

Analyzing Differentiable Fuzzy Logic Operators

Emile van Krieken^{a,*}, Erman Acar^{a,b}, Frank van Harmelen^a

^a Vrije Universiteit Amsterdam, Netherlands

^b Civic AI Lab, Amsterdam, Netherlands

ARTICLE INFO

Article history:

Received 10 June 2020

Received in revised form 27 September 2021

Accepted 28 September 2021

Available online 7 October 2021

Keywords:

Fuzzy logic

Neural-symbolic AI

Learning with constraints

ABSTRACT

The AI community is increasingly putting its attention towards combining symbolic and neural approaches, as it is often argued that the strengths and weaknesses of these approaches are complementary. One recent trend in the literature is weakly supervised learning techniques that employ operators from fuzzy logics. In particular, these use prior background knowledge described in such logics to help the training of a neural network from unlabeled and noisy data. By interpreting logical symbols using neural networks, this background knowledge can be added to regular loss functions, hence making reasoning a part of learning.

We study, both formally and empirically, how a large collection of logical operators from the fuzzy logic literature behave in a differentiable learning setting. We find that many of these operators, including some of the most well-known, are highly unsuitable in this setting. A further finding concerns the treatment of implication in these fuzzy logics, and shows a strong imbalance between gradients driven by the antecedent and the consequent of the implication. Furthermore, we introduce a new family of fuzzy implications (called sigmoidal implications) to tackle this phenomenon. Finally, we empirically show that it is possible to use Differentiable Fuzzy Logics for semi-supervised learning, and compare how different operators behave in practice. We find that, to achieve the largest performance improvement over a supervised baseline, we have to resort to non-standard combinations of logical operators which perform well in learning, but no longer satisfy the usual logical laws.

© 2021 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

In recent years, integrating symbolic and statistical approaches to Artificial Intelligence (AI) gained considerable attention [22,6]. This research line has gained further traction due to recent influential critiques on purely statistical deep learning [47,59], which has been the focus of the AI community in the last decade. While deep learning has brought many important breakthroughs in computer vision [8], natural language processing [61] and reinforcement learning [68], the concern is that progress will be halted if its shortcomings are not dealt with. Among these is the massive amounts of data that deep learning models need to learn even a simple concept. In contrast, symbolic AI can easily reuse concepts and can express domain knowledge using only a single logical statement. Finally, it is much easier to integrate background knowledge using symbolic AI.

* Corresponding author.

E-mail addresses: e.van.krieken@vu.nl (E. van Krieken), erman.acar@vu.nl (E. Acar), Frank.van.Harmelen@vu.nl (F. van Harmelen).

However, symbolic AI has issues with scalability: dealing with large amounts of data while performing complex reasoning task, and is not able to deal with the noise and ambiguity of e.g. sensory data. The latter is related to the well-known *symbol grounding problem* which Harnad [30] defines as how “the semantic interpretation of a formal symbol system can be made intrinsic to the system, rather than just parasitic on the meanings in our heads”. In particular, symbols refer to concepts that have an intrinsic meaning to us humans, but computers manipulating these symbols cannot *understand* (or *ground*) this meaning. On the other hand, a properly trained deep learning model excels at modeling complex sensory data. These models could bridge the gap between symbolic systems and the real world. Therefore, several recent approaches [18,23,67,46,20] aim at interpreting symbols that are used in logic-based systems using deep learning models. These are among the first systems to implement “a hybrid nonsymbolic/symbolic system (...) in which the elementary symbols are grounded in (...) non-symbolic representations that pick out, from their proximal sensory projections, the distal object categories to which the elementary symbols refer.” Harnad [30].

1.1. Reasoning and learning using gradient descent

We introduce *Differentiable Fuzzy Logics* (DFL) which aims to integrate reasoning and learning by using logical formulas expressing background knowledge. The symbols in these formulas are interpreted using a deep learning model of which the parameters are to be learned. DFL constructs differentiable loss functions based on these formulas that can be minimized using gradient descent. This ensures that the deep learning model acts in a manner that is consistent with the background knowledge as we can backpropagate towards the parameters of the deep learning model.

To ensure loss functions are differentiable, DFL uses fuzzy logic semantics [41]. Predicates, functions and constants are interpreted using the deep learning model. By maximizing the degree of truth of the background knowledge using gradient descent, both learning and reasoning are performed in parallel. We can apply the loss functions constructed using DFL for more challenging machine learning tasks than purely supervised learning. These methods fall under the umbrella of weakly supervised learning [76]. For example, it can be used for semi-supervised learning [74,32] or to detect noisy or inaccurate supervision [19]. For such problems, DFL corrects the predictions of the deep learning model when it is logically inconsistent with the background knowledge.

To further our understanding of such losses, we present in this paper an analysis of the choice of operators used to compute the logical connectives in DFL. For example, functions called *t-norms* are used to connect two fuzzy propositions [41]. Because they return the degree of truth of the event that both propositions are true, such t-norms generalize the classical conjunction. Similarly, a fuzzy implication generalizes the classical implication. Most of these operators are differentiable, which enable their use in DFL. Interestingly, the derivatives of these operators determine how DFL corrects the deep learning model when its predictions are inconsistent with the background knowledge. We show that the qualitative properties of these derivatives are integral to both the theory and practice of DFL. We approach this problem both from the view of symbolic and of statistical approaches to AI, to bridge the conceptual gap between those views. This provides insights that otherwise would be overlooked.

1.2. Contributions

The main contribution of this article is to answer the following question: “*What fuzzy logic operators for existential quantification, universal quantification, conjunction, disjunction and implication have convenient theoretical properties when using them in gradient descent?*” We analyze both theoretically and empirically the effect of the choice of operators used to compute the logical connectives in Differentiable Fuzzy Logics on the learning behavior of a DFL system. To this end,

- We introduce Differentiable Fuzzy Logics (Section 4) which combines fuzzy logic and gradient-based learning, and analyze its behavior over different choices of fuzzy logic operators (Section 3).
- We analyze the theoretical properties of aggregation functions, which are used to compute the universal quantifier \forall and the existential quantifier \exists , t-norms and t-conorms which are used to compute the connectives \wedge and \vee , and fuzzy implications which are used to compute the connective \rightarrow .
- We introduce a new family of fuzzy implications called sigmoidal implications (Section 5) using the insights from these analyses.
- We perform experiments to compare fuzzy logic operators in a semi-supervised experiment (Section 9).
- We give several recommendations for choices of operators.

2. Differentiable Logics

Loss functions are real-valued functions that represent a cost and must be minimized. Differentiable Logics (DL) are logics for which differentiable loss functions are constructed that compute the truth value of given formulas using the semantics of the logic. These logics use background knowledge to deduce the truth value of statements in unlabeled or poorly labeled data, allowing us to use such data during learning, possibly together with normal labeled data. This can be beneficial as unlabeled, poorly labeled and partially labeled data is cheaper and easier to come by. This approach differs from Inductive Logic Programming [54] which derives formulas from data. DL instead informs what the data could have been.




Fig. 1. In this running example, we have an image with two objects on it, o_1 and o_2 .

We motivate the use of DL with the following classification scenario we consider throughout our analysis. Assume we have an agent A whose goal is to describe the scene on an image. It gets feedback from a supervisor S , who does not have an exact description of these images available. However, S does have a background knowledge base \mathcal{K} about the concepts contained on the images. The intuition behind Differentiable Logics is that S can correct A 's descriptions of scenes when they are not consistent with the knowledge base \mathcal{K} .

Example 1. We illustrate this idea with the following example. Agent A has to describe image I in Fig. 1 that contains objects o_1 and o_2 . A and the supervisor S consider the unary class predicates $\{\text{chair}, \text{cushion}, \text{armRest}\}$ and the binary predicate $\{\text{partOf}\}$. Since S does not have a description of I , it will have to correct A based on the knowledge base \mathcal{K} . A describes the image by using a confidence value in $[0, 1]$ for each observation. For instance, $p(\text{chair}(o_1))$ indicates the confidence A assigns to $\text{chair}(o_1)$, i.e., whether o_1 is a chair or not.

$$\begin{array}{ll}
 p(\text{chair}(o_1)) = 0.9 & p(\text{chair}(o_2)) = 0.4 \\
 p(\text{cushion}(o_1)) = 0.05 & p(\text{cushion}(o_2)) = 0.5 \\
 p(\text{armRest}(o_1)) = 0.05 & p(\text{armRest}(o_2)) = 0.1 \\
 p(\text{partOf}(o_1, o_1)) = 0.001 & p(\text{partOf}(o_2, o_2)) = 0.001 \\
 p(\text{partOf}(o_1, o_2)) = 0.01 & p(\text{partOf}(o_2, o_1)) = 0.95
 \end{array}$$

Suppose that \mathcal{K} contains the following logic formula which says parts of a chair are either cushions or armrests:

$$\forall x, y \text{ chair}(x) \wedge \text{partOf}(y, x) \rightarrow \text{cushion}(y) \vee \text{armRest}(y).$$

S might reason that since A is relatively confident of $\text{chair}(o_1)$ and $\text{partOf}(o_2, o_1)$ that the antecedent of this formula is satisfied, and either $\text{cushion}(o_2)$ or $\text{armRest}(o_2)$ has to hold. Since $p(\text{cushion}(o_2)) > p(\text{armRest}(o_2))$, a possible correction would be to tell A to increase its degree of belief in $\text{cushion}(o_2)$.

We would like to automate the kind of supervision S performs in the previous example. To this end, we study what we call as *Differentiable Fuzzy Logics* (DFL), a family of Differentiable Logics in the literature based on fuzzy logic. Here, the term ‘‘Differentiable Logics’’ refers to a logic together with a translation scheme from logical expressions to differentiable loss functions. Then, ‘‘Differentiable Fuzzy Logics’’ stands for the case where the logic is a fuzzy logic and the translation scheme applies to logical expressions which include fuzzy operators. Note that due to its numerical character, fuzzy logic is an obvious candidate for a differentiable logic. Therefore, in DFL, truth values of ground atoms are numbers in $[0, 1]$, and logical connectives are interpreted using some function over these truth values. Examples of logics in this family are Logic Tensor Networks [67], the similarly named Deep Fuzzy Logic [50], and the logics underlying Semantic Based Regularization [18], LYRICS [48] and KALE [26], which we compare in Section 11.1. This family stands orthogonal to the well-studied mathematical classification of the fuzzy logics landscape [12]. Instead, in our analysis, we use a variety of individual T-norms with different properties combined with a variety of aggregation functions.

3. Background

We assume basic familiarity on the syntax and the semantics of first-order logic. We shall denote predicates using the sans serif font (e.g., cushion), a set V of variables denoted by x, y, z, \dots or x_1, x_2, x_3, \dots , and a set O of domain objects denoted by o_1, o_2, \dots , and constants a, b, c, \dots . We limit ourselves to function-free formulas in *prenex normal form* that start with quantifiers followed by a quantifier-free subformula. An example of a formula in prenex form is $\forall x, y P(x, y) \wedge Q(x) \rightarrow R(y)$. An *atom* is $P(t_1, \dots, t_m)$ where t_1, \dots, t_m are terms. If t_1, \dots, t_m are all constants, we say it is a *ground atom*.

Table 1
The t-norms of interest.

Name	T-norm	Properties
Gödel (minimum)	$T_G(a, b) = \min(a, b)$	idempotent, continuous
Product	$T_P(a, b) = a \cdot b$	strict
Łukasiewicz	$T_{LK}(a, b) = \max(a + b - 1, 0)$	continuous
Drastic product	$T_D(a, b) = \begin{cases} \min(a, b), & \text{if } a = 1 \text{ or } b = 1 \\ 0, & \text{otherwise} \end{cases}$	
Nilpotent minimum	$T_{nM}(a, b) = \begin{cases} 0, & \text{if } a + b \leq 1 \\ \min(a, b), & \text{otherwise} \end{cases}$	left-continuous
Yager	$T_Y(a, b) = \max(1 - ((1 - a)^p + (1 - b)^p)^{\frac{1}{p}}, 0), p \geq 1$	continuous

Table 2
The t-conorms of interest.

Name	T-conorm	Properties
Gödel (maximum)	$S_G(a, b) = \max(a, b)$	idempotent, continuous
Product (probabilistic sum)	$S_P(a, b) = a + b - a \cdot b$	strict
Łukasiewicz	$S_{LK}(a, b) = \min(a + b, 1)$	continuous
Drastic sum	$S_D(a, b) = \begin{cases} \max(a, b), & \text{if } a = 0 \text{ or } b = 0 \\ 1, & \text{otherwise} \end{cases}$	
Nilpotent maximum	$S_{nM}(a, b) = \begin{cases} 1, & \text{if } a + b \geq 1 \\ \max(a, b), & \text{otherwise} \end{cases}$	right-continuous
Yager	$S_Y(a, b) = \min((a^p + b^p)^{\frac{1}{p}}, 1), p \geq 1$	continuous

Table 3
Some common aggregation operators.

Name	Generalizes	Aggregation operator
Minimum	T_G	$A_{T_G}(x_1, \dots, x_n) = \min(x_1, \dots, x_n)$
Product	T_P	$A_{T_P}(x_1, \dots, x_n) = \prod_{i=1}^n x_i$
Łukasiewicz	T_{LK}	$A_{T_{LK}}(x_1, \dots, x_n) = \max(\sum_{i=1}^n x_i - (n - 1), 0)$
Maximum	S_G	$E_{S_G}(x_1, \dots, x_n) = \max(x_1, \dots, x_n)$
Probabilistic sum	S_G	$E_{S_P}(x_1, \dots, x_n) = 1 - \prod_{i=1}^n (1 - x_i)$
Bounded sum	S_{LK}	$E_{S_{LK}}(x_1, \dots, x_n) = \min(\sum_{i=1}^n x_i, 1)$

Fuzzy logic is a many-valued logic where truth values are real numbers in $[0, 1]$ where 0 denotes completely false and 1 denotes completely true. It is often used to model reasoning in the presence of *vagueness* i.e., without sharp boundaries or to imprecisely classify concepts such as a *tall person* or a *small number* [28,56]. We will look at predicate fuzzy logics in particular, which extend propositional fuzzy logics with universal and existential quantification. For a brief background on fuzzy logic operators, see Appendix A, and for an extensive treatment on mathematical fuzzy logic we refer the reader to standard textbooks which include [28], [56] and [12].

In this paper we exclusively use the *strong negation* $N(a) = 1 - a$. Conjunctions $a \wedge b$ are generalized using *T-norms*, which are any function $T : [0, 1]^2 \rightarrow [0, 1]$ that is commutative, associative, monotonic, and has $T(1, a) = a$. In fuzzy logic, conjunctions using the t-norm are often denoted on the logical formula level as $a \otimes b$. However, on the semantic level (i.e. $T(a, b)$), a and b will refer to the truth values of the corresponding logical formulas a and b . An overview of the most common T-norms is given in Table 1 alongside common properties that are defined in Appendix A.3. *T-conorms* generalize disjunctions and are often denoted $a \oplus b$ in fuzzy logic literature. They are formed from a t-norm T using $S(a, b) = 1 - T(1 - a, 1 - b)$, following DeMorgans law $a \oplus b = \neg(-a \otimes -b)$. The most common t-conorms are given in Table 2.

Universal and existential quantification are generalized using increasing *aggregation operators* $A, E \in \bigcup_{n \in \mathbb{N}} [0, 1]^n \rightarrow [0, 1]$, where $A(1, \dots, 1) = E(1, \dots, 1) = 1$ and $A(0, \dots, 0) = E(0, \dots, 0) = 0$ [10]. Universal aggregators can be constructed from a t-norm T and existential aggregators from a t-conorm S :

$$A_T() = 1 \qquad E_S() = 0 \tag{1}$$

$$A_T(x_1, \dots, x_n) = T(x_1, A_T(x_2, \dots, x_n)) \qquad E_S(x_1, \dots, x_n) = S(x_1, E_S(x_2, \dots, x_n)) \tag{2}$$

An overview of common aggregation operators is given in Table 3. A formal introduction is in Appendix A.4.

Finally, implications are generalized using *fuzzy implications* [38], which are functions $I : [0, 1]^2 \rightarrow [0, 1]$ that are increasing with respect to the first argument, and decreasing with respect to the second. Furthermore, we assume boundary conditions $I(0, 0) = I(1, 1) = 1$, and $I(1, 0) = 0$, which follows from the classical implication. Fuzzy implications can have additional properties that we discuss in Appendix A.5. We will discuss two classes of fuzzy implications. The first is *S-implications* (Appendix A.5.1, examples in Table 4), which are formed using a t-conorm by generalizing the material impli-

Table 4
S-implications formed from $\neg a \oplus c$ with the common t-conorms from Table 2.

Name	T-conorm	S-implication	Properties
Gödel (Kleene-Dienes)	S_G	$I_{KD}(a, c) = \max(1 - a, c)$	All but IP
Product (Reichenbach)	S_P	$I_{RC}(a, c) = 1 - a + a \cdot c$	All but IP
Łukasiewicz	S_{LK}	$I_{LK}(a, c) = \min(1 - a + c, 1)$	All
Dubouis-Prade	S_D	$I_{DP}(a, c) = \begin{cases} c, & \text{if } a = 1 \\ 1 - a, & \text{if } c = 0 \\ 1, & \text{otherwise} \end{cases}$	All
Nilpotent (Fodor)	S_{Nm}	$I_{FD}(a, c) = \begin{cases} 1, & \text{if } a \leq c \\ \max(1 - a, c), & \text{otherwise} \end{cases}$	All

Table 5
The R-implications constructed using the t-norms from Table 1.

Name	T-norm	R-implication	Properties
Gödel	T_G	$I_G(a, c) = \begin{cases} 1, & \text{if } a \leq c \\ c, & \text{otherwise} \end{cases}$	LN, EP, IP
product (Goguen)	T_P	$I_{GG}(a, c) = \begin{cases} 1, & \text{if } a \leq c \\ \frac{c}{a}, & \text{otherwise} \end{cases}$	LN, EP, IP
Łukasiewicz	T_{LK}	$I_{LK}(a, c) = \min(1 - a + c, 1)$	All
Weber	T_D	$I_{WB}(a, c) = \begin{cases} 1, & \text{if } a < 1 \\ c, & \text{otherwise} \end{cases}$	LN, EP, IP
Nilpotent (Fodor)	T_{Nm}	$I_{FD}(a, c) = \begin{cases} 1, & \text{if } a \leq c \\ \max(1 - a, c), & \text{otherwise} \end{cases}$	All

ation $a \rightarrow c = \neg a \oplus c$ using $I_S(a, c) = S(N(a), c)$. The second is *R-implications*, which is the standard choice in t-norm fuzzy logics (Appendix A.5.2, examples in Table 5). These are constructed from t-norms using $I_T(a, c) = \sup_{b \in [0,1]} T(a, b) \leq c$.

4. Differentiable Fuzzy Logics

As mentioned earlier, *Differentiable Fuzzy Logics* (DFL) are Differentiable Logics based on fuzzy logic. Truth values of ground atoms are continuous, and logical connectives are interpreted using differentiable fuzzy operators. In principle, DFL can handle both predicates and functions. To ease the discussion, we will not analyze functions and constants and leave them out of the discussion.¹

4.1. Semantics

DFL defines a new semantics using vector embeddings and functions on such vectors in place of classical semantics. In classical logic, a *structure* consists of a domain of discourse and an interpretation function, and is used to give meaning to the predicates. DFL defines *structures* using *embedded interpretations*² instead:

Definition 1. A *Differentiable Fuzzy Logics structure* is a tuple $S = \langle O, \eta, \theta \rangle$, where O is a finite but unbounded set called *domain of discourse* and every $o \in O$ is a d -dimensional³ vector, $\eta : \mathcal{P} \times \mathbb{R}^W \rightarrow (O^m \rightarrow [0, 1])$ is an (*embedded*) *interpretation*, and $\theta \in \mathbb{R}^W$ are *parameters*. η maps predicate symbols $P \in \mathcal{P}$ with arity m to a function of m objects to a truth value $[0, 1]$. That is, $\eta(P, \theta) : O^m \rightarrow [0, 1]$. We will use the notation $\eta_\theta(P)$ to denote $\eta(P, \theta)$.

To address the *symbol grounding problem* [30], objects in the domain of discourse are d -dimensional vectors of reals. Their semantics come from the underlying semantics of the vector space as terms are interpreted in a real (valued) world [67]. Predicates are interpreted as functions mapping these vectors to a fuzzy truth value. Embedded interpretations can be implemented using neural network models⁴ with trainable network parameters θ . Note that different values for the parameters θ will produce different DFL structures. Next, we define the truth values of formulas in DFL.

¹ Functions and constants are modeled in [67] and [48].

² Serafini and Garcez [67] uses the term “(semantic) grounding” or “symbol grounding” [51] instead of ‘embedded interpretation’, “to emphasize the fact that \mathcal{L} is interpreted in a ‘real world’” but we find this potentially confusing as this could also refer to groundings in Herbrand semantics. Furthermore, by using the word ‘interpretation’ we highlight the parallel with classical logical interpretations.

³ Without loss of generality we fix the dimensionality of the vectors representing the objects. Extensions to a varying number of dimensions are straightforward by introducing types, such as done in [2].

⁴ We use ‘models’ to refer to deep learning models like neural networks, and not to models from model theory.

Definition 2. Let $\mathcal{S} = \langle O, \eta, \theta \rangle$ be a DFL structure, N a fuzzy negation, T a t-norm, S a t-conorm, I a fuzzy implication, and A and E universal and existential aggregators respectively. Furthermore, let $\mu : V \rightarrow O$ be a *variable assignment*, where we use $\mu[x : o]$ to refer to the new assignment where x is mapped to domain object o , that is, $\mu[x : o](x) = o$ and $\mu[x : o](x') = \mu(x')$ for $x \neq x'$. Then we say that \mathcal{S} satisfies the formula $\varphi \in \mathcal{L}$ w.r.t. μ (i.e., $\mathcal{S}, \mu \models \varphi$) in the degree of $e_\theta(\varphi)$ (i.e., the truth value of φ) where $e_\theta : (V \rightarrow O) \times \mathcal{L} \rightarrow [0, 1]$ is the *valuation function* defined inductively on the structure of φ as follows:

$$e_\theta(\mu, P(x_1, \dots, x_m)) = \eta_\theta(P)(\mu(x_1), \dots, \mu(x_m)) \quad (3)$$

$$e_\theta(\mu, \neg\phi) = N(e_\theta(\mu, \phi)) \quad (4)$$

$$e_\theta(\mu, \phi \otimes \psi) = T(e_\theta(\mu, \phi), e_\theta(\mu, \psi)) \quad (5)$$

$$e_\theta(\mu, \phi \oplus \psi) = S(e_\theta(\mu, \phi), e_\theta(\mu, \psi)) \quad (6)$$

$$e_\theta(\mu, \phi \rightarrow \psi) = I(e_\theta(\mu, \phi), e_\theta(\mu, \psi)) \quad (7)$$

$$e_\theta(\mu, \forall x \phi) = \bigwedge_{o \in O} e_\theta(\mu[x : o], \phi) \quad (8)$$

$$e_\theta(\mu, \exists x \phi) = \bigvee_{o \in O} e_\theta(\mu[x : o], \phi) \quad (9)$$

Equation (3) defines the fuzzy truth value of an atomic formula. μ assigns objects to the terms x_1, \dots, x_m resulting in a list of d -dimensional vectors. These are the inputs to the interpretation η_θ of the predicate symbol P (i.e., $\eta_\theta(P)$) to get a fuzzy truth value. Equations (4) - (7) define the truth values of the connectives using the operators N, T, S and I . Equations (8) and (9) define the truth value of universally quantified formulas $\forall x \phi$ and existentially quantified formulas $\exists x \phi$. This is done by enumerating the domain of discourse $o \in O$, evaluating the truth value of ϕ with o assigned to x in μ , and combining the truth values using aggregation operators A and E .

Note that our assumption on the finiteness of the domain is pragmatic: It reflects the finiteness of the data in machine learning settings. Hence, many fundamental results in the realm of mathematical (fuzzy) logic will not hold in general for the logic we defined [12].

4.2. Relaxing quantifiers

For infinite domains, or for domains that are so large that we cannot compute the full semantics of the \forall and \exists quantifiers, we can choose to sample a batch b of objects from O to approximate the computation of the valuation. This can be done by replacing Equation (8) with

$$e_\theta(\mu, \forall x \phi) = \bigwedge_{i=1}^b e_\theta(\mu[x : o_i], \phi), \quad o_1, \dots, o_b \text{ chosen from } O. \quad (10)$$

Choosing the batch of objects can be done in several ways. One approach would be to sample from a real-world distribution over the domain of discourse O , if available. For example, the domain of discourse might be the natural images, and the real-world distribution would be the distribution over natural images. A more common approach is to assume access to a dataset \mathcal{D} of independent samples from such a distribution [25](p.109) and to choose minibatches from this dataset. Note that by relaxing quantifiers using sampling we lose the soundness of our computation, as different batches will have different truth values for the formulas.

4.3. Learning using fuzzy maximum satisfiability

In DFL, *fuzzy maximum satisfiability* [19] is the problem of finding parameters θ that maximize the valuation of the knowledge base \mathcal{K} .

Definition 3. Let \mathcal{K} be a knowledge base of formulas, \mathcal{S} a DFL structure for \mathcal{K} and e_θ a valuation function. Then the *Differentiable Fuzzy Logics loss* \mathcal{L}_{DFL} of a knowledge base of formulas \mathcal{K} is computed using

$$\mathcal{L}_{DFL}(\mathcal{S}, \mathcal{K}) = \mathcal{L}_{DFL}(\langle \mathcal{O}, \eta, \theta \rangle, \mathcal{K}) = - \sum_{\varphi \in \mathcal{K}} e_\theta(\{\}, \varphi). \quad (11)$$

The *fuzzy maximum satisfiability problem* is the problem of finding parameters θ^* that minimize Equation (11):

$$\theta^* = \operatorname{argmin}_\theta \mathcal{L}_{DFL}(\mathcal{S}, \mathcal{K}). \quad (12)$$

This optimization problem can be solved using a gradient descent method. If the operators N, T, S, I, A and E are all differentiable, we can repeatedly apply the chain rule, i.e. reverse-mode differentiation, on the DFL loss $\mathcal{L}_{DFL}(\mathcal{S}, \mathcal{K})$. This procedure finds the derivative with respect to the truth values of the ground atoms $\frac{\partial \mathcal{L}_{DFL}(\mathcal{S}, \mathcal{K})}{\partial \eta_{\theta}(\mathcal{P})(o_1, \dots, o_m)}$. We can use these partial derivatives to update the parameter θ_n at iteration n of the optimization process again using the chain rule, resulting in a different embedded interpretation $\eta_{\theta_{n+1}}$. This procedure is computed as follows for the i th parameter:

$$\theta_{n+1,i} = \theta_{n,i} - \epsilon \cdot \frac{\partial \mathcal{L}_{DFL}(\mathcal{S}_n, \mathcal{K})}{\partial \theta_{n,i}} = \theta_{n,i} - \epsilon \cdot \sum_{\mathcal{P}(o_1, \dots, o_m)} \frac{\partial \mathcal{L}_{DFL}(\mathcal{S}_n, \mathcal{K})}{\partial \eta_{\theta}(\mathcal{P})(o_1, \dots, o_m)} \cdot \frac{\partial \eta_{\theta}(\mathcal{P})(o_1, \dots, o_m)}{\partial \theta_{n,i}}, \quad (13)$$

where ϵ is the learning rate. Note that the parameters θ_n are implicitly passed to \mathcal{L} through the structure $\mathcal{S}_n = \langle O, \eta, \theta_n \rangle$. We refer for implementation details to Appendix B.

