Model-based Defeasible Reasoning Literature Review

Carl Combrinck
University of Cape Town
Cape Town, South Africa
CMBCAR007@myuct.ac.za

ABSTRACT

Defeasible reasoning attempts to formalize aspects of human reasoning in which prior conclusions can be retracted with the addition of new information. Classical forms of reasoning are limited in expressing uncertainty and do not produce reasonable conclusions for exceptional knowledge. Although much work has been done in defining algorithms for computing several forms of defeasible consequence, these approaches are primarily based on the ranking of formulas in a knowledge base. Recent work has shown that particular models of a knowledge base characterize reasonable forms of defeasible entailment. However, there is work to be done in constructing, representing and using such models for defeasible entailment checking. In this review, we provide a brief history of, and motivation for, defeasible reasoning and its current trajectory, followed by an analysis of two principle patterns of reasoning (rational and lexicographic closure). Our focus, in defining these, will be on their model-based semantics, noting how this relates to existing formula-based approaches.

CCS CONCEPTS

• **Theory of computation → Automated reasoning**; • **Computing methodologies → Nonmonotonic, default reasoning and belief revision**.

KEYWORDS

artificial intelligence, knowledge representation and reasoning, defeasible reasoning, rational closure, lexicographic closure

1 INTRODUCTION AND MOTIVATION

At its core, *knowledge representation* can be understood as the expression of information about an environment through some language [5]. “It is sunny today" is a simple example of how knowledge is expressed via natural language. The manipulation of such knowledge for the purposes of deriving new knowledge defines the notion of *reasoning*. As humans, the ability to reason with existing knowledge comes naturally, however, in the context of artificial intelligence, we require more rigorous and structured methods to represent and manipulate knowledge. As such, knowledge is typically encoded symbolically and manipulated to produce new symbolic representations (conclusions) [10].

A simple, yet expressive means of achieving this is through *classical propositional logic*. Propositions (simple statements about the world) are assigned truth values and are combined using several operators to form more complex logical statements [1]. Semantics for several reasoning services allow for the conclusion of new propositional statements from a set of existing statements. Known as knowledge bases, these sets represent a collection of knowledge with which to reason.

Propositional logic has a number of desirable characteristics, some of which include its simplicity and ties to Boolean algebra [11]. However, it has limited expressively on account of its straightforward syntax and semantics. In particular, propositional logic cannot express the notion of *typicality*, where it is understood that a particular implication usually holds, but that there are possibly exceptions to this.

Another issue worth considering is that of *monotonicity*. The semantics of propositional logic define *entailment* (a fundamental reasoning service) which allows for reasonable conclusions to be drawn from existing knowledge. Entailment resolves what should follow as conclusions from a set of statements. Its definition, in propositional logic, is said to be monotonic, meaning that conclusions drawn from some knowledge will never be retracted with the addition of new knowledge [7]. This property, while desirable in certain contexts, imposes limitations on the capabilities of propositional logic. Consequently, defeasible approaches to reasoning have been, and continue to be, explored as nonmonotonic alternatives to the classical approaches.

Defeasible reasoning is a term coined and used in the field of philosophy and refers to nondeductive forms of reasoning in which conclusions are not final (as in propositional logic). These contingent assertions may, therefore, be retracted in certain circumstances. This is consistent with aspects of how humans reason. For example, suppose one sees an object which appears red. A natural conclusion is that the object is red. However, upon learning that the object is being illuminated by red lights, one may retract the prior conclusion that the object is red. Roughly speaking, defeasible reasoning is the same as nonmonotonic reasoning in artificial intelligence [12].

There are many approaches to defeasible reasoning as there is simply not one clear way in which to reason defeasibly. This has produced a number of competing formalisms, some of which have been shown to be equivalent in terms of expressivity and represent different perspectives of the same underlying pattern of reasoning[6].

Our focus in this review will be on forms of defeasible reasoning which are part of the framework proposed by Kraus, Lehmann and Magidor (henceforth referred to as KLM) [7] and its extensions. Using the preferential approach and extending, in parallel, the semantics defined by Shoham [15, 16], and the associated proof-theoretic system defined by Gabbay [3], KLM [7] defined a framework for defeasible reasoning referred to as the KLM framework. This framework is of particular interest due to its having both a model and proof theory, as well as computationally efficient algorithms for the associated reasoning services. [6].

The KLM framework extends the language of propositional logic by adding a preferential consequence relation at the meta-level and defining satisfaction and entailment using preferential interpretations [7]. This consequence relation is defined by several postulates which are formulated such that if a pair (α, β) is in the relation, this could be understood as meaning that "from α, I am willing to jump to conclude β unless I have information to the contrary" [6]. Lehmann and Magidor, in their paper titled "What does a conditional knowledge base entail?" [9], went on to further refine this framework, defining a more restricted class of "reasonable" consequence relations termed rational consequence relations. They defined these relations semantically using ranked interpretations (a subclass of preferential interpretations). There is also a shift to using an object-level connective to represent defeasible implication with a focus on entailment using propositional logic extended with this new connective [2]. They define a key form of entailment in *rational closure* and argue that a reasonable pattern of defeasible entailment should at least permit any assertion that rational closure endorses [9]. Rational closure can be viewed as a formalism of prototypical reasoning, a conservative form of reasoning that, intuitively, only makes assertions of typicality in "normal" cases (such assertions would not apply to abnormal cases) [8, 9, 14].

With this in mind, Lehmann [8] formalizes a default pattern of reasoning termed *lexicographic closure* based on Reiter's work on default logics [13]. Lexicographic closure, a subsumption of rational closure, is consistent with the description of presumptive reasoning. This represents a less conservative form of entailment which permits assertions, even in abnormal cases, unless there is evidence to the contrary [8].

Suppose, for example, we have that "birds fly", "birds have wings", "penguins are birds" and "penguins do not fly" [7]. The conservative approach, characterized by prototypical reasoning, would suggest that it cannot be concluded that "penguins have wings", since penguins are clearly atypical birds (and so we do not assume of penguins any typical characteristics of birds). On the other hand, a presumptive approach would have us conclude that indeed "penguins have wings" since we know penguins are birds, albeit atypical birds, and have no evidence that would suggest penguins do not have wings [8]. It is also important to notice that in the case of propositional logic, we would have that penguins cannot exist due to the contradiction that birds fly and penguins, which are birds, do not fly (recall that we have no means of expressing typicality as we have in defeasible approaches). In such cases, classical entailment would permit any penguin-based conclusions (since there would never exist a penguin to disprove any such assertions).

