Extending Defeasible Reasoning beyond Rational Closure Literature Review

Thabo Vincent Moloi
University of Cape Town
Cape Town, South Africa
mlxtha036@myuct.ac.za

ABSTRACT

Logic plays a pivotal role in artificial intelligence (AI), serving as a foundational framework for representing and reasoning about knowledge. Classical reasoning, grounded in classical logic, employs deductive reasoning to derive conclusions from a set of premises, adhering to the principle of bivalence, which posits that propositions are either true or false. In contrast, defeasible reasoning accommodates the revision or override of conclusions in response to new information or exceptions, allowing for tentative and adaptable conclusions. This review aims to explore key facets of defeasible reasoning and assess existing user interface implementations in this context.

KEYWORDS

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

1 INTRODUCTION

Artificial Intelligence (AI), refers to the development of computer systems that can perform tasks that typically require human intelligence. These tasks include things like image recognition, and making decisions which depend on the existing knowledge. This knowledge needs to be structured in a way that computers are able to understand, represent and manipulate. Simple statement like *"Last year had 365 days"* is very easy for humans to understand which is not the case for computers.

Knowledge Representation and Reasoning (KRR) forms a crucial domain withing the field of AI, focusing on the symbolic representation of knowledge and its automated manipulation through reasoning programs [2]. Knowledge base is a collection of knowledge (real-world information) and it serves as a foundation for storing and managing, and retrieving the knowledge. The knowledge-based system can use reasoning to draw conclusions from the knowledge base.

Reasoning can contain exceptions where conclusions are drawn from uncertain or incomplete knowledge base. Using logic can help identify the exceptions and avoid contradictions within the knowledge base. This is where defeasible reasoning comes in. In defeasible reasoning, adding new information can lead to the retraction of previously inferred conclusions. The most widely recognized framework for defeasible reasoning is the KLM framework developed by Kraus, Lehmann and Magidor [4]. In the KLM approach, defeasible reasoning is modeled using non-monotonic logics, which allow for conclusions to be drawn tentatively, subject to revision in light of new information. The key idea behind the KLM approach is the use of preference orderings on rules to determine the strength of conclusions.

This literature review begins by delving into propositional logic and provides insigts into its syntax and semantics. It further discusses defeasible reasoning, and discussing the framework. Finally, it discusses the existing implementations.

2 PROPOSITIONAL LOGIC

Propositional logic is a formal system used to reason about knowledge or information and it abstracts concepts from natural language into a formal language [10]. It is a branch of logic that deals with propositions, which are declarative statements that are either true or false but not both.

The language can typically be constructed from a set of propositional variables or atoms, which represent statements that can either be true or false and the boolean connectives (such as AND, OR, NOT) are used to combine these variables to form more complex formulas [6]. The statements represent the real-world information, for example, "humans are mammals" or "an ostrich can fly". Sometimes a statement may have one or more statements, an example would be "the sky is blue and the sun is shining".

2.1 Syntax

Each propositional atom is associated with a Boolean value which can either be true, denoted as T, or false, denoted as F. The indivisible atoms are denoted by lowercase letter $p, q, r, \ldots$. The set of all of propositional atoms will be denoted as $\mathcal{P}$ which consists of all the statements such that $\mathcal{P} = \{p, q, r, \ldots\}$.

Atoms can be combined to form formulas using Boolean operators, also known as logical connectives, each with specific rules for their combinations [1]. The key connectives are negation, conjunction, disjunction, implication, and equivalence which are shown in Table 1.

