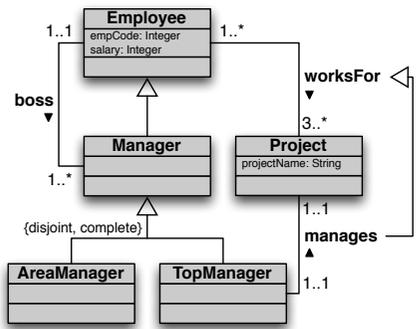
is a conjunctive query $q(\vec{x}) \leftarrow \vec{y}. conj(\vec{x}, \vec{y})$ where each atom in the body $conj(\vec{x}, \vec{y})$:

- 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} .

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



Chap. 2: Description Logics and the DL-Lite family
Queries over Description Logics ontologies – Example



Conjunctive query over the above ontology:

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

Chap. 2: Description Logics and the DL-Lite family
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 $q(\vec{x}) \leftarrow \exists \vec{y}. \text{conj}(\vec{x}, \vec{y})$ a CQ.

Def.: The **answer** to $q(\vec{x})$ over \mathcal{I} , denoted $q^{\mathcal{I}}$
 ... is the set of **tuples \vec{c} of constants of \mathcal{A}** such that the formula $\exists \vec{y}. \text{conj}(\vec{c}, \vec{y})$ evaluates to true in \mathcal{I} .

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

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

Chap. 2: Description Logics and the DL-Lite family
Query answering over ontologies

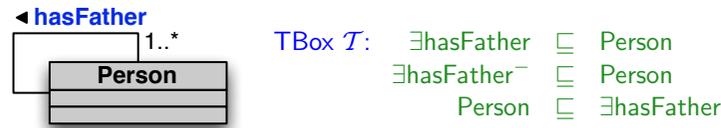
Def.: **Query answering** over an ontology \mathcal{O}
 Is the problem of **computing the certain answers** to a query over \mathcal{O} .

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: A special case of query answering is **instance checking**: 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).

Chap. 2: Description Logics and the DL-Lite family
Query answering over ontologies – Example



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 over an ontology.

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

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

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

Example

$$q(x) \leftarrow \text{Manager}(x), \text{worksFor}(x, \text{tones})$$

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

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:

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

- 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:

$$\begin{aligned} Cl &\longrightarrow A \mid \exists Q \\ Cr &\longrightarrow A \mid \exists Q \mid \neg A \mid \neg \exists Q \end{aligned}$$

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

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

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	$\neg Q$	$\neg \text{manages}$	$(\Delta^I \times \Delta^I) \setminus Q^I$
conc. incl.	$Cl \sqsubseteq Cr$	$\text{Father} \sqsubseteq \exists \text{child}$	$Cl^I \subseteq Cr^I$
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.	$P(c_1, c_2)$	$\text{child}(\text{bob}, \text{ann})$	$(c_1^I, c_2^I) \in P^I$

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: $A_1 \sqsubseteq \exists Q_{A_2}$
 $\exists Q_{A_2} \sqsubseteq A_2$

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 following, 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.



Chapter III

Linking ontologies to data



Outline

- 7 The impedance mismatch problem
- 8 Ontology-Based Data Access Systems
- 9 Query answering in Ontology-Based Data Access Systems



Outline

- 7 The impedance mismatch problem
- 8 Ontology-Based Data Access Systems
- 9 Query answering in Ontology-Based Data Access Systems



Managing ABoxes

In the traditional DL setting, it is 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.
 - The 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 an ABox becomes large, managing it 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 we 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 a 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.

Solution:

- We need to specify how to construct from the data values in the relational sources the (abstract) objects that populate the ABox of the ontology.
- This specification is embedded in the mappings between the data sources and the ontology.

Note: the **ABox** is only **virtual**, and the objects are not materialized.



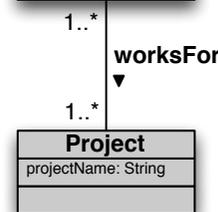
Solution to the impedance mismatch problem

We need to define a **mapping language** that allows for specifying how to transform data into abstract objects:

- Each mapping assertion maps:
 - a query that retrieves values from a data source to . . .
 - a set of atoms specified over the ontology.
- Basic idea: use **Skolem functions** in the atoms over the ontology to “generate” the objects from the data values.
- Semantics of mappings:
 - Objects are denoted by terms (of exactly one level of nesting).
 - Different terms denote different objects (i.e., we make the unique name assumption on terms).



