Non-Classical Reasoning: A Focus on Belief Revision

Claire Denny
dnncla004@myuct.ac.za
University of Cape Town

ABSTRACT

In reality, decision-making is commonly performed without complete information and certainty, exceptions exist and knowledge is not static. Classical reasoning, being monotonic i.e. deductive, is too inflexible to facilitate the style of reasoning this necessitates. In this literature review, we investigate non-classical forms of reasoning that model the non-monotonic nature of human reasoning. In particular, we investigate belief revision. Other non-classical forms of reasoning reviewed as related work are defeasible reasoning and belief update. Each can be described in two ways: using semantics and using postulates. We give descriptions using postulates, but include a brief overview of the semantics in aid of a cohesive understanding of the three forms of reasoning presented in this paper. We find belief revision, belief update and defeasible reasoning to be linked to each other. We also find that belief revision reflects more than one characteristic of human reasoning, but the extent to which these two reasoning systems resonate is not known.

CCS CONCEPTS

• **Theory of computation** → **Logic**; • **Computing methodologies** → **Nonmonotonic, default reasoning and belief revision**; *Reasoning about belief and knowledge*;

KEYWORDS

propositional logic, defeasible reasoning, cognitive defeasible reasoning, belief revision, belief update, propositional logic, non-monotonicity

1 INTRODUCTION

In their everyday lives, humans are faced with incomplete knowledge but still must act [24]. As a result of their incomplete explicit knowledge, humans rely on background knowledge, heuristics and their logic's capacity to be flexible in what they believe about the world, allowing themselves to withdraw previously made conclusions, if necessary, and make new ones, given new evidence [19]. This flexibility in their reasoning classifies it as non-monotonic. Non-monotonicity is defined formally in [19]. Consider a medical practitioner or a mechanic - people whose professions require diagnoses of problems such that the appropriate action can be performed. In contexts where an action is urgent, the time required to reach deductive certainty may be too long [19], necessitating non-monotonic logic such as default logic. Lehmann and Magidor [16] acknowledge that humans are remarkably good at making and correcting inferences using their knowledge bases. Noting the importance of the property of flexibility in intelligence [23], which humans demonstrate daily, progress in artificial intelligence hinges on non-monotonic reasoning and its interactions with other intelligent activities being further researched [24]. Using classical propositional logic and its notation as a starting point, this paper reviews what is known about three types of non-classical reasoning, namely Belief Revision, Belief Update and Defeasible Reasoning. It goes into more depth regarding Belief Revision, with an aim of enabling investigation into whether it is a form of reasoning that people employ - and if it is, looking at the extent to which it corresponds.

2 CLASSICAL PROPOSITIONAL LOGIC

Propositional logic uses a formal language based on an alphabet of propositional variables [27]. Ragni [27] tells us that propositions are statements with a Boolean value (either true or false), the simplest of which are referred to as atoms, as they are indivisible. Propositional logic focuses on the ways in which statements can be combined or modified and the properties and relationships that arise from doing so.

There is a property called monotonicity that all classical propositional operators must satisfy [12]. Makinson [19] defines monotonocity as the principle that, given a set A, of propositions, if β follows from A, then β follows from any set B, where B $\supseteq$ A.

2.1 Notation and terminology

Reference will be made to interpretations, worlds or states of the world. These terms are equivalent, and refer to assignments of truth values for the relevant propositional alphabet [16]. A set of statements explicitly known about the world, known to the reasoner or agent, is called a knowledge base. The symbol $\models$ represents entailment or logical consequence. A model of α refers to a state of a world, ψ, where α is true, which is to say $\psi \models \alpha$ [27]. The notation m(α) can also be written as mod(α) and means the models of α, which is to say that $mod(\alpha) = \{\psi | \psi \models \alpha\}$ [27]. Saying that knowledge base K entails some statement α i.e. $K \models \alpha$, we have intuitively that m(K)$\subseteq$m(α), as if K contains statements in addition to α, the set of models of K would be more constrained than the set of all worlds where it is just α that must hold. In classical logic, there are also the following connectives with which all truth-functional connectives can be expressed [19]: $\neg$ (negation), $\wedge$ (conjunction), $\vee$ (disjunction). Material implication (if) is denoted as $\rightarrow$ and equivalence (iff) is denoted as $\leftrightarrow$. Classical consequence is denoted $\vdash$ and is considered such if it is consistent with what is known about the world.

The logic systems reviewed in later sections will assume an underlying logic that includes classical propositional logic.

3 BELIEF REVISION

In belief change, we have a classical knowledge base assumed to be correct. A propositional language is assumed, with a Tarskian consequence relation $C_n(\cdot)$ [2]. In the event of new information α, inconsistent with the knowledge base K, the situation must be handled such that consistency is restored, or preserved. The defeasible reasoning approach would be to flag the clash as defeasible, weakening the propositions previously in K that conflict; the *belief revision* approach, within belief change, would be to revise K by invalidating worlds which are sufficiently far from α [9], which is to say that the beliefs are modified such that there is consistency.

