

The hardness of revising defeasible preferences

Guido Governatori^{1,3,4}, Francesco Olivieri^{1,2,3},
Simone Scannapieco^{1,2,3}, and Matteo Cristani²

¹ NICTA, Queensland Research Laboratory, Australia^{*}

² Department of Computer Science, University of Verona, Italy

³ Institute for Integrated and Intelligent Systems, Griffith University, Australia

⁴ Queensland University of Technology, Australia

Abstract. Non-monotonic reasoning typically deals with three kinds of knowledge. *Facts* are meant to describe immutable statements of the environment. *Rules* define relationships among elements. Lastly, an ordering among the rules, in the form of a *superiority relation*, establishes the relative strength of rules. To revise a non-monotonic theory, we can change either one of these three elements. We prove that the problem of revising a non-monotonic theory by only changing the superiority relation is a NP-complete problem.

1 Introduction

Preferences are a powerful tool agents use to make decisions. Given a knowledge base, agents are able to set an (partial) ordering among the elements of such a base by stating that they prefer an element better than another.

The number of possible applications of preferences is vast. Given a goal and a set of actions, an agent can choose a particular course of action instead of another in order to achieve the goal. Even in many legal contexts, where the agent typically has neither the power to change the normative system, nor to decide what norms are effective, the agent may argue that one norm applies instead of another. This case is peculiar: we have two norms (rules) stating opposite conclusions, but the apparent conflict can be solved through the preference mechanism.

Through the years, non-monotonic reasoning has been advanced for reasoning with partial, and possibly conflicting information. Settled in the context of logics representing non-monotonic reasoning, we usually deal with three types of knowledge. First, we have *facts*. Facts are meant to describe simple pieces of information which are considered to be always true in the environment the agent acts in. Then we have *rules*. A rule describes how given elements interact with each other in order to obtain some conclusions. Lastly, we have a mechanism to solve conflicts (we assume to work within a skeptical system). Typically, this is presented in the form of a binary relation among pairs of rules. Such a *superiority* (or *preference*) *relation* states a (partial) ordering among the rules: when two rules for opposite conclusions are “applicable” at the same time, the superiority relation solves the conflict in favour of one conclusion upon the other. In fact, the superiority

^{*} NICTA is funded by the Australian Government as represented by the Department of Communications and the Australian Research Council through the ICT Centre of Excellence program.

relation expresses more than that; it actually reflects preferences of the agent on the inner structure of the theory, becoming a fundamental mechanism in many real life scenarios. In this paper we use Defeasible Logic, as presented in [1], to formally encode facts, rules and the superiority relation.

Naturally, there is more. The information at hand may change and when this happens, we need to revise our knowledge base or theory. Much work has been done to understand either how to properly change a theory, or which conditions the “new theory” should meet. When updating the knowledge base of a defeasible theory, revising operators may act on either one of the three constitutive elements, being that the facts, the rules [2], the superiority relation, or a combination of those.

Revising facts can be seen as changing the operational environment, or just simply examining a different factual scenario on which rules and the superiority relation are applied to. In a legal proceedings this corresponds to change the evidence of the case. Changing rules expresses the ability in adding and/or removing the dependencies among the atoms of the theory. In a legal setting this means to create, delete or modify some norms. Finally, there are some contexts where the above two approaches are not possible and the only solution left is to change the relative strength between pairs of rules. Always referring to legal systems, this may be the case in situations like that of an average citizen appealing to a court. He has no power to change the Law, and has no power on what norms are effective in the jurisdiction he is situated in. These powers instead are reserved to persons, entities and institutions specifically designated to do so, for example the parliament and, under some given constraints, also by judges (in Common Law juridical system, especially). However, a citizen can argue that one norm takes precedence over another in a specific case. This amounts to saying that one norm is to be preferred to the other in the case.

The main results of the present paper come from a thorough investigation on revising a defeasible theory where only changing the superiority relation is allowed. [3,4] study “patterns”, or conditions, on the elements of a defeasible theory that would distinguish situations where a revision is possible, against situations where it is not. Scope of this investigation is to make a step further. We prove that the problem of revising a theory by only changing the superiority relation is NP-complete.

