

ONTOLOGY-MEDIATED QUERY ANSWERING

Part 3: Extending the Applicability of OMQA Approach

Meghyn Bienvenu (*CNRS & Université de Bordeaux*)

Discuss some results along three directions:

Understanding limits / possibilities of query rewriting

OMQA with other query languages

Handling inconsistent KBs

CLOSER LOOK AT QUERY REWRITING

Many of the proposed rewriting algorithms produce **unions of conjunctive queries** (UCQs = \vee of CQs)

Many of the proposed rewriting algorithms produce **unions of conjunctive queries (UCQs = \vee of CQs)**

Not hard to see **smallest UCQ-rewriting may be exponentially large:**

- Query: $A_1^0(x) \wedge \dots \wedge A_n^0(x)$
- Ontology: $A_1^1 \sqsubseteq A_1^0 \quad A_2^1 \sqsubseteq A_2^0 \quad \dots \quad A_n^1 \sqsubseteq A_n^0$
- Rewriting: $\bigvee_{(i_1, \dots, i_n) \in \{0,1\}^n} A_1^{i_1}(x) \wedge A_2^{i_2}(x) \wedge \dots \wedge A_n^{i_n}(x)$

Many of the proposed rewriting algorithms produce **unions of conjunctive queries** (UCQs = \vee of CQs)

Not hard to see **smallest UCQ-rewriting may be exponentially large**:

- Query: $A_1^0(x) \wedge \dots \wedge A_n^0(x)$
- Ontology: $A_1^1 \sqsubseteq A_1^0 \quad A_2^1 \sqsubseteq A_2^0 \quad \dots \quad A_n^1 \sqsubseteq A_n^0$
- Rewriting: $\bigvee_{(i_1, \dots, i_n) \in \{0,1\}^n} A_1^{i_1}(x) \wedge A_2^{i_2}(x) \wedge \dots \wedge A_n^{i_n}(x)$

But: **simple polysize FO-rewriting does exist!** $\bigwedge_{i=1}^n (A_i^0(x) \vee A_i^1(x))$

Many of the proposed rewriting algorithms produce **unions of conjunctive queries (UCQs = \vee of CQs)**

Not hard to see **smallest UCQ-rewriting may be exponentially large:**

- Query: $A_1^0(x) \wedge \dots \wedge A_n^0(x)$
- Ontology: $A_1^1 \sqsubseteq A_1^0 \quad A_2^1 \sqsubseteq A_2^0 \quad \dots \quad A_n^1 \sqsubseteq A_n^0$
- Rewriting: $\bigvee_{(i_1, \dots, i_n) \in \{0,1\}^n} A_1^{i_1}(x) \wedge A_2^{i_2}(x) \wedge \dots \wedge A_n^{i_n}(x)$

But: **simple polysize FO-rewriting does exist!** $\bigwedge_{i=1}^n (A_i^0(x) \vee A_i^1(x))$

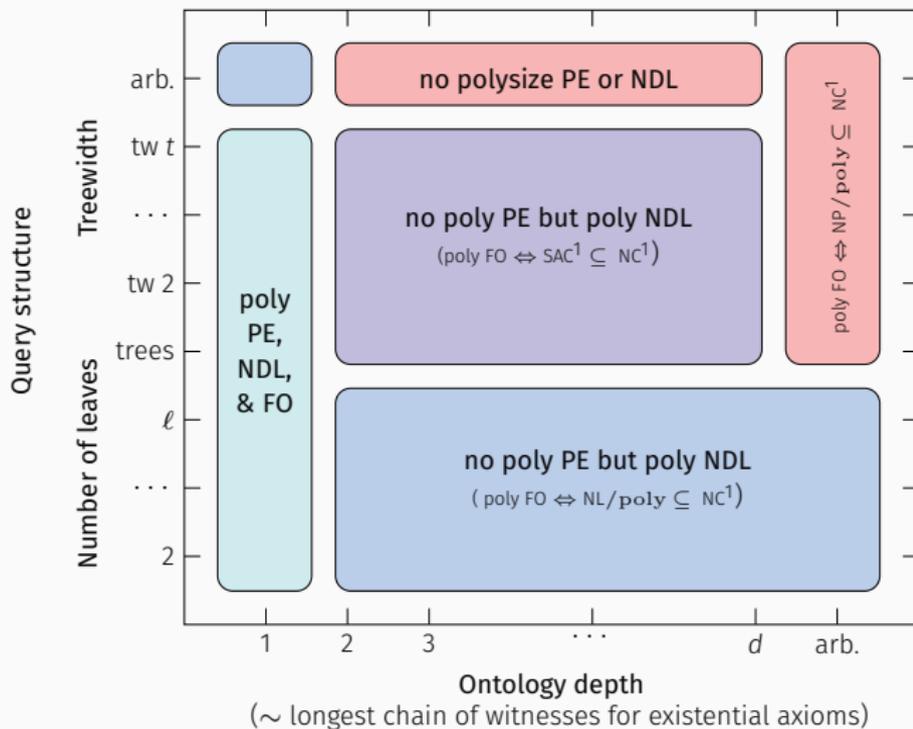
What happens if we adopt other representations?

