

Extracting Formulae in Many-Valued Logic from Deep Neural Networks

Yani Zhang and Helmut Bölcskei

Abstract

We propose a new perspective on deep rectified linear unit (ReLU) networks, namely as circuit counterparts of Łukasiewicz infinite-valued logic—a many-valued (MV) generalization of Boolean logic. An algorithm¹ for extracting formulae in MV logic from trained deep ReLU networks is presented. The algorithm respects the network architecture, in particular compositionality, thereby honoring algebraic information present in the training data. We also establish the representation benefits of deep networks from a mathematical logic perspective.

Index Terms

Mathematical logic, many-valued logic, McNaughton functions, deep neural networks

I. INTRODUCTION AND CONTRIBUTIONS

State-of-the-art deep neural networks often exhibit remarkable reasoning capabilities, e.g., in mathematical tasks [1], program synthesis [2], and algorithmic reasoning [3]. In an attempt to develop an understanding of this phenomenon, the present paper establishes a systematic connection between neural networks and mathematical logic. Specifically, we shall be interested in reading out logical formulae from (trained) deep neural networks, with the underlying idea that these formulae yield a logical description of the data the network was trained on.

Consider a neural network that realizes a map $f : [0, 1]^n \rightarrow [0, 1]$. Taking a step back, we note that for input variables and the output taking on the values 0 and 1 only, f reduces to a Boolean

The authors are with the Chair for Mathematical Information Science, ETH Zurich, Switzerland. H. Bölcskei gratefully acknowledges support by the Lagrange Mathematics and Computing Research Center, Paris, France.

¹A Python implementation for ReLU networks with integer, rational, or real weights is available at <https://www.mins.ee.ethz.ch/research/downloads/NN2MV.html>

function and can hence be studied by means of Boolean algebra, see e.g., [4]. The algebraic expression of this Boolean function—in terms of AND, OR, and NOT operations—provides a logical description of the underlying functional relationship. Here, we shall be concerned with the generalization of this correspondence from Boolean functions $f : \{0, 1\}^n \rightarrow \{0, 1\}$ to general functions $f : [0, 1]^n \rightarrow [0, 1]$. This immediately leads to the following questions:

- 1) What is the logical system replacing Boolean logic?
- 2) How can formulae in this logical system be extracted from neural networks realizing $f : [0, 1]^n \rightarrow [0, 1]$?

As for the first question, we shall show that the theory of infinite-valued Łukasiewicz logic [5], sometimes also referred to as Łukasiewicz-Tarski logic, provides a suitable framework for characterizing general functions $f : [0, 1]^n \rightarrow [0, 1]$ from a logical perspective. With slight abuse of terminology, we shall refer to infinite-valued Łukasiewicz logic as many-valued (MV) logic throughout the paper. Based on a fundamental result [6], which characterizes the class of truth functions in MV logic—also called McNaughton functions—as continuous piecewise linear functions with integer coefficients, we show that neural networks with the ReLU nonlinearity $\rho(x) = \max\{0, x\}$ and integer weights² naturally implement statements in MV logic and vice versa. We develop an algorithm for extracting formulae in MV logic from McNaughton functions specified through deep ReLU networks. There are two procedures in the literature for extracting MV formulae from McNaughton functions, namely the Schauder hat method [7] and the hyperplane method [8], both developed in the course of alternative and constructive proofs of the McNaughton theorem [6] and applying to the one-dimensional case $n = 1$ only, while our algorithm works for arbitrary dimension n . Moreover, for $n = 1$, by virtue of honoring the compositional structure of deep ReLU networks, the algorithm we propose can result in significantly shorter formulae than [7] and [8]. This will be established analytically and illustrated through numerical results.

In practice, trained neural networks will not exhibit integer weights, unless explicitly enforced in the training process. Extensions of MV logic, namely Rational Łukasiewicz logic [9] and $\mathbb{R}\mathcal{L}$ [10], have truth functions that are again continuous piecewise linear, but with rational and real coefficients, respectively. Such functions are likewise naturally realized by ReLU networks, but correspondingly with rational and real weights. For pedagogical reasons and to render the

²By weights, we mean the entries of the weight matrices and bias vectors associated with the network.

presentation more accessible, we first present the entire program outlined above, including the algorithm for extracting MV formulae from ReLU networks, for the case of integer weights and then describe the extensions to the rational and real cases.

