New Tractable Cases in Default Reasoning from Conditional Knowledge Bases

Thomas Eiter and Thomas Lukasiewicz

Institut und Ludwig Wittgenstein Labor für Informationssysteme, TU Wien
Favoritenstraße 9–11, A-1040 Wien, Austria
{eiter, lukasiewicz}@kr.tuwien.ac.at

Abstract. We present new tractable cases for default reasoning from conditional knowledge bases. In detail, we introduce q-Horn conditional knowledge bases, which allow for a limited use of disjunction. We show that previous tractability results for ε-entailment, proper ε-entailment, and z- and z^+-entailment in the Horn case can be extended to the q-Horn case. Moreover, we present feedback-free-Horn conditional knowledge bases, which constitute a new, meaningful class of conditional knowledge bases. We show that the maximum entropy approach and lexicographic entailment are tractable in the feedback-free-Horn case. Our results complement and extend previous results, and contribute in refining the tractability/intractability frontier of default reasoning from conditional knowledge bases.

1 Introduction

A conditional knowledge base consists of a collection of strict statements in classical logic and a collection of defeasible rules (also called defaults). The former are statements that must always hold, while the latter are rules $\phi \rightarrow \psi$ that read as "generally, if ϕ then ψ." For example, the knowledge "penguins are birds" and "penguins do not fly" can be represented by strict sentences, while the knowledge "birds fly" should be expressed by a defeasible rule (since penguins are birds that do not fly).

The semantics of a conditional knowledge base KB is given by the set of all defaults that are plausible consequences of KB. The literature contains several different proposals for plausible consequence relations and extensive work on their desired properties. The core of these properties are the rationality postulates proposed by Kraus, Lehmann, and Magidor [17], which constitute a sound and complete axiom system for several classical model-theoretic entailment relations under uncertainty measures on worlds. More precisely, they characterize classical model-theoretic entailment under preferential structures, infinitesimal probabilities, possibility measures, and world rankings. Moreover, they characterize an entailment relation based on conditional objects. A survey of all these relationships is given in [4]. We will use the notion of ε-entailment to refer to these equivalent entailment relations.

Mainly to solve problems with irrelevant information, the notion of rational closure as a more adventurous notion of entailment has been introduced by Lehmann [20]. It is equivalent to entailment in system Z by Pearl [22] (which is generalized to variable strength defaults in system Z^+ by Goldszmidt and Pearl [15, 16]), to the least specific possibility entailment by Benferhat et al. [4], and to a conditional (modal) logic-based entailment by Lamarre [18]. Finally, mainly to solve problems with property inheritance from classes to exceptional subclasses, the maximum entropy approach was proposed by Goldszmidt et al. [13] (and recently generalized to variable strength defaults by Bourne and Parsons [7]); lexicographic entailment was introduced by Lehmann [19] and Benferhat et al. [3]; and conditional entailment was proposed by Geffner [11, 12].

However, while the semantic aspects of these formalisms are quite well understood, about their computational properties only partial results have been known so far. In previous work [9], we were filling some of these gaps by drawing a precise picture of the complexity of major formalisms for default reasoning from conditional knowledge bases. The main goal of this paper now is to complement this work by finding meaningful cases in which default reasoning from conditional knowledge bases is tractable. In particular, we aim at identifying nontrivial restrictions that can be checked efficiently and that guarantee sufficient expressiveness.

The main contributions of this paper can be summarized as follows:

- We introduce q-Horn conditional knowledge bases, which enrich in the spirit of [5] Horn conditional knowledge bases by allowing limited use of disjunction in both strict statements and defeasible rules. For example, a default $saturday \rightarrow hiking \vee shopping$, which informally expresses that on Saturday, someone is normally out for hiking or shopping, can be expressed in a q-Horn KB, but not in a Horn KB.
- We show that previous tractability results for ε-entailment [20, 16], proper ε-entailment [14], and z- and z^{+}-entailment [16] in the Horn case can be extended to the q-Horn case. Thus, in all these approaches, tractability is retained under a limited use of disjunction.
- We present feedback-free-Horn conditional knowledge bases, which restrict the literal-Horn case (where default rules are Horn-like) by requesting that, roughly speaking, default consequents do not fire back into the classical knowledge of KB and that the defaults can be grouped into non-interfering clusters of bounded size. We give some examples from the literature that underline the importance of the feedback-free-Horn case. In particular, we show that taxonomic hierarchies that are augmented by default knowledge can be expressed in the feedback-free-Horn case.
- We show that in the feedback-free-Horn case, default reasoning under $z^{\star}$-entailment [13], $z_{\mathrm{s}}^{\star}$-entailment [7], lex-entailment [3], and $\mathrm{lex}_{\mathrm{p}}$-entailment [19] is tractable. To our knowledge, no or only limited tractable cases [8] for these notions of entailment from conditional knowledge bases have been identified so far.
- Our tractability results for the feedback-free-Horn case are complemented by our proof that without a similar restriction on literal-Horn defaults, all the respective semantics remain intractable. In particular, this applies to the 1-literal-Horn case, in which each default is literal-Horn and has at most one atom in its antecedent, and the strict knowledge consists of Horn-clauses having at most two literals.

Note that detailed proofs of all results in this extended abstract are given in [10].

2 Preliminaries

2.1 Conditional Knowledge Bases

We assume a set of propositional *atoms* $At = \{p_1, \ldots, p_n\}$ with $n \geq 1$. We use $\perp$ and $\top$ to denote the propositional constants *false* and *true*, respectively. The set of *classical formulas* is the closure of $At \cup \{\perp, \top\}$ under the Boolean operations $\neg$ and $\wedge$. We use $(\phi \Rightarrow \psi)$ and $(\phi \vee \psi)$ to abbreviate $\neg(\phi \wedge \neg\psi)$ and $\neg(\neg\phi \wedge \neg\psi)$, respectively, and adopt the usual conventions to eliminate parentheses. A *literal* is an atom $p \in At$ or its negation $\neg p$. A *Horn clause* is a classical formula $\phi \Rightarrow \psi$, where ϕ is either $\top$ or a conjunction of atoms, and ψ is either $\perp$ or an atom. A *definite Horn clause* is a Horn clause $\phi \Rightarrow \psi$, where ψ is an atom.

A *conditional rule* (or *default*) is an expression $\phi \rightarrow \psi$, where ϕ and ψ are classical formulas. A *conditional knowledge base* is a pair $KB = (L, D)$, where L is a finite set of classical formulas and D is a finite set of defaults. Informally, L contains facts and rules that are certain, while D contains defeasible rules. In case $L = \emptyset$, we call KB a *default knowledge base*. A default $\phi \rightarrow \psi$ is *Horn* (resp., *literal-Horn*), if ϕ is either $\top$ or a conjunction of atoms, and ψ is a conjunction of Horn clauses (resp., ψ is a literal). A *definite literal-Horn default* is a literal-Horn default $\phi \rightarrow \psi$, where ψ is an atom.

Given a conditional knowledge base $KB = (L, D)$, a *strength assignment* σ on KB is a mapping that assigns each $d \in D$ an integer $\sigma(d) \geq 0$. A *priority assignment* on KB is a strength assignment π on KB with $\{\pi(d) \mid d \in D\} = \{0, 1, \ldots, k\}$ for some $k \geq 0$.

