

Deduction in Existential Conjunctive First-Order Logic: An Algorithm and Experiments

Khalil Ben Mohamed, Michel Leclère, and Marie-Laure Mugnier

LIRMM, CNRS - Université Montpellier 2, France
{benmohamed,leclere,mugnier}@lirmm.fr

Abstract. We consider the deduction problem in the existential conjunctive fragment of first-order logic with atomic negation. This problem can be recast in terms of other database and artificial intelligence problems, namely query containment, clause entailment and boolean query answering. We refine an algorithm scheme that was proposed for query containment, which itself improves other known algorithms in databases. To study it experimentally, we build a random generator and analyze the influence of several parameters on the problem instance difficulty. Using this methodology, we experimentally compare several heuristics.

Keywords: Deduction, Negation, Conjunctive Queries with Negation, Graphs, Homomorphism, Algorithm, Heuristics, Experiments.

1 Introduction

We consider deduction checking in the fragment of first-order logic (FOL) composed of *existentially closed conjunctions of literals* (without functions). The deduction problem, called DEDUCTION in this paper, takes as input two formulas f and g in this fragment and asks if f can be deduced from g (noted $g \vdash f$). DEDUCTION is representative of several fundamental artificial intelligence and database problems. It can be immediately recast as a *query containment* checking problem, which is a fundamental database problem. This problem takes two queries q_1 and q_2 as input, and asks if q_1 is contained in q_2 , i.e. if the set of answers to q_1 is included in the set of answers to q_2 for all databases (e.g. [1]). Algorithms based on query containment can be used to solve various problems, such as query evaluation and optimization [2][3] or rewriting queries using views [4]. So-called (positive) *conjunctive queries* form a class of natural and frequently used queries and are considered as basic queries in databases [2][5] and more recently in the semantic web. Conjunctive queries with negation extend this class with negation on atoms. Query containment checking for conjunctive queries with negation is essentially the same problem as DEDUCTION, in the sense that there are natural polynomial reductions from one problem to another, which preserve the structure of the objects. Another related problem in artificial intelligence is the *clause entailment* problem, i.e. a basic problem in inductive logic programming [6]: given two clauses C_1 and C_2 , does C_1 entail C_2 ? If we consider first-order clauses, i.e. universally closed disjunctions of literals, without function symbols, by contraposition we obtain an instance of DEDUCTION.

Query answering is a key problem in the domain of knowledge representation and reasoning. Generally speaking, it takes a knowledge base and a query as input and asks for the set of answers to the query that can be retrieved from the knowledge base. When the query is boolean, i.e. with a yes/no answer, the problem can be recast as checking whether the query can be deduced from the knowledge base. When the knowledge base is simply composed of a set of positive and negative factual assertions, i.e. existentially closed conjunctions of literals (possibly stored in a relational database), and the query is a boolean conjunctive query with negation, we again obtain DEDUCTION. Integration of an ontology in this knowledge base is discussed in the conclusion of this paper.

If the considered fragment is restricted to positive literals, deduction checking is “only” NP-complete and this has been intensively studied from an algorithm viewpoint, in particular in the form of the equivalent constraint satisfaction problem [7]. In contrast, when atomic negation is considered, deduction checking becomes Π_2^P -complete¹(e.g. [8]) and very few algorithms for solving it can be found in the literature. Several algorithms have been proposed for the database query containment problem [9][10][11]. They all use homomorphism as a core notion. In this paper, we refine the algorithm scheme introduced in [11] for query containment checking, which itself improves other algorithms proposed in databases. To study it experimentally, we build a random generator and analyze the influence of several parameters on the problem instance difficulty. Then we experimentally compare several heuristics using this methodology. A research report with more figures, algorithms, as well as preliminary results on the comparison with logical provers is available [12]².

Paper layout. Section 2 introduces the framework. In Section 3, we present our experimental methodology and choices. Section 4 is devoted to the comparison of several heuristics, which leads to refine the algorithm. Section 5 outlines the prospects of this work.