The ideas presented in this paper are inspired by [11]–[13]. Specifically, Amato et al. [11], [12] pointed out that neural networks, with the clipped ReLU (CReLU) nonlinearity $\sigma(x) = \min\{1, \max\{0, x\}\}$ and rational weights, realize truth functions in Rational Łukasiewicz logic. Di Nola et al. [13] proved that CReLU networks with real weights realize truth functions in $\mathbb{R}\mathcal{L}$ logic. The universal correspondence between ReLU networks and MV logic, Rational Łukasiewicz logic, and $\mathbb{R}\mathcal{L}$ along with the algorithm for extracting logical formulae from ReLU networks appear to be entirely new.

II. BOOLEAN LOGIC AND MV LOGIC

We start with a brief review of the basic concepts in Boolean and MV logic. Boolean algebra on the set $\{0, 1\}$ consists of the application of the logical operations AND, OR, and NOT, denoted by \odot , \oplus , and \neg , respectively, to Boolean propositional variables taking value in $\{0, 1\}$. The algebra is fully characterized through the following identities on Boolean variables x, y, z :

$$\begin{aligned}
x \oplus 0 &= x & x \odot 1 &= x \\
x \oplus \neg x &= 1 & x \odot \neg x &= 0 \\
x \oplus y &= y \oplus x & x \odot y &= y \odot x \\
(x \oplus y) \oplus z &= x \oplus (y \oplus z) & & (1) \\
(x \odot y) \odot z &= x \odot (y \odot z) \\
x \oplus (y \odot z) &= (x \oplus y) \odot (x \oplus z) \\
x \odot (y \oplus z) &= (x \odot y) \oplus (x \odot z)
\end{aligned}$$

Note that one can define the operation \odot in terms of \oplus and \neg according to $x \odot y := \neg(\neg x \oplus \neg y)$.

MV logic [14] generalizes Boolean logic by allowing for propositional variables that take truth values in the interval $[0, 1]$. The corresponding algebraic counterpart is known as Chang’s MV algebra [15] (see Definition A.1). We proceed to the definition of the so-called standard MV algebra.

Definition 2.1: Consider the unit interval $[0, 1]$, and define $x \oplus y = \min\{1, x + y\}$ and $\neg x = 1 - x$, for $x, y \in [0, 1]$. It can be verified that the structure $\mathcal{I} = \langle [0, 1], \oplus, \neg, 0 \rangle$ is an MV

algebra [5]. In particular, \mathcal{I} constitutes the algebraic counterpart of Łukasiewicz infinite-valued logic [15]. We further define the operation $x \odot y := \neg(\neg x \oplus \neg y) = \max\{0, x + y - 1\}$.

It can be shown that the Boolean algebra $\mathcal{B} := \langle \{0, 1\}, \text{OR}, \text{NOT}, 0 \rangle$ is a special case of MV algebras. The MV algebra \mathcal{I} in Definition 2.1 is referred to as standard because an equation holds in every MV algebra iff it holds in \mathcal{I} [15], [16]. Additional relevant material on MV algebras is provided in Appendix A.

III. EXTRACTION OF FORMULAE IN MV LOGIC FROM RELU NETWORKS

MV terms (see Definition A.2) are finite strings composed of propositional variables x_1, x_2, \dots connected by \oplus , \odot , and \neg operations and brackets $()$, such as e.g., $(x_1 \oplus \neg x_2) \odot x_3$. We define the length of an MV term as the total number of occurrences of propositional variables, e.g., $(x_1 \odot \neg x_1) \oplus x_2$ is of length 3. Term functions (see Definition A.3) are the corresponding truth functions obtained by interpreting the logical operations according to how they are specified in the concrete MV algebra used, e.g., $x \oplus y = \min\{1, x + y\}$ in the standard MV algebra \mathcal{I} . In the Boolean algebra \mathcal{B} , term functions are binary tables $\{\{0, 1\}^n \rightarrow \{0, 1\} : n \in \mathbb{N}\}$. In the standard MV algebra \mathcal{I} , term functions are characterized by continuous piecewise linear functions with integer coefficients as formalized by the McNaughton theorem [6].