An *interpretation* (or *world*) is a truth assignment $I \colon At \rightarrow \{\mathbf{true}, \mathbf{false}\}$, which is extended to classical formulas as usual. We use $\mathcal{I}_{At}$ to denote the set of all worlds for At. The world I *satisfies* a classical formula ϕ, or I is a *model* of ϕ, denoted $I \models \phi$, iff $I(\phi) = \mathbf{true}$. I *satisfies* a default $\phi \rightarrow \psi$, or I is a *model* of $\phi \rightarrow \psi$, denoted $I \models \phi \rightarrow \psi$, iff $I \models \phi \Rightarrow \psi$. I *satisfies* a set K of classical formulas and defaults, or I is a *model* of K, denoted $I \models K$, iff I satisfies every member of K. The world I *verifies* a default $\phi \rightarrow \psi$ iff $I \models \phi \wedge \psi$. I *falsifies* a default $\phi \rightarrow \psi$, iff $I \models \phi \wedge \neg\psi$ (that is, $I \not\models \phi \rightarrow \psi$). A set of defaults D *tolerates* a default d *under* a set of classical formulas L iff $D \cup L$ has a model that verifies d. A set of defaults D is *under* L *in conflict* with a default $\phi \rightarrow \psi$ iff all models of $D \cup L \cup \{\phi\}$ satisfy $\neg\psi$.

A *world ranking* κ is a mapping $\kappa \colon \mathcal{I}_{At} \rightarrow \{0, 1, \ldots\} \cup \{\infty\}$ such that $\kappa(I) = 0$ for at least one world I. It is extended to all classical formulas ϕ as follows. If ϕ is satisfiable, then $\kappa(\phi) = \min\{\kappa(I) \mid I \in \mathcal{I}_{At}, I \models \phi\}$; otherwise, $\kappa(\phi) = \infty$. A world ranking κ is *admissible* with a conditional knowledge base (L, D) iff $\kappa(\neg\phi) = \infty$ for all $\phi \in L$, and $\kappa(\phi) < \infty$ and $\kappa(\phi \wedge \psi) < \kappa(\phi \wedge \neg\psi)$ for all defaults $\phi \rightarrow \psi \in D$. A *default ranking* σ on D maps each $d \in D$ to a nonnegative integer.

2.2 Semantics for Conditional Knowledge Bases

ε-Semantics (Adams [1] and Pearl [21]). We describe the notions of ε-consistency, ε-entailment, and proper ε-entailment in terms of world rankings.

A conditional knowledge base KB is *ε-consistent* iff there exists a world ranking that is admissible with KB. It is *ε-inconsistent* iff no such a world ranking exists.

A conditional knowledge base KB *ε-entails* a default $\phi \rightarrow \psi$ iff either $\kappa(\phi) = \infty$ or $\kappa(\phi \wedge \psi) < \kappa(\phi \wedge \neg\psi)$ for all world rankings κ admissible with KB. Moreover, KB *properly ε-entails* $\phi \rightarrow \psi$ iff KB ε-entails $\phi \rightarrow \psi$ and KB does not ε-entail $\phi \rightarrow \perp$.

The next theorem is a simple generalization of a result by Adams [1].

Theorem 2.1 (essentially [1]). *A conditional knowledge base (L, D)* ε-entails *a default $\phi \rightarrow \psi$ iff the conditional knowledge base $(L, D \cup \{\phi \rightarrow \neg\psi\})$ is ε-inconsistent.*

Systems Z and Z^+ (Pearl [22] and Goldszmidt and Pearl [15,16]). Entailment in system Z^+ applies to ε-consistent conditional knowledge bases $KB = (L, D)$ with strength assignment σ on KB. It is linked to a default ranking z^+ and a world ranking κ^+, which are the unique solution of the following system of equations:

$$z^+(d) = \sigma(d) + \kappa^+(\phi \wedge \psi) \qquad \text{(for all } d = \phi \rightarrow \psi \in D) \quad (1)$$

$$\kappa^+(I) = \begin{cases} \infty & \text{if } I \not\models L \\ 0 & \text{if } I \models L \cup D \\ 1 + \max\limits_{d \in D \,:\, I \not\models d} z^+(d) & \text{otherwise} \end{cases} \qquad \text{(for all } I \in \mathcal{I}_{At}) \quad (2)$$

A default $\phi \rightarrow \psi$ is z^+-*entailed* by (KB, σ) *at strength* τ iff either $\kappa^+(\phi) = \infty$ or $\kappa^+(\phi \wedge \psi) + \tau < \kappa^+(\phi \wedge \neg\psi)$.

Entailment in system Z is a special case of entailment in system Z^+. It applies to ε-consistent conditional knowledge bases KB. A default $\phi \rightarrow \psi$ is z-*entailed* by KB iff $\phi \rightarrow \psi$ is z^+-entailed by (KB, σ) at strength 0, where $\sigma(d) = 0$ for all $d \in D$.

Maximum Entropy (Goldszmidt et al. [13] and Bourne and Parsons [7]). The notion of $z^\star_\mathrm{s}$-entailment applies to ε-consistent conditional knowledge bases $KB = (L, D)$ with positive strength assignment σ. It is defined whenever the following system of equations (3) and (4) has a unique solution $z^\star_\mathrm{s}, \kappa^\star_\mathrm{s}$ with positive $z^\star_\mathrm{s}$:

$$\kappa^\star_\mathrm{s}(\phi \wedge \neg\psi) = \sigma(\phi \rightarrow \psi) + \kappa^\star_\mathrm{s}(\phi \wedge \psi) \qquad \text{(for all } d = \phi \rightarrow \psi \in D) \quad (3)$$

$$\kappa^\star_\mathrm{s}(I) = \begin{cases} \infty & \text{if } I \not\models L \\ 0 & \text{if } I \models L \cup D \\ \sum\limits_{d \in D \,:\, I \not\models d} z^\star_\mathrm{s}(d) & \text{otherwise} \end{cases} \qquad \text{(for all } I \in \mathcal{I}_{At}) \quad (4)$$

The uniqueness of $z^\star_\mathrm{s}$ and $\kappa^\star_\mathrm{s}$ is guaranteed by assuming that $\kappa^\star_\mathrm{s}$ is *robust* [7], which is the following property: for all distinct defaults $d_1, d_2 \in D$, it holds that all models I_1 and I_2 of L having smallest ranks in $\kappa^\star_\mathrm{s}$ such that $I_1 \not\models d_1$ and $I_2 \not\models d_2$, respectively, are different. That is, d_1 and d_2 do not have a common minimal falsifying model under L. We say KB is *robust* iff the system of equations given by (3) and (4) has a unique solution $z^\star_\mathrm{s}, \kappa^\star_\mathrm{s}$ such that $z^\star_\mathrm{s}$ is positive and $\kappa^\star_\mathrm{s}$ is robust. A default $\phi \rightarrow \psi$ is $z^\star_\mathrm{s}$-*entailed* by (KB, σ) *at strength* τ iff either $\kappa^\star_\mathrm{s}(\phi) = \infty$ or $\kappa^\star_\mathrm{s}(\phi \wedge \psi) + \tau \leq \kappa^\star_\mathrm{s}(\phi \wedge \neg\psi)$.

The notion of $z^\star$-entailment is a special case of $z^\star_\mathrm{s}$-entailment. It applies to ε-consistent minimal-core conditional knowledge bases $KB = (L, D)$ without strength assignment, where KB is *minimal-core* iff for each default $d \in D$ there exists a model I of $L \cup (D - \{d\})$ that falsifies d. A default $\phi \rightarrow \psi$ is $z^\star$-*entailed* by KB iff $\phi \rightarrow \psi$ is $z^\star_\mathrm{s}$-entailed by (KB, σ) at strength 1, where $\sigma(d) = 1$ for all $d \in D$.