2 Framework

We note $FOL\{\exists, \wedge, \neg_a\}$ the fragment of FOL composed of existentially closed conjunctions of literals, with constants but without other function symbols. A formula in $FOL\{\exists, \wedge, \neg_a\}$ can also be seen as a *set* of (positive and negative) literals.

In [11], queries (i.e. formulas in the present paper) are seen as labeled graphs. This allows us to rely on graph notions that have no simple equivalent in logic (such as pure subgraphs, see later). More precisely, a formula f is represented as a bipartite, undirected and labeled graph F , called *polarized graph (PG)*, with two kinds of nodes: term nodes and predicate nodes. Each term of the formula is associated to a term node, that is unlabeled if it is a variable, otherwise it is labeled by the constant itself. A positive (resp. negative) literal with predicate r is associated to a predicate node labeled $+r$ (resp. $-r$) and it is linked to the nodes assigned to its terms. The labels on edges correspond to the position of each term in the literal (see Figure 1 for an example). For simplicity, the subgraph corresponding to a literal, i.e. induced by a predicate node and

¹ $\Pi_2^P = (co-NP)^{NP}$.

² <http://www.lirmm.fr/~benmohamed/Publications/RR.pdf>

its neighbors, is also called a *literal*. We note it $+r(t_1, \dots, t_n)$ (resp. $-r(t_1, \dots, t_n)$) if the predicate node has label $+r$ (resp. $-r$) and list of neighbors t_1, \dots, t_n . The notation $\sim r(t_1, \dots, t_n)$ indicates that the literal with predicate r may be positive or negative. Literals $+r(t_1, \dots, t_n)$ and $-r(u_1, \dots, u_n)$ with the same predicate but different signs are said to be *opposite*. Literals $+r(t_1, \dots, t_n)$ and $-r(t_1, \dots, t_n)$ with the same list of arguments are said to be *contradictory*. Given a predicate node label (resp. literal) l , \bar{l} denotes the complementary predicate label (resp. literal) of l , i.e. it is obtained from l by reversing its sign. Formulas are denoted by small letters (f and g) and the associated graphs by the corresponding capital letters (F and G). We note $G \vdash F$ iff $g \vdash f$. A PG is *consistent* if it does not contain two contradictory literals (i.e. the associated formula is satisfiable).




Fig. 1. Polarized graphs associated with f and g

Homomorphism is a core notion in this work. A *homomorphism* h from a PG F to a PG G is a mapping from nodes of F to nodes of G , which preserves bipartition (the image of a term - resp predicate - node is a term - resp. predicate - node), preserves edges (if rt is an edge with label i in F then $h(r)h(t)$ is an edge with label i in G), preserves predicate node labels (a predicate node and its image have the same label) and can instantiate term node labels (if a term node is labeled by a constant, its image has the same label, otherwise the image can be any label). On the associated formulas f and g , it corresponds to a *substitution* h from f to g such that $h(f) \subseteq g$. When there is a homomorphism h from F to G , we say that F *maps* to G by h . F is called the *source* graph and G the *target* graph.

Definition 1 (Complete graph and completion). Let G be a consistent PG. It is complete w.r.t. a set of predicates \mathcal{P} , if for each $p \in \mathcal{P}$ with arity k , for each k -tuple of term nodes (not necessarily distinct) t_1, \dots, t_k in G , it contains $+p(t_1, \dots, t_k)$ or $-p(t_1, \dots, t_k)$. A completion G' of G is a PG obtained from G by repeatedly adding new predicate nodes (on term nodes present in G) without yielding inconsistency of G . Each addition is a completion step. A completion of G is called total if it is a complete graph w.r.t. the set of predicates considered, otherwise it is called partial.

If F and G have only positive literals, $G \vdash F$ iff F maps to G . When we consider positive and negative literals, only one side of this property remains true: if F maps to G then $G \vdash F$; the converse is false, as shown in Example 1.