Theorem 3.1 ([6]): Consider the MV algebra \mathcal{I} in Definition 2.1. Let $n \in \mathbb{N}$. For a function $f : [0, 1]^n \rightarrow [0, 1]$ to have a corresponding MV term τ such that the associated term function $\tau^{\mathcal{I}}$ satisfies $\tau^{\mathcal{I}} = f$ on $[0, 1]^n$, it is necessary and sufficient that f satisfy the following conditions:

- 1) f is continuous with respect to the natural topology on $[0, 1]^n$,
- 2) there exist linear polynomials p_1, \dots, p_ℓ with integer coefficients, i.e.,

$$p_j(x_1, \dots, x_n) = m_{j1}x_1 + \dots + m_{jn}x_n + b_j,$$

for $j = 1, \dots, \ell$, with $m_{j1}, \dots, m_{jn}, b_j \in \mathbb{Z}$, such that for every $x \in [0, 1]^n$, there is a $j \in \{1, \dots, \ell\}$ with $f(x) = p_j(x)$.

Functions satisfying these conditions are called McNaughton functions.

ReLU networks (see Definition B.1) are compositions of affine transformations and the ReLU nonlinearity $\rho(x) = \max\{0, x\}$ (applied element-wise) and as such realize continuous piecewise linear functions. Specifically, the class of ReLU networks with integer weights is equivalent to the class of formulae in MV logic. The corresponding formal statement is as follows.

Theorem 3.2: For $n \in \mathbb{N}$, let $\tau(x_1, \dots, x_n)$ be an MV term in n variables with $\tau^{\mathcal{I}} : [0, 1]^n \rightarrow [0, 1]$ the associated term function in \mathcal{I} . There exists a ReLU network Φ with integer weights, satisfying

$$\Phi(x_1, \dots, x_n) = \tau^{\mathcal{I}}(x_1, \dots, x_n),$$

for all $(x_1, \dots, x_n) \in [0, 1]^n$. Conversely, for every ReLU network $\Phi : [0, 1]^n \rightarrow [0, 1]$ with integer weights, there exists an MV term $\tau(x_1, \dots, x_n)$ whose associated term function in \mathcal{I} satisfies

$$\tau^{\mathcal{I}}(x_1, \dots, x_n) = \Phi(x_1, \dots, x_n),$$

for all $(x_1, \dots, x_n) \in [0, 1]^n$.

The remainder of this section is devoted to the proof of Theorem 3.2, along with the development of an algorithm for extracting formulae in MV logic from ReLU networks with integer weights.

First, we show how, for a given MV term τ , a ReLU network with integer weights realizing the associated term function $\tau^{\mathcal{I}}$ can be constructed. To this end, we note that the operation $\neg x = 1 - x$, by virtue of being affine, is trivially realized by a ReLU network. Further, there exist ReLU networks Φ^{\oplus} and Φ^{\odot} , with integer weights, realizing the \oplus and \odot operations in \mathcal{I} , i.e.,

$$\Phi^{\oplus}(x, y) = \min\{1, x + y\}$$

$$\Phi^{\odot}(x, y) = \max\{0, x + y - 1\},$$

for all $x, y \in [0, 1]$. Detailed constructions of Φ^{\oplus} and Φ^{\odot} are provided in Lemma B.3. According to Lemma B.2 [17], compositions of ReLU networks are again ReLU networks. The ReLU network realizing the term function associated with the MV term τ can hence be obtained by concatenating ReLU networks implementing the operations \oplus, \odot , and \neg as they appear in the expression for τ . What is more, inspection of the proof of Lemma B.2 reveals that the integer-valued nature of the weights is preserved under composition. Therefore, the resulting overall ReLU network has integer weights. We illustrate this procedure by way of the simple example $\tau = (x \oplus x) \odot \neg y$, which yields the associated ReLU network (see Appendix F for details)

