



Negation of graded beliefs

Bénédicte Legastelois, Marie-Jeanne Lesot, Adrien Revault d'Allonnes

► To cite this version:

Bénédicte Legastelois, Marie-Jeanne Lesot, Adrien Revault d'Allonnes. Negation of graded beliefs. Int. Conf. on Information Processing and Management of Uncertainty in Knowledge-Based Systems, IPMU 2016, Jun 2016, Eindhoven, Netherlands. hal-01340493

HAL Id: hal-01340493

<https://hal.sorbonne-universite.fr/hal-01340493>

Submitted on 5 Jun 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Negation of Graded Beliefs

Bénédicte Legastelois¹, Marie-Jeanne Lesot¹, and Adrien Revault d'Allonnes²

¹ Sorbonne Universités, UPMC Univ Paris 06,
CNRS, LIP6 UMR 7606, 4 place Jussieu 75005 Paris, France
{Benedicte.Legastelois,Marie-Jeanne.Lesot}@lip6.fr

² Université Paris 8 - EA 4383 - LIASD, FR-93526, Saint-Denis, France
Adrien.Revault-d_Allonnes@paris8.fr

Abstract. Negation is a key element in the construction of logical systems and plays a central role in reasoning and information manipulation tools. This paper considers the issue of negating graded beliefs, in the framework of a graded doxastic logic. It studies three interpretations of negation for these high level pieces of information, where negation is transferred to the three components of graded beliefs: the formula about which a belief is expressed, the belief modality and the belief level. The paper discusses the choice of appropriate formal frameworks for each of them, considering modal, fuzzy and many-valued logics; it characterises their use and underlines their relations, in particular regarding their effects on the belief degrees.

Keywords: belief reasoning, modal logic, weighted logic, doxastic logic

1 Introduction

One of the authors of this paper firmly believes unicorns exist, the other two do not. One does *not firmly* believe they exist, the other firmly does *not believe* they exist. Most people can tell the difference between these two points of view and can even tell them apart from someone who firmly believes unicorns do *not exist*. All three assertions are distinct negations of the original belief, which serves to show how tricky it is to negate a high level concept like graded beliefs.

The very notion of graded belief is a complex one, for which several interpretations can be considered: first, degrees of beliefs can be interpreted in terms of certainty. They can for instance be related to the -subjective- certainty an agent associates with the fact about which he expresses a belief [1]. Another type of uncertainty arises when an agent reasons about the beliefs of another agent, as revealed by the latter [2]: this uncertainty can be modelled in a possibilistic framework, as offered by Generalised Possibilistic Logic [8, 9] a graded version of Meta-Epistemic Logic [2].

This paper takes a different point of view: interpreting degrees of beliefs in terms of belief strength. Graded beliefs are used to express partial beliefs, i.e. beliefs that are more or less true, but all attached to the same level of certainty.

For this general type of graded beliefs, the paper offers a discussion of the negation issue, proposing to model the three interpretations of the negation given

in the introduction and, in particular, their cognitive underpinnings. It focuses on choosing appropriate formal paradigms to model all three, specifically looking at modal, many-valued and fuzzy logics: the former with its doxastic variant for belief representation, the latter two weighted logics to help represent and manipulate gradual concepts, with distinct truth degrees, beyond the classic binary case. In these high level logics, negation models notions which cannot be addressed in classical logics.

After setting a formal representation of the tackled issues in Section 2, the paper reviews the principles of negation in these logics in Section 3. The three interpretations of negation for graded beliefs are then discussed in turn in the following three sections. Conclusion and future works are presented in Section 7.

2 Formalising the Graded Belief Negation Issue

In order to formalise the issue of graded belief negation, this section first introduces the notation proposed to represent graded beliefs and then expresses the three interpretations of their negation outlined in the introduction.

2.1 Notations for Graded Beliefs

A graded belief extends the notion of belief by introducing a measure of the extent to which something is believed. We propose to denote this $B(\varphi, \alpha)$, read ‘ φ is believed to a degree α ’: this notation simultaneously represents (i) the considered formula, φ , about which a modal assertion is expressed; (ii) the type of non factual modality, belief, B ; (iii) the degree to which belief is partially held, α .

Using this notation, a sentence such as ‘I firmly believe that unicorns exist’ can be represented as $B(\varphi, \alpha)$ for which φ is the formula ‘unicorns exist’ and α represents the belief degree ‘firmly’, e.g. 0.8 if a numerical transposition were chosen, the choice of a correspondence between a numerical value and a linguistic term being out of the scope of this paper, 0.8 is given here as a mere example.