The structure of the paper is as follows. In Section 2 we recall the basics of Defeasible Logic and we introduce the notion of what means for a formula to be “tautological” in the context of revision in Defeasible Logic, where revision operations are limited to the superiority relation. In Section 3 we formally set up the decision problem related to the type of revision we are interested in this paper. We conclude the paper in Section 4 with a summary and a discussion of some related work and possible future work.

2 Defeasible Logic

We shall describe the structure of a defeasible theory, and proceed by reporting definitions of a theory *based on a specific set of rules*, of a theory being *decisive* (along with two preliminary results), and of a literal being *tautological*. Admittedly, this section is dense but such are the necessary means to prove the NP-completeness result of the next section.

A defeasible theory consists of five different kinds of knowledge: facts, strict rules, defeasible rules, defeaters, and a superiority relation [1].

Let PROP be a set of propositional atoms, Lbl be a set of arbitrary labels. The set Lit = PROP \cup $\{\neg p \mid p \in \text{PROP}\}$ denotes the set of *literals*. The *complement* of a literal q is denoted by $\sim q$; if q is a positive literal p , then $\sim q$ is $\neg p$, and if q is a negative literal $\neg p$ then $\sim q$ is p .

Definition 1 A defeasible theory D is a structure $(F, R, >)$, where

1. $F \subseteq \text{Lit}$ denote simple pieces of information that are considered always to be true.
For example, a fact is that “Sylvester is a cat”, formally $\text{cat}(\text{Sylvester})$;
2. R contains three types of rules: strict rules, defeasible rules, and defeaters.
3. $> \subseteq R \times R$ is a binary relation whose transitive closure is acyclic.

A theory is finite if the set of facts and rules are finite.

A rule is an expression $r : A(r) \hookrightarrow C(r)$ and consists of: (i) A unique name $r \in \text{Lbl}$, (ii) the *antecedent* $A(r)$ which is a finite subset of Lit, (iii) an *arrow* $\hookrightarrow \in \{\rightarrow, \Rightarrow, \rightsquigarrow\}$ denoting, respectively, a strict rule, a defeasible rule and a defeater, and (iv) its *consequent* (or *head*) $C(r) \in \text{Lit}$, which is a single literal. A *strict rule* is a rule in which whenever the premises are indisputable (e.g., facts), then so is the conclusion. For example,

$$\text{cat}(X) \rightarrow \text{mammal}(X)$$

means that “every cat is a mammal”. On the other hand, a *defeasible rule* is a rule that can be defeated by contrary evidence; for example, “cats typically eat birds”:

$$\text{cat}(X) \Rightarrow \text{eatBirds}(X).$$

The underlying idea is that if we know that something is a cat, then we may conclude that it eats birds, unless there is evidence proving otherwise. *Defeaters* are rules that cannot be used to draw any conclusion. Their only use is to prevent some conclusions, i.e., to defeat defeasible rules by producing evidence to the contrary. An example is “if a cat has just fed itself, then it might not eat birds”:

$$\text{justFed}(X) \rightsquigarrow \neg \text{eatBirds}(X).$$

The *superiority relation* $>$ among rules is used to define where one rule may override the (opposite) conclusion of another one, e.g., given the defeasible rules

$$\begin{aligned} r : \text{cat}(X) &\Rightarrow \text{eatBirds}(X) \\ r' : \text{domesticCat}(X) &\Rightarrow \neg \text{eatBirds}(X) \end{aligned}$$

which would contradict one another if Sylvester is both a cat and a domestic cat, they do not in fact contradict if we state that r' wins against r , leading to conclude that Sylvester does not to eat birds.

Like in [1], we consider only a propositional version of this logic, and we do not take into account function symbols. Every expression with variables represents the finite set of its variable-free instances.

We use the infix notation $r > s$ to mean that $(r, s) \in >$. The set of strict rules in R is denoted by R_s , and the set of strict and defeasible rules by R_{sd} . We name $R[q]$ the set of rules in R whose head is q . A *conclusion* of D is a tagged literal and can have one of the following forms:

- $+\Delta q$, which means that q is definitely provable in D , i.e., there is a definite proof for q , that is a proof using facts, and strict rules only;
- $-\Delta q$, which means that q is definitely not provable, or refuted, in D (i.e., a definite proof for q does not exist);
- $+\partial q$, which means that q is defeasibly provable in D ;
- $-\partial q$, which means that q is not defeasibly provable, or refuted, in D .