Example 1. See Figure 1: F does not map to G but $g \vdash f$. Indeed, if we complete g w.r.t. predicate p , we obtain the formula g' (equivalent to g): $g' = (g \wedge p(b) \wedge p(c)) \vee (g \wedge \neg p(b) \wedge p(c)) \vee (g \wedge p(b) \wedge \neg p(c)) \vee (g \wedge \neg p(b) \wedge \neg p(c))$. Each of the four conjunctions

of g' is a way to complete g w.r.t. p . F maps to each of the graphs associated with them. Thus f is deductible from g' .

One way to solve DEDUCTION is therefore to generate all total completions obtained from G using predicates appearing in G , and then to test if F maps to each of these graphs.

Theorem 1. [11] *Let F and G be two PGs (with G consistent), $G \vdash F$ iff for all G^c , total completion of G w.r.t. the set of predicates appearing in G , F maps to G^c .*

We can restrict the set of predicates considered to those appearing in opposite literals both in F and in G [11]. In the sequel, this set is called the *completion vocabulary* of F and G and denoted \mathcal{V} .

A brute-force approach, introduced in [9], consists of computing the set of total completions of G and checking the existence of a homomorphism from F to each of them. However, the complexity of this algorithm is prohibitive: $\mathcal{O}(2^{(n_G)^k \times |\mathcal{V}|} \times \text{hom}(F, G^c))$, where n_G is the number of term nodes in G , k is the maximum arity of a predicate, \mathcal{V} is the completion vocabulary and $\text{hom}(F, G^c)$ is the complexity of checking the existence of a homomorphism³ from F to G^c .

Two types of improvements of this method are proposed in [11]. First, let us consider the space leading from G to its total completions and partially ordered by the relation “subgraph of”. This space is explored as a binary tree with G as root. The children of a node are obtained by adding, to the graph associated with this node (say G'), a predicate node in positive and negative form (each of the two new graphs is thus obtained by a completion step from G'). The aim is to find a set of partial completions covering the set of total completions of G , i.e. the question becomes: “Is there a set of partial completions $\{G_1, \dots, G_n\}$ such that (1) F maps to each G_i for $i = 1 \dots n$; (2) each total completion G^c of G is covered by a G_i (i.e. G_i is a subgraph of G^c)?” After each completion step, we check whether F maps to the current partial completion: if yes, this completion is one of the sought G_i , otherwise the exploration continues.

Figure 2 illustrates this method on the very easy case of Example 1. Two graphs G_1 and G_2 are built from G , respectively by adding $+p(c)$ and $-p(c)$. F maps to G_1 , thus there is no need to complete G_1 . F does not map to G_2 : two graphs G_3 and G_4 are built from G_2 , by adding $+p(b)$ and $-p(b)$ to G_2 . F maps to G_3 and to G_4 , respectively. Finally, the set proving that F is deductible from G is $\{G_1, G_3, G_4\}$ (and there are four total completions of G w.r.t. p). Algorithm 1 implements this method (the numbers in the margin are relative to the refinements studied in Section 4).

The second kind of improvement consists of identifying subgraphs of F for which there must be a homomorphism to G when $G \vdash F$. Such a subgraph, say F' , can be used as a filter to detect a failure before entering the completion process: if F' does not map to G , then $G \not\vdash F$. In [10], this property is exhibited for F^+ , which is the set of all positive literals in F . This result is generalized in [11] with the notions of *pure subgraph* and *compatible* homomorphism.

Definition 2 (pure subgraph). *A PG is said to be pure if it does not contain opposite literals (i.e. each predicate appears only in one form, positive or negative). A pure*

³ Homomorphism checking is NP-complete. A brute-force algorithm solves it in $\mathcal{O}(n_G^{n_F})$, where n_F is the number of term nodes in F .




Fig. 2. The search tree of Example 1. Each black dot represents a G^c and each square a G_i .