Example 2. To illustrate the computation of the valuation function e_{θ} , we return to the problem in Example 1. The *domain of discourse* is the set of objects on natural images. We have access to a dataset of two objects $\mathcal{D} = \{o_1, o_2\}$. The valuation of the formula $\varphi = \forall x, y \text{ chair}(x) \otimes \text{partOf}(y, x) \rightarrow \text{cushion}(y) \oplus \text{armRest}(y)$ is

$$\begin{aligned} e_{\theta}(\mu, \varphi) = & A(A(I(T(\eta_{\theta}(\text{chair})(o_1), \eta_{\theta}(\text{partOf})(o_1, o_1)), S(\eta_{\theta}(\text{cushion})(o_1), \eta_{\theta}(\text{armRest})(o_1))), \\ & I(T(\eta_{\theta}(\text{chair})(o_1), \eta_{\theta}(\text{partOf})(o_2, o_1)), S(\eta_{\theta}(\text{cushion})(o_2), \eta_{\theta}(\text{armRest})(o_2))), \\ & A(I(T(\eta_{\theta}(\text{chair})(o_2), \eta_{\theta}(\text{partOf})(o_1, o_2)), S(\eta_{\theta}(\text{cushion})(o_1), \eta_{\theta}(\text{armRest})(o_1))), \\ & I(T(\eta_{\theta}(\text{chair})(o_2), \eta_{\theta}(\text{partOf})(o_2, o_2)), S(\eta_{\theta}(\text{cushion})(o_2), \eta_{\theta}(\text{armRest})(o_2))))). \end{aligned}$$

Next, we choose the operators as $T = T_P$, $S = S_P$, $A = A_{T_P}$ and $I = I_{RC}$, such that

$$e_{\theta}(\mu, \varphi) = \prod_{x, y \in \mathcal{C}} 1 - \eta_{\theta}(\text{chair})(x) \cdot \eta_{\theta}(\text{partOf})(y, x) \cdot (1 - \eta_{\theta}(\text{cushion})(y))(1 - \eta_{\theta}(\text{armRest})(y)).$$

If we interpret the predicate functions using the confidence values from Example 1 so that $\eta_{\theta}(\mathcal{P}(x)) = p(\mathcal{P}(x))$, we find that $e_{\theta}(\varphi) = 0.612$. Taking $\mathcal{K} = \{\varphi\}$, we find using the chain rule that

$$\begin{array}{ll} \frac{\partial \mathcal{L}_{DFL}(\mathcal{S}, \mathcal{K})}{\partial \eta_{\theta}(\text{chair})(o_1)} = -0.4261 & \frac{\partial \mathcal{L}_{DFL}(\mathcal{S}, \mathcal{K})}{\partial \eta_{\theta}(\text{chair})(o_2)} = -0.0058 \\ \frac{\partial \mathcal{L}_{DFL}(\mathcal{S}, \mathcal{K})}{\partial \eta_{\theta}(\text{cushion})(o_1)} = 0.0029 & \frac{\partial \mathcal{L}_{DFL}(\mathcal{S}, \mathcal{K})}{\partial \eta_{\theta}(\text{cushion})(o_2)} = 0.7662 \\ \frac{\partial \mathcal{L}_{DFL}(\mathcal{S}, \mathcal{K})}{\partial \eta_{\theta}(\text{armRest})(o_1)} = 0.0029 & \frac{\partial \mathcal{L}_{DFL}(\mathcal{S}, \mathcal{K})}{\partial \eta_{\theta}(\text{armRest})(o_2)} = 0.4257 \\ \frac{\partial \mathcal{L}_{DFL}(\mathcal{S}, \mathcal{K})}{\partial \eta_{\theta}(\text{partOf})(o_1, o_1)} = -0.4978 & \frac{\partial \mathcal{L}_{DFL}(\mathcal{S}, \mathcal{K})}{\partial \eta_{\theta}(\text{partOf})(o_2, o_2)} = -0.1103 \\ \frac{\partial \mathcal{L}_{DFL}(\mathcal{S}, \mathcal{K})}{\partial \eta_{\theta}(\text{partOf})(o_1, o_2)} = -0.2219 & \frac{\partial \mathcal{L}_{DFL}(\mathcal{S}, \mathcal{K})}{\partial \eta_{\theta}(\text{partOf})(o_2, o_1)} = -0.4031. \end{array}$$

We can now do a gradient update step to update the confidence values from Example 1, or find what the partial derivative of the parameters θ of some deep learning model p_{θ} should be using Equation (13).

One particularly interesting property of Differentiable Fuzzy Logics is that the partial derivatives of the subformulas with respect to the satisfaction of the knowledge base have a somewhat explainable meaning. For example, as hypothesized in Example 1, the computed partial derivatives reflect whether we should increase $p(\text{cushion}(o_2))$, as it is indeed the (**absolute**) largest partial derivative.

5. Derivatives of operators

We will now show that the choice of operators that are used for the logical connectives actually determines the inferences that are done when using DFL. If we used a different set of operators in Example 2, we would have gotten very different derivatives. These could in some cases make more sense, and in some other cases less. Furthermore, it is much easier to find a global minimum of the fuzzy maximum satisfiability problem (Equation (12)) for some operators than for others. This is often because of the smoothness of the operators. In this section, we analyze a wide variety of functions that can be used for logical reasoning and present some of their properties that determine how useful they are in inferences such as those illustrated above.

We will not discuss any varieties of fuzzy negations since the strong negation $N_C(a) = 1 - a$ is already continuous, intuitive and has simple derivatives.

Definition 4. A function $f : \mathbb{R} \rightarrow \mathbb{R}$ is said to be *vanishing* if there are $a, b \in \mathbb{R}$, $a < b$ such that for all $c \in (a, b)$, $f(c) = 0$, i.e. there is an interval for which the function is 0. Otherwise, the function is *nonvanishing*.

A function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ has a *vanishing derivative* if for all $a_1, \dots, a_n \in \mathbb{R}$ there is some $1 \leq i \leq n$ such that $\frac{\partial f(a_1, \dots, a_n)}{\partial a_i}$ is vanishing.

Whenever the derivative of an operator vanishes, it loses its learning signal. This definition does not include functions that only pass through 0, such as when using the product t-conorm for $a \oplus -a$, where we find that the derivative of $S_P(a, 1 - a)$ is 0 only at $\frac{1}{2}$. Furthermore, all the partial derivatives of the connectives used in the backward pass from the valuation function to the ground atoms have to be multiplied. If the partial derivatives are less than 1, their product will also approach 0. This can happen for instance with a large sequence of conjunctions using the product t-norm.

The drastic product T_D and operators derived from it such as the drastic sum S_D and the Dubois-Prade and Weber implications (I_{DP} and I_{WB}) have vanishing derivatives almost everywhere. The output confidence values of deep learning models are the result of transformations on real numbers using functions like the sigmoid or softmax that result in truth values in $(0, 1)$. The operators derived from T_D only have nonvanishing derivatives when their inputs are exactly 0 or 1, invalidating their use in this application

Definition 5. A function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is said to be *single-passing* if it has nonzero derivatives on at most one input argument. That is, for all $x_1, \dots, x_n \in [0, 1]$ it holds that $\left| \left\{ i \mid \frac{\partial f(x_1, \dots, x_n)}{\partial x_i} \neq 0, i \in \{1, \dots, n\} \right\} \right| \leq 1$.

Using just single-passing Fuzzy Logic operators can be inefficient, since then at most one input will have a nonzero derivative (i.e. a learning signal), yet the complete forward pass still has to be computed to find this input. In particular, this will hold when choosing operators based on the Gödel t-norm.

Proposition 1. Any composition of single-passing functions is also single-passing.

For the proof, see Appendix C.1.

Concluding, for any logical operator to be usable in the learning task, it will need to have a nonvanishing derivative at the majority of the input signals, so it can contribute to the learning signal at all, and ideally not be single-passing so that it can contribute effectively to the learning signal.

6. Aggregation

After the global considerations from the previous section, we next analyze in detail aggregation operators for universal and existential quantification separately and outline their benefits and disadvantages in DFL.

6.1. Minimum and maximum aggregators

The minimum aggregator is given as $A_{T_G}(x_1, \dots, x_n) = \min(x_1, \dots, x_n)$. The partial derivatives are

$$\frac{\partial A_{T_G}(x_1, \dots, x_n)}{\partial x_i} = \begin{cases} 1 & \text{if } i = \operatorname{argmin}_{j \in \{1, \dots, n\}} x_j \\ 0 & \text{otherwise.} \end{cases} \quad (14)$$

It is single-passing with the only nonzero gradient being on the input with the lowest truth value. Many practical formulas have exceptions. An exception to a formula like $\forall x \text{ Raven}(x) \rightarrow \text{Black}(x)$ would be a raven over which a bucket of red paint is thrown. The minimum aggregator would have a derivative on that exception when ‘red raven’ is correctly predicted. Additionally, it is inefficient, as we still have to compute the forward pass for inputs that do not get a feedback signal.

The partial derivatives of the maximum aggregator $E_{S_G}(x_1, \dots, x_n) = \max(x_1, \dots, x_n)$ are similar, but increase the input with the highest truth value instead. This can be a reasonable aggregator for existential quantification, as it will reinforce the belief that the input we are most confident about is correct will make the existential quantifier true. A downside is that it can only consider one such input, despite the fact that the condition might hold for multiple inputs.

6.2. Łukasiewicz aggregator

The Łukasiewicz aggregator is given as $A_{T_{LU}}(x_1, \dots, x_n) = \max(\sum_{i=1}^n x_i - (n - 1), 0)$. The partial derivatives are given by

$$\frac{\partial A_{T_{LU}}(x_1, \dots, x_n)}{\partial x_i} = \begin{cases} 1 & \text{if } \sum_{i=1}^n x_i > n - 1 \\ 0 & \text{otherwise.} \end{cases} \quad (15)$$

The gradient is nonvanishing only when $\sum_{i=1}^n x_i > n - 1$, i.e., when the average value of x_i is larger than $\frac{n-1}{n}$ [57]. As $\lim_{n \rightarrow \infty} \frac{n-1}{n} = 1$, for larger values of n , all inputs have to be high for this to hold.

For the next proposition, we refer to the *fraction of inputs* for which some condition holds. The probability that the condition holds for a point uniformly sampled from $[0, 1]^n$ is this fraction.

Proposition 2. The fraction of inputs $x_1, \dots, x_n \in [0, 1]$ for which the derivative of A_{TLU} is nonvanishing is $\frac{1}{n!}$.

For proof, see Appendix C.2.1. Clearly, for the majority of inputs there is a vanishing gradient, implying that this universal aggregator would not be useful in a DFL learning setting.

A similar argument can be made for the existential Łukasiewicz aggregator, the bounded sum $E_{SLK}(x_1, \dots, x_n) = \min(\sum_{i=1}^n x_i, 1)$, which will only have nonvanishing derivatives of 1 to all argument when the average value of x_i is smaller than $\frac{1}{n}$. Like the Łukasiewicz aggregator, it also has nonvanishing derivatives only on a fraction $\frac{1}{n!}$ of its domain. This is therefore not a useful existential aggregator: The agent will learn nothing unless all inputs are close to 0.

6.3. Yager aggregator

The Yager universal aggregator is given by

$$A_{TY}(x_1, \dots, x_n) = \max \left(1 - \left(\sum_{i=1}^n (1 - x_i)^p \right)^{\frac{1}{p}}, 0 \right), \quad p > 0 \tag{16}$$

Here, $p = 1$ corresponds to the Łukasiewicz aggregator, $p \rightarrow \infty$ corresponds to the minimum aggregator, and $p \rightarrow 0$ corresponds to the aggregator formed from the drastic product A_{TD} . The derivative of the Yager aggregator is

$$\frac{\partial A_{TY}(x_1, \dots, x_n)}{\partial x_i} = \begin{cases} \left(\sum_{j=1}^n (1 - x_j)^p \right)^{1 - \frac{1}{p}} \cdot (1 - x_i)^{p-1} & \text{if } \left(\sum_{j=1}^n (1 - x_j)^p \right)^{\frac{1}{p}} < 1 \\ 0 & \text{if } \left(\sum_{j=1}^n (1 - x_j)^p \right)^{\frac{1}{p}} > 1. \end{cases} \tag{17}$$

This derivative vanishes whenever $\sum_{j=1}^n (1 - x_j)^p \geq 1$. As $1 - x_i \in [0, 1]$, $(1 - x_i)^p$ is a decreasing function with respect to p . Therefore, $\sum_{i=1}^n (1 - x_i)^p < 1$ holds for a larger fraction of inputs when p increases, with the fraction being 0 for $p = 0$ as it corresponds to the drastic aggregator, and 1 for $p = \infty$.

The exact fraction of inputs with a nonvanishing derivative is hard to express in the general case.⁵ However, we can find a closed-form expression for the Euclidean case $p = 2$.

Proposition 3. The fraction of inputs $x_1, \dots, x_n \in [0, 1]$ for which the derivative of A_{TY} with $p = 2$ is nonvanishing is $\frac{\pi^{\frac{n}{2}}}{2^n \cdot \Gamma(\frac{1}{2}n + \frac{1}{2})}$, where Γ is the Gamma function.

See Appendix C.2.2 for the proof. We plot the fraction of nonvanishing derivatives for several values of p in Fig. 2. For fairly small p , the vast majority of the inputs will have a vanishing derivative, and similar for high n , showing that this aggregator is also of little use in a learning context.

Similarly, the derivatives of the Yager existential aggregator $E_{SY}(x_1, \dots, x_n) = \min \left(\left(\sum_{i=1}^n x_i^p \right)^{\frac{1}{p}}, 1 \right)$ are nonvanishing in the same fraction of inputs as the Yager universal aggregator.

6.4. Generalized Mean and Generalized Mean Error

If we are concerned only in maximizing the truth value of A_{TY} , we can simply remove the *max* constraint, resulting in an aggregator that has a nonvanishing derivative everywhere. However, then the co-domain of the function is no longer $[0, 1]$. We can do an affine transformation on this function to ensure this is the case (Appendix D.1), after which we obtain the *Generalized Mean Error*.

⁵ Assume that $x_1, \dots, x_n \sim U(0, 1)$ are independently and standard uniformly distributed. Note that $z_i = 1 - x_i$ is also standard uniformly distributed. z_i^p is distributed by the beta distribution $Beta(1/p, 1)$ [27]. $Y = \sum_{i=1}^n z_i$ is the sum of n such beta-distributed variables. Unfortunately, there is no closed-form expression for the probability density function of sums of independent beta random variables [60]. A suitable approximation would be to use the central limit theorem as z_1, \dots, z_n are identically and independently distributed.

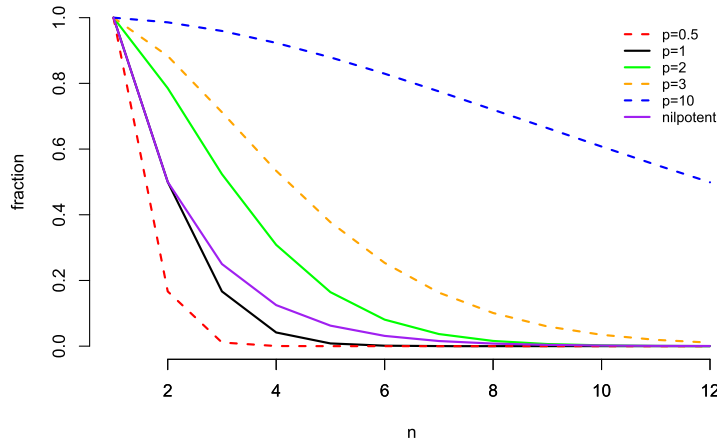


Fig. 2. The fraction of inputs for which the Yager aggregator A_{T_Y} for several values of p and the nilpotent minimum aggregator $A_{T_{nM}}$ have nonvanishing derivatives. The values for dotted lines are estimated using Monte Carlo simulation.

Definition 6. For any $p > 0$, the Generalized Mean Error A_{GME} is defined as

$$A_{GME}(x_1, \dots, x_n) = 1 - \left(\frac{1}{n} \sum_{i=1}^n (1 - x_i)^p \right)^{\frac{1}{p}}. \tag{18}$$

The ‘error’ here is difference between the predicted value x_i and the ‘ground truth’ value, a truth value of 1. This function has the following derivative:

$$\frac{\partial A_{GME}(x_1, \dots, x_n)}{\partial x_i} = (1 - x_i)^{p-1} \frac{1}{n} \left(\sum_{j=1}^n (1 - x_j)^p \right)^{\frac{1}{p}-1}. \tag{19}$$

When $p > 1$, this derivative is greatest for the inputs that are lowest, which can speed up the optimization by being sensitive to outliers. For $p < 1$, the opposite is true. A special case is $p = 1$:

$$A_{MAE}(x_1, \dots, x_n) = 1 - \frac{1}{n} \sum_{i=1}^n (1 - x_i) \tag{20}$$

having the simple derivative $\frac{\partial A_{MAE}(x_1, \dots, x_n)}{\partial x_i} = \frac{1}{n}$. This measure is equal to 1 minus the *mean absolute error* (MAE) and is associated with the Łukasiewicz t-norm. Another special case is $p = 2$:

$$A_{RMSE}(x_1, \dots, x_n) = 1 - \sqrt{\frac{1}{n} \sum_{i=1}^n (1 - x_i)^2}. \tag{21}$$

This function is equal to 1 minus the *root-mean-square error* (RMSE) which is commonly used for regression tasks and heavily weights outliers. We can do the same for the Yager existential aggregator (Appendix D.1):

Definition 7. For any $p > 0$, the Generalized Mean is defined as

$$E_{GM}(x_1, \dots, x_n) = \left(\frac{1}{n} \sum_{i=1}^n x_i^p \right)^{\frac{1}{p}}. \tag{22}$$

$p = 1$ corresponds to the arithmetic mean and $p = 2$ to the geometric mean. In contrast to the Generalized Mean Error, its derivative $\frac{1}{n} \left(\frac{1}{n} \sum_{j=1}^n x_j^p \right)^{\frac{1}{p}-1} x_i^{p-1}$ has greater values for smaller inputs when $p < 1$, and lower values when $p > 1$. Since we want to ensure the derivative is high only for inputs with high truth values to reinforce that those are likely the inputs that confirm the formula, we will want to use $p > 1$. Note that the arithmetic mean E_{GM} has the same derivative as the mean absolute error A_{MAE} , meaning that with $p = 1$ universal and existential quantification cannot be distinguished. Furthermore, increasing all inputs equally is not a great idea for the existential quantifier, as there are likely only a few inputs for which the formula holds. Also note that, unlike existential aggregators directly formed from t-conorms, the only maximum of this aggregator is $x_1, \dots, x_n = 1$. However, it will take long until one can reach this optimum for low inputs.

6.5. Product aggregator and probabilistic sum

The product aggregator is given as $A_{T_p}(x_1, \dots, x_n) = \prod_{i=1}^n x_i$. This is also the probability of the intersection of n independent events. It has the following partial derivatives:

$$\frac{\partial A_{T_p}(x_1, \dots, x_n)}{\partial x_i} = \prod_{j=1, j \neq i}^n x_j. \tag{23}$$

This derivative vanishes if there are at least two i so that $x_i = 0$. Furthermore, the derivative for x_i will be decreased if some other input x_j is low. Finally, we cannot compute this aggregator in practice due to numerical underflow when multiplying many small numbers. Noting that $\operatorname{argmax} f(x) = \operatorname{argmax} \log(f(x))$, we observe that the *log-product aggregator*

$$A_{\log T_p}(x_1, \dots, x_n) = (\log \circ A_{T_p})(x_1, \dots, x_n) = \sum_{i=1}^n \log(x_i) \tag{24}$$

can be used for formulas in prenex normal form, as then the truth value of the universal quantifiers is not used for another connective. Unlike the other aggregators, its codomain is the non-positive numbers instead of $[0, 1]$. Furthermore, the log-product aggregator can be seen as the log-likelihood function where we take the correct label to be 1, and thus this is similar to cross-entropy minimization. The partial derivatives are

$$\frac{\partial A_{\log T_p}(x_1, \dots, x_n)}{\partial x_i} = \frac{1}{x_i}. \tag{25}$$

In contrast to Equation (23), the values of the other inputs are irrelevant, and derivatives with respect to lower-valued inputs will be far greater as there is a singularity at $x = 0$ (i.e. the value becomes infinite). We can conclude therefore that the product aggregator is particularly promising as it is nonvanishing and can handle outliers. The log-product aggregator also combines well with the generalized mean aggregator (Equation 7) for formulas of the form $\forall x \exists y$. The logarithm reduces the outer exponentiation, resulting in the derivative $\frac{x_i^{p-1}}{\sum_i x_i^p}$.

The probabilistic sum aggregator $E_{S_p}(x_1, \dots, x_n) = 1 - \prod_{i=1}^n (1 - x_i)$ is trickier. Its derivatives are

$$\frac{\partial E_{S_p}}{\partial x_i} = \prod_{j=1, j \neq i}^n (1 - x_j). \tag{26}$$

This derivative is quite intuitive: It increases x_i if the other inputs x_j are low. However, this does not take into account the value of x_i itself. If all inputs are low, this will increase all inputs equally. Since the logarithm does not distribute over addition, we cannot use the same trick here as for the product aggregator, so care has to be taken when computing $\log \circ E_{S_p}$ to prevent numerical underflow errors.

6.6. Nilpotent aggregators

The Nilpotent t-norm is given by $T_{nm}(a, b) = \begin{cases} \min(a, b), & \text{if } a + b > 1 \\ 0, & \text{otherwise.} \end{cases}$ In Appendix D.3 we show that the Nilpotent aggregator $A_{T_{nm}}$ is equal to

$$A_{T_{nm}}(x_1, \dots, x_n) = \begin{cases} \min(x_1, \dots, x_n), & \text{if } x_i + x_j > 1; x_i \text{ and } x_j \text{ are the two lowest values in } x_1, \dots, x_n \\ 0, & \text{otherwise} \end{cases}. \tag{27}$$

The derivative is found as follows:

$$\frac{\partial A_{T_{nm}}(x_1, \dots, x_n)}{\partial x_i} = \begin{cases} 1, & \text{if } i = \operatorname{argmin}_j x_j \text{ and } x_i + x_j > 1 \text{ where } x_j \text{ is the second lowest value in } x_1, \dots, x_n \\ 0, & \text{otherwise} \end{cases}. \tag{28}$$

Like the minimum aggregator it is single-passing, and like the Łukasiewicz aggregator it has a derivative that vanishes for the majority of the input space, namely when the sum of the two smallest values is lower than 1.

Proposition 4. The fraction of inputs for which the derivative of $A_{T_{nm}}$ is nonvanishing is $\frac{1}{2^{n-1}}$.

For the proof, see Appendix C.2.3. The fraction of inputs for which there is a nonvanishing derivative is plotted in Fig. 2. Again, this means that for larger numbers of inputs n , this aggregator will vanish on almost every input and is not a useful construction in a learning context.

Similarly, the Nilpotent existential aggregator is given as

$$E_{S_{nM}}(x_1, \dots, x_n) = \begin{cases} \max(x_1, \dots, x_n), & \text{if } x_i + x_j < 1; x_i \text{ and } x_j \text{ are the two largest values in } x_1, \dots, x_n \\ 1, & \text{otherwise.} \end{cases} \quad (29)$$

It has a similar derivative that increases the largest input if the two largest inputs together are lower than 1. This is somewhat similar to how the maximum aggregator behaves, but with an additional condition that will stop increasing the largest input if there is another that is also quite high.

6.7. Summary

The minimum aggregator is computationally inefficient and cannot handle exceptions well. Universal aggregation operators that vanish when receiving a large amount of inputs will not scale well, and these include operators based on the Yager family of t-norms and the nilpotent aggregator. Removing the bounds from the Yager aggregators introduces interesting connections to loss functions from the classical machine learning literature. This is also the case for the logarithmic version of the product aggregator, which corresponds to the cross-entropy loss function. They have natural means for dealing with outliers, and thus are promising for practical use. We have more options for existential quantification, as problems with vanishing gradients are not as important since we only care about ensuring the formula is true for at least one input, instead of all of them.

7. Conjunction and disjunction

Next, we analyze the partial derivatives of t-norms and t-conorms, which are used as conjunction and disjunction in Fuzzy Logics. In t-norm Fuzzy Logics, the weak disjunction $\max(a, b)$, or the Gödel t-conorm is used instead of the dual t-conorm.

Suppose that we have a t-norm T and a t-conorm S . We define the following two quantities, where the choice of taking the partial derivative to a is without loss of generality, since T and S are commutative by definition:

$$d_T(a, b) = \frac{\partial T(a, b)}{\partial a}, \quad d_S(a, b) = \frac{\partial S(a, b)}{\partial a} \quad (30)$$

It should be noted that by Definition 14, $\frac{\partial T(a, 1)}{\partial a} = 1$ as $T(a, 1) = a$ for any t-norm T , and by Definition 16, $\frac{\partial S(a, 0)}{\partial a} = 1$ as $S(a, 0) = a$ for any t-conorm S . Furthermore, we note that if S is a t-conorm and the N_C -dual of the t-norm T , then $\frac{\partial 1-T(1-a, 1-b)}{\partial a} = -\frac{\partial 1-T(1-a, 1-b)}{\partial 1-a} = \frac{\partial T(1-a, 1-b)}{\partial 1-a}$.

The main difference in analyzing t-norms and t-conorms is that the maximum of $T(a, b)$ (namely 1) is when both arguments a and b are 1. In contrast, in t-conorms, an infinite number of maxima exist. Some of these maxima might be more desirable than others. Referring back to the formula in Example 1, we showed that it is preferable to increase the truth value of $\text{cushion}(y)$ and not of $\text{armRest}(y)$. Similarly, when a conjunct is negated, or when it appears in the antecedent of an implication (like in the aforementioned formula) we have to choose which of the two conjuncts to decrease. By noting that $\frac{\partial T(a, b)}{\partial a} = \frac{\partial S(1-a, 1-b)}{\partial 1-a}$, we find that the t-norm “chooses” in the same way its dual t-conorm would “choose”. Similarly, if a disjunction is negated, it will minimize both its arguments in the way that its dual t-norm would maximize its arguments.

Example 3. We introduce a running example to analyze the behavior of different t-norms. We will optimize $(a \otimes b) \oplus (c \otimes \neg a)$ using gradient descent. The truth value of this expression is computed using $f(a, b, c) = S(T(a, b), T(1 - a, c))$. Using the boundary conditions from Definition 14 and 16, we find the global optima $a = 1.0$, $b = 1.0$ and $a = 0.0$, $c = 1.0$. The derivative to this function is, using the chain rule,

$$\frac{\partial f(a, b, c)}{\partial a} = \frac{\partial S(T(a, b), T(1 - a, c))}{\partial T(a, b)} \cdot \frac{\partial T(a, b)}{\partial a} + \frac{\partial S(T(a, b), T(1 - a, c))}{\partial T(1 - a, c)} \cdot \frac{\partial T(1 - a, c)}{\partial a}. \quad (31)$$

7.1. Gödel t-norm

The Gödel t-norm is $T_G(a, b) = \min(a, b)$ and the Gödel t-conorm is $S_G(a, b) = \max(a, b)$. We find

$$\frac{\partial T_G(a, b)}{\partial a} = \begin{cases} 1, & \text{if } a < b \\ 0, & \text{if } a > b \end{cases}, \quad \frac{\partial S_G(a, b)}{\partial a} = \begin{cases} 1, & \text{if } a > b \\ 0, & \text{if } a < b \end{cases}. \quad (32)$$

Both T_G and S_G are single-passing, but their derivatives are not defined $a = b$. A benefit of the magnitude of the derivative nearly always being 1 is that there will not be any exploding or vanishing gradients caused by multiple repeated applications of the chain rule.

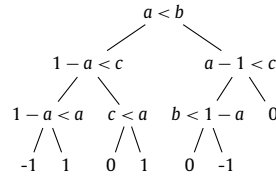


Fig. 3. Decision tree for the derivative of $S_G(T_G(a, b), T_G(1 - a, c))$ with respect to a .

Example 4. Filling in Equation (31) representing $(a \otimes b) \oplus (c \otimes \neg a)$ with T_G and S_G gives

$$\begin{aligned} \frac{\partial f(a, b, c)}{\partial a} &= \mathbf{1}_{\min(a,b) > \min(1-a,c)} \cdot \mathbf{1}_{a < b} - \mathbf{1}_{\min(1-a,c) > \min(a,b)} \cdot \mathbf{1}_{1-a < c} \\ &= \mathbf{1}_{a > \min(1-a,c) \wedge a < b} - \mathbf{1}_{1-a > \min(a,b) \wedge 1-a < c} \end{aligned}$$

where $\mathbf{1}_c$ is the indicator function. This corresponds to the decision tree in Fig. 3. The value of a can be modified to increase the truth of either one of the conjunctions. In order to choose which of the two should be true, it compares a with $1 - a$. If $1 - a < a$, it will increase the first conjunct by increasing a . Gradient ascent always finds a global optimum for this formula.

A small perturbation in the truth values of the inputs can flip the derivative around. For instance, if $a < b$ and $1 - a < c$, then it will increase a if its value is 0.501 and decrease it if it is 0.499. Furthermore, it can cause gradient ascent to get stuck in local optima. For instance, if $\varphi = (a \oplus b) \otimes (\neg a \oplus c)$ and $a = 0.4, b = 0.2$ and $c = 0.1$, gradient ascent increases a until $a > 0.5$, at which point the gradient flips and it decreases a until $a < 0.5$. Experiments with optimizing this through gradient descent show that we can only find a global optimum in 88.8% of random initializations of a, b and c .

7.2. Łukasiewicz t-norm

The Łukasiewicz t-norm is $T_{LK}(a, b) = \max(a + b - 1, 0)$ and the Łukasiewicz t-conorm is $S_{LK}(a, b) = \min(a + b, 1)$. The partial derivatives are:

$$\frac{\partial T_{LK}(a, b)}{\partial a} = \begin{cases} 1, & \text{if } a + b > 1 \\ 0, & \text{if } a + b < 1 \end{cases}, \quad \frac{\partial S_{LK}(a, b)}{\partial a} = \begin{cases} 1, & \text{if } a + b < 1 \\ 0, & \text{if } a + b > 1 \end{cases} \tag{33}$$

These derivatives vanish on as much as half of their domain (Proposition 2). However, like the Gödel t-norm, when there is a gradient, it is large and will not cause vanishing or exploding gradients.