Alchourrón, Gärdenfors and Makinson (AGM) concern themselves with three operations on the knowledge base, which is to say, three types of theory changes: expansion, contraction and revision [1]. Expansion is performed where new, consistent information is added to theory or knowledge base K and the expanded set is closed under entailment [1, 2]. Contraction is performed where a proposition *alpha* previously in a theory or knowledge base K is rejected [1]. The operation yields a K' that does not entail α [2] - we say the contraction of K by α. It must also, however, be closed under entailment, so other propositions than α may need to be rejected along with α [1]. Revision, as discussed earlier, is performed on K in the case of proposition α such that the resulting K' is both consistent and closed under entailment [1, 2]. It can be seen as the composition of two sub-operations: expansion of the contraction of K by $\neg\alpha$, by α. Using notation introduced in [2], this can be written as $K^*_\alpha = (K^-_{\neg\alpha})^+_\alpha$, which is termed the *Levi identity* in [7]. In the AGM model [1], the operator $*$ is that of *partial meet* revision.

3.1 Notation and terminology

$C_n(\cdot)$ is an operation that takes a set of propositions K and returns a set of propositions K' [1]. In terms of notation, we can write $\beta \in C_n(A)$ as $A \vdash \beta$. This consequence relation includes classical tautology, is compact and satisfies *introduction of disjunction in the premises* [1]. A theory refers to a set of propositions closed under C_n. A set of propositions or beliefs, A, is consistent modulo $C_n \leftrightarrow \beta \wedge \neg\beta \notin C_n(A)$ for any proposition β.

3.2 Properties

The following are the properties or postulates of Belief Revision as defined by [1]. Any function $*$ that satisfies them is considered an AGM revision function. Properties 1-6 are the basic AGM properties specifically for belief revision, and properties 7-8 are supplementary AGM properties [2, 25].

1. Closure

$\text{K} * \alpha = C_n(\text{K} * \alpha)$

This implies logical omniscience on the part of the ideal agent or reasoner, including after revision of their belief set [25].

2. Success

$\text{K} * \alpha \models \alpha$

This expresses that the new information should always be part of the new belief set [25]. Peppas [25] also considers ways to relax this property, given that it places substantial trust in the reliability of α.

3. Inclusion

$\text{K} * \alpha \subseteq C_n(\text{K} \vee \{\alpha\})$

4. Vacuity

If $\neg\alpha \notin \text{K}$ then $C_n(\text{K} \vee \{\alpha\}) \subseteq \text{K} * \alpha$

Properties 3 and 4 are motivated by the principle of minimum change [25]. Together, they express that in the case of information α, consistent with belief set or knowledge base K, belief revision involves performing expansion on K by α i.e. none of the original beliefs need to be withdrawn.

5. Consistency

$\text{K} * \alpha = C_n(\alpha \wedge \neg\alpha)$ only if $\models \neg\alpha$

This expresses that the agent should prioritise consistency, where the only acceptable case of not doing so is if the new information, α, is inherently inconsistent - in which case, *success* overrules *consistency* [25].

6. Extensionality

If $\alpha \equiv \phi$ then $\text{K} * \alpha = \text{K} * \phi$

This is also known as the *irrelevance of syntax* postulate. It effectively expresses that the content i.e. the belief represented, and not the syntax, affects the revision process, in that logically equivalent sentences or beliefs will cause logically equivalent changes to the belief set or knowledge base [25]. This property would not hold without the notion of *epistemic entrenchment* (degree of resistance to change [25]) or Katsuno and Mendelzon's treatment of integrity constraints [9].

7. Superexpansion

$\text{K} * (\alpha \wedge \phi) \subseteq C_n(\text{K} * \alpha \vee \{\phi\})$

8. Subexpansion

If $\neg\phi \notin \text{K} *$ then $C_n(\text{K} * \alpha \vee \{\phi\}) \subseteq \text{K} * (\alpha \wedge \phi)$

Properties 7 and 8 are motivated by the principle of minimal change [25]. Together, they express that for two propositions α and ϕ, if in revising belief set K by α one obtains belief set K' consistent with ϕ, then to obtain the effect of revising K with $\alpha \wedge \phi$, simply perform expansion on K' with ϕ. In short, $K * (\alpha \wedge \phi) = (K * \alpha) + \phi$.

3.3 Discussion

A belief state can be modelled by a belief set i.e. by a set K of sentences that is closed under logical consequences [6]. There is an argument for the need of a *belief base* B_K for a belief set K [6], where the base contains the explicit beliefs or beliefs of independent standing and the belief set K comprises $C_n(B_K)$. In this way, a distinction is made between basic and derived beliefs. The idea is that revisions are performed on the finite belief base, as opposed to the infinite belief set [6].

Rational consequence relations and revision operators are linked - the rules on the former can be interpreted in terms of the latter [10]. Moreover, belief change has connections to non-monotonic inference [18]. Makinson and Gärdenfors [20] study these on a syntactic level. Casini et al. [2] take this idea and previous results from a paper by Casini and Meyer, and explore integrating belief change and non-monotonic inference. They do this by looking at belief change for a preferential non-monotonic framework. They are not the first to study belief revision in a conditional framework [2] - previous approaches, giving the conditionals a subjunctive interpretation and using known connections between the conditionals and belief revision operators, have been taken by Kern-Isberner and Wobcke [2]. These approaches faced a problem of defining revision operators that avoid Gärdenfors' *impossibility result* [5], that arises due to the *Ramsey Test* ($\phi > \alpha \in K \leftrightarrow \alpha \in K + \phi$) and the *preservation criterion* (effectively equivalent to Vacuity) being inconsistent with each other for non-trivial cases [6]. Cross and Thomason restrict the revision procedure and show a theory of conditionals that satisfy the Ramsey Test can be found [6]. In contrast, Casini et al. [2] do not give the conditionals a subjunctive interpretation, and use conditional knowledge bases. Having the conditionals as the objects of the belief change implies that [2] does not have the impossibility problem. Other approaches to revision operators include a system of spheres [10].