$$\Phi^{\tau} = W_3 \circ \rho \circ W_2 \circ \rho \circ W_1, \tag{2}$$

where

$$W_1(x, y) = \begin{pmatrix} -2 & 0 \\ 0 & 1 \\ 0 & -1 \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} + \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}, \quad x, y \in \mathbb{R},$$

$$W_2(x) = (-1 \quad -1 \quad 1) x + 1, \quad x \in \mathbb{R}^3,$$

$$W_3(x) = x, \quad x \in \mathbb{R}.$$

The proof of the converse statement in Theorem 3.2 will be effected in a constructive manner, in the process developing an algorithm for extracting the MV term encoded by a given ReLU network with integer weights. For convenience of exposition, henceforth we apply ρ to the output of ReLU networks (see Definition B.1), e.g., we write $\rho \circ W_2 \circ \rho \circ W_1$ instead of $W_2 \circ \rho \circ W_1$; this does not affect the network input-output relation as we only consider networks realizing McNaughton functions and McNaughton functions always map to $[0, 1]$.

Honoring the compositional structure of deep ReLU networks, our algorithm proceeds in a compositional manner. Specifically, it consists of the following three steps.

Step 1: Transform the ReLU network (with integer weights) into an equivalent network employing the CReLU nonlinearity $\sigma(x) = \min\{1, \max\{0, x\}\}$. This is done by exploiting the fact that the domain of the ReLU network is the unit cube $[0, 1]^n$ and, consequently, with finite-valued weights, the input to all neurons in the network is bounded. Concretely, if the input of a given ρ -neuron is contained in the interval $[-A, B]$, with $A, B \in \mathbb{R}_{\geq 0}$ and $B \geq A$, we replace this ρ -neuron by one or multiple σ -neurons (shifted by integer values) according to

$$\rho(t) = \begin{cases} \sigma(t), & B \in (0, 1] \\ \sigma(t) + \sigma(t - 1), & B \in (1, 2] \\ \sigma(t) + \sigma(t - 1) + \cdots + \sigma(t - \lfloor B \rfloor), & B > 2. \end{cases} \quad (3)$$

Neurons that have $B = 0$ are deleted. Applying these replacements to all neurons in the ρ -network, we obtain a σ -network with integer weights and the same input-output relation as the ρ -network.

Step 2: Extract MV terms from individual σ -neurons, which are of the form

$$\sigma(m_1 x_1 + \cdots + m_n x_n + b), \quad (4)$$

with $m_1, \dots, m_n, b \in \mathbb{Z}$. The following lemma, a proof of which can be found in Appendix C, forms the basis for accomplishing this in an iterative manner.

Lemma 3.3 ([7], [18]): Consider the function $f(x_1, \dots, x_n) = m_1x_1 + \dots + m_nx_n + b$, $(x_1, \dots, x_n) \in [0, 1]^n$, with $m_1, \dots, m_n, b \in \mathbb{Z}$. W.l.o.g. assume that $\max_{i=1}^n |m_i| = m_1$. Let $f_\circ(x_1, \dots, x_n) = (m_1 - 1)x_1 + m_2x_2 + \dots + m_nx_n + b$. Then,

$$\sigma(f) = (\sigma(f_\circ) \oplus x_1) \odot \sigma(f_\circ + 1). \quad (5)$$

Before proceeding to the next step, we demonstrate the iterative application of Lemma 3.3 to the example σ -neuron $\sigma(x_1 - x_2 + x_3 - 1)$. First, we eliminate the variable x_1 according to

$$\sigma(x_1 - x_2 + x_3 - 1) = (\sigma(-x_2 + x_3 - 1) \oplus x_1) \odot \sigma(-x_2 + x_3). \quad (6)$$

Next, we remove x_3 inside $\sigma(-x_2 + x_3 - 1)$ and $\sigma(-x_2 + x_3)$ according to

$$\sigma(-x_2 + x_3 - 1) = (\sigma(-x_2 - 1) \oplus x_3) \odot \sigma(-x_2) \quad (7)$$

$$\sigma(-x_2 + x_3) = (\sigma(-x_2) \oplus x_3) \odot \sigma(-x_2 + 1). \quad (8)$$