Name	Meaning	Symbol	Example
negation	not	$\neg$	$\neg p$
conjunction	and	$\wedge$	$p \wedge q$
disjunction	or	$\vee$	$p \vee q$
implication	if then	$\rightarrow$	$p \rightarrow q$
equivalence	if and only if	$\leftrightarrow$	$p \leftrightarrow q$

Table 1: Boolean operators in propositional logic

All of the operators in Table 1 are binary, except negation. This means that they take two operands while negation only takes one operand. The order of precedence for logical operators is as follows: negation has the highest precedence, followed by conjunction, then disjunction, and finally implication and equivalence, which have the lowest precedence [8]. The propositional language $\mathcal{L}$ is formed by recursively defining propositional formulas using the binary operators, and the negation operator $\neg$ over the set $\mathcal{P}$. The set of all propositional formulas $\mathcal{L}$ can be defined as $\forall p \in \mathcal{P}, p \in \mathcal{L}$ and if $\alpha, \beta \in \mathcal{L}$ then $\neg\alpha, \alpha \wedge \beta, \alpha \vee \beta, \alpha \rightarrow \beta, \alpha \leftrightarrow \beta \in \mathcal{L}$.

2.2 Semantics

In propositional logic, the meaning of formulas is defined by assigning truth values to the atoms of a formula, similar to evaluating arithmetic expressions by assigning values to variables [1]. A valuation u is a function that maps propositions P to truth values expressed as $u : \mathcal{P} \rightarrow \{T, F\}$ where and T and F represent true and false respectively [10]. One random valuation that could be $p\bar{q}r$, read as p is true, q is false and r is true. This can be represented in a truth table.

p	q	$\neg p$	$p \wedge q$	$p \vee q$	$p \rightarrow q$	$p \leftrightarrow q$
T	T	F	T	T	T	T
T	F	F	F	T	F	F
F	T	T	F	T	T	F
F	F	T	F	F	T	T

Table 2: Truth values for p and q

3 DEFEASIBLE REASONING

Defeasible reasoning is a type of reasoning where conclusions are derived from information that is incomplete or uncertain, and these conclusions can be revised or overridden when new information becomes available. Defeasible reasoning differs from classical reasoning primarily in how it handles conclusions. In classical reasoning, once a conclusion is derived from a set of premises, it is considered definitive and does not change unless there is a logical error or new information that contradicts it. On the other hand, in defeasible reasoning, conclusions are tentative and can be overridden or revised based on new information or exceptions to the general rules. This flexibility allows defeasible reasoning to model real-world scenarios more accurately, where certainty is often limited and subject to change.

Let $\mathcal{K}$ be a knowledge base with the following statements.

- $b \rightarrow f$
- $b \rightarrow w$
- $p \rightarrow b$
- $p \rightarrow \neg f$

Which means birds can fly, birds wings, penguin is a bird, penguins cannot fly respectively. In classical reasoning, we can conclude that $p \rightarrow f$ since $b \rightarrow f$ and $p \rightarrow b$. Penguins, however, cannot fly ($p \rightarrow \neg f$) according to our knowledge base. Since $p \rightarrow f$ and $p \neg f$, then it means that p is false, meaning that penguins do not exist which is not the case. In defeasible reasoning, this exception is acknowledged.

3.1 KLM framework

The KLM approach introduced a conditional assertion $\alpha \mathrel{|\!\sim} \beta$, if α and β are formulas. $\alpha \mathrel{|\!\sim} \beta$ is read as "α typically implies β" [11, 12]. Unlike a strict logical implication, this statement allows for exceptions or cases where the association does not hold true

Revisiting the our knowledge base, we can replace the symbol $\rightarrow$ with $\mathrel{|\!\sim}$ in statements where we can have exceptions. So, the new knowledge base would be

$$\mathcal{K} = \{b \mathrel{|\!\sim} f, b \rightarrow w, p \rightarrow b, p \mathrel{|\!\sim} \neg f\}$$

From the new knowledge, it can be seen that "birds typically fly", "birds have wings", "penguin is a bird", and "penguins typically don't fly". This allows us to draw conclusions from uncertain information.