Belief revision systems are defined in [21] as Artificial Intelligence programs dealing with contradictions. Both theoretical studies and practical implementations have been performed [21]. Research in this area has to address several problems: the inference problem, the non-monotonicity problem, dependency recording, disbelief propagation and the revision of beliefs. Martins and Shapiro [21] elaborate on these and explain the issues involved in each.

Regarding the belief revision postulates or properties, there exists more than one function $*$ that satisfies the AGM properties of belief revision [25]. This is not, however, a weakness - Peppas [25] argues that it simply expresses that people may change their minds in different ways to one another.

3.4 Examples

Gärdenfors [6] notes three main methodological questions to resolve: the representation of beliefs in the knowledge base, the relation between explicitly represented elements and derived beliefs, and the decision process regarding what to retract. He illustrates the importance of these with an example [6]: consider a knowledge base that includes information α (All European swans are white), β (The bird caught in the trap is a swan), γ (The bird caught in the trap comes from Sweden) and δ (Sweden is part of Europe). A logical inference from this information is belief ϵ that the bird caught in the trap is white. Suppose we receive the fact that the bird caught in the trap is black. We would want to add $\neg\epsilon$ to the knowledge base, but this would result in an inconsistent collection of information, so a decision is necessitated regarding choosing what propositions to retract prior to adding $\neg\epsilon$. In revision situations, one idea is that information loss from revisions should be minimal whereas another idea is that some beliefs are deemed more entrenched than others and so the least important ones should be retracted [6]. In this example, we can retract α, but then must decide which of its logical consequences we wish to retain.

Another example is found in [2]. In this example, we have an alphabet A = {a, m, n, v}. Respectively, these propositions represent being an avian red-blood cell, being a mammalian red-blood cell, being a vertebrate red-blood cell, and having a nucleus. Consider the situation where knowledge base K = $\{v \vdash n, a \vdash v, m \vdash v, m \vdash \neg n\}$ and $\neg m \in C_n(K)$. The presence of $\neg m$ in $C_n(K)$ when mammalian red-blood cells exist, leads to a conflict. The response is to revise the knowledge base. In the framework that Casini et al. [2] propose, the conflict can resolved by weakening [27] the conflicting proposition(s) already in the knowledge base. In propositional belief change, the conflict would be resolved by eliminating some information, likely either $m \vdash v$, $m \vdash \neg n$ or $v \vdash n$.

Consider a murder trial. α and β are our primary suspects. Initial belief base $K = \{(\alpha \wedge \neg\beta) \vee (\neg\alpha \wedge \beta)\}$. We believe one person committed the crime, and we believe it was either α or β i.e. if one of them is innocent, then the other is guilty. During the trial, testimonies are received that incriminate first α and then β. Given that we believe $\alpha \implies \neg\beta$ and $\beta \implies \neg\alpha$, the testimonies yield $(K * \alpha) * \beta \models \neg\alpha$ i.e. we believe β committed the murder. In this example, the order in which the information is received affects our final beliefs. This is a problem noted in cases of belief revision iteration [3].

3.5 Semantics Overview

We recall that $\Vdash$ and $\models$ denote satisfaction and entailment respectively. In belief revision, we consider an ordered or ranked preferential structure of worlds. Ordering is by typicality or normality, where the most typical interpretation is minimal [18, 31]. For belief base K, $m(K)$ having minimal rank in the structure, we have that $m(K * \alpha) = \min \leq (\alpha)$ and say that $\leq$ is a pre-order i.e. a reflexive and transitive relation [10]. We say that this pre-order is K-faithful, as the three conditions of faithfulness in [10] are satisfied and the models of K are at the lowest rank or row in the structure. In belief revision, we are essentially comparing distances between theories, wanting the theories the minimum distance away in which the new information holds [18]. Lewis terms a probabilistic approach in this vein as *imaging* [30]. The result of the revision operation depends on the ordering - semantically, this reflects the plurality of orderings of closeness that exist between candidate or possible worlds [8].

4 BELIEF UPDATE

Belief update is a belief change operation, as is belief revision [7, 9, 10, 13]. As such, in belief update as in belief revision, we assume a classical knowledge base and a classical propositional language, with a Tarskian consequence relation $C_n(\cdot)$. In the event of new information or input corresponding to a change in the world, an update operation is performed [8]. Else, a revision operation is performed [8]. This can be further clarified: the agent's interpretation of the new information is what determines the choice of operation to perform. If the new information is interpreted as indication the world has changed i.e. there is a dynamic state of affairs, then the choice is belief update. If the new information is interpreted to indicate that the information previously known must be incorrect or flawed i.e. there is a static state of affairs, then the choice is belief revision. New information μ can thus understood as an action effect [13], of which there are two possible types: an *ontic (physical) effect* and an *epistemic effect*. Belief update as a operation can be understood as a form of action progression [13]. The belief change or theory change operation called erasure is to belief update as contraction is to belief revision.

In the event of new information μ, the belief update approach would be to take each model of knowledge base or belief base ψ and update it to be a model of μ by as minimal a change as possible [9]. We consider the properties characterising belief update operators as proposed by Katsuno and Mendelzon [9] in Section 4.2, noting that these properties characterise more than one such operator [8].

