

Quantifiers for Differentiable Logics in Rocq (Extended Abstract)

Jairo Miguel Marulanda-Giraldo¹, Ekaterina Komendantskaya^{1,2}, Alessandro
Bruni³, Reynald Affeldt⁴, Matteo Capucci⁵, and Enrico Marchioni¹

¹ University of Southampton, UK

² Heriot-Watt University, UK

³ IT-University of Copenhagen, Denmark

⁴ National Institute of Advanced Industrial Science and Technology (AIST), Japan

⁵ Independent Researcher, Italy

Abstract. The interpretation of logical expressions into loss functions has given rise to so-called differentiable logics. They function as a bridge between formal logic and machine learning, offering a novel approach for property-driven training. The added expressiveness of these logics comes at the price of a more intricate semantics for first-order quantifiers. To ease their integration into machine-learning backends, we explore how to formalize semantics for first-order differentiable logics using the Mathematical Components library in the Rocq proof assistant. We seek to give rigorous semantics for quantifiers, verify their properties with respect to other logical connectives, as well as prove the soundness and completeness of the resulting logics.

Keywords: Neural Network Verification · Formal Specifications · Loss Functions · Differentiable Logics · Interactive Theorem Proving.

1 Introduction

Quantitative logics, i.e. logics that have semantics over the real numbers instead of over $\{0, 1\}$ have been studied for decades, and date back to the ideas of Kleene, Gödel, and Łukasiewicz at the start of the 20th century [10, 22]. Fuzzy logics [22], and the logics of the Lawvere quantale [22, 5, 15] are important examples of quantitative logics. To illustrate, let us have a toy syntax with atomic propositions and conjunction, such as

$$\Phi \ni \phi := A \mid \phi \wedge \phi \tag{1}$$

where A is interpreted in a domain $D \subseteq [-\infty, \infty]$. D varies among logics and restricts the interpretation of connectives. For example, the Gödel logic has a

⁵ Komendantskaya and Capucci are funded by the Advanced Research + Invention Agency (ARIA).

standard semantics over $[0, 1]$ where the conjunction is interpreted as the minimum function.

Recently, there was a surge of interest in quantitative logics, stimulated by the growing interest in *AI safety* [13, 14]. Differentiable Logics (DLs) form a family of methods that applies key insights from quantitative logics to this domain for property-driven learning[25]. Generally, it is considered desirable to be able to use machine learning algorithms in a way that imposes certain logical specifications during training [20, 27]. Differentiable logics have been shown to effectively translate arbitrary logical specifications into real-valued and differentiable functions that, in turn, can be used as *loss functions* in standard gradient-descent algorithms [8]. Such loss functions help improve the adherence of the resulting neural networks to specifications [16]. At the same time, DLs have proven to be useful in compiling specifications for the back-ends of neural network verifiers [12], a process necessary to provide programming language support to property-driven training [11]. This calls for stronger guarantees about the correctness of such compilers, and rigorous semantics for DLs, as well as their soundness, completeness, and compositionality [1, 8, 25].

Nevertheless, there is one fundamental problem that differentiable logics face. Many specifications of interest for machine learning involve quantifiers, yet the majority of quantitative logics is propositional [5, 22, 25]. A canonical specification of this kind is *robustness* [9], i.e. small perturbations to the inputs of a neural network should result in small changes to its output, formally:

Definition 1 (ϵ - δ -Robustness). *Let $\epsilon, \delta \in \mathbb{R}^+$, $\|\cdot\|$ be a norm, and $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$ be a measurable function. One says f is ϵ - δ -robust around $\bar{x} \in \mathbb{R}^n$ if*

$$\forall x \in \mathbb{R}^n, \|x - \bar{x}\| \leq \epsilon \Rightarrow \|f(x) - f(\bar{x})\| \leq \delta \quad (2)$$

Expanding some sound and complete propositional quantitative logics to first-order logic often comes at the expense of either completeness or continuity. For example, the first-order extension of Gödel logic is the only one, among the most prominent fuzzy logics [22, 25], that is sound and complete w.r.t. models with values in $[0, 1]$ and with universal and existential quantifiers interpreted as infima and suprema [3]. However, connectives of this logic are continuous and therefore not suitable for gradient-descent algorithms.