- **positive existential queries (PE)**, **non-recursive Datalog (NDL)**,
first-order queries (FO)

SUCCINCTNESS LANDSCAPE FOR DL-LITE

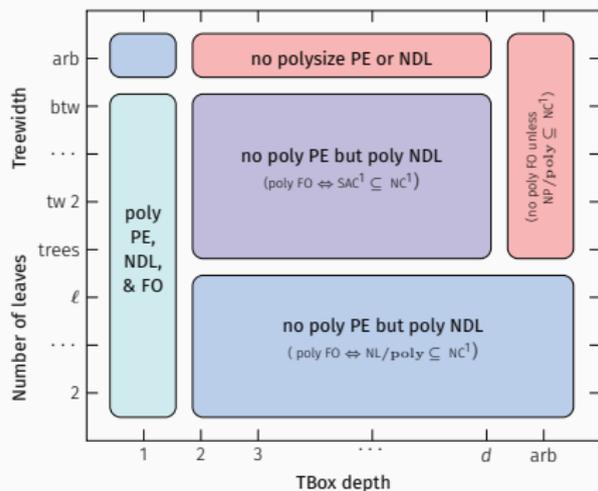
(for DL-Lite_R ontologies, so-called 'pure' rewritings)

no poly PE but poly NDL
 no poly PE or NDL

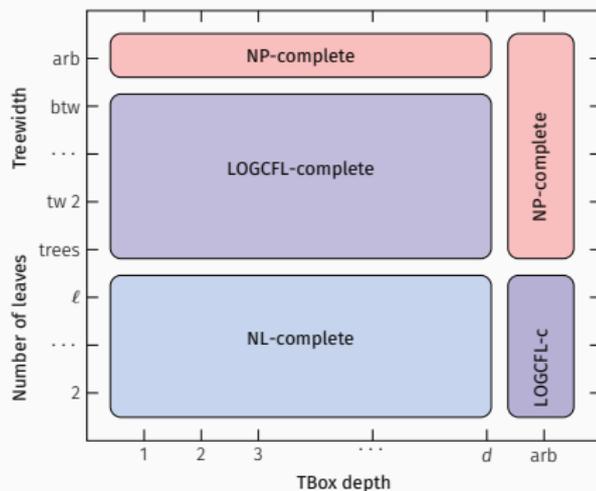


COMPARING SUCCINCTNESS & COMPLEXITY LANDSCAPES

Size of rewritings



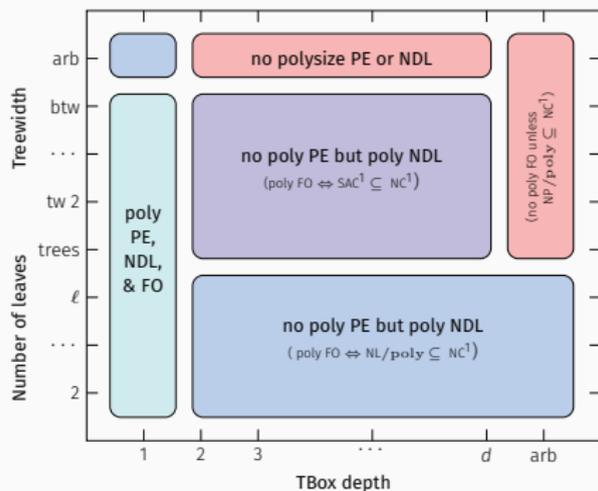
Combined complexity of OMQA



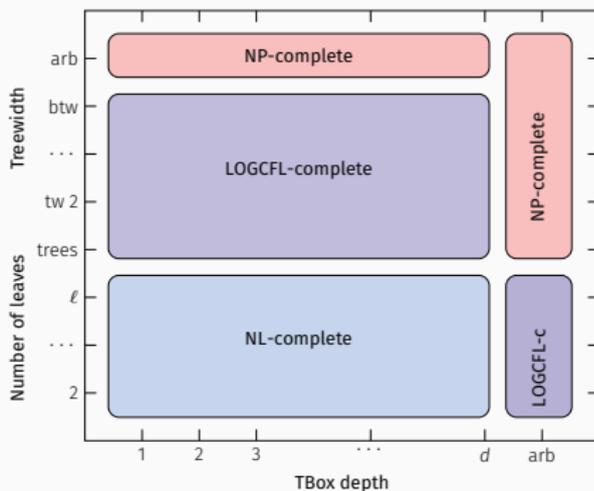
polysize NDL-rewritings \sim polynomial (LOGCFL / NL) complexity

COMPARING SUCCINCTNESS & COMPLEXITY LANDSCAPES

Size of rewritings



Combined complexity of OMQA



polysize NDL-rewritings \sim polynomial (LOGCFL / NL) complexity

Can we marry the positive succinctness & complexity results?