Example 5. Using the Łukasiewicz t-norm and t-conorm in Equation (31) gives rise to the following computation

$$\frac{\partial f(a, b, c)}{\partial a} = \mathbf{1}_{\max(a+b-1,0) + \max(c-a,0) < 1} \cdot (\mathbf{1}_{a+b > 1} - \mathbf{1}_{c-a > 0}).$$

Choosing random values to initialize a, b and c , gradient descent is able to find a (global) optimum in about 83.5% of the initializations.

7.3. Yager t-norm

The family of Yager t-norms [75] is $T_Y(a, b) = \max(1 - ((1 - a)^p + (1 - b)^p)^{\frac{1}{p}}, 0)$ and the family of Yager t-conorms is $S_Y(a, b) = \min((a^p + b^p)^{\frac{1}{p}}, 1)$ for $p \geq 0$. We plot these for $p = 2$ in Fig. 4. The derivatives are given by

$$\frac{\partial T_Y(a, b)}{\partial a} = \begin{cases} ((1 - a)^p + (1 - b)^p)^{\frac{1}{p}-1} \cdot (1 - a)^{p-1} & \text{if } (1 - a)^p + (1 - b)^p < 1, \\ 0 & \text{if } (1 - a)^p + (1 - b)^p > 1, \end{cases} \tag{34}$$

$$\frac{\partial S_Y(a, b)}{\partial a} = \begin{cases} (a^p + b^p)^{\frac{1}{p}-1} \cdot a^{p-1} & \text{if } a^p + b^p < 1, \\ 0 & \text{if } a^p + b^p > 1 \end{cases} \tag{35}$$

We plot these derivatives in Fig. 5, showing for each a vanishing derivative on a non-negligible section of the domain. Using the method described in footnote 5 (Section 6.3), Mathematica finds a closed form expression for the fraction of inputs for which the Yager t-norm is nonvanishing as $\frac{\sqrt{\pi} 4^{-1/p} \Gamma(\frac{1}{p})}{p \Gamma(\frac{1}{2} + \frac{1}{p})}$. Observe that when $p \neq 1$, the derivative of T_Y is undefined at $a = b = 1$ and the derivative of S_Y is undefined at $a = b = 0$. This requires care in the implementation to prevent numerical issues. For $p > 1$, the lower of the two truth values has a higher derivative for the t-norm, while for the t-conorm, the

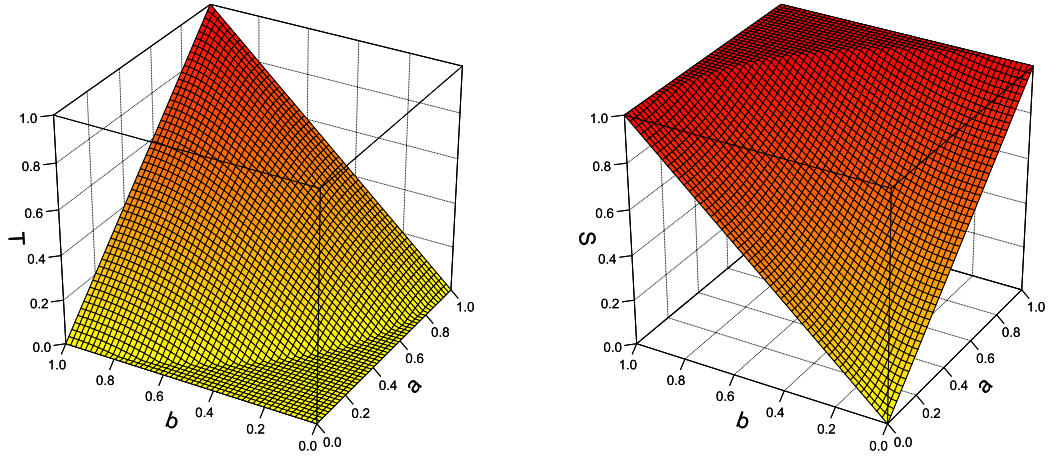


Fig. 4. Left: The Yager t-norm. Right: The Yager t-conorm. For both, $p = 2$.

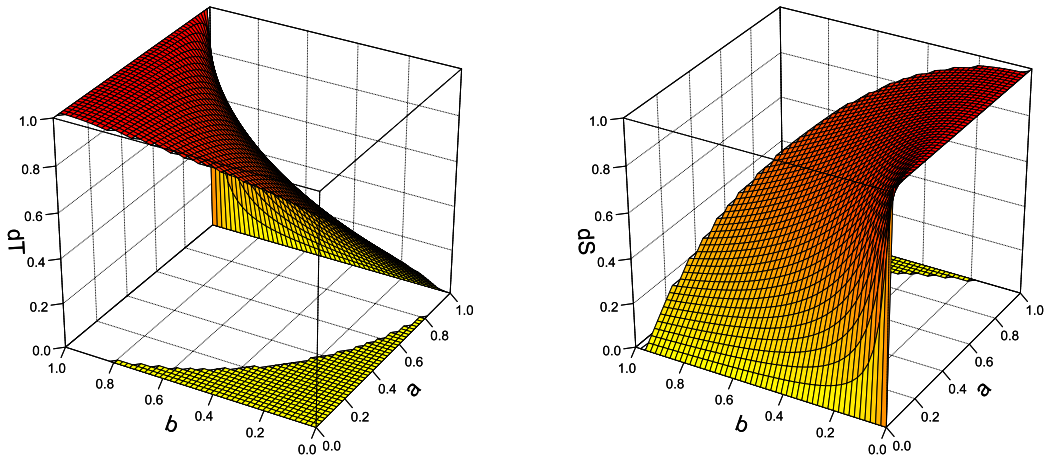


Fig. 5. Left: The derivative of the Yager t-norm. Right: The derivative of the Yager s-norm. For both, $p = 2$.

higher of the two truth values has a higher derivative. As p increases, T_Y and S_Y will behave more like T_G and S_G . Note that when $p < 1$, the t-norm will have higher derivatives for higher inputs as the derivative has a singularity at $\lim_{a \rightarrow 1} = \infty$ ($b < 1$).

7.4. Product t-norm

The product t-norm and t-conorm, visualized in Fig. 6, are $T_P(a, b) = a \cdot b$ and $S_P(a, b) = a + b - a \cdot b$. Their derivatives are

$$\frac{\partial T_P(a, b)}{\partial a} = b, \quad \frac{\partial S_P(a, b)}{\partial a} = 1 - b. \tag{36}$$

The derivative of the t-norm is 0 only when $a = b = 0$, and similarly when $a = b = 1$ for the t-conorm. The derivative of the t-norm can be interpreted as follows: ‘If we wish to increase $a \otimes b$, a should be increased in proportion to b .’ This is not a sensible learning strategy: If both a and b are small, in which case the conjunction is most certainly not satisfied, the derivative will be low instead of high. The derivative of the t-conorm is more intuitive, as it says ‘If we wish to increase $a \oplus b$, a should be increased in proportion to $1 - b$ ’. If b is not yet true, we definitely want at least a to be true.

Example 6. By using the product t-norm and t-conorm in Equation (31), we get

$$\frac{\partial f(a, b, c)}{\partial a} = (1 - (1 - a) \cdot c) \cdot b - (1 - a \cdot b) \cdot c$$

As explained, increase a in proportion to b if it is not true that c and $\neg a$ are true, and decrease a in proportion to c if it is not true that a and b are true.

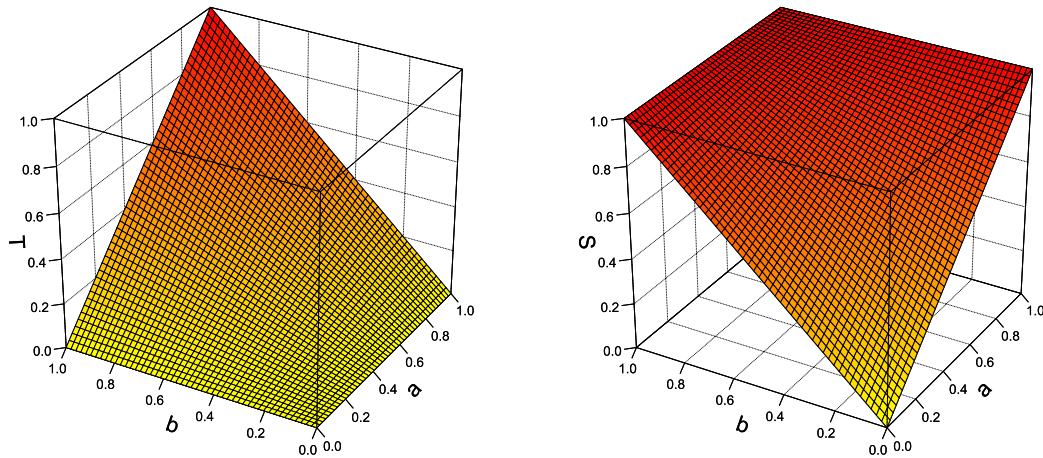


Fig. 6. Left: The product t-norm. Right: The product t-conorm.

7.5. Summary

The Gödel t-norm and t-conorm are simple and effective, having strong derivatives almost everywhere. However, they can be quite brittle by making very binary choices. The Łukasiewicz t-norm and t-conorm also have strong derivatives, but vanish on half of the domain. The Yager family of t-norms and t-conorms also vanish on a significant part of its domain. The derivative of the t-norm is larger for lower values, which is a sensible learning strategy. This is not the case for the product t-norm, where the derivative is dependent on the other input value. However, the product t-conorm is intuitive, and corresponds to the intuition that if one input is not true, the other one should be.

8. Implication

Finally, we consider what functions are suitable for modeling the implication. We will start by discussing the particular challenges associated with the implication operator.

8.1. Challenges of the material implication

A significant proportion of background knowledge is written as universally quantified implications. Examples of such statements are ‘all humans are mortal’, ‘laptops consist of a screen, a processor and a keyboard’ and ‘only humans wear clothes’. These formulas are of the form $\forall x \phi(x) \rightarrow \psi(x)$, where we call $\phi(x)$ the antecedent and $\psi(x)$ the consequent.

The implication is used in two well known rules of inference from classical logic. *Modus ponens* inference says that if $\forall x \phi(x) \rightarrow \psi(x)$ and we know that $\phi(x)$ is true, then $\psi(x)$ should also be true. *Modus tollens* inference, or *contraposition*, says that if $\forall x \phi(x) \rightarrow \psi(x)$ and we know that $\psi(x)$ is false, then $\phi(x)$ should also be false, as otherwise $\psi(x)$ should also have been.

Unlike sequences of conjunctions where each of the formulas should be true, when the agent predicts a scene in which an implication is false, the supervisor has multiple choices. Consider the implication ‘all ravens are black’. There are 4 categories for this formula: *black ravens* (BR), *non-black non-ravens* (NBNR), *black non-ravens* (BNR) and *non-black ravens* (NBR). Assume our agent observes an NBR, then there are four options to consider.

1. *Modus Ponens* (MP): The antecedent is true, so by modus ponens, the consequent is also true. We trust the agent’s observation of a raven and believe it was a black raven (BR).
2. *Modus Tollens* (MT): The consequent is false, so by modus tollens, the antecedent is also false. We trust the agent’s observation of a non-black object and believe that it was not a raven (NBNR).
3. *Distrust*: We think the agent is wrong in both observations, and conclude it was a black non-raven (BNR).
4. *Exception*: We trust the agent in observing a non-black raven (NBR) and ignore the fact that its observation goes against the background knowledge that ravens are black.⁶

The distrust option seems somewhat useless. The exception option can be correct, but we cannot know when there is an exception from the agent’s observations alone.

⁶ This option is not completely ludicrous as white ravens do in fact exist. However, they are rare.

We can assume there are far more non-black objects which are not ravens, than there are ravens. Thus, from a statistical perspective, it is most likely that the agent observed an NBNR. This shows the imbalance associated with the implication, which was first noted in [42] for the Reichenbach implication. It is quite similar to the *class imbalance problem* in Machine Learning [37] in the sense that one has far more *contrapositive* examples than positive examples of the background knowledge.

This problem is closely related to the Raven paradox [31,72,42] from the field of confirmation theory which ponders what evidence can confirm a statement like ‘ravens are black’. It is usually stated as follows:

1. Premise 1: Observing examples of a statement contributes positive evidence towards that statement.
2. Premise 2: Evidence for some statement is also evidence for all logically equivalent statements.
3. Conclusion: Observing examples of non-black non-ravens is evidence for ‘all ravens are black’.

The conclusion follows since ‘non-black objects are non-ravens’ is logically equivalent to ‘ravens are black’. For DFL a similar thing happens. When we correct the observation of an NBR to a BR, the difference in truth value is equal to when we correct it to NBNR. More precisely, representing ‘ravens are black’ as $I(a, b)$, where, for example, $I(1, 1)$ corresponds to BR:

$$A(x_1, \dots, I(1, 0), \dots, x_n) - A(x_1, \dots, I(1, 1), \dots, x_n) = A(x_1, \dots, I(1, 0), \dots, x_n) - A(x_1, \dots, I(0, 0), \dots, x_n)$$

as $I(0, 0) = I(1, 1) = 1$. When one agent observes a thousand BR’s and a single NBR, and another agent observes a thousand NBNR’s and a single NBR, their truth value for ‘ravens are black’ is equal. The first agent has seen many ravens of which only a single was not black. The second only observed non ravens, and a single raven that was not black. Intuitively, the first agent’s beliefs seem more in line with the background knowledge. We will now proceed to analyze a number of implication operators in light of this discussion.

8.2. Analyzing the implication operators

In this section, we choose to take the negation of the derivative with respect to the antecedent as it makes it easier to compare them: all fuzzy implications are monotonically decreasing with respect to the antecedent.

Definition 8. A fuzzy implication I is *contrapositive differentiable symmetric* if $\frac{\partial I(a,c)}{\partial c} = -\frac{\partial I(1-c,1-a)}{\partial 1-c}$ for all $a, c \in [0, 1]$.

A consequence of contrapositive differentiable symmetry is that if $c = 1 - a$, then derivatives with respect to the antecedent and consequent are each other’s negation since $\frac{\partial I(a,c)}{\partial c} = -\frac{\partial I(1-c,1-a)}{\partial 1-c} = -\frac{\partial I(1-(1-a),c)}{\partial 1-(1-a)} = -\frac{\partial I(a,c)}{\partial a}$. This could be seen as the ‘distrust’ option which increases the consequent and negated antecedent equally.

Proposition 5. If a fuzzy implication I is N -contrapositive symmetric (that is, for all $a, c \in [0, 1]$, $I(a, c) = I(N(c), N(a))$), where N is the classical negation, it is also contrapositive differentiable symmetric.

By this proposition all S -implications are contrapositive differentiable symmetric. This property says there is no difference in how the implication handles the derivatives with respect to the consequent and antecedent.

Proposition 6. If an implication I is left-neutral (that is, for all $c \in [0, 1]$, $I(1, c) = c$), then $\frac{\partial I(1,c)}{\partial c} = 1$. If, in addition, I is contrapositive differentiable symmetric, then $-\frac{\partial I(a,0)}{\partial a} = 1$.

The proofs of these two propositions are in Appendix C.3. All S -implications and R -implications are left-neutral, but only S -implications are all also contrapositive differentiable symmetric. The derivatives of R -implications vanish when $a \leq c$, that is, on no less than half of the domain. This is not necessarily a bad property, although this depends highly on the sort of application we use DFL for. If, for example, $a = 0.499$ and $c = 0.5$ for the implication ‘ravens are black’, this expresses a state of uncertainty, and there is probably more that can be learned! However, this will not be possible since the implication vanishes.

8.3. Gödel-based implications

Implications based on the Gödel t-norm make strong discrete choices and are single-passing. As $I_{KD}(a, c) = \max(1 - a, c)$, the derivatives are

$$\frac{\partial I_{KD}(a, c)}{\partial c} = \begin{cases} 1, & \text{if } 1 - a < c \\ 0, & \text{if } 1 - a > c \end{cases}, \quad -\frac{\partial I_{KD}(a, b)}{\partial a} = \begin{cases} 1, & \text{if } 1 - a > c \\ 0, & \text{if } 1 - a < c \end{cases} \tag{37}$$

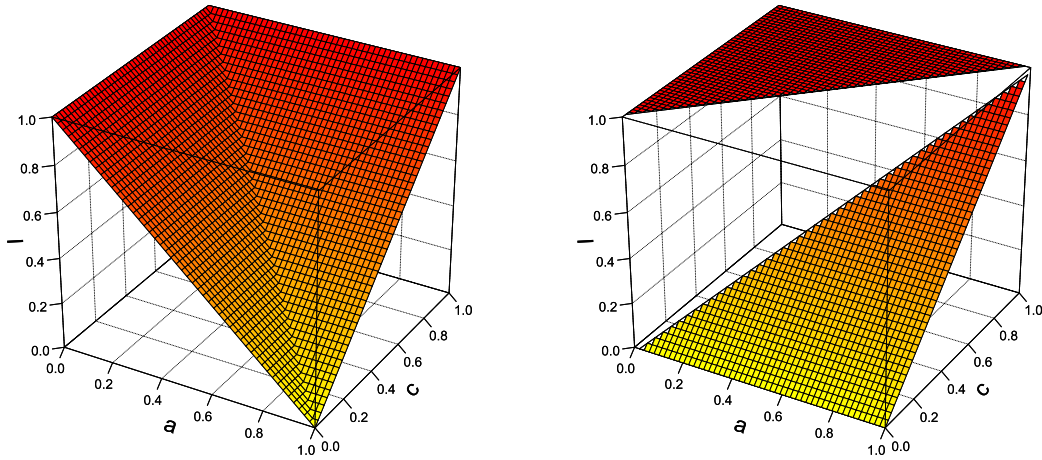


Fig. 7. Left: The Kleene Dienes implication. Right: The Gödel implication. Plots in this section are rotated so that the smallest value is in the front to help understand the shape of the functions. In particular, plots of the derivatives of the implications are rotated 180 degrees compared to the implications themselves.

Or, simply put, if we are more confident in the truth of the consequent than in the truth of the negated antecedent, increase the truth of the consequent. Otherwise, decrease the truth of the antecedent. This decision can be somewhat arbitrary and does not take into account the imbalance of modus ponens and modus tollens.

The Gödel implication is a simple R-implication: $I_G(a, c) = \begin{cases} 1, & \text{if } a \leq c \\ c, & \text{otherwise} \end{cases}$. Its derivatives are:

$$\frac{\partial I_G(a, c)}{\partial c} = \begin{cases} 1, & \text{if } a > c \\ 0, & \text{otherwise} \end{cases}, \quad -\frac{\partial I_G(a, c)}{\partial a} = 0. \tag{38}$$

These two implications are shown in Fig. 7. The Gödel implication increases the consequent whenever $a > c$, and the antecedent is never changed. This makes it a poorly performing implication in practice. For example, consider $a = 0.1$ and $c = 0$. Then the Gödel implication increases the consequent, even if the agent is fairly certain that neither is true. Furthermore, as the derivative with respect to the negated antecedent is always 0, it can never choose the modus tollens correction, which, as we argued, is actually often the best choice.

8.4. Łukasiewicz and Yager-based implications

The Łukasiewicz implication is both an S- and an R-implication. It is given by $I_{LK}(a, c) = \min(1 - a + c, 1)$ and has the simple derivatives

$$\frac{\partial I_{LK}(a, c)}{\partial c} = -\frac{\partial I_{LK}(a, c)}{\partial a} = \begin{cases} 1, & \text{if } a > c \\ 0, & \text{otherwise.} \end{cases} \tag{39}$$

Whenever the implication is not satisfied because the antecedent is higher than the consequent, simply increase the negated antecedent and the consequent until it is lower. This could be seen as the ‘distrust’ choice as both observations of the agent are equally corrected, and so does not take into account the imbalance between modus ponens and modus tollens cases. The derivatives of the Gödel implication I_G are equal to those of I_{LK} except that I_G always has a zero derivative for the negated antecedent.

The Yager S-implication is given as $I_Y(a, c) = \min\left(\left((1 - a)^p + c^p\right)^{\frac{1}{p}}, 1\right)$, $p > 0$. We plot I_Y for $p = 2$ in Fig. 8. $p = 1$ is I_{LK} , $p = 0$ is I_{DP} , and $p = \infty$ is I_{KD} . The derivatives are computed as

$$\frac{\partial I_Y(a, c)}{\partial c} = \begin{cases} \left(\left((1 - a)^p + c^p\right)^{\frac{1}{p}-1} \cdot c^{p-1}\right), & \text{if } (1 - a)^p + c^p \leq 1, \\ 0, & \text{otherwise} \end{cases} \tag{40}$$

$$-\frac{\partial I_Y(a, c)}{\partial a} = \begin{cases} \left(\left((1 - a)^p + c^p\right)^{\frac{1}{p}-1} \cdot (1 - a)^{p-1}\right), & \text{if } (1 - a)^p + c^p \leq 1, \\ 0, & \text{otherwise.} \end{cases} \tag{41}$$

We plot these derivatives for $p = 2$ in Fig. 9. For all p , $\lim_{c \rightarrow 0} \frac{\partial I_Y(1, c)}{\partial c} = 1$. Furthermore, for $p > 1$, $\lim_{a \rightarrow 1} \frac{\partial I_Y(a, c)}{\partial c} \Big|_{c=0} = 0$

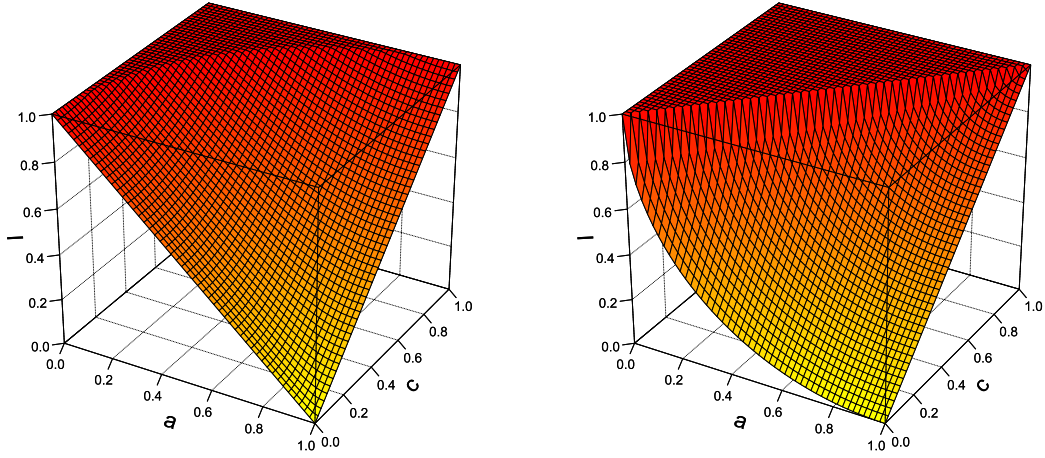


Fig. 8. Left: The Yager S-implication. Right: The Yager R-implication. For both, $p = 2$.

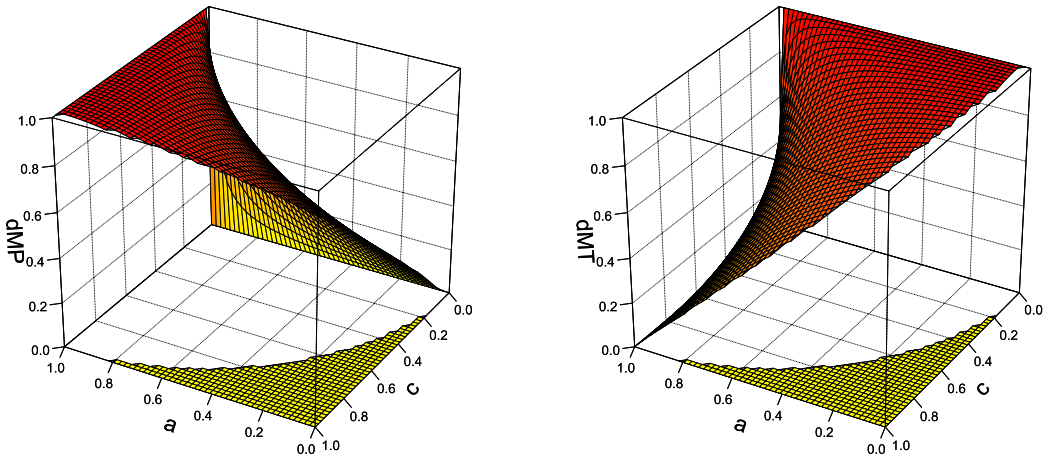


Fig. 9. Plots of the derivatives of the Yager S-implication for $p = 2$.

and for $p < 1$, $\lim_{a \rightarrow 1} \frac{\partial I_Y(a, 0)}{\partial a} \Big|_{c=0} = \infty$. For $p > 1$, I_Y can be understood as an increasingly less smooth version of the Kleene-Dienes implication I_{KD} . Lastly, this derivative, like those for T_Y and S_Y (Section 7.3), is nonvanishing for a fraction of $\frac{\sqrt{\pi} 4^{-1/p} \Gamma(\frac{1}{p})}{p \Gamma(\frac{1}{2} + \frac{1}{p})}$ of the input space.

The Yager R-implication is found (Appendix D.4) as $I_{T_Y}(a, c) = \begin{cases} 1, & \text{if } a \leq c \\ 1 - ((1-c)^p - (1-a)^p)^{\frac{1}{p}}, & \text{otherwise.} \end{cases}$ We plot I_{T_Y} for $p = 2$ in Fig. 8. As expected, $p = 1$ reduces to I_{LK} , $p = 0$ reduces to I_{WB} and $p = \infty$ reduces to I_C . It is contrapositive symmetric only for $p = 1$. The derivatives of this implication are

$$\frac{\partial I_{T_Y}(a, c)}{\partial c} = \begin{cases} ((1-c)^p - (1-a)^p)^{\frac{1}{p}-1} \cdot (1-c)^{p-1}, & \text{if } a > c \\ 0, & \text{otherwise,} \end{cases} \tag{42}$$

$$-\frac{\partial I_{T_Y}(a, c)}{\partial a} = \begin{cases} ((1-c)^p - (1-a)^p)^{\frac{1}{p}-1} \cdot (1-a)^{p-1}, & \text{if } a > c \\ 0, & \text{otherwise.} \end{cases} \tag{43}$$

We plot these in Fig. 10. Note that if $p > 1$, for all $c < 1$ it holds that $\lim_{a \downarrow c} \frac{\partial I_{T_Y}(a, c)}{\partial c} = \lim_{a \downarrow c} -\frac{\partial I_{T_Y}(a, c)}{\partial a} = \infty$ as when a approaches c from above, $(1-c)^p - (1-a)^p$ approaches 0, giving a singularity as $0^{\frac{1}{p}-1}$ is undefined. This collection of singularities makes the training unstable in practice.

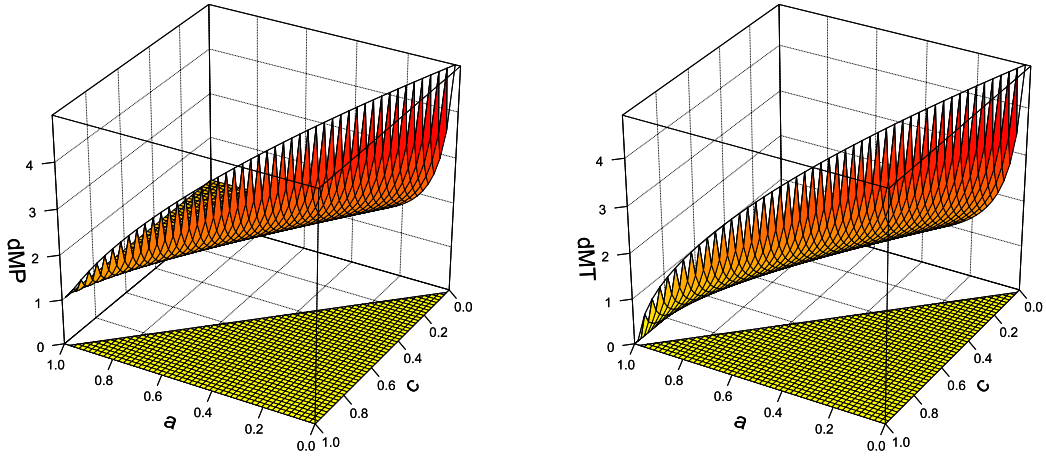


Fig. 10. Plots of the derivatives of the Yager R-implication for $p = 2$.