We then note that $\sigma(-x_2) = 0$ as $x_2 \in [0, 1]$, and owing to $x \odot 0 = 0$, for $x \in [0, 1]$, (7) reduces to $\sigma(-x_2 + x_3 - 1) = 0$. Likewise, in (8) we get $\sigma(-x_2) \oplus x_3 = x_3$. We can now further simplify (8) according to

$$x_3 \odot \sigma(-x_2 + 1) = x_3 \odot (1 - \sigma(x_2)) \quad (9)$$

$$= x_3 \odot \neg x_2, \quad (10)$$

where in (9) we used $\sigma(x) = 1 - \sigma(-x + 1)$, for $x \in \mathbb{R}$, and (10) is by $\sigma(x) = x$ and $\neg x = 1 - x$, both for $x \in [0, 1]$. Substituting the simplified results of (7) and (8) back into (6), we obtain the MV term corresponding to $\sigma(x_1 - x_2 + x_3 - 1)$ as $x_1 \odot (x_3 \odot \neg x_2)$.

For later use, we also state a result on the length of MV terms obtained through iterative application of Lemma 3.3.

Lemma 3.4: Consider $f(x_1, \dots, x_n) = m_1x_1 + \dots + m_nx_n + b$, $(x_1, \dots, x_n) \in [0, 1]^n$, with $m_1, \dots, m_n, b \in \mathbb{Z}$. The MV term corresponding to the function $\sigma(f)$, obtained by iteratively applying Lemma 3.3 has length at most $2^m - 1$, where $m := \sum_{i=1}^n |m_i|$.

The proof of Lemma 3.4 can be found in Appendix E.

Step 3: Compose the MV terms corresponding to the individual σ -neurons according to the layered structure of the CReLU network to get the MV term associated with the ReLU network. To illustrate this step, suppose that the neurons $\sigma^{(1)}(\cdot)$ and $\sigma^{(2)}(\cdot)$ have associated MV terms $\tau^{(1)}$ and $\tau^{(2)}$, respectively, and a third neuron $\sigma^{(3)}(m_1x_1 + m_2x_2 + b)$ has associated MV

term $\tau^{(3)}(x_1, x_2)$. The MV term corresponding to the CReLU network $\sigma^{(3)}(m_1\sigma^{(1)} + m_2\sigma^{(2)} + b)$ is obtained by replacing all occurrences of x_1 in $\tau^{(3)}$ by $\tau^{(1)}$ and all occurrences of x_2 by $\tau^{(2)}$. This finalizes the proof of Theorem 3.2.

The essence of the proof of Theorem 3.2 resides in a strong algebraic property shared by the standard MV algebra \mathcal{I} and ReLU networks. Concretely, compositions of ReLU networks with integer weights again yield ReLU networks with integer weights, and compositions of formulae in MV logic result in formulae in MV logic. As we shall see in Section VI, this parallelism extends to the cases of ReLU networks with rational weights and Rational Łukasiewicz logic as well as ReLU networks with real weights and $\mathbb{R}\mathcal{L}$.

IV. EXISTING RESULTS ON MV TERM EXTRACTION FROM McNAUGHTON FUNCTIONS

Recalling that a ReLU network with integer weights realizes a continuous piecewise linear function with integer coefficients, the algorithm devised in the previous section can equivalently be seen as extracting an MV term from the McNaughton function realized by the network. As mentioned in the introduction, there are two algorithms in the literature for extracting MV terms from McNaughton functions, namely the Schauder hat method [7] and the hyperplane method [8], both of which apply, however, only to the one-dimensional case $n = 1$. The basic tenets of these two algorithms can be summarized as follows.