Lexicographic Entailment (Lehmann [19] and Benferhat et al. [3]). Lexicographic entailment in [3] applies to conditional knowledge bases $KB = (L, D)$ with a priority assignment π on KB, which defines an ordered partition $(D_0, \ldots, D_k)$ of D by $D_i = \{d \in D \mid \pi(d) = i\}$, for all $i \leq k$. It is used to define a *preference ordering* on worlds as follows. A world I is π-*preferable* to a world I' iff there exists some $i \in \{0, \ldots, k\}$ such that $|\{d \in D_i \mid I \models d\}| > |\{d \in D_i \mid I' \models d\}|$ and $|\{d \in D_j \mid I \models d\}| = |\{d \in D_j \mid I' \models d\}|$ for all $i < j \leq k$. A model I of a set of classical formulas $\mathcal{F}$ is a π-*preferred* model of $\mathcal{F}$ iff no model of $\mathcal{F}$ is π-preferable to I. A default $\phi \rightarrow \psi$ is lex_p-*entailed* by (KB, π) iff ψ is satisfied in every π-preferred model of $L \cup \{\phi\}$.

The notion of lexicographic entailment in [19] is a special case of lexicographic entailment as above. It applies to ϵ-consistent conditional knowledge bases KB, and uses the default ranking z of KB in system Z as priority assignment. That is, a default $\phi \rightarrow \psi$ is *lex-entailed* by KB iff $\phi \rightarrow \psi$ is *lex*$_{\mathrm{p}}$*-entailed* by (KB, z).

2.3 Example

Consider the following conditional knowledge base $KB = (L, D)$, adapted from [16], which represents the strict knowledge "all penguins are birds", and the defeasible rules "generally, birds fly", "generally, penguins do not fly", "generally, birds have wings", "generally, penguins live in the arctic", and "generally, flying animals are mobile".

$$L = \{penguin \Rightarrow bird\}\ ,$$
$$D = \{bird \rightarrow fly,\ penguin \rightarrow \neg fly,\ bird \rightarrow wings,\ penguin \rightarrow arctic,\ fly \rightarrow mobile\}\ .$$

We would like KB to entail "generally, birds are mobile" (as birds generally fly, and flying animals are generally mobile) and "generally, red birds fly" (as the property "red" is not mentioned at all in KB and should thus be considered irrelevant to the flying ability of birds). Moreover, KB should entail "generally, penguins have wings" (as the set of all penguins is a subclass of the set of all birds, and thus penguins should inherit all properties of birds), and "generally, penguins do not fly" (as properties of more specific classes should override inherited properties of less specific classes).

The corresponding behavior of ε-, z-, $z^{\star}$-, and lex-entailment is shown in Table 1. In detail, *bird* $\rightarrow$ *mobile* is a plausible consequence of KB under all notions of entailment except for ε-entailment. Moreover, in this example, every notion of entailment except for ε-entailment ignores irrelevant information, while every notion of entailment except for ε- and z-entailment shows property inheritance from the class of all birds to the exceptional subclass of all penguins. Finally, the default *penguin* $\rightarrow \neg$*fly* is entailed by KB under all notions of entailment.

For instance, let us verify that *penguin* $\rightarrow \neg$*fly* is ε-entailed by KB. By Theorem 2.1, we have to check that $(L, D \cup \{penguin \rightarrow fly\})$ is ε-inconsistent. But this is indeed the case, since there is no world ranking κ that satisfies $\kappa(penguin) < \infty$ as well as $\kappa(penguin \wedge \neg fly) < \kappa(penguin \wedge fly)$ and $\kappa(penguin \wedge fly) < \kappa(penguin \wedge \neg fly)$.

Table 1. Plausible consequences of KB under different semantics

	bird $\rightarrow$ *mobile*	*red* $\wedge$ *bird* $\rightarrow$ *fly*	*penguin* $\rightarrow$ *wings*	*penguin* $\rightarrow \neg$*fly*
ε-entailment	$-$	$-$	$-$	$+$
z-entailment	$+$	$+$	$-$	$+$
$z^{\star}$-entailment	$+$	$+$	$+$	$+$
lex-entailment	$+$	$+$	$+$	$+$

3 Overview of Tractability Results

3.1 Problem Statements

A *default reasoning problem* is a pair (KB, d), where $KB = (L, D)$ is a conditional knowledge base and d is a default. It is *Horn* (resp., *literal-Horn*) iff L is a finite set of Horn clauses, D is a finite set of Horn (resp., literal-Horn) defaults, and d is a Horn (resp., literal-Horn) default. In case of z^+- and $z^\star_{\mathrm{s}}$-entailment, we assume that KB and d have additionally a strength assignment $\sigma(KB)$ and a strength $\tau(d)$, respectively. In case of $\mathrm{lex}_{\mathrm{p}}$-entailment, KB has in addition a priority assignment $\pi(KB)$.

Informally, a default reasoning problem represents the input for the entailment problem under a fixed semantics $\mathcal{S}$. We tacitly assume that KB satisfies any preconditions that the definition of $\mathcal{S}$-entailment in the previous section may request.

We consider the following problems:

- ENTAILMENT**:** Given a default reasoning problem (KB, d), decide whether KB entails d under some fixed semantics $\mathcal{S}$. In case of z^+- and $z^\star_{\mathrm{s}}$-entailment, decide whether d is z^+- and $z^\star_{\mathrm{s}}$-entailed, respectively, by $(KB, \sigma(KB))$ at strength $\tau(d)$. In case of $\mathrm{lex}_{\mathrm{p}}$-entailment, we are asked whether d is $\mathrm{lex}_{\mathrm{p}}$-entailed by $(KB, \pi(KB))$.
- RANKING**:** Given a conditional knowledge base KB, compute the default ranking R of KB according to some fixed semantics $\mathcal{S}$ (that is, the rank of each $d \in D$).
- RANK-ENTAILMENT**:** Same as entailment, but the (unique) default ranking R of KB according to some fixed semantics $\mathcal{S}$ is part of the problem input.

3.2 Previous Tractability Results

Previous results on the tractability/intractability frontier can be described as follows.

Deciding ε-entailment is intractable in the general case [20] and tractable in the Horn case [20, 16]. Similarly, deciding proper ε-entailment is intractable in the general case [9] and tractable in the Horn case [14]. Moreover, the problems ENTAILMENT, RANKING, and RANK-ENTAILMENT for systems Z and Z^+ are intractable in the general case [9] and tractable in the Horn case [16].

The problems ENTAILMENT, RANKING, and RANK-ENTAILMENT for the semantics $z^\star$ and $z^\star_{\mathrm{s}}$ are intractable even in the literal-Horn case [9]. Moreover, also deciding lex- and $\mathrm{lex}_{\mathrm{p}}$-entailment is intractable in the literal-Horn case [9]. To our knowledge, no or only limited tractable cases for these notions of entailment have been identified so far (a limited tractable case for $\mathrm{lex}_{\mathrm{p}}$-entailment has been presented in [8]).

3.3 New Tractability Results

It would be interesting to know whether the tractability results for ε-entailment, proper ε-entailment, and z- and z^+ entailment in the Horn case can be extended to more expressive classes of problems. Moreover, it would be interesting to know whether there are meaningful tractable classes of problems for $z^\star$-, $z^\star_{\mathrm{s}}$-, lex-, and $\mathrm{lex}_{\mathrm{p}}$-entailment.

Concerning the first issue, we introduce the class of q-Horn conditional knowledge bases. This class generalizes Horn conditional knowledge bases syntactically by allowing a restricted use of disjunction, and contains instances that cannot be represented in Horn conditional knowledge bases. As we show, the tractability results for Horn conditional knowledge bases extend to q-Horn conditional knowledge bases.

Finding meaningful tractable cases for the more sophisticated semantics for conditional knowledge bases is more challenging. A natural attempt is to show that a further restriction of the literal-Horn case leads to tractability. An obvious candidate restriction is bounding the size of the bodies in the strict and classical rules to at most one atom.