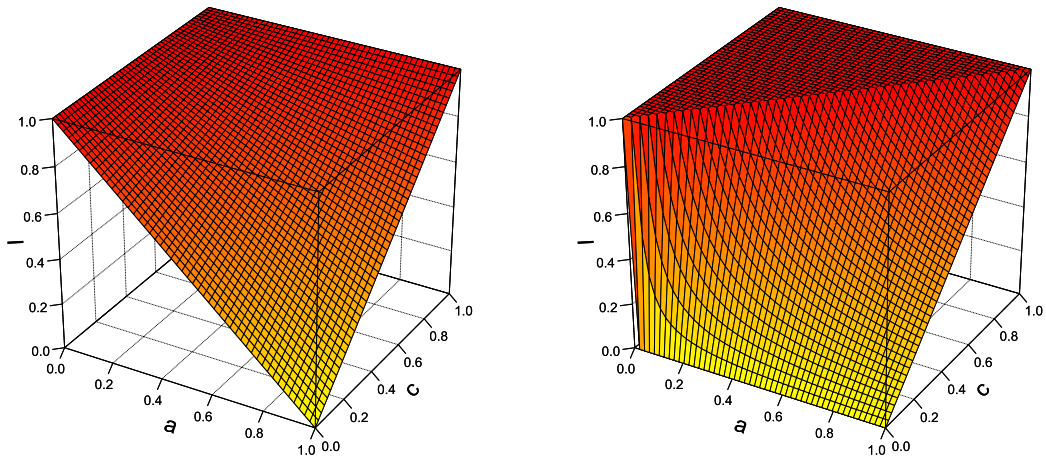


Fig. 11. Left: The Reichenbach implication. Right: The Goguen implication.

8.5. Product-based implications

The product S-implication, also known as the Reichenbach implication, is given by $I_{RC}(a, c) = 1 - a + a \cdot c$. We plot it in Fig. 11. Its derivatives are given by:

$$\frac{\partial I_{RC}(a, c)}{\partial c} = a, \quad -\frac{\partial I_{RC}(a, c)}{\partial a} = 1 - c. \tag{44}$$

These derivatives closely follow the modus ponens and modus tollens rules. When the antecedent is high, increase the consequent, and when the consequent is low, decrease the antecedent. However, around $(1 - a) = c$, the derivative is equal and the ‘distrust’ option is chosen. This can result in counter-intuitive behavior. For example, if the agent predicts 0.6 for raven and 0.5 for black and we use gradient descent until we find a maximum, we could end up at 0.3 for raven and 1 for black. We would end up increasing our confidence in black as raven was high. However, because of additional modus tollens reasoning, raven is barely true.

Furthermore, if the agent most of the time predicts values around $a = 0, c = 0$ as a result of the modus tollens case being the most common, then a majority of the gradient decreases the antecedent as $-\frac{\partial I_{RC}(a,0)}{\partial a} \Big|_{a=0} = 1$. We identify two methods that counteract this behavior. We introduce the second method in Section 8.5.1. The first method for counteracting the ‘corner’ behavior notes that different aggregators change how the derivatives of the implications behave when their truth value is high. For instance, we find that the derivatives with respect to the negated antecedent when using the log-product aggregator and RMSE aggregator are

$$\frac{\partial \log \circ A_P(I_{RC}(a_1, c_1), \dots, I_{RC}(a_n, c_n))}{\partial 1 - a_i} = \frac{1 - c_i}{1 - a_i + a_i \cdot c_i} = \frac{\neg c}{a \rightarrow c}, \tag{45}$$

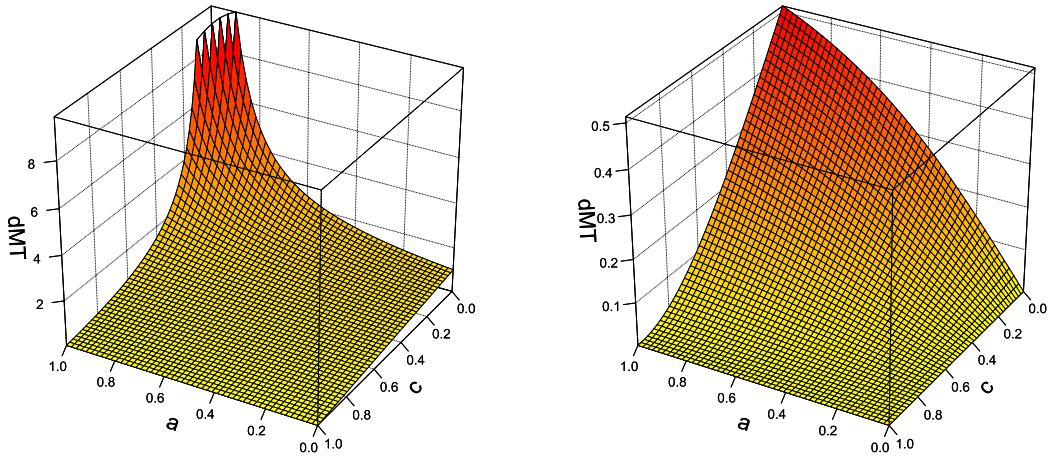


Fig. 12. Left: The antecedent derivative of the Reichenbach implication with the log-product aggregator. Right: The antecedent derivative of the Reichenbach implication with the RMSE aggregator.

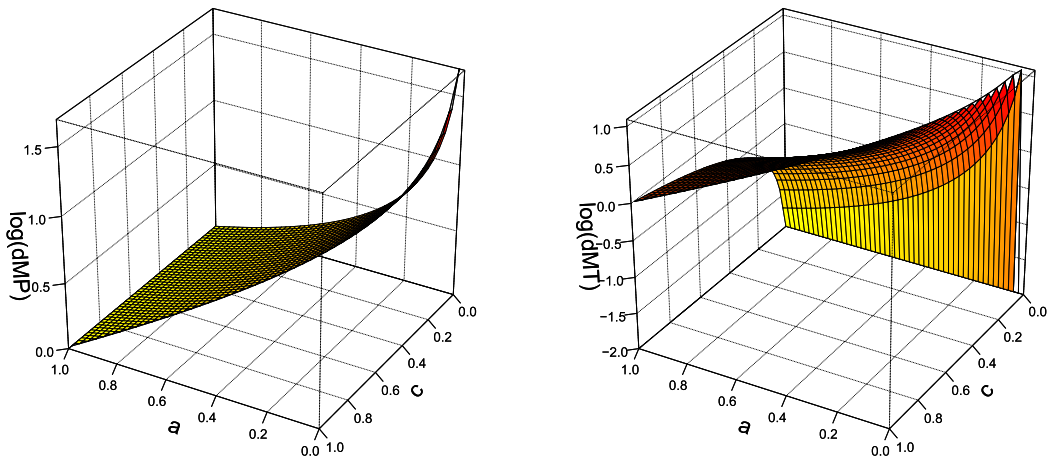


Fig. 13. The derivatives of the Goguen implication. Note that we plot these in log scale.

$$\frac{\partial A_{RMSE}(I_{RC}(a_1, c_1), \dots, I_{RC}(a_n, c_n))}{\partial 1 - a_i} = \frac{(1 - c_i)(a_i - a_i \cdot c_i)}{\sqrt{n \sum_{j=1}^n (a_j - a_j \cdot c_j)^2}} = \frac{-c_i(\neg(a_i \rightarrow c_i))}{\sqrt{n \sum_{j=1}^n (\neg(a_j \rightarrow c_j))^2}}. \tag{46}$$

We plot these functions in Fig. 12. For the RMSE aggregator we choose $n = 2$ and a_1, c_1 so that $(a_1 - a_1 \cdot c_1)^2 = 0.9$. Note that the derivative with respect to the negated antecedent using the RMSE aggregator is 0 in $a_i = 0, c_i = 0$ as then $a_i - a_i \cdot c_i = 0$, and using the log-product aggregator, the derivative is 1. By differentiable contrapositive symmetry, the consequent derivative is 0 when using both aggregators. This shows that when using the RMSE aggregator, the derivatives will vanish at the corners $a = 0, c = 0$ and $a = 1, c = 1$, while when using the log-product aggregator, one of a and c will still have a gradient.

The R-implication of the product t-norm is the Goguen implication and given by $I_{GG}(a, c) = \begin{cases} 1, & \text{if } a \leq c \\ \frac{c}{a}, & \text{otherwise} \end{cases}$. We plot this implication in Fig. 11. The derivatives of I_{GG} are

$$\frac{\partial I_{GG}(a, c)}{\partial c} = \begin{cases} 0, & \text{if } a \leq c \\ \frac{1}{a}, & \text{otherwise} \end{cases}, \quad -\frac{\partial I_{GG}(a, c)}{\partial a} = \begin{cases} 0, & \text{if } a \leq c \\ \frac{c}{a^2}, & \text{otherwise} \end{cases}. \tag{47}$$

We plot these in Fig. 13. This derivative is not very useful. First of all, both the modus ponens and modus tollens derivatives increase with $\neg a$. This is opposite of the modus ponens rule as when the antecedent is low, it increases the consequent most. For example, if raven is 0.1 and black is 0, then the derivative with respect to black is 10, because of the singularity when a approaches 0.

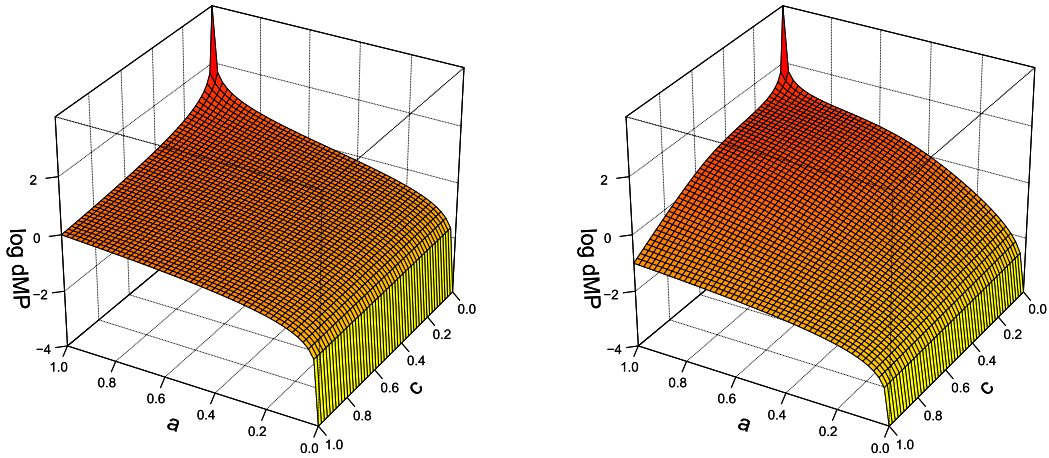


Fig. 14. The consequent derivatives of the log-Reichenbach and log-Reichenbach-sigmoidal (with $s = 9$) implications. The figure is plotted in log scale.

8.5.1. Sigmoidal implications

For the second method for tackling the corner problem, we introduce a new class of fuzzy implications formed by transforming other fuzzy implications using the sigmoid function and translating it so that the boundary conditions still hold. The derivation, along with several proofs of properties, can be found in Appendix D.2.

Definition 9. If I is a fuzzy implication, then the I -sigmoidal implication σ_I is given for $s > 0$ and $b_0 \in \mathbb{R}$ as

$$\sigma_I(a, c) = \frac{1 + e^{-s(1+b_0)}}{e^{-b_0s} - e^{-s(1+b_0)}} \cdot \left((1 + e^{-b_0s}) \cdot \sigma(s \cdot (I(a, c) + b_0)) - 1 \right) \tag{48}$$

where $\sigma(x) = \frac{1}{1+e^x}$ denotes the sigmoid function.

Here b_0 is a parameter that controls the position of the sigmoidal curve and s controls the ‘spread’ of the curve. σ_I is the function $\sigma(s \cdot (I(a, c) + b_0))$ linearly transformed so that its codomain is the closed interval $[0, 1]$. For the common value of $b_0 = -\frac{1}{2}$, a simpler form exists:

$$\sigma_I(a, c) = \frac{1}{e^{\frac{s}{2}} - 1} \cdot \left((1 + e^{\frac{s}{2}}) \cdot \sigma\left(s \cdot \left(I(a, c) - \frac{1}{2}\right)\right) - 1 \right). \tag{49}$$

Next, we give the derivative of σ_I . Substituting $d = \frac{1+e^{-s(1+b_0)}}{e^{-s \cdot b_0} - e^{-s(1+b_0)}}$ and $h = (1 + e^{-s \cdot b_0})$, we find

$$\frac{\partial \sigma_I(a, c)}{\partial I(a, c)} = d \cdot h \cdot s \cdot \sigma(s \cdot (I(a, c) + b_0)) \cdot (1 - \sigma(s \cdot (I(a, c) + b_0))). \tag{50}$$

This keeps the properties of the original function but smoothens the gradient for higher values of s . As the derivative of the sigmoid function is positive, this derivative vanishes only when the derivative of I vanishes.

We plot the derivatives for the Reichenbach-sigmoidal implication $\sigma_{I_{RC}}$ in Fig. 15. As expected by Proposition 16, it is differentiable contrapositive symmetric. Compared to the derivatives of the Reichenbach implication it has a small gradient in all corners. When using the log-product aggregator, the derivative of the antecedent with respect to the total valuation is divided by the truth of the implication. In Fig. 14 we compare the consequent derivative of the normal Reichenbach implication with the Reichenbach-sigmoidal implication when using the log function. Clearly, for both there is a singularity at $a = 1, c = 0$, as then the implication is 0 and so the derivative of the log function becomes infinite. A significant difference is that the sigmoidal variant is less ‘flat’ than the normal Reichenbach implication. This can be useful, as this means there is a larger gradient for values of c that make the implication less true. In particular, the gradient at the modus ponens case ($a = 1, c = 1$) and the modus tollens case ($a = 0, c = 0$) are far smaller, which could help balancing the effective total gradient by solving the ‘corner’ problem of the Reichenbach implication we brought up in Section 8.5. These derivatives are smaller for higher values of s .

In Fig. 16 we plot the Reichenbach-sigmoidal implication for different values of the hyperparameters b_0 and s . Comparing 16a and 16b we see that larger values of b_0 move the sigmoidal shape so that its center is at lower input values. Note that for $s = 0.01$ in Fig. 16c, the plotted function is indiscernible from the plot of the Reichenbach implication in Fig. 11 as the interval on which the sigmoid acts is extremely small and the sigmoidal transformation is almost linear. For very high values of s like in 16d we see that the ‘S’ shape is much thinner, and a larger part of the domain has a low derivative.

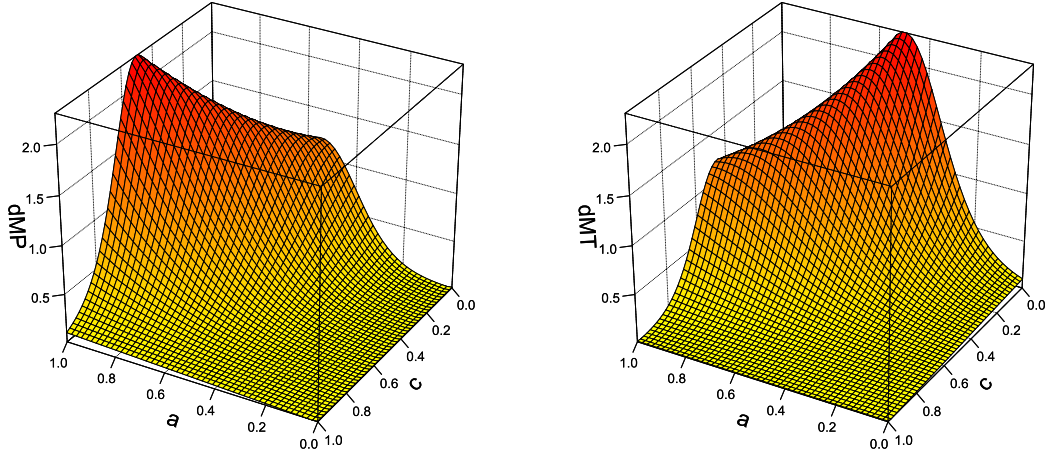


Fig. 15. The derivatives of the Reichenbach-sigmoidal implication for $s = 9$.

8.6. Summary

We analyzed several fuzzy implications from a theoretical perspective, while keeping the challenges caused by the material implication in mind. As a result of this analysis, we find that popular R-implications, in particular the Gödel implication, the Yager R-implication and the Goguen implication, will not work well in a differentiable setting. The other analyzed implications seem to have more intuitive derivatives, but may have other practical issues like non-smoothness.

9. Experimental setup

To get insights in the behavior of these operators in practice, we next perform a series of simple experiments to analyze them. We discuss experiments using the MNIST dataset of handwritten digits [43] to investigate the behavior of different fuzzy operators introduced in this paper. The goal of the experiments is not to show that our method is state of the art for the problem of semi-supervised learning on MNIST, but rather to be able to get insights into how fuzzy operators behave in a differentiable setting.⁷

9.1. Measures

To investigate the performance of the different configurations of DFL, we first introduce several useful metrics. These give us insight into how different operators behave. In this section, we assume we are dealing with formulas of the form $\varphi = \forall x_1, \dots, x_m \phi \rightarrow \psi$.

Definition 10. The *consequent magnitude* $|\text{cons}|$ and the *antecedent magnitude* $|\text{ant}|$ for a knowledge base \mathcal{K} is defined as the sum of the partial derivatives of the consequent and antecedent with respect to the DFL loss:

$$|\text{cons}| = \sum_{\varphi \in \mathcal{K}} \sum_{\mu \in M_\varphi} \frac{\partial e_\theta(\mu, \varphi)}{\partial e_\theta(\mu, \psi)}, \quad |\text{ant}| = \sum_{\varphi \in \mathcal{K}} - \sum_{\mu \in M_\varphi} \frac{\partial e_\theta(\mu, \varphi)}{\partial e_\theta(\mu, \phi)}, \quad (51)$$

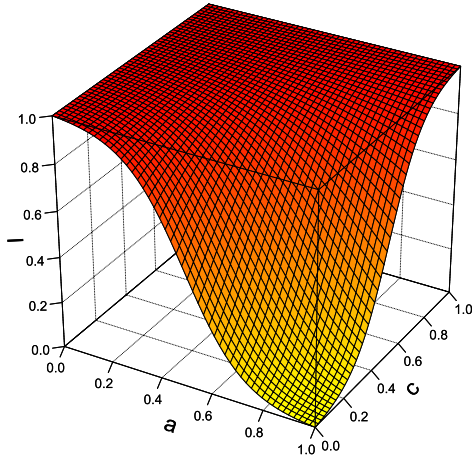
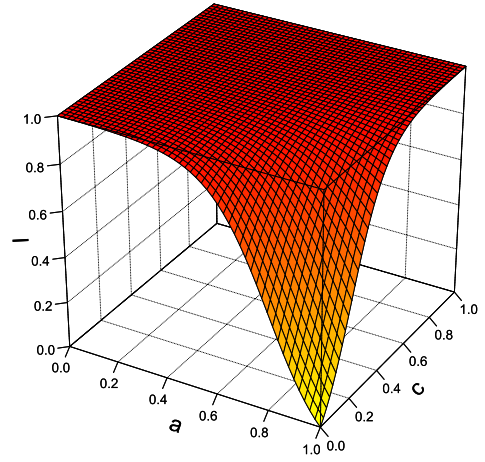
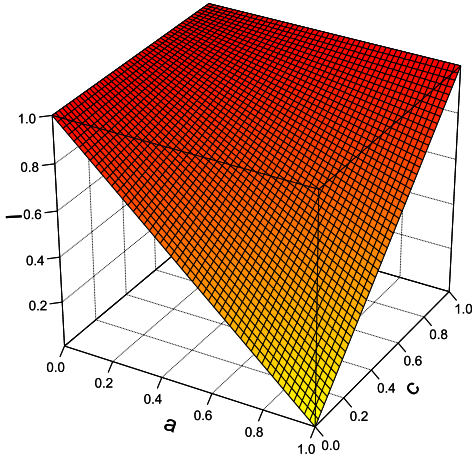
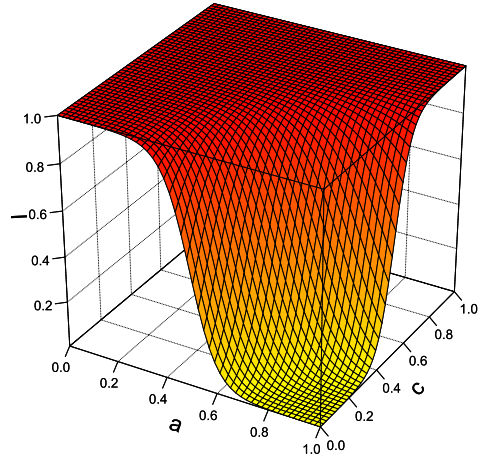
where M_φ is the set of instances of the universally quantified formula φ and ψ and ϕ are evaluated under instantiation μ .

The *consequent ratio* $\text{cons}\%$ is the sum of consequent magnitudes divided by the sum of consequent and antecedent magnitudes: $\text{cons}\% = \frac{|\text{cons}|}{|\text{cons}| + |\text{ant}|}$.

Definition 11. Given a *labeling function* l that returns the truth value of a formula according to the data for instance μ , the *consequent and antecedent correctly updated magnitudes* are the sum of partial derivatives for which the consequent or the negated antecedent is true:

$$cu_{\text{cons}} = \sum_{\varphi \in \mathcal{K}} \sum_{\mu \in M_\varphi} l(\psi, \mu) \cdot \frac{\partial e_\theta(\mu, \varphi)}{\partial e_\theta(\mu, \psi)}, \quad cu_{\text{ant}} = \sum_{\varphi \in \mathcal{K}} - \sum_{\mu \in M_\varphi} l(\neg\phi, \mu) \cdot \frac{\partial e_\theta(\mu, \varphi)}{\partial e_\theta(\mu, \phi)}. \quad (52)$$

⁷ Code is available at <https://github.com/HEmile/differentiable-fuzzy-logics>.

(a) $b_0 = -0.5, s = 9$ (b) $b_0 = -0.2, s = 9$ (c) $b_0 = -0.5, s = 0.01$ (d) $b_0 = -0.5, s = 20$ **Fig. 16.** The Reichenbach-sigmoidal implication for different values of b_0 and s .

That is, if the consequent is true in the data, we measure the magnitude of the derivative with respect to the consequent. To evaluate these quantities, we define ratios similar to a precision metric:

Definition 12. The *correctly updated ratios* for consequent and antecedent are defined as

$$cu_{\text{cons}}\% = \frac{\sum_{\varphi \in \mathcal{K}} cu_{\text{cons}}\varphi}{\sum_{\varphi \in \mathcal{K}} |\text{cons}|_{\varphi}}, \quad cu_{\text{ant}}\% = \frac{\sum_{\varphi \in \mathcal{K}} cu_{\text{ant}}\varphi}{\sum_{\varphi \in \mathcal{K}} |\text{ant}|_{\varphi}}. \quad (53)$$

These quantify what fraction of the updates are going in the right direction. When these ratios approach 1, DFL will always increase the truth value of the consequent or negated antecedent correctly.⁸ Otherwise, we are increasing truth values of subformulas that are wrong. Ideally, we want these measures to be high.

⁸ It can still change the truth value of a ground atom wrongly if ϕ or ψ are not atomic formulas.

9.2. Formulas

We use a knowledge base \mathcal{K} of universally quantified logic formulas. There is a predicate for each digit, that is zero, one, ..., eight and nine. For example, $\text{zero}(x)$ is true whenever x is a handwritten digit labeled with 0. We have two sets of formulas where we learn an additional binary predicate.

9.2.1. The same problem

The same problem is a simple problem to test different operators for implication and universal aggregation. We use the binary predicate `same` that is true whenever both its arguments are the same digit. We next describe the formulas we use.

1. $\forall x, y \text{ zero}(x) \otimes \text{zero}(y) \rightarrow \text{same}(x, y), \dots, \forall x, y \text{ nine}(x) \otimes \text{nine}(y) \rightarrow \text{same}(x, y)$. If both x and y are handwritten zeros, for example, then they represent the same digit.
2. $\forall x, y \text{ zero}(x) \otimes \text{same}(x, y) \rightarrow \text{zero}(y), \dots, \forall x, y \text{ nine}(x) \otimes \text{same}(x, y) \rightarrow \text{nine}(y)$. If x and y represent the same digit and one of them represents zero, then the other one does as well.
3. $\forall x, y \text{ same}(x, y) \rightarrow \text{same}(y, x)$. This formula encodes the symmetry of the `same` predicate.

We find in Appendix F.1 that a set of operators is better than random guessing for the consequent updates if $cu_{\text{cons}}\% > 0.1$, and that we know with confidence a set of operators to be better than random if $cu_{\text{ant}}\% > 0.99$.

9.2.2. The sum9 problem

In the second problem we use the binary predicate `sum9` that is true whenever its arguments sum to 9. We use this problem to test existential quantification, conjunction and disjunction. The formulas are

1. $\forall x \exists y \text{ sum9}(x, y)$. For each digit, there is another such that their sum is 9.⁹
2. $\forall x, y \text{ sum9}(x, y) \rightarrow (\text{zero}(x) \otimes \text{nine}(y)) \oplus (\text{one}(x) \otimes \text{eight}(y)) \oplus \dots \oplus (\text{nine}(x) \otimes \text{zero}(y))$: This formula defines the `sum9` predicate.

9.3. Experimental methodology

We split the MNIST dataset so that 1% of it is labeled and 99% is unlabeled. Given a handwritten digit \mathbf{x} labeled with digit y , $p_\theta(y|\mathbf{x})$ computes the distribution over the 10 possible labels. We use 2 convolutional layers with max pooling, the first with 10 and the second with 20 filters. Then follows two fully connected hidden layers with 320 and 50 nodes and a softmax output layer. The probability that `same`($\mathbf{x}_1, \mathbf{x}_2$) for two handwritten digits \mathbf{x}_1 and \mathbf{x}_2 holds is modeled by $p_\theta(\text{same}|\mathbf{x}_1, \mathbf{x}_2)$. This takes the 50-dimensional embeddings of \mathbf{x}_1 and \mathbf{x}_2 of the fully connected hidden layer $e_{\mathbf{x}_1}$ and $e_{\mathbf{x}_2}$. These are used in a network architecture called a Neural Tensor Network [69]:

$$p_\theta(\text{same}|\mathbf{x}_1, \mathbf{x}_2) = \sigma \left(u^\top \tanh \left(e_{\mathbf{x}_1}^\top W^{[1:k]} e_{\mathbf{x}_2} + V \begin{bmatrix} e_{\mathbf{x}_1} \\ e_{\mathbf{x}_2} \end{bmatrix} + b \right) \right). \tag{54}$$

$W^{[1:k]} \in \mathbb{R}^{d \times d \times k}$ is used for the bilinear tensor product, $V \in \mathbb{R}^{k \times 2d}$ is used for the concatenated embeddings and $b \in \mathbb{R}^k$ is used as a bias vector. We use $k = 50$ for the size of the hidden layer. $u \in \mathbb{R}^k$ is used to compute the output logit, which goes through the sigmoid function σ to get the confidence value.

The loss function we use is split up in three parts, the first over the unlabeled dataset \mathcal{D}_u , and the two others over the labeled dataset \mathcal{D}_l :

$$\mathcal{L}(\theta) = w_{DFL} \cdot \mathcal{L}_{DFL}(\langle \mathcal{D}_u, \eta, \theta \rangle, \mathcal{K}) - \sum_{\mathbf{x}, y \in \mathcal{D}_l} \log p_\theta(y|\mathbf{x}) - \sum_{\substack{\mathbf{x}_1, y_1, \mathbf{x}_2, y_2 \\ \in \mathcal{D}_l \times \mathcal{D}_l}} \log p_\theta(\text{same} = \mathbf{1}_{y_1=y_2} | \mathbf{x}_1, \mathbf{x}_2) \tag{55}$$

The first term is the DFL loss which is weighted by the *DFL weight* w_{DFL} . The second is the supervised cross entropy loss with a batch size of 64. The third is the supervised binary cross entropy loss used to learn to recognize `same`(x, y).¹⁰ This loss is $\log p_\theta(\text{sum9} = \mathbf{1}_{y_1+y_2=9} | \mathbf{x}_1, \mathbf{x}_2)$ for the `sum9` problem. As there are far more negative examples than positive examples, we undersample the negative examples. Note that the two supervised losses can also be seen as a universal aggregation over logical facts using the log-product aggregator: $-A_{\log T_P}(p_\theta(y_1|\mathbf{x}_1), \dots, p_\theta(y_j|\mathcal{D}_l) | \mathbf{x}|\mathcal{D}_l)$. For optimization, we used ADAM [39] with a learning rate of 0.001.

⁹ We sample minibatches of 64 digits, which means there is a negligible probability that there exists a digit in the minibatch for which there is no match (0.0117, to be precise).

¹⁰ It is possible to not use this loss term and learn the `same` predicate using just the formulas, although this is more challenging and only works with a good set of fuzzy operators.

Table 6

Results on the same problem for several symmetric configurations with either S-implications or R-implications. For all, $w_{DFL} = 1$ except for T_P , for which $w_{DFL} = 10$.