Recently, a promising solution was proposed by Capucci: interpreting quantifiers as *p-means* [7], a generalization of *p*-norms over a probability space [6]. This new semantics gives hope that the open problem of finding a suitable approach to quantification in DLs will find its resolution, and we can soon find a logic that is sound and complete relative to this new quantitative semantics.

With rigorous semantics for quantifiers, first-order DLs could be integrated into verifier back-ends. We must hence provide guarantees of the resulting logics, as well as of quantifiers with respect the other logical connectives. Rigorous computer formalizations of propositional semantics for DLs have been used to this end [1]. Extending these formalizations to first-order logics is a non trivial challenge that is yet to be overcome. Furthermore, the new semantics proposed by Capucci presents a particular challenge for formal verification, since, unlike

the previous formalizations of DLs [1], it now also involves results from real analysis and probability. Most notably, it involves formalisations of measure spaces, probability spaces, and Lebesgue integrals, as well as the use of results such as Jensen’s and Hölder’s inequalities [19].

Rocq’s Mathematical Components library (*MathComp*) [26], is a particularly good fit for this task, due to its extensive mathematical libraries. Many of the aforementioned standard results from measure theory are formalized in the library modules on algebra and analysis. However, some, such as the encoding of extended real numbers, still require further development.

In this extended abstract, we first quickly review the approach to quantification proposed by Capucci, explain its relation to the available mathematical libraries in Rocq, and report on our current work on formalizing the novel semantics. With this formalization, we contribute towards developing the semantics for quantifiers in DLs. Tangentially, we extend *MathComp* as necessary. In the long term, this formalization is expected to become part of a larger collaborative project [2], that develops a novel first-order quantitative logic and provides its full formalization in Rocq, including, when and if that will be proven, the formalisation of the soundness and completeness results for the logic. Our work seeks to aid in the development of programming language support for property-driven development of neural networks, as well as influence machine learning research in general [12, 2].

2 Preliminaries

We introduce preliminaries from the extended arithmetic of the reals. They are an abridged version of [7], specifically we do not address the ‘non-linear’ fragment therein. We also diverge from *ibid.* in notation, preferring standard linear logic notation.

Our base setting are the positive extended reals $[0, \infty]$, considered as sup-lattice with the usual order \leq . The topology on \mathbb{R}^+ is extended to $[0, \infty]$ by adding to the opens all the intervals $(a, \infty]$. As a measure space, $[0, \infty]$ is considered equipped with completion of its Borel σ -field (i.e. the Lebesgue σ -field); and then further equipped with the obvious extension of the Lebesgue measure given by setting $\lambda((a, \infty]) = \infty$ for $a < \infty$ and $\lambda(\{\infty\}) = 0$.

Definition 2 (Multiplication). *On $[0, \infty]$, **conjunctive multiplication** and **disjunctive multiplication** are, respectively, the following operations:*

$$\begin{array}{c|ccc}
 a \otimes b & 0 & a \in (0, \infty) & \infty \\
 \hline
 0 & 0 & 0 & 0 \\
 b \in (0, \infty) & 0 & ab & \infty \\
 \infty & 0 & \infty & \infty
 \end{array}
 \qquad
 \begin{array}{c|ccc}
 a \wp b & 0 & a \in (0, \infty) & \infty \\
 \hline
 0 & 0 & 0 & \infty \\
 b \in (0, \infty) & 0 & ab & \infty \\
 \infty & \infty & \infty & \infty
 \end{array}
 \tag{3}$$

Notice \otimes and \wp differ only when a is 0 and b is ∞ , or *vice versa*. Often we write ab instead of $a \otimes b$.

Definition 3 (Duality Operator). *Let $a \in [0, \infty]$. Then the **dual** of a is*

$$a^\perp = \begin{cases} 1/a & a \in (0, \infty) \\ \infty & a = 0 \\ 0 & a = \infty \end{cases}$$

Note $a \wp b = (a^\perp \otimes b^\perp)^\perp$. Moreover we define $a \dashv b = a^\perp \wp b$, which extends the definition of b/a .

2.1 p -Means

The following definitions relate specifically to the new quantifier semantics. They are what are classically known as generalized weighted means [23], though geometric mean, much like multiplication above, bifurcates into a conjunctive and a disjunctive version.