4.1 Notation and terminology

We denote an update operation as $\diamond$, defining it as a function accepting input μ, to be applied to a belief base ψ to yield a new belief base $\psi' = \psi \diamond \mu$. The terms belief base and knowledge base are used interchangeably, as are the terms belief and theory. Updates are performed world by world [8]. We defined world in Section 2.1 and gave notation for models. In belief change literature, there is a preferred notation of $[[\delta]]$ for the models of sentence δ, where belief bases can be represented by sentences, as can new information. Using this notation, we have $\psi \in [[\delta]]$ to communicate that ψ is a model of δ i.e. a world in which δ holds. The term *ontic* means feedback-free [13] and *epistemic* means relating to knowledge or to the degree of validation of knowledge.

4.2 Properties

The following are the properties of Belief Update as defined by [9].
1. Success
$\psi \diamond \mu \models \mu$
This property expresses that priority is given to the input [8]. An update to ψ must satisfy the input μ, irrespective of the content of ψ.
2. Vacuity
If $\psi \models \mu$ then $\psi \diamond \mu \equiv \psi$
This expresses that if the input is vacuous, meaning that if $\mu \in C_n(\psi)$ (μ derivable from ψ [9]), then no change is necessary [8].
3. Consistency
$\not\models \neg\psi$ and $\not\models \neg\mu$ then $\not\models \neg(\psi \diamond \mu)$
In words, this reads as: if ψ and μ are consistent, then the update of ψ by μ is consistent too. Effectively, it expresses that if there exists inconsistency in the updated belief base, it is because there is inconsistency in the original base or in the input [8].
4. Irrelevance of Syntax
$\psi_1 \equiv \psi_2$ and $\mu_1 \equiv \mu_2$ then $\psi_1 \diamond \mu_1 \equiv \psi_2 \diamond \mu_2$
This expresses that update operations should not be syntax-sensitive regarding the belief base and the input [8]. This property corresponds to the Extensionality property for belief revision.
Properties 5-8 do not have titles.
5. $(\psi \diamond \mu) \wedge \phi \models \psi \diamond (\mu \wedge \phi)$
This expresses that an update by the conjunction of two sentences is weaker than updating by one sentence and then just adding the other [8].
6. If $\psi \diamond \mu_1 \models \mu_2$ and $\psi \diamond \mu_2 \models \mu_1$ then $\psi \diamond \mu_1 \equiv \psi \diamond \mu_2$
This expresses that sentences which are equivalent under the belief base, when received as new information, lead to equivalent updates [8].
7. If ψ is complete, $(\psi \diamond \mu_1) \wedge (\psi \diamond \mu_2) \models \psi \diamond (\mu_1 \vee \mu_2)$
This property only makes sense if the language is finite [8], as we say that a base ψ is complete $\leftrightarrow$ for every sentence μ, either $\psi \rightarrow \mu$ or $\psi \rightarrow \neg\mu$ [8]. A complete knowledge or belief base therefore contains no uncertainty as to what the possible worlds are [8]. This property expresses that if updating by μ_1 yields a possible world that can also be produced by updating by μ_2, then updating by the disjunction of μ_1 and μ_2 should yield that same world.
8. $(\psi_1 \vee \psi_2) \diamond \mu \equiv (\psi_1 \diamond \mu) \vee (\psi_2 \diamond \mu)$
This property corresponds to model-wise updating [8], linking it to the semantic definition of the update operation in that an update is performed on a belief base world by world or model by model [8].

4.3 Discussion

The difference between the belief update and belief revision operations is first noted in [9], in the context of extended relational databases [8]. Keller and Winslett termed the belief update operation as *change-recording* and the belief revision operation as *knowledge-adding* [9]. Katsuno and Mendelzon [9] extend Keller and Winslett's work [11], formalising their informal distinction of the two operations, working from a more generalised setting and considering a more extensive set of cases. Katsuno and Mendelzon then take it further, propounding a way to combine belief update and belief revision into one operator parameterised by time [9]. Lang [13] uses this idea of time as a parameter in his proposal for belief update as an action progression and new information as an action effect. He explores the question of whether there can be an operator for action regression as interpreted in [14], given that he believes there is one for action progression. He explores this, terming it *reverse update* and propounding properties that should characterise such an operator. Boutilier also combines belief update and belief revision, based on a propositioal framework [8]. Other approaches to belief update include frameworks using situation calculus and theories of action as the basis [4, 8, 13].
Lang [13] raises the point that property 8 characterising it restricts the agent from using past beliefs to infer new ones given new information at a later point in time. Papers such as [8] argue that properties or postulates 2, 5, 6, and 7 are controversial and should not necessarily be required of belief update operators, and further explore an additional property. We thus have literature that refers to *KM update operators* (those that satisfy properties 1-8), those that refer to *basic update operators* (those that satisfy properties 1, 3, 4, 8 - despite being controversial, property 4 is argued still to be desirable [8]) and those that are classified *inertial* (those that satisfy properties 1, 2, 3, 4, 8 - despite being controversial, property 2 is argued still to be desirable [8]) [13]. Herzig and Rifi [8] break down the KM properties into sub-postulates to identify from where the controversy stems. They also examine ten update operations from the literature, characterising them in terms of strength and computational complexity for comparison purposes and evaluating them with regard to the KM properties. Update operations in literature, that are not based on orderings, are: the Possible Models Approach (PMA) i.e. minimal change, FORBUS i.e. numeric minimal change, Minimal Change with Maximal Disjunctive Inclusion (MCD), Winslett's Standard Semantics (WSS), WSS↓ i.e. making WSS insensitive to syntax, Minimal Change with Exceptions (MCE) i.e. taking WSS↓ and making it conservative [8].
Herzig and Rifi [8] study the role and issues of integrity constraints in belief change. In a critique of [9], they find that two of the ten update operations in literature satisfy their criteria. Components of the belief update operation such as *belief erasure* are discussed in [4, 9, 22]. In [9], contraction and erasure are compared. The semantic perspective on belief update is given in [9], which requires property 5 and 6 [8].

