EXTENDING DEFEASIBLE REASONING BEYOND RATIONAL CLOSURE

Maqhobosheane Mohlerepe
University of Cape Town
Cape Town, South Africa
mhlmaq001@myuct.ac.za

ABSTRACT

Defeasible reasoning offers a means to navigate the complexities of real-world decision-making by accommodating exceptions and uncertainties. This paper explores the theoretical underpinnings and current practical implementations of defeasible reasoning within the context of artificial intelligence (AI). Theoretical discussions delve into the limitations of classical propositional logic and the emergence of defeasible reasoning as a solution to address these shortcomings. Furthermore, the implementation aspect reviews existing defeasible reasoning tools, focusing on algorithms, optimizations, and their implications for enhancing performance and efficacy. The project outlined herein aims to develop an intuitive defeasible reasoning tool, leveraging insights from both theoretical and implementation-focused research.

CCS CONCEPTS

Theory of computation → Automated reasoning; • Computing methodologies → Nonmonotonic, Rational closure; Propositional logic.

KEYWORDS

Artificial Intelligence, Knowledge Representation and Reasoning, Defeasible Reasoning, Propositional Logic, Rational Closure, KLM.

1 INTRODUCTION

The question of how human knowledge can be represented by computer language is a pertinent one in the field of artificial intelligence. Classical reasoning, characterized by deductive and inductive processes, provides a structured framework for drawing conclusions from premises. However, the strict adherence to deductive validity fails to account for exceptions and uncertainties encountered in real-world scenarios. Defeasible reasoning offers a solution by allowing for rational but not necessarily deductively valid inferences. This departure from classical logic is essential for modeling human-like cognition in AI systems, where decision-making often involves navigating complex, uncertain environments. In this project, we aim to develop a defeasible reasoning tool that incorporates insights from theoretical research on nonmonotonic reasoning and practical implementations of defeasible reasoners.

2 PROPOSITIONAL LOGIC

Propositional logic is a language that formalises knowledge representation and reasoning. It is a language that deals with indivisible statements that can either be true or false known as *atomic propositions*[1].The language makes use of these atomic propositions and logical operators(connectives) to build up propositional formulas[2]. This section formally outlines the propositional logic language.

2. 1 SYNTAX

A *proposition* is a declarative statement about the world which can either be true or false. An *atomic proposition* is one which does not depend on the truth or falsify of any other proposition.
Variables are usually assigned to atomic proposition such that a *vocabulary* $\mathcal{P}$ consists of a set of atomic propositions $\{p, q, r, \ \ldots\}$. For example, we may represent the atomic proposition, "it is raining in Cape Town", as r.
These atomic propositions can be joined using logical connectives to create $propositional\ formulas$. Logical connectives represent the relation between propositions and can either be unary (applied to a single proposition) or binary (applied to two propositions). There are several connectives, however, we will focus on the classical propositional connectives shown in the table below [3]:

Operator Name	**Operator Symbol**	**Operator Type**
Negation	¬	Unary
Conjunction	∧	Binary
Disjunction	∨	Binary
Implication	→	Binary
Equivalence	↔	Binary

2.1 Propositional formulas Propositional formulas are statements built from propositional atoms. They are often denoted using Greek letters (α, β, γ ...).

A set of propositional formulas $\mathcal{L}$ is inductively defined by the following rules:

i. Every atom $p \in \mathcal{P}$ is in $\mathcal{L}$.
ii. The negation of any atom $p \in \mathcal{P}$, i.e. $\neg p$, is in $\mathcal{L}$
iii. If α and β are in $\mathcal{L}$, then $\neg\alpha$, $\alpha \wedge \beta$, $\alpha \vee \beta$, $\alpha \rightarrow \beta$ and $\alpha \leftrightarrow \beta$ are all in $\mathcal{L}$.

Constants $\top$ and $\bot$, which denote a formula that is always true and one that is always false respectively, are also in $\mathcal{L}$[13]. A finite set of propositional formulas is known as a knowledge base $\boldsymbol{\mathcal{K}}$.