For the three well-behaved classes of OMQs, can define

NDL-rewritings of optimal complexity:

- rewriting can be constructed by L^C transducer
- evaluating the rewriting can be done in C

with $C \in \{\text{NL}, \text{LOGCFL}\}$ the complexity of the OMQ class

OPTIMAL NDL-REWRITINGS

For the three well-behaved classes of OMQs, can define

NDL-rewritings of optimal complexity:

- rewriting can be constructed by L^C transducer
- evaluating the rewriting can be done in C

with $C \in \{\text{NL}, \text{LOGCFL}\}$ the complexity of the OMQ class

Upper bound on **time needed to evaluate our NDL-rewritings:**

- **depth d / number of leaves ℓ occur in the exponent**

For the three well-behaved classes of OMQs, can define

NDL-rewritings of optimal complexity:

- rewriting can be constructed by L^C transducer
- evaluating the rewriting can be done in C

with $C \in \{\text{NL}, \text{LOGCFL}\}$ the complexity of the OMQ class

Upper bound on **time needed to evaluate our NDL-rewritings:**

- **depth d / number of leaves ℓ occur in the exponent**

Is it **possible to do better?**

- formally: **fixed-parameter tractable (FPT)?** $f(d, \ell) \cdot p(|q|, |\mathcal{T}|, |\mathcal{A}|)$

For the three well-behaved classes of OMQs, can define

NDL-rewritings of optimal complexity:

- rewriting can be constructed by L^C transducer
- evaluating the rewriting can be done in C

with $C \in \{\text{NL}, \text{LOGCFL}\}$ the complexity of the OMQ class

Upper bound on **time needed to evaluate our NDL-rewritings:**

- depth d / number of leaves ℓ occur in the exponent

Is it **possible to do better?**

- formally: **fixed-parameter tractable (FPT)?** $f(d, \ell) \cdot p(|q|, |\mathcal{T}|, |\mathcal{A}|)$

Parameterized complexity of answering tree-shaped OMQs (\mathcal{T}, q) :

- parameters: depth d of \mathcal{T} , number ℓ of leaves in CQs
- **not FPT** if **depth d** taken as parameter W[2]-hard
- **not FPT** if number of **leaves ℓ** taken as parameter W[1]-hard

Succinctness and combined complexity landscapes:

Bienvenu, Kikot, Kontchakov, Podolskii, and Zakharyashev:
[Ontology-Mediated Queries: Combined Complexity and Succinctness of Rewritings via Circuit Complexity](#). Journal of the ACM (JACM), 2018.

Earlier work on succinctness, also covering 'impure' rewritings:

Gottlob, Kikot, Kontchakov, Podolskii, Schwentick, and Zakharyashev:
[The Price of Query Rewriting in Ontology-based Data Access](#). Artificial Intelligence (AIJ), 2014.

Optimal rewritings and parameterized complexity results:

Bienvenu, Kikot, Kontchakov, Podolskii, Ryzhikov and Zakharyashev:
[The Complexity of Ontology-Based Data Access with OWL 2 QL and Bounded Treewidth Queries](#). PODS 2017.

We have seen that:

- for \mathcal{EL} ontologies, **FO-rewritings need not exist**
- for \mathcal{ALC} ontologies, **FO- and Datalog rewritings may not exist**

But these are **worst-case results**

- only say that some OMQ that does not have a rewriting
- **possible** that **rewritings exist** for many **OMQs encountered in practice**

To **extend the applicability of query rewriting** beyond DL-Lite:

- devise **ways of identifying 'good cases'**
- **construct rewritings** when they exist

Use $(\mathcal{L}, \mathcal{Q})$ to denote set of **ontology-mediated queries** (\mathcal{T}, q) where:

- \mathcal{T} is an \mathcal{L} -TBox
- q is a query from \mathcal{Q} $\mathcal{Q} \in \{\text{IQ}, \text{CQ}\}$

For example: $(\mathcal{EL}, \text{CQ})$, $(\mathcal{ALC}, \text{IQ})$

FO-rewritability in $(\mathcal{L}, \mathcal{Q})$

- Input: OMQ (\mathcal{T}, q) from $(\mathcal{L}, \mathcal{Q})$
- Problem: **decide whether (\mathcal{T}, q) has an FO-rewriting**

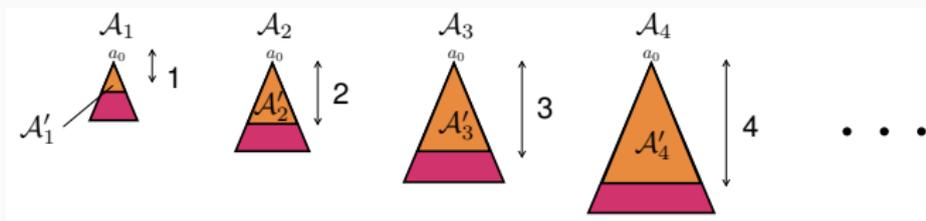
Datalog-rewritability decision problem can be **defined analogously**

FO-rewritability is **EXPTIME-complete** in $(\mathcal{EL}, \text{IQ})$ and $(\mathcal{EL}, \text{CQ})$

FO-rewritability is **EXPTIME-complete** in $(\mathcal{EL}, \text{IQ})$ and $(\mathcal{EL}, \text{CQ})$

Characterization of non-existence of FO-rewriting

OMQ $(\mathcal{T}, A(x))$ is not FO-rewritable iff there exist tree-shaped ABoxes