Algorithm 1. recCheck(G)

Input: a consistent PG G

Data: F, \mathcal{V}

Result: true if $G \vdash F$, false otherwise

begin

if there is a homomorphism from F to G **then return true** ;

if G is complete w.r.t. \mathcal{V} **then return false** ;

(3) *** Filtering step ***\

(1) Choose $r \in \mathcal{V}$ and t_1, \dots, t_n in G such that $+r(t_1, \dots, t_n) \notin G$ and $-r(t_1, \dots, t_n) \notin G$;

 Let G' be obtained from G by adding $+r(t_1, \dots, t_n)$;

 Let G'' be obtained from G by adding $-r(t_1, \dots, t_n)$;

(2) **return recCheck(G') AND recCheck(G'')** ;

end

subgraph of F is a subgraph of F that contains all term nodes in F (but not necessarily all predicate nodes)⁴ and is pure.

We will use the following notations for pure subgraphs of F : F^{max} denotes a pure subgraph that is maximal for inclusion; F^+ is the F^{max} with all positive predicate nodes in F ; F^- is the F^{max} with all negative predicate nodes in F ; F^{Max} denotes a F^{max} of maximal cardinality.

Moreover, a homomorphism from a pure subgraph of F to G has to be “compatible” with a homomorphism from F to a total completion of G .

Definition 3 (Border, Compatible homomorphism). Let F and G be two PGs and F' be a pure subgraph of F . The predicate nodes of $F \setminus F'$ are called border predicate nodes of F' w.r.t. F . A homomorphism h from F' to G is said to be compatible w.r.t. F if, for each border predicate node inducing the literal $\sim p(t_1, \dots, t_k)$, the opposite literal $\overline{\sim p}(h(t_1), \dots, h(t_k))$ is not in G .

Theorem 2. [11] If $G \vdash F$ then, for each pure subgraph F' of F , there is a compatible homomorphism from F' to G w.r.t. F .

The following filtering step can thus be performed before the recursive algorithm: select some F^{max} ; if there is no compatible homomorphism from F^{max} to G then return false.

⁴ Note that this subgraph does not necessarily correspond to a set of literals because some term nodes may be isolated.

3 Experimental Methodology

Due to the lack of benchmarks or real-world data available for the studied problem, we built a random generator of polarized graphs. The chosen parameters are as follows:

- the number of *term* nodes (i.e. the number of terms in the associated formula)⁵;
- the number of distinct *predicates*;
- the *arity* of these predicates (set at 2 in the following experiments);
- the *density* per predicate, which is, for each predicate p , the ratio of the number of literals with predicate p in the graph to the number of literals with predicate p in a total completion of this graph w.r.t. $\{p\}$.
- the *percentage of negation* per predicate, which is, for each predicate p , the percentage of *negative* literals with predicate p among all literals with predicate p in the graph.

An instance of DEDUCTION is obtained by generating a pair (F, G) of polarized graphs. In this paper, we chose the same number of term nodes for both graphs. The difficulty of the problem led us to restrict this number to between 5 and 8 term nodes. In the sequel we adopt the following notations: nbT represents the number of term nodes, $nbPred$ the number of distinct predicates, SD (resp. TD) the Source (resp. Target) graph Density per predicate and neg the percentage of negation per predicate.

In order to discriminate between different techniques, we first experimentally studied the influence of the parameters on the “difficulty” of instances. We measured the difficulty in three different ways: the running time, the size of the search tree and the number of homomorphism checks (see [12] for more details). Concerning the negation percentage, we checked that the maximal difficulty is obtained when there are as many negative relation nodes as positive relation nodes.