In [9], Lehmann and Magidor provides an algorithm for computing rational closure that involves ranking knowledge base statements in terms of their exceptionality (and hence specificity). Giordano et al. [4] provide a semantic characterisation of rational closure using minimal ranked entailment. It is shown that for a given knowledge base, there is a unique ranked interpretation (specifically the minimal model) that is able to completely define entailment with respect to this knowledge base as stipulated by rational closure.

Finally, rational defeasible entailment, a further refinement of the KLM framework [7] was proposed by Casini et al. [2]. This framework, in line with the conclusions of [9], has rational closure as its most conservative form of entailment and includes lexicographic closure. Importantly, it is shown that any form of rational defeasible entailment can be characterized by a ranked model that respects the corresponding ranking for rational closure. This provides a different perspective on computing entailment for a given knowledge base, compared to the provided algorithms in [2]. Such model-based approaches will form a significant part of our work, in defining algorithms to represent and construct ranked models corresponding to rational and lexicographic closure for the purposes of entailment checking.

2 PROPOSITIONAL LOGIC

2.1 Language

We define a set $\mathcal{P}$ containing all atomic propositions (represented using lower-case letters) [1]. These represent the most basic form of knowledge (they are indivisible). We may attribute some meaning to each propositional atom (for example, we may want p to represent whether a particular bird is a penguin) but propositional logic allows for the analysis of statements independent of their intuitive meaning [6].

In propositional logic, the meaning of these atoms is enriched via the addition of several logical operators in order to construct formulas (which define what can be expressed in the logic). These include $\{\neg, \wedge, \vee, \rightarrow, \leftrightarrow\}$ with $\neg$, a unary operator and the rest, binary operators.

Formulas can now be defined using the following grammar [1]:

$$\begin{aligned} fml &::= p \in \mathcal{P} \\ fml &::= \neg fml \\ fml &::= fml\ op\ fml \\ op &::= \ \wedge \mid \vee \mid \rightarrow \mid \leftrightarrow \end{aligned}$$

This says that formulas can be atoms, the negations ($\neg$) of other formulas, or the combination of two other formulas using one of the binary connectives. The set of all possible formulas is often referred to as $\mathcal{L}$ (the language of propositional logic).

2.2 Semantics

2.2.1 Interpretations. The formulas described have not yet been assigned truth values. A formula such as $\mathsf{p} \wedge \mathsf{q}$ is neither true nor false, until truth values are assigned to the atoms of which the formula is comprised. This is analogous to the idea that the value of $a + b$ can only be determined once values have been assigned to a and b [1]. In propositional logic, the assignment of truth values to atoms is fulfilled by interpretations.

An interpretation $\mathcal{I}$ is defined as a function $\mathcal{I} : \mathcal{P} \mapsto \{T, F\}$ which maps each propositional atom to a value of T or F (true and false respectively). For example, if we had $\mathcal{P} = \{\mathsf{p}, \mathsf{q}\}$, a valid interpretation may map p to true and q to false. That is: $\mathcal{I}(\mathsf{p}) = T$ and $\mathcal{I}(\mathsf{q}) = F$.

We are now able to define the semantics of the logical operators, which will allow us to assign truth values to formulas.

2.2.2 Logical Operators. Let A, B be any formulas containing only atoms from a set $\mathcal{P}$. We denote the value of each formula under a given interpretation $\mathcal{I}$ (of the atoms in $\mathcal{P}$) as $v_{\mathcal{I}}(A)$ and $v_{\mathcal{I}}(B)$ respectively. The way in which $v_{\mathcal{I}}$ assigns truth values can be defined inductively as follows [1]:

$$v_{\mathcal{I}}(A) = \mathcal{I}(A) \quad \text{if } A \text{ is an atom}$$

$$v_{\mathcal{I}}(\neg A) = \begin{cases} T & v_{\mathcal{I}}(A) = F \\ F & \text{otherwise} \end{cases}$$

$$v_{\mathcal{I}}(A \land B) = \begin{cases} T & v_{\mathcal{I}}(A) = T \text{ and } v_{\mathcal{I}}(B) = T \\ F & \text{otherwise} \end{cases}$$

$$v_{\mathcal{I}}(A \lor B) = \begin{cases} F & v_{\mathcal{I}}(A) = F \text{ and } v_{\mathcal{I}}(B) = F \\ T & \text{otherwise} \end{cases}$$

$$v_{\mathcal{I}}(A \to B) = \begin{cases} F & v_{\mathcal{I}}(A) = T \text{ and } v_{\mathcal{I}}(B) = F \\ T & \text{otherwise} \end{cases}$$

$$v_{\mathcal{I}}(A \leftrightarrow B) = \begin{cases} T & v_{\mathcal{I}}(A) = v_{\mathcal{I}}(B) \\ F & v_{\mathcal{I}}(A) \neq v_{\mathcal{I}}(B) \end{cases}$$

For example, say $\mathcal{P} = \{p, q\}$. Consider the formula $(p \lor q) \to \neg q$ and the interpretation $\mathcal{I}$ where $\mathcal{I}(p) = T$ and $\mathcal{I}(q) = F$. $v_{\mathcal{I}}(p \lor q) = T$ since $v_{\mathcal{I}}(p) = \mathcal{I}(p) = T$. $v_{\mathcal{I}}(\neg q) = T$ since $v_{\mathcal{I}}(q) = F$. Finally, $v_{\mathcal{I}}((p \lor q) \to \neg q) = T$ since $v_{\mathcal{I}}(p \lor q) = T$ and $v_{\mathcal{I}}(\neg q) = T$.

2.2.3 Satisfaction. Given any interpretation for any formula, we are now able to evaluate the formula with respect to the interpretation. In most cases, we are interested in interpretations which satisfy a particular formula or set of formulas, meaning that the formula, or each formula in the set, evaluates to true under the given interpretation. Formally, we define satisfaction using the symbol $\Vdash$ as follows: $\mathcal{I} \Vdash A$ (where $\mathcal{I}$ is an interpretation of a formula $A \in \mathcal{L}$) if and only if $v_{\mathcal{I}}(A) = T$. This is extended to sets of formulas (also known as knowledge bases): $\mathcal{I} \Vdash \mathcal{K}$ (where $\mathcal{I}$ is an interpretation of the formulas in the set $\mathcal{K}$) if and only if $v_{\mathcal{I}}(A) = T$ for every formula $A \in \mathcal{K}$. Interpretations which satisfy a formula or knowledge base are referred to as models of that formula or knowledge base. We use the notation $Mod(A)$ (sometimes $[\![A]\!]$) and $Mod(\mathcal{K})$ to refer to the set of models of a formula A and knowledge base $\mathcal{K}$, respectively. If this set is non-empty, we say the corresponding formula or knowledge base is *satisfiable* (i.e. that there is at least one interpretation which satisfies it).