	S-Implications				R-Implications			
	Accuracy	cons%	$cu_{cons}\%$	$cu_{ant}\%$	Accuracy	cons%	$cu_{cons}\%$	$cu_{ant}\%$
T_G	95.3	0.32	0.31	0.83	95.0	1	0.11	-
T_P	96.5	0.08	0.72	0.99	94.8	0.62	0.04	0.96
T_{LK}	94.9	0.5	0.86	0.12	94.9	0.5	0.86	0.12
$T_Y, p = 1.5$	95.2				95.2			
$T_Y, p = 2$	77.7	0.20	0.51	0.75	95.0	0.62	0.61	0.46
$T_Y, p = 20$	95.6	0.02	0.54	0.75	95.5	0.53	0.01	0.99
T_{Nm}	95.2				95.2			

10. Results

We ran experiments for many combinations of operators with the aim of showing that the discussed insights are present in practice. We report the accuracy of recognizing digits in the test set, the consequent ratio cons%, and the consequent and antecedent correctly updated ratios $cu_{cons}\%$ and $cu_{ant}\%$. We train for at most 70.000 iterations (or until convergence). The purely supervised baseline has a test accuracy of $95.18\% \pm 0.204$, and runs for about 35 minutes. Semi-supervised methods should improve upon this baseline to be useful. Our implementation including DFL runs for 1 hour and 52 minutes.

10.1. Symmetric configurations

First, we consider several *symmetric* configurations, where the conjunction is a t-norm T , disjunction the dual t-conorm of T , universal aggregation the extended t-norm A_T , existential aggregation the extended t-conorm E_S and the implication either is the S-implication based on the t-conorm or the R-implication based on the t-norm. For example, for T_P we use T_P for conjunction, S_P for disjunction, $A_{\log T_P} = \log \circ A_{T_P}$ for aggregation and I_{RC} for implication. Symmetric configurations will retain many equivalence relations in fuzzy logic, unlike when one would choose arbitrary configuration of operators.

10.1.1. Symmetric configurations on same problem

All configurations are run with $w_{DFL} = 1$ except for T_P which is run using $w_{DFL} = 10$. The results on the same problem can be found in Table 6. One general observation that can be made is that S-implications seem to work much better than R-implications. The only configuration with R-implications that outperform the supervised baseline is $T_Y, p = 20$, but here the S-implication performs similar to it. We hypothesize this is because the derivatives of R-implications vanish whenever $a \leq c$.

The Gödel t-norm performs on par with the supervised baseline. This is because the min aggregator is single-passing, like the configuration as a whole. The single instance which receives a derivative might just be an exception as argued in Section 6.1 and evident from the low values of $cu_{cons}\%$ and $cu_{ant}\%$.

The Łukasiewicz t-norm performs worse than the supervised baseline. Since A_{LK} either has a derivative of 0 or 1 everywhere, the total gradient is very large when it does not vanish. By the definition of I_{LK} , $cons\% = \frac{1}{2}$ as the consequent and negated antecedent derivatives are equal (see Equation (39)). $cu_{cons}\%$ is very low with only 0.01, which is worse than random guessing. As half of the gradient is MP reasoning, that half is nearly always incorrect. The performance of the Yager t-norm seems highly dependent on the choice of the parameter p . For $p = 20$ the top performance is quite a bit higher than the baseline. The lower the value of p , the more likely it is that the derivative of the universal aggregator vanishes. However, for $p = 2$, the results are even worse than the Łukasiewicz t-norm, which corresponds to $p = 1$, while the derivative with $p = 1.5$ simply vanishes throughout the whole run.

The product t-norm performs best and also has the highest values for $cu_{cons}\%$ and $cu_{ant}\%$. To a large extent this is because the log-product aggregator is very effective, as other symmetric configurations also perform much better with it Appendix F.2. Finally, the Nilpotent t-norm performs exactly like the supervised baseline since the derivative of the universal aggregator vanished during the complete training run.

10.1.2. Symmetric configurations on the sum9 problem

In addition to the same problem, we also run the symmetric configurations on the sum9 problem to be able to also take into account how existential quantification and disjunction behave. The results are in Table 7. These closely reflect the results for the same problem. Again, the only configurations that clearly outperform the supervised baseline are the product t-norm and the Yager t-norm with $p = 20$, with the product t-norm being the most promising candidate. Furthermore, in addition to the Nilpotent t-norm, the derivatives of the Łukasiewicz t-norm and Yager t-norm with $p = 2$ vanish throughout the whole run.

10.2. Individual operators

We also perform several experiments where we investigate the contribution of specific fuzzy operators without regard to whether the resulting configurations are sensible in a logical sense. We do this to better understand how each operator

Table 7
Results on the sum9 problem for several symmetric configurations using the S-implication. $w_{DFL} = 1$ except for T_P with $w_{DFL} = 10$.

	Accuracy	cons%	cucons%	cuant%
T_G	95.2	0.31	0.44	0.78
T_P	96.1	0.13	0.80	0.95
T_{LK}	95.2			
$T_Y, p = 1.5$	95.2			
$T_Y, p = 2$	95.2			
$T_Y, p = 20$	95.5	0.99	0.82	0.71
T_{Nm}	95.2			

Table 8

Left: Results on the same problem, varying the universal aggregator. For all, $w_{DFL} = 10$. Right: Results on the sum9 problem, varying the existential aggregator.

	Universal aggregation				Existential aggregation				
	Accuracy	cons%	cucons%	cuant%	Accuracy	cons%	cucons%	cuant%	
A_{T_G}	86.6	0.37	0.65	0.45	E_{S_G}	95.3	0.38	0.60	0.66
$A_{\log T_P}$	96.3	0.14	0.53	0.96	E_{S_P}	96.1	0.15	0.68	0.94
$A_{T_{LK}}$	78.2	0.00	0.94	0.99	$E_{S_{LK}}$	95.9	0.17	0.76	0.87
$A_{T_Y, p = 1.5}$	79.5	0.00	0.77	0.98	E_{S_Y}	95.9	0.21	0.64	0.85
$A_{T_Y, p = 2}$	83.3	0.00	0.83	0.98	E_{S_Y}	96.3	0.27	0.59	0.81
$A_{T_Y, p = 20}$	84.0	0.00	0.81	0.98	E_{S_Y}	96.3	0.37	0.62	0.70
$A_{GME, p = 1.5}$	96.1	0.43	0.43	0.76	E_{G_M}	96.9	0.29	0.13	0.95
A_{RMSE}	96.2	0.46	0.42	0.72	E_{G_M}	96.7	0.29	0.14	0.94
$A_{GME, p = 20}$	95.5	0.45	0.38	0.70	E_{G_M}	96.4	0.39	0.64	0.70
$A_{T_{Nm}}$	79.8	0.34	0.50	0.58	$E_{S_{Nm}}$	95.5	0.31	0.58	0.78

Table 9

Results on the sum9 problem, varying the t-norm and t-conorm together.

	Accuracy	cons%	cucons%	cuant%
T_G	20.4	0.47	0.15	0.90
T_P	20.3	0.46	0.16	0.92
T_{LK}	97.0	0.31	0.15	0.93
$T_Y, p = 1.5$	97.0	0.29	0.11	0.96
$T_Y, p = 2$	96.8	0.30	0.11	0.96
$T_Y, p = 20$	94.9	0.66	0.16	0.86
T_{Nm}	96.9	0.31	0.17	0.93

contributes to the learning process. Throughout this section, we fix the universal aggregation operator to the log-product aggregator $A_{\log T_P}$, the existential aggregation operator to the generalized mean E_{G_M} with $p = 1.5$, the conjunction and disjunction to T_Y and S_Y , also with $p = 1.5$, and the implication to the Reichenbach-sigmoidal implication with $s = 9$ and $b = -0.5$. We select these because of their promise in initial experiments.

10.2.1. Aggregation

Table 8 shows the results when varying the universal aggregator in the same problem, and when varying the existential aggregator in the sum9 problem. The log-product operator with $w_{DFL} = 10$ is the best universal aggregator, with A_{RMSE} trailing behind it slightly. Other generalized mean errors are also effective. We also see that the single-passing aggregators (minimum and Nilpotent minimum aggregators) and aggregators that vanish on a large part of their domain (Yager-based and Nilpotent minimum) all perform poorly. However, curiously the Yager-based aggregators have very high $cucons\%$ and $cuant\%$.

The generalized means have the best result for existential aggregation, with $p = 1.5$ performing best. They manage to properly select the inputs that make the existential quantifier true by softly increasing the largest inputs. Unlike with universal aggregation, the Yager existential aggregator is also a decent choice. The maximum aggregator and Nilpotent maximum aggregator only somewhat outperform the supervised baseline, however, which again is due to them being single-passing.

10.2.2. Conjunction and disjunction

In Table 9, we compare different t-norms together with their corresponding t-conorms. Here, it is the operators that vanish on a large part of their domain that work best, namely Yager t-norms and the Nilpotent minimum. These seem to work much better than when used in aggregation since the amount of inputs is much smaller, reducing the probability that the derivative vanishes. The product t-norm and Gödel t-norm, which corresponds to weak conjunction and disjunction, seem to do perform very poorly. In this table, there is a clear relation where lower cons% seem to perform better. It is likely

Table 10
Results on the same problem for several symmetric configurations with either S-implications or R-implications.

	S-Implications				R-Implications			
	Accuracy	cons%	$cu_{\text{cons}}\%$	$cu_{\text{ant}}\%$	Accuracy	cons%	$cu_{\text{cons}}\%$	$cu_{\text{ant}}\%$
I_{KD}	95.7	0.08	0.82	0.98	I_G	90.5	1	0.05
I_{RC}	96.3	0.09	0.73	0.98	I_{GG}	93.6	0.99	0.00
I_{LK}	96.0	0.5	0.07	0.93	I_{LK}	96.0	0.5	0.07
$I_Y, p = 1.5$	96.1	0.14	0.77	0.96	I_{RY}	96.1	0.14	0.14
$I_Y, p = 2$	96.2	0.13	0.82	0.97	I_{RY}	95.7	0.13	0.64
$I_Y, p = 20$	95.1	0.57	0.48	0.98	I_{RY}	96.0	0.14	0.65
I_{FD}	96.1	0.19	0.68	0.91	I_{FD}	96.1	0.19	0.68
$\sigma_{I_{RC}}$	96.3	0.14	0.53	0.96				

lower in the Yager t-norm since the disjunction in the consequent will very often be 1, which happens when the sum in the Yager t-norm hits the boundary.

10.2.3. Implications

In Table 10, we compare different fuzzy implications on the same problem. The Reichenbach implication and Yager S-implication work well, both having an accuracy around 97%. The Kleene Dienes and Yager R-implications surpass the baseline as well. In these experiments, the sigmoidal-Reichenbach implication, which we run with $s = 9$ and $b = -\frac{1}{2}$ performs as well as the normal Reichenbach implication. However, we find in Appendix F.3 that with the SGD algorithm, this implication outperforms the Reichenbach implication, reaching 97.3 accuracy.

As argued in Sections 8.3 and 8.5, the Gödel implication and Goguen implication have worse performance than the supervised baseline by making many incorrect modus ponens inferences. While the derivatives of I_{LK} and I_G only differ in that I_G disables the derivatives with respect to negated antecedent, I_{LK} performs among the better while I_G is the worst test implication, suggesting that the derivatives with respect to the negated antecedent are required to successfully applying DFL. Note that S-implications tend to perform better than R-implications, in particular for the Gödel t-norm and the product t-norm. This could be because they inherently balance derivatives with respect to the consequent and negated antecedent by being contrapositive differentiable symmetric.

10.2.4. Additional experiments

In Appendix F.3 we investigate the parameters s and b_0 of the sigmoidal-Reichenbach implication. We find here that it is the best performing implication on the same problem when using the vanilla SGD optimizer, reaching 97.3% accuracy. Furthermore, in Appendix F.4 we investigate for the same problem what the influence of each rule is to the learning process.

We also ran the sum9 problem on the vanilla SGD optimizer with the fixed configuration from Section 10.2, which reaches 97.7% accuracy. This is significantly higher than with ADAM, confirming our findings for same problem in Section 10.2.3. Finally, we ran with the same settings using both the formulas from the same problem and the sum9 problem. This has the best accuracy we find in our experiments with **98.0%**, and confirms that adding more background knowledge increases the final performance.

10.3. Analysis

We plot the accuracy of the different configurations with respect to cu_{cons} and cu_{ant} in Figs. 17a and 17b. The blue dots represent runs on the same problem, while the red dots represent runs on the sum9 problem. Fig. 17b shows a positive correlation, suggesting that it is vital to the learning process that updates going into the antecedent are correct. Although there seems to be a slight positive correlation in Fig. 17a for the same problem, it is not as pronounced. Furthermore, it seems that for the sum9, this correlation is negative instead, as the configurations with the highest accuracy have low values of cu_{cons} .

We plot all experimental values of cons% to the values of $cu_{\text{cons}}\%$ and $cu_{\text{ant}}\%$ in Figs. 17c and 17d. For both, there seems to be a negative correlation. Apparently, a larger consequent ratio decreases the correctness of the updates. In Appendix F.3 we find, when experimenting with the value of s , that this could be because for lower values of cons%, a smaller portion of the reasoning happens in the corners around $a = 0, c = 0$ and $a = 1, c = 1$, and more for instances that the agent is less certain about. Since all S-implications have strong derivatives at both these corners (Proposition 6), this phenomenon is likely present in other S-implications.

This all suggests we need to properly balance the contribution of updates to the antecedent and consequent. Since usually, as reasoned in Section 8.1, derivatives with respect to the antecedent are more common, this balance should be reflected in the experimental ratio between these updates.

10.4. Conclusions

We have run experiments on many configurations of hyperparameters to explore what works and what does not. The only well performing fully symmetric option is the product t-norm with the Reichenbach implication. If we are willing to

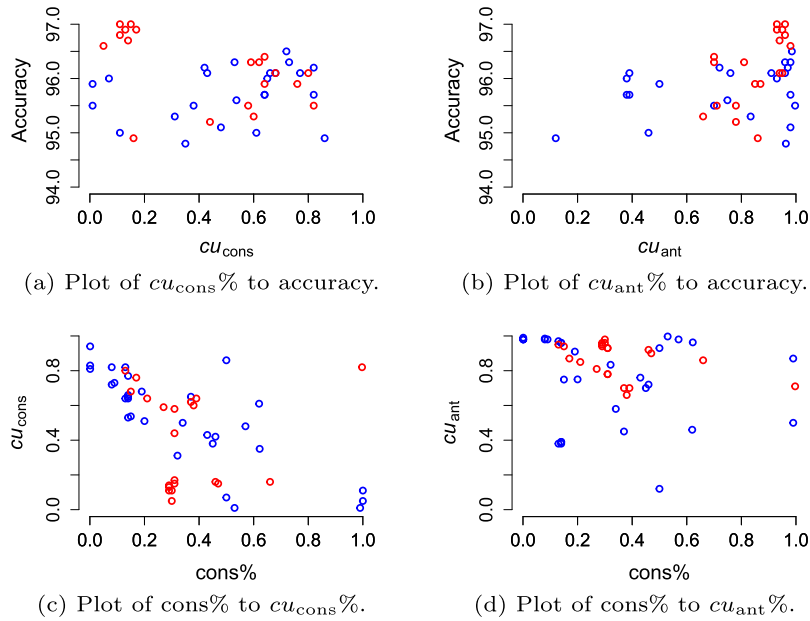


Fig. 17. We plot several of the analytical measures to find their relations. Blue dots represent runs on the same problem while red dots represent runs on the sum9 problem. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

forego symmetry, we find that the choice of the aggregators is the most important factor for performance. For universal aggregation, we recommend the log-product aggregator, while for existential quantification we recommend the generalized mean with a value of p of somewhere between 1 and 2. We found that it is especially important to choose aggregation operators that do not vanish on a large part of their domain, and that are not single-passing. In our experiments, a well tuned sigmoidal-Reichenbach implication coupled with vanilla SGD proved to be the most effective fuzzy implication. In general, we recommend choosing S-implications above R-implications. For conjunction and disjunction, we recommend tuning the Yager t-norm, although this value can be dependent on the complexity of the formulas to prevent the derivative from vanishing during the whole run.

Although Differentiable Fuzzy Logics significantly improves on the supervised baseline and is thus suited for semi-supervised learning, it is not currently competitive with state-of-the-art methods like Ladder Networks [62] which has an accuracy of 98.9% for 100 labeled pictures and 99.2% for 1000.

11. Related work

Differentiable Fuzzy Logics falls into the discipline of Statistical Relational Learning [24], which concerns models that can reason under uncertainty and learn relational structures like graphs.

11.1. Differentiable Fuzzy Logics

Special cases of DFL have been researched in several papers under different names. Logic Tensor Networks (LTN) [2,67] implements function symbols and uses neural model to interpret predicates. LTN is applied to weakly supervised learning on Scene Graph Parsing [19] and transfer learning in Reinforcement Learning [3].

Semantic-based regularization (SBR) [17] applies DFL to kernel machines. They use R-implications and the mean aggregator. Sen et al. [66] applies SBR to collective classification by predicting using a trained deep learning model, and then optimizes the DFL loss to find new truth values. This ensures predictions are consistent with the formulas during test-time.

Marra et al. [50] uses t-norm Fuzzy Logics, where the R-implication is used alongside weak disjunction. By using t-norms based on generator functions, the satisfiability computation can be simplified and generalizations of common loss functions can be found. Marra et al. [48] applies DFL to image generation. It uses the product t-norm, the log-product aggregator and the Goguen implication. By using function symbols that represent generator neural networks, they create constraints that are used to create a semantic description of an image generation problem. Rocktäschel et al. [65] uses the product t-norm and Reichenbach implication for relation extraction by using an efficient matrix embedding of the rules. Guo et al. [26] extends this to link prediction and triple classification by using a margin-based ranking loss for implications.

Demeester et al. [16] uses a regularization technique equivalent to the Łukasiewicz implication. Instead of using existing data, it finds a loss function which does not iterate over objects, yet can guarantee that the rules hold. This is very scalable, but can only model simple implications. A promising approach is using adversarial sets [52], which is a set of objects

from the domain that do not satisfy the knowledge base. These are probably the most informative objects. It uses gradient descent to find objects that *minimize* the satisfiability. The parameters of the deep learning model are then updated so that it predicts consistent with the knowledge base on this adversarial set. A benefit of this approach is that it does not have to iterate over instances that already satisfy the constraints. Adversarial sets are applied to natural language interpretation in [53]. Both papers use the Łukasiewicz implication and Gödel t-norm and t-conorm. They are not able to infer new labels on existing unlabeled data as they use artificial data, but these methods are not orthogonal and can be used jointly.

11.2. Neuro-symbolic methods using fuzzy logic operators

Posterior regularization [21,32] is a framework for weakly-supervised learning on structured data. It projects the output of a deep learning model to a ‘rule-regularized subspace’ to make it consistent with the knowledge base. This output is used as a label for the deep learning model to imitate. Unlike this paper, it does not compute derivatives over the computation of the satisfaction of the knowledge base. Marra et al. [49] and Daniele and Serafini [13] instead use gradient descent for the projection. Therefore, unlike earlier methods for posterior regularization, derivatives with respect to the operators are used. They learn relative formula weights jointly with the parameters of the deep learning model.

Arakelyan et al. [1] uses t-norms, t-conorms and existential quantification to answer queries by finding what entity embedding has the highest truth value of a given query. This search is done using gradient descent. By comparing what entity embedding best fits the optimized entity embedding, the authors can answer complex FOL queries. The authors either use the product or the Gödel t-norms.

Another recent work which employs fuzzy logic operators in a neuro-symbolic setting is Logical Neural Networks [63]. This work stands orthogonal to our work, as the foremost distinction is that they employ logics on the low-level (i.e., logical connectives as neurons and neural activation functions) while we employ it on the higher level (i.e., in defining the loss function). They limit their work to the propositional level for simplification purposes, although they argue that extending it to relational level is straightforward.

∂ ILP [20] is a differentiable inductive logic programming that uses the product t-norm and t-conorm to do differentiable inference. The Neural Theorem Prover [64] does differentiable proving of queries and combines different proof paths using the Gödel t-norm and t-conorm. Šourek et al. [70] also introduces a method for differentiable query proving, with learnable weights for formulas. They use operators inspired by fuzzy logic and transformed by the sigmoid function.

There is a vast literature on Fuzzy Neural Networks [35,36,44] that replace standard neural network neurons with neurons based on fuzzy logic. Some neurons use fuzzy logic operators which are differentiated through if the networks are trained using backpropagation.

11.3. Differentiable probabilistic logics

Some approaches use probabilistic logics instead of fuzzy logics and interpret predicates probabilistically. As deep learning classifiers can model probability distributions, probabilistic logics could be a more natural choice than fuzzy logics. DeepProbLog [46] is a probabilistic logic programming language with neural predicates that compute the probabilities of ground atoms. It supports automatic differentiation which can be used to back-propagate from the loss at a query predicate to the deep learning models that implement the neural predicates, similar to DFL. It also supports probabilistic rules which can handle exceptions to rules. We compare another differentiable probabilistic logic called Semantic Loss [74] in Appendix E and show similarities between it and DFL using operators based on the product t-norm. This similarity suggests that many practical problems that DPFL has are also present in Semantic Loss. They apply Semantic Loss to MNIST semi-supervised learning with a different knowledge base than ours. As inference is exponential in the size of the grounding for probabilistic logics, both approaches use an advanced compilation technique [14] to make inference feasible for larger problems.

12. Discussion

This paper presented theoretical results of Differentiable Fuzzy Logics operators and then evaluated their behavior on semi-supervised learning. We now discuss some problems with deploying solutions using DFL.

DFL can be seen as a form of multi-objective optimization [33]. In the DFL loss (Equation (11)) we sum up the valuations of different formulas, each of which is a separate objective. Each of these objectives can be weighted differently, resulting in wildly varying loss landscapes. Having so many objectives requires significant hyperparameter tuning. A method capable of learning relative formula weights jointly like [49,13,70], could solve this problem.

A second challenge is related to the class imbalance problem [37,9]. We argued in Section 8.1 that for a significant portion of common-sense background knowledge, the modus tollens case is by far the most common. Our same problem indeed showed that most well-performing implications have a far larger derivative with respect to the negated antecedent than to the consequent. This imbalance will only increase for more complex problems. However, simply removing derivatives with respect to the antecedent does not seem to be the solution. A reason for this could be that those are usually correct, unlike derivatives with respect to the consequent. In fact, we found in Appendix F.4 that the formula in which the digits are

in the antecedent performs better on its own than the formula in which the digits are in the consequent, even though the model could not learn from any new positive examples.

Although we have focused on experimenting with the accuracy of the derivatives of the implication, it should be noted that the derivatives of the disjunction operator make a choice as well. For example, if the agent observes a walking object and the supervisor knows that only humans and animals can walk, how is the supervisor supposed to choose whether it is a human or an animal? Here, similar imbalances exist in the different possible classes: There might be more images of humans than of animals.

Further, we pose whether it is more important that we choose operators based on the performance on the task at hand, or based on its logical properties. The best configuration uses operators based on both the product and Yager t-norms. The product t-norm is the only viable symmetric choice in our experiments. The largest benefit of a ‘symmetric’ choice of operators is that the truth value of formulas that are logically equivalent in classical logic will be equal. This makes it easier to analyze how the background knowledge will behave and does not require putting it in a particular form.

As a final remark, noteworthy is the interpretation of truth values. As aforementioned, the logic we use is fuzzy logic which was originally aimed to address logical reasoning in the presence of vagueness rather than probabilistic uncertainty. The truth values derived using fuzzy operators, therefore, are not probabilistic (see, for instance, [28, p. 4]).¹¹

However, since a considerably large amount of problems addressed by machine learning literature is probabilistic (as it has mathematical origins in statistics), the classification task used in our running example is also of probabilistic origin. With this choice we also aimed to respect the recent literature: Applications of fuzzy operators on a general set of problems which are not necessarily fuzzy are not uncommon in neuro-symbolic AI. Examples include [67], [63], [1], among others cited in Section 11.

13. Conclusion

We analyzed Differentiable Fuzzy Logics in order to understand how reasoning using logical formulas behaves in a differentiable setting. We examined how the properties of a large amount of different operators affect DFL. We have found substantial differences between the properties of a large number of such Differentiable Fuzzy Logics operators, and we showed that many of them, including some of the most popular operators, are highly unsuitable for use in a differentiable learning setting. By analyzing aggregation functions, we found that the log-product aggregator and the RMSE aggregator have convenient connections to both fuzzy logic and machine learning and can deal with outliers. Next, we analyzed conjunction and disjunction operators and found several strong candidates. In particular, the Gödel t-norm and t-conorm are a simple choice, and that the Yager t-norm and the product t-conorm have intuitive derivatives.

We noted an interesting imbalance between derivatives with respect to the negated antecedent and the consequent of the implication. Because the modus tollens case is much more common, we conclude that a large part of the useful inferences on the MNIST experiments is made by decreasing the antecedent, or by ‘modus tollens reasoning’. Furthermore, we found that derivatives with respect to the consequent often increase the truth value of something that is false as the consequent is false in the majority of times. Therefore, we argue that ‘modus tollens reasoning’ should be embraced in future research. As a possible solution to problems caused by this imbalance, we introduced a smoothed fuzzy implication called the Reichenbach-sigmoidal implication.

Experimentally, we found that the product t-norm is the only t-norm that can be used as a base for all choices of operators. The product t-conorm and the Reichenbach implication have intuitive derivatives that correspond to inference rules from classical logic, and the log-product aggregator is the most effective universal aggregation operator.

In order to gain the largest improvements over a supervised baseline however, we had to abandon the normal symmetric configurations of norms, where t-norms, t-conorms, implications and the aggregation operators satisfy the usual algebraic relations. Instead, we had to resort to non-symmetric configurations where operators based on different t-norms are combined. The Reichenbach-sigmoidal implication performs best in our experiments. Its hyperparameters can be tweaked to decrease the imbalance of the derivatives with respect to the negated antecedent and consequent. For existential quantification, we found that the general mean error performs best, and for conjunction and disjunction the family of Yager t-norms and the Nilpotent minimum has the highest final accuracy.

We believe a proper empirical comparison of different methods that introduce background knowledge through logic could be useful to properly understand the details, performance, possible applications and challenges of each method. Secondly, we believe more work is required in using background knowledge to help deep models train on real-world problems. One research direction would be to develop methods that can properly deal with exceptions. An approach in which formula importance weights can be learned could be used to distinguish between relevant and irrelevant formulas in the background knowledge, and probabilistic instead of fuzzy logics could be a more natural fit. Lastly, additional research on the vast space of fuzzy logic operators might find more properties that are useful in DFL.

¹¹ Indeed, when reasoning about belief, using fuzzy logic semantics instead of probabilistic logic semantics straight out-of-the-box, can yield undesirable results: Consider an event a where $p(a)$ (probability of a) is 0.5. Now consider a disjunction, where $p(a \vee a)$ has the value 0.5. However, in Łukasiewicz logic, $S(a, a)$ would yield 1.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

We sincerely thank all the anonymous reviewers whose comments substantially improved the content and the quality of this manuscript. This work is partly funded by the MaestroGraph research programme with project number 612.001.552, which is financed by the Dutch Research Council (NWO). Erman Acar is generously funded by the Hybrid Intelligence Project which is financed by the Dutch Ministry of Education, Culture and Science with project number 024.004.022. This work is also supported by the DAS-5 distributed supercomputer [4].

Appendix A. Background on fuzzy logic operators

In this section, we will introduce the semantics of the fuzzy operators \otimes (t-norm), \oplus (t-conorm) and \neg (negation) that are used to connect truth values of fuzzy predicates, and the semantics of the \forall quantifier. We follow [38] in this section and refer to it for proofs and additional results.

A.1. Fuzzy negation

The functions that are used to compute the negation of a truth value of a formula are called *fuzzy negations*.

Definition 13. A *fuzzy negation* is a decreasing function $N : [0, 1] \rightarrow [0, 1]$ so that $N(1) = 0$ and for all x , $N(N(x)) \geq x$ [11]. N is called *strict* if it is strictly decreasing and continuous, and *strong* if for all $a \in [0, 1]$, $N(N(a)) = a$.

A consequence of these conditions is that $N(0) = 1$. Throughout the paper we also use N to refer to the classical negation $N(a) = 1 - a$.

A.2. Triangular norms

The functions that are used to compute the conjunction of two truth values are called *t-norms*. For a rigorous overview, see [40].

Definition 14. A *t-norm* (triangular norm) is a function $T : [0, 1]^2 \rightarrow [0, 1]$ that is commutative and associative, and

1. *Monotonicity*: For all $a \in [0, 1]$, $T(a, \cdot)$ is increasing and
2. *Neutrality*: For all $a \in [0, 1]$, $T(1, a) = a$.

The phrase ' $T(a, \cdot)$ is increasing' means that whenever $0 \leq b_1 \leq b_2 \leq 1$, then $T(a, b_1) \leq T(a, b_2)$.

Definition 15. A t-norm T can have the following properties:

- a) *Continuity*: A continuous t-norm is continuous in both arguments.
- b) *Left-continuity*: A left-continuous t-norm is left-continuous in both arguments. That is, for all $a, b \in [0, 1]$, $\lim_{x \rightarrow a^-} T(x, b) = T(a, b)$ (the limit of $T(x, b)$ as x increases and approaches a is a).
- c) *Idempotency*: An idempotent t-norm has the property that for all $a \in [0, 1]$, $T(a, a) = a$.
- d) *Strict-monotony*: A strictly monotone t-norm has the property that for all $a \in (0, 1)$, $T(a, \cdot)$ is strictly increasing.
- e) *Strict*: A strict t-norm is continuous and strictly monotone.