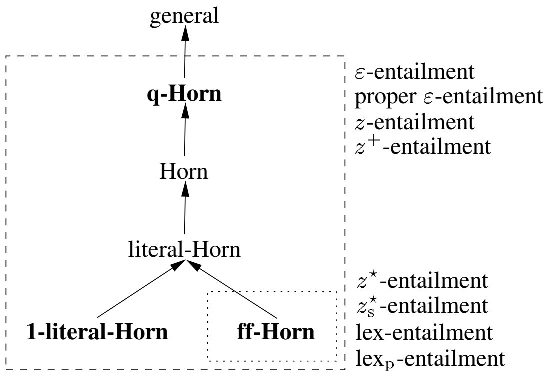

Fig. 1. Tractability of ENTAILMENT

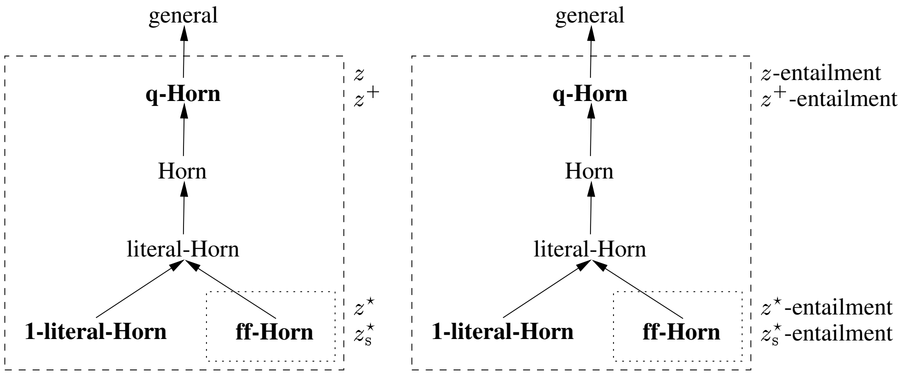

Fig. 2. Tractability of RANKING and RANK-ENTAILMENT

Unfortunately, this does not buy tractability. An analysis of our proof reveals that the interaction of the defaults among each other and with the classical background knowledge must be controlled such that interferences have a local effect. This leads us to the class of feedback-free-Horn (ff-Horn) default reasoning problems. As we show, tractability is gained on this class for all the intractable semantics discussed here.

The hierarchy of all classes of conditional knowledge bases considered in this paper, along with the corresponding tractability results, is given in Figures 1–2.

4 Q-Horn

4.1 Motivating Example

Q-Horn conditional knowledge bases generalize Horn conditional knowledge bases by allowing a limited form of disjunction, which is illustrated by the following example.

Example 4.1. Assume that John is looking for Mary. Unfortunately, he did not find her at home. So, he is wondering where she might be. He knows that Mary might have tea with her friends, that she might be in the library, or that she might play tennis. He also knows that these scenarios are pairwise exclusive and not exhaustive. Moreover, John knows that "generally, in the afternoon, Mary is having tea with her friends or she is in the library" and that "generally, on Friday afternoon, Mary plays tennis". This knowledge can be expressed by the following $KB = (L, D)$:

$$L = \{\neg tea \vee \neg library,\ \neg tea \vee \neg tennis,\ \neg library \vee \neg tennis\}\ ,$$
$$D = \{afternoon \rightarrow tea \vee library,\ Friday \wedge afternoon \rightarrow tennis\}\ .$$

Assume that it is Friday afternoon and that John is wondering whether he should go to the library to look for Mary. That is, does KB entail $Friday \wedge afternoon \rightarrow library$?

4.2 Definitions

A *clause* is a disjunction of literals. A default $\phi \rightarrow \psi$ is *clausal* iff ϕ is either $\top$ or a conjunction of literals, and ψ is a conjunction of clauses. A conditional knowledge base $KB = (L, D)$ is *clausal* iff L is a finite set of clauses and D is a finite set of clausal defaults. A default reasoning problem (KB, d) is *clausal* iff both KB and d are clausal.

A classical formula ϕ is in *conjunctive normal form* (or *CNF*) iff ϕ is either $\top$ or a conjunction of clauses. We use the operator $\sim$ to map each atom a to its negation $\neg a$, and each negated atom $\neg a$ to a. We define a mapping $\mathcal{N}$ that associates each clausal default d with a classical formula in CNF as follows. If d is of the form $\top \rightarrow c_1 \wedge \cdots \wedge c_n$ with clauses $c_1, \ldots, c_n$, then $\mathcal{N}(d) = c_1 \wedge \cdots \wedge c_n$. If d is of the form $l_1 \wedge \cdots \wedge l_m \rightarrow c_1 \wedge \cdots \wedge c_n$ with literals $l_1, \ldots, l_m$ and clauses $c_1, \ldots, c_n$, then $\mathcal{N}(d)$ is the conjunction of all $\sim l_1 \vee \cdots \vee \sim l_m \vee c_i$ with $i \in \{1, \ldots, n\}$. We extend $\mathcal{N}$ to classical formulas in CNF ϕ by $\mathcal{N}(\phi) = \phi$. We extend $\mathcal{N}$ to finite sets K of classical formulas in CNF and clausal defaults as follows. Let K' denote the set of all $k \in K$ with $\mathcal{N}(k) \neq \top$. If $K' \neq \emptyset$, then $\mathcal{N}(K)$ is the conjunction of all $\mathcal{N}(k)$ with $k \in K'$. Otherwise, $\mathcal{N}(K) = \top$.

A *partial assignment* S is a set of literals such that for every atom $a \in At$ at most one of the literals a and $\neg a$ is in S. A classical formula in CNF ϕ is *q-Horn* [5] iff there exists a partial assignment S such that (i) each clause in ϕ contains at most two literals outside of S, and (ii) if a clause in ϕ contains exactly two literals $u, v \notin S$, then neither $\sim u$ nor $\sim v$ belongs to S. Note that every conjunction of Horn clauses is q-Horn.

A finite set K of classical formulas in CNF and clausal defaults is *q-Horn* iff $\mathcal{N}(K)$ is q-Horn. A conditional knowledge base $KB = (L, D)$ is *q-Horn* iff KB is clausal and $L \cup D$ is q-Horn. Clearly, every Horn KB is q-Horn, but not vice versa. A default reasoning problem (KB, d) is *q-Horn*, if KB is q-Horn and d is a clausal default.

Example 4.2. The conditional knowledge base $KB = (L, D)$ of Example 4.1 is q-Horn. In detail, the classical formula $\mathcal{N}(L \cup D)$ associated with KB is given by:

$$\begin{aligned}\mathcal{N}(L \cup D) = {} & (\neg tea \vee \neg library) \wedge (\neg tea \vee \neg tennis) \wedge (\neg library \vee \neg tennis) \wedge \\ & (\neg afternoon \vee tea \vee library) \wedge (\neg Friday \vee \neg afternoon \vee tennis)\ .\end{aligned}$$

A partial assignment that satisfies (i) and (ii) is given by $\{\neg Friday, \neg afternoon\}$. That is, $\mathcal{N}(L \cup D)$ is q-Horn. Since KB is also clausal, it thus follows that KB is q-Horn.

Note that KB can be made Horn by "renaming" atoms, in particular, by replacing the atom *library* by a negated new atom $\overline{library}$, where $\overline{library}$ stands for $\neg library$. However, if the scenarios were exhaustive and thus the clause $library \vee tea \vee tennis$ is in KB, then no Horn renaming of KB is possible. But, the resulting KB is still q-Horn.