KLM outlines a set of criteria that an entailment relation must satisfy to be considered rational. Specifically, an entailment relation $\approx$ is deemed rational if it adheres to the following principles [9, 11, 12].

$$\text{(Refleivity)} \qquad \mathcal{K} \approx \alpha \mathrel{|\!\sim} \alpha$$

$$\text{(Left Logical Equivalence)} \qquad \frac{\mathcal{K} \approx \alpha \leftrightarrow \beta, \mathcal{K} \approx \alpha \mathrel{|\!\sim} \gamma}{\mathcal{K} \approx \beta \mathrel{|\!\sim} \gamma}$$

$$\text{(Right Weakening)} \qquad \frac{\mathcal{K} \approx \alpha \rightarrow \beta, \mathcal{K} \approx \gamma \mathrel{|\!\sim} \alpha}{\mathcal{K} \approx \gamma \mathrel{|\!\sim} \beta}$$

$$\text{(And)} \qquad \frac{\mathcal{K} \approx \alpha \mathrel{|\!\sim} \beta, \mathcal{K} \approx \alpha \mathrel{|\!\sim} \gamma}{\mathcal{K} \approx \alpha \mathrel{|\!\sim} \beta \wedge \gamma}$$

$$\text{(Or)} \qquad \frac{\mathcal{K} \approx \alpha \mathrel{|\!\sim} \gamma, \mathcal{K} \approx \beta \mathrel{|\!\sim} \gamma}{\mathcal{K} \approx \alpha \vee \beta \mathrel{|\!\sim} \gamma}$$

$$\text{(Continuous Monotonicity)} \qquad \frac{\mathcal{K} \approx \alpha \mathrel{|\!\sim} \gamma, \mathcal{K} \approx \alpha \mathrel{|\!\sim} \beta}{\mathcal{K} \approx \alpha \wedge \beta \mathrel{|\!\sim} \gamma}$$

$$\text{(Rational Monotonicity)} \qquad \frac{\mathcal{K} \approx \alpha \mathrel{|\!\sim} \gamma, \mathcal{K} \not\approx \alpha \not\mathrel{|\!\sim} \beta}{\mathcal{K} \approx \alpha \wedge \beta \mathrel{|\!\sim} \gamma}$$

Lehmann and Magidor suggest that defeasible entailment should adhere to all the KLM properties mentioned earlier [3]. This concept is referred to as LM-rationality.

3.1.1 ***Preferential Interpretations.*** Preferential semantics establishes an order among valuations, indicating that if an agent favors one valuation u over another v, the agent is likely to prioritize u over v [10]. The concept of preference can be interpreted in various ways, with one interpretation being that more typical valuations are favored.

3.1.2 ***Ranked intepretations.*** The concept of preferential interpretations is refined to create ranked interpretations by introducing conditions that restrict the partial order for the interpretation [12]. These restrictions lead to an ordering that establishes a series of non-empty levels, where interpretations on the lower levels are considered more "typical" or preferred.

Moreover, they extend the definition of a preferential consequence relation to define a rational consequence relation, which must satisfy all seven of the KLM properties. They demonstrate that ranked interpretations can define rational consequence relations, and conversely, every rational consequence relation can be defined by a ranked interpretation.

Formally, a ranked interpretation is a function $\mathcal{R} : \mathcal{U} \rightarrow \mathcal{N} \cup \{\infty\}$ which must satify the following conveity property: $\forall i \in \mathcal{N}, \exists u \in \mathcal{U}$ such taht $\mathcal{R}(u) = i$, then there must be a $v \in \mathcal{U}$ such that $\mathcal{R}(v) = j$ where $0 \leq j < i$ [10]. Defeasibility is addressed by providing semantics for the operator $\mathrel{|\!\sim}$ using ranked interpretations. In a ranked interpretation $\mathcal{R}$, a defeasible implication $\alpha \mathrel{|\!\sim} \beta$ holds if, for all the most typical interpretations in which α is true (i.e., the models of α on the lowest level containing an interpretation that satisfies α), β is also true.

Consider a set of propositions $\mathcal{P} = \{p, q, r\}$, the ranked interpretation would be

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

Table 3: Ranked interpretations for $\mathcal{P} = \{p, q, r\}$

Table 3 shows that $\mathcal{R}(\overline{pqr}) = 0$ which shows that $\overline{pqr}$ is most typical while $\overline{pq}r$ and $\bar{p}q\bar{r}$ are impossible as their rank is ∞.