Table 1 shows the four basic t-norms and two other t-norms of interest alongside their properties.

A.3. Triangular conorms

The functions that are used to compute the disjunction of two truth values are called *t-conorms* or *s-norms*.

Definition 16. A *t-conorm* (triangular conorm, also known as s-norm) is a function $S : [0, 1]^2 \rightarrow [0, 1]$ that is commutative and associative, and

1. *Monotonicity*: For all $a \in [0, 1]$, $S(a, \cdot)$ is increasing and

2. *Neutrality*: For all $a \in [0, 1]$, $S(0, a) = a$.

T-conorms are obtained from t-norms using De Morgan's laws from classical logic, i.e. $p \vee q = \neg(\neg p \wedge \neg q)$. Therefore, if T is a t-norm and N_C the strong negation, T 's N_C -dual S is calculated using

$$S(a, b) = 1 - T(1 - a, 1 - b) \tag{A.1}$$

Table 2 shows several common t-conorms derived using Equation (A.1) and the t-norms from Table 1, alongside the same optional properties as those for t-norms in Definition 15.

A.4. Aggregation operators

The functions that are used to compute quantifiers like \forall and \exists are aggregation operators [45].

Definition 17. An *aggregation operator* [10] is a function $A : \bigcup_{n \in \mathbb{N}} [0, 1]^n \rightarrow [0, 1]$ that is non-decreasing with respect to each argument, and for which $A(0, \dots, 0) = 0$ and $A(1, \dots, 1) = 1$.

Aggregation operators are *variadic functions* which are functions that are defined for any sequence of arguments. For this reason we will often use the notation $\mathbf{A}_{i=1}^n x_i := A(x_1, \dots, x_n)$. Table 3 shows some common aggregation operators that we will talk about. Furthermore, we will only consider *symmetric* aggregation operators, that are invariant to permutation of the sequence.

The \forall quantifier is interpreted as the conjunction over all arguments x . Therefore, we can extend a t-norm T from 2-dimensional inputs to n -dimensional inputs as they are commutative and associative [40]:

$$\begin{aligned} A_T() &= 0 \\ A_T(x_1, x_2, \dots, x_n) &= T(x_1, A_T(x_2, \dots, x_n)) \end{aligned} \tag{A.2}$$

These operators are a straightforward choice for modeling the \forall quantifier, as they can be seen as a series of conjunctions. All operators constructed in this way are *symmetric* aggregation operators, for which the output value is the same for every ordering of its arguments. This generalizes commutativity.

We can do the same for a t-conorm S to model the \exists quantifier:

$$\begin{aligned} E_S() &= 0 \\ E_S(x_1, x_2, \dots, x_n) &= S(x_1, E_S(x_2, \dots, x_n)) \end{aligned} \tag{A.3}$$

A.5. Fuzzy implications

The functions that are used to compute the truth value of $p \rightarrow q$ are called fuzzy implications. p is called the *antecedent* and q the *consequent* of the implication. We follow [38] and refer to it for details and proofs.

Definition 18. A *fuzzy implication* is a function $I : [0, 1]^2 \rightarrow [0, 1]$ so that for all $a, c \in [0, 1]$, $I(\cdot, c)$ is decreasing, $I(a, \cdot)$ is increasing and for which $I(0, 0) = 1$, $I(1, 1) = 1$ and $I(1, 0) = 0$.

From this definition follows that $I(0, 1) = 1$.

Definition 19. Let N be a fuzzy negation. A fuzzy implication I satisfies

- a) *left-neutrality (LN)* if for all $c \in [0, 1]$, $I(1, c) = c$;
- b) the *exchange principle (EP)* if for all $a, b, c \in [0, 1]$, $I(a, I(b, c)) = I(b, I(a, c))$;
- c) the *identity principle (IP)* if for all $a \in [0, 1]$, $I(a, a) = 1$;
- d) *N-contrapositive symmetry (CP)* if for all $a, c \in [0, 1]$, $I(a, c) = I(N(c), N(a))$;
- e) *N-left-contrapositive symmetry (L-CP)* if for all $a, c \in [0, 1]$, $I(N(a), c) = I(N(c), a)$;
- f) *N-right-contrapositive symmetry (R-CP)* if for all $a, c \in [0, 1]$, $I(a, N(c)) = I(c, N(a))$.

All these statements generalize a law from classical logic. *Left neutrality* generalizes $(1 \rightarrow p) \equiv p$, the *exchange principle* generalizes $p \rightarrow (q \rightarrow r) \equiv q \rightarrow (p \rightarrow r)$, and the *identity principle* generalizes that $p \rightarrow p$ is a tautology. Furthermore, *N-contrapositive symmetry* generalizes $p \rightarrow q \equiv \neg q \rightarrow \neg p$, *N-left-contrapositive symmetry* generalizes $\neg p \rightarrow q \equiv \neg q \rightarrow p$ and *N-right-contrapositive symmetry* generalizes $p \rightarrow \neg q \equiv q \rightarrow \neg p$.

A.5.1. S-implications

In classical logic, the (material) implication is defined as follows:

$$p \rightarrow q = \neg p \vee q$$

Using this definition, we can use a t-conorm S and a fuzzy negation N to construct a fuzzy implication.

Definition 20. Let S be a t-conorm and N a fuzzy negation. The function $I_{S,N} : [0, 1]^2 \rightarrow [0, 1]$ is called an (S, N) -implication and is defined for all $a, c \in [0, 1]$ as

$$I_{S,N}(a, c) = S(N(a), c). \quad (\text{A.4})$$

If N is a strong fuzzy negation, then $I_{S,N}$ is called an S -implication (or strong implication).

As we only consider the strong negation N_C , we omit the N and use I_S to refer to I_{S,N_C} .

All S -implications I_S are fuzzy implications and satisfy LN, EP and R-CP. Additionally, if the negation N is strong, it satisfies CP and if, in addition, it is strict, it also satisfies L-CP. In Table 4 we show several S -implications that use the strong fuzzy negation N_C and the common t-conorms (Table 2). Note that S -implications are rotations of the t-conorms.

A.5.2. R-implications

R-implications are another way of constructing implication operators. They are the standard choice in t-norm fuzzy logics.

Definition 21. Let T be a t-norm. The function $I_T : [0, 1]^2 \rightarrow [0, 1]$ is called an R -implication and defined as

$$I_T(a, c) = \sup\{b \in [0, 1] \mid T(a, b) \leq c\} \quad (\text{A.5})$$

The *supremum* of a set A , denoted $\sup\{A\}$, is the lowest upper bound of A . All R -implications are fuzzy implications, and all satisfy LN, IP and EP. T is a left-continuous t-norm if and only if the supremum can be replaced with the maximum function. Note that if $a \leq c$ then $I_T(a, c) = 1$. We can see this by looking at Equation (A.5). The largest value for b possible is 1, since then, using the *neutrality* property of t-norms, $T(a, 1) = a \leq c$.

Table 5 shows the R -implications created from the common T-norms. Note that I_{LK} and I_{FD} appear in both tables: They are both S -implications and R -implications.

Appendix B. Implementation of Differentiable Fuzzy Logics

The computation of the satisfaction is shown in pseudocode form in Algorithm 1. By first computing the dictionary g that contains truth values for all ground atoms,¹² we can reduce the amount of forward passes through the computations of the truth values of the ground atoms that are required to compute the satisfaction.

This algorithm can fairly easily be parallelized for efficient computation on a GPU by noting that the individual terms that are aggregated over in lines 12 and 14 (the different *instances* of the quantifiers) are not dependent on each other. By noting that formulas are in prenex normal form, we can set up the dictionary g using tensor operations so that the recursion has to be done only once for each formula. This can be done by applying the fuzzy operators elementwise over vectors of truth values instead of a single truth value, where each element of the vector represents a variable assignment.

The complexity of this computation then is $O(|\mathcal{K}| \cdot P \cdot b^d)$, where \mathcal{K} is the set of formulas, P is the amount of predicates used in each formula and d is the maximum depth of nesting of universal quantifiers in the formulas in \mathcal{K} (known as the *quantifier rank*). This is exponential in the amount of quantifiers, as every object from the constants C has to be iterated over in lines 12 and 14, although as mentioned earlier this can be mitigated somewhat using efficient parallelization. Still, computing the valuation for transitive rules (such as $\forall x y, z \ Q(x, z) \otimes R(z, y) \rightarrow P(x, y)$) will for example be far more demanding than for antisymmetry formulas (such as $\forall x, y \ P(x, y) \rightarrow \neg P(y, x)$).

Appendix C. Proofs

C.1. Single-passing

Proposition 7. Any composition of single-passing functions is also single-passing.

¹² The dictionary g could be seen as a ‘fuzzy Herbrand interpretation’, in that it assigns a truth value to all ground atoms.

Algorithm 1 Computation of the Differentiable Fuzzy Logics loss. First it computes the fuzzy Herbrand interpretation η_θ given the current embedded interpretation η_θ . This performs a forward pass through the neural networks that are used to interpret the predicates. Then it computes the valuation of each formula φ in the knowledge base \mathcal{K} , implementing Equations (3)-(8).

```

1: function  $e(\varphi, g, C, \mu)$  ▷ The valuation function computes the Fuzzy truth value of  $\varphi$ .
2:   if  $\varphi = P(x_1, \dots, x_m)$  then ▷ Find the truth value of a ground atom using the dictionary  $g$ .
3:     return  $g[P, (\mu(x_1), \dots, \mu(x_m))]$ 
4:   else if  $\varphi = \neg\phi$  then
5:     return  $N(e(\phi, g, C, \mu))$ 
6:   else if  $\varphi = \phi \otimes \psi$  then
7:     return  $T(e(\phi, g, C, \mu), e(\psi, g, C, \mu))$ 
8:   else if  $\varphi = \phi \oplus \psi$  then
9:     return  $S(e(\phi, g, C, \mu), e(\psi, g, C, \mu))$ 
10:  else if  $\varphi = \phi \rightarrow \psi$  then
11:    return  $I(e(\phi, g, C, \mu), e(\psi, g, C, \mu))$ 
12:  else if  $\varphi = \forall x \phi$  then ▷ Apply the universal aggregation operator.
13:    return  $\bigwedge_{o \in C} e(\phi, g, C, \mu \cup \{(x, o)\})$  ▷ Each assignment can be seen as an instance of  $\varphi$ .
14:  else if  $\varphi = \exists x \phi$  then
15:    return  $\bigvee_{o \in C} e(\phi, g, C, \mu \cup \{(x, o)\})$ 
16:  end if
17: end function
18:
19: procedure  $\text{DFL}(\eta_\theta, \mathcal{P}, \mathcal{K}, O, N, T, S, I, A, E)$  ▷ Computes the Differentiable Fuzzy Logics loss.
20:    $C \leftarrow o_1, \dots, o_b$  sampled from  $O$  ▷ Sample  $b$  constants to use this pass.
21:    $g \leftarrow \text{dict}()$  ▷ Collects truth values for ground atoms.
22:   for  $P \in \mathcal{P}$  do
23:     for  $o_1, \dots, o_{\alpha(P)} \in C$  do
24:        $g[P, (o_1, \dots, o_{\alpha(P)})] \leftarrow \eta_\theta(P)(o_1, \dots, o_{\alpha(P)})$  ▷ Calculate the truth values of the ground atoms.
25:     end for
26:   end for
27:   return  $\bigwedge_{\varphi \in \mathcal{K}} w_\varphi \cdot e_\theta(\varphi, g, C, \emptyset)$  ▷ Calculate valuation of the formulas  $\varphi$ . Start with an empty variable assignment. This implements Equation (11).
28: end procedure

```

Proof. We will prove this by structural induction. Let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be a single-passing function and let $x_1, \dots, x_n \in \mathbb{R}$. Then clearly $f(x_1, \dots, x_n)$ is single-passing.

Next, assume by induction that $g : \mathbb{R}^n \rightarrow \mathbb{R}$ is a composition of single-passing functions that we assume is single-passing. Let $y : \mathbb{R}^m \rightarrow \mathbb{R}$ be a single-passing function. Let Z be the set of inputs to y and define $x_i = y(Z)$. We show that the composition $g(x_1, \dots, y(Z), \dots, x_n)$ is also single-passing. For any $z \in Z$ holds that

$$\frac{\partial g(x_1, \dots, x_n)}{\partial z} = \frac{\partial g(x_1, \dots, x_n)}{\partial x_i} \frac{\partial y(Z)}{\partial z}. \quad (\text{C.1})$$

As g is single-passing, there is at most 1 number $j \in 1, \dots, n$ so that $\frac{\partial g(x_1, \dots, x_n)}{\partial x_j} \neq 0$. If there are 0, then there can also be no $k \in 1, \dots, m$ such that $\frac{\partial g(x_1, \dots, x_n)}{\partial z} \neq 0$ as $\frac{\partial g(x_1, \dots, x_n)}{\partial x_i} = 0$. If there is 1, then either $j \neq i$, which is the direct input x_j . If $j = i$, then by the assumption of $y(Z)$ being single-passing, there is at most 1 $k \in 1, \dots, m$ so that $\frac{\partial y(Z)}{\partial z} \neq 0$ and by Equation (C.1) there is at most 1 input such that $\frac{\partial g(x_1, \dots, x_n)}{\partial x} \neq 0$. We conclude that the composition $g(x_1, \dots, y(Z), \dots, x_n)$ is single-passing. \square

C.2. Nonvanishing fractions

C.2.1. Łukasiewicz aggregator

Proposition 8. The fraction of inputs $x_1, \dots, x_n \in [0, 1]$ for which the derivative of A_{TLU} is nonvanishing is $\frac{1}{n!}$.

Proof. Consider standard uniformly distributed random variables $x_1, \dots, x_n \sim U(0, 1)$. The sum $Y = \sum_{i=1}^n x_i$ is Irwin-Hall distributed [34,29]. The cumulative density function of this distribution is

$$F_Y(y) = \frac{1}{n!} \sum_{k=0}^{\lfloor y \rfloor} (-1)^k \binom{n}{k} (y-k)^{n-1}, \quad (\text{C.2})$$

where $\lfloor \cdot \rfloor$ is the floor function. The derivative of A_{TLU} is nonvanishing when $Y > n-1$, or equivalently as the Irwin-Hall distribution is symmetric, when $Y < 1$. Using F_Y gives $F_Y(1) = \frac{1}{n!} ((-1)^0 \binom{n}{0} (1-0)^n + (-1)^1 \binom{n}{1} 1(1-1)^n) = \frac{1}{n!}$. \square

This result is the same for the bounded sum aggregator, as for that the condition is that $Y < 1$.

C.2.2. Yager aggregator

Proposition 9. The fraction of inputs $x_1, \dots, x_n \in [0, 1]$ for which the derivative of A_{T_Y} with $p = 2$ is nonvanishing is $\frac{\pi^{\frac{n}{2}}}{2^n \cdot \Gamma(\frac{1}{2}n + \frac{1}{2})}$, where Γ is the Gamma function.

Proof. The points $x_1, \dots, x_n \in \mathbb{R}$ for which $\sum_{i=1}^n (1 - x_i)^2 < 1$ hold describes the volume of an n -ball¹³ with radius 1. This volume is found by [5](p.5):

$$V(n) = \frac{\pi^{\frac{n}{2}}}{\Gamma(\frac{1}{2}n + 1)}. \tag{C.3}$$

We are interested in the part of this volume where $x_1, \dots, x_n \in [0, 1]$, that is, those in a single orthant¹⁴ of the n -ball. The amount of orthants in which an n -ball lies is 2^n .¹⁵ Thus, the volume of the part of the n -ball where $x_1, \dots, x_n \in [0, 1]$ is $\frac{V(n)}{2^n} = \frac{\pi^{\frac{n}{2}}}{2^n \cdot \Gamma(\frac{1}{2}n + \frac{1}{2})}$. As the total volume of points lying in $[0, 1]^n$ is 1, this is also the fraction of points for which the derivative of A_{T_Y} with $p = 2$ is nonvanishing. \square

C.2.3. Nilpotent aggregator

Proposition 10. The fraction of inputs $x_1, \dots, x_n \in [0, 1]$ for which the derivative of $A_{T_{nM}}$ is nonvanishing is $\frac{1}{2^{n-1}}$.

Proof. Consider n standard uniformly distributed random variables $x_1, \dots, x_n \sim U(0, 1)$. We are interested in the probability that $x_{(1)} + x_{(2)} > 1$, where $x_{(k)}$ is the k -th smallest sample (known as the k -th order statistic [15]). Weisberg [73] derives the cumulative density function for linear combinations of standard uniform order statistics. Let $k_0 = 0 < k_1 < \dots < k_S \leq k_n$ be S integers indicating the coefficients $d_i > 0$. We aim to calculate the probability that $\sum_{i=1}^S d_i x_{(i)} > v$. Let $r_s = k_s - k_{s-1}$ for all $1, \dots, S$ and let $r_{S+1} = n - k_S$. Finally, let $c_{S+1} = 0$ and $c_{(s)} = c_{(s+1)} + d_i$. m is the largest integer so that $v \leq c_{(m)}$. Then, Weisberg [73] finds that

$$P \left\{ \sum_{s=1}^S d_s x_{(k_s)} > v \right\} = \sum_{s=1}^m \frac{g_s^{(r_s-1)}(c_{(s)})}{(r_s - 1)!} \tag{C.4}$$

where $g_s^{(i)}$ is the i -th order derivative of

$$g_s(c) = \frac{(c - v)^n}{c \prod_{i=1, i \neq s}^{S+1} (c - c_{(i)})^{r_i}} \tag{C.5}$$

Filling this in for our case, we find that $S = 2$, where $k_1 = 1, k_2 = 2$ as $d_1 = d_2 = 1$. Therefore, $r_1 = r_2 = 1$ and $r_3 = n - 2$ and $c_{(1)} = 2, c_{(2)} = 1$ and $c_{(3)} = 0$. The largest integer m so that $1 \leq c_{(m)}$ is 2. Filling this in, we find that

$$P \{x_{(1)} + x_{(2)} > 1\} = \frac{g_1^{(1-1)}(2)}{(1 - 1)!} + \frac{g_2^{(1-1)}(1)}{(1 - 1)!} \tag{C.6}$$

$$= \frac{(2 - 1)^n}{2(2 - 1)^1(2 - 0)^{n-2}} + \frac{(1 - 1)^n}{1(1 - 2)(1 - 0)^{n-2}} = \frac{1}{2^{n-1}} \quad \square \tag{C.7}$$

C.3. Implications

Proposition 11. If a fuzzy implication I is N_C -contrapositive symmetric, where N_C is the strong negation, it is also contrapositive differentiable symmetric.

Proof. Say we have an implication I that is N_C -contrapositive symmetric. Because I is N_C -contrapositive symmetric, $I(1 - c, 1 - a) = I(a, c)$. Thus, $-\frac{\partial I(1-c, 1-a)}{\partial 1-c} = -\frac{\partial I(a, c)}{\partial 1-c} = \frac{\partial I(a, c)}{\partial c}$. \square

¹³ An n -ball is the generalization of the concept of a ball to any dimension and is the region enclosed by a $n - 1$ hypersphere. For example, the 3-ball (or ball) is surrounded by a sphere (or 2-sphere). Similarly, the 2-ball (or disk) is surrounded by a circle (or 1-sphere). A hypersphere with radius 1 is the set of points which are at a distance of 1 from its center.

¹⁴ An orthant in n dimensions is a generalization of the quadrant in two dimensions and the octant in three dimensions.

¹⁵ To help understand this, consider $n = 2$. The 1-ball is the circle with center $(0, 0)$. The area of this circle is evenly distributed over the four quadrants.

Proposition 12. If an implication I is left-neutral, then $\frac{\partial I(1,c)}{\partial c} = 1$. If, in addition, I is contrapositive differentiable symmetric, then $-\frac{\partial I(a,0)}{\partial a} = 1$.

Proof. First, assume I is left-neutral. Then for all $c \in [0, 1]$, $I(1, c) = c$. Taking the derivative with respect to c , it is clear that $\frac{\partial I(1,c)}{\partial c} = 1$. Next, assume I is contrapositive differentiable symmetric. Then, $\frac{\partial I(1,c)}{\partial c} = -\frac{\partial I(1-c,1-1)}{\partial 1-c} = -\frac{\partial I(1-c,0)}{\partial 1-c} = 1$. As $1 - c \in [0, 1]$, $-\frac{\partial I(a,0)}{\partial a} = 1$. \square

Appendix D. Derivations of used functions

D.1. p -error aggregators

The unbounded Yager aggregator is

$$A_{UY}(x_1, \dots, x_n) = 1 - \left(\sum_{i=1}^n (1 - x_i)^p \right)^{\frac{1}{p}}, \quad p \geq 0. \tag{D.1}$$

We can do an affine transformation $w \cdot A_{UY}(x_1, \dots, x_n) - h$ on this function to ensure the boundary conditions in Definition 17 hold.

$$w \cdot A_{UY}(0, \dots, 0) - h = 0 \tag{D.2}$$

$$w \cdot A_{UY}(1, \dots, 1) - h = 1 \tag{D.3}$$

Solving Equation (D.2) and (D.3) for h , we find

$$\begin{aligned} w \cdot \left(1 - \left(\sum_{i=1}^n (1 - 0)^p \right)^{\frac{1}{p}} \right) - h &= 0 \\ w \cdot \left(1 - n^{\frac{1}{p}} \right) - h &= 0 \\ h &= w - w \cdot \sqrt[p]{n}, \end{aligned} \tag{D.4}$$

$$\begin{aligned} w \cdot \left(1 - \left(\sum_{i=1}^n (1 - 1)^p \right)^{\frac{1}{p}} \right) - h &= 1 \\ w \cdot (1 - 0) - h &= 1 \\ h &= w - 1. \end{aligned} \tag{D.5}$$

Equating (D.4) and (D.5) and solving for h , we find

$$\begin{aligned} w - w \cdot \sqrt[p]{n} &= w - 1 \\ w &= \frac{1}{\sqrt[p]{n}}. \end{aligned} \tag{D.6}$$

And so $h = \frac{1}{\sqrt[p]{n}} - 1$. Filling in and simplifying we find

$$\begin{aligned} A_{pE}(x_1, \dots, x_n) &= \frac{1}{\sqrt[p]{n}} \cdot \left(1 - \left(\sum_{i=1}^n (1 - x_i)^p \right)^{\frac{1}{p}} \right) - \left(\frac{1}{\sqrt[p]{n}} - 1 \right) \\ &= 1 - \frac{1}{\sqrt[p]{n}} \cdot \left(\sum_{i=1}^n (1 - x_i)^p \right)^{\frac{1}{p}} \\ &= 1 - \left(\frac{1}{n} \sum_{i=1}^n (1 - x_i)^p \right)^{\frac{1}{p}} \end{aligned} \tag{D.7}$$

Similarly for the t-conorm $A_{UY\mathcal{S}}(x_1, \dots, x_n) = \left(\sum_{i=1}^n x_i^p \right)^{\frac{1}{p}}$, $p \geq 0$.

$$\begin{aligned}
 w \cdot \left(\sum_{i=1}^n 0^p \right)^{\frac{1}{p}} - h &= 0 \\
 w \cdot 0 - h &= 0 \\
 h &= 0
 \end{aligned} \tag{D.8}$$

$$\begin{aligned}
 w \cdot \left(\sum_{i=1}^n 1^p \right)^{\frac{1}{p}} - h &= 1 \\
 w \cdot n^{\frac{1}{p}} - h &= 1 \\
 w &= \frac{1}{\sqrt[p]{n}}
 \end{aligned} \tag{D.9}$$

$$A_{p\text{-MEAN}}(x_1, \dots, x_n) = \left(\frac{1}{n} \sum_{i=1}^n x_i^p \right)^{\frac{1}{p}}, \quad p \geq 0. \tag{D.10}$$

D.2. Sigmoidal functions

In Machine Learning, the logistic function or sigmoid function $\sigma(x) = \frac{1}{1+e^{-x}}$ is a common activation function [25](p.65-66). This inspired [71] to introduce parameterized families of aggregation functions they call Max-Sigmoid activation functions:

$$A'_{\sigma+\wedge}(x_1, \dots, x_n) = \sigma \left(s \cdot \left(\sum_{i=1}^n x_i - n + 1 + b_0 \right) \right), \quad A'_{\sigma+\vee}(x_1, \dots, x_n) = \sigma \left(s \cdot \left(\sum_{i=1}^n x_i + b_0 \right) \right) \tag{D.11}$$

We generalize this transformation for any function $f : [0, 1]^n \rightarrow \mathbb{R}$ that is symmetric and increasing:

$$A'_{\sigma_f}(x_1, \dots, x_n) = \sigma(s \cdot (f(x_1, \dots, x_n) + b_0)) \tag{D.12}$$

This cannot be an aggregation function according to Definition 17 as $\sigma \in (0, 1)$ and so the boundary conditions $A'_\sigma(1, \dots, 1) = 1$ and $A'_\sigma(0, \dots, 0)$ do not hold. We can solve this by adding two linear parameters w and h , redefining σ_f as

$$\sigma_f(x_1, \dots, x_n) = w \cdot \sigma(s \cdot (f(x_1, \dots, x_n) + b_0)) - h \tag{D.13}$$

For this, we need to make sure the lowest value of f on the domain $[0, 1]^n$ maps to 0 and the highest to 1. For this, we define $\inf_f = \inf\{f(x_1, \dots, x_n) | x_1, \dots, x_n \in [0, 1]\}$ and $\sup_f = \sup\{f(x_1, \dots, x_n) | x_1, \dots, x_n \in [0, 1]\}$. This gives the following system of equations

$$A_\sigma(0, \dots, 0) = w \cdot \sigma(s \cdot (\inf_f + b_0)) - h = 0 \tag{D.14}$$

$$A_\sigma(1, \dots, 1) = w \cdot \sigma(s \cdot (\sup_f + b_0)) - h = 1 \tag{D.15}$$

First solve both equations for w , starting with Equation (D.14):

$$\begin{aligned}
 w \cdot \sigma(s \cdot (\inf_f + b_0)) - h &= 0 \\
 \frac{1}{1 + e^{-s \cdot (\inf_f + b_0)}} &= \frac{h}{w} \\
 w &= h \cdot (1 + e^{-s \cdot (\inf_f + b_0)})
 \end{aligned} \tag{D.16}$$

Likewise for Equation (D.15):

$$\begin{aligned}
 w \cdot \sigma(s \cdot (\sup_f + b_0)) - h &= 1 \\
 \frac{1}{1 + e^{-s \cdot (\sup_f + b_0)}} &= \frac{1 + h}{w} \\
 w &= (1 + h) \cdot (1 + e^{-s \cdot (\sup_f + b_0)})
 \end{aligned} \tag{D.17}$$

Now we can solve for h by equating Equations (D.17) and (D.16). Then

$$(1 + h) \cdot 1 + e^{-s \cdot (\sup f + b_0)} = h \cdot 1 + e^{-s \cdot (\inf f + b_0)}$$

$$h = \frac{1 + e^{-s \cdot (\sup f + b_0)}}{1 + e^{-s \cdot (\inf f + b_0)} - 1 + e^{-s \cdot (\sup f + b_0)}}$$

We thus get the following formula:

$$A_\sigma(x_1, \dots, x_n) = \frac{1 + e^{-s \cdot (\sup f + b_0)}}{e^{-s \cdot (\inf f + b_0)} - e^{-s \cdot (\sup f + b_0)}} \cdot \left((1 + e^{-s \cdot (\inf f + b_0)}) \cdot \sigma(s \cdot (f(x_1, \dots, x_n) + b_0)) - 1 \right) \tag{D.18}$$

If f is a fuzzy logic operator of which the outputs are all in $[0, 1]$, the most straightforward choice of b_0 is $-\frac{1}{2}$. This translates the outputs of f to $[-\frac{1}{2}, \frac{1}{2}]$ and so it uses a symmetric part of the sigmoid function. For some function $f \in [0, 1]^n \rightarrow [0, 1]$, we find the following simplification, noting that the supremum of f is 1 and the infimum is 0:

$$\sigma_f(x_1, \dots, x_n) = \frac{1 + e^{-s(1-\frac{1}{2})}}{e^{s(0-\frac{1}{2})} - e^{-s(1-\frac{1}{2})}} \cdot \left((1 + e^{-s(0-\frac{1}{2})}) \cdot \sigma\left(s \cdot \left(f(x_1, \dots, x_n) - \frac{1}{2}\right)\right) - 1 \right) \tag{D.19}$$

$$= \frac{1 + e^{-\frac{s}{2}}}{e^{\frac{s}{2}} - e^{-\frac{s}{2}}} \cdot \frac{e^{\frac{s}{2}} - 1}{e^{\frac{s}{2}} - 1} \cdot \left((1 + e^{\frac{s}{2}}) \cdot \sigma\left(s \cdot \left(f(x_1, \dots, x_n) - \frac{1}{2}\right)\right) - 1 \right) \tag{D.20}$$

$$= \frac{e^{\frac{s}{2}} - e^{-\frac{s}{2}}}{(e^{\frac{s}{2}} - e^{-\frac{s}{2}})(e^{\frac{s}{2}} - 1)} \cdot \left((1 + e^{\frac{s}{2}}) \cdot \sigma\left(s \cdot \left(f(x_1, \dots, x_n) - \frac{1}{2}\right)\right) - 1 \right) \tag{D.21}$$

$$= \frac{1}{e^{\frac{s}{2}} - 1} \cdot \left((1 + e^{\frac{s}{2}}) \cdot \sigma\left(s \cdot \left(f(x_1, \dots, x_n) - \frac{1}{2}\right)\right) - 1 \right) \tag{D.22}$$

Next, we prove several properties of the sigmoidal implication.