2.2 SEMANTICS

Having established the grammar of the language, we will now formally define how meaning is drawn from symbols used in propositional logic.
Each formula α in $\mathcal{L}$ has a truth value (a value of True or False). This truth value of any formula α is dependent on its constituent parts, that is, it is dependent on the meaning of logical operators in the formula and the truth value of the atoms it consists of.
Valuations are all the logically possible states of the world based on a set of propositional atoms $\mathcal{P}$ in some language $\mathcal{L}$[5].
Formally, a *valuation* is a function $u : \mathcal{P} \mapsto \{T, F\}$ where T and F are True and False respectively. Consider a set of propositional atoms $\mathcal{P} = \{r, s\}$, the set of all possible valuations $\mathcal{U} = \{rs, r\bar{s}, \bar{r}s, \bar{r}\bar{s}\}$ assigns a false value to each atom with a bar over it and true otherwise. If an atom is true in a valuation, then the valuation is said to satisfy the atom. Therefore, the valuation $u = \bar{r}s$ satisfies r, formally denoted as $u \Vdash r$. This applies to any formula $\alpha\ and\ \beta$ in the language $\mathcal{L}$ such that u can be said to satisfy any formula in $\mathcal{L}$ if one the following criteria holds:
Note that the truth value of a formula α under a valuation u is written as $u(\alpha)$.

i. α is an atom, e.g. r, and $u(r)$ = T
ii. α is the negation of another formula, e.g. $\neg\beta$, and $u(\beta)$ = F
iii. α is the conjunction of two other formulas, e.g. $\beta \wedge \gamma$, and $u(\gamma)$ = T
iv. α is the disjunction of two other formulas, e.g. $\beta \vee \gamma$, and at least one of $u(\beta) = T$, or $u(\gamma)$= T, holds.
v. $\alpha = \beta \rightarrow \gamma$, and at least one of $u(\neg\beta) = T$ or $u(\gamma)$ = T holds.
vi. $\alpha = \beta \leftrightarrow \gamma$, and $u(\beta) = u(\gamma)$

Any valuation u that satisfies α is a model of α, denoted as $Mod(\alpha)$. The set of valuations that satisfy all the formulas in a knowledge base $\boldsymbol{\mathcal{K}}$ are the models of $\boldsymbol{\mathcal{K}}$, $Mod(\boldsymbol{\mathcal{K}})$. Entailment($\Vdash$) is a fundamental concept in reasoning, it formally expresses the notion of logical consequence. It allows us to make conclusions about statements that follow from other statements. $\boldsymbol{\mathcal{K}}$ entails α, which means α follows from a knowledge base $\boldsymbol{\mathcal{K}}$ and is denoted as $\boldsymbol{\mathcal{K}} \Vdash \alpha$, if and only if $\boldsymbol{Mod}(\boldsymbol{\mathcal{K}}) \subseteq Mod(\alpha)$.

3 DEFEASIBLE REASONING

Defeasible reasoning, unlike classical logic such as propositional logic, offers framework capable of accommodating the complexities inherent in real-world scenarios[3]. Unlike classical reasoning systems, which adhere strictly to the principles of monotonicity, defeasible reasoning introduces flexibility by acknowledging the presence of typicality and exceptionality in human cognition and the world at large.

In classical logic, the addition of new information to a knowledge base invariably leads to an expansion of the set of conclusions entailed by that knowledge base, without the possibility of retracting previously drawn inferences [11]. This rigidity poses significant limitations when confronted with scenarios where typical cases do not universally apply, or exceptions challenge established norms.

Defeasible reasoning, on the other hand, allows for the representation of typical cases as well as exceptions within its logical framework. There are many different attempts at formalising defeasible reasoning[4] which have been outlined in [5] One prominent approach to defeasible reasoning is encapsulated within the KLM (Kraus, Lehmann, and Magidor) framework, which uses conditional logic and a variant of preferential logic [6].