This notation highlights the relation between graded beliefs and the doxastic variant of modal logic [12], which aims at representing and reasoning about beliefs, as well as to weighted logics, such as many-valued or fuzzy logics, which manipulate degrees, respectively with ordinal and numerical values.

The negation of a graded belief is linguistically expressed as ‘*it does not hold that* φ is believed to a degree α ’. We formally denote this $\mathcal{N}(B(\varphi, \alpha))$ where \mathcal{N} is a general negation operator whose meaning is the topic of this paper: we propose a special notation to distinguish it from the classical negation, so as not to confuse it with $\neg B(\varphi, \alpha)$ and avoid implicitly transposing classical results known for \neg to the case of graded belief negation.

2.2 Objects of the Negation of Graded Beliefs

The challenge of negating graded beliefs comes from the fact that several interpretations can be considered, depending on which of their components is seen as

the negated object: formally, negation can be considered as transferred to each of them, leading to examine the following formulae, illustrated below:

global negation	transferred on	modality	formula	degree
$\mathcal{N}(B(\varphi, \alpha))$		$\mathcal{N}(B)(\varphi, \beta)$	$B(\mathcal{N}(\varphi), \gamma)$	$B(\varphi, \mathcal{N}(\alpha))$

The purpose of this work is to discuss them, questioning the choice of appropriate formal frameworks and definitions for the three types of negation operators \mathcal{N} . Note that, in the general notation above, different degrees, α , β , γ and $\mathcal{N}(\alpha)$, are used to avoid implicitly imposing *a priori* constraints. Establishing their relations to one another, including the possibility of their being equal, is an integral part of the problem of negation interpretation: the three object transpositions above show a simplified view of the problem. The three components of graded beliefs are actually intertwined and their combination and mutual influence matters.

First, note that, when considering the transfer to the believed formula φ , i.e. $B(\mathcal{N}(\varphi), \gamma)$, the general negation \mathcal{N} applies to a logical formula and is therefore naturally interpreted as the logical negation, i.e. $\mathcal{N} \equiv \neg$. The question is then whether the statement $\mathcal{N}(B(\varphi, \alpha))$, read ‘*it does not hold that φ is believed to a degree α* ’, allows to draw some conclusion about the belief in $\neg\varphi$, formally written $B(\neg\varphi, \gamma)$: if *it does not hold that* ‘I firmly believe that unicorns exist’, do I believe that unicorns do *not exist* and, if so, how much?

Second, transferring the negation to the belief degree means that ‘I believe φ ’ does hold, but to a degree other than α : using the running unicorn example, the question is to specify the meaning of the level for which it holds that ‘I *not firmly* believe that unicorns exist’. The general negation \mathcal{N} then applies to the belief degrees and is denoted N , some suitable negation operator for degrees, to be defined. This interpretation questions the relationship between the global negation $\mathcal{N}(B(\varphi, \alpha))$ and $B(\varphi, N(\alpha))$.

Third, the modality itself comes into question, considering that the statement ‘*it does not hold that φ is believed to a degree α* ’ gives some information about a ‘non-belief’ of φ , written $\mathcal{N}(B)(\varphi, \alpha)$, to indicate that the negation applies to the modality. Depending on how it is taken, this interpretation may lead to the introduction and manipulation of a second, opposite modality, on top of belief, something along the lines of disbelief.

3 Literature Review of Negation Principles

Negation is a central component of reasoning and information manipulation tools and constitutes an essential part of any logical system: its specification, together with, e.g., that of implication and the definition of an inference rule, suffices to define fully such a logical system. If, in classical logic, it establishes simple relations from true to false formulae, for systems with higher expressive power, it offers more complex, and richer, options and behaviours.

Since graded beliefs are related to both the formal frameworks of modal and weighted logics, the interpretation of their negation is related to the negation

they respectively define: after reviewing the reference case of classical logic, this section discusses the richness of negation meaning when applied to high level notions e.g. representing knowledge or beliefs in the case of modal logics, and to truth degrees for many-valued and fuzzy logics.

3.1 Negation in Classical Logic

In classical logic, the manipulation of negation relies on the laws of excluded middle (EM) and non-contradiction (NC) which respectively state that $\varphi \vee \neg\varphi$ is a tautology and that $\varphi \wedge \neg\varphi$ is a contradiction. (EM) imposes that any formula is either true or its negation is; (NC) imposes that they cannot be true simultaneously. Together they mean that exactly one among φ and $\neg\varphi$ is true.