Throughout the following, fix a probability space $(S, \Sigma_S, \mathbb{P})$.

Definition 4 (p -Means). *Let $f : S \rightarrow [0, \infty]$ be a measurable function. For $p \in (0, \infty)$, the (**generalized weighted**) p -**mean** of f is*

$$\langle f \rangle_{S,p} := \left(\int_S f(s)^p \, d\mathbb{P}(s) \right)^{1/p} \quad (4)$$

where we extended the functions $(-)^p$ as follows

$$\infty^p = \begin{cases} 1 & p = 0 \\ \infty & p > 0 \end{cases} \quad 0^p = 0. \quad (5)$$

Dually, the (**generalized weighted**) **harmonic** p -**mean** of f is

$$\langle f \rangle_{S,-p} := (\langle f^\perp \rangle_{S,p})^\perp. \quad (6)$$

When S can be inferred from the context, we write $\langle f \rangle_p$.

The definition of p -means can be extended to $p = 0$ and $p = \infty$ by taking limits [7]. First we have

Lemma 1. *As $p \rightarrow +\infty$,*

$$\langle f \rangle_{+p} \rightarrow \text{ess sup}(f) =: \langle f \rangle_{+\infty}, \quad \langle f \rangle_{-p} \rightarrow \text{ess inf}(f) =: \langle f \rangle_{-\infty}. \quad (7)$$

These quantities are so defined:

Definition 5 (Essential Extrema). *Let (S, Σ_S, μ) be a measure space and $f : S \rightarrow [0, \infty]$ a measurable function.*

1. Let $U = \{a \in [0, \infty] : \mu(\{x \in X : a < f(x)\}) = 0\}$ and $\inf(U)$ be the infimum of U . The **essential supremum** of f is

$$\text{ess sup}(f) = \inf U \quad (8)$$

recalling that $\inf \emptyset = \infty$.

2. The **essential infimum** of f is

$$\text{ess inf}(f) = -\text{ess sup}(-f) \quad (9)$$

On the other end of the spectrum, we have:

Lemma 2. As $p \rightarrow 0$, both $\langle f \rangle_{+p}$ and $\langle f \rangle_{-p}$ converge to a limit, thus defining **disjunctive** and **conjunctive geometric means**:

$$\langle f \rangle_{+p} \rightarrow \langle f \rangle_{+0}, \quad \langle f \rangle_{-p} \rightarrow \langle f \rangle_{-0}. \quad (10)$$

For bounded functions, these quantities coincide with the classical (weighted) geometric mean:

Definition 6 (Geometric Mean). Let $f : S \rightarrow [0, \infty)$ be a measurable function and (S, Σ_S, μ) a measure space. The **geometric mean** of f is

$$GM[f] = \exp\left(\frac{1}{\mu(S)} \int_S \ln f(s) d\mu(s)\right) \quad (11)$$

For unbounded functions, conjunctive and disjunctive geometric means may differ in the same way as \otimes and \wp , namely in the way they handle 0 and ∞ . See [7] for clarifications.

3 Proposed Language and its Semantics

We introduce the main ideas for first-order quantitative logic following Capucci [7], where the case is made that the positive reals support a family of substructural logics where the multiplicative connectives are interpreted as actual multiplication, and the additives as the p -norm and converge to the *actual* additives as $p \rightarrow \infty$. We stress that only a language (i.e. a syntax for formulae), and not a logic (i.e. an entailment relation), are defined therein. Here we propose a simplified version of that language which features only multiplicative connectives (in the style of classical multiplicative linear logic [17]).

For simplicity, we use the same symbols of Section 2 for our *language*.

A *first-order theory* over this language is given by a fixed set of sorts \mathcal{S} and a family of atomic predicates for each context, denoted as $\{\mathcal{A}(\mathbf{X})\}_{\mathbf{X} \in \text{List } \mathcal{S}}$. Recall a context is a finite (and possibly empty) list of typed variables $\mathbf{X} = (x_1 : X_1, \dots, x_n : X_n)$, where $X_i \in \mathcal{S}$. Then, for each context, and simultaneously over all contexts, we inductively define the set of formulae of the theory $\Phi(\mathbf{X})$

by closing the atomic predicates under duality, multiplicative conjunction, and universal and existential quantification over *fresh* variables:

$$\begin{aligned}
\Phi(\mathbf{X}) \ni \phi(\mathbf{x}) := & A(\mathbf{x}) \in \mathcal{A}(\mathbf{X}) \\
& | \phi(\mathbf{x})^\perp \\
& | \phi(\mathbf{x}) \otimes \phi(\mathbf{x}) \\
& | \forall^p(y \in Y). \psi(y, \mathbf{x}) \\
& | \exists^p(y \in Y). \psi(y, \mathbf{x})
\end{aligned} \tag{12}$$

where $p \in [0, \infty]$ and $\psi \in \Phi(Y, \mathbf{X})$, where Y is another sort. Let $\Phi = \bigcup_{\mathbf{X}} \Phi(\mathbf{X})$ be the set of formulae over arbitrary contexts. We encode multiplicative disjunction and linear implication respectively as

$$\phi_1 \wp \phi_2 := (\phi_1^\perp \otimes \phi_2^\perp)^\perp \quad \phi_1 \multimap \phi_2 := \phi_1^\perp \wp \phi_2$$

An interpretation of such a theory is given by (1) a choice of probability space $\llbracket X \rrbracket$ for each sort $X \in \mathcal{S}$, where, for a context \mathbf{X} as above, we let $\llbracket \mathbf{X} \rrbracket = \llbracket X_1 \rrbracket \times \cdots \times \llbracket X_n \rrbracket$ (as well as $\llbracket () \rrbracket = 1$); and (2) a given measurable function $\llbracket A \rrbracket : \llbracket \mathbf{X} \rrbracket \rightarrow [0, \infty]$ for each atomic predicate $A \in \mathcal{A}(\mathbf{X})$. Then the translation function (corresponding to *multiplicative semantics* in [7]) $\llbracket \cdot \rrbracket : \Phi \rightarrow [0, \infty]$ is defined inductively on the structure of formulae as follows:

$$\begin{aligned}
\llbracket \phi(\mathbf{x})^\perp \rrbracket & := \llbracket \phi(\mathbf{x}) \rrbracket^\perp \\
\llbracket \phi_1(\mathbf{x}) \otimes \phi_2(\mathbf{x}) \rrbracket & := \llbracket \phi_1(\mathbf{x}) \rrbracket \llbracket \phi_2(\mathbf{x}) \rrbracket \\
\llbracket \forall^p(y \in Y). \psi(y, \mathbf{x}) \rrbracket & := \langle \llbracket \psi(\cdot, \mathbf{x}) \rrbracket \rangle_{\llbracket Y \rrbracket, -p} \\
\llbracket \exists^p(y \in Y). \psi(y, \mathbf{x}) \rrbracket & := \langle \llbracket \psi(\cdot, \mathbf{x}) \rrbracket \rangle_{\llbracket Y \rrbracket, p}
\end{aligned} \tag{13}$$

Hence the semantics is defined w.r.t. the quantale $[0, \infty]_\otimes$ we described above (definition 2, [7]), which we note is isomorphic to the *Lawvere quantale* introduced in [21] and central in [5, 4].

As an example of the usefulness of this semantics, we can use it to construct the *softmax operator* [24], using the same logical formulae used for *argmax*, as shown in [7]. Indeed, suppose $f : S \rightarrow [0, \infty)$ is a measurable function we want to express the softmax of. The first-order theory of softmax has one sort X and a single atomic predicate $\phi(x) \in \mathcal{A}(X)$. We target f by interpreting X as S and

set $\llbracket \phi(\bar{x}) \rrbracket = f(\bar{s})$. Then the softmax of f is obtained as follows:

$$\begin{aligned}
(\text{softmax } f)(\bar{s}) &= \llbracket (\exists^1(x \in X). \phi(x)) \multimap \phi(\bar{x}) \rrbracket \\
&= \llbracket (\exists^1(x \in X). \phi(x)) \rrbracket^\perp \wp \llbracket \phi(\bar{x}) \rrbracket \\
&= (\llbracket \phi(x) \rrbracket_{S,1})^\perp \wp \llbracket \phi(\bar{x}) \rrbracket \\
&= \left(\int_S \llbracket \phi \rrbracket(s) d\mathbb{P}(s) \right)^\perp \wp \llbracket \phi(\bar{x}) \rrbracket \\
&= \left(\int_S f(s) d\mathbb{P}(s) \right)^\perp \wp f(\bar{s}) \\
&= \frac{f(\bar{s})}{\int_S f(s) d\mathbb{P}(s)}
\end{aligned} \tag{14}$$