One can expect that increasing the number of predicates occurring in graphs increases the difficulty, in terms of running time as well as the size of the searched space. Indeed, the number of completions increases exponentially (there are $(2^{n_C})^{nbPred}$ total completions for $nbPred$ predicates). These intuitions are only partially validated by the experiments: see Table 1, which shows, for each number of predicates, the density values at the difficulty peak. We observe that the difficulty increases up to a certain number of predicates (3 here, with a CPU time of 4912 and a Tree size of 72549) and beyond this value, it continuously decreases. Moreover, the higher the number of predicates, the lower SD which entailed the greatest difficulty peak, and the higher the difference between TD and SD at the difficulty peak.

For each value of the varying parameter, we considered 500 instances and computed the mean search cost of the results on these instances (with a timeout set at 5 minutes). The program is written in Java. The experiments were performed on a Sun fire X4100 Server AMD Opteron 252, equipped with a 2.6 GHz Dual-Core CPU and 4G of RAM, under Linux. In the sequel we only show the CPU time when the three difficulty measures are correlated.

⁵ We do not generate constants; indeed, constants tend to make the problem easier to solve because there are fewer potential homomorphisms; moreover, this parameter does not influence the studied heuristics.

Table 1. Influence of the number of predicates: $nbT=5$, $neg=50\%$

NbPred	SD	TD	CPU time (ms)	Tree size
1	0.24	0.28	20	68
2	0.08	0.16	3155	53225
3	0.08	0.4	4912	72549
4	0.08	0.68	3101	44089
5	0.08	0.76	683	8483

4 Refinement of the Algorithm

In this section, we analyze three refinements of Algorithm 1, which concern the following aspects: (1) the choice of the next literal to add; (2) the choice of the child to explore first; (3) dynamic filtering at each node of the search tree.

1. Since the search space is explored in a depth-first manner, the choice of the next literal to add, i.e. $\sim r(t_1, \dots, t_n)$ in Algorithm 1 (Point 1), is crucial. A brutal technique consists of choosing r and t_1, \dots, t_n randomly. Our proposal is to guide this choice by a compatible homomorphism, say h , from a F^{max} to the current G . More precisely, the border predicate nodes $\sim r(e_1, \dots, e_n)$ w.r.t. this F^{max} can be divided into two categories. In the first category, we have the border nodes s.t. $\sim r(h(e_1) \dots h(e_n)) \in G$, which can be used to *extend* h ; if all border nodes are in this category, h can be extended to a homomorphism from F to G . The choice of the literal to add is based on a node $\sim r(e_1, \dots, e_n)$ in the second category: r is its predicate symbol and $t_1, \dots, t_n = h(e_1) \dots h(e_n)$ are its neighbors (note that neither $\sim r(h(e_1) \dots h(e_n))$ nor $\overline{\sim r}(h(e_1) \dots h(e_n))$ is in G since $\sim r(e_1, \dots, e_n)$ is in the second category and h is compatible). Intuitively, the idea is to give priority to predicate nodes potentially able to transform this compatible homomorphism into a homomorphism from F to a (partial) completion of G , say G' . If so, all completions including G' are avoided.

Figure 3 shows the results obtained with the following choices:

- *random choice*;
- *random choice + filter*: random choice and F^+ as filter (i.e. at each recCheck step a compatible homomorphism from F^+ to G is looked for: if none exists, the false value is returned);
- *guided choice*: F^+ used both as a filter and as a guide.

As expected, the guided choice is always much better than the random choice (with or without filter): on the guided choice peaks ($TD=0.15$ and $TD=0.2$), it is almost 11 and 8 times better than the random choice with filter. The greatest difference is for $TD=0.25$ with the guided choice almost 116 times better than the random choice with filter.