3.2 Rational Closure

The rational closure algorithm enables us to conduct a type of rational inference on defeasible knowledge bases [5]. To do so, each statement in the relevant knowledge base must be given a ranking. Rational closure is the most conservative form of defeasible entailment. This means that anything entailed by rational closure will also be entailed by the other common forms of defeasible entailment [14].

3.2.1 Minimal Ranked Entailment. Involves establishing a partial ordering $\preceq_{\mathcal{K}}$ of all ranked models of knowledge base $\mathcal{K}$. The ordering is defined such that for any two ranked interpretations $\mathcal{R}_1$ and $\mathcal{R}_1$, $\mathcal{R}_1 \preceq_{\mathcal{K}} \mathcal{R}_2$ if for every valuation $v \in \mathcal{U}$, $\mathcal{R}_1(v) \leq \mathcal{R}_2(v)$ [4]. The most typical models are ranked lower which leads to the identification of a minimal element [6]. This minimal element is denoted as $\mathcal{R}^{\mathcal{K}}_{RC}$. Minimal ranked entailment, denoted by $\approx$, is then formulated based on a defeasible knowledge base $\mathcal{K}$ and its corresponding ranked interpretation $\mathcal{R}^{\mathcal{K}}_{RC}$. For any defeasible implication $\alpha \mid\sim \beta$, $\mathcal{K} \approx \alpha \mid\sim \beta$ holds if and only if $\mathcal{R}^{\mathcal{K}}_{RC}$ entails $\alpha \mid\sim \beta$. This approach to entailment, based on ranked interpretations, is shown to have similarities to preferential entailment.

3.2.2 Materialisation. To ascertain the exceptional nature of a statement, it is imperative to materialize defeasible statements [5]. This involves defining the counterpart of each defeasible implication $\alpha \mid\sim \beta$ as the propositional formula $\alpha \rightarrow \beta$. The collection of all such counterparts for $\mathcal{K}$ is represented as $\overrightarrow{\mathcal{K}}$. A propositional statement α is deemed exceptional in relation to a knowledge base $\mathcal{K}$ if and only if $\mathcal{K} \vDash \neg\alpha$, meaning α is false in all the most typical interpretations across every ranked model in $\mathcal{K}$.

4 EXTENDING IMPLEMENTATION

This section will review the existing implementations of a user interface of defeasible reasoning.

4.1 A SILK Graphical UI

Extending the user interface (UI) in defeasible reasoning involves building upon existing implementations, as exemplified by the work presented in "A SILK Graphical UI for Defeasible Reasoning, with a Biology Causal Process Example" by Benjamin Grosof et al [7]. The paper introduces a graphical user interface (GUI) for knowledge entry, query answering, and justification browsing within the SILK system, aiming to facilitate user specification and comprehension of advanced courteous prioritized defeasible reasoning. The GUI, implemented as a plug-in to the Eclipse Integrated Development Environment (IDE), offers features such as syntactic validation, color-coded components for clarity, and automatic tracking of query results against user changes to the rules. It also provides query result justification trees for exploration, aiding in user debugging and what-if explorations. This work represents a significant step towards enhancing user-friendliness and functionality in defeasible reasoning systems.

4.2 Defeasible-Inference Platform

The Defeasible-Inference Platform (DIP) is introduced as an ontology engineering tool with acceptable performance, particularly the Rational Closure algorithm, which is suitable for current use [13]. The Lexicographic and Relevant Closures, while less promising, are still useful for experimental purposes with smaller ontologies. Integration of these algorithms into ontology editing tools, such as Protégé, is highlighted as a potential improvement. Protégé, an open-source ontology editor and knowledge base framework, offers a tabbed UI displaying different ontology views, including Classes, Object Properties, Individuals, and Entities. It supports various ontology reasoners as plugins, highlighting inferred information with color-coded visual cues for better user understanding. The section concludes with instructions on accessing and utilizing DIP within Protégé.