Returning to the previous example, we notice that $\mathcal{I} \Vdash (p \lor q) \to \neg q$ and so $\mathcal{I} \in Mod((p \lor q) \to \neg q)$.

2.2.4 Entailment. Using the above model-based semantics, we define entailment (or logical consequence), denoted using the $\models$ symbol. A formula A entails a formula B, written as $A \models B$, if and only if $Mod(A) \subseteq Mod(B)$. Intuitively, whenever A is true under a given interpretation, such will be the case for B and so we are able to conclude B whenever we have A. This is extended to knowledge bases as expected: $\mathcal{K} \models A$ if and only if $Mod(\mathcal{K}) \subseteq Mod(A)$.

Consider a knowledge base $\mathcal{K} = \{p \lor q, \neg p\}$. We have that $Mod(\mathcal{K}) = \{\bar{p}q\}$ ($\bar{p}q$ is shorthand for an interpretation that maps p to false and q to true). Consequently, we have that $\mathcal{K} \models p \to q$ since $\bar{p}q \Vdash p \to q$ and so every model of $\mathcal{K}$ is also a model of $p \to q$.

2.2.5 Object and Meta Levels. Worth noting is the distinction between two key levels in the logic, the object level and the meta level. The object level refers to anything that is used to explicitly model knowledge (e.g. the propositional formulas), whereas the meta level refers to those things which operate above the object level [6].

An example of this distinction is highlighted in the differences between the $\to$ connective and the $\models$ operator. While both suggest that the truth of the left argument implies the truth of the right hand argument (representing an "if then" relationship), the $\to$ connective achieves this by relating formulas on the object level as knowledge (and hence its truth may vary between interpretations). Entailment, on the other hand, suggests that this relationship holds across all interpretations (and hence operates on the meta level). In this case, there is a simple connection between these two operators: $A \models B$ if and only if $A \to B$ is true in every interpretation.

3 DEFEASIBLE REASONING

3.1 Motivation

Consider, again, the penguin triangle example [7]: "birds fly", "birds have wings", "penguins are birds" and "penguins do not fly". We may represent this as a knowledge base $\mathcal{K}$ in propositional logic with $\mathcal{K} = \{b \to f, b \to w, p \to w, p \to \neg f\}$. The issue here is than there are no models of $\mathcal{K}$ in which p is true (penguins cannot exist) and so we cannot meaningfully reason about penguins. We would rather have some means of handling the fact that penguins do exist but are an exceptional type of bird that does not fly, and allow for reasonable forms on entailment in such cases.

3.2 The KLM Framework and Extensions

Initially, KLM [7] extended propositional logic to include a meta-level consequence relation $\mid\!\sim$ to represent defeasible implications (where $\alpha \mid\!\sim \beta$, with propositional formulas α, β, is read as "typically, if α, then β" [2]). They provide several postulates defining this relation in an attempt to reasonably represent the notion of *typicality.* The semantics of $\mid\!\sim$ are then defined using preferential interpretations, however, for the purposes of this review, we are more interested in a subclass of preferential interpretations termed *ranked interpretations* [9].

3.2.1 Ranked Interpretations. A ranked interpretation is a function $\mathcal{R} : \mathcal{U} \mapsto \mathcal{N} \cup \{\infty\}$, such that for every $i \in \mathcal{N}$, if there exists a $u \in \mathcal{U}$ such that $\mathcal{R}(u) = i$, then there must be a $v \in \mathcal{U}$ such that $\mathcal{R}(v) = j$ with $0 \leq j < i$, where $\mathcal{U}$ is the set of all possible propositional interpretations [6]. Ranked interpretations, therefore, assign to each propositional interpretation, a rank (with lower ranks corresponding, semantically, with more typical interpretations and higher ranks with less typical "worlds"). Worlds with a rank of ∞, according to the ranked interpretation, are impossible, whereas worlds with finite ranks are possible.

For $\mathcal{P} = \{p, q, r\}$, a possible ranked interpretation would be:

∞	$\bar{p}\bar{q}r$
2	$\bar{p}qr$ $p\bar{q}r$
1	pqr $pq\bar{r}$ $\bar{p}q\bar{r}$
0	$\bar{p}\bar{q}\bar{r}$ $p\bar{q}\bar{r}$

The fact that $\mathcal{R}(\bar{p}\bar{q}\bar{r}) = 0 < 1 = \mathcal{R}(pqr)$, intuitively represents that, under this ranking of interpretations, $\bar{p}\bar{q}\bar{r}$ is a more typical world than pqr. Since $\mathcal{R}(\bar{p}\bar{q}r) = \infty$, this world is deemed impossible by $\mathcal{R}$.

In line with the refinements made by Lehmann and Magidor [9], we add $\mathrel{|\!\sim}$ to the object level, allowing formulas of the form $\alpha \mathrel{|\!\sim} \beta$ with $\alpha, \beta \in \mathcal{L}$, in addition to the usual set of propositional formulas. This extended language is defined as $\mathcal{L}_P := \mathcal{L} \cup \{\alpha \mathrel{|\!\sim} \beta \mid \alpha, \beta \in \mathcal{L}\}$ [6]. Defeasible knowledge (that which may contain exceptions) can be modelled explicitly in the language with this extension, allowing for the formulation of nonmonotonic entailment at the meta level [6]. What remains is to define how formulas in this extended language are satisfied with respect to ranked interpretations, as well as the semantics of defeasible entailment.

3.2.2 *Satisfaction.* Given that ranked interpretations indicate the relative typicality of worlds, it makes sense to define whether a ranked interpretation satisfies a defeasible implication based on the most typical worlds in that interpretation. In order to define the "most typical worlds", a definition of minimal worlds with respect to a formula in $\mathcal{L}$ is required.

Given a ranked interpretation $\mathcal{R}$ and any formula $\alpha \in \mathcal{L}$, it holds that $u \in [\![\alpha]\!]^{\mathcal{R}}$ (the models of α in $\mathcal{R}$) is minimal if and only if there is no $v \in [\![\alpha]\!]^{\mathcal{R}}$ such that $\mathcal{R}(v) < \mathcal{R}(u)$ [6]. This defines the concept of the "best α worlds" (i.e. the lowest ranked, or most typical, of the worlds in which α is true).

Given a ranked interpretation $\mathcal{R}$ and a defeasible implication $\alpha \mathrel{|\!\sim} \beta$, $\mathcal{R}$ satisfies $\alpha \mathrel{|\!\sim} \beta$, written $\mathcal{R} \Vdash \alpha \mathrel{|\!\sim} \beta$ if and only if for every s minimal in $[\![\alpha]\!]^{\mathcal{R}}$, $s \Vdash \beta$. If $\mathcal{R} \Vdash \alpha \mathrel{|\!\sim} \beta$ then $\mathcal{R}$ is said to be a *model* of $\alpha \mathrel{|\!\sim} \beta$ [6].