2. Experiments have shown that the order in which the children of a node, i.e. G' and G'' in Algorithm 1 (Point 2), are explored is important. Assume that Point 1 in Algorithm 1 relies on a guiding subgraph. Consider Figure 4, where F^+ is the guiding subgraph (hence the border is composed of negative predicate nodes): we see that it is




Fig. 3. Influence of the completion choice: $nbT=7$, $nbPred=1$, $SD=0.14$, $neg=50\%$

always better to explore G' before G'' . If we take F^- as the guiding subgraph, then the inverse order is better. More generally, let $\sim r(e_1 \dots e_n)$ be the border node that defines the literal to add. Let us call h -extension (resp. h -contradiction) the graph built from G by adding $\sim r(h(e_1) \dots h(e_n))$ (resp. $\overline{\sim r}(h(e_1) \dots h(e_n))$). See Example 2. It is better to first explore the child corresponding to the h -contradiction. Intuitively, by contradicting the compatible homomorphism found, this gives priority to failure detection.

Example 2. See Figure 1. $F^+ = \{+p(x), +s(x, y)\}$. Let F^+ be the guiding subgraph. The only border node of F^+ w.r.t. F is $-p(y)$. $h = \{(x, a), (y, b)\}$ is the only compatible homomorphism from F^+ to G . The h -extension (resp. h -contradiction) is obtained by adding $+p(b)$ (resp. $-p(b)$).

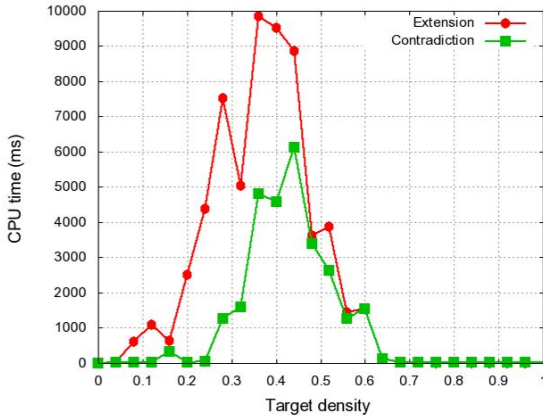


Fig. 4. Influence of the exploration order: $nbT=7$, $nbPred=3$, $SD=0.06$, $neg=50\%$

Table 2. Influence of the dynamic filtering: $nbT=8$, $nbPred=3$, $SD=0.03$, $TD=0.16$, $neg=50\%$

Configuration	CPU time (ms)	Tree size	Hom check
<i>Max</i>	3939	2423	3675
<i>Max-Max</i>	3868	2394	3858
<i>Max-all</i>	3570	1088	6338

3. The last improvement consists of performing dynamic filtering at each node of the search tree. Once again, the aim is to detect a failure sooner. More precisely, we consider a set of F^{max} and check if there is a compatible homomorphism from each element in this set to the newly generated graph. Table 2 shows the results obtained at the difficulty peak with the following configurations: *Max*: F^{Max} as guide and no filter; *Max-Max*: F^{Max} as guide and F^{Max} (the subgraph on the predicate nodes in $F \setminus F^{Max}$) as filter; *Max-all*: F^{Max} as guide and all other F^{max} as filters.

Unsurprisingly, the stronger the dynamic filtering, the smaller the size of the search tree. The CPU time is almost the same for all configurations (and all TD values) though *Max-all* checks much more homomorphisms than the others. Since our current algorithm for homomorphism checking can be improved, these results show that *Max-all* is the best choice. The algorithm finally obtained is given in [12].

5 Perspectives

Our experiments show that the problem is really complex in practice in the difficult area. This may be an argument in favor of restricting conjunctive queries with negation (to pure queries for instance) or using alternative kinds of negation. Closed-world negation is often used. However, even in this case, good algorithms are required for comparing queries (cf. the query containment problem, which is the basis of many mechanisms in databases) and we again find the deduction problem with classical negation.

Ontologies play a central role in knowledge bases, and this role is increasing in databases. A lightweight ontology, in the form of a partial order, or more generally a preorder, on predicates (i.e. on the concepts and predicates of the ontology) can be taken into account without increasing complexity. Note that, in this case, we obtain exactly the deduction problem in a fragment of conceptual graphs [13][14]. Homomorphism is extended in a straightforward way to take the partial order into account. The heuristics studied here still work with an extension of opposite literals: $+r(t_1, \dots, t_n)$ and $-s(u_1, \dots, u_n)$ are opposite if $r \geq s$ (then, a pure subgraph is defined as previously with this extended definition). How this work can be extended to more complex ontologies is an open issue.