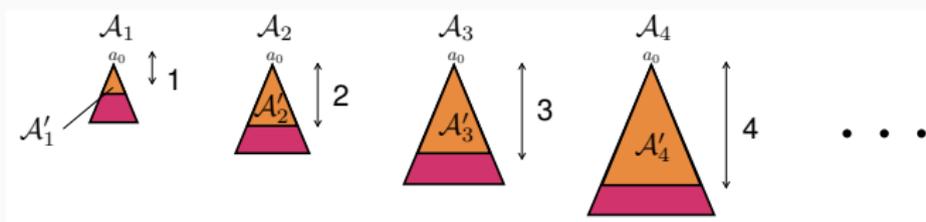


such that for all $i \geq 1$: $\mathcal{T}, \mathcal{A}_i \models A(a_0)$ and $\mathcal{T}, \mathcal{A}'_i \not\models A(a_0)$

FO-rewritability is **EXPTIME-complete** in $(\mathcal{EL}, \text{IQ})$ and $(\mathcal{EL}, \text{CQ})$

Characterization of non-existence of FO-rewriting

OMQ $(\mathcal{T}, A(x))$ is **not FO-rewritable** iff there exist tree-shaped ABoxes



such that for all $i \geq 1$: $\mathcal{T}, \mathcal{A}_i \models A(a_0)$ and $\mathcal{T}, \mathcal{A}'_i \not\models A(a_0)$

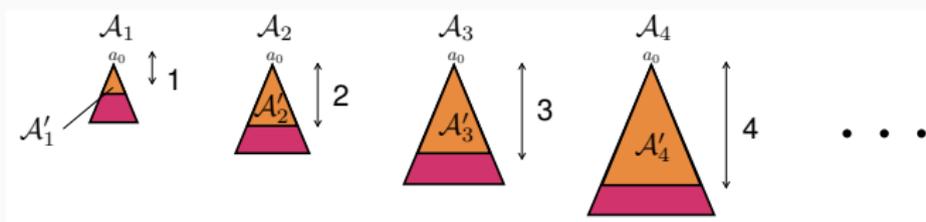
Pumping argument: enough to find ABox of **particular finite size k_0**

- desired ABox \mathcal{A}_{k_0} exists \Rightarrow can construct full sequence of ABoxes

FO-rewritability is **EXPTIME-complete** in $(\mathcal{EL}, \text{IQ})$ and $(\mathcal{EL}, \text{CQ})$

Characterization of non-existence of FO-rewriting

OMQ $(\mathcal{T}, A(x))$ is **not FO-rewritable** iff there exist tree-shaped ABoxes



such that for all $i \geq 1$: $\mathcal{T}, \mathcal{A}_i \models A(a_0)$ and $\mathcal{T}, \mathcal{A}_i' \not\models A(a_0)$

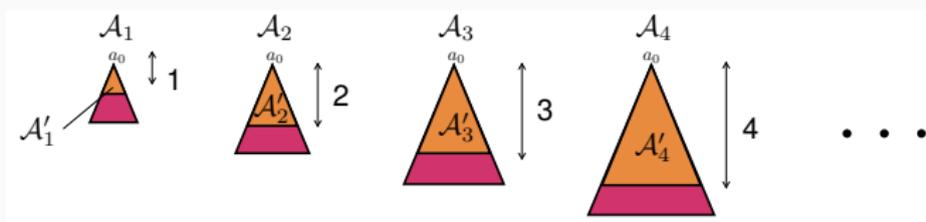
Pumping argument: enough to find ABox of **particular finite size k_0**
 · desired ABox \mathcal{A}_{k_0} exists \Rightarrow can construct full sequence of ABoxes

Use **tree automata** to check whether **such a witness ABox exists**

FO-rewritability is **EXPTIME-complete** in $(\mathcal{EL}, \text{IQ})$ and $(\mathcal{EL}, \text{CQ})$

Characterization of non-existence of FO-rewriting

OMQ $(\mathcal{T}, A(x))$ is **not FO-rewritable** iff there exist tree-shaped ABoxes



such that for all $i \geq 1$: $\mathcal{T}, \mathcal{A}_i \models A(a_0)$ and $\mathcal{T}, \mathcal{A}'_i \not\models A(a_0)$

Pumping argument: enough to find ABox of **particular finite size** k_0
 · desired ABox \mathcal{A}_{k_0} exists \Rightarrow can **construct full sequence** of ABoxes

Use **tree automata** to check whether **such a witness ABox exists**

Can **generalize this technique** to handle CQs as well

Idea for IQs: use **existing backwards-chaining rewriting procedure**

- if **FO-rewriting** does exist, **terminates**
- to ensure **termination in general**: use **characterization result**

To make practical: **decomposed algorithm**

- allows for **structure sharing**
- produces **(succinct) NDL-rewriting** instead of UCQ-rewriting

Experimental results are **very encouraging**:

- **terminates quickly**, produced **rewritings** are **typically small**
- suggests that in practice **FO-rewritings do exist** for **majority of IQs**

Recently **extended to handle CQs** with promising results

FO-rewritability and Datalog-rewritability of $(\mathcal{ALC}, \text{IQ})$ are both NEXPTIME-complete.

FO-rewritability and **Datalog-rewritability** of $(\mathcal{ALC}, \text{IQ})$ are both **NEXPTIME-complete**.