Proposition 13. For all $a_1, c_1, a_2, c_2 \in [0, 1]$,

1. if $I(a_1, c_1) < I(a_2, c_2)$, then also $\sigma_I(a_1, c_1) < \sigma_I(a_2, c_2)$;
2. if $I(a_1, c_1) = I(a_2, c_2)$, then also $\sigma_I(a_1, c_1) = \sigma_I(a_2, c_2)$.

Proof. 1. We note that σ_I can be written as $\sigma_I(a, c) = w \cdot \sigma(s \cdot (I(a, c) + b_0)) - h$ for constants $w = \frac{(1 + e^{-s \cdot (1 + b_0)})^2}{e^{-s \cdot b_0} - e^{-s \cdot (1 + b_0)}}$ and $h = \frac{1 + e^{-s \cdot (1 + b_0)}}{e^{-s \cdot b_0} - e^{-s \cdot (1 + b_0)}}$. As $s > 0$, $-s \cdot b_0 > -s \cdot (1 + b_0)$. Therefore, $e^{-s \cdot b_0} - e^{-s \cdot (1 + b_0)} > 0$. Furthermore, as $e^{-\frac{s}{2}} > 0$ then certainly $(1 + e^{-\frac{s}{2}})^2 > 0$. As both $e^{-s \cdot b_0} - e^{-s \cdot (1 + b_0)} > 0$ and $(1 + e^{-\frac{s}{2}})^2 > 0$, then also $w > 0$. As $s > 0$, $s \cdot (I(a_1, c_1) + b_0) < s \cdot (I(a_2, c_2) + b_0)$ as $I(a_1, c_1) < I(a_2, c_2)$. Next, note that the sigmoid function σ is a monotonically increasing function. Using $w > 0$ we find that $\sigma_I(a_1, c_1) = w \cdot \sigma(s \cdot (I(a_1, c_1) + b_0)) < w \cdot \sigma(s \cdot (I(a_2, c_2) + b_0)) = \sigma_I(a_2, c_2)$.

2.

$$\sigma_I(a_1, c_1) = w \cdot \sigma(s \cdot (I(a_1, c_1) + b_0)) - h = w \cdot \sigma(s \cdot (I(a_2, c_2) + b_0)) - h = \sigma_I(a_2, c_2) \quad \square$$

Proposition 14. $\sigma_I(a, c)$ is 1 if and only if $I(a, c) = 1$. Similarly, $\sigma_I(a, c)$ is 0 if and only if $I(a, c) = 0$.

Proof. Assume there is some $a, c \in [0, 1]$ so that $I(a, c) = 1$. By construction, $\sigma_I(a, c)$ is 1 (see Appendix D.2).

Now assume there is some $a_1, c_1 \in [0, 1]$ so that $\sigma_I(a_1, c_1) = 1$. Now consider some a_2, c_2 so that $I(a_2, c_2) = 1$. By the construction of σ_I , $\sigma_I(a_2, c_2) = 1$. For the sake of contradiction assume $I(a_1, c_1) < 1$. However, by Proposition 13 as $I(a_1, c_1) < I(a_2, c_2)$ then $\sigma_I(a_1, c_1) < \sigma_I(a_2, c_2)$ has to hold. This is in contradiction with $\sigma_I(a_1, c_1) = \sigma_I(a_2, c_2) = 1$ so the assumption that $I(a_1, c_1) < 1$ has to be wrong and $I(a_1, c_1) = 1$.

The proof for $I(a, c) = 0$ is analogous. \square

Proposition 15. For all fuzzy implications I , σ_I is also a fuzzy implication.

Proof. By Definition 18 $I(\cdot, c)$ is decreasing and $I(a, \cdot)$ is increasing. Therefore, by Proposition 13.1, $\sigma_I(\cdot, c)$ is also decreasing and $\sigma_I(a, \cdot)$ is also increasing. Furthermore, $I(0, 0) = 1$, $I(1, 1) = 1$ and $I(1, 0) = 0$. We find by Proposition 14 that then also $\sigma_I(0, 0) = 1$, $\sigma_I(1, 1) = 1$ and $\sigma_I(1, 0) = 0$. \square

I -sigmoidal implications only satisfy left-neutrality if I is left-neutral and s approaches 0.

Proposition 16. If a fuzzy implication I is contrapositive symmetric with respect to N , then σ_I also is.

Proof. Assume we have an implication I that is contrapositive symmetric and so for all $a, c \in [0, 1]$, $I(a, c) = I(N(c), N(a))$. By Proposition 13.2, $\sigma_I(a, c) = \sigma_I(N(c), N(a))$. Thus, σ_I is also contrapositive symmetric with respect to N . \square

By this proposition, if I is an S-implication, σ_I is contrapositive symmetric and thus also contrapositive differentiable symmetric.

Proposition 17. If I satisfies the identity principle, then σ_I also satisfies the identity principle.

Proof. Assume we have a fuzzy implication I that satisfies the identity principle. Then $I(a, a) = 1$ for all a . By Proposition 14 it holds that $\sigma_I(a, a)$ is also 1. \square

D.3. Nilpotent aggregator

Proposition 18. Equation (A.2) is equal for the Nilpotent t-norm to

$$A_{T_{NM}}(x_1, \dots, x_n) = \begin{cases} \min(x_1, \dots, x_n), & \text{if } x_i + x_j > 1; x_i \text{ and } x_j \text{ are the two lowest values in } x_1, \dots, x_n \\ 0, & \text{otherwise.} \end{cases} \quad (D.23)$$

Proof. We will prove this by induction. Base case: Assume $n = 2$. Then $A_{T_{NM}}(x_1, x_2) = T_{NM}(x_1, x_2)$. x_1 and x_2 are the two lowest values of x_1, x_2 , so the condition in Equation (D.23) would change to $x_1 + x_2 > 1$.

Inductive step: We assume Equation (D.23) holds for some $n \geq 2$. Then by Equation (A.2)

$A_{T_{NM}}(x_1, \dots, x_{n+1}) = T_{NM}(A_{T_{NM}}(x_1, \dots, x_n), x_{n+1})$. Note that if $A_{T_{NM}}(x_1, \dots, x_n) = 0$ then $A_{T_{NM}}(x_1, \dots, x_{n+1})$ is also 0 as $x_{n+1} \in [0, 1]$ and so $0 + x_{n+1} > 1$ can never hold. We identify three cases:

1. If x_{n+1} is the lowest value in x_1, \dots, x_{n+1} , then $A_{T_{NM}}(x_1, \dots, x_n)$ is either the second lowest value in x_1, \dots, x_{n+1} or 0. If it is 0, the sum of the second and third lowest values is not greater than 1, and so the sum of the two lowest values can neither be. If it is not, then $T_{NM}(A_{T_{NM}}(x_1, \dots, x_n), x_{n+1})$ first compares if $A_{T_{NM}}(x_1, \dots, x_n) + x_{n+1} > 1$, that is, if the sum of the two lowest values in x_1, \dots, x_{n+1} is higher than 1, and returns x_{n+1} if this holds and 0 otherwise.
2. If x_{n+1} is the second lowest value in x_1, \dots, x_{n+1} , then $A_{T_{NM}}(x_1, \dots, x_n)$ is either the lowest value in x_1, \dots, x_{n+1} or 0. If it is 0, the sum of the first and third lowest values is not greater than 1, and so the sum of the two lowest values can neither be. If it is not, then $T_{NM}(A_{T_{NM}}(x_1, \dots, x_n), x_{n+1})$ first compares if $A_{T_{NM}}(x_1, \dots, x_n) + x_{n+1} > 1$, that is, if sum of the two lowest values in x_1, \dots, x_{n+1} is higher than 1, and returns $A_{T_{NM}}(x_1, \dots, x_n)$ if this holds and 0 otherwise.
3. If x_{n+1} is neither the lowest nor second lowest value in x_1, \dots, x_{n+1} , then the sum s of the two lowest values in x_1, \dots, x_n is also the sum of the two lowest values in x_1, \dots, x_{n+1} . If $A_{T_{NM}}(x_1, \dots, x_n) = 0$, then s can not have been greater than 1 and so $A_{T_{NM}}(x_1, \dots, x_{n+1})$ is also 0. If it is not, then $s > 1$ and $A_{T_{NM}}(x_1, \dots, x_n)$ is the lowest value and surely $A_{T_{NM}}(x_1, \dots, x_n) + x_{n+1} > 1$ as x_{n+1} is at least as large as the second lowest value. \square

By considering $E_{S_{NM}}(x_1, \dots, x_n) = 1 - A_{T_{NM}}(1 - x_1, \dots, 1 - x_n)$, it is easy to see that

$$E_{S_{NM}}(x_1, \dots, x_n) = \begin{cases} \max(x_1, \dots, x_n), & \text{if } x_i + x_j < 1; x_i \text{ and } x_j \text{ are the two largest values in } x_1, \dots, x_n \\ 1, & \text{otherwise.} \end{cases} \quad (D.24)$$

D.4. Yager R-implication

The Yager t-norm is defined as $T_Y(a, b) = 1 - ((1 - a)^p + (1 - b)^p)^{\frac{1}{p}}$. The Yager R-implication then is defined (see Definition 21) as

$$I_{T_Y}(a, c) = \sup\{b \in [0, 1] | T_Y(a, b) \leq c\} \quad (D.25)$$

When $a \leq c$, $I_{T_Y} = 1$ as $T_Y(a, 1) = a \leq c$. Assuming $a > c$, we find by filling in

$$I_{T_Y}(a, c) = \sup\{b \in [0, 1] | 1 - ((1 - a)^p + (1 - b)^p)^{\frac{1}{p}} \leq c, \quad a > c\} \quad (D.26)$$

To get a closed-form solution of I_{T_Y} we have to find the largest b for which $1 - ((1 - a)^p + (1 - b)^p)^{\frac{1}{p}} \leq c$. Solving this inequality for b , we find

$$\begin{aligned}
 c &\geq 1 - ((1-a)^p + (1-b)^p)^{\frac{1}{p}} \\
 (1-c)^p &\leq (1-a)^p + (1-b)^p \\
 1-b &\geq ((1-c)^p - (1-a)^p)^{\frac{1}{p}} \\
 b &\leq 1 - ((1-c)^p - (1-a)^p)^{\frac{1}{p}}.
 \end{aligned} \tag{D.27}$$

If $a > c$, then $(1-c)^p > (1-a)^p$ and thus $(1-c)^p - (1-a)^p > 0$. Furthermore, as $a, c \in [0, 1]$, $(1-c)^p - (1-a)^p \leq 1$. Therefore, it has to be true that $1 - ((1-c)^p - (1-a)^p)^{\frac{1}{p}} \in [0, 1]$. The largest value $b \in [0, 1]$ for which the condition holds is then equal to $1 - ((1-c)^p - (1-a)^p)^{\frac{1}{p}}$ as it is in $[0, 1]$ and satisfies the inequality.

Combining this with the earlier observation that when $a \leq c$, $I_{T_Y} = 1$, we find the following R-implication:

$$I_{T_Y}(a, c) = \begin{cases} 1, & \text{if } a \leq c \\ 1 - ((1-c)^p - (1-a)^p)^{\frac{1}{p}}, & \text{otherwise.} \end{cases} \tag{D.28}$$

We plot I_{T_Y} for $p = 2$ in Fig. 8. As expected, $p = 1$ reduces to the Łukasiewicz implication. The derivatives of this implication are

$$\frac{\partial I_{T_Y}((, a), c)}{\partial a} = \begin{cases} ((1-c)^p - (1-a)^p)^{\frac{1}{p}-1} \cdot (1-c), & \text{if } a > c \\ 0, & \text{otherwise,} \end{cases} \tag{D.29}$$

$$-\frac{\partial I_{T_Y}((, a), c)}{\partial c} = \begin{cases} ((1-c)^p - (1-a)^p)^{\frac{1}{p}-1} \cdot (1-a), & \text{if } a > c \\ 0, & \text{otherwise.} \end{cases} \tag{D.30}$$

Appendix E. Differentiable Product Fuzzy Logic

We compare Differentiable Product Fuzzy Logic (DPFL), which uses the product t- and t-conorm T_P , S_P , the Reichenbach implication I_{RC} and the log-product aggregator $A_{\log T_P}$, and a probabilistic logic method called Semantic Loss [74].

Definition 22. Let \mathcal{P} be a set of predicates, \mathcal{O} the domain of discourse, η_θ an embedded interpretation of \mathcal{L} and \mathcal{K} a knowledge base of background knowledge. The *Semantic Loss* is defined as

$$\mathcal{L}_S(\theta; \mathcal{K}) = -\log \sum_{\mathbf{w} \models \mathcal{K}} \prod_{\mathbf{w} \models \mathbf{P}(o_1, \dots, o_m)} \eta_\theta(\mathbf{P})(o_1, \dots, o_m) \prod_{\mathbf{w} \models \neg \mathbf{P}(o_1, \dots, o_m)} (1 - \eta_\theta(\mathbf{P})(o_1, \dots, o_m)), \tag{E.1}$$

where \mathbf{w} is a *world* (also known as a *Herbrand interpretation*) that assigns a binary truth value to every ground atom and where $\eta_\theta(\mathbf{P})(o_1, \dots, o_m)$ is the probability of a ground atom.

This computes the logarithm of the sum of probabilities of all worlds for which \mathcal{K} holds, or the probability of sampling a world that is consistent with \mathcal{K} . The probability that a ground atom $\mathbf{P}(o_1, \dots, o_m)$ is true is $\eta_\theta(\mathbf{P})(o_1, \dots, o_m)$. The different ground atoms are assumed to be independent. By marginalizing out the world \mathbf{w} , it can be used for injecting background knowledge in unsupervised or semi-supervised learning. Compared to DPFL, Semantic Loss is exponential in the size of the amount of ground atoms as the sum over valid worlds has to be computed. Equivalent formulas have equal Semantic Loss, and a knowledge base consisting of a conjunction of facts is equal to the cross-entropy loss function.

DPFL is connected in an interesting way to Semantic Loss. It corresponds to a single iteration of the *loopy belief propagation* algorithm [55].

Proposition 19. Let φ be a closed formula so that if $\mathbf{P}_1(o_{11}, \dots, o_{1m})$ and $\mathbf{P}_2(o_{21}, \dots, o_{2m})$ are both ground atoms appearing in φ , then $\mathbf{P}_1(o_{11}, \dots, o_{1m}) = \mathbf{P}_2(o_{21}, \dots, o_{2m})$. Then it holds that $\mathcal{L}_S(\theta; \varphi) = -e_\theta(\{\}, \varphi)$ when using $T = T_P$, $S = S_P$, $I = I_{RC}$ and $A = A_{\log T_P}$.

As there are no loops when each ground atom appears uniquely in φ , the factor graph over which loopy belief propagation is done is a tree. As $e_\theta(\{\}, \varphi)$ corresponds to a single iteration of loopy belief propagation, this is equal to regular belief propagation which is an exact method for computing queries on probabilistic models [58]. Clearly, this condition on φ is very strong. Although loopy belief propagation is known to often be a good approximation empirically [55], the degree to which DPFL approximates Semantic Loss requires further research as this is not a guarantee. However, if DPFL approximates Semantic Loss well, it can be a strong alternative as it is not an exponential computation. However, it also means that most problems of DPFL will also be present in Semantic Loss. For example, if we just have the formula $\forall \text{raven}(x) \rightarrow \text{black}(x)$, the grounding of the knowledge base will not contain repeated ground atoms, and thus Semantic Loss and DPFL are equivalent and share difficulties related to the imbalance of modus ponens and modus tollens.

E.1. Proof

In this section, we prove Proposition 19. Without loss of generality we assume that $\mathcal{K} = \varphi = \forall x_1, \dots, x_n \phi$. Slightly rewriting Equation (E.1), we find that the probability distribution of Semantic Loss is

$$p(\varphi|\eta_\theta) = \sum_{\mathbf{w}} p(\varphi|\mathbf{w})p(\mathbf{w}|\eta_\theta) \quad (\text{E.2})$$

where we define the valuation probability $p(\varphi|\mathbf{w}) = I[\mathbf{w} \models \varphi]$ and the world probability

$$p(\mathbf{w}|\eta_\theta) = \prod_{\mathbf{w} \models \mathbf{P}(o_1, \dots, o_k)} \eta_\theta(\mathbf{P}(o_1, \dots, o_k)) \prod_{\mathbf{w} \not\models \mathbf{P}(o_1, \dots, o_k)} (1 - \eta_\theta(\mathbf{P}(o_1, \dots, o_k))).$$

A Bayesian network is a joint distribution factorized as $p(x) = \prod_{i=1}^n p(x_i|x_{pa(i)})$ where $x_{pa(i)}$ is the set of random variables that are parents of x_i . In particular, we are interested in the joint distribution $p(\varphi, \mathbf{w}|\eta_\theta)$. We use the compositional structure of φ to expand $p(\varphi|\mathbf{w})$.

Let $p(\phi|ch(\phi))$ be the probability that ϕ , a subformula of φ , is true according to the binary truth table conditioned on the truth value of the direct subformulas $ch(\phi)$ of ϕ . For example, if $\phi = \alpha \otimes \beta$, then $p(\phi|\alpha, \beta) = 1$ if α and β are 1, and otherwise $p(\phi|\alpha, \beta) = 0$. For an atomic formula $\mathbf{P}(o_1, \dots, o_k)$, we do a lookup in the world \mathbf{w} , $p(\mathbf{P}(o_1, \dots, o_k)|\mathbf{w}_{\mathbf{P}(o_1, \dots, o_k)}) = \mathbf{w}_{\mathbf{P}(o_1, \dots, o_k)}$. Let Φ be the set of all subformulas of φ . We express the joint distribution as

$$p(\Phi, \mathbf{w}|\eta_\theta) = p(\mathbf{w}|\eta_\theta) \prod_{\phi \in \Phi} p(\phi|ch(\phi)). \quad (\text{E.3})$$

A specific world \mathbf{w} uniquely determines a single Φ so that $\prod_{\phi \in \Phi} p(\phi|ch(\phi)) = 1$. Note that the distribution $p(\varphi|\mathbf{w}) = \prod_{\phi \in \Phi} p(\phi|ch(\phi))$ forms a polytree (or directed tree), as a logical expression is formed as a tree. From this Bayesian network, we define the factor graph over which we do the belief propagation. For brevity, we denote a specific ground atom $\mathbf{P}(o_1, \dots, o_k)$ as \mathbf{P}_O .

We note that we use m to denote an instantiation instead of μ in this section. μ is used to denote the messages as customary in belief propagation.

- There is a variable node $w_{\mathbf{P}_O}$ for every ground atom \mathbf{P}_O appearing in the grounding of φ . Additionally, there is a factor node $f_{w_{\mathbf{P}_O}}$

$$f_{w_{\mathbf{P}_O}}(w_{\mathbf{P}_O}) = \eta_\theta(\mathbf{P}(o_1, \dots, o_k))^{w_{\mathbf{P}_O}} \cdot (1 - \eta_\theta(\mathbf{P}(o_1, \dots, o_k)))^{1-w_{\mathbf{P}_O}}.$$
- There is a variable node ϕ and a factor node f_ϕ for every subformula $\phi \in \Phi$.
- For every $\phi = \mathbf{P}_O$, $\phi \in \Phi$, $f_\phi(\phi, w_{\mathbf{P}_O}) = I[\phi = w_{\mathbf{P}_O}]$.
- For every $\phi = \neg\alpha$, $\phi \in \Phi$, $f_\phi(\phi, \alpha) = I[\phi = 1 - \alpha]$.
- For every $\phi = \alpha \otimes \beta$, $\phi \in \Phi$, $f_\phi(\phi, \alpha, \beta) = I[\phi = \alpha \cdot \beta]$.
- Let $\varphi = \forall x_1, \dots, x_n \phi$ be the top node. Denote the set of all instances of φ is M . Then $f_\varphi(\varphi, m_1, \dots, m_{|M|}) = I[\varphi = \prod_{m \in M} \alpha_m]$ where e_μ is the random variable corresponding to the instantiation of m in ϕ .

We ignore the other connectives as they can be formed from \neg and \otimes , both in classical logic as in DPFL. Next, we compute the messages in belief propagation. We start from the world variable nodes $w_{\mathbf{P}_O}$ and move up through the computation tree to φ . The messages for factors to variables are given as [7] $\mu_{f_s \rightarrow x}(x) = \sum_X f_s(x, X) \prod_{y \in ne(f_s) \setminus x} \mu_{y \rightarrow f_s}(y)$ where $ne(x)$ is the set of neighbors of node x . The messages for variables to factors are given as $\mu_{x \rightarrow f_s}(x) = \prod_{l \in ne(x) \setminus f_s} \mu_{f_l \rightarrow x}(x)$.

- a. $\mu_{f_{w_{\mathbf{P}_O}} \rightarrow w_{\mathbf{P}_O}}(w_{\mathbf{P}_O}) = \eta_\theta(\mathbf{P}(o_1, \dots, o_k))^{w_{\mathbf{P}_O}} \cdot (1 - \eta_\theta(\mathbf{P}(o_1, \dots, o_k)))^{1-w_{\mathbf{P}_O}}$, factor to variable for ground atom.
- b. $\mu_{w_{\mathbf{P}_O} \rightarrow f_\phi}(w_{\mathbf{P}_O}) = \mu_{f_{w_{\mathbf{P}_O}} \rightarrow w_{\mathbf{P}_O}}(w_{\mathbf{P}_O}) \prod_{\psi: \psi = \mathbf{P}_O, \phi \neq \psi} \mu_{f_\psi \rightarrow w_{\mathbf{P}_O}} = \mu_{f_{w_{\mathbf{P}_O}} \rightarrow w_{\mathbf{P}_O}}(w_{\mathbf{P}_O})$, ground atom variable to atomic formula factor $\phi = \mathbf{P}(o_1, \dots, o_k)$. We here assume that the incoming messages $\mu_{f_\psi \rightarrow w_{\mathbf{P}_O}}$ from other atomic formulas using ground atom \mathbf{P}_O are initialized with 1. We have not yet, and are not able to, compute these as the graph might contain loops.
- c. $\mu_{f_\phi \rightarrow \phi}(\phi) = I[\phi = 1] \mu_{w_{\mathbf{P}_O} \rightarrow f_\phi}(1) + I[\phi = 0] \mu_{w_{\mathbf{P}_O} \rightarrow f_\phi}(0) = \eta_\theta(\mathbf{P}(o_1, \dots, o_k))^\phi \cdot (1 - \eta_\theta(\mathbf{P}(o_1, \dots, o_k)))^{1-\phi}$, factor to variable for atomic formulas.
- d. $\mu_{\phi \rightarrow f_\alpha}(\alpha) = \mu_{f_\alpha \rightarrow \alpha}(\alpha)$ for subformula variables to factors of other subformulas α . As ϕ is only used in two factors, this simply passes the downstream message through.
- e. $\mu_{f_\phi \rightarrow \phi}(\phi) = \mu_{f_\alpha \rightarrow \alpha}(1)^{1-\phi} \cdot \mu_{f_\alpha \rightarrow \alpha}(0)^\phi$ factor to variable for negated subformulas $\phi = \neg\alpha$.
- f. $\mu_{f_\phi \rightarrow \phi}(\phi) = (\mu_{f_\alpha \rightarrow \alpha}(1) \cdot \mu_{f_\beta \rightarrow \beta}(1))^\phi \cdot (\mu_{f_\alpha \rightarrow \alpha}(0) \cdot \mu_{f_\beta \rightarrow \beta}(0) + \mu_{f_\alpha \rightarrow \alpha}(1) \cdot \mu_{f_\beta \rightarrow \beta}(0) + \mu_{f_\alpha \rightarrow \alpha}(0) \cdot \mu_{f_\beta \rightarrow \beta}(1))^{1-\phi}$ factor to variable for conjunctions $\phi = \alpha \otimes \beta$.
- g. $\mu_{f_\varphi \rightarrow \varphi}(\varphi) = (\prod_{m \in M} \mu_{f_{\alpha_m} \rightarrow \alpha_m}(1))^\varphi \cdot (\sum_{\alpha_1, \dots, \alpha_{|M|}: \exists i: \alpha_i = 0} \prod_{m \in M} \mu_{f_{\alpha_m} \rightarrow \alpha_m}(\alpha_m))^{1-\varphi}$ for the factor of the universally quantified formula $\varphi = \forall x_1, \dots, x_n \phi$ where α_m is the subformula corresponding to the instantiation m . The second term is the sum over all cases so that there is a subformula with truth value 0.

We wish to know what the marginal probability $p(\varphi = 1|\eta_\theta)$ is. A marginal of a variable ϕ in a factor graph is found as $p(\phi) = \prod_{s \in ne(\phi)} \mu_{f_s \rightarrow \phi}$. The variable node φ only has the factor node f_φ as a neighbor, so using (g.) we find¹⁶

$$p(\varphi = 1|\eta_\theta) \approx \prod_{m \in M} \mu_{\alpha_m \rightarrow f_\varphi}(1). \tag{E.4}$$

Next, we use induction to prove that the computation of $\mu_{f_{\alpha_m} \rightarrow \alpha_m}(1)$ is equal to Differentiable Product Fuzzy Logic.

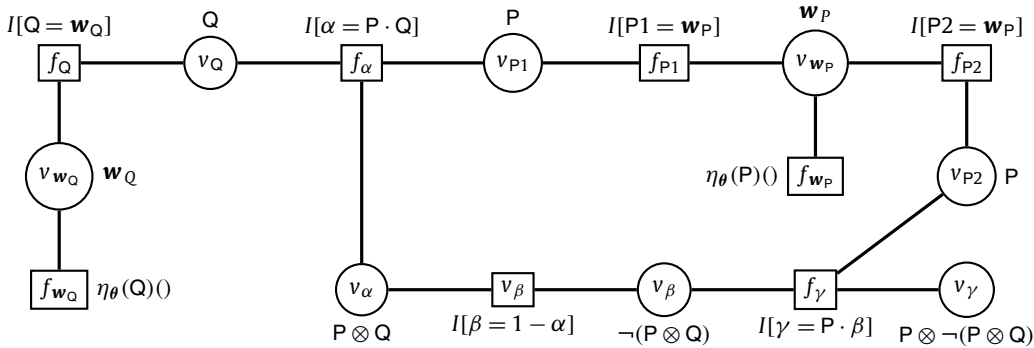
- Let ϕ be any subformula of α_m for some instantiation m of φ . Let e_θ be the valuation function (Definition 2) with $T = T_P$ and $N = N_C$. We prove that $\mu_{f_\phi \rightarrow \phi}(1) = e_\theta(m, \phi)$ and $\mu_{f_\phi \rightarrow \phi}(0) = 1 - e_\theta(m, \phi)$.
- Base case $\phi = P_0$. By (c.), $\mu_{f_\phi \rightarrow \phi}(1) = \eta_\theta(P)(o_1, \dots, o_k)$ and $\mu_{f_\phi \rightarrow \phi}(0) = 1 - \eta_\theta(P)(o_1, \dots, o_k)$. By Equation (3),¹⁷ $e_\theta(m, \phi) = \eta_\theta(P)(o_1, \dots, o_k)$.
- Inductive step $\phi = \neg\alpha$. By (e.) and using the inductive hypothesis, $\mu_{f_\phi \rightarrow \phi}(1) = \mu_{f_\alpha \rightarrow \alpha}(0) = 1 - e_\theta(m, \alpha)$ and $\mu_{f_\phi \rightarrow \phi}(0) = \mu_{f_\alpha \rightarrow \alpha}(1) = e_\theta(m, \alpha)$. By Equation (4), $e_\theta(m, \phi) = 1 - e_\theta(m, \alpha)$.
- Inductive step $\phi = \alpha \otimes \beta$. By (f.) and using the inductive hypothesis, $\mu_{f_\phi \rightarrow \phi}(1) = \mu_{f_\alpha \rightarrow \alpha}(1) \cdot \mu_{f_\beta \rightarrow \beta}(1) = e_\theta(m, \alpha) \cdot e_\theta(m, \beta)$ and $\mu_{f_\phi \rightarrow \phi}(0) = \mu_{f_\alpha \rightarrow \alpha}(0) \cdot \mu_{f_\beta \rightarrow \beta}(0) + \mu_{f_\alpha \rightarrow \alpha}(1) \cdot \mu_{f_\beta \rightarrow \beta}(0) + \mu_{f_\alpha \rightarrow \alpha}(0) \cdot \mu_{f_\beta \rightarrow \beta}(1) = (1 - e_\theta(m, \alpha))(1 - e_\theta(m, \beta)) + e_\theta(m, \alpha)(1 - e_\theta(m, \beta)) + (1 - e_\theta(m, \alpha))e_\theta(m, \beta) = 1 - e_\theta(m, \alpha) \cdot e_\theta(m, \beta)$. By Equation (5) and $T_P(a, c) = a \cdot c$, $e_\theta(m, \phi) = e_\theta(m, \alpha) \cdot e_\theta(m, \beta)$.