The *size* of a classical formula in CNF ϕ, denoted $\|\phi\|$, is defined as the number of occurrences of literals in ϕ. We use $|\phi|$ to denote the number of clauses in ϕ. The *size* of a clausal default $d = \phi \to \psi$, denoted $\|d\|$, is defined as $\|\phi\| + \|\psi\|$. The *size* of a finite set of clauses L, denoted $\|L\|$, is defined as the size of $\mathcal{N}(L)$. The *size* of a clausal $KB = (L, D)$, denoted $\|KB\|$, is defined as the size of $\mathcal{N}(L \cup D)$. We use $|D|$ to denote the cardinality of D.

4.3 Q-Horn Formulas

The problems of deciding whether a q-Horn formula is satisfiable and of recognizing q-Horn formulas are both tractable and can in fact be solved in linear time.

Proposition 4.3 (see [5, 6]). *a) Given a q-Horn formula ϕ, deciding whether ϕ is satisfiable can be done in time $O(\|\phi\|)$. b) Given a classical formula in CNF ϕ, deciding whether ϕ is q-Horn can be done in time $O(\|\phi\|)$.*

By this result, it follows easily that also q-Horn conditional knowledge bases can be recognized in linear time.

Theorem 4.4. *Given a clausal conditional knowledge base $KB = (L, D)$, deciding whether KB is q-Horn can be done in time $O(\|KB\|)$.*

4.4 ε-Semantics

The following theorem shows that deciding whether a q-Horn KB is ε-consistent is tractable. The proof of this result is based on the fact that checking the ε-consistency of KB is reducible to a polynomial number of classical satisfiability tests. By closure properties of q-Horn formulas, it then follows that for q-Horn KB, each satisfiability test is done on a q-Horn formula and thus possible in polynomial time.

Theorem 4.5. *Given a q-Horn conditional knowledge base $KB = (L, D)$, deciding whether KB is ε-consistent is possible in time $O(|D|^2\, \|KB\|)$.*

The next result shows that deciding ε-entailment is tractable in the q-Horn case.

Theorem 4.6. *Given a q-Horn default reasoning problem $(KB, d) = ((L, D), \phi \to \psi)$, deciding whether KB ε-entails d is possible in time $O((\|\phi\| + |\psi|)\, |D|^2\, (\|KB\| + \|\psi\|))$.*

Finally, deciding proper ε-entailment is also tractable in the q-Horn case.

Theorem 4.7. *Given a q-Horn default reasoning problem $(KB, d) = ((L, D), \phi \to \psi)$, deciding whether KB properly ε-entails d is possible in time $O((\|\phi\| + |\psi|)\, |D|^2$ $(\|KB\| + \|\psi\|))$.*

4.5 Systems Z and Z^+

We next focus on entailment in systems Z and Z^+. The following result, which can be proved in a similar way as Theorems 4.5–4.7, shows that computing the default ranking z^+ is tractable in the q-Horn case. Since system Z^+ properly generalizes system Z, this result shows also that computing the default ranking z is tractable in the q-Horn case.

Theorem 4.8. *Given an ε-consistent q-Horn conditional knowledge base $KB = (L, D)$ with strength assignment σ, the default ranking z^+ can be computed in polynomial time.*

Finally, the following theorem shows that deciding z^+-entailment is tractable in the q-Horn case. Again, since system Z^+ properly generalizes system Z, this result shows also that deciding z-entailment is tractable in the q-Horn case. Trivially, these tractability results remain true when z^+ and z, respectively, are part of the input, that is, for RANK-ENTAILMENT.

Theorem 4.9. *Given a q-Horn default reasoning problem $(KB, d) = ((L, D), \phi \rightarrow \psi)$, where KB is ε-consistent and has a strength assignment σ, deciding whether (KB, σ) z^+-entails d at a given strength $\tau \geq 0$ can be done in polynomial time.*

5 Feedback-free-Horn

5.1 Intractability Results for 1-Literal-Horn Case

How do we obtain tractability of deciding s-entailment, where $s \in \{z^\star, z^\star_{\mathrm{s}}, \mathrm{lex}, \mathrm{lex}_{\mathrm{p}}\}$? In particular, are there any syntactic restrictions on default reasoning problems that give tractability? We could, for example, further restrict literal-Horn defaults by limiting the number of atoms in the antecedent of each default as follows. A default $\phi \rightarrow \psi$ is *1-literal-Horn* iff ϕ is either $\top$ or an atom, and ψ is a literal. A *1-Horn clause* is a classical formula $\phi \Rightarrow \psi$, where ϕ is either $\top$ or an atom, and ψ is a literal. A conditional knowledge base $KB = (L, D)$ is *1-literal-Horn* iff L is a finite set of 1-Horn clauses and D is a finite set of 1-literal-Horn defaults. A default reasoning problem (KB, d) is 1-literal-Horn iff both KB and d are 1-literal-Horn.

Unfortunately, the following theorem shows that deciding $z^\star$-entailment is still (presumably) intractable even for this very restricted kind of default reasoning problems.

Theorem 5.1. *Given a 1-literal-Horn KB, which is ε-consistent and minimal-core, and a 1-literal-Horn default d, deciding whether KB $z^\star$-entails d is* co-NP-*hard.*

Informally, this intractability is due to the fact that the default knowledge generally does not fix a unique instantiation of the atoms to truth values, in particular, when defaults "fire back" into the bodies of other defaults, and when defaults are logically related through their heads.

Since $z^\star_{\mathrm{s}}$-entailment is a proper generalization of $z^\star$-entailment, it immediately follows that deciding $z^\star_{\mathrm{s}}$-entailment is (presumably) intractable in the 1-literal-Horn case.

Corollary 5.2. *Given a 1-literal-Horn conditional knowledge base KB, which is ε-consistent and robust, a strength assignment σ on KB, a 1-literal-Horn default d, and a strength τ, deciding whether (KB, σ) $z^\star_{\mathrm{s}}$-entails d at strength τ is* co-NP-*hard.*

The following theorem shows that also deciding lex- and $\mathrm{lex}_{\mathrm{p}}$-entailment is (presumably) intractable in the 1-literal-Horn case.

Theorem 5.3. *a) Given an ε-consistent 1-literal-Horn conditional knowledge base KB and a 1-literal-Horn default d, deciding whether KB lex-entails d is* co-NP-*hard.*

b) Given a 1-literal-Horn conditional knowledge base KB with priority assignment π and a 1-literal-Horn default d, deciding whether (KB, π) lex_{p}-entails d is co-NP-*hard.*

5.2 Motivating Examples

We will see that deciding s-entailment, where $s \in \{z^\star, z^\star_{\mathrm{s}}, \mathrm{lex}, \mathrm{lex_p}\}$, becomes tractable, if we assume that the default reasoning problems can be sensibly decomposed into smaller problems of size bounded by a constant. We now give some examples to illustrate the main ideas behind this kind of decomposability. In the following examples, we assume that conditional knowledge bases are implicitly associated with a strength assignment σ (resp., priority assignment π), when $s = z^\star_{\mathrm{s}}$ (resp., $s = \mathrm{lex_p}$).