Upper bound: connection to **constraint satisfaction problems (CSPs)**

- **CSP(\mathfrak{B})**: decide if **homomorphism** from input structure \mathcal{D} into \mathfrak{B}
- **(Boolean) OMQs in $(\mathcal{ALC}, \text{IQ}) \sim$ (complement of) CSPs**
- **exponential reduction** to problem of **deciding whether a CSP is definable in FO / Datalog**
- use **NP upper bounds** for latter problems

FO-rewritability of $(\mathcal{ALC}, \text{UCQ})$: **2NEXPTIME-complete**

FO-rewritability of instance queries

Bienvenu, Lutz, Wolter: [First Order-Rewritability of Atomic Queries in Horn Description Logics](#). IJCAI 2013.

FO-rewritability of conjunctive queries

Bienvenu, Hansen, Lutz, Wolter: [First Order-Rewritability and Containment of Conjunctive Queries in Horn Description Logics](#). IJCAI 2016.

Practical algorithms and implementations

Hansen, Lutz, Seylan, Wolter: [Efficient Query Rewriting in the Description Logic EL and Beyond](#). IJCAI 2015.

Hansen & Lutz: [Computing FO-Rewritings in \$\mathcal{EL}\$ in Practice: from Atomic to Conjunctive Queries](#). ISWC 2017.

CSP connection & FO- and Datalog-rewritability of IQs

Bienvenu, ten Cate, Lutz, and Wolter: [Ontology-based Data Access: A Study through Disjunctive Datalog, CSP, and MMSNP](#). TODS, 2014.

FO-rewritability of UCQs

Feier, Kuusisto, and Lutz: [Rewritability in Monadic Disjunctive Datalog, MMSNP, and Expressive Description Logics](#). ICDT 2017.

Practical, incomplete approach for Datalog rewritings

Kaminski, Nenov, and Cuenca Grau: [Datalog Rewritability of Disjunctive Datalog Programs and its Applications to Ontology Reasoning](#). AAAI 2014.

BEYOND CONJUNCTIVE QUERIES

Conjunctive query with safe negation ($CQ^{\neg S}$):

- like a CQ, but **can also have negated atoms**
- **safety condition: every variable occurs in some positive atom**
- example: find **menus whose main course is not spicy**

$\exists y \text{Menu}(x) \wedge \text{hasMain}(x, y) \wedge \neg \text{Spicy}(y)$

Conjunctive query with safe negation ($CQ^{\neg s}$):

- like a CQ, but **can also have negated atoms**
- **safety condition: every variable occurs in some positive atom**
- example: find **menus whose main course is not spicy**

$\exists y \text{Menu}(x) \wedge \text{hasMain}(x, y) \wedge \neg \text{Spicy}(y)$

Conjunctive query with inequalities (CQ^{\neq})

- like a CQ, but **can also have atoms $t_1 \neq t_2$** (t_1, t_2 vars or individuals)
- example: find **restaurant offering two menus having different dessert courses**

$\exists y_1 y_2 z_1 z_2$ $\text{offers}(x, y_1) \wedge \text{Menu}(y_1) \wedge \text{hasDessert}(y_1, z_1) \wedge$
 $\text{offers}(x, y_2) \wedge \text{Menu}(y_2) \wedge \text{hasDessert}(y_2, z_2) \wedge z_1 \neq z_2$

Conjunctive query with safe negation ($CQ^{\neg s}$):

- like a CQ, but **can also have negated atoms**
- **safety condition: every variable occurs in some positive atom**
- example: find **menus whose main course is not spicy**

$$\exists y \text{Menu}(x) \wedge \text{hasMain}(x, y) \wedge \neg \text{Spicy}(y)$$

Conjunctive query with inequalities (CQ^{\neq})

- like a CQ, but **can also have atoms $t_1 \neq t_2$** (t_1, t_2 vars or individuals)
- example: find **restaurant offering two menus having different dessert courses**

$$\begin{aligned} \exists y_1 y_2 z_1 z_2 \quad & \text{offers}(x, y_1) \wedge \text{Menu}(y_1) \wedge \text{hasDessert}(y_1, z_1) \wedge \\ & \text{offers}(x, y_2) \wedge \text{Menu}(y_2) \wedge \text{hasDessert}(y_2, z_2) \wedge z_1 \neq z_2 \end{aligned}$$

Note: can define $UCQ^{\neg s}$ s and UCQ^{\neq} s in the obvious way

Adding **negation** leads to undecidability even in very restricted settings.

Theorem The following problems are **undecidable**:

- CQ^{\neg} answering in DL-Lite_R
- UCQ^{\neg} answering in \mathcal{EL}_{\perp}
- CQ^{\neq} answering in DL-Lite_R
- CQ^{\neq} answering in \mathcal{EL}_{\perp}

Adding **negation leads to undecidability even in very restricted settings.**

Theorem The following problems are **undecidable**:

- $CQ^{\neg S}$ answering in $DL\text{-Lite}_R$
- $UCQ^{\neg S}$ answering in \mathcal{EL}_{\perp}
- CQ^{\neq} answering in $DL\text{-Lite}_R$
- CQ^{\neq} answering in \mathcal{EL}_{\perp}

Possible solution: adopt alternative (epistemic) semantics

Significant interest in combining DLs with Datalog rules

Unfortunately, this **almost always leads to undecidability**:

Theorem **Datalog query answering is undecidable** in every DL that **can express** (directly or indirectly) $A \sqsubseteq \exists r.A$