Impedance mismatch – Example



1..* worksFor 1..*

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 from an alphabet Γ_V .
- To denote objects, we use **object terms** instead of object constants. An object term has the form $f(d_1, \dots, d_n)$, with $f \in \Lambda$, 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.

Def.: **Mapping assertion** between a database and a TBox

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

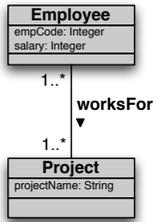
$$\Phi \rightsquigarrow \Psi$$

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 D1` \rightsquigarrow `Employee(pers(SSN)), Project(proj(PrName)), projectName(proj(PrName), PrName), worksFor(pers(SSN), proj(PrName))`

m_2 : `SELECT SSN, Salary FROM D2, D3 WHERE D2.Code = D3.Code` \rightsquigarrow `Employee(pers(SSN)), salary(pers(SSN), Salary)`



Outline

- 7 The impedance mismatch problem
- 8 **Ontology-Based Data Access Systems**
- 9 Query answering in Ontology-Based Data Access Systems



Ontology-Based Data Access System

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

Def.: 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} .

We need to specify the syntax and semantics of mapping assertions.



Mapping assertions

A mapping assertion in \mathcal{M} has the form

$$\Phi(\vec{x}) \rightsquigarrow \Psi(\vec{t}, \vec{y})$$

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**;
- \vec{x}, \vec{y} are variables, with $\vec{y} \subseteq \vec{x}$;
- \vec{t} are variable terms of the form $f(\vec{z})$, with $f \in \Lambda$ and $\vec{z} \subseteq \vec{x}$.

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



Semantics of mappings

To define the semantics of an OBDA system $\mathcal{O} = \langle \mathcal{T}, \mathcal{M}, \mathcal{D} \rangle$, we first need to define the semantics of mappings.

Def.: Satisfaction of a mapping assertion with respect to a database

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 $P(s_1, s_2)$, then $(s_1^{\mathcal{I}}, s_2^{\mathcal{I}}) \in P^{\mathcal{I}}$.

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

Def.: 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.



Outline

- 7 The impedance mismatch problem
- 8 Ontology-Based Data Access Systems
- 9 Query answering in Ontology-Based Data Access Systems



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.



Splitting of mappings

A mapping assertion $\Phi \rightsquigarrow \Psi$, where the TBox query Ψ is constituted by the atoms X_1, \dots, X_k , can be split into several mapping assertions:

$$\Phi \rightsquigarrow X_1 \quad \dots \quad \Phi \rightsquigarrow X_k$$

This is possible, since all variables in Ψ are distinguished.

Example

```
m1: SELECT SSN, PrName FROM D1 ~> Employee(pers(SSN)),
      Project(proj(PrName)),
      projectName(proj(PrName), PrName),
      worksFor(pers(SSN), proj(PrName))
```

is split into

```
m11: SELECT SSN, PrName FROM D1 ~> Employee(pers(SSN))
m12: SELECT SSN, PrName FROM D1 ~> Project(proj(PrName))
m13: SELECT SSN, PrName FROM D1 ~> projectName(proj(PrName), PrName)
m14: SELECT SSN, PrName FROM D1 ~> worksFor(pers(SSN), proj(PrName))
```

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 source updates (cf. Data Warehousing).



Top-down approach to query answering

Consists of three steps:

- 1 **Reformulation:** Compute the perfect reformulation $q_{pr} = \text{PerfectRef}(q, \mathcal{T}_P)$ of the original query q , using the inclusion assertions of the TBox \mathcal{T} (see later).
- 2 **Unfolding:** Compute from q_{pr} a new query q_{unf} by unfolding q_{pr} using (the split version of) the mappings \mathcal{M} .
 - Essentially, each atom in q_{pr} that unifies with an atom in Ψ is substituted with the corresponding query Φ over the database.
 - The unfolded query is such that $\text{Eval}(q_{unf}, \mathcal{D}) = \text{Eval}(q_{pr}, \mathcal{A}_{\mathcal{M},\mathcal{D}})$.
- 3 **Evaluation:** Delegate the evaluation of q_{unf} to the relational DBMS managing \mathcal{D} .