This principle can also be expressed in terms of truth values: denoting by \mathcal{F} the set of all well-formed formulae and $d : \mathcal{F} \rightarrow \mathbb{B}$, the function that computes the truth value of any formula, the negation principle can be expressed as $d(\neg\varphi) = n(d(\varphi))$, where $n : \{0, 1\} \rightarrow \{0, 1\}$ is the function $n(x) = 1 - x$.

From an informational point of view, $\neg\varphi$ is usually taken as the opposite piece of information with respect to φ : if, e.g., φ is ‘unicorns exist’, $\neg\varphi$ usually means ‘unicorns do not exist’.

3.2 Negation in Modal Logic

Modal logics [7, 3] manipulate non factual pieces of information, such as knowledge and beliefs, through the modal operator \Box . Their combination with negation then raises the issue of the relations between the formulae $F_1 = \Box\varphi$, $F_2 = \neg\Box\varphi$, $F_3 = \Box\neg\varphi$ and, applying negation both before and after the modal operator, $F_4 = \neg\Box\neg\varphi = \Diamond\varphi$. The mere existence of these four variants underlines the richness of meaning negation can express in modal logics.

The negation behaviours expressed in classical logic still hold in modal logics. In particular, applying (EM) to the modal formula $\Box\varphi$ implies that $\Box\varphi \vee \neg\Box\varphi$ is a tautology, which establishes a relation between F_1 and F_2 , and similarly between F_3 and F_4 .

Relations between F_1 and F_4 are established by axioms (D) and (CD), reciprocal of each other, which state that:

$$(D) \quad \vdash \Box\varphi \rightarrow \neg\Box\neg\varphi \quad (CD) \quad \vdash \neg\Box\neg\varphi \rightarrow \Box\varphi \quad (1)$$

Since φ is any formula, taking it to be a negated one and using the double negation property $\neg\neg\varphi = \varphi$ makes it possible to establish relations between F_3 and F_2 , when similarly applying (D) and (CD) to $\Box\neg\varphi$.

Among the variants of modal logics, doxastic logic [12] interprets the modal operator as a belief operator, offering a formalism to represent and reason with beliefs. The doxastic reading of axiom (D) expresses that if φ is believed, then its contrary is not: it conveys the impossibility of believing both a formula and its contrary. (D) is thus considered as modelling one facet of the assumed rationality

of the agent whose beliefs are represented. It is known as the consistency axiom and is included in the usual axiomatic definition of doxastic logic, viz. $KD45$.

Its reciprocal, (CD) , conveys a similar, if complementary, consistency: in a doxastic reading, it would state that if a formula is not believed, then its contrary is. Its premise applies to an absence of belief and its conclusion to a belief. Rewriting the implication with negation and disjunction, it, therefore, imposes that an agent either believe a formula or its contrary, leading to a modal form of (EM) , stating that $\Box\neg\varphi \vee \Box\varphi$ is a tautology. This constraint, which requires he take a stance on his belief in φ , excludes modelling a neutral, perplexed or undetermined frame of mind, i.e. the absence of belief. It is considered as too restrictive and generally not included in the axiomatic definition of doxastic logic.

These axioms illustrate the fact that processing non factual pieces of information leads to discussions and allows the expression of negations with specific behaviours, as opposed to what classical logic offers.

3.3 Negation in Many-Valued Logic

Graded logics, including many-valued [18] and fuzzy [20] logics, respectively reviewed in this subsection and the next, consider factual pieces of information, like classical logic does, but extend the set of admissible truth values beyond the binary case $\{0, 1\}$. They thus manipulate truth degrees, with major consequences on the behaviour of the negation: they relax the law of non-contradiction, allowing both φ and $\neg\varphi$ to be partially true together. The negation principles are thus redefined to set how far a formula and its negation can simultaneously hold.

Formally, for many-valued logics [18] the set of admissible truth values is the totally ordered set $\mathcal{L}_M = \{\tau_0, \dots, \tau_{M-1}\}$, where M is a predefined positive integer and $\forall \alpha, \beta \in \{0, \dots, M-1\}, (\alpha \leq \beta \Leftrightarrow \tau_\alpha \leq \tau_\beta)$. \mathcal{L}_M represents intermediary truth values between ‘false’, written τ_0 , and ‘true’, τ_{M-1} , at various levels of granularity, depending on the total number of levels, M .