Central to the concept of defeasible reasoning is the notion that knowledge is inherently incomplete and subject to revision in light of new information. This recognition aligns with human cognition, where individuals routinely adjust their beliefs and conclusions based on evolving understanding and contradictory evidence.

In this section, we comprehensively explore the KLM framework as algorithms based on this framework have been proven to be computable and efficient[7].

3.1 The KLM Framework

Defeasible reasoning within the KLM-style framework introduces a syntax and semantics that enriches classical propositional logic, allowing for the representation of typicality and exceptionality in logical inference. At the syntactic level, the framework incorporates defeasible implications (DIs), denoted as $\alpha \mid\sim \beta$, *where* α *typically implies* β[8]. The nesting of this connective is not allowed.

Preferential interpretations, as articulated by KLM [9] and Lehmann and Magidor [10], are rooted in the preferential semantics initially formulated by Shoham[5], [9] . This semantic framework introduces a structured ordering of valuations, where a preferred valuation is deemed more typical than others. Central to this concept is the notion of "states," each associated with a classical valuation. It is important to distinguish states from valuations, as multiple states may correspond to the same valuation. Consequently, a single interpretation may encompass an infinite number of states. Within this framework, a specific subset of preferential interpretations known as *ranked interpretations* is of particular interest.

Semantically, defeasible reasoning is grounded in structures known as *ranked interpretations*. These interpretations establish a hierarchical ordering of interpretations based on their degree of 'typicality'. Each interpretation is assigned a rank within this hierarchy, with ranks ranging from 0 (most typical) to n (least typical), and a reserved rank of ∞ denoting impossibility. A ranked interpretation is formally defined as a function, denoted as $R : \mathcal{U} \rightarrow \mathbb{N} \cup \{\infty\}$, where $\mathcal{U}$ represents the set of interpretations. For an element u in the set $\mathcal{U}$, if $R(u) = 0$, it indicates that u is assigned the highest degree of typicality. Moreover, this interpretation must satisfy the convexity property, which states that for every natural number i, if $R(v) = i$, then for any natural number j such that $0 \leq j < i$, there exists an element u in the set $\mathcal{U}$ such that $R(u) = j$[8]. The convexity property ensures a structured ordering of interpretations, where each rank builds upon the previous one without gaps, reflecting the relative strength of logical entailments. [7]$\mathcal{R}^{\mathcal{K}}$ is sometimes used to denote a ranking interpretation $\mathcal{R}$ with respect to knowledge base $\mathcal{K}$[8].

Within this semantic framework, a ranked interpretation, $\mathcal{R}$ is said to satisfy a defeasible implication $\alpha \mid\sim \beta$ if, in the lowest rank where α is satisfied, β holds in all corresponding models within that rank. This is formally denoted as d R ⊩ $\alpha \mid\sim \beta$[8]. In this scenario, R is considered a model of $\alpha \mid\sim \beta$ [6]. Essentially, R entails $\alpha \mid\sim \beta$ if and only if the implication $\alpha \rightarrow \beta$ holds true in all the 'best' or most typical interpretations of α [11].

Furthermore, interpretations satisfying standard propositional formulas are defined as those interpretations where the formula holds true across all possible interpretations of a given rank, reflecting the absence of typicality considerations in classical propositional logic.

3.2 KLM: Defeasible Entailment

In defeasible reasoning, unlike propositional logic, there exists no predetermined approach for determining defeasible entailment (denoted as $\mid\approx$). The KLM framework proposed a set of properties that any defeasible entailment relation must satisfy[10], leading to the characterization of such forms of defeasible entailment as LM-Rational[12]. These properties include reflexivity, left logical equivalence, right weakening, conjunction, disjunction, cautious monotonicity, and rational monotonicity and are formally denoted as follows[5]:

For any knowledge base, $\mathcal{K}$ and propositional formulas α, β, γ

$$(Reflexivity)\ \mathcal{K} \mid\approx \alpha \mid\sim \alpha$$

$$(Left\ logical\ equivalence)\ \alpha \equiv \frac{\beta, \mathcal{K} \mid\approx \alpha \mid\sim \gamma}{\mathcal{K} \mid\approx \beta \mid\sim \gamma}$$

$$(Right\ weakening)\ \frac{\mathcal{K} \mid\approx \alpha \mid\sim \beta, \beta \mid= \gamma}{\mathcal{K} \mid\approx \alpha \mid\sim \gamma}$$

$$(And)\ \frac{\mathcal{K} \mid\approx \alpha \mid\sim \beta, \mathcal{K} \mid\approx \alpha \mid\sim \gamma}{\mathcal{K} \mid\approx \alpha \mid\sim \beta \wedge \gamma}$$

$$(Or)\ \frac{\mathcal{K} \mid\approx \alpha \mid\sim \gamma, \mathcal{K} \mid\approx \beta \mid\sim \gamma}{\mathcal{K} \mid\approx \alpha \vee \beta \mid\sim \gamma}$$

$$(Cautious\ monotonicity)\ \frac{\mathcal{K} \mid\approx \alpha \mid\sim \beta, \mathcal{K} \mid\approx \alpha \mid\sim \gamma}{\mathcal{K} \mid\approx \alpha \wedge \beta \mid\sim \gamma}$$

$$(Rational\ monotonicity)\ \frac{\mathcal{K} \mid\approx \alpha \mid\sim \gamma, \mathcal{K} \not\mid\approx \alpha \mid\sim \neg\beta}{\mathcal{K} \mid\approx \alpha \wedge \beta \mid\sim}$$

Among the methods proposed for achieving LM-rational defeasible entailment are rational closure and lexicographic closure, both of which align with the principles of LM-rationality[8].

3.2 Rational Closure

Rational closure provides a comprehensive approach to defeasible entailment by capturing the notion that general rules hold unless specific exceptions or contradictions exist. It is an extension of preferential reasoning [13]. Formally, rational closure entails determining whether a defeasible implication is entailed by a knowledge base using a method adhering to rationality properties proposed by Lehmann and Magidor. It ensures that conclusions derived from defeasible implications remain consistent and coherent, even in the presence of exceptions. Rational closure can be defined semantically or algorithmically as outlined in [8].

3.2.1 Minimal Ranked Entailment The semantic definition is based on *ranked interpretations* as detailed in section 3.1. There exists a singular minimal entity $\mathcal{R}_{\mathcal{K}}$ within the collection of all ranked interpretations corresponding to a specific knowledge base $\mathcal{K}$. Should $\mathcal{R}_{\mathcal{K}}$ entail $\alpha \mid\sim \beta$, it is denoted as K $\mid\approx$ $\alpha \mid\sim \beta$, indicating that $\mathcal{K}$ reasonably entails the defeasible implication $\alpha \mid\sim \beta$, or equivalently, that $\alpha \mid\sim \beta$ is encompassed within the rational closure of $\mathcal{K}$. Recognizing that this form of defeasible entailment stems from a ranked interpretation renders it LM-rational, a determination established by Lehmann and Magidor in [10].

3.2.2 Base Rank Algorithm The base rank is the initial step to algorithmically determining whether a defeasible implication is encompassed within the rational closure of a knowledge base. As detailed in [8], the base rank is the first of the two sub-algorithms used for rational closure. Firstly, we will outline some preliminary definitions necessary to understand the algorithms. For any

knowledge base $\mathcal{K}$, $\vec{\mathcal{K}}$ is used to denote $\mathcal{K}$'s *material counterpart* which is identical to $\mathcal{K}$ but with defeasible implications replaced with material implications, ie $|\sim$ replaced with $\rightarrow$ [14].
This is formally denoted as $\vec{\mathcal{K}} = \{\alpha \rightarrow \beta : \alpha \mid \sim \beta \in K\}$.