Note that often $f = \exp(-\beta u)$ for some scoring function $u : S \rightarrow [-\infty, \infty]$ and inverse temperature $\beta \in (0, \infty]$ —this is the form most common in machine learning [18]. Similarly, $\bar{s} \in \text{argmax } f \iff \llbracket (\exists^\infty(x \in X). \phi(x)) \multimap \phi(\bar{x}) \rrbracket \geq 1$.

For a second example, we show how the robustness property of Definition 1 can be encoded in such a language. Since this is usually a ‘hard’ predicate, we have many choices on how to approach it as a soft predicate, here we give a very crude such encoding, parametrised by the given constants $\epsilon, \delta \in \mathbb{R}$, the function $f : \mathbb{R}^m \rightarrow \mathbb{R}^n$, the point $\bar{x} \in \mathbb{R}^n$, as well as by a ‘softness degree’ $p \in [0, \infty]$. Thus we look at a first-order theory with one sort X and predicates $E, D \in \mathcal{A}(X)$, and we interpret it by setting

$$\llbracket X \rrbracket = \mathbb{R}^m, \quad \llbracket E \rrbracket = \mathbf{1}_{\{x \in \mathbb{R}^n \mid \|x - \bar{x}\| \leq \epsilon\}}, \quad \llbracket D \rrbracket = \mathbf{1}_{\{x \in \mathbb{R}^n \mid \|f(x) - f(\bar{x})\| \leq \delta\}}. \tag{15}$$

where $\mathbf{1}_A$ denotes the indicator function of a measurable set A . Then (2) is

$$\llbracket \forall^p(x \in X).(E(x) \multimap D(x)) \rrbracket = \left(\int_{\mathbb{R}^m} \left(\frac{\llbracket E \rrbracket(s)}{\llbracket D \rrbracket(s)} \right)^p d\mathbb{P}(s) \right)^{-1/p} \tag{16}$$

4 Properties of Quantifiers

In the machine learning community there is a general consensus on the desirable properties of loss functions—convexity or continuity are widely considered desirable [20]. From a programming language perspective, there is no consensus as to how to define soundness for quantitative logics. In the future, we intend to follow the general approach applied by Slusarz *et al.* Moreover, Varnai and Dimarogonas suggest characterizing quantitative logics in terms of their *geometric properties*, valuable for optimization tasks [27]. As for quantifiers, we wish that our formulation possesses good numerical properties, as well as behave similarly to quantifiers in classical logic. Currently we are working to formalize and prove the following properties in Rocq, which were presented by Capucci [7].

Through the following, Let \mathbf{X} be a context, Y a sort, and $\phi(\mathbf{x}) \in \Phi(\mathbf{X})$, $\psi_i(y, \mathbf{x}) \in \Phi(Y, \mathbf{X})$.

Lemma 3 (Duality).

1. $\llbracket \forall^p(y \in Y).\psi(y, \mathbf{x}) \rrbracket = \llbracket \exists^p(y \in Y).\psi(y, \mathbf{x})^\perp \rrbracket^\perp$
2. $\llbracket \exists^p(y \in Y).\psi(y, \mathbf{x}) \rrbracket = \llbracket \forall^p(y \in Y).\psi(y, \mathbf{x})^\perp \rrbracket^\perp$

Lemma 4 (Distributivity over Implication).

1. $\llbracket \phi(\mathbf{x}) \multimap \forall^p(y \in Y).\psi(y, \mathbf{x}) \rrbracket = \llbracket \forall^p(y \in Y).(\phi(\mathbf{x}) \multimap \psi(y, \mathbf{x})) \rrbracket$
2. $\llbracket \forall^p(y \in Y).(\psi(y, \mathbf{x}) \multimap \phi(\mathbf{x})) \rrbracket = \llbracket (\exists^p(y \in Y).\psi(y, \mathbf{x})) \multimap \phi(\mathbf{x}) \rrbracket$.

These lemmas will be potentially useful to prove the *residuation property*, an important feature of many quantitative logics [17].