Given a defeasible theory D , a proof P of length n in D is a finite sequence $P(1), \dots, P(n)$ of tagged literals of the type $+\Delta q$, $-\Delta q$, $+\partial q$ and $-\partial q$, where the proof conditions defined in the rest of this section hold. $P(1..n)$ denotes the first n steps of proof P .

Given $\# \in \{\Delta, \partial\}$ and a proof P in D , a literal q is $\#$ -*provable* in D if there is a line $P(m)$ of P such that $P(m) = \#q$. A literal q is $\#$ -*refuted* in D if there is a line $P(m)$ of P such that $P(m) = \#q$.

The definition of Δ describes just forward chaining of strict rules.

- $+\Delta$: If $P(n+1) = +\Delta q$ then
- (1) $q \in F$ or
 - (2) $\exists r \in R_s[q] \forall a \in A(r) : +\Delta a \in P(1..n)$.

Literal q is definitely provable if either (1) is a fact, or (2) there is a strict rule for q , whose antecedents have all been definitely proved.

- $-\Delta$: If $P(n+1) = -\Delta q$ then
- (1) $q \notin F$ and
 - (2) $\forall r \in R_s[q] \exists a \in A(r) : -\Delta a \in P(1..n)$.

Literal q cannot be definitely proven ($-\Delta q$) if (1) is not a fact and (2) every strict rule for q has at least one definitely refuted antecedent.

The following definition states notions of being applicable and discarded.

Definition 2 *In the proof condition for $\pm\partial$, a rule $r \in R_{sd}$ is (i) applicable iff $\forall a \in A(r)$, $+\partial a \in P(1..n)$; (ii) discarded iff $\exists a \in A(r)$ such that $-\partial a \in P(1..n)$.*

We now introduce the proof conditions to show that a literal is defeasibly provable.

- $+\partial$: If $P(n+1) = +\partial q$ then
- (1) $+\Delta q \in P(1..n)$ or
 - (2) (2.1) $-\Delta \sim q \in P(1..n)$ and
 - (2.2) $\exists r \in R_{sd}[q]$ s.t. r is applicable, and
 - (2.3) $\forall s \in R[\sim q]$. either s is discarded, or
 - (2.3.1) $\exists t \in R[q]$ s.t. t is applicable and $t > s$.

Literal q is defeasibly provable if (1) q is already definitely provable, or (2) we argue using the defeasible part of the theory. For (2), $\sim q$ is not definitely provable (2.1), and there exists an applicable strict or defeasible rule for q (2.2). Every attack s is either discarded (2.3), or defeated by a stronger rule t (2.3.1).

On the other hand, to prove the a literal is defeasibly refuted ($-\partial$) we have to show that all possible ways to prove it fail. This is encoded by the following proof conditions that correspond to a constructive negation of the conditions for $+\partial$.

- $-\partial$: If $P(n+1) = -\partial q$ then
- (1) $-\Delta q \in P(1..n)$ and either
 - (2) (2.1) $+\Delta \sim q \in P(1..n)$ or
 - (2.2) $\forall r \in R_{sd}[q]$. either r is discarded, or
 - (2.3) $\exists s \in R[\sim q]$ s.t. s is applicable, and
 - (2.3.1) $\forall t \in R[q]$. either t is discarded, or $t \not\prec s$.

As usual, given $\# \in \{\Delta, \partial\}$, a literal p and a theory D , we use $D \vdash \pm\#p$ to denote that there is a proof P in D where for some line i , $P(i) = \pm\#p$. Alternatively, we say that $\pm\#p$ holds in D , or simply $\pm\#p$ holds when the theory is clear from the context. The set of positive and negative conclusions is called *extension*. Formally,

Definition 3 Given a defeasible theory D , the (defeasible) extension of D is defined as $E(D) = (+\Delta, -\Delta, +\partial, -\partial)$, where $\pm\# = \{l : l \text{ appears in } D \text{ and } D \vdash \pm\#l\}$, $\# \in \{\Delta, \partial\}$.