4.4 Examples

[9] offers us the following example: the agent has an initial knowledge base ψ comprising that there is either a book (b) or a magazine (m) on the table, but not both. New information is later received - there is a book on the table (b). The revision approach would be for the agent to interpret the information simply as more information about a fixed state of events, and thus for the agent to conclude that there is thus no magazine on the table ($\neg m$). The update approach would be for the agent to interpret the information as that, since receiving the initial information, a book has been placed on the table, The agent thus cannot conclude that, given there is a book on the table, there is no magazine on the table ($\not\models \neg m$). This example highlights that when performing revision, it is because new information about the world has been received, whereas when performing update,

it is because the world has changed, casting prior knowledge of the world into uncertainty.

Using this example, we adjust it to fit a unifying approach suggested by Katsuno and Mendelzon [9]. This is done by adding a time parameter to each sentence. The belief base now comprises pairs of the form <sentence μ, time t> that are read as μ *holds at point in time* t. Let us denote the time that the initial information is received by t_1 and the time that the new information is received, for the update operation, by t_2. The initial knowledge base ψ contains $< \neg(b \leftrightarrow m), t_1 >$. Performing revision of ψ by b yields a knowledge base that incorporates $< b, t_1 >$. Performing update of ψ by b yields a knowledge base that incorporates $< b, t_2 >$. It follows from this that in the case of revision, the $< \neg m, t_1 >$ is implied, but is not in the case of update. This is clear, as the new information interpreted by the agent for the update operation has a different time-stamp to that of the initial information. Explicitly considering the passing of time opens up the state of affairs to changes, which is to say that it inserts the concept of a dynamic environment as opposed to the static one assumed by belief revision.

4.5 Semantics Overview

In belief update, we consider a set of worlds, ordered $<_w$ of W, indexed by worlds [18]. This type of indexing facilitates the update of a set of worlds, performed independently. Another way of saying this is that for each world, there exists a total pre-order. In the event of new information μ, we must ensure that in our candidate worlds [8], the truth of μ is reflected. To do this, the minimal worlds in each world's total pre-order where μ holds are selected and altogether these represent the new belief base. The result of the operation depends on the ordering - semantically, this reflects the plurality of orderings of closeness that exist between candidate or possible worlds [8].

5 DEFEASIBLE REASONING

Defeasible reasoning is a form of non-monotonic reasoning. Ragni [27] differentiates strict and defeasible knowledge as the latter allowing for exceptions where the former does not, be it an explicit or implicit exception. Regarding defeasible reasoning, this is possible because the preconditions in a defeasible knowledge base are assumed to hold in the absence of explicit contradictory knowledge [17], which is to say that they are not explicitly true.

5.1 Notation and terminology

In defeasible logic, the monotonicity principle of classical logic is weakened [27] and used as one of the properties which all defeasible operations must satisfy. There are eight main properties, and these are formalised and propounded by Kraus, Lehmann and Magidor. [12]. The weakening of the principle of monotonicity is reflected in the notation used for defeasible logic.

The concepts of entailment and consequence adjust to accommodate conclusions being retractable: $\models$ becomes $\approx$ and $\vdash$ becomes $\mid\sim$, where $\alpha \mid\sim \beta$ reads α "typically $\rightarrow$ " β.

5.2 Properties

The following are the properties of Defeasible Reasoning as defined by Kraus, Lehmann and Magidor. [12].

1. Reflexivity (Ref)
K $\approx \alpha \mid\sim \alpha$
2. Right weakening (RW)
K $\approx \alpha \mid\sim \beta$ and $\alpha \models \gamma \therefore$ K $\approx \alpha \mid\sim \gamma$
This expresses that plausible consequences are closed under logical consequences.
3. Left logical equivalents (LLE)
K $\approx \alpha \mid\sim \gamma$, $\beta \models \alpha$ and $\alpha \models \beta \therefore$ K $\approx \beta \mid\sim \gamma$
This expresses that logically-equivalent propositions have the same consequences.
4. And
K $\approx \alpha \mid\sim \beta$ and K $\approx \alpha \mid\sim \gamma \therefore$ K $\approx \alpha \mid\sim \beta \wedge \gamma$
5. Or
K $\approx \alpha \mid\sim \gamma$ and K $\approx \beta \mid\sim \gamma \therefore$ K $\approx \alpha \vee \beta \mid\sim \gamma$
This is a classical principle that does not imply monotonocity, but rather expresses that a consequence of two different formulae ought to be a consequence of their disjunction.
6. Monotonicity
Classical propositional logic is monotonic, where monotonicity is as follows: K $\approx \alpha \mid\sim \beta \therefore$ K $\approx \alpha \wedge \gamma \mid\sim \beta$ Lehmann and Magidor [16] argue that defeasible entailment ($\approx$) ought to require weaker forms of monotonicity, namely Cautious Monotonicity and Rational Monotonicity.
Cautious Monotonicity (CM)

$$\text{K} \approx \alpha \mid\sim \beta \text{ and K} \approx \alpha \mid\sim \gamma \therefore \text{K} \approx \alpha \wedge \gamma \mid\sim \beta$$