As in classical logic, the negation principle can be written $d(\neg\varphi) = n(d(\varphi))$. The truth negation function n is modified, compared to the classic negation (see Section 3.1), to process truth degrees and no longer just binary truth values. The usual definition is:

$$\begin{aligned} n : \mathcal{L}_M &\rightarrow \mathcal{L}_M \\ \tau_i &\mapsto \tau_{M-1-i} \end{aligned} \tag{2}$$

This extension preserves compatibility with classical logic: the negation of ‘false’, $n(\tau_0) = \tau_{M-1}$, is still ‘true’, and vice-versa. For intermediary truth degrees, this negation operator can be seen as computing the symmetrical value with respect to the scale’s middle value $\tau_{\frac{M-1}{2}}$. In the case where the number of truth levels M is odd, i.e. when this middle value indeed belongs to the scale, it is its own negation $n(\tau_{\frac{M-1}{2}}) = \tau_{\frac{M-1}{2}}$. This means that a formula whose truth degree is $\tau_{\frac{M-1}{2}}$ has the same truth degree as its negation. This specific behaviour of the negation, allowing a formula and its negation to be somewhat true together, can be interpreted in terms of a generalised law of non-contradiction.

Many-valued logic is further motivated by its capacity to provide linguistic representations of truth degrees, using a correspondence between the discrete ordered scale \mathcal{L}_M and a set of linguistic labels, e.g. based on adverbs qualifying the truth degree: for instance, \mathcal{L}_5 is in bijection with the set {'false', 'rather false', 'neither', 'rather true', 'true'}. Taken thus, the negation operation expressed in Eq. (2) can be interpreted in terms of the linguistic notion of antonymy [17].

3.4 Negation in Fuzzy Logic

Fuzzy logic [20] is an infinitely many-valued logic, for which truth values are not defined on a discrete scale but on the real interval $[0, 1]$. As detailed below, two levels account for the richness of its negation behaviour: fuzzy logic offers many negation operators on truth degree, as well as two negation types for predicates.

As in classical and many-valued logics, the negation principle can be expressed as $d(\neg\varphi) = n(d(\varphi))$, where the truth negation function n is adapted to manipulate values in $[0, 1]$. Such a function is called a *fuzzy negator* [19] and is defined by three properties: (i) domain and co-domains: n is a function $n : [0, 1] \rightarrow [0, 1]$; (ii) monotonicity: n is non-increasing; (iii) boundary conditions: $n(0) = 1$ and $n(1) = 0$.

Note that the general definition does not impose involutivity, which makes a major difference with the classic case when processing double negations: the truth degree of $\neg\neg\varphi$ can be different from φ 's.

However, the most usual definition of a fuzzy negator [20],

$$\begin{aligned} n : [0, 1] &\rightarrow [0, 1] \\ x &\mapsto 1 - x \end{aligned} \tag{3}$$

is both involutive and strict (i.e. strictly decreasing and continuous). Note that it is a straightforward generalisation of the negation operator defined for classic logic (see Section 3.1) from the domain $\{0, 1\}$ to $[0, 1]$; the preservation of the cases when the truth value is 0 or 1 are guaranteed by the boundary conditions. It also generalises the many-valued negation operator (see Eq. (2)) when the degrees τ_i are mapped to a discretisation of the interval $[0, 1]$, $\tau_i \rightarrow i/(M - 1)$.

Other fuzzy negators include (see e.g. [19]) $n(x) = \sqrt{1 - x^2}$ or $n(x) = 1 - x^2$ which is strict but not involutive. The Gödel operator, defined by $n(x) = 0$ if $x > 0$ and $n(0) = 1$, and its dual, defined by $n(x) = 1$ if $x < 1$ and $n(1) = 0$, are neither continuous nor involutive. These two examples are drastic choices: their use restricts the truth values to the Boolean $\{0, 1\}$. The variety of functions satisfying the general definition of fuzzy negators allows to define complex negation stances, richer than the behaviours the classic case can express.

Further, fuzzy logic also offers two types of negations for predicates defined as sub-intervals of a fuzzy partition over a numerical universe: complements and antonyms (see [16], for instance). Formally, for a predicate with membership function $A : [a^-, a^+] \rightarrow [0, 1]$, its complement, based on a fuzzy negator, and its antonym are respectively defined by the membership functions $\bar{A}(x) = n(A(x))$ and $\hat{A}(x) = A(a^+ - a^- - x)$. They usually differ one from another. Observe

that, if the predicate is a modality of a Ruspini partition, its antonym can be interpreted as the symmetrical modality with respect to the central one, establishing a relationship with the principle of the many-valued negation operator defined in Eq. (2). It is then also related to the already mentioned linguistic definition of antonymy [17]: fuzzy predicates thus allow, for instance, to represent and distinguish between the three notions ‘hot’, ‘not hot’ and ‘cold’.