The inference mechanism of DL does not allow to derive inconsistencies unless the monotonic part of the logic is inconsistent, as clarified by the following definition.

Definition 4 A defeasible theory D is inconsistent iff there exists a literal p such that $((D \vdash +\partial p \text{ and } D \vdash +\partial \sim p) \text{ iff } (D \vdash +\Delta p \text{ and } D \vdash +\Delta \sim p))$.

In the rest of the paper, we shall make use neither of strict rules, nor defeaters. The restriction does not result in any loss of generality: (1) the superiority relation does not play any role in proving definite conclusions, and (2) for defeasible conclusions [1] proves that it is always possible to remove (a) strict rules from the superiority relation and (b) defeaters from the theory to obtain an equivalent theory without defeaters and where the strict rules are not involved in the superiority relation. A consequence of this assumption is that the theories we work with in this paper are consistent.

Before reporting the transformation used in the proof of NP-completeness, we need to introduce some additional terminology. Definition 5 constructs a defeasible theory starting from a fixed set of rules and an empty set of facts. This formulation limits the revision problem to preference changes, notwithstanding the particular instance of the superiority relation.

Definition 5 Given a set of defeasible rules R , a defeasible theory D is based on R iff

$$D = (\emptyset, R, >).$$

The aim of Definition 6 is to specify the possible relationships between a literal and all theories based on a set of rules R .

Definition 6 Given a set of defeasible rules R , a literal p is

1. $>-R$ -tautological iff for all theories D based on R , $D \vdash +\partial p$.
2. $>-R$ -non-tautological iff there exists a theory D based on R such that $D \not\vdash +\partial p$.
3. $>-R$ -refutable iff there exists a theory D based on R such that $D \vdash -\partial p$.

4. \succ - R -irrefutable iff for all theories D based on R , $D \not\vdash \neg \partial p$.

The notion of \succ - R -irrefutable represents the negative counterpart of \succ - R -tautological; the same holds for \succ - R -refutable and \succ - R -non-tautological.

If p is \succ - R -tautological, then, in every theory based on the set of rules R , an instance of the superiority relation such that p is defeasibly refuted does not exist. Accordingly, if a literal is \succ - R -tautological, then we cannot revise it.

On the contrary, if an instance of the superiority relation where p is no longer provable exists, then p is \succ - R -refutable.

Definition 7 A defeasible theory D is decisive iff for every literal p in D either $D \vdash +\partial p$, or $D \vdash -\partial p$.

Proposition 8 ([5]) Given a defeasible theory D , if the atom dependency graph of D is acyclic, then D is decisive.

3 Computational cost analysis

The processes of revising a defeasible theory by modifying only the superiority relation was first analysed in [3]. The focus of this section is to show that deciding whether such a revision is possible is NP-complete. To this end, we provide a reduction from 3-SAT to the following decision problem.

INSTANCE: A decisive defeasible theory D based on R , i.e., $D = (\emptyset, R, \succ)$, and a literal p .

QUESTION: Is it possible to change \succ into \succ' such that $D' = (\emptyset, R, \succ')$ and either

1. If $D \vdash +\partial p$, then $D' \vdash -\partial p$?
2. If $D \vdash -\partial p$, then $D' \vdash +\partial p$?

Definition 9 below exhibits a transformation from a 3-SAT formula into a set of defeasible rules.

Definition 9 Given a 3-SAT formula $\Gamma = \bigwedge_{i=1}^n C_i$ such that $C_i = \bigvee_{j=1}^3 a_j^i$, we define the Γ -transformation as the operation that maps Γ into the following set of defeasible rules

$$R_\Gamma = \{r_{ij}^a : \Rightarrow a_j^i \\ r_{ij} : a_j^i \Rightarrow c_i \\ r_{\sim i} : \Rightarrow \sim c_i \\ r_i : \sim c_i \Rightarrow p\}.$$

The above definition clearly shows that the mapping is polynomial in the number of literals appearing in the 3-SAT formula Γ .

The third step of the proof construction is to ensure that the proposed mapping allows the revision problem to give a correct answer (either positive, or negative) for every 3-SAT formula.

Lemma 10 Any defeasible theory D based on R_Γ of Definition 9 (for any Γ) is decisive.