On the experimental side, further work includes precise comparison with techniques used by logical solvers. We have not found logical algorithms dedicated to the DEDUCTION problem addressed in this paper. For the moment, we have considered the free tools *Prover9* (the successor of the Otter prover [15]) and *Mace4* [16], which consider full first-order logic. Prover9 is based on the resolution method and Mace4 enumerates models by domains of increasing size. These tools are complementary: Prover9 looks for a proof and Mace4 looks for a model that is a counterexample. On preliminary

results, we show that our refined algorithm is better than the combination Prover9-Mace4. However the comparison needs to be refined and other provers have to be considered.

References

1. Abiteboul, S., Hull, R., Vianu, V.: *Foundations of Databases: The Logical Level*. Addison-Wesley, Reading (1995)
2. Chandra, A.K., Merlin, P.M.: *Optimal Implementation of Conjunctive Queries in Relational Databases*. In: 9th ACM Symposium on Theory of Computing, pp. 77–90 (1977)
3. Aho, A.V., Sagiv, Y., Ullman, J.D.: *Equivalences Among Relational Expressions*. *SIAM J. Comput.* 8(2), 218–246 (1979)
4. Halevy, A.Y.: *Answering Queries Using Views: a Survey*. *VLDB Journal* 10(4), 270–294 (2001)
5. Ullman, J.D.: *Principles of Database and Knowledge-Base Systems*, vol. 2. Computer Science Press, Rockville (1989)
6. Muggleton, S., De Raedt, L.: *Inductive Logic Programming: Theory and Methods*. *J. Log. Program.* 19/20, 629–679 (1994)
7. Rossi, F., van Beek, P., Walsh, T.: *Handbook of Constraint Programming*. Elsevier, Amsterdam (2006)
8. Farré, C., Nutt, W., Teniente, E., Urpí, T.: *Containment of Conjunctive Queries over Databases with Null Values*. In: Schwentick, T., Suciu, D. (eds.) *ICDT 2007*. LNCS, vol. 4353, pp. 389–403. Springer, Heidelberg (2007)
9. Ullman, J.D.: *Information Integration Using Logical Views*. In: Afrati, F.N., Kolaitis, P.G. (eds.) *ICDT 1997*. LNCS, vol. 1186, pp. 19–40. Springer, Heidelberg (1997)
10. Wei, F., Lausen, G.: *Containment of Conjunctive Queries with Safe Negation*. In: Calvanese, D., Lenzerini, M., Motwani, R. (eds.) *ICDT 2003*. LNCS, vol. 2572, pp. 343–357. Springer, Heidelberg (2003)
11. Leclère, M., Mugnier, M.-L.: *Some Algorithmic Improvements for the Containment Problem of Conjunctive Queries with Negation*. In: Schwentick, T., Suciu, D. (eds.) *ICDT 2007*. LNCS, vol. 4353, pp. 401–418. Springer, Heidelberg (2007)
12. Ben Mohamed, K., Leclère, M., Mugnier, M.-L.: *Deduction in existential conjunctive first-order logic: an algorithm and experiments*. RR-10010, LIRMM (2010)
13. Kerdiles, G.: *Saying it with Pictures: a Logical Landscape of Conceptual Graphs*. Univ. Montpellier II, Amsterdam (2001)
14. Mugnier, M.-L., Leclère, M.: *On Querying Simple Conceptual Graphs with Negation*. *Data Knowl. Eng.* 60(3), 468–493 (2007)
15. McCune, W.: *OTTER 3.3 Reference Manual*. CoRR, vol. cs.SC/0310056 (2003)
16. McCune, W.: *Mace4 Reference Manual and Guide*. CoRR, vol. cs.SC/0310055 (2003)