In particular: **undecidable in both DL-Lite and \mathcal{EL}**

Significant interest in combining DLs with Datalog rules

Unfortunately, this almost always leads to undecidability:

Theorem Datalog query answering is undecidable in every DL that can express (directly or indirectly) $A \sqsubseteq \exists r.A$

In particular: undecidable in both DL-Lite and \mathcal{EL}

Possible solutions:

- use restricted classes of Datalog queries (e.g. path queries)
- DL-safe rules: can only apply rules to (named) individuals

Prominent navigational query languages (graph databases):

Prominent navigational query languages (graph databases):

Regular Path Queries (RPQs): find pairs of objects that are connected by a chain of roles that comply with a given regular language

$$(\text{hasCourse} \cup \text{courseOf}^-) \cdot (\text{hasIngred} \cup \text{ingredOf}^-)^* \cdot \text{Spicy?}(x, y)$$

Prominent navigational query languages (graph databases):

Regular Path Queries (RPQs): find pairs of objects that are connected by a chain of roles that comply with a given regular language

$$(\text{hasCourse} \cup \text{courseOf}^-) \cdot (\text{hasIngred} \cup \text{ingredOf}^-)^* \cdot \text{Spicy?}(x, y)$$

Conjunctive RPQs: allow to join RPQs conjunctively

- similar to CQs, but each atom is an RPQ
- extend CQs with the navigational power of RPQs

Both languages have **1-way and 2-way variants**

COMPLEXITY LANDSCAPE FOR PATH QUERIES

	2RPQ		C2RPQ	
	data	combined	data	combined
Graph DBs & RDFS	NL-c	NL-c	NL-c	NP-c
DL-Lite	NL-c	PTIME-c	NL-c	PSPACE-c
\mathcal{EL} ... Horn- <i>SHIQ</i>	PTIME-c	PTIME-c	PTIME-c	PSPACE-c
\mathcal{ALC} ... <i>SHIQ</i>	coNP-h	EXP-c	coNP-h	2EXP-c

Remarks:

- **positive results for Horn DLs**: low complexity for 2RPQs, remains PTIME data for C2RPQs
- **tractable data complexity**: shown by adapting CQ answering approach presented in this course

Aggregate queries (COUNT, SUM, AVG) very common in databases, but not well explored for OMQA

Counting conjunctive queries: how many ways to map the CQ?

Cardinality queries: how many A ? how many r ?

Aggregate queries (COUNT, SUM, AVG) very common in databases, but not well explored for OMQA

Counting conjunctive queries: how many ways to map the CQ?

Cardinality queries: how many A ? how many r ?

Challenging to handle:

- coNP-c data complexity even for cardinality queries in DL-Lite
- 2EXPTIME-c in combined complexity (counting CQs + DL-Lite_R / \mathcal{EL})
- cannot rely upon canonical model to get optimal value

Some tractable / lower complexity cases have been identified

Queries with Inequalities or Negation

Gutiérrez-Basulto, Ibáñez-García, Kontchakov, Kostylev: [Queries with Negation and Inequalities over Lightweight Ontologies](#). Journal of Web Semantics, 2015.

Regular Path Queries

Bienvenu, Ortiz, Simkus: [Regular Path Queries in Lightweight Description Logics: Complexity and Algorithms](#). JAIR, 2015.

Counting Queries

Bienvenu, Manière, Thomazo: [Answering Counting Queries over DL-Lite Ontologies](#). IJCAI 2020.

HANDLING INCONSISTENT KBS

In realistic settings, can expect some **errors in the data**

- ABox likely to be **inconsistent** with the TBox (ontology)

Standard semantics: everything is implied - **not informative!**

- when \mathcal{K} unsatisfiable, $\text{cert}(q, \mathcal{K})$ contains all possible tuples

In realistic settings, can expect some **errors in the data**

- ABox likely to be **inconsistent** with the TBox (ontology)

Standard semantics: everything is implied - **not informative!**

- when \mathcal{K} unsatisfiable, $\text{cert}(q, \mathcal{K})$ contains all possible tuples

Two approaches to inconsistency handling:

- **resolve the inconsistencies**
 - preferable, but not always applicable!
- live with the inconsistencies - **adopt alternative semantics**
 - **meaningful answers** to queries despite inconsistencies

Note: focus on case where errors in ABox (**assume TBox reliable**)

EXAMPLE: REASONABLE INFERENCES

TBox $\mathcal{T}_{\text{univ}}$:

$\text{Prof} \sqsubseteq \text{Fac}$	$\text{Prof} \sqsubseteq \exists \text{teaches}$	$\text{Prof} \sqsubseteq \neg \text{Lect}$	$\text{Fac} \sqsubseteq \neg \text{Course}$
$\text{Lect} \sqsubseteq \text{Fac}$	$\text{Lect} \sqsubseteq \exists \text{teaches}$	$\text{Prof} \sqsubseteq \neg \text{Fellow}$	
$\text{Fellow} \sqsubseteq \text{Fac}$	$\exists \text{teaches}^- \sqsubseteq \text{Course}$	$\text{Lect} \sqsubseteq \neg \text{Fellow}$	

Consider following ABoxes:

\mathcal{A}_1	=	$\{\text{Prof}(\text{anna}), \text{Lect}(\text{anna}), \text{Fellow}(\text{alex})\}$
\mathcal{A}_2	=	$\{\text{Prof}(\text{anna}), \text{Fellow}(\text{alex}), \text{Lect}(\text{alex})\}$

Which assertions would be reasonable to infer from these two KBs?