This says that in order for a ranked interpretation $\mathcal{R}$ to satisfy a defeasible implication $\alpha \mathrel{|\!\sim} \beta$, it need only satisfy $\alpha \rightarrow \beta$ in the most typical (lowest ranked) α worlds of $\mathcal{R}$. The intuition here is that α typically implies β, if the best α worlds also satisfy β.

In the case of a propositional formula $\alpha \in \mathcal{L}$, it is required that every finitely-ranked interpretation in $\mathcal{R}$ satisfies α in order for $\mathcal{R}$ to satisfy α. This is consistent with idea that propositional formulas, which do not permit exceptionality, should be satisfied in every plausible world of a ranked interpretation, if such a ranking is to satisfy the formula. A useful result is that $\mathcal{R} \Vdash \alpha \in \mathcal{L}$ if and only if $\mathcal{R} \Vdash \neg\alpha \mathrel{|\!\sim} \bot$ ($\bot$ represents a contradiction such as $\mathsf{p} \wedge \neg\mathsf{p}$). This corresponds to the case in which all the best $\neg\alpha$ worlds have infinite rank, which must mean that there are no $\neg\alpha$ worlds with finite rank, and so all finitely-ranked worlds are α worlds. In this way, it is possible to express propositional statements using defeasible implications and so it is not necessary to discriminate between these two cases (we may deal, exclusively, with defeasible implications knowing that these subsume propositional formulas).

3.2.3 *Entailment.* It is now possible to model knowledge that expresses typicality, and thus handles exceptional cases more reasonably, however, there remains the concern of monotonicity with respect to entailment. We seek a reasonable form of entailment that permits the retraction of conclusions in cases where knowledge is added that contradicts these conclusions. There is no set way of defining such a form of defeasible entailment, however, it is argued that certain entailment relations are more reasonable than others.

KLM [7], as part of their framework for defeasible reasoning, first proposed 6 postulates defining what they argued to be a class of reasonable consequence relations (based on semantics defined by Shoham [15, 16]). Lehmann and Magidor [9] then define a class of entailment relations based on these postulates, with the addition of a 7th, to produce the postulates which define LM-rational entailment relations:

1. (LLE) Left logical equivalence: $\frac{\mathcal{K} \mathrel{|\!\approx} \alpha \leftrightarrow \beta,\ \mathcal{K} \mathrel{|\!\approx} \alpha \mathrel{|\!\sim} \gamma}{\mathcal{K} \mathrel{|\!\approx} \beta \mathrel{|\!\sim} \gamma}$
2. (RW) Right weakening: $\frac{\mathcal{K} \mathrel{|\!\approx} \alpha \rightarrow \beta,\ \mathcal{K} \mathrel{|\!\approx} \gamma \mathrel{|\!\sim} \alpha}{\mathcal{K} \mathrel{|\!\approx} \gamma \mathrel{|\!\sim} \beta}$
3. (Ref) Reflexivity: $\mathcal{K} \mathrel{|\!\approx} \alpha \mathrel{|\!\sim} \alpha$
4. And: $\frac{\mathcal{K} \mathrel{|\!\approx} \alpha \mathrel{|\!\sim} \beta,\ \mathcal{K} \mathrel{|\!\approx} \alpha \mathrel{|\!\sim} \gamma}{\mathcal{K} \mathrel{|\!\approx} \alpha \mathrel{|\!\sim} \beta \wedge \gamma}$
5. Or: $\frac{\mathcal{K} \mathrel{|\!\approx} \alpha \mathrel{|\!\sim} \gamma,\ \mathcal{K} \mathrel{|\!\approx} \beta \mathrel{|\!\sim} \gamma}{\mathcal{K} \mathrel{|\!\approx} \alpha \vee \beta \mathrel{|\!\sim} \gamma}$
6. (CM) Cautious Monotonicity: $\frac{\mathcal{K} \mathrel{|\!\approx} \alpha \mathrel{|\!\sim} \gamma,\ \mathcal{K} \mathrel{|\!\approx} \alpha \mathrel{|\!\sim} \beta}{\mathcal{K} \mathrel{|\!\approx} \alpha \wedge \beta \mathrel{|\!\sim} \gamma}$
7. (RM) Rational Monotonicity: $\frac{\mathcal{K} \mathrel{|\!\approx} \alpha \mathrel{|\!\sim} \gamma,\ \mathcal{K} \not\mathrel{|\!\approx} \alpha \not\mathrel{|\!\sim} \beta}{\mathcal{K} \mathrel{|\!\approx} \alpha \wedge \beta \mathrel{|\!\sim} \gamma}$

KLM defined preferential entailment using the semantics of preferential interpretations [7] and Lehmann and Magidor, similarly, defined ranked entailment with ranked interpretations [9]. The semantics in both cases are very similar to that of classical entailment. Consider the definition of ranked entailment:

Given a knowledge base $\mathcal{K}$, and a defeasible implication $\alpha \mathrel{|\!\sim} \beta$, $\mathcal{K} \mathrel{|\!\approx}_{\mathcal{R}} \alpha \mathrel{|\!\sim} \beta$, read as $\mathcal{K}$ rank entails $\alpha \mathrel{|\!\sim} \beta$, if and only if for every ranked interpretation, $\mathcal{R}$, such that $\mathcal{R} \Vdash \mathcal{K}$, it is the case that $\mathcal{R} \Vdash \alpha \mathrel{|\!\sim} \beta$ [9].

It was shown [9] that preferential and ranked entailment (while defined, originally, using different semantics) represent the same entailment relation and are both LM-rational (conform to the 7 postulates).

While this seems a reasonable way of defining entailment (requiring that all models satisfy a formula), it is still monotonic and hence does not address the issue at hand. We will now define minimal ranked entailment and show how it is a potential solution to our problem of monotonic entailment.

A partial order over all ranked models of a knowledge base $\mathcal{K}$, denoted $\preceq_{\mathcal{K}}$, is defined as follows [2]:

Given a knowledge base, $\mathcal{K}$, and $\mathcal{R}^{\mathcal{K}}$ the set of all ranked interpretations of $\mathcal{K}$, it holds for every $\mathcal{R}_1^{\mathcal{K}}, \mathcal{R}_2^{\mathcal{K}} \in \mathcal{R}^{\mathcal{K}}$ that $\mathcal{R}_1^{\mathcal{K}} \preceq_{\mathcal{K}} \mathcal{R}_2^{\mathcal{K}}$ if and only if for every $u \in \mathscr{U}$, $\mathcal{R}_1^{\mathcal{K}}(u) \leq \mathcal{R}_2^{\mathcal{K}}(u)$.