The Schauder hat method [7] constructs so-called Schauder hats, i.e., functions of pyramidal shape, supported on unions of simplices. More concretely, a simplicial complex over $[0, 1]$ obtained by splitting the unit cube according to different permutations of the linear pieces of the McNaughton function f is subdivided into a unimodular simplicial complex. Thanks to unimodularity, each Schauder hat can then be expressed in terms of “min” and “max” operations, which are, in turn, realizable by the operations \oplus , \odot , and \neg according to

$$\begin{aligned} \min\{x, y\} &= \neg(\neg x \odot y) \odot y := x \wedge y \\ \max\{x, y\} &= \neg(\neg x \oplus y) \oplus y := x \vee y. \end{aligned} \tag{11}$$

The overall MV term corresponding to f is finally obtained by combining the MV terms associated with the individual Schauder hats through \oplus operations. We refer the reader to [7] for a detailed account of the algorithm.

The hyperplane method [8] expresses the McNaughton function f , with linear pieces p_1, \dots, p_ℓ , in terms of the truncated linear polynomials $\sigma(p_1), \dots, \sigma(p_\ell)$, where $\sigma(x) = \min\{1, \max\{0, x\}\}$, for $x \in \mathbb{R}$, according to

$$f = \min_I \max_J \sigma(p_i), \quad (12)$$

where $I, J \subset \{1, \dots, \ell\}$ are index sets. Next, the MV terms corresponding to $\sigma(p_1), \dots, \sigma(p_\ell)$ are determined by repeated application of Lemma 3.3. These MV terms are finally combined into the MV term corresponding to f by using (12) and (11).

The only result on MV term extraction in the multi-dimensional case we are aware of is [7, Theorem 3], which establishes the existence of a Turing machine delivering the MV term underlying a given McNaughton function.

V. DEEP NETWORKS YIELD SHORTER FORMULAE

In this section, we perform an in-depth comparison between our algorithm for the univariate case, the Schauder hat method, and the hyperplane method. Specifically, we shall be interested in the length of the logical formulae produced by the different algorithms. We will establish, in the case of our method, that deep networks tend to produce shorter formulae. For illustration purposes, we let our discussion be guided by the following example. Consider the hat function $g : [0, 1] \rightarrow [0, 1]$ in Fig. 1,

$$\begin{aligned} g(x) &= \rho(2x) - 2\rho(2x - 1) \\ &= \begin{cases} 2x, & 0 \leq x \leq \frac{1}{2} \\ 2 - 2x, & \frac{1}{2} < x \leq 1. \end{cases} \end{aligned} \quad (13)$$

Let $g_1(x) = g(x)$, and define the s -th order sawtooth function g_s as the s -fold composition of g with itself, i.e.,

$$g_s := \underbrace{g \circ \dots \circ g}_s, \quad s \geq 2. \quad (14)$$

We note that g_s consists of the 2^s linear regions

$$\left[0, \frac{1}{2^s}\right], \left[\frac{1}{2^s}, \frac{2}{2^s}\right], \dots, \left[\frac{2^s - 1}{2^s}, 1\right],$$

with slope on each of these regions given by either 2^s or -2^s . As all the coefficients of g_s are integer, it constitutes a valid McNaughton function.




Fig. 1: Left: The function g , right: The function g_2 .

A. Analytical characterization of MV term lengths

We start by analytically characterizing the length of the MV terms corresponding to g_s produced by the different algorithms.

The Schauder hat method. Permutations of the linear pieces of g_s subdivide the interval $[0, 1]$ into 2^s simplices, namely $[0, \frac{1}{2^s}]$, $[\frac{1}{2^s}, \frac{2}{2^s}]$, \dots , $[\frac{2^s-1}{2^s}, 1]$. For $s \geq 2$, these simplices are further subdivided to ensure the unimodularity condition. The number of Schauder hats is³ $\Omega(2^s)$. As the MV term associated with each Schauder hat is of length at least 2, the overall MV term corresponding to g_s has length at least $\Omega(2^{s+1})$. As illustrated by the numerical results in Fig. 2, the actual length grows (as a function of s) significantly faster than this loose lower bound suggests.