Unfolding

To unfold a query q_{pr} with respect to a set of mapping assertions:

- 1 For each non-split mapping assertion $\Phi_i(\vec{x}) \rightsquigarrow \Psi_i(\vec{t}, \vec{y})$:
 - 1 Introduce a **view symbol** Aux_i of arity equal to that of Φ_i .
 - 2 Add a **view definition** $\text{Aux}_i(\vec{x}) \leftarrow \Phi_i(\vec{x})$.
- 2 For each split version $\Phi_i(\vec{x}) \rightsquigarrow X_j(\vec{t}, \vec{y})$ of a mapping assertion, introduce a **clause** $X_j(\vec{t}, \vec{y}) \leftarrow \text{Aux}_i(\vec{x})$.
- 3 Obtain from q_{pr} in all possible ways queries q_{aux} defined over the view symbols Aux_i as follows:
 - 1 Find a most general unifier ϑ that unifies each atom $X(\vec{z})$ in the body of q_{pr} with the head of a clause $X(\vec{t}, \vec{y}) \leftarrow \text{Aux}_i(\vec{x})$.
 - 2 Substitute each atom $X(\vec{z})$ with $\vartheta(\text{Aux}_i(\vec{x}))$, i.e., with the body the unified clause to which the unifier ϑ is applied.
- 4 The unfolded query q_{unf} is the **union** of all queries q_{aux} , together with the view definitions for the predicates Aux_i appearing in q_{aux} .



Query reformulation

Consider the query $q(x) \leftarrow \text{Professor}(x)$

Intuition: Use the PIs as basic rewriting rules:

$\text{AssistantProf} \sqsubseteq \text{Professor}$
as a logic rule: $\text{Professor}(z) \leftarrow \text{AssistantProf}(z)$

Basic rewriting step:

when an atom in the query unifies with the **head** of the rule,
substitute the atom with the **body** of the rule.

Towards the computation of the perfect reformulation, we add to the input query above the query

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

We say that the PI $\text{AssistantProf} \sqsubseteq \text{Professor}$ applies to the atom $\text{Professor}(x)$.



Query reformulation (cont'd)

Similarly, 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

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

Consider now the query $q(x) \leftarrow \text{teaches}(x, y)$

and the PI $\text{Professor} \sqsubseteq \exists \text{teaches}$
as a logic rule: $\text{teaches}(z, f(z)) \leftarrow \text{Professor}(z)$

We add to the reformulation the query

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



Query reformulation – Constants

Conversely, for the query $q(x) \leftarrow \text{teaches}(x, \text{kdbb})$

and the same PI as before $\text{Professor} \sqsubseteq \exists \text{teaches}$
as a logic rule: $\text{teaches}(z, f(z)) \leftarrow \text{Professor}(z)$

$\text{teaches}(x, \text{kdbb})$ does not unify with $\text{teaches}(z, f(z))$, since the skolem term $f(z)$ in the head of the rule does not unify with the constant kdbb .

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

The same holds for the following query, where y is distinguished, since unifying $f(z)$ with y would correspond to returning a skolem term as answer to the query:

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



Query reformulation – Join variables

An analogous behavior with join variables. Consider the query

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

and the PI $\text{Professor} \sqsubseteq \exists \text{teaches}$
as a logic rule: $\text{teaches}(z, f(z)) \leftarrow \text{Professor}(z)$

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



Query reformulation – Reduce step

We now have the query

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

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

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

does not apply to $\text{teaches}(x, y)$ or $\text{teaches}(z, y)$, since y is in join.

However, we can transform the above query by **unifying** the atoms $\text{teaches}(x, y)$, $\text{teaches}(z, y)$. 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 reformulation – Summary

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

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

($-$ 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 rules applicable that were not so before.

The UCQ resulting from this process is the **perfect reformulation** $r_{q, \mathcal{T}}$



Query reformulation algorithm

Algorithm $PerfectRef(q, \mathcal{T}_P)$

Input: conjunctive query q , set of $DL-Lite_{\mathcal{R}}$ PIs \mathcal{T}_P

Output: union of conjunctive queries PR

$PR := \{q\};$

repeat

$PR' := PR;$

for each $q \in PR'$ **do**

for each g in q **do**

for each PI I in \mathcal{T}_P **do**

if I is applicable to g

then $PR := PR \cup \{q[g/(g, I)]\}$

for each g_1, g_2 in q **do**

if g_1 and g_2 unify

then $PR := PR \cup \{\tau(\text{reduce}(q, g_1, g_2))\};$

until $PR' = PR;$

return PR

Notice that NIs do not play any role in the reformulation of the query



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 $\langle c \rangle$ 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 $\langle c_1, c_2 \rangle$ such that $P(c_1, c_2) \in \mathcal{A}$

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



Query evaluation

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

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



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

Theorem

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

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 if NIs were not specified in \mathcal{T} .