Intuitively, this partial order favours ranked interpretations that have their worlds "pushed down" as far as possible [6].

This partial order has a unique minimal element, $\mathcal{R}_{RC}^{\mathcal{K}}$, as shown by Giordano et al. [4]. We now define minimal ranked entailment using this minimal element as follows:

Given a defeasible knowledge base $\mathcal{K}$, the minimal ranked interpretation satisfying $\mathcal{K}$, $\mathcal{R}_{RC}^{\mathcal{K}}$, defines an entailment relation, $\mathrel{|\!\approx}$, called minimal ranked entailment, such that for any defeasible implication $\alpha \mathrel{|\!\sim} \beta$, $\mathcal{K} \mathrel{|\!\approx} \alpha \mathrel{|\!\sim} \beta$ if and only if $\mathcal{R}_{RC}^{\mathcal{K}} \Vdash \alpha \mathrel{|\!\sim} \beta$ [6]. Minimal ranked entailment is not only LM-rational, but also nonmonotonic.

In this form of entailment, a particular ranked interpretation is used to characterize the behaviour of the entailment relation. That is, the formulas that a specific ranked interpretation satisfies are precisely those entailed by the relation. It is worth noting that such an entailment relation defined by a ranked interpretation will exhibit the 7 KLM postulates (and is thus part of the KLM framework) [9]. More refined classes of entailment can be defined by referring to subsets of ranked interpretations which provide the semantics for these relations.

One such refinement of the KLM framework is rational defeasible entailment [2]. We will see that minimal ranked entailment is one formalism of a conservative reasoning pattern and, as argued by Casini et al. [2], should be the “nonmonotonic core" of any reasonable entailment relation. Rational deafeasible entailment relations are thus, much like LM-rational entailment relations, defined by extending the original KLM postulates as well as via a class of ranked interpretations. We will look at two specific patterns of entailment which fall within this extension of the KLM framework, namely *rational closure* and *lexicographic closure*. In particular, we focus on their model-based semantics as our work will involve computing entailment from this perspective.

3.3 Rational Closure

Rational closure was first described as a potential answer to the question of what a defeasible knowledge base should entail [9]. It represents a prototypical pattern (one that is extremely conservative in abnormal cases) of defeasible reasoning in the KLM framework and is defined in terms of what assertions may follow from a given knowledge base $\mathcal{K}$. Lehmann and Magidor [9] conclude that rational closure represents a reasonable solution to the problem of defeasible entailment and that it represents the most conservative form of reasonable entailment. As such, they propose that any other reasonable form of entailment, while possibly being more “adventurous" in its conclusions, should endorse at least those assertions in the rational closure of the corresponding knowledge base.

There are 2 principle ways in which to compute the rational closure of a given knowledge base. The first, which has been previously described, is minimal ranked entailment. This approach defines rational closure and the semantics of the associated entailment relation using a unique ranked interpretation for a given knowledge base. The second, which was described in [9] when Lehmann and Magidor first defined rational closure, is an algorithmic approach involving the ranking of statements in the knowledge base. Of these two methods, we will focus on the first due to its model-based approach.

3.3.1 Minimal Ranked Entailment. As discussed, minimal ranked entailment involves finding the minimal ranked model of a given knowledge base, where minimal refers to the property of having its worlds “pushed down" as far as possible [6]. Consider the following example which illustrates how such a model can be used to compute entailment.

Consider the following knowledge base: $\mathcal{K} := \{\mathtt{bird} \mid\!\sim \mathtt{fly}, \mathtt{bird} \mid\!\sim \mathtt{wings}, \neg(\mathtt{kiwi} \rightarrow \mathtt{bird}) \mid\!\sim \bot\}$.

Intuitively, $\mathcal{K}$ suggests that birds usually fly, birds usually have wings, and kiwis are birds (note that kiwi here refers to the national bird of New Zealand). At this point, the only information in $\mathcal{K}$ directly concerning kiwis, is that they are birds (and so, reasonably, they are assumed to be typical birds). Using the partial order of ranked interpretations defined previously, the minimal ranked model, $\mathcal{R}_{RC}^{\mathcal{K}}$, of $\mathcal{K}$ is:

∞	$\mathtt{\overline{b}fkw}$ $\mathtt{\overline{bf}kw}$ $\mathtt{\overline{b}fk\overline{w}}$ $\mathtt{\overline{bf}k\overline{w}}$
1	$\mathtt{b\overline{f}kw}$ $\mathtt{bfk\overline{w}}$ $\mathtt{b\overline{fk}w}$ $\mathtt{b\overline{f}k\overline{w}}$ $\mathtt{bf\overline{kw}}$ $\mathtt{b\overline{fkw}}$
0	$\enclose{circle}{\mathtt{bfkw}}$ $\mathtt{bf\overline{k}w}$ $\mathtt{\overline{b}f\overline{k}w}$ $\mathtt{\overline{bfk}w}$ $\mathtt{\overline{b}f\overline{kw}}$ $\mathtt{\overline{bfkw}}$

For brevity, the propositions have been shortened: `b` is `bird`, `f` is `fly`, `w` is `wings` and `k` is `kiwi`. Notice that moving any world in this ranked interpretation to a lower rank would result in an interpretation that is no longer a model (i.e. all worlds have been assigned the lowest possible rank while preserving the truth of each statement in the knowledge base).

We see that $\mathcal{R}_{RC}^{\mathcal{K}} \Vdash \mathtt{kiwi} \mid\!\sim \mathtt{wings}$ since the circled minimal `kiwi` world has that `wings` is true, i.e. it follows that kiwis typically have wings (since it is known that kiwis are birds, and birds typically have wings).