Proof. It is easy to verify by case inspection that the atom dependency graph is acyclic.

Proposition 8 and Lemma 10 guarantee that any theory obtained by the Γ -transformation provides an answer. These results are also intended to establish relationships between the notions of tautological and refutable given in Definition 6.

We shall now propose the main result of NP-completeness. First, we shall prove that the revision problem is in NP. Second, we shall demonstrate the NP-hardness by exploiting the formulation of the 3-SAT problem and the transformation proposed in Definition 9.

Theorem 11 *The problem of determining the revision of a defeasible literal by changing the preference relation is NP-complete.*

Proof. The proof that the problem is in NP is straightforward. Given a defeasible theory $D = (F, R, >)$ and a literal p to be revised, an oracle guesses a revision (in terms of a new preference relation $>'$ applied to D) and checks if the state of p has changed based on the extensions of D and $D' = (F, R, >')$. The complexity of this check is bound to the calculus of $E(D)$ and $E(D')$, which [6] proves to be linear in the order of the theory.

To prove the NP-hardness, given a 3-SAT formula $\Gamma = \bigwedge_{i=1}^n C_i$ such that $C_i = \bigvee_{j=1}^3 a_j^i$, a defeasible theory D based on the set of defeasible rules R_Γ obtained by Γ -transformation, and a literal p in D , we show that:

1. if p is $>-R_\Gamma$ -tautological, then Γ is not satisfiable;
2. if p is $>-R_\Gamma$ -non-tautological, then Γ is satisfiable.

1. Lemma 10 allows us to reformulate the contrapositive using $>-R_\Gamma$ -refutable. If Γ is satisfiable, then there exists an interpretation I such that

$$\begin{aligned} I \models \Gamma &\iff I \models \bigwedge_{i=1}^n C_i \\ &\iff I \models C_1 \text{ and } \dots \text{ and } I \models C_n \\ &\iff I \models \bigvee_{j=1}^3 a_j^1 \text{ and } \dots \text{ and } I \models \bigvee_{j=1}^3 a_j^n. \end{aligned}$$

Thus, for each i , there exists j such that $I \models a_j^i$.

We build a defeasible theory $D' = (\emptyset, R_\Gamma, >')$ as follows. If there exists a literal b_k^i such that $b_k^i = \sim a_j^i$, then (r_{ij}^a, r_{ik}^b) is in $>'$. It follows that, by construction, D' proves $+\partial a_j^i$. This means that every rule r_{ij} is applicable and it is not weaker than the corresponding rule $r_{\sim i}$. Hence, we have $-\partial \sim c_i$, for all i . Consequently, each rule r_i for p is discarded and we conclude $-\partial p$. Accordingly, p is $>-R_\Gamma$ -refutable.

2. Again, by Lemma 10, every theory based on R_Γ is decisive. Thus, p is $>-R_\Gamma$ -refutable. Accordingly, there exists a theory $D' = (\emptyset, R_\Gamma, >')$ such that $D' \vdash -\partial p$. Given that $R_\Gamma[p] \neq \emptyset$ and $R_\Gamma[\sim p] = \emptyset$ by construction, every rule for p is discarded. Namely, we have $-\partial \sim c_i$, for all i .

Each rule $r_{\sim i}$ is vacuously applicable. To have $-\partial \sim c_i$, there must exist a rule r_{ij} that is applicable. Therefore, for each i there is at least one j such that $+\partial a_j^i$.

We build an interpretation I as follows:

$$I(a_j^i) = 1 \text{ iff } D \vdash +\partial a_j^i.$$

Since for each $1 \leq i \leq n$, we have $I(a_j^i) = 1$ for at least one j , then also $I \models C_i$ for all i , and we conclude that $I \models \Gamma$.

In addition, Theorem 11 specifies that there are situations where it is not possible to revise a literal by only using the superiority relation. For example, if we take a 3-SAT formula which is a tautology, the Γ -transformation generates a theory that cannot be revised only using the superiority relation.