$\text{Prof}(\text{anna})$	$\text{Lect}(\text{anna})$	$\text{Fac}(\text{anna})$
$\text{Fellow}(\text{alex})$	$\text{Lect}(\text{alex})$	$\text{Fac}(\text{alex})$

Many proposed semantics are based upon the notion of repair

Repair of knowledge base $(\mathcal{T}, \mathcal{A})$

= inclusion-maximal subset of \mathcal{A} s.t. $(\mathcal{T}, \mathcal{A})$ satisfiable

Intuition: different **ways of achieving consistency** while **retaining as much of the original data as possible**

Many proposed semantics are based upon the notion of repair

Repair of knowledge base $(\mathcal{T}, \mathcal{A})$

= inclusion-maximal subset of \mathcal{A} s.t. $(\mathcal{T}, \mathcal{A})$ satisfiable

Intuition: different **ways of achieving consistency** while **retaining as much of the original data as possible**

For example, the following knowledge base $\mathcal{K}^* = \langle \mathcal{T}^*, \mathcal{A}^* \rangle$ with

$$\mathcal{T}^* = \{\text{Prof} \sqsubseteq \text{PhD}, \text{Postdoc} \sqsubseteq \text{PhD}, \text{Prof} \sqsubseteq \neg \text{Postdoc}\}$$

$$\mathcal{A}^* = \{\text{Postdoc}(\text{ann}), \text{Prof}(\text{ann}), \text{Teach}(\text{ann}, c)\}$$

has two repairs:

$$\mathcal{R}_1 = \{\text{Postdoc}(\text{ann}), \text{Teach}(\text{ann}, c)\}$$

$$\mathcal{R}_2 = \{\text{Prof}(\text{ann}), \text{Teach}(\text{ann}, c)\}$$

AR semantics ('AR' for ABox Repair):

query each repair separately, intersect results

$$\mathcal{T}, \mathcal{A} \models_{\text{AR}} q(\vec{a}) \Leftrightarrow \mathcal{T}, \mathcal{B} \models q(\vec{a}) \text{ for every repair } \mathcal{B} \in \text{Rep}(\mathcal{K})$$

Plausible answers: hold no matter which repair is chosen

Generally viewed as natural semantics:

- cautious inference (KR), **consistent query answering** (DBs)

AR semantics ('AR' for ABox Repair):

query each repair separately, intersect results

$$\mathcal{T}, \mathcal{A} \models_{\text{AR}} q(\vec{a}) \Leftrightarrow \mathcal{T}, \mathcal{B} \models q(\vec{a}) \text{ for every repair } \mathcal{B} \in \text{Rep}(\mathcal{K})$$

Plausible answers: hold no matter which repair is chosen

Generally viewed as natural semantics:

- cautious inference (KR), **consistent query answering** (DBs)

On our example KB:

- $\mathcal{K}^* \not\models_{\text{AR}} \text{Postdoc}(ann)$
- $\mathcal{K}^* \models_{\text{AR}} \text{PhD}(ann)$
- $\mathcal{K}^* \models_{\text{AR}} \text{Teach}(ann, c)$

$$\begin{aligned} \langle \mathcal{T}^*, \mathcal{R}_2 \rangle &\not\models \text{Postdoc}(ann) \\ \text{Postdoc}(ann) \in \mathcal{R}_1 &\ \& \ \text{Prof}(ann) \in \mathcal{R}_2 \\ \text{Teach}(ann, c) &\in \mathcal{R}_1 \cap \mathcal{R}_2 \end{aligned}$$

AR semantics ('AR' for ABox Repair):

query each repair separately, intersect results

$$\mathcal{T}, \mathcal{A} \models_{\text{AR}} q(\vec{a}) \Leftrightarrow \mathcal{T}, \mathcal{B} \models q(\vec{a}) \text{ for every repair } \mathcal{B} \in \text{Rep}(\mathcal{K})$$

Plausible answers: hold no matter which repair is chosen

Generally viewed as natural semantics:

- cautious inference (KR), **consistent query answering** (DBs)

On our example KB:

- $\mathcal{K}^* \not\models_{\text{AR}} \text{Postdoc}(ann)$ $\langle \mathcal{T}^*, \mathcal{R}_2 \rangle \not\models \text{Postdoc}(ann)$
- $\mathcal{K}^* \models_{\text{AR}} \text{PhD}(ann)$ $\text{Postdoc}(ann) \in \mathcal{R}_1$ & $\text{Prof}(ann) \in \mathcal{R}_2$
- $\mathcal{K}^* \models_{\text{AR}} \text{Teach}(ann, c)$ $\text{Teach}(ann, c) \in \mathcal{R}_1 \cap \mathcal{R}_2$

Bad news: querying under AR semantics is **intractable**

- **coNP-hard** in data complexity, **even in simple settings** (DL-Lite)

Brave semantics: answers that hold in **at least one repair**

$$\mathcal{T}, \mathcal{A} \models_{\text{brave}} q(\vec{a}) \Leftrightarrow \mathcal{T}, \mathcal{R} \models q(\vec{a}) \text{ for some repair } \mathcal{R} \in \text{Rep}(\mathcal{K})$$