Suppose the statement $\neg(\mathtt{kiwi} \rightarrow \neg\mathtt{fly}) \mid\!\sim \bot$ (that kiwis do not fly) was added to $\mathcal{K}$. The minimal ranked model ($\mathcal{R}_{RC}^{\mathcal{K}'}$) of $\mathcal{K} \cup \{\neg(\mathtt{kiwi} \rightarrow \neg\mathtt{fly}) \mid\!\sim \bot\}$, is:

∞	$\mathtt{\overline{b}fkw}$ $\mathtt{\overline{bf}kw}$ $\mathtt{\overline{b}fk\overline{w}}$ $\mathtt{\overline{bf}k\overline{w}}$ $\mathtt{bfkw}$ $\mathtt{bfk\overline{w}}$
1	$\enclose{circle}{\mathtt{b\overline{f}kw}}$ $\mathtt{b\overline{fk}w}$ $\enclose{circle}{\mathtt{b\overline{f}k\overline{w}}}$ $\mathtt{bf\overline{kw}}$ $\mathtt{b\overline{fkw}}$
0	$\mathtt{bf\overline{k}w}$ $\mathtt{\overline{b}f\overline{k}w}$ $\mathtt{\overline{bfk}w}$ $\mathtt{\overline{b}f\overline{kw}}$ $\mathtt{\overline{bfkw}}$

Now, notice that $\mathcal{R}_{RC}^{\mathcal{K}'} \nVdash \mathtt{kiwi} \mid\!\sim \mathtt{wings}$, since the circled minimal `kiwi` worlds do not both have `wings` being true. This demonstrates that minimal ranked entailment, and hence rational closure, is indeed nonmonotonic, since a previous conclusion was retracted with the addition of new information. Importantly, it also demonstrates the conservative nature of prototypical reasoning, formalized in rational closure. In $\mathcal{K}'$, kiwis are atypical birds since they are birds that do not fly. Since kiwis, therefore, no longer conform to the prototype of birds, prototypical reasoning would not assume of kiwis any typical characteristics of typical birds. In this case, the conclusion $\mathtt{kiwi} \mid\!\sim \mathtt{wings}$ is retracted, as a result.

3.3.2 RationalClosure. The other approach will be described as an algorithm, `RationalClosure` [2], which assigns a rank to each statement in a knowledge base, representing the defeasibility of that statement [6]. A lower rank corresponds to a more general statement and so, in the case of queries involving exceptional statements, these are the statements that are more readily ignored.

In order to rank statements in a knowledge base $\mathcal{K}$, the materialization $\overrightarrow{\mathcal{K}}$ of the knowledge base is first defined such that $\overrightarrow{\mathcal{K}} = \{\alpha \rightarrow \beta \mid \alpha \mid\!\sim \beta \in \mathcal{K}\}$. Secondly, exceptionality of a statement $\alpha \mid\!\sim \beta$ with respect to a knowledge base $\mathcal{K}$, is defined according to whether $\overrightarrow{\mathcal{K}} \models \neg\alpha$. Such a statement is exceptional since the opposite of its antecedent can be derived from the knowledge base (analogous to the penguin example, in which the propositional knowledge base would resolve that no penguins can exist).

The `BaseRank` [2] algorithm starts with $E_0 = \overrightarrow{\mathcal{K}}$ and determines which statements are exceptional in E_0. The exceptional statements are carried forward and placed in a set E_1. This is then repeated until there is no change between E_{i-1} and E_i. A statement is assigned the rank i corresponding to the the last E_i in which it appears (this is termed the base rank of the formula with respect to the knowledge base and is notated $br_{\mathcal{K}}(\alpha)$ - note $br_{\mathcal{K}}(\alpha \mathrel{|\!\sim} \beta) \equiv br_{\mathcal{K}}(\alpha)$. The last set of exceptional statements is assigned a rank of ∞ (i.e. those statements which would appear in every subsequent E_i if the algorithm was allowed to continue). Any propositional formulas in the original knowledge base are also assigned rank ∞ since these are maximally specific.

Finally, for a given query $\alpha \mathrel{|\!\sim} \beta$, the `RationalClosure` [2] algorithm checks if $\overrightarrow{\mathcal{K}} \models \neg\alpha$ and removes the statements having the lowest finite rank until this is no longer holds (i.e. until it is possible for α to be true). The answer to the query is provided by checking whether this final set of statements entails $\alpha \rightarrow \beta$.

In our kiwi example, the following ranking of formulas would be computed:

0	$\texttt{bird} \rightarrow \texttt{fly}, \texttt{bird} \rightarrow \texttt{wings}$
∞	$\texttt{kiwi} \rightarrow \texttt{bird}, \texttt{kiwi} \rightarrow \neg\texttt{fly}$

In answering the query of $\texttt{kiwi} \mathrel{|\!\sim} \texttt{wings}$, the algorithm finds that $\{\texttt{bird} \rightarrow \texttt{fly}, \texttt{bird} \rightarrow \texttt{wings}, \texttt{kiwi} \rightarrow \texttt{bird}, \texttt{kiwi} \rightarrow \neg\texttt{fly}\} \models \neg\texttt{kiwi}$ and removes the statements with the lowest rank (i.e. rank 0). $\{\texttt{kiwi} \rightarrow \texttt{bird}, \texttt{kiwi} \rightarrow \neg\texttt{fly}\} \models \neg\texttt{kiwi}$ does not hold and so the algorithm checks whether $\{\texttt{kiwi} \rightarrow \texttt{bird}, \texttt{kiwi} \rightarrow \neg\texttt{fly}\} \models \texttt{kiwi} \rightarrow \texttt{wings}$. The answer here, as expected, is no and thus it is not that case that $\mathcal{K} \mathrel{|\!\approx} \texttt{kiwi} \mathrel{|\!\sim} \texttt{wings}$.

An essential observation, which links these two approaches, is that $br_{\mathcal{K}}(\alpha)$ =min$\{i \mid \exists u \in Mod(\alpha)$ with $\mathcal{R}^{\mathcal{K}}_{RC}(u) = i\}$ [4]. Therefore, it is possible to compute the base rank of any formula with respect to a knowledge base using the minimal model of that knowledge base. Consequently, minimal ranked entailment and the `RationalClosure` algorithm compute the same entailment relation.

3.4 Lexicographic Closure

Another approach to defeasible entailment within the KLM framework, lexicographic closure, is described by Lehmann [8]. This is a formalism of the presumptive pattern of reasoning introduced by Reiter [13] in the context of default logics. Presumptive reasoning, unlike prototypical reasoning, is more "adventurous" and willing to conclude statements so long as their is no evidence to the contrary (even in atypical cases). The semantics of lexicographic closure depends on a "seriousness" ordering that is defined based on two criteria: specificity and cardinality.

In defining lexicographic closure for a knowledge base, Lehmann defines a basis for a formula as a subset of the knowledge that has a material counterpart consistent with the formula and maximal with respect to the seriousness ordering [8]. We will, instead, focus on how the seriousness order can be used to define a model-based semantics for lexicographic closure (as was done for rational closure). A formula-based approach will also be outlined.