Corollary 12 *There are theories and literals for which a revision by only modifying the superiority relation is not possible.*

The notion itself of “being a tautology in defeasible logic” may assume unexpected features. [4] provides a thorough characterisation of the problem where this is out of the scope of the present paper, and as such is not reported here. Nonetheless, we give the examples used in [4] to illustrate that finding patterns to determine whether a literal is tautological – in the sense of a literal that is always in the positive extension of theory based on a set of rules – is hard.

The following example shows a theory wherein literal p cannot be contracted by only modifying the superiority relation.

Example 1. Let D be a consistent, defeasible theory based on R where

$$R = \{ \Rightarrow_{r_1} l \Rightarrow_{r_2} \neg a \\ \Rightarrow_{r_3} a \Rightarrow_{r_4} p \\ \Rightarrow_{r_5} b \Rightarrow_{r_6} p \\ \Rightarrow_{r_7} \neg l \Rightarrow_{r_8} \neg b \}.$$

The rules above are depicted in a graphical way, where the subscripts attached to the arrows are the rule labels, and rules are chained. Thus, for example a is the consequent of r_3 as well as the antecedent of r_4 .

To contract p , we must block both the chains proving p , those being r_3r_4 and r_5r_6 . In order to do so, we should have that $D \vdash +\partial l$ as well as $D \vdash +\partial \neg l$. This is not possible being D consistent.

Unfortunately, recognising patterns as those shown in Example 1 does not guarantee to identify literals that cannot be contracted. In fact, let us consider Example 2.

Example 2. Let $D = (\emptyset, R = \{r_1, \dots, r_{18}\}, \emptyset)$ be a defeasible theory where

$$R = \{ \Rightarrow_{r_1} a \Rightarrow_{r_2} p \\ \Rightarrow_{r_3} b \Rightarrow_{r_4} p \\ \Rightarrow_{r_5} c \Rightarrow_{r_6} p \\ \Rightarrow_{r_7} l \Rightarrow_{r_8} \neg a \\ \Rightarrow_{r_9} \neg l \Rightarrow_{r_{10}} \neg b \\ \Rightarrow_{r_{11}} m \Rightarrow_{r_{12}} \neg b \\ \Rightarrow_{r_{13}} \neg m \Rightarrow_{r_{14}} \neg c \\ \Rightarrow_{r_{15}} n \Rightarrow_{r_{16}} \neg c \\ \Rightarrow_{r_{17}} \neg n \Rightarrow_{r_{18}} \neg a \}.$$

To contract p , we must block derivations of $+\partial a$, $+\partial b$ and $+\partial c$. This can be obtained by adding the following tuples to the superiority relation: (r_7, r_9) , (r_{11}, r_{13}) and (r_{15}, r_{17}) .

This is a counter-example to the pattern proposed in Example 1 but, once again, there are counter-examples to counter-examples (Example 3), and so on.

Example 3. Let $D = (\emptyset, R = \{r_1, \dots, r_{25}\}, \emptyset)$ be a defeasible theory where

$$\begin{aligned} & \Rightarrow_{r_{19}} e \Rightarrow_{r_{20}} p \\ & \Rightarrow_{r_{21}} f \Rightarrow_{r_{22}} p \\ n & \Rightarrow_{r_{23}} \neg e \\ \neg n & \Rightarrow_{r_{24}} \neg f \\ \neg m & \Rightarrow_{r_{25}} \neg f. \end{aligned}$$

To contract p , we must now block derivations also of $+\partial e$, and $+\partial f$. Derivation of e can be blocked only if we prove the antecedent of r_{23} , that is n (the derivation of c is blocked as well). This implies that the derivation of f is blocked only if $+\partial \neg m$ holds (the only antecedent of rule r_{25}). We can now operate only on the provability of either l , or $\neg l$. In both cases, one between a or b cannot be refuted.

4 Summary and related works

We started this paper by setting the problem of preference revision within the non-monotonic sceptical formalism of Defeasible Logic. After having introduced the logical mechanism, we proved the main result of NP-completeness for the decisional version of the problem. We have also proved that this type revision is not guaranteed to be successful.

To the best of our knowledge, revision of non-monotonic theories based on modifications of the underlying superiority relation has been neglected so far. Even if preferences have been widely studied in many different application areas (we cite, above others, database management [7,8] and legal reasoning [9]), our analysis is the first formal study to computationally characterise the problem of defeasible preference revision.