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

TBox: $Professor \sqsubseteq \exists teaches$
 $\exists teaches^- \sqsubseteq Course$

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

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

ABox: $teaches(john, kbdb)$
 $Professor(mary)$

It is easy to see that $Eval(SQL(r_{q,\mathcal{T}}), DB(\mathcal{A}))$ in this case produces as answer $\{john, mary\}$.



Query answering in $DL-Lite_{\mathcal{R}}$ – An interesting example

TBox: $Person \sqsubseteq \exists hasFather$ ABox: $Person(mary)$
 $\exists hasFather^- \sqsubseteq Person$

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

$q(x) \leftarrow Person(x), hasFather(x, y_1), hasFather(y_1, y_2), hasFather(y_2, -)$
 \Downarrow Apply $Person \sqsubseteq \exists hasFather$ to the atom $hasFather(y_2, -)$

$q(x) \leftarrow Person(x), hasFather(x, y_1), hasFather(y_1, y_2), Person(y_2)$
 \Downarrow Apply $\exists hasFather^- \sqsubseteq Person$ to the atom $Person(y_2)$

$q(x) \leftarrow Person(x), hasFather(x, y_1), hasFather(y_1, y_2), hasFather(-, y_2)$
 \Downarrow Unify atoms $hasFather(y_1, y_2)$ and $hasFather(-, y_2)$

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

...

$q(x) \leftarrow Person(x), hasFather(x, -)$
 \Downarrow Apply $Person \sqsubseteq \exists hasFather$ to the atom $hasFather(x, -)$

$q(x) \leftarrow Person(x)$



Complexity of TBox reasoning

Theorem

TBox reasoning over both $DL-Lite_{\mathcal{R}}$ and $DL-Lite_{\mathcal{F}}$ ontologies is **PTIME** in the size of the **TBox** (schema complexity).

Proof (sketch).

Follows from the previous theorem, and from the reduction of TBox reasoning to ontology satisfiability. Indeed, the size of the ontology constructed in the reduction is polynomial in the size of the input TBox. □



Beyond $DL-Lite$

Can we further extend these results to more expressive ontology languages?

Essentially NO!
(unless we take particular care)



Beyond $DL-Lite$

We now consider DL languages that allow for constructs not present in $DL-Lite$ or for combinations of constructs that are not legal in $DL-Lite$.

We recall here syntax and semantics of constructs used in what follows.

Construct	Syntax	Example	Semantics
conjunction	$C_1 \sqcap C_2$	Doctor \sqcap Male	$C_1^I \cap C_2^I$
disjunction	$C_1 \sqcup C_2$	Doctor \sqcup Lawyer	$C_1^I \cup C_2^I$
qual. exist. restr.	$\exists Q.C$	$\exists \text{child.Male}$	$\{a \mid \exists b. (a, b) \in Q^I \wedge b \in C^I\}$
qual. univ. restr.	$\forall Q.C$	$\forall \text{child.Male}$	$\{a \mid \forall b. (a, b) \in Q^I \rightarrow b \in C^I\}$



Summary of results on data complexity

	C_l	C_r	\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	$A \mid \exists P \mid \exists P^-$	$A \mid \exists P \mid \exists P^-$	✓	✓	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.



TBox reasoning ○○○○○○○○ TBox & ABox reasoning ○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○ Complexity of reasoning in DLs ○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○ The Description Logic *DL-Lite_A* ●○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○ References ○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○

coNP-hard DLs Chap. 4: Reasoning in the DL-Lite family

Reduction from 2+2-SAT (cont'd)

Lemma

$\langle \mathcal{T}, A_\varphi \rangle \not\models q()$ iff φ is satisfiable.

Proof (sketch).

“ \Rightarrow ” If $\langle \mathcal{T}, A_\varphi \rangle \not\models q()$, then there is a model \mathcal{I} of $\langle \mathcal{T}, A_\varphi \rangle$ s.t. $\mathcal{I} \not\models q()$. We define a truth assignment $\alpha_{\mathcal{I}}$ by setting $\alpha_{\mathcal{I}}(v_i) = \text{true}$ iff $v_i^{\mathcal{I}} \in T^{\mathcal{I}}$. Notice that, since $L \sqsubseteq T \sqcup F$, if $v_i^{\mathcal{I}} \notin T^{\mathcal{I}}$, then $v_i^{\mathcal{I}} \in F^{\mathcal{I}}$. It is easy to see that, since $q()$ asks for a false clause and $\mathcal{I} \not\models q()$, for each clause c_j , one of the literals in c_j evaluates to *true* in $\alpha_{\mathcal{I}}$.