Initially, we establish a series of materializations, denoted as $E_0, E_1, \ldots, E_{n-1}, E_\infty$, where E_0 starts with the materialization of the knowledge base $\mathcal{K}$ itself. Subsequent materializations, Ei, are derived by filtering out statements of the form $\alpha \rightarrow \beta$ from the previous materialization E_{i-1}, where α can be demonstrated as false based on the existing statements. This iterative process continues until no new statements can be derived, signaling the completion of the sequence, with E_∞ representing the final materialization.

Following this, we construct a ranking based on the difference between consecutive materializations, where each rank delineates the statements that become more specific as we ascend in the ranks. The statements from the original knowledge base $\mathcal{K}$ occupy the bottom rank (referred to as the infinite rank), while statements become progressively more general in higher ranks. A statement $\alpha \rightarrow \beta$ is assigned a base rank i if it resides in rank i, signifying that in any ranked interpretation derived from the corresponding materialization E_i, $\alpha \rightarrow \beta$ will be valid in at least one interpretation belonging to the most typical rank.

The second sub-algorithm is the process of checking defeasible entailment and is detailed below. Given a knowledge base $\mathcal{K}$ and a defeasible implication statement $\alpha \mid\sim \beta$:

1. Initially, we verify if $\neg\alpha$ is entailed by the statements in the infinite rank.
2. Subsequently, we systematically eliminate sets of classical implications, starting from the highest rank and progressing downwards, until identifying a rank where $\neg\alpha$ is not entailed by the remaining statements. This establishes the satisfiability of α with respect to the current and remaining ranks.
3. We then assess whether the remaining statements entail $\alpha \rightarrow \beta$.
4. If, after this process, only the bottom rank remains, and the statements therein entail $\neg\alpha$, we conclude that K defeasibly entails $\alpha \mid\sim \beta$.

This algorithm ensures that a statement is deemed defeasibly entailed by a knowledge base only if it holds true in accordance with the minimal ranked interpretation of the knowledge base[15].

4 CURRENT IMPLEMENTATIONS

4.1 Insights from Literature In an investigation into the Scalability of Rational Closure (SCADR), the author emphasizes the significance of query set characteristics in determining the efficacy of optimization techniques for rational closure. Specifically, it highlights the impact of repeated antecedents and the average rank number at which queries become consistent with the knowledge base on the performance of optimization algorithms [16].

Moreover, in an extension of SCADR, the authors delve into computational complexity and algorithmic efficiency, stressing the need for strategic evaluation methods beyond simple random knowledge base generation. They emphasize that the performance of optimization methodologies depends on a complex interplay of factors such as the number of atoms, complexity of formulas, and redundancy in knowledge [17].

Lastly, work that has been done by [13] introduces Bayesian refinement and its implications for reduced minimal ranked entailment. This work underscores the importance of considering various perspectives to yield novel methodologies and insights, particularly in identifying redundancy and refining base rank approaches.

4.2 Discussions Building upon the insights from the literature, the following research directions emerge:

Optimization Strategies: Investigate dynamic algorithmic strategies for rational closure, considering the findings from query set characteristics and optimization techniques [16].

Integration and Refinement: Explore the integration of Bayesian refinement techniques with existing optimization methodologies, such as reduced minimal ranked entailment, to enhance the efficiency of rational closure systems [13].

Strategic Evaluation: Develop strategic evaluation methods for computational complexity, focusing on real-world scenarios and problem sets that capture the diverse factors influencing algorithmic efficiency [17].

Tool Development: Utilizing alternative tools and libraries to improve the performance of rational closure systems, beyond the ones explored in the existing literature.

5 CONCLUSIONS

In summary, this literature review explored how defeasible reasoning can address the limitations of classical logic in artificial intelligence. It started by discussing classical reasoning's approach and its drawbacks. Then introduced defeasible reasoning as solution to the limitations of classical reasoning. Theoretical discussions highlighted the shortcomings of classical logic, while practical implementations focused how we could improve algorithms for real-world decision-making. Insights from existing research underscored the importance of optimization strategies and algorithmic efficiency. Overall, this review lays the groundwork for the aim of the project which is to implement either non-problem specific or problem specific optimisations of existing implementations (or both).