As far as we are aware of, the work most closely related to ours is that of [10], where the authors study how to abduct preference relations to support the derivation of a specific conclusion in a given theory. Therefore, the problem they address is conceptually different from what we presented in this paper, given that we focus on modifying an existing superiority relation instead of generating a new one that guarantees the derivation of a specific conclusion.

In [7,8], authors present algorithms for computing minimal and preference-protecting minimal contractions for finite as well as finitely representable preference relations. Nevertheless, also this work has several dissimilarities with ours. First, they focus only on contraction operators, and do not give a complexity analysis of preference revision in general. Second, the analysis carried out is meant for applications where the preference relation is required to be transitive and irreflexive (that is to say, a strict partial order).

As a matter of fact, our work comes full circle and show the computational limits of revision in defeasible contexts. Revision based on change of factual knowledge, which

essentially corresponds to an update operation, has been analysed in [11]. Other methods to revise, contract, or expand a defeasible theory were proposed in [2], studying how to revise the belief set of a theory based on the introduction of new rules. It is easy to verify that such revisions operates in polynomial time.

Preference handling in Defeasible Logic can gain much from typisation of preferences themselves. The notion of preference type and its algebraic structure has been studied previously and can be applied directly here [12]. Analogously, one of the possible directions of generalisation for the notion of preference is the notion of partial order, investigated at a combinatorial level by [13] and then studied from a computational viewpoint in [14].

References

1. Antoniou, G., Billington, D., Governatori, G., Maher, M.J.: Representation results for defeasible logic. *ACM Transactions on Computational Logic* **2**(2) (2001) 255–287
2. Billington, D., Antoniou, G., Governatori, G., Maher, M.J.: Revising nonmonotonic theories: The case of defeasible logic. In Burgard, W., Christaller, T., Cremers, A.B., eds.: *KI*. Volume 1701 of *Lecture Notes in Computer Science.*, Springer (1999) 101–112
3. Governatori, G., Olivieri, F., Scannapieco, S., Cristani, M.: Superiority based revision of defeasible theories. In Dean, M., Hall, J., Rotolo, A., Tabet, S., eds.: *RuleML*. Volume 6403 of *Lecture Notes in Computer Science.*, Springer (2010) 104–118
4. Governatori, G., Olivieri, F., Scannapieco, S., Cristani, M.: Revision of defeasible logic preferences. *CoRR* **abs/1206.5833** (2012)
5. Antoniou, G., Billington, D., Governatori, G., Maher, M.J.: Embedding defeasible logic into logic programming. *TPLP* **6**(6) (2006) 703–735
6. Maher, M.J.: Propositional defeasible logic has linear complexity. *Theory and Practice of Logic Programming* **1**(6) (2001) 691–711
7. Mindolin, D., Chomicki, J.: Minimal contraction of preference relations. In Fox, D., Gomes, C.P., eds.: *AAAI*, AAAI Press (2008) 492–497
8. Mindolin, D., Chomicki, J.: Contracting preference relations for database applications. *Artif. Intell.* **175**(7-8) (2011) 1092–1121
9. Governatori, G., Rotolo, A., Olivieri, F., Scannapieco, S.: Legal contractions: a logical analysis. In Francesconi, E., Verheij, B., eds.: *ICAIL*, ACM (2013) 63–72
10. Inoue, K., Sakama, C.: Abducing priorities to derive intended conclusions. In Dean, T., ed.: *IJCAI*, Morgan Kaufmann (1999) 44–49
11. Katsuno, H., Mendelzon, A.O.: Propositional knowledge base revision and minimal change. *Artificial Intelligence* **52**(3) (1991) 263–294
12. Cristani, M.: Many-sorted preference relations. In Fensel, D., Giunchiglia, F., McGuinness, D.L., Williams, M.A., eds.: *KR*, Morgan Kaufmann (2002) 265–276
13. Düntsch, I.: A microcomputer based system for small relation algebras. *J. Symb. Comput.* **18**(1) (1994) 83–86
14. Cristani, M., Hirsch, R.: The complexity of constraint satisfaction problems for small relation algebras. *Artif. Intell.* **156**(2) (2004) 177–196