An overview of defeasible entailment

Adam Kaliski[1] and Thomas Meyer[1]

CAIR and University of Cape Town, South Africa
KLSADA002@myuct.ac.za tmeyer@cs.uct.ac.za

1 Introduction and Background

The process of determining whether or not a statement, or *query*, is inferred by a knowledge base is entailment. Classical entailment is regular deduction, which holds that if

1. All humans are mortal
2. Socrates is a human

then logically Socrates is mortal. This form of reasoning is, however, completely *monotonic*, which does not allow for exceptions to stated rules, and therefore also does not allow for the concept of *typicality*. This inherently limits the expressivity of a given expert system, since reality often contains exceptions. Consider now the addition of a *defeasible* rule (along with a regular rule)

3. Humans are *typically* not philosophers
4. Socrates is a philosopher

Defeasibility is the introduction of typicality, and the associated concept of exceptionality. A defeasible rule is one that allows itself to be broken, by stating that the rule is typically the case, rather than always the case. However, adding this semantic extension to classical logic raises an important problem regarding reasoning, because a knowledge base with defeasible rules cannot be reasoned about in the same way as a classical knowledge base. Furthermore, whereas classical, monotonic entailment is unique, there are arbitrarily many different defeasible entailments, and therefore what a defeasible entailment relation will infer from a given defeasible knowledge base is dependent on which entailment it is. Given the above knowledge base regarding Socrates, the question is can we conclude that Socrates is mortal? The answer is dependent on whether or not we are using a *prototypical* or *presumptive* entailment. A prototypical entailment is conservative, and draws fewer inferences, and will therefore conclude that Socrates is an *atypical* human, and therefore cannot conclude that he is mortal. On the other hand, a presumptive entailment, which draws as many inferences as it reasonably can, will not see any clashes regarding mortality, and infer as much as possible, including that he is mortal.
There have been a number of approaches to characterizing defeasible reasoning. Circumscription [9, 10] was one of the earliest attempts at reasoning about typicality and exceptions. Belief revision [1] is another method that allows for a form of defeasible reasoning. Other approaches to defeasible reasoning include default logic [11, 12], as well as auto-epistemic logic [8, 4].

2 Problem Statement

The defeasible reasoning framework that this paper uses was defined by Kraus, Lehmann, and Magidor [5, 7, 6], as well as extensions by Casini et. al. [3, 2]. The KLM framework is a very strong candidate for defining defeasible entailment relations, because it defines a number of properties and postulates that isolate a subset of defeasible entailments as *rational*. This provides a theoretical bedrock which further work can use as a foundation to extend or modify the properties to describe a defeasible entailment useful to them. Another argument in favour of this framework is that it is computationally well behaved, reducible to classical entailment. However, there are many defeasible entailments that make little sense from a human reasoning standpoint that nevertheless satisfy the KLM properties. This implies that the postulates defining a rational defeasible entailment are too permissive to isolate only those entailments that are useful. Extending the KLM postulates to isolate only those defeasible entailments that have meaning would allow for far better analysis of entailments, and allow for a range of extensions and applications using this framework.
The KLM framework is defined by three foundational papers, and then further extended by later work. A significant barrier of entry to those who wish to either apply or extend this framework is to understand it. The initial KLM papers are both dense and also conflict with one another in certain respects. For example, the meaning of symbols were changed between papers without comment. There is therefore an argument to be made that an overview that compiles the current state of the framework in a single document would enable future research groups to far more easily utilize the KLM framework.

3 Aims

A main aim of this paper is how best to extend these properties to isolate exactly the defeasible entailments that infer in a useful, or understandable way. Work on this has already been done by Casini et. al. [3], and is included in this paper. Further extensions in this paper include methods for defining any defeasible entailment relation based on extending or constraining the set of inferences, and the properties and distinctions between what is termed basic defeasible entailment, and rational defeasible entailment.
An overarching aim of this work is to provide a single point of reference for any project that wishes to extend, constrain, or otherwise modify the framework to define a defeasible entailment of interest, or any project that wishes to implement an already defined defeasible entailment into an application. In particular, the paper details the semantics, properties, and the algorithms associated with various monotonic and nonmonotonic entailments for defeasible knowledge bases, precise definitions for defeasibility and non-monotonicity, and syntax-sensitive entailments.

References

1. Alchourrón, C.E., Gärdenfors, P., Makinson, D.: On the logic of theory change: Partial meet contraction and revision functions. The journal of symbolic logic **50**(2), 510–530 (1985)
2. Casini, G., Meyer, T., Moodley, K., Nortjé, R.: Relevant closure: A new form of defeasible reasoning for description logics. In: Fermé, Eduardoand Leite, J. (ed.) Logics in Artificial Intelligence. pp. 92–106. Springer International Publishing, Cham (2014)
3. Casini, G., Meyer, T., Varzinczak, I.: Taking defeasible entailment beyond rational closure. In: European Conference on Logics in Artificial Intelligence. pp. 182–197. Springer (2019)
4. Gelfond, M.: On stratified autoepistemic theories. In: AAAI. vol. 87, pp. 207–211 (1987)
5. Kraus, S., Lehmann, D., Magidor, M.: Nonmonotonic reasoning, preferential models and cumulative logics. Artificial intelligence **44**(1-2), 167–207 (1990)
6. Lehmann, D.: Another perspective on default reasoning. Annals of Mathematics and Artificial Intelligence **15**(1), 61–82 (1995)
7. Lehmann, D., Magidor, M.: What does a conditional knowledge base entail? Artificial intelligence **55**(1), 1–60 (1992)
8. Marek, W., Truszczyński, M.: Autoepistemic logic. Journal of the ACM (JACM) **38**(3), 587–618 (1991)
9. McCarthy, J.: Circumscription—a form of non-monotonic reasoning. Artificial intelligence **13**(1-2), 27–39 (1980)
10. McCarthy, J.: Applications of circumscription to formalizing common sense knowledge. Artificial Intelligence **28**, 89–116 (1986)
11. Reiter, R.: A logic for default reasoning. Artificial intelligence **13**(1-2), 81–132 (1980)
12. Reiter, R., Criscuolo, G.: Some representational issues in default reasoning. Computers & Mathematics with Applications **9**(1), 15–27 (1983)