Example 5.4. Take again $KB = (L, D)$ of Section 2.3. Assume that we are wondering whether KB s-entails *penguin* $\rightarrow$ *fly*, *red* $\wedge$ *bird* $\rightarrow$ *fly*, *bird* $\rightarrow$ *mobile*, *penguin* $\rightarrow$ *arctic*, or *penguin* $\rightarrow$ *wings*, where $s \in \{z^\star, z^\star_{\mathrm{s}}, \mathrm{lex}, \mathrm{lex_p}\}$. As it turns out, each of these problems can be reduced to one classical reasoning problem and one default reasoning problem. More precisely, the former is done w.r.t. the set of atoms $\{$*penguin*, *bird*, *red*$\}$, which refers to the atoms in L and the antecedent of the query default, while the latter is done w.r.t. the sets of atoms $\{$*fly*, *mobile*$\}$, $\{$*arctic*$\}$, and $\{$*wings*$\}$, respectively, by sensibly eliminating irrelevant defaults and simplifying the remaining defaults by instantiating atoms to truth values. For instance, deciding whether KB s-entails *red* $\wedge$ *bird* $\rightarrow$ *fly* is reduced to the classical reasoning problem of computing the least model of $L \cup \{$*red* $\wedge$ *bird*$\}$ and the default reasoning problem of deciding whether $(\{$*red*, *bird*, $\neg$*penguin*$\}, \{$*bird* $\rightarrow$ *fly*, *fly* $\rightarrow$ *mobile*$\})$ s-entails *red* $\wedge$ *bird* $\rightarrow$ *fly*.

We next consider a taxonomic hierarchy adorned with some default knowledge [2].

Example 5.5. The strict knowledge "all birds and fish are animals", "all penguins and sparrows are birds", "no bird is a fish", "no penguin is a sparrow", and the defeasible knowledge "generally, animals do not swim", "generally, fish swims", and "generally, penguins swim" can be represented by the following $KB = (L, D)$:

$$\begin{aligned} L = \{&\mathit{bird} \Rightarrow \mathit{animal},\ \mathit{fish} \Rightarrow \mathit{animal},\ \mathit{penguin} \Rightarrow \mathit{bird}, \\ &\mathit{sparrow} \Rightarrow \mathit{bird},\ \mathit{bird} \Rightarrow \neg\mathit{fish},\ \mathit{penguin} \Rightarrow \neg\mathit{sparrow}\}\,, \\ D = \{&\mathit{animal} \rightarrow \neg\mathit{swims},\ \mathit{fish} \rightarrow \mathit{swims},\ \mathit{penguin} \rightarrow \mathit{swims}\}\,. \end{aligned}$$

Do sparrows generally swim? That is, does KB s-entail *sparrow* $\rightarrow$ *swims*, where $s \in \{z^\star, z^\star_{\mathrm{s}}, \mathrm{lex}, \mathrm{lex_p}\}$? This default reasoning problem can be reduced to one classical reasoning problem w.r.t. the set of atoms $\{$*animal*, *bird*, *fish*, *sparrow*, *penguin*$\}$ and one default reasoning problem w.r.t. the set of atoms $\{$*swims*$\}$. In detail, we first compute the least model of $L \cup \{$*sparrow*$\}$ and then decide whether $(\{$*sparrow*, *bird*, *animal*, $\neg$*fish*, $\neg$*penguin*$\}$, $\{$*animal* $\rightarrow \neg$*swims*$\})$ s-entails *sparrow* $\rightarrow$ *swims*.

5.3 Definitions

Suppose that for a literal-Horn conditional knowledge base $KB = (L, D)$, there exists a set of atoms $At_a \subseteq At$ such that L is defined over At_a and that all consequents of definite literal-Horn defaults in D are defined over $At - At_a$. The greatest such At_a, which clearly exists, is called the *activation set* of KB. Intuitively, in any "context" given by L and ϕ, where ϕ is either $\top$ or a conjunction of atoms from At, all those atoms in At_a that are not logically entailed by $L \cup \{\phi\}$ can be safely set to **false** in the preferred models of $L \cup \{\phi\}$.

For At_a, there is a greatest partition $\{At_1, \ldots, At_n\}$ of $At - At_a$ such that every $d \in D$ is defined over some $At_a \cup At_i$ with $i \in \{1, \ldots, n\}$, which we call the *default partition* of KB. We say $KB = (L, D)$ is *k-feedback-free-Horn* (or *k-ff-Horn*) iff it is literal-Horn, it has an activation set At_a, and it has a default partition $\{At_1, \ldots, At_n\}$ such that every At_i with $i \in \{1, \ldots, n\}$ has a cardinality of at most k.

Example 5.6. The conditional knowledge base KB of Example 5.5 is 1-ff-Horn. More precisely, its activation set (resp., default partition) is given by $\{$*animal*, *bird*, *fish*, *sparrow*, *penguin*$\}$ (resp., $\{\{$*swims*$\}\}$).

Moreover, KB of Example 5.4 is 2-ff-Horn. Its activation set (resp., default partition) is given by $\{$*penguin*, *bird*, *red*$\}$ (resp., $\{\{$*fly*, *mobile*$\}$, $\{$*arctic*$\}$, $\{$*wings*$\}\}$).

For sets of Horn clauses L, we use L^+ to denote the set of all definite Horn clauses in L. For sets of literal-Horn defaults D, we use D^+ to denote the set of all definite literal-Horn defaults in D. Assume additionally that $d = \phi \to \psi$ is a literal-Horn default. Then, a literal-Horn default $\alpha \to a$ (resp., $\alpha \to \neg a$) with $a \in At$ is *active* w.r.t. (L, D) and d iff $L^+ \cup D^+ \cup \{\phi\} \models \alpha$ (resp., $L^+ \cup D^+ \cup \{\phi\} \models \alpha \wedge a$).

A default reasoning problem $(KB, d) = ((L, D), \phi \to \psi)$ is *k-ff-Horn*, where $k \geq 1$, iff (i) it is literal-Horn, and (ii) $(L, D_a \cup \{d\})$ has an activation set At_a and a default partition $\{At_1, \ldots, At_n\}$ such that d is defined over some $At_a \cup At_j$ with $|At_j| \leq k$, where D_a is the set of all active defaults in D w.r.t. KB and d. The class *k-ff-Horn* consists of all k-ff-Horn default reasoning problems; we define the class *feedback-free-Horn* (or *ff-Horn*) by ff-Horn = $\bigcup_{k \geq 1}$ k-ff-Horn.

Example 5.7. Consider the literal-Horn default reasoning problem (KB, d) with $KB = (L, D)$ as in Example 5.4 and $d =$ *red* $\wedge$ *bird* $\to$ *fly*. The set D_a of active defaults in D w.r.t. KB and d is given by $D_a = \{$*bird* $\to$ *fly*, *bird* $\to$ *wings*, *fly* $\to$ *mobile*$\}$.

Now, $(L, D_a \cup \{d\})$ has the activation set $At_a = \{$*penguin*, *bird*, *red*, *arctic*$\}$ and the default partition $\{At_1, At_2\}$, where $At_1 = \{$*fly*, *mobile*$\}$ and $At_2 = \{$*wings*$\}$. Moreover, d is defined over $At_a \cup At_1$ with $|At_1| = 2$. That is, (KB, d) is 2-ff-Horn.

For Horn conditional knowledge bases $KB = (L, D)$ with activation set At_a, and classical formulas α that are either $\top$ or conjunctions of atoms from At, we define the classical formula $\alpha^\star$ as follows. If $L \cup \{\alpha\}$ is satisfiable, then $\alpha^\star$ is the conjunction of all $b \in At$ with $L \cup \{\alpha\} \models b$ and all $\neg b$ with $b \in At_a$ and $L \cup \{\alpha\} \not\models b$. Otherwise, we define $\alpha^\star = \bot$. Moreover, for satisfiable $L \cup \{\alpha\}$, we define the world $I^\star_\alpha$ over the activation set At_a by $I^\star_\alpha(b) = \mathbf{true}$ iff $L \cup \{\alpha\} \models b$, for all $b \in At_a$.

5.4 Recognizing Feedback-free-Horn

Both recognizing k-ff-Horn conditional knowledge bases, and computing their activation set and default partition are efficiently possible using standard methods.

Theorem 5.8. *a) Given a literal-Horn conditional knowledge base KB and an integer $k \geq 1$, deciding whether KB is k-ff-Horn can be done in linear time.*

b) Given a k-ff-Horn conditional knowledge base KB, computing the activation set At_a and the default partition $\{At_1, \ldots, At_n\}$ can be done in linear time.