Lemma 5 (Abductive). *Let $\llbracket Z \rrbracket \subseteq \llbracket Y \rrbracket$ then*

1. $\llbracket \exists^p(z \in Z).\psi(z, \mathbf{x}) \rrbracket \leq \llbracket \exists^p(y \in Y).\psi(y, \mathbf{x}) \rrbracket$
2. $\llbracket \forall^p(y \in Y).\psi(y, \mathbf{x}) \rrbracket \leq \llbracket \forall^p(z \in Z).\psi(z, \mathbf{x}) \rrbracket$.

Intuitively, confidence depends on the amount of evidence.

The following are often desirable properties of loss functions.

Lemma 6 (Monotonic). *If $\llbracket \psi_1 \rrbracket \leq \llbracket \psi_2 \rrbracket$ then*

1. $\llbracket \exists^p(y \in Y).\psi_1(y, \mathbf{x}) \rrbracket \leq \llbracket \exists^p(y \in Y).\psi_2(y, \mathbf{x}) \rrbracket$
2. $\llbracket \forall^p(y \in Y).\psi_1(y, \mathbf{x}) \rrbracket \leq \llbracket \forall^p(y \in Y).\psi_2(y, \mathbf{x}) \rrbracket$.

Lemma 7 (p -Monotonic and Bounded). *If $0 \leq q \leq p$ then*

1. $\llbracket \exists^q(y \in Y).\psi(y, \mathbf{x}) \rrbracket \leq \llbracket \exists^p(y \in Y).\psi(y, \mathbf{x}) \rrbracket \leq \llbracket \exists^\infty(y \in Y).\psi(y, \mathbf{x}) \rrbracket$
2. $\llbracket \forall^\infty(y \in Y).\psi(y, \mathbf{x}) \rrbracket \leq \llbracket \forall^p(y \in Y).\psi(y, \mathbf{x}) \rrbracket \leq \llbracket \forall^q(y \in Y).\psi(y, \mathbf{x}) \rrbracket$.

Hence we can approximate the quantifier semantics of Gödel logic while maintaining differentiability.

5 Work in progress on the Rocq formalization

In “Taming Differentiable Logics with Coq Formalisation” a formalization for several quantitative logics was developed [1]. We seek to expand this formalization so that it is suitable for reasoning about first-order DLs, with p -means as the semantics for quantifiers. So far we have formalized the semantics presented in section 3, and some basic properties of the p -means. To illustrate, we present the encodings needed for lemma 3. Note the following implementations have been simplified for clarity.

To encode the p -mean, we make use of the `Lnorm`, `MathComp`’s encoding of the p -norm [6], and add an encoding for the geometric mean.


```

Definition Lnorm P p f :=
  match p with
  | p%:E => (\int[μ]_x `|f x| ^ p) ^ p^-1
  | +oo => ess_sup P (abse \o f)
  | -oo => ess_inf P (abse \o f)
end.

```

```

Definition geo_mean P f :=
  expeR \int[P]_x (lne (f x)).

```

```

Definition pmean P p f :=
  if p == 0 then geo_mean P f else Lnorm P p f.

```

Where `ess_sup`, `ess_inf`, `geo_mean`, and `pmean` correspond respectively to the essential supremum, essential infimum, geometric mean and p -mean. For the dual, we use *MathComp*'s power function.

```

Definition dual a := if a == 0 then +oo else x ^ -1.

```

We can represent quantifiers in terms of the previous encodings, and add notations for clarity.

```

Notation "x ^'" := (cdual x).
Notation "'forall_ p f " := (pmean P p f).
Notation "'exists_ p f " := (('forall_p (fun y => (f y) ^')) ^').

```

Lastly, lemma 3 is encoded as Lemma Duality, using the facts that the dual is idempotent and the harmonic p -mean non-negative, encoded as Lemma `idem_dual` and Lemma `forall_gt0`, respectively.

```

Lemma Duality p x :
  (0 < p) ->
  'forall_p (psi x) = ('exists_p (fun y => (psi x y) ^')) ^'.

```

Proof.

```

by move=> ?; rewrite (*this is true since*)
  idem_dual // = (*the dual is idempotent and*)
  ?forall_gt0 //; (*the harmonic p-mean is non-negative and*)
  under eq_fun do rewrite (*in the body of the harmonic p-mean*)
  idem_dual //. (*the dual is idempotent.*)

```

Qed.