“ \Leftarrow ” From a truth assignment α that satisfies φ , we construct an interpretation \mathcal{I}_α with $\Delta^{\mathcal{I}_\alpha} = \{c_1, \dots, c_k, v_1, \dots, v_n, t, f\}$, and:

- $c_j^{\mathcal{I}_\alpha} = c_j$, $v_i^{\mathcal{I}_\alpha} = v_i$, $\text{true}^{\mathcal{I}_\alpha} = t$, $\text{false}^{\mathcal{I}_\alpha} = f$
- $T^{\mathcal{I}_\alpha} = \{v_i \mid \alpha(v_i) = \text{true}\} \cup \{t\}$, $F^{\mathcal{I}_\alpha} = \{v_i \mid \alpha(v_i) = \text{false}\} \cup \{f\}$

It is easy to see that \mathcal{I}_α is a model of $\langle \mathcal{T}, A_\varphi \rangle$ and that $\mathcal{I}_\alpha \not\models q()$. □

D. Calvanese Part 2: Ontology-Based Access to Inform. KBDB – 2007/2008 (189/215)

TBox reasoning ○○○○○○○○ TBox & ABox reasoning ○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○ Complexity of reasoning in DLs ○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○ The Description Logic *DL-Lite_A* ●○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○ References ○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○

Chap. 4: Reasoning in the DL-Lite family

Outline

- 10 TBox reasoning
- 11 TBox & ABox reasoning
- 12 Complexity of reasoning in Description Logics
- 13 The Description Logic *DL-Lite_A*
 - Missing features in *DL-Lite*
 - Combining functionality and role inclusions
 - Syntax and semantics of *DL-Lite_A*
 - Reasoning in *DL-Lite_A*
- 14 References

D. Calvanese Part 2: Ontology-Based Access to Inform. KBDB – 2007/2008 (190/215)

TBox reasoning ○○○○○○○○ TBox & ABox reasoning ○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○ Complexity of reasoning in DLs ○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○ The Description Logic *DL-Lite_A* ●○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○ References ○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○

Missing features in *DL-Lite* Chap. 4: Reasoning in the DL-Lite family

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.

D. Calvanese Part 2: Ontology-Based Access to Inform. KBDB – 2007/2008 (191/215)

TBox reasoning ○○○○○○○○ TBox & ABox reasoning ○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○ Complexity of reasoning in DLs ○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○ The Description Logic *DL-Lite_A* ●○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○ References ○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○

Combining functionality and role inclusions Chap. 4: Reasoning in the DL-Lite family

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

D. Calvanese Part 2: Ontology-Based Access to Inform. KBDB – 2007/2008 (192/215)

Restriction on TBox assertions in $DL-Lite_A$ ontologies

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

Restriction on $DL-Lite$ TBoxes

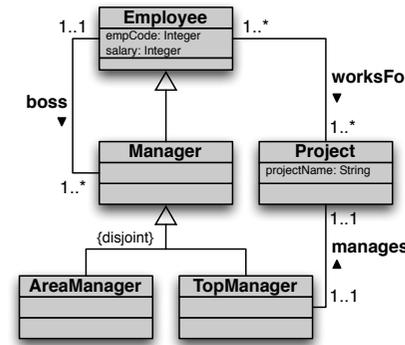
No functional role or attribute can be specialized by using it in the right-hand side of a role or attribute inclusion assertions.

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}** .



$DL-Lite_A$ – 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
- (funct salary)
- \exists worksFor \sqsubseteq Employee
- \exists worksFor $^-$ \sqsubseteq Project
- Employee \sqsubseteq \exists worksFor $^-$
- Project \sqsubseteq \exists worksFor $^-$
- (funct manages)
- (funct manages $^-$)
- manages \sqsubseteq worksFor
- ⋮

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



Reasoning in $DL-Lite_A$ – 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 $DL-Lite_A$ 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 $DL-Lite_A$ 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 $DL-Lite_A$

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 FOL query $q_F()$
- 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 $DL-Lite_A$ 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 $Eval(SQL(PerfectRef(q_N, \mathcal{T}_P)), DB(\mathcal{A}))$ returns true.
- There exists a functionality assertion $F \in \mathcal{T}$ such that $Eval(SQL(q_F), DB(\mathcal{A}))$ returns true.