5 CONCLUSIONS

Artificial Intelligence (AI) has significantly advanced with the development of computer systems capable of performing tasks that traditionally required human intelligence. Knowledge Representation and Reasoning (KRR) within AI focuses on symbolically representing and manipulating knowledge. Defeasible reasoning, a form of reasoning that handles uncertain or incomplete information, plays a crucial role in KRR. The KLM framework, developed by Kraus, Lehmann, and Magidor, provides a foundational framework for defeasible reasoning, utilizing non-monotonic logics to draw tentative conclusions subject to revision.

Propositional logic serves as the foundation for understanding defeasible reasoning, providing syntax and semantics for logical statements. Defeasible reasoning diverges from classical reasoning by allowing conclusions to be tentative, subject to change based on new information or exceptions.

The KLM framework sets criteria for rational entailment, emphasizing properties like reflexivity, logical equivalence, and monotonicity. Preferential and ranked interpretations refine the framework, introducing notions of preference and typicality in reasoning.

Rational closure emerges as a conservative form of defeasible entailment, ensuring that any conclusion entailed by rational closure is also entailed by other forms of defeasible entailment. The algorithm involves assigning rankings to statements in a knowledge base, enabling rational inference based on these rankings.

Extensions in implementation, such as graphical user interfaces (GUIs) like SILK and the Defeasible-Inference Platform (DIP), enhance the usability and performance of defeasible reasoning systems. These implementations offer intuitive interfaces and integration with ontology editing tools like Protégé, making them accessible for practical applications in ontology engineering environments.

Overall, defeasible reasoning, with its foundations in logical frameworks like the KLM framework and implementations like DIP, presents a powerful approach to handling uncertainty and incomplete information in AI systems, with promising applications in various domains.

REFERENCES

[1] Mordechai Ben-Ari. 2012. *Mathematical Logic for Computer Science*. Springer Science & Business Media.

[2] Ronald J Brachman and Hector J Levesque. 2004. *Knowledge Representation and Reasoning*. Elsevier.

[3] Giovanni Casini, Thomas Meyer, and Ivan Varzinczak. 2018. Defeasible entailment: From rational closure to lexicographic closure and beyond. In *Proceeding of the 17th International Workshop on Non-Monotonic Reasoning (NMR 2018)*. 109–118.

[4] Giovanni Casini, Thomas Meyer, and Ivan Varzinczak. 2019. Taking defeasible entailment beyond rational closure. In *European Conference on Logics in Artificial Intelligence*. 182–197.

[5] Victoria Chama. 2019. *Explanation For Defeasible Entailment*. Master's Thesis. Faculty of Science, University of Cape Town.

[6] Laura Giordano, Valentina Gliozzi, Nicola Olivetti, and Gian Luca Pozzato. 2015. Semantic characterization of rational closure: From propositional logic to description logics. *Artificial Intelligence* 226 (2015), 1–33.

[7] Benjamin Grosof, Mark Burstein, Mike Dean, Carl Andersen, Brett Benyo, William Ferguson, Daniela Inclezan, and Richard Shapiro. 2010. A SILK graphical UI for defeasible reasoning, with a biology causal process example. In *Proceedings of the 2010 International Conference on Posters & Demonstrations Track - Volume 658* (Shanghai, China) *(ISWC-PD'10)*. CEUR-WS.org, Aachen, DEU, 113–116.

[8] Cordelia Hall and John O'Donnell. 2000. *Propositional Logic*. Springer London, London. 35–87 pages. https://doi.org/10.1007/978-1-4471-3657-6_2

[9] Michael Harrison and Thomas Meyer. 2020. DDLV: A system for rational preferential reasoning for datalog. *South African Computer Journal* 32, 2 (2020). https://doi.org/10.18489/sacj.v32i2.850

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

[11] Sarit Kraus, Daniel Lehmann, and Menachem Magidor. 1990. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence* 44, 1-2 (1990), 167–209.

[12] 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

[13] Kody Moodley. 2015. *Practical Reasoning for Defeasible Description Logics*. Ph. D. Dissertation. University of KwaZulu-Natal, Durban, South Africa.

[14] Matthew Morris, Tala Rossa, and Thomas Meyer. 2020. Algorithmic definitions for KLM-style defeasible disjunctive Datalog. *SACJ* 32, 2 (December 2020). Research Article.