Kraus, Lehmann and Magidor [12] explain that this property expresses that adding new information, the truth of which could have been plausibly concluded, should not conflict with prior conclusions.
Rational Monotonicity (RM)

$$\text{K} \approx \alpha \mid\sim \beta \text{ and K} \not\approx \alpha \mid\sim \neg\gamma \therefore \text{K} \approx \alpha \wedge \gamma \mid\sim \beta$$

5.3 Discussion

Within defeasible reasoning, there are two different stances on drawing conclusions in the case of exceptions: prototypical reasoning and presumptive reasoning [15]. In prototypical reasoning, given a typical situation, inheriting properties is fine. This type of reasoning is formalised by Lehmann and Magidor [15] using the rational closure defined in [16]. In presumptive reasoning, if there is no evidence indicating otherwise i.e. if not explicitly negated, inheriting properties is fine. This type of reasoning is the one intended by default logic and is formalised in [15].

Logic formalisations such as in [15, 16] focus on the form of the propositions over their meaning, in acknowledgement that the reader's knowledge of the world may influence the study of the formal properties [15]. Pelletier and Elio [24], in contrast, focus on the meanings over the forms. The latter approach does enhance readability, but loses the concise precision of the former and introduces potential variation in the impact on the reader. Regarding knowledge representation, Lehmann [15] explores defeasible reasoning from a different perspective to Reiter [28]. Reiter originally propounded to represent default information as normal defaults, defining a set of such to be extensions provided to a set of propositions. He later critiques himself in [29], and proposes semi-normal defaults as an extension of the class of defaults, as his approach in [28] led to undesirable consequences when normal default*s* were involved. Lehmann further explores possibilities of formalising the logic by considering normal defaults and assigning a meaning to sets of normal defaults such that the consideration of non-normal defaults are unnecessary because the interactions between defaults meet expectations [15].

5.4 Examples

An illustration of the need for defeasible reasoning is given in [12], by means of the following example: (i) Birds fly (b → f), (ii) penguins are birds (p → b) and (iii) penguins do not fly (p → ¬ f). Suppose we have a knowledge base K=(i),(ii). Receiving the additional information that penguins do not fly introduces an inconsistency, or rather an incoherence [2], in that if birds fly and penguins are birds but do not fly, penguins cannot exist (¬p). The solution proposed by [12] is to weaken the proposition that *birds fly* (b → f) to *typically, birds fly* (b |∼ f).

There are properties in addition to those given in Section 3.2, which are also desirable [15] and feature in several of the examples here presented. Four such properties that pertain to closure are of particular interest: the presumption of typicality, the presumption of independence, priority to typicality, and respect for specificity. Explanations and justifications of the above properties are given in [15].

Consider the following example: (i) if Alice is tired, typically, she will sleep ($t \mid\sim s$), and (ii) Alice is not a student (n). Rational monotonicity tells us that the closure of K will contain either $t \wedge n \mid\sim s$ or $t \mid\sim \neg n$. Applying the principle of the presumption of typicality, given there is no convincing reason that K should contain $t \mid\sim \neg n$, it should therefore contain $t \wedge n \mid\sim s$.

Suppose that knowledge base K consists of (i) if Alice is tired, typically, she will sleep ($t \mid\sim s$), and (ii) if Alice is tired, typically, she is a student ($t \mid\sim \neg n$). We cannot use the principle of the presumption of typicality to support including $t \wedge n \mid\sim s$ in the closure of K. The presumption of independence, however, expresses that if we lose typicality with respect to one consequent (e.g. n), we may still presume typicality with respect to another (e.g. s), unless there is reason to believe otherwise, as defined in [15]. In this scenario, we therefore accept $t \wedge n \mid\sim s$.

Reasons need not be beliefs. For example, perceptual states can be reasons [26]. Pollock [26] illustrates this point with an example from the perspective of the agent: sentence β that *X looks red to me* offers a reason for the belief ϕ that X is red. This is, however, a defeasible reason - for example, consider new information α delivered to the agent by a trustworthy person, that *X is not really red, but it appears so due to the lighting conditions*. The agent's reason to believe that X is red (ϕ) no longer justifies it to hold - α is consistent with β, but $\alpha \wedge \beta$ does not offer a reason to believe ϕ. In this example, β is what is called a *prima facie* reason and α is called a *defeater*. Pollock defines and explores different types of defeaters in [26], arguing that prima facie reasons and defeaters are the primary cause of the non-monotonic nature of human reasoning.