4 From $\mathcal{N}(B(\varphi, \alpha))$ to $B(\neg\varphi, \beta)$: Negating Formulae

This paper addresses the issue of negation in the case of graded beliefs, at the crossroads of doxastic and weighted logics: this issue can thus be discussed as extensions of the manipulation rules reviewed in the previous section.

As described in Section 2, the general form of negation, written $\mathcal{N}(B(\varphi, \alpha))$, reads ‘*it does not hold that φ is believed to a degree α* ’. A first interpretation of this negation considers its transposition to the formula about which the belief is expressed, leading to examine the relation between $\mathcal{N}(B(\varphi, \alpha))$ and $B(\neg\varphi, \beta)$.

After discussing the general principles underlying this interpretation, this section reads it as a graded variant of the consistency axiom (*D*).

4.1 General Principles

Establishing a relation between $\mathcal{N}(B(\varphi, \alpha))$ and $B(\neg\varphi, \beta)$ raises the question of drawing conclusions about a belief in $\neg\varphi$ from the negation of a belief in φ to a degree α : in the running example, this interpretation aims at establishing a relation between the facts ‘*it does not hold that I firmly believe that unicorns exist*’ and ‘*I \sharp believe that unicorns *do not exist**’, where \sharp represents an appropriate modulating adverb, yet to be determined.

Let us underline that, even if this interpretation is expressed as a shifting of the negation to the formula, the associated belief degree, β , will likely also be impacted, allowing for a discussion: imposing it *a priori* to be equal to any value, in particular to α , would limit the expressiveness of the considered interpretation.

4.2 Graded Variant of the Consistency Axiom

We propose to study the transfer of the global negation to the formula in the framework of doxastic logic, respectively interpreting \mathcal{N} and \neg as an outer and an inner negation with respect to the belief operator: setting the degrees aside, for a moment, this introduces a correspondence between the formulae $\mathcal{N}(B(\varphi, \alpha))$ and $B(\neg\varphi, \beta)$, on the one hand, and $\neg\Box\varphi$ and $\Box\neg\varphi$, on the other.

This interpretation of \mathcal{N} requires that it satisfies the basic properties of a negation operator and, in particular, involutivity. From its very definition, this is indeed the case, as can be informally illustrated by its proposed reading: $\mathcal{N}(\mathcal{N}(B(\varphi, \alpha)))$ is read ‘*it does not hold that *it does not hold that φ is believed to a degree α** ’. It can be argued that this awkward expression is expected to be equivalent to ‘ φ is believed to a degree α ’.

In doxastic logic, a relation is established between formulae $\neg\Box\varphi$ and $\Box\neg\varphi$, denoted F_2 and F_3 in Section 3.2, as the contrapositive of axiom (D) , according to which $\vdash \Box\neg\varphi \rightarrow \neg\Box\varphi$. This implication is equivalently obtained when applying (D) to formula $\neg\varphi$.

Studying this negation for graded beliefs further leads to question a graded equivalent for axiom (D) . There exist several graded extensions of modal logics, mainly specified from a semantic point of view: they consider enriched Kripke frame definitions [6, 5, 14, 4] or introduce counting functions in the semantic definition of the modal operator [11, 10, 13, 15]. Using a relative counting approach to introduce a weighted modality \Box_α , we established in previous work [15] a graded variant of (D) stating that $\Box_\alpha\varphi \rightarrow \neg\Box_\beta\neg\varphi$ is a tautology, for any $\beta > 1 - \alpha$, on the same semantic frame hypotheses as (D) .

This weighted extension makes a relevant candidate for the desired relation between the global belief negation $\mathcal{N}(B(\varphi, \alpha))$ and the belief in the negated formula $B(\neg\varphi, \beta)$: applying (D_α) to $\neg\varphi$ and still considering \mathcal{N} as the outer and \neg as the inner negation w.r.t. the modal operator: $B(\neg\varphi, \alpha) \rightarrow \mathcal{N}(B(\varphi, \beta))$ for all $\alpha, \beta \in [0, 1]$ such that $\beta > 1 - \alpha$. This formula can be read ‘if $\neg\varphi$ is believed at degree α , then it does not hold that φ is believed at degree β ’. As a consequence, φ and $\neg\varphi$ can both be partially believed together, so long as their respective degrees satisfy the inequality constraint.