Moreover, recognizing k-ff-Horn default reasoning problems is also efficiently possible.

Theorem 5.9. *a) Given a literal-Horn default reasoning problem* (KB, d)*, and an integer* $k \geq 1$*, deciding whether* (KB, d) *is k-ff-Horn can be done in linear time.*

b) Given a k-ff-Horn default reasoning problem (KB, d) *with* $KB = (L, D)$*, computing the set* D_a *of active defaults in* D *w.r.t.* KB *and* d *can be done in linear time.*

5.5 Maximum Entropy Semantics

In the sequel, let $KB = (L, D)$ be an ε-consistent k-ff-Horn conditional knowledge base with positive strength assignment σ. Let At_a denote the activation set of KB, and let $(At_1, \ldots, At_n)$ be the default partition of KB. Let $z^\star_{\mathrm{s}}$ be a ranking that maps each $d \in D$ to a positive integer, and let $\kappa^\star_{\mathrm{s}}$ be defined by (4).

For each $i \in \{1, \ldots, n\}$, let D_i denote the set of all defaults in D that are defined over $At_a \cup At_i$. Let the function $\kappa^\star_{\mathrm{s},i}$ on worlds I over At be defined as follows:

$$\kappa^\star_{\mathrm{s},i}(I) = \begin{cases} \infty & \text{if } I \not\models L \\ 0 & \text{if } I \models L \cup D_i \\ \sum\limits_{d \in D_i \,:\, I \not\models d} z^\star_{\mathrm{s}}(d) & \text{otherwise.} \end{cases} \tag{5}$$

In order to compute the default ranking $z^\star_{\mathrm{s}}$, we have to compute ranks of the form $\kappa^\star_{\mathrm{s}}(\alpha \wedge \beta_1 \wedge \cdots \wedge \beta_n)$, where α is either $\top$ or a conjunction of atoms from At_a, and each β_i is either $\top$ or a conjunction of literals over At_i. It can now be shown that such $\kappa^\star_{\mathrm{s}}(\alpha \wedge \beta_1 \wedge \cdots \wedge \beta_n)$ coincide with $\sum_{i=1}^{n} \kappa^\star_{\mathrm{s},i}(\alpha^\star \wedge \beta_i)$.

Using this result, it can be shown that computing the default ranking $z^\star_{\mathrm{s}}$ is tractable in the k-ff-Horn case. Since $z^\star_{\mathrm{s}}$ is a proper generalization of $z^\star$, this result shows also that computing the default ranking $z^\star$ is tractable in the k-ff-Horn case.

Theorem 5.10. *Let* $k > 0$ *a fixed integer. Given an* ε*-consistent k-ff-Horn* $KB = (L, D)$ *with positive strength assignment* σ*, computing the default ranking* $z^\star_{\mathrm{s}}$ *for KB, if KB is robust, and returning nil otherwise, can be done in polynomial time.*

In the sequel, let $(KB, d) = ((L, D), \phi \rightarrow \psi)$ be a k-ff-Horn default reasoning problem with ε-consistent and robust KB. Let σ be a positive strength assignment on KB. Let D_a be the set of active defaults in D w.r.t. KB and d, let At_a be the activation set of $(L, D_a \cup \{d\})$, and let $(At_1, \ldots, At_n)$ be the default partition of $(L, D_a \cup \{d\})$. Let $z^\star_{\mathrm{s}}, \kappa^\star_{\mathrm{s}}$ be the unique solution of (3) and (4).

For every $i \in \{1, \ldots, n\}$, let $\overline{D}_i$ denote the set of all defaults in D_a that are defined over $At_a \cup At_i$, and let σ_i be the restriction of σ to $\overline{D}_i$. Let $\overline{z}^\star_{\mathrm{s},i}$ map each default in $\overline{D}_i$ to a positive integer, and let the function $\overline{\kappa}^\star_{\mathrm{s},i}$ on worlds I over At be defined by:

$$\overline{\kappa}^\star_{\mathrm{s},i}(I) = \begin{cases} \infty & \text{if } I \not\models L \\ 0 & \text{if } I \models L \cup \overline{D}_i \\ \sum\limits_{d \in \overline{D}_i \,:\, I \not\models d} \overline{z}^\star_{\mathrm{s},i}(d) & \text{otherwise.} \end{cases} \tag{6}$$

It can be shown that in order to decide whether KB $z^\star_{\mathrm{s}}$-entails d at given strength $\tau > 0$, it is sufficient to know all $z^\star_{\mathrm{s}}(d)$ with $d \in \overline{D}_j$, where $j \in \{1, \ldots, n\}$ such that d is defined over $At_a \cup At_j$. Moreover, it can be shown that the restriction of $z^\star_{\mathrm{s}}$ to $\overline{D}_j$ coincides with the default ranking for $(L, \overline{D}_j)$ under the strength assignment σ_j.

Using these results, it can be shown that deciding $z_s^\star$-entailment is tractable in the k-ff-Horn case. Again, since $z_s^\star$ properly generalizes $z^\star$, this result shows also that deciding $z^\star$-entailment is tractable in the k-ff-Horn case. Trivially, these tractability results remain true when $z_s^\star$ and $z^\star$, respectively, are part of the input.

Theorem 5.11. *Let $k > 0$ be fixed. Given a k-ff-Horn default reasoning problem $(KB, d) = ((L, D), \phi \to \psi)$, where KB is ε-consistent and robust, and a positive strength assignment σ on KB, deciding whether (KB, σ) $z_s^\star$-entails d at given strength $\tau > 0$ can be done in polynomial time.*

Example 5.12. Let the 2-ff-Horn default reasoning problem (KB, d) be given by $KB = (L, D)$ of Example 5.4 and $d = \mathit{red} \wedge \mathit{bird} \to \mathit{fly}$. Let $\sigma(\delta) = 1$ for all $\delta \in D$.

Now, d is $z_s^\star$-entailed by (KB, σ) at strength τ iff either (i) $L \cup \{\mathit{red} \wedge \mathit{bird}, \neg \mathit{fly}\}$ is unsatisfiable, or (ii) both $L \cup \{\mathit{red} \wedge \mathit{bird}, \mathit{fly}\}$ and $L \cup \{\mathit{red} \wedge \mathit{bird}, \neg\mathit{fly}\}$ are satisfiable, and $\kappa_s^\star(\mathit{red} \wedge \mathit{bird} \wedge \mathit{fly}) + \tau \leq \kappa_s^\star(\mathit{red} \wedge \mathit{bird} \wedge \neg\mathit{fly})$. It can be shown that the latter is equivalent to $\overline{\kappa}_{s,1}^\star((\mathit{red} \wedge \mathit{bird})^\star \wedge \mathit{fly}) + \tau \leq \overline{\kappa}_{s,1}^\star((\mathit{red} \wedge \mathit{bird})^\star \wedge \neg\mathit{fly})$, that is, $\overline{\kappa}_{s,1}^\star(\mathit{red} \wedge \mathit{bird} \wedge \mathit{fly}) + \tau \leq \overline{\kappa}_{s,1}^\star(\mathit{red} \wedge \mathit{bird} \wedge \neg\mathit{fly})$, where $\overline{\kappa}_{s,1}^\star$ is defined through the default ranking $\overline{z}_{s,1}^\star$ for $(L, D_1) = (L, \{\mathit{bird} \to \mathit{fly}, \mathit{fly} \to \mathit{mobile}\})$ under $\sigma_1 = \sigma|_{D_1}$.