To formalize the rest of section 4 in Rocq, as well as the lemmas in section 2.1, we are currently working on extending the analysis module of *MathComp*. In particular, Hölder's inequalities must be generalized to functions that go to the extended reals. In this process, we noticed the original encoding of the power function over extended real numbers incorrectly assumed its exponent is a real number greater than or equal to zero. The implementation has now been generalized for negative exponents.

6 Conclusions and Future Work

In this extended abstract we described our work in progress. We presented the main ideas behind a first-order quantitative logic to be applied in AI verification. We presented a promising translation for quantifiers and introduced some desirable properties for this translation, following closely [7]. We argued for the usefulness of a computer formalization to provide compilation guarantees. Lastly, we presented some preliminary progress in formalization of these results in Rocq.

In the future we hope to:

1. Develop a Hilbert and Sequent Calculus for the language.
2. Prove soundness and completeness for the resulting logic.
3. Formalize the properties mentioned in section 4 and the resulting proofs of soundness and completeness.
4. Test the performance of the logic for property-driven training.
5. Integrate our results into verification back-ends such as that of Vehicle [12].

7 Acknowledgements

J. Marulanda-Giraldo and E. Komendantskaya acknowledge the partial support of the EPSRC grant AISEC: AI Secure and Explainable by Construction (EP/T026960/1). M. Capucci and E. Komendantskaya were supported by ARIA: Mathematics for Safe AI grant. J. Marulanda-Giraldo received PhD Scholarship from the University of Southampton.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Affeldt, R., Bruni, A., Komendantskaya, E., Ślusarz, N., Stark, K.: Taming Differentiable Logics with Coq Formalisation. In: Bertot, Y., Kutsia, T., Norrish, M. (eds.) 15th International Conference on Interactive Theorem Proving (ITP 2024). Leibniz International Proceedings in Informatics (LIPIcs), 4:1–4:19. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany (2024). <https://doi.org/10.4230/LIPIcs.ITP.2024.4>. <https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.ITP.2024.4>
2. Atkey, R., Capucci, M., Komendantskaya, E., Mardare, R.: Quantitative Predicate Logic as a Foundation for Verified ML. ARIA grant (2024)
3. Baaz, M., Preining, N., Zach, R.: First-order Gödel logics. *Annals of Pure and Applied Logic* **147**(1), 23–47 (2007). <https://doi.org/10.1016/j.apal.2007.03.001>. <https://www.sciencedirect.com/science/article/pii/S016800720700019X>
4. Bacci, G., Mardare, R., Panangaden, P., Plotkin, G.: Polynomial Lawvere Logic, (2024). arXiv: 2402.03543. <http://arxiv.org/abs/2402.03543> (visited on 11/01/2024). Pre-published. <https://doi.org/10.48550/arXiv.2402.03543>