Using Equation (E.4) we then find that $p(\varphi = 1|\eta_\theta) \approx \prod_{m \in M} e_\theta(m, \phi_m)$, which is equal to the Differentiable Product Fuzzy Logic computation of the universal quantifier in Equation (8).

Importantly, as the computation of the logic is itself a tree, the only loops are caused through ground atoms appearing in multiple subformulas. Therefore, when each ground atom only appears in a single formula, Differentiable Product Fuzzy Logic computes the same probability as Semantic Loss.

E.2. Example

For example, a formula P corresponds to a variable node P with two possible values 1 and 0, along with the factor node with factor $\eta_\theta(P)()$ if P is true and $1 - \eta_\theta(P)()$ otherwise. If we consider the formula $\gamma = P \otimes \neg(P \otimes Q)$ we find the following factor graph:



Here, box nodes correspond to factor nodes and circle nodes correspond to variable nodes. As γ is the top formula, this is where the messages get passed to. Note that there is a single loop, which is present because the atom P is used twice in the formula. This causes two incorrect messages: $\mu_{v_{w_P} \rightarrow f_{P1}}$ and $\mu_{v_{w_P} \rightarrow f_{P2}}$. The first is incorrect as it does not have access to the incoming message $\mu_{f_{P2} \rightarrow v_{w_P}}$ and puts it to 1.

With Differentiable Product Fuzzy Logic, we find the expression $\eta_\theta(P)() \cdot (1 - \eta_\theta(P)() \cdot \eta_\theta(Q)())$, while the correct probability is where P is 1 and Q is 0, that is $\eta_\theta(P)() \cdot (1 - \eta_\theta(Q)())$.

Appendix F. Additional experiments

In this Appendix, we report additional experiments on the same problem. Throughout this Appendix we use vanilla stochastic gradient descent with a learning rate of 0.01 and 0.5 momentum instead of ADAM. The reason for this is that the best configurations seem to perform better with this optimizer, while the other configurations seem to perform better with ADAM.

¹⁶ We use approximation equality for simplification. Note that this is not necessarily true in loopy belief propagation.

¹⁷ Equation (3) refers to the ungrounded case, but here the full grounding is already done so the lookup function l is not required.

Table F.11
Configurations using the RMSE aggregator with $w_{DFL} = 1$ and the log product aggregator with $w_{dfi} = 10$.

	$A_{\log T_P, w_{DFL} = 10}$				$A_{RMSE, w_{DFL} = 1}$			
	Accuracy	cons%	$cu_{cons}\%$	$cu_{ant}\%$	Accuracy	cons%	$cu_{cons}\%$	$cu_{ant}\%$
T_G	96.3	0.10	0.89	0.97	96.2	0.10	0.89	0.97
T_{LK}	96.6	0.5	0.33	0.67	96.9	0.5	0.06	0.95
T_P	96.7	0.44	0.48	0.69	97.0	0.08	0.81	0.98
$T_Y, p = 2$	96.4	0.40	0.54	0.74	96.6	0.12	0.87	0.97
$T_Y, p = 20$	96.0	0.29	0.51	0.78	95.9	0.18	0.82	0.98
T_{Nm}	95.4	0.29	0.44	0.84	95.4	0.03	0.62	0.99

F.1. Comparing different formulas of the same problem

We note what the values of $cu_{cons}\%$ and $cu_{ant}\%$ roughly are for each of the used formulas if we were to pick at random.

- $\forall x, y \text{ zero}(x) \otimes \text{zero}(y) \rightarrow \text{same}(x, y), \dots, \forall x, y \text{ nine}(x) \otimes \text{nine}(y) \rightarrow \text{same}(x, y)$. For this formula, $cu_{cons}\% \geq \frac{1}{10}$ as it is the distribution of $\text{same}(x, y)$ ¹⁸ and $cu_{ant}\% \leq \frac{99}{100}$ as it is 1 minus the probability that both x and y are zero. The modus ponens case is true in more than $\frac{1}{100}$ cases, the modus tollens case in less than $\frac{9}{10}$ cases and the ‘distrust’ option in more than $\frac{9}{100}$ cases.
- $\forall x, y \text{ zero}(x) \otimes \text{same}(x, y) \rightarrow \text{zero}(y), \dots, \forall x, y \text{ nine}(x) \otimes \text{same}(x, y) \rightarrow \text{nine}(y)$. For this formula, $cu_{cons}\% = \frac{1}{10}$ as it is the probability that a digit represents zero and $cu_{ant}\% \leq \frac{99}{100}$. The modus ponens cases are true in more than $\frac{1}{100}$ cases, the modus tollens in $\frac{9}{10}$ cases and the ‘distrust’ option in $\frac{9}{100}$ cases.
- $\forall x, y \text{ same}(x, y) \rightarrow \text{same}(y, x)$. As this is a bi-implication, $cu_{cons}\% \geq \frac{1}{10}$ and $cu_{ant}\% \leq \frac{9}{10}$. The ‘distrust’ option is not possible in this formula.

From this, we can see that a set of operators is better than random guessing for the consequent updates if $cu_{cons}\% > 0.1$. It is more difficult to say what the value of $cu_{ant}\%$ should be as good as random guessing, as the probabilities are upper bounded with the lowest bound at 0.9. We can only say that we know a set of operators to be better than random if $cu_{ant}\% > 0.99$.

F.2. Varying the aggregators

In this section, we analyze symmetric configurations in the same problem, except that we use aggregators other than the one formed by extending the t-norm. In particular, we will consider the RMSE aggregator (A_{GME} with $p = 2$) and the log-product aggregator $A_{\log T_P}$.

Table F.11 shows the results when using the RMSE aggregator and a DFL weight of 1 and the log product aggregator and a DFL weight of 10. Nearly all configurations perform significantly better using these aggregators than when using their ‘symmetric’ aggregator. In particular, the Gödel, Łukasiewicz and Yager t-norms all outperform the baseline with both aggregators as they are differentiable everywhere and can handle outliers.

The product t-norm seems to do slightly worse with the RMSE aggregator than with the log-product aggregator. Like we discussed in Section 8.5, $cons\%$ is higher using this aggregator because the corners $a_i = 0, c_i = 0$ and $a_i = 1, c_i = 1$ will have no gradient when using the RMSE aggregator. However, the values of $cu_{cons}\%$ and $cu_{ant}\%$ are much lower than when using the log-product aggregator. This could have to do with the previously made point: As it no longer has a gradient of 1 at the corners $a = 0, c = 0$ and $a = 1, c = 1$, the large gradients are only when the agent is not yet confident about some prediction. This case is inherently ‘riskier’, but also contributes more information. It is not as informative to increase the confidence of $a = 0$ if a is already very low.

The Łukasiewicz t-norm has a particularly high accuracy of 96.9% with the log product and is on the level of performance of the product t-norm. However, it has a very low value for $cu_{cons}\%$ of 0.06 and a relatively low value for $cu_{ant}\%$. Interestingly, it is also the only configuration for which $cu_{cons}\%$ is higher when using the RMSE aggregator than the log-product aggregator.

F.3. Reichenbach-sigmoidal implication

The newly introduced Reichenbach-sigmoidal implication σ_{IRC} is a promising candidate for the choice of implication as we have argued in Section 8.5.1. To get a better understanding of this implication, we investigate the effect of its parameters in the same problem. We fix the aggregator to the log-product, the conjunction operator to the Yager t-norm with $p = 2$, and use a DFL weight of $w_{DFL} = 10$.

¹⁸ It is slightly more than $\frac{1}{10}$ because we are using a minibatch of examples. Therefore, the reflexive pairs (i.e., $\text{same}(x, x)$) are common.

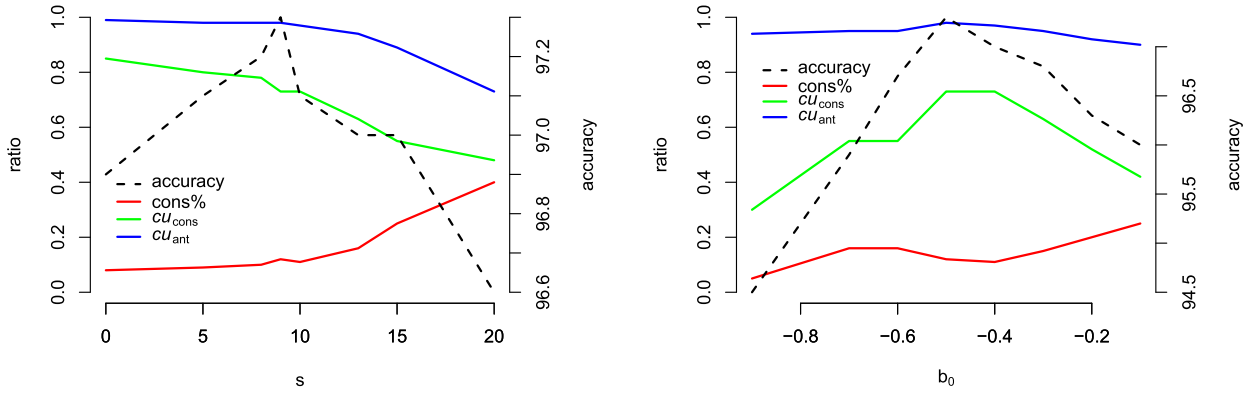


Fig. F.18. The results using the Reichenbach-sigmoidal implication σ_{IRC} , the $A_{\log T_p}$ aggregator, T_Y with $p = 2$ and $w_{DFL} = 10$. Left shows the results for various values of s , keeping b_0 fixed to -0.5 , and right shows the results for various values of b_0 , keeping s fixed to 9 .

Table F.12

The results using T_Y , $p = 2$ for the conjunction, $A_{\log T_p}$ for the aggregator with $w_1 = 10$ and several S-implications and R-implications. These are run using vanilla SGD instead of ADAM.

	Accuracy	cons%	$cu_{cons}\%$	$cu_{ant}\%$		Accuracy	cons%	$cu_{cons}\%$	$cu_{ant}\%$
I_{KD}	96.1	0.10	0.88	0.97	I_G	90.6	1	0.07	
I_{LK}	97.0	0.5	0.03	0.97	I_{LK}	97.0	0.5	0.03	0.97
I_{RC}	96.9	0.08	0.85	0.99	I_{GG}	94.0	0.86	0.01	0.97
$I_Y, p = 2$	96.6	0.12	0.87	0.97	I_{TY}	95.4	0.58	0.03	0.99
I_{FD}	96.3	0.07	0.65	0.96	I_{FD}	96.3	0.07	0.65	0.96
					$I_{TY}, p = 0.5$	96.3	0.18	0.65	0.37

Table F.13

The results using σ_{IRC} for the implication with $s = 9$ and $b_0 = -\frac{1}{2}$, $T_Y, p = 2$ for the conjunction and $A_{\log T_p}$ for the aggregator with $w_{DFL} = 10$, leaving some formulas of the same problem out. The numbers indicated the formulas that are present during training.

Formulas	Accuracy	cons%	$cu_{cons}\%$	$cu_{ant}\%$
(1) & (2)	97.1	0.05	0.54	0.99
(2) & (3)	95.9	0.12	0.75	0.95
(1) & (3)	96.3	0.15	0.52	0.98
(1)	95.6	0.05	0.59	1.00
(2)	95.2	0.03	0.78	0.99
(3)	95.8	0.19	0.64	0.95

On the left plot of Fig. F.18 we find the results when we experiment with the parameter s , keeping b_0 fixed to $-\frac{1}{2}$. Note that when s approaches 0 the Reichenbach-sigmoidal implication is I_{RC} . The value of 9 gives the best results, with 97.3% accuracy. This outperforms the other implications, also when using the vanilla SGD optimizer, as displayed in Table F.12. Interestingly enough, there seem to be clear trends in the values of $cons\%$, $cu_{cons}\%$ and $cu_{ant}\%$. Increasing s seems to increase $cons\%$. This is because the antecedent derivative around the corner $a = 0$, $c = 0$ will be low, as argued in Section 8.5.1. When s increases, the corners will be more smoothed out. Furthermore, both $cu_{cons}\%$ and $cu_{ant}\%$ decrease when s increases. This could again be because around the corners the derivatives become small. Updates in the corner will likely be correct as the model is already confident about those. For a higher value of s , most of the gradient magnitude is at instances on which the model is less confident. We note that the same happened when using the RMSE aggregator and the product t-norm. Regardless, the best parameter value clearly is not the one for which the values of $cu_{cons}\%$ and $cu_{ant}\%$ are highest, namely the Reichenbach implication itself.

On the right plot of Fig. F.18 we experiment with the value of b_0 . Clearly, $-\frac{1}{2}$ works best, having the highest accuracy and $cu_{cons}\%$.

F.4. Influence of individual formulas

Finally, we compare what the influence of the different formulas of the same problem is in Table F.13. Removing the reflexivity formula (3) does not largely impact the performance. The biggest drop in performance is by removing formula (1) that defines the same predicate. Using only formula (1) gets slightly better performance than only using formula (2), despite the fact that no positive labeled examples can be found using formula (1) as the predicates zero to nine are not in its consequent. Since 95% of the derivatives are with respect to the negated antecedent, this formula contributes by finding

additional counterexamples. Furthermore, improving the accuracy of the same predicate improves the accuracy on digit recognition: just using the reflexivity formula (3) has the highest accuracy when used individually, even though it does not use the digit predicates.

References

- [1] E. Arakelyan, D. Daza, P. Minervini, M. Cochez, Complex query answering with neural link predictors, in: *International Conference on Learning Representations*, 2021.
- [2] S. Badreddine, A. d'Avila Garcez, L. Serafini, M. Spranger, Logic tensor networks, arXiv:2012.13635 [cs], <http://arxiv.org/abs/2012.13635>, 2020.
- [3] S. Badreddine, M. Spranger, Injecting prior knowledge for transfer learning into reinforcement learning algorithms using logic tensor networks, arXiv preprint arXiv:1906.06576, 2019.
- [4] H. Bal, D. Epema, C. de Laat, R. van Nieuwpoort, J. Romein, F. Seinstra, C. Snoek, H. Wijshoff, A medium-scale distributed system for computer science research: infrastructure for the long term, *Computer* 49 (2016) 54–63.
- [5] K. Ball, An elementary introduction to modern convex geometry, in: *Flavors of Geometry*, vol. 31, 1997, pp. 1–58.
- [6] T.R. Besold, A.d. Garcez, S. Bader, H. Bowman, P. Domingos, P. Hitzler, K.U. Kuehnberger, L.C. Lamb, D. Lowd, P.M.V. Lima, L. de Penning, G. Pinkas, H. Poon, G. Zaverucha, Neural-symbolic learning and reasoning: a survey and interpretation, arXiv preprint arXiv:1711.03902, <http://arxiv.org/abs/1711.03902>, 2017.
- [7] C.M. Bishop, *Pattern Recognition and Machine Learning*, 2006.
- [8] A. Brock, J. Donahue, K. Simonyan, Large scale GAN training for high fidelity natural image synthesis, arXiv preprint arXiv:1809.11096, 2018.
- [9] M. Buda, A. Maki, M.A. Mazurowski, A systematic study of the class imbalance problem in convolutional neural networks, *Neural Netw.* 106 (2018) 249–259.
- [10] T. Calvo, A. Kolesárová, M. Komorníková, R. Mesiar, Aggregation operators: properties, classes and construction methods, in: T. Calvo, G. Mayor, R. Mesiar (Eds.), *Aggregation Operators: New Trends and Applications*, Physica-Verlag HD, Heidelberg, 2002, pp. 3–104.
- [11] R. Cignoli, F. Esteve, L. Godo, F. Montagna, On a class of left-continuous t-norms, *Fuzzy Sets Syst.* 131 (2002) 283–296, [https://doi.org/10.1016/S0165-0114\(01\)00215-9](https://doi.org/10.1016/S0165-0114(01)00215-9).
- [12] P. Cintula, P. Hájek, C. Nogura, *Handbook of Mathematical Fuzzy Logic*, vol. 1, College Publications, 2011.
- [13] A. Daniele, L. Serafini, Knowledge enhanced neural networks, in: A.C. Nayak, A. Sharma (Eds.), *PRICAI 2019: Trends in Artificial Intelligence*, Springer International Publishing, Cham, 2019, pp. 542–554.
- [14] A. Darwiche, SDD: a new canonical representation of propositional knowledge bases, in: *IJCAI International Joint Conference on Artificial Intelligence*, 2011, pp. 819–826.
- [15] H.A. David, H.N. Nagaraja, *Order Statistics*, third edition, *Encyclopedia of Statistical Sciences*, 2003.
- [16] T. Demeester, T. Rocktäschel, S. Riedel, Lifted rule injection for relation embeddings, in: *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, Association for Computational Linguistics, 2016, pp. 1389–1399, <http://aclweb.org/anthology/D16-1146>.
- [17] M. Diligenti, M. Gori, C. Sacca, Semantic-based regularization for learning and inference, *Artif. Intell.* 244 (2017) 143–165.
- [18] M. Diligenti, S. Roychowdhury, M. Gori, Integrating prior knowledge into deep learning, in: *Machine Learning and Applications (ICMLA)*, 2017 16th IEEE International Conference on, IEEE, 2017, pp. 920–923.
- [19] I. Donadello, L. Serafini, A.d. Garcez, Logic tensor networks for semantic image interpretation, in: *IJCAI International Joint Conference on Artificial Intelligence*, 2017, pp. 1596–1602, <http://arxiv.org/abs/1705.08968>.
- [20] R. Evans, E. Grefenstette, Learning explanatory rules from noisy data, *J. Artif. Intell. Res.* 61 (2018) 65–170, <https://doi.org/10.1613/jair.5477>.
- [21] K. Ganchev, J. Gillenwater, Posterior regularization for structured latent variable models, *J. Mach. Learn. Res.* 11 (2010) 2001–2049, <http://dl.acm.org/citation.cfm?id=1756006.1859918>.
- [22] A.d. Garcez, K.B. Broda, D.M. Gabbay, *Neural-Symbolic Learning Systems: Foundations and Applications*, Springer Science & Business Media, 2012.
- [23] M. Garnelo, K. Arulkumar, M. Shanahan, Towards deep symbolic reinforcement learning, arXiv preprint arXiv:1609.05518, 2016.
- [24] L. Getoor, B. Taskar, *Introduction to Statistical Relational Learning*, vol. 1, MIT Press, Cambridge, 2007.
- [25] I. Goodfellow, Y. Bengio, A. Courville, Y. Bengio, *Deep Learning*, vol. 1, MIT Press, Cambridge, 2016.
- [26] S. Guo, Q. Wang, L. Wang, B. Wang, L. Guo, Jointly embedding knowledge graphs and logical rules, in: *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, 2016, pp. 192–202.
- [27] A.K. Gupta, S. Nadarajah, *Handbook of Beta Distribution and Its Applications*, CRC Press, 2004.
- [28] P. Hájek, *The Metamathematics of Fuzzy Logic*, Kluwer, 1998.
- [29] P. Hall, The distribution of means for samples of size N drawn from a population in which the variate takes values between 0 and 1, all such values being equally probable, *Biometrika* 19 (1927) 240–245, <http://www.jstor.org/stable/2331961>.
- [30] S. Harnad, The symbol grounding problem, *Physica D* 42 (1990) 335–346.
- [31] C.G. Hempel, Studies in the Logic of Confirmation (II.), *Mind*, vol. 54, 1945, pp. 97–121, <http://www.jstor.org/stable/2250948>.
- [32] Z. Hu, X. Ma, Z. Liu, E. Hovy, E. Xing, Harnessing deep neural networks with logic rules, in: *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Association for Computational Linguistics, 2016, pp. 2410–2420, <http://aclweb.org/anthology/P16-1228>.
- [33] C.L. Hwang, A.S.M. Masud, *Multiple Objective Decision Making—Methods and Applications: A State-of-the-Art Survey*, vol. 164, Springer Science & Business Media, 2012.
- [34] J.O. Irwin, On the frequency distribution of the means of samples from a population having any law of frequency with finite moments, with special reference to Pearson's type II, *Biometrika* (1927) 225–239.
- [35] J.S.R. Jang, ANFIS: adaptive-network-based fuzzy inference system, *IEEE Trans. Syst. Man Cybern.* (1993), <https://doi.org/10.1109/21.256541>.
- [36] J.S.R. Jang, C.T. Sun, E. Mizutani, *Neuro-Fuzzy and Soft Computing: A Computational Approach to Learning and Machine Intelligence*, 1997.
- [37] N. Japkowicz, S. Stephen, The class imbalance problem: a systematic study, *Intell. Data Anal.* 6 (2002) 429–449.
- [38] B. Jayaram, M. Baczyński, *Fuzzy Implications*, vol. 231, Springer, Berlin, Heidelberg, 2008, <http://link.springer.com/10.1007/978-3-540-69082-5>.
- [39] D.P. Kingma, J. Ba, Adam: a method for stochastic optimization, arXiv:1412.6980 [cs], <http://arxiv.org/abs/1412.6980>, 2017.
- [40] E.P. Klement, R. Mesiar, E. Pap, *Triangular Norms*, vol. 8, Springer Science & Business Media, 2013.
- [41] G. Klir, B. Yuan, *Fuzzy Sets and Fuzzy Logic*, vol. 4, Prentice Hall, New Jersey, 1995.
- [42] E. van Krieken, E. Acar, F. van Harmelen, Semi-supervised learning using differentiable reasoning, *IfCoLog J. Log. Appl.* 6 (2019) 633–653.
- [43] Y. LeCun, C. Cortes, MNIST handwritten digit database, <http://yann.lecun.com/exdb/mnist/>, 2010.
- [44] C.T. Lin, G.C. Lee, Neural-network-based fuzzy logic control and decision system, *IEEE Trans. Comput.* (1991), <https://doi.org/10.1109/12.106218>.
- [45] Y. Liu, E. Kerre, An overview of fuzzy quantifiers. (I). Interpretations, *Fuzzy Sets Syst.* 95 (1998) 1–21.
- [46] R. Manhaeve, S. Dumančić, A. Kimmig, T. Demeester, L. De Raedt, DeepProbLog: neural probabilistic logic programming, in: S. Bengio, H.M. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, R. Garnett (Eds.), *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, 3–8 December 2018, Montréal, Canada*, 2018, <http://arxiv.org/abs/1805.10872>, <http://papers.nips.cc/book/advances-in-neural-information-processing-systems-31-2018>, 2018.

- [47] G. Marcus, Deep learning: a critical appraisal, arXiv preprint arXiv:1801.00631, 2018.
- [48] G. Marra, F. Giannini, M. Diligenti, M. Gori, Constraint-based visual generation, arXiv preprint arXiv:1807.09202, 2018.
- [49] G. Marra, F. Giannini, M. Diligenti, M. Gori, Integrating learning and reasoning with deep logic models, arXiv preprint arXiv:1901.04195, pp. 1–17, <https://arxiv.org/pdf/1901.04195v1.pdf>, 2019.
- [50] G. Marra, F. Giannini, M. Diligenti, M. Maggini, M. Gori, Learning and T-norms theory, arXiv preprint arXiv:1907.11468, 2019.
- [51] M. Mayo, Symbol grounding and its implications for artificial intelligence BT, in: Twenty-Sixth Australasian Computer Science Conference (ACSC2003), 16, 2003, pp. 55–60, <http://crpit.com/confpapers/CRPITV16Mayo.pdf>.
- [52] P. Minervini, T. Demeester, T. Rocktäschel, S. Riedel, Adversarial sets for regularising neural link predictors, in: Uncertainty in Artificial Intelligence- Proceedings of the 33rd Conference, UAI 2017, 2017.
- [53] P. Minervini, S. Riedel, Adversarially regularising neural NLI models to integrate logical background knowledge, in: Proceedings of the 22nd Conference on Computational Natural Language Learning, 2018, pp. 65–74.
- [54] S. Muggleton, L. de Raedt, Inductive logic programming: theory and methods, J. Log. Program. 19–20 (1994) 629–679, [https://doi.org/10.1016/1066-9490\(94\)90035-3](https://doi.org/10.1016/1066-9490(94)90035-3), <http://www.sciencedirect.com/science/article/pii/0743106694900353>.
- [55] K.P. Murphy, Y. Weiss, M.I. Jordan, Loopy belief propagation for approximate inference: an empirical study, in: Proceedings of the Fifteenth Conference on Uncertainty in Artificial Intelligence, Morgan Kaufmann Publishers Inc., 2013, pp. 467–475, <http://arxiv.org/abs/1301.6725>.
- [56] V. Novák, I. Perfilieva, J. Močkoř, Mathematical Principles of Fuzzy Logic, Springer, US, 1999, <https://doi.org/10.1007/978-1-4615-5217-8>.
- [57] H. Páll Jónsson, Real Logic and Logical Tensor Networks, University of Amsterdam, 2018, MSc thesis, 45.
- [58] J. Pearl, Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference, Elsevier, 1988.
- [59] J. Pearl, Theoretical impediments to machine learning with seven sparks from the causal revolution, in: Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining, ACM, 2018, p. 3.
- [60] T.G. Pham, N. Turkkan, Reliability of a standby system with beta-distributed component lives, IEEE Trans. Reliab. (1994) 285114, <https://doi.org/10.1109/24.285114>.
- [61] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, I. Sutskever, Language Models Are Unsupervised Multitask Learners, 2019.
- [62] A. Rasmus, M. Berglund, M. Honkala, H. Valpola, T. Raiko, Semi-supervised learning with ladder networks, in: Advances in Neural Information Processing Systems, 2015, pp. 3546–3554.
- [63] R. Riegel, A. Gray, F. Luus, N. Khan, N. Makondo, I.Y. Akhalwaya, H. Qian, R. Fagin, F. Barahona, U. Sharma, S. Ikbal, H. Karanam, S. Neelam, A. Likhyan, S. Srivastava, Logical neural networks, arXiv:2006.13155, 2020.
- [64] T. Rocktäschel, S. Riedel, End-to-end differentiable proving, <https://arxiv.org/pdf/1705.11040.pdf>, <http://arxiv.org/abs/1705.11040>, 2017.
- [65] T. Rocktäschel, S. Singh, S. Riedel, Injecting logical background knowledge into embeddings for relation extraction, in: Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, 2015, pp. 1119–1129, <https://rockt.github.io/pdf/rocktaschel2015injecting.pdf>, <http://www.aclweb.org/anthology/N15-1118>, <http://aclweb.org/anthology/N15-1118>.
- [66] P. Sen, G. Namata, M. Bilgic, L. Getoor, B. Galligher, T. Eliassi-Rad, Collective classification in network data, AI Mag. 29 (2008) 93.
- [67] L. Serafini, A.D. Garcez, Logic tensor networks: deep learning and logical reasoning from data and knowledge, in: CEUR Workshop Proceedings 1768, 2016.
- [68] D. Silver, J. Schrittwieser, K. Simonyan, I. Antonoglou, A. Huang, A. Guez, T. Hubert, L. Baker, M. Lai, A. Bolton, et al., Mastering the game of Go without human knowledge, Nature 550 (2017) 354.
- [69] R. Socher, D. Chen, C.D. Manning, A.Y. Ng, Reasoning with neural tensor networks for knowledge base completion, in: Proc. of NIPS'13, 2013, pp. 1–10.
- [70] G. Šourek, V. Aschenbrenner, F. Železný, O. Kuželka, Lifted relational neural networks, in: CEUR Workshop Proceedings, 2015, <https://arxiv.org/pdf/1508.05128.pdf>.
- [71] G. šourek, V. Aschenbrenner, F. Železný, S. Schockaert, O. Kuželka, Lifted relational neural networks: efficient learning of latent relational structures, J. Artif. Intell. Res. 62 (2018) 69–100, <https://doi.org/10.1613/jair.1.11203>.
- [72] P.B. Vranas, Hempel's raven paradox: a lacuna in the standard Bayesian solution, Br. J. Philos. Sci. 55 (2004) 545–560, <https://doi.org/10.1093/bjps/55.3.545>.
- [73] H. Weisberg, The distribution of linear combinations of order statistics from the uniform distribution, Ann. Math. Stat. 42 (1971) 704–709, <https://doi.org/10.1214/aoms/1177693419>, <https://doi.org/10.1214/aoms/1177693419>.
- [74] J. Xu, Z. Zhang, T. Friedman, Y. Liang, G. den Broeck, A semantic loss function for deep learning with symbolic knowledge, in: J. Dy, A. Krause (Eds.), Proceedings of the 35th International Conference on Machine Learning, PMLR, Stockholm, Sweden, 2018, pp. 5502–5511, <http://proceedings.mlr.press/v80/xu18h.html>.
- [75] R.R. Yager, On a general class of fuzzy connectives, Fuzzy Sets Syst. 4 (1980) 235–242.
- [76] Z.H. Zhou, A brief introduction to weakly supervised learning, Nat. Sci. Rev. 5 (2017) 44–53.