5 From $\mathcal{N}(B(\varphi, \alpha))$ to $B(\varphi, \mathcal{N}(\alpha))$: Negating Degrees

Considering now that the negation bears not on the formula but on the degree, a new set of possibilities arises. Indeed, interpreting ‘*it does not hold that φ is believed to a degree α* ’ as ‘ φ is believed to a degree $\mathcal{N}(\alpha)$ ’, where \mathcal{N} remains to be defined, offers new interpretations, both on the choice of operator and on the meaning these impose on the ensuing doxastic reading. At a fundamental level, this question depends on the interpretation of the degrees.

After detailing this interpretation, this section proposes to read the degrees as partial membership to a belief set, in the fuzzy set formalism, and studies the relevance of fuzzy negators, from a doxastic point of view.

5.1 General Principle

Note, first, that this section focuses on the case where the belief degrees are represented as numerical values, more specifically in the interval $[0, 1]$. A discretisation of this range could be considered, but it makes the interpretation harder to grasp. Indeed, it blurs the difference with the case where linguistic labels are used to express belief degrees but these hold no meaning when interpreted alone, apart from the modality. These will be considered when the negation is applied to the modality, as discussed in Section 6.

A literal interpretation of the negation of the degree considers a negative statement ‘*it does not hold that φ is believed to a degree α* ’ as meaning ‘ φ is believed to a degree which is *not* α ’, suggesting that any value other than α

is suitable: it may lead to define $N(\alpha)$ as a set and not a value, e.g. $N(\alpha) = [0, 1] \setminus \{\alpha\}$. This interpretation, although acceptable to the letter, leads to a highly imprecise and uninformative understanding of the statement: we do not examine it further and propose to establish N as a function from $[0, 1]$ to $[0, 1]$.

5.2 Belief Degree as Fuzzy Membership to a Belief Set

Considering the universe of all well-formed formulae \mathcal{F} , the set of gradual beliefs \mathcal{B} of an agent can be defined as a fuzzy subset of \mathcal{F} : each formula more or less belongs to \mathcal{B} and the belief degrees are interpreted as membership degrees.

The boundary case, where the degree is $B(\varphi, 1)$, that is, a formula φ representing a fully believed fact, believed without any restriction, is interpreted as totally belonging to the belief set \mathcal{B} . From a defuzzified point of view, it can be seen as equivalent to $\Box\varphi$ in doxastic logic. A formula with a lower membership degree to \mathcal{B} is believed less: defuzzifying the belief set through an α -cut, with $\alpha < 1$, can then be interpreted as enlarging the set of beliefs taken into account for reasoning, to extend beyond the full-fledged ones. The other boundary case appears to be somewhat harder to interpret. $B(\varphi, 0)$, means that φ does not belong to the belief set at all and represents a formula about which no belief holds. This interpretation raises the question of introducing an additional modality, to represent ‘disbelief’, and is discussed in the next section.

5.3 Relevance of Fuzzy Negators

Considering belief degrees as fuzzy membership degrees to a belief set suggests to use a fuzzy negator for the negation operator N on degrees, as reviewed in Section 3.4. This subsection examines their properties, regarding boundary conditions, monotonicity and involutivity, and their effect on a doxastic reading.

Boundary Conditions Using the running example, the relevance of fuzzy negator boundary conditions questions the relevance of equating ‘I believe to a degree $N(1)$ that unicorns exist’ to ‘I believe to a degree 0 that unicorns exist’ –and vice-versa– and therefore depends on the interpretation of these degrees. Considering $B(\varphi, 1)$ as equivalent to $\Box\varphi$, as suggested above, the boundary conditions can be interpreted as preserving the compatibility with the classic case of the modal expression of (EM), which states that $\Box\varphi \vee \neg\Box\varphi$ is a tautology.

Monotonicity From a doxastic point of view, the monotonicity of fuzzy negators implies comparing beliefs –or rather their levels– and their negations.

Let us extend the running example slightly and suppose we hold the following graded beliefs: B_1 : ‘I believe to a degree $N(0.8)$ that unicorns exist’ and B_2 : ‘I believe to a degree $N(0.6)$ that leprechauns exist’. Negating both these beliefs with a fuzzy, and therefore non-decreasing, negator entails that $\mathcal{N}(B_1)$ is believed, at most, as much as $\mathcal{N}(B_2)$.

Additionally, any increasing operator would contradict the scale's structure, e.g. creating situations where $N(\alpha) < N(1) = 0$. The different fuzzy negators, be they symmetric, drastic or if they offer some form of compromise, allow ways of modelling different belief stances, none of which contradict monotonicity.