5.5 Semantics Overview

We recall that $\Vdash$ and $\models$ denote satisfaction and entailment respectively. In defeasible reasoning, we consider an ordered or ranked preferential [32] structure of interpretations. Ordering is by typicality or normality, where the most typical interpretation is minimal [18, 31] i.e. $\psi_1 < \psi_2$ reads as ψ_1 is more typical than ψ_2. Intuitively, this tells us that if $\alpha \mid\sim \beta$, then $\alpha \vdash \beta$ is in the minimal, i.e. lowest-ranked, interpretations. Ordering relation $<$ is constrained to be irreflexive and transitive [18] i.e. strict partial [12]. It also satisfies the *stoppering* or *smoothness* or *limit assumption* [18]: for proposition α and interpretation $\psi \in W$, where $\psi \models \alpha$, either ψ is minimal of the interpretations where $W \vDash \alpha$ or $\exists$ some $\psi' \in W$, such that $\psi' < \psi$, that is minimal in the sense of the former. In the event of a knowledge base K and new information received α, we transition from our minimal interpretations $m(K)$ to minimal, in the sense of $<$, interpretations that entail α [18]. The idea is to have the sets of interpretations transitioned between be as close in rank as possible.

6 DISCUSSION

The literature proposes systems that perform non-monotonic inferences [12], and evaluates these systems according to formalised characteristics noted as desirable [1, 9, 12]. Reference is made to non-monotonic systems modelling characteristics of the way in which humans reason in the face of incomplete information and uncertainty [25, 27]. Whether such formalisations match human reasoning is an open question.

7 CONCLUSIONS

Belief revision allows for the consideration that to the average human, one belief may have a different level of importance than another belief [25]. The process of the operation reflects that, in the face of new information, the agent may need to change their beliefs and accept the implications the new information may have for their beliefs. We have an interest in how non-classical forms of reasoning relate to human reasoning, given that human reasoning prompted the existence of research in this area. For this purpose, what is known about classical propositional reasoning and three forms of non-classical reasoning has been explored. Our focus has been on belief revision, with the aim of informing for later investigation into the extent to which it corresponds to cognitive reasoning. Defeasible reasoning and belief update were included for context, as related work.

Classical reasoning is deductive i.e. monotonic, whereas non-classical reasoning showcases a flexibility, characteristic of cognitive reasoning, which can be classified as non-monotonicity. For each of the non-classical forms of reasoning we considered, we assumed an underlying propositional logic. Conflicting information is interpreted, and thus resolved to preserve consistency, differently in each. In defeasible reasoning, conflicting information generally indicates exceptions - facilitated by means of weakening [27] notation, reflecting a more relaxed monotonicity principle than that prevalent in classical reasoning. This manifests through the notion of *typicality*, or, semantically speaking, it manifests in that if we are given that the most normal, or lowest-ranked, models of α are also models of β, we can say that $\alpha \mid\sim \beta$ and the other way around [12]. In this way, defeasible reasoning can allow exceptions. Both belief revision and belief update are considered to be forms of belief change. Belief change involves a belief base and a belief set, where explicit knowledge resides in the base and inferences or knowledge derived from that in the base resides in the belief set. In belief revision, conflicting information indicates flawed prior knowledge on the part of the agent, defeating, i.e. forcing the retraction of, conclusions drawn from it. Such information is referred to as a defeater, either of type *rebutting* or *undercutting* [26], and is taken into account by selecting the models of the new information closest to the models of the base [9]. In belief update, the notion of time enters the scenario, contrasting the static environment assumed in belief revision to a dynamic environment [9] assumed in belief update. This means that conflicting information is taken to indicate or correspond to a change in the world or real state of affairs. Such information is taken into account by taking each model ψ of the base, and selecting the set of models of the new information that is closest to ψ [9].

REFERENCES

[1] Carlos E. Alchourrón, Peter Gärdenfors, and David Makinson. 1985. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic* 50 (1985), 510–530. https://doi.org/10.2307/2274239

[2] Giovanni Casini, Eduardo Fermé, Thomas Meyer, and Ivan Varzinczak. 2018. A Semantic Perspective on Belief Change in a Preferential Non-Monotonic Framework. In *Proceedings of the 16th International Conference on Principles of Knowledge Representation and Reasoning (KR)*, M. Thielscher and F. Toni (Eds.). AAAI.

[3] Adnan Darwiche and Judea Pearl. 1997. On the Logic of Iterated Belief Revision. *Artificial Intelligence* 89 (01 1997), 1–29. https://doi.org/10.1016/S0004-3702(96)00038-0

[4] Alvaro del Val and Yoav Shoham. 1992. Deriving Properties of Belief Update from Theories of Action. In *Proceedings of the Tenth National Conference on Artificial Intelligence (AAAI'92)*. AAAI Press, San Jose, California, 584–589. http://dl.acm.org/citation.cfm?id=1867135.1867225

[5] Peter Gärdenfors. 1988. *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. MIT Press.

[6] Peter GÃď'rdenfors. 1992. *Belief Revision: An Introduction*. Cambridge University Press, 1–26. https://doi.org/10.1017/CBO9780511526664.001

[7] Sven Ove Hansson. 2017. Logic of Belief Revision. In *The Stanford Encyclopedia of Philosophy* (winter 2017 ed.), Edward N. Zalta (Ed.). Metaphysics Research Lab, Stanford University.