It is now easy to verify that $\overline{z}_{s,1}^\star(d_1) = 1$ for all $d_1 \in D_1$, that both $L \cup \{\mathit{red} \wedge \mathit{bird}, \mathit{fly}\}$ and $L \cup \{\mathit{red} \wedge \mathit{bird}, \neg\mathit{fly}\}$ are satisfiable, that $\overline{\kappa}_{s,1}^\star(\mathit{red} \wedge \mathit{bird} \wedge \mathit{fly}) = 0$, and that $\overline{\kappa}_{s,1}^\star(\mathit{red} \wedge \mathit{bird} \wedge \neg\mathit{fly}) = 1$. Thus, (KB, σ) $z_s^\star$-entails $\mathit{red} \wedge \mathit{bird} \to \mathit{fly}$ at strength 1.

5.6 Lexicographic Entailment

We now focus on lexicographic entailment. In the sequel, let $(KB, d) = ((L, D), \phi \to \psi)$ be a k-ff-Horn default reasoning problem. Let π be a priority assignment on KB. Let D_a denote the set of all active defaults w.r.t. KB and d, let At_a be the activation set of $(L, D_a \cup \{d\})$, and let $(At_1, \ldots, At_n)$ be the default partition of $(L, D_a \cup \{d\})$.

For every $i \in \{1, \ldots, n\}$, let D_i denote the set of all defaults in D_a that are defined over $At_a \cup At_i$, and let $KB_i = (L, D_i)$. Let π_i be the unique priority assignment on KB_i that is consistent with π on KB (that is, $\pi_i(d) < \pi_i(d)$ iff $\pi(d) < \pi(d)$, for all $d \in D_i$). Let $j \in \{1, \ldots, n\}$ such that d is defined over $At_a \cup At_j$.

In order to decide whether (KB, π) lex_p-entails d, we must check whether every π-preferred model of $L \cup \{\phi\}$ satisfies ψ. It can now be shown that we can equivalently check whether every π_j-preferred model of $L \cup \{\phi^\star\}$ satisfies ψ.

Using this result, it can be shown that deciding lex_p-entailment is tractable in the k-ff-Horn case. Moreover, as computing the z-partition for ε-consistent conditional knowledge bases KB is tractable in the Horn case [16], this result shows also that deciding lex-entailment is tractable in the k-ff-Horn case.

Theorem 5.13. *Let $k > 0$ be fixed. Given a k-ff-Horn default reasoning problem $(KB, d) = ((L, D), \phi \to \psi)$ and a priority assignment π on KB, deciding whether (KB, π) lex_p-entails d can be done in linear time.*

Example 5.14. Let the 2-ff-Horn default reasoning problem (KB, d) be given by $KB = (L, D)$ of Example 5.4 and $d = \mathit{red} \wedge \mathit{bird} \to \mathit{fly}$. Let $\pi(\delta) = 0$, if $\delta \in \{\mathit{bird} \to \mathit{fly}, \mathit{bird} \to \mathit{wings}, \mathit{fly} \to \mathit{mobile}\}$, and $\pi(\delta) = 1$, if $\delta \in \{\mathit{penguin} \to \neg\mathit{fly}, \mathit{penguin} \to \mathit{arctic}\}$.

It can be shown that (KB, π) lex_p-entails $\mathit{red} \wedge \mathit{bird} \to \mathit{fly}$ iff either $L \cup \{\mathit{red} \wedge \mathit{bird}\}$ is unsatisfiable, or all π_1-preferred models of $L \cup \{(\mathit{red} \wedge \mathit{bird})^\star\} = L \cup \{\mathit{red} \wedge \mathit{bird}\}$

satisfy *fly*, where π_1 is the priority assignment on $KB_1 = (L, D_1) = (L, \{bird \rightarrow fly, fly \rightarrow mobile\})$ that maps each element of D_1 to 0. It is now easy to verify that this is indeed the case. That is, (KB, π) lex$_\mathrm{p}$-entails $red \wedge bird \rightarrow fly$.

Acknowledgments. This work has been partially supported by the Austrian Science Fund Project N Z29-INF and a DFG grant.

References

1. E. W. Adams. *The Logic of Conditionals*, volume 86 of *Synthese Library*. D. Reidel, Dordrecht, Netherlands, 1975.
2. F. Bacchus, A. Grove, J. Y. Halpern, and D. Koller. From statistical knowledge bases to degrees of beliefs. *Artif. Intell.*, 87:75–143, 1996.
3. S. Benferhat, C. Cayrol, D. Dubois, J. Lang, and H. Prade. Inconsistency management and prioritized syntax-based entailment. In *Proceedings IJCAI-93*, pages 640–645, 1993.
4. S. Benferhat, D. Dubois, and H. Prade. Nonmonotonic reasoning, conditional objects and possibility theory. *Artif. Intell.*, 92(1–2):259–276, 1997.
5. E. Boros, Y. Crama, and P. L. Hammer. Polynomial-time inference of all valid implications for Horn and related formulae. *Ann. Math. Artif. Intell.*, 1:21–32, 1990.
6. E. Boros, P. L. Hammer, and X. Sun. Recognition of q-Horn formulae in linear time. *Discrete Appl. Math.*, 55:1–13, 1994.
7. R. A. Bourne and S. Parsons. Maximum entropy and variable strength defaults. In *Proceedings IJCAI-99*, pages 50–55, 1999.
8. C. Cayrol, M.-C. Lagasquie-Schiex, and T. Schiex. Nonmonotonic reasoning: From complexity to algorithms. *Ann. Math. Artif. Intell.*, 22(3–4):207–236, 1998.
9. T. Eiter and T. Lukasiewicz. Complexity results for default reasoning from conditional knowledge bases. In *Proceedings KR-2000*, pages 62–73, 2000.
10. T. Eiter and T. Lukasiewicz. Default reasoning from conditional knowledge bases: Complexity and tractable cases, 2000. Manuscript.
11. H. Geffner. *Default Reasoning: Causal and Conditional Theories*. MIT Press, 1992.
12. H. Geffner and J. Pearl. Conditional entailment: Bridging two approaches to default reasoning. *Artif. Intell.*, 53(2–3):209–244, 1992.
13. M. Goldszmidt, P. Morris, and J. Pearl. A maximum entropy approach to nonmonotonic reasoning. *IEEE Trans. Pattern Anal. Mach. Intell.*, 15(3):220–232, 1993.
14. M. Goldszmidt and J. Pearl. On the consistency of defeasible databases. *Artif. Intell.*, 52(2):121–149, 1991.
15. M. Goldszmidt and J. Pearl. System Z^+: A formalism for reasoning with variable strength defaults. In *Proceedings AAAI-91*, pages 399–404, 1991.
16. M. Goldszmidt and J. Pearl. Qualitative probabilities for default reasoning, belief revision, and causal modeling. *Artif. Intell.*, 84(1–2):57–112, 1996.
17. S. Kraus, D. Lehmann, and M. Magidor. Nonmonotonic reasoning, preferential models and cumulative logics. *Artif. Intell.*, 14(1):167–207, 1990.
18. P. Lamarre. A promenade from monotonicity to non-monotonicity following a theorem prover. In *Proceeding KR-92*, pages 572–580, 1992.
19. D. Lehmann. Another perspective on default reasoning. *Ann. Math. Artif. Intell.*, 15(1):61–82, 1995.
20. D. Lehmann and M. Magidor. What does a conditional knowledge base entail? *Artif. Intell.*, 55(1):1–60, 1992.
21. J. Pearl. Probabilistic semantics for nonmonotonic reasoning: A survey. In *Proceedings KR-89*, pages 505–516, 1989.
22. J. Pearl. System Z: A natural ordering of defaults with tractable applications to default reasoning. In *Proceedings TARK-90*, pages 121–135, 1990.