Involutivity Finally, a property that not all fuzzy negators share but which is held by the most usual ones is involutivity. Choosing to believe $\mathcal{N}(\mathcal{N}(B_1))$ as much as B_1 can be understood as refusing to gain or lose any belief by *not not accepting 0.8* for a belief level. Even if the expression is somewhat farfetched, the underlying notion seems desirable.

6 From $\mathcal{N}(B(\varphi, \alpha))$ to $\mathcal{N}(B)(\varphi, \beta)$: Negating Modalities

The third and final interpretation of negated graded beliefs transfers negation to the modality, where the negative statement '*it does not hold that φ is believed to a degree α* ' is seen as providing information about a disbelief, therefore requiring an additional modality, besides belief. At a more intuitive, yet nuanced, level, in the case where belief degrees are linguistically expressed by adverbs, we propose to consider several modalities, intrinsically combining the modality with the level, e.g. distinguishing between the modalities 'weakly believe' and 'firmly believe'. This section considers these two cases in turn.

The difference with the transfer of negation to the formula should be underlined: the latter discusses how far one can believe a fact and its contrary simultaneously, with different degrees. This transfer considers a single fact and questions the links between believing and disbelieving it.

6.1 Belief and Disbelief

The introduction of a disbelief modality proposes to interpret '*it does not hold that φ is believed to a degree α* ' as ' *φ is *disbelieved* to a degree β* ', transferring the global negation to the modal operator, and simultaneously allowing an effect on the associated degrees, so as to avoid limiting the expressiveness of the considered negation interpretation *a priori*.

Note that the β coefficient describes a degree of disbelief, and not a belief degree. This induces a major change as opposed to the previously discussed interpretations, where degrees are measured on a single scale and all have the same meaning. As, nevertheless, the belief and disbelief modalities are related to one another, the associated degrees open the way to a signed scale rather than two independent ones. This interpretation was mentioned in Section 5.2 where understanding a 0-belief as utter disbelief was suggested. It then begs the question of at which level, i.e. around which α , does belief become disbelief.

The fuzzy interpretation discussed in the previous section can be considered as an extreme case, where the open interval $(0,1]$ represents various belief degrees and a single value, 0, is considered for disbelief. Such an interpretation can

Framework	Negation expression	Degree constraint
Graded modal	$B(\neg\varphi, \alpha) \rightarrow \mathcal{N}(B(\varphi, \beta))$	$\alpha, \beta \in [0, 1]$ and $\beta > 1 - \alpha$
Fuzzy	$\mathcal{N}(B(\varphi, \alpha)) = B(\varphi, N(\alpha))$	e.g. $N(\alpha) = 1 - \alpha$
Bimodal	$\mathcal{N}(B(\varphi, \alpha)) = Db(\varphi, f(\alpha))$	f isotone function
Multimodal	$\mathcal{N}(B_t(\varphi)) = B_{n(t)}(\varphi)$	$t \in \mathcal{L}_M, n(\tau_i) = \tau_{M-1-i}$

Table 1. Discussed interpretations of graded belief negation

be considered as bridging the gap between the interpretation of graded belief negation in terms of degree and in terms of modalities.

A natural value for changing from belief to disbelief considering a signed scale is probably 0.5, the scale's middle value: $(0.5, 1]$ then represent intermediate degrees of belief and $[0, 0.5)$ intermediate degrees of disbelief and 0.5 a neutral value. However, one could choose another value, to emphasise different attitudes with regards to belief, e.g. to have more positive values than negative ones, even if both sides are uncountable.

6.2 Finite Set of Belief Modalities

When considering the transfer of negation to the modality, a variant of special interest is the case where a belief level is considered as integrated in the modality, e.g. when ‘firmly believe’ and ‘weakly believe’ are considered to be two separate, yet comparable, modalities, instead of a single, continuous, belief modality nuanced with a degree: this interpretation is equivalent to considering a finite set of modalities with more than the two levels, belief and disbelief, discussed in the previous subsection. One can for instance consider five levels of belief {‘low’, ‘weak’, ‘moderate’, ‘strong’ and ‘high’} or more, with more linguistic labels.

This choice naturally leads to exploit a formalisation in many-valued logic, introducing several modalities formally defined as $B_t, t \in \mathcal{L}_M = \{\tau_0, \dots, \tau_{M-1}\}$. Negation then requires to define $\mathcal{N}(B_t), t \in \mathcal{L}_M$, which can be translated to $B_{n(t)}$ with the many-valued negation operator n (see Eq. (2)). The negation can then be considered as being applied to the adverb, which represents the degree-modality combination, and therefore interpreted in terms of linguistic antonyms.