REFERENCES

[1] 'Propositional Logic | Internet Encyclopedia of Philosophy'. Accessed: Mar. 26, 2024. [Online]. Available: https://iep.utm.edu/propositional-logic-sentential-logic/

[2] H. K. Büning and T. Lettmann, *Propositional Logic: Deduction and Algorithms*. Cambridge University Press, 1999.

[3] G. Antoniou and M.-A. Williams, *Nonmonotonic Reasoning*. MIT Press, 1997.

[4] 'Nonmonotonic Reasoning', in *Exploring Artificial Intelligence*, Morgan Kaufmann, 1988, pp. 439–481. doi: 10.1016/B978-0-934613-67-5.50016-2.

[5] A. Kaliski, 'An Overview of KLM-Style Defeasible Entailment. 2020. An overview of KLM-style defeasible entailment. Master's thesis. University of Cape Town, Cape Town, South Africa', University of Cape Town, Cape Town, 2020.

[6] Y. Shoham, 'A semantical approach to nonmonotonic logics', presented at the Logic in Computer Science, Oct. 1987. Accessed: Mar. 26, 2024. [Online]. Available: https://www.semanticscholar.org/paper/A-semantical-approach-to-nonmonotonic-logics-Shoham/931d6c614d5f74edfa6b2f44436449981dbdeea2

[7] D. Lehmann, 'Another perspective on default reasoning', *Ann. Math. Artif. Intell.*, vol. 15, no. 1, pp. 61–82, Mar. 1995, doi: 10.1007/BF01535841.

[8] G. Casini, T. Meyer, and I. Varzinczak, 'Taking Defeasible Entailment Beyond Rational Closure', in *Logics in Artificial Intelligence*, F. Calimeri, N. Leone, and M. Manna, Eds., Cham: Springer International Publishing, 2019, pp. 182–197. doi: 10.1007/978-3-030-19570-0_12.

[9] S. Kraus, D. Lehmann, and M. Magidor, 'Nonmonotonic reasoning, preferential models and cumulative logics', *Artif. Intell.*, vol. 44, no. 1, pp. 167–207, Jul. 1990, doi: 10.1016/0004-3702(90)90101-5.

[10] D. Lehmann and M. Magidor, 'What does a conditional knowledge base entail?', in *Artificial Intelligence*, May 1992, pp. 1–60. doi: 10.1016/0004-3702(92)90041-U.

[11] Cohen Jaron, 'Model-Based Defeasible Reasoning', Thesis, University of Cape Town, Cape Town, 2022.

[12] G. Casini, T. Meyer, and I. Varzinczak, 'Defeasible Entailment: from Rational Closure to Lexicographic Closure and Beyond', Oct. 2018. Accessed: Mar. 26, 2024. [Online]. Available: https://www.semanticscholar.org/paper/Defeasible-Entailment%3A-from-Rational-Closure-to-and-Casini-Meyer/1a39abedbc1db1ae86036eafb3925300f2c938c7

[13] Slater Luke, 'Extending Defeasible Reasoning Beyond Rational Closure', Thesis, University of Cape Town, 2022.

[14] Bailey Aidan, 'Scalable Defeasible Reasoning', Thesis, University of Cape Town, Cape Town, 2021.

[15] M. Freund, 'Preferential reasoning in the perspective of Poole default logic', *Artif. Intell.*, vol. 98, no. 1, pp. 209–235, Jan. 1998, doi: 10.1016/S0004-3702(97)00053-2.

[16] Hamilton Joel, 'An Investigation into the Scalability of Rational Closure', Thesis, University of Cape Town, Cape Town, 2021.

[17] Evashna Pillay, 'An Investigation into the Scalability of Rational Closure V2', Thesis, University of Cape Town, Cape Town, 2022.