5. Bacci, G., Mardare, R., Panangaden, P., Plotkin, G.: Propositional logics for the Lawvere quantale. *Electronic Notes in Theoretical Informatics and Computer Science* **3** (2023)
6. Bourgain, J.: *New classes of Lp-spaces*. Springer (2006)
7. Capucci, M.: On Quantifiers for Quantitative Reasoning. arXiv preprint arXiv:2406.04936 (2024)
8. Casadio, M., Komendantskaya, E., Daggitt, M.L., Kokke, W., Katz, G., Amir, G., Refaeli, I.: Neural network robustness as a verification property: a principled case study. In: *International conference on computer aided verification*, pp. 219–231 (2022)
9. Casadio, M., Komendantskaya, E., Daggitt, M.L., Kokke, W., Katz, G., Amir, G., Refaeli, I.: Neural Network Robustness as a Verification Property: A Principled Case Study, (2022). arXiv: 2104.01396 [cs.LG]. <https://arxiv.org/abs/2104.01396>.
10. Cintula, P., Hájek, P., Noguera, C., *et al.*: *Handbook of Mathematical Fuzzy Logic. Volume 1*. College Publications (2011)
11. Daggitt, M., Kokke, W., Komendantskaya, E., Atkey, R., Arnaboldi, L., Slusarz, N., Casadio, M., Coke, B., Lee, J.: The Vehicle Tutorial: Neural Network Verification with Vehicle. In: *Narodytska, N., Amir, G., Katz, G., Isac, O. (eds.) Proceedings of the 6th Workshop on Formal Methods for ML-Enabled Autonomous Systems*. Kalpa Publications in Computing, pp. 1–5. EasyChair (2023). <https://doi.org/10.29007/5s2x>. /publications/paper/Rkrv
12. Daggitt, M.L., Kokke, W., Atkey, R., Slusarz, N., Arnaboldi, L., Komendantskaya, E.: Vehicle: Bridging the Embedding Gap in the Verification of Neuro-Symbolic Programs. arXiv preprint arXiv:2401.06379 (2024)
13. Dalrymple, D.: Safeguarded AI: constructing guaranteed safety, (2024). <https://www.aria.org.uk/media/3nhijno4/aria-safeguarded-ai-programme-thesis-v1.pdf>. Programme Thesis.
14. Dalrymple, D.", Skalse, J., Bengio, Y., Russell, S., Tegmark, M., Seshia, S., Omohundro, S., Szegedy, C., Goldhaber, B., Ammann, N., Abate, A., Halpern, J., Barrett, C., Zhao, D., Zhi-Xuan, T., Wing, J., Tenenbaum, J.: Towards Guaranteed Safe AI: A Framework for Ensuring Robust and Reliable AI Systems, (2024). arXiv: 2405.06624 [cs.AI]. <https://arxiv.org/abs/2405.06624>.
15. Fischer, M., Balunovic, M., Drachler-Cohen, D., Gehr, T., Zhang, C., Vechev, M.: DL2: training and querying neural networks with logic. In: *International Conference on Machine Learning*, pp. 1931–1941 (2019)
16. Flinkow, T., Pearlmutter, B.A., Monahan, R.: Comparing Differentiable Logics for Learning with Logical Constraints. arXiv preprint arXiv:2407.03847 (2024)
17. Galatos, N., Jipsen, P., Kowalski, T., Ono, H.: *Residuated lattices: an algebraic glimpse at substructural logics*. Elsevier (2007)
18. Goodfellow, I., Bengio, Y., Courville, A.: *Deep Learning*. MIT Press (2016)
19. Kazarinoff, N.D.: *Analytic inequalities*. Courier Corporation (2014)
20. Klebanov, L.B., Rachev, S.T., Fabozzi, F.J.: *Robust and non-robust models in statistics*. Nova Science Publishers Hauppauge (2009)
21. Lawvere, F.W.: Metric Spaces, Generalized Logic, and Closed Categories. *Rendiconti del seminario matematico e fisico di Milano* **43**, 135–166 (1973)
22. Metcalfe, G., Olivetti, N., Gabbay, D.M.: *Proof theory for fuzzy logics*. Springer Science & Business Media (2008)
23. Mitrinovic, D.S., Vasic, P.M.: *Analytic inequalities*. Springer (1970)
24. Ren, J., Wang, H.: Chapter 3 - Calculus and optimization. In: *Mathematical Methods in Data Science*. Ed. by J. Ren and H. Wang, pp. 51–89. Elsevier (2023). <https://>

- [//doi.org/10.1016/B978-0-44-318679-0.00009-0](https://doi.org/10.1016/B978-0-44-318679-0.00009-0). <https://www.sciencedirect.com/science/article/pii/B9780443186790000090>
25. Slusarz, N., Komendantskaya, E., Daggitt, M.L., Stewart, R.J., Stark, K.: Logic of Differentiable Logics: Towards a Uniform Semantics of DL. In: Piskac, R., Voronkov, A. (eds.) LPAR 2023: Proceedings of 24th International Conference on Logic for Programming, Artificial Intelligence and Reasoning, Manizales, Colombia, 4-9th June 2023. EPiC Series in Computing, pp. 473–493. EasyChair (2023). <https://doi.org/10.29007/C1NT>
 26. Team, M.C.: Mathematical components library, <https://github.com/math-comp/math-comp> (2007).
 27. Varnai, P., Dimarogonas, D.V.: On robustness metrics for learning STL tasks. In: 2020 American Control Conference (ACC), pp. 5394–5399 (2020)