7 Conclusion and Future Works

This paper proposed a discussion of the problem of negating graded beliefs. In the formal frameworks of modal, fuzzy and many-valued logics, it examined three main interpretations of this negation, depending on what it bears on: the formula about which a modal assertion is expressed; the modality, so as to express belief; or the degree, which represents the level of belief. Table 1 gathers the given interpretations and the associated frameworks. Although the objects of the negation have been considered separately, the results show them to be closely related to one another, in particular in their influence on the associated belief degree. This suggests that the core component of graded beliefs is the

level and, thus, their manipulation revolves around the construction of the rules governing its evolution.

Beyond negation, graded belief manipulation, in order to allow the combination of available beliefs into new ones, requires the definition of reasoning tools, in particular regarding their conjunction and disjunction. Future work will aim at extending the principles established for the negation operation to other connectives, leading to a general formal framework to reason about graded beliefs.

References

1. Bacchus, F., Grove, A.J., Halpern, J.Y., Koller, D.: From statistical knowledge bases to degrees of belief. *Artificial intelligence* 87(1), 75–143 (1996)
2. Banerjee, M., Dubois, D.: A simple modal logic for reasoning about revealed beliefs. In: *Proc. of the Int. Conf. on Symbolic and Quantitative Approaches to Reasoning with Uncertainty*, pp. 805–816. Springer (2009)
3. Blackburn, P., De Rijke, M., Venema, Y.: *Modal Logic*. Camb. Univ. Press (2001)
4. Bou, F., Esteve, F., Godo, L., Rodriguez, R.: Characterizing fuzzy modal semantics by fuzzy multimodal systems with crisp accessibility relations. In: *Proc. of the Joint 2009 IFSA World Congress and EUSFLAT Conference*. pp. 1541–1546 (2009)
5. Boutilier, C.: Modal logics for qualitative possibility theory. *International Journal of Approximate Reasoning* 10(2), 173–201 (1994)
6. Fariñas del Cerro, L., Herzig, A.: A modal analysis of possibility theory. In: *Fundamentals of Artificial Intelligence Research*. pp. 11–18. Springer (1991)
7. Chellas, B.F.: *Modal logic: an introduction*, vol. 316. Cambridge Univ Press (1980)
8. Dubois, D., Prade, H.: Generalized possibilistic logic. In: *Proc. of the 5th Int. Conf. on Scalable Uncertainty Management*. pp. 428–432. Springer (2011)
9. Dubois, D., Prade, H., Schockaert, S.: Rules and meta-rules in the framework of possibility theory and possibilistic logic. *Scientia Iranica* 18, 566–573 (2011)
10. Fattorosi-Barnaba, M., Cerrato, C.: Graded modalities I. *Studia Logica* 47, 99–110 (1988)
11. Fine, K.: In so many possible worlds. *Notre Dame Journal of Formal Logic* 13(4), 516–520 (1972)
12. Hintikka, J.: *Knowledge and belief: An introduction to the logic of the two notions*, vol. 181. Cornell University Press Ithaca (1962)
13. van der Hoek, W., Meyer, J.J.: Graded modalities in epistemic logic. In: *Proc. of the 2nd Int. Symposium on Logical Foundations of Computer Science, Tver’92*. LNCS, vol. 620, pp. 503–514. Springer (1992)
14. Laverny, N., Lang, J.: From knowledge-based programs to graded belief-based programs. Part I: On-line reasoning. In: *Proc. of ECAI*. IOS Press (2004)
15. Legastelois, B., Lesot, M.J., Revault d’Allonnes, A.: Typology of axioms for a weighted modal logic. *Proc. of the IJCAI Workshop on Weighted Logics for Artificial Intelligence WL4AI* pp. 40–48 (2015)
16. Moyse, G., Lesot, M.J., Bouchon-Meunier, B.: Oppositions in fuzzy linguistic summaries. In: *Proc of the IEEE Int. Conf. on Fuzzy Systems*. IEEE (2015)
17. Muehleisen, V.L.: *Antonymy and semantic range in english*. Ph.D. thesis, Northwestern University (1997)
18. Rescher, N.: *Many-valued logic*. Springer (1968)
19. Weber, S.: A general concept of fuzzy connectives, negations and implications based on t-norms and t-conorms. *Fuzzy Sets and Systems* 11(1), 103–113 (1983)
20. Zadeh, L.: Fuzzy sets. *Information and Control* 8, 338–353 (1965)