3.4.1 Ranked Model. Given a knowledge base $\mathcal{K}$, the seriousness ordering associated with lexicographic closure can be defined on the subsets of $\mathcal{K}$. The first of the seriousness criteria, specificity, is determined by the base rank of the formulas in the subset (a higher base rank is considered more specific, as discussed in rational closure). The second, cardinality, is, naturally, determined by the number of statements. Combining these two criteria lexicographically, we obtain a partial order $<_S$ in which $A <_S B$ if A has statements with lower specificity than B, and in the case that both have statements of the same specificity, if A has fewer such statements. Associated with A, B are tuples containing the number of statements of each base rank in the respective subset of $\mathcal{K}$, starting with the ∞ base rank and then working down through the finite ranks (e.g. $(1, 0, 2)$ and $(1, 1, 4)$). Comparing these tuples using the natural lexicographic ordering of natural numbers in tuples defines formally the seriousness ordering.

Constructing the ranked model corresponding to lexicographic closure can be done as follows [2, 8]:

For a knowledge base $\mathcal{K}$, and worlds $m, n \in \mathscr{U}$, the preference order $\prec_{LC}$ over $\mathscr{U}$ is defined as: $m \prec_{LC} n$ if and only if $V(m) <_S V(n)$ where $V(m) \subseteq \mathcal{K}$ is the set of defeasible implications violated by $m \in \mathcal{U}$. $\prec_{LC}$ is a modular partial order over $\mathscr{U}$, and so defines a ranked interpretation, denoted $\mathcal{R}^{\mathcal{K}}_{LC}$ [6]. Hence, worlds violating fewer specific formulas in $\mathcal{K}$ are preferred and considered more typical.

Casini et al. define this preference order similarly but in terms of formulas satisfied as opposed to violated. Thus, $m \prec^{\mathcal{K}}_{LC} n$ if and only if $\mathcal{R}^{\mathcal{K}}_{RC}(n) = \infty$, or $\mathcal{R}^{\mathcal{K}}_{RC}(m) < \mathcal{R}^{\mathcal{K}}_{RC}(n)$, or $\mathcal{R}^{\mathcal{K}}_{RC}(m) = \mathcal{R}^{\mathcal{K}}_{RC}(n)$ and m satisfies more formulas than n in $\mathcal{K}$ [2]. This definition also highlights the fact that lexicographic closure is a refinement of rational closure, in that it respects the ranking of rational closure but refines preference for statements with the same rank.

Returning to our kiwi example where $\mathcal{K} := \{\texttt{bird} \mathrel{|\!\sim} \texttt{fly}, \texttt{bird} \mathrel{|\!\sim} \texttt{wings}, \neg(\texttt{kiwi} \rightarrow \texttt{bird}) \mathrel{|\!\sim} \bot, \neg(\texttt{kiwi} \rightarrow \neg\texttt{fly}) \mathrel{|\!\sim} \bot\}$, we can construct the model $\mathcal{R}^{\mathcal{K}}_{LC}$ corresponding to lexicographic closure by "lifting up" worlds that satisfy fewer statements while preserving the original rational closure ordering.

∞	$\overline{\texttt{b}}\texttt{fkw}$ $\overline{\texttt{bf}}\texttt{kw}$ $\overline{\texttt{b}}\texttt{fk}\overline{\texttt{w}}$ $\overline{\texttt{bf}}\texttt{k}\overline{\texttt{w}}$ $\texttt{bfkw}$ $\texttt{bfk}\overline{\texttt{w}}$
2	$\texttt{b}\overline{\texttt{f}}\texttt{k}\overline{\texttt{w}}$ $\texttt{b}\overline{\texttt{fkw}}$
1	$\enclose{circle}{\texttt{b}\overline{\texttt{f}}\texttt{kw}}$ $\texttt{b}\overline{\texttt{fk}}\texttt{w}$ $\texttt{bf}\overline{\texttt{kw}}$
0	$\texttt{bf}\overline{\texttt{k}}\texttt{w}$ $\overline{\texttt{b}}\texttt{f}\overline{\texttt{k}}\texttt{w}$ $\overline{\texttt{bfk}}\texttt{w}$ $\overline{\texttt{b}}\texttt{f}\overline{\texttt{kw}}$ $\overline{\texttt{bfkw}}$

Checking if the formula, $\texttt{kiwi} \mathrel{|\!\sim} \texttt{wings}$, is satisfied by $\mathcal{R}^{\mathcal{K}}_{LC}$, and hence in the lexicographic closure of $\mathcal{K}$, we find that $\mathcal{R}^{\mathcal{K}}_{LC} \Vdash \texttt{kiwi} \mathrel{|\!\sim} \texttt{wings}$ and hence that $\mathcal{K} \mathrel{|\!\approx}_{LC} \texttt{kiwi} \mathrel{|\!\sim} \texttt{wings}$.

Notice that lexicographic closure endorses that kiwis have wings whereas rational closure would not. This speaks to the presumptive nature of lexicographic closure, as it is willing to assert that kiwis have wings despite the fact that kiwis are atypical birds (since there is nothing that suggests that kiwis do not have wings in the knowledge base).

This observation also demonstrates the point made earlier that there are multiple valid solutions to the problem of defeasible entailment. In our kiwi example, the expected behaviour (based on our understanding of kiwis) is that a defeasible entailment relation should not conclude that kiwis typically have wings, since they are one of the few flightless birds that do not possess wings. This form of reasoning is most consistent with rational closure (i.e. since kiwis are atypical, they should not inherit other properties of typical birds, such as having wings). However, if kiwis were substituted for penguins in the example (appropriate since penguins too are birds that cannot fly), we would expect to conclude that penguins typically have wings (since we understand that penguins are flightless birds that do have wings). In this case, lexicographic closure may be more appropriate since it is willing to conclude that penguins typically have wings.

3.4.2 DefeasibleEntailment. As is the case for rational closure, there is an algorithm for computing lexicographic closure based on ranking formulas. This algorithm, `DefeasibleEntailment` [2], can be used to compute any basic defeasible entailment relation. Defined by Casini et al., basic defeasible entailment represents a refinement of the KLM framework. The algorithm can be thought of as a generalisation of the `RationalClosure` algorithm that accepts, as a parameter, a ranking function (in the case of rational closure, this would be `BaseRank`).

For computing lexicographic closure, we can define a rank function (in the same way in which a formula's base rank was be defined using the ranked model for rational closure): The lexicographic rank function, $r_{\mathcal{K}}^{LC}$, with respect to a knowledge base $\mathcal{K}$, is defined as $r_{\mathcal{K}}^{LC}(\alpha) := \min\{\mathcal{R}_{LC}^{\mathcal{K}}(v) \mid v \in Mod(\alpha)\}$ [2]. This ranking function can be used for computing lexicographic closure with the `DefeasibleEntailment` algorithm.