Answers **supported by consistent part of data** \rightsquigarrow **possible answers**

More permissive than AR semantics: $\mathcal{K}^* \models_{\text{brave}} \text{Postdoc}(ann)$

Brave semantics: answers that hold in **at least one repair**

$$\mathcal{T}, \mathcal{A} \models_{\text{brave}} q(\vec{a}) \Leftrightarrow \mathcal{T}, \mathcal{R} \models q(\vec{a}) \text{ for some repair } \mathcal{R} \in \text{Rep}(\mathcal{K})$$

Answers **supported by consistent part of data** \rightsquigarrow **possible answers**

More permissive than AR semantics: $\mathcal{K}^* \models_{\text{brave}} \text{Postdoc}(ann)$

IAR semantics: query the **intersection of the repairs**

$$\mathcal{T}, \mathcal{A} \models_{\text{IAR}} q(\vec{a}) \Leftrightarrow \mathcal{T}, \mathcal{D} \models q(\vec{a}) \text{ where } \mathcal{D} = \bigcap_{\mathcal{R} \in \text{Rep}(\mathcal{K})} \mathcal{R}$$

Idea: **focus on surest facts, not involved in any contradiction**

Stricter than AR semantics: $\mathcal{K}^* \not\models_{\text{IAR}} \text{PhD}(ann)$

Brave semantics: answers that hold in **at least one repair**

$$\mathcal{T}, \mathcal{A} \models_{\text{brave}} q(\vec{a}) \Leftrightarrow \mathcal{T}, \mathcal{R} \models q(\vec{a}) \text{ for some repair } \mathcal{R} \in \text{Rep}(\mathcal{K})$$

Answers **supported by consistent part of data** \rightsquigarrow **possible answers**

More permissive than AR semantics: $\mathcal{K}^* \models_{\text{brave}} \text{Postdoc}(ann)$

IAR semantics: query the **intersection of the repairs**

$$\mathcal{T}, \mathcal{A} \models_{\text{IAR}} q(\vec{a}) \Leftrightarrow \mathcal{T}, \mathcal{D} \models q(\vec{a}) \text{ where } \mathcal{D} = \bigcap_{\mathcal{R} \in \text{Rep}(\mathcal{K})} \mathcal{R}$$

Idea: **focus on surest facts, not involved in any contradiction**

Stricter than AR semantics: $\mathcal{K}^* \not\models_{\text{IAR}} \text{PhD}(ann)$

Good news: brave and IAR semantics **tractable for DL-Lite**

- still possible to use **query rewriting**

Introductory chapter

Bienvenu, Bourgaux: [Inconsistency-Tolerant Querying of Description Logic Knowledge Bases](#). Lecture Notes of the 12th International Reasoning Web Summer School, 2016.

Recent survey article

Bienvenu: [A Short Survey on Inconsistency Handling in Ontology-Mediated Query Answering](#). Künstliche Intelligenz, 2020.

SAT-based approach to AR semantics

Bienvenu, Bourgaux, Goasdoué: [Computing and Explaining Query Answers over Inconsistent DL-Lite Knowledge Bases](#). JAIR, 2019.

OMQA RESEARCH

Ontology-mediated query answering:

- new paradigm for intelligent information systems
- offers many **advantages**, but also **computational challenges**
- active area with **lots left to explore!**

Ontology-mediated query answering:

- new paradigm for intelligent information systems
- offers many **advantages**, but also **computational challenges**
- active area with **lots left to explore!**

Efficient OMQA algorithms:

- optimized rewriting algorithms: **compact rewritings**, exploit **mapping structure**, **cost-based rewriting selection**
- tackling more expressive DLs: **identify easier cases** (existence of rewritings), use upper + lower **approximations**

Ontology-mediated query answering:

- new paradigm for intelligent information systems
- offers many **advantages**, but also **computational challenges**
- active area with **lots left to explore!**

Efficient OMQA algorithms:

- optimized rewriting algorithms: **compact rewritings**, exploit **mapping structure**, **cost-based rewriting selection**
- tackling more expressive DLs: **identify easier cases** (existence of rewritings), use upper + lower **approximations**

Support for **building and maintaining** OMQA systems

- ontology + mapping **bootstrapping**, **module extraction**, **debugging**, ontology **evolution** and **versioning**
- **learning ontologies, queries, mappings**
- inspired **new reasoning tasks**: query inseparability, query emptiness, justification finding, logical difference, ...

Improving the **usability** of OMQA systems

- **interfaces** and support for **query formulation**
- **explaining** query (non-)answers

Improving the **usability** of OMQA systems

- **interfaces** and support for **query formulation**
- **explaining** query (non-)answers

Broadening the scope:

- **new data formats:** graph data, key-value stores, temporal data
- **further query languages:** regular path queries, streaming queries

Improving the **usability** of OMQA systems

- **interfaces** and support for **query formulation**
- **explaining** query (non-)answers

Broadening the scope:

- **new data formats:** graph data, key-value stores, temporal data
- **further query languages:** regular path queries, streaming queries

Beyond classical OMQA:

- **inconsistency-tolerant** query answering
- **probabilistic** query answering
- **privacy-aware** query answering