[8] Andreas Herzig and Omar Rifi. 1999. Propositional belief base update and minimal change. *Artificial Intelligence* 115, 1 (1999), 107 – 138. https://doi.org/10.1016/S0004-3702(99)00072-7

[9] Hirofumi Katsuno and Alberto O. Mendelzon. 1991. On the Difference Between Updating a Knowledge Base and Revising It. In *Proceedings of the Second International Conference on Principles of Knowledge Representation and Reasoning (KR'91)*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 387–394. http://dl.acm.org/citation.cfm?id=3087158.3087197

[10] Hirofumi Katsuno and Alberto O. Mendelzon. 1991. Propositional Knowledge Base Revision and Minimal Change. *Artif. Intell.* 52, 3 (Dec. 1991), 263–294. https://doi.org/10.1016/0004-3702(91)90069-V

[11] Arthur M. Keller and Marianne W. Wilkins. 1985. On the Use of an Extended Relational Model to Handle Changing Incomplete Information. *IEEE Transactions on Software Engineering* SE-11, 7 (July 1985), 620–633. https://doi.org/10.1109/TSE.1985.232506

[12] Sarit Kraus, Daniel Lehmann, and Menachem Magidor. 1990. Nonmonotonic reasoning, preferential models and cumulative logics. *Artif. Intell.* 44 (1990), 167–207.

[13] Jérôme Lang. 2007. Belief Update Revisited. In *Proceedings of the 20th International Joint Conference on Artifical Intelligence (IJCAI'07)*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2517–2522. http://dl.acm.org/citation.cfm?id=1625275.1625681

[14] Jérôme Lang, Fangzhen Lin, and Pierre Marquis. 2003. Causal Theories of Action: A Computational Core., In Reasoning About Actions and Change. *IJCAI International Joint Conference on Artificial Intelligence*, 1073–1078.

[15] Daniel Lehmann. 1995. Another perspective on default reasoning. *Annals of Mathematics and Artificial Intelligence* 15, 1 (1995), 61–82.

[16] Daniel Lehmann and Menachem Magidor. 1992. What does a conditional knowledge base entail? *Artif. Intell.* 55 (1992), 1–60.

[17] Luca Longo and Pierpaolo Dondio. 2014. Defeasible Reasoning and Argument-Based Systems in Medical Fields: An Informal Overview. In *2014 IEEE 27th International Symposium on Computer-Based Medical Systems*. IEEE, 376–381. https://doi.org/10.1109/CBMS.2014.126

[18] David Makinson. 1993. Five Faces of Minimality. *Studia Logica* 52 (09 1993), 339–380. https://doi.org/10.1007/BF01057652

[19] David Makinson. 2005. *Bridges from Classical to Nonmonotonic Logic*. King's College Publications, King's College London.

[20] David Makinson and Peter Gärdenfors. 1991. The relations between the logic of theory change and nonmonotonic logic. In *The Logic of Theory Change*, André Fuhrmann and Michael Morreau (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 183–205. https://doi.org/10.1007/BFb0018421

[21] João Martins and Stuart Shapiro. 1988. A Model for Belief Revision. *Artif. Intell.* 35 (01 1988), 25–79. https://doi.org/10.1016/0004-3702(88)90031-8

[22] Tim Miller and Christian Muise. 2016. Belief Update for Proper Epistemic Knowledge Bases. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence (IJCAI'16)*. AAAI Press, New York, New York, USA, 1209–1215. http://dl.acm.org/citation.cfm?id=3060621.3060789

[23] Francis Pelletier and Renee Elio. 1997. What Should Default Reasoning Be, by Default? *Computational Intelligence* 13 (05 1997), 165–187. https://doi.org/10.1111/0824-7935.00037

[24] Francis Pelletier and Renee Elio. 2005. The Case for Psychologism in Default and Inheritance Reasoning. *Synthese* 146 (08 2005), 7–35. https://doi.org/10.1007/s11229-005-9063-z

[25] Pavlos Peppas. 2008. Belief Revision. In *Handbook of Knowledge Representation*, F. van Harmelen, V. Lifschitz, and B. Porter (Eds.). Elsevier Science. https://doi.org/10.1016/S1574-6526(07)03008-8

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

[27] Marco Ragni, Christian Eichhorn, Tanja Bock, Gabriele Kern-Isberner, and Alice Ping Ping Tse. 2017. Formal Nonmonotonic Theories and Properties of Human Defeasible Reasoning. *Minds and Machines* 27 (02 2017). https://doi.org/10.1007/s11023-016-9414-1

[28] Raymond Reiter. 1980. A Logic for Default Reasoning. *Artif. Intell.* 13 (04 1980), 81–132. https://doi.org/10.1016/0004-3702(80)90014-4

[29] Raymond Reiter and Giovanni Criscuolo. 1983. Some representational issues in default reasoning. *Computers & Mathematics With Applications - COMPUT MATH APPL* 9 (12 1983), 15–27. https://doi.org/10.1016/0898-1221(83)90004-4

[30] Gavin Rens, Thomas Meyer, Gabriele Kern-Isberner, and Abhaya Nayak. 2018. Probabilistic Belief Revision via Similarity of Worlds Modulo Evidence. 343–356. https://doi.org/10.1007/978-3-030-00111-7_29

[31] Y. Shoham. 1987. Readings in Nonmonotonic Reasoning. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, Chapter A Semantical Approach to Nonmonotonic Logics, 227–250. http://dl.acm.org/citation.cfm?id=42641.42659

[32] Frans Voorbraak. 1993. Preference-based Semantics for Nonmonotonic Logics. In *Proceedings of the 13th International Joint Conference on Artifical Intelligence - Volume 1 (IJCAI'93)*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 584–589. http://dl.acm.org/citation.cfm?id=1624025.1624107