Using this algorithm, we can describe an algorithm for computing lexicographic closure in terms of the ranking of formulas computed for rational closure. Beginning with the ranked formulas computed for rational closure of $\mathcal{K}$, the algorithm starts by checking whether $\overrightarrow{\mathcal{K}} \models \neg\alpha$. If this holds, instead of discarding all the statements of the lowest finite rank, the algorithm weakens this lowest rank by replacing it with one formula logically equivalent to the disjunction of all subsets of this lowest rank of size $x - 1$ where x is the number of formulas in the lowest finite rank. This is achieved by combining the formulas in each of these subsets conjunctively, and then combining all the resulting formulas disjunctively (e.g. $\{\mathsf{k} \rightarrow \mathsf{b}, \mathsf{k} \rightarrow \neg\mathsf{f}, \mathsf{k} \rightarrow \neg\mathsf{w}\}$ would be weakened to $\{(\mathsf{k} \rightarrow \mathsf{b} \land \mathsf{k} \rightarrow \neg\mathsf{f}) \lor (\mathsf{k} \rightarrow \mathsf{b} \land \mathsf{k} \rightarrow \neg\mathsf{w}) \lor (\mathsf{k} \rightarrow \neg\mathsf{f} \land \mathsf{k} \rightarrow \neg\mathsf{w})\}$). The motivation for this, is that instead of removing all formulas, the algorithm weakens the formulas by considering the removal of one formula at a time, and if it is the case that at least one such removal prevents the entailment of $\neg\alpha$, it can begin checking entailment of the query. If the entailment still holds, the rank is weakened further by considering the formula produced using subsets of size $x-2$ from the original formulas. This is repeated until either, the entailment doesn't hold and so the entailment of the query can be checked, or until the algorithm is forced to check subsets of size 0. In this latter case, it is clear that regardless of the number of formulas removed, the entailment still holds and so the last weakening of the rank is discarded completely (effectively equivalent to having removed the lowest rank entirely as in rational closure) and the process is repeated on the next lowest rank of formulas.

Although a valid solution to computing lexicographic closure, it should be noted that this approach has some efficiency concerns as the size of the weakened formulas grows exponentially with each weakening.

4 CONCLUSIONS

We have demonstrated that there exist many approaches to defeasible reasoning. In particular, our kiwi and penguin example illustrates how more conservative forms of reasoning may be preferable in certain contexts, but may, in other contexts, prevent useful conclusions from being endorsed.

It was shown that entailment with respect to a specific pattern of reasoning, can be approached from a number of perspectives. Specifically, both rational closure and lexicographic closure may be computed using model- and formula-based algorithms. Our future work will focus on the former in leveraging model-based characterizations of rational and lexicographic closure to formulate reasonable algorithms for entailment checking.

To this end, we also described the relationship between rational and lexicographic closure as one of refinement, based on Casini et al.'s [2] identifying rational closure as the nonmonotonic core of rational defeasible entailment. This allowed us to describe model construction for computing lexicographic closure based on the model for rational closure.

An important consideration in our work will be the relative efficiency of constructing and using models for entailment and the existing algorithms for entailment. We found that defeasible entailment checking, in the case of rational closure, could be reduced to classical entailment checking and thus may use existing, highly efficient implementations of propositional reasoning services. What remains is to investigate whether model-based implementations will be able to compete with these existing solutions.

REFERENCES

[1] Mordechai Ben-Ari. 2012. *Propositional Logic: Formulas, Models, Tableaux.* Springer London, London, 1, 7–47.

[2] Giovanni Casini, Thomas Meyer, and Ivan Varzinczak. 2019. Taking Defeasible Entailment Beyond Rational Closure. In *Logics in Artificial Intelligence*. Springer International Publishing, Cham, 182–197.

[3] Dov M. Gabbay. 1985. Theoretical Foundations for Non-Monotonic Reasoning in Expert Systems. In *Logics and Models of Concurrent Systems*, Krzysztof R. Apt (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 439–457.

[4] L. Giordano, V. Gliozzi, N. Olivetti, and G.L. Pozzato. 2015. Semantic characterization of rational closure: From propositional logic to description logics. *Artificial Intelligence* 226 (2015), 1–33. https://doi.org/10.1016/j.artint.2015.05.001

[5] Crina Grosan and Ajith Abraham. 2011. *Knowledge Representation and Reasoning.* Springer Berlin Heidelberg, Berlin, Heidelberg, 131–147. https://doi.org/10.1007/978-3-642-21004-4_6

[6] Adam Kaliski. 2020. *An Overview of KLM-Style Defeasible Entailment.* Master's thesis. Faculty of Science, University of Cape Town, Rondebosch, Cape Town, 7700.

[7] Sarit Kraus, Daniel Lehmann, and Menachem Magidor. 1990. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence* 44, 1 (1990), 167–207. https://doi.org/10.1016/0004-3702(90)90101-5

[8] Daniel Lehmann. 1999. Another perspective on Default Reasoning. *Annals of Mathematics and Artificial Intelligence* 15 (11 1999). https://doi.org/10.1007/BF01535841

[9] Daniel Lehmann and Menachem Magidor. 1992. What does a conditional knowledge base entail? *Artificial Intelligence* 55, 1 (1992), 1–60. https://doi.org/10.1016/0004-3702(92)90041-U

[10] Hector J. Levesque. 1986. Knowledge Representation and Reasoning. *Annual Review of Computer Science* 1, 1 (1986), 255–287. https://doi.org/10.1146/annurev.cs.01.060186.001351 arXiv:https://doi.org/10.1146/annurev.cs.01.060186.001351

[11] James Donald Monk. 2012. *Mathematical logic.* Vol. 37. Springer Science & Business Media.

[12] John L. Pollock. 1987. Defeasible reasoning. *Cognitive Science* 11, 4 (1987), 481–518. https://doi.org/10.1016/S0364-0213(87)80017-4

[13] Raymond Reiter. 1980. A logic for default reasoning. *Artificial Intelligence* 13, 1 (1980), 81–132. https://doi.org/10.1016/0004-3702(80)90014-4 Special Issue on Non-Monotonic Logic.

[14] Raymond Reiter and Giovanni Criscuolo. 1983. Some Representational issues in default reasoning. In *Computational Linguistics*, Nick Cercone (Ed.). Pergamon, 15–27. https://doi.org/10.1016/B978-0-08-030253-9.50007-X

[15] Yoav Shoham. 1987. Nonmonotonic Logics: Meaning and Utility. In *Proceedings of the 10th International Joint Conference on Artificial Intelligence - Volume 1* (Milan, Italy) *(IJCAI'87)*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 388–393.

[16] Yoav Shoham. 1987. *A Semantical Approach to Nonmonotonic Logics.* Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 227–250.