The hyperplane method. The MV term associated with g_s obtained from the hyperplane method is of the form

$$\wedge_I \vee_J \sigma(p_i), \quad (15)$$

where I and J are index sets. The cardinality of I is equal to the number of linear pieces of g_s , which is 2^s and the cardinality of J is at least one (see [8, Sec. 1.4] for the details). For each $\sigma(p_i)$, we can choose to apply either Lemma 3.3 [7], [18], which delivers an MV term of length $\Theta(2^{2^s})$, or the optimized procedure [19] which generates an MV term of length $\Theta((2^s)^2)$. The former choice results in overall length of at least $\Omega(2^{s+2^s})$, while the latter one yields at least $\Omega(8^s)$. In the simulation results below, we consistently apply the procedure in [19].

³We write $f = \Omega(g)$ to signify that $\limsup_{t \rightarrow \infty} \left| \frac{f(t)}{g(t)} \right| > 0$.

Our algorithm with deep networks. We first note that g can be realized by a ReLU network Φ_g according to $\Phi_g = \rho \circ W_2 \circ \rho \circ W_1 = g$ with

$$W_1(x) = \begin{pmatrix} 2 \\ 2 \end{pmatrix} x - \begin{pmatrix} 0 \\ 1 \end{pmatrix}, \quad W_2(x) = \begin{pmatrix} 1 & -2 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix}.$$

Next, we run our algorithm described in Steps 1-3 in the proof of Theorem 3.2 on Φ_g . First, Φ_g is converted into the equivalent CReLU network $\Psi_g = \sigma \circ W_2^* \circ \sigma \circ W_1^* = \Phi_g$, with

$$W_1^*(x) = \begin{pmatrix} 2 \\ 2 \end{pmatrix} x - \begin{pmatrix} 0 \\ 1 \end{pmatrix}, \quad W_2^*(x) = \begin{pmatrix} 1 & -1 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix}.$$

Next, we apply Lemma 3.3 to the individual σ -neurons occurring in Ψ_g to get

$$\sigma(2x) = x \oplus x \tag{16}$$

$$\sigma(2x - 1) = x \odot x \tag{17}$$

$$\sigma(x_1 - x_2) = x_1 \odot \neg x_2. \tag{18}$$

Finally, the overall MV term corresponding to the function g is obtained by replacing x_1 in (18) by (16) and x_2 by (17) resulting in

$$\Psi_g(x) = (x \oplus x) \odot \neg(x \odot x). \tag{19}$$

We next establish that the self-compositions g_s can be realized by ReLU networks and can hence be converted into equivalent CReLU networks. To this end, we first note that the ReLU network implementing g_2 is given by

$$\Phi_g^2 := \Phi_g \circ \Phi_g = \rho \circ W_2 \circ \rho \circ W_1 \circ \rho \circ W_2 \circ \rho \circ W_1,$$

which has the equivalent CReLU network $\Psi_g^2 := \Psi_g \circ \Psi_g$. Likewise, the s -th order sawtooth function g_s , for $s \geq 3$, can be realized by CReLU networks according to

$$\Psi_g^s := \underbrace{\Psi_g \circ \dots \circ \Psi_g}_s, \quad s \geq 3.$$

We can now conclude that the MV term corresponding to Ψ_g^s , for $s \geq 2$, is obtained from the right-hand-side of (19) by iteratively applying the following procedure $s - 1$ times: Replace every occurrence of x by the right-hand-side of (19). The length of the resulting MV term is given by 4^s . Note that as opposed to the order-wise lower bounds above, here we compute the length of the extracted MV term precisely. We conclude that our algorithm, working off deep

networks realizing g_s , provably produces MV terms of length comparable to or shorter than the other two methods.

Our algorithm with shallow networks. Our algorithm offers flexibility the other two algorithms do not have, namely it can start from different ReLU network realizations of a given McNaughton function. We now investigate the impact of this flexibility on the result produced by the algorithm. Specifically, we shall be interested in understanding how the network architecture, notably deep vs. shallow, influences the length of the MV term delivered. This will be done by first converting the ReLU network Φ_g^s realizing g_s into an equivalent shallow network ϕ_g^s , given by

$$\phi_g^s = \rho \circ w_2 \circ \rho \circ w_1,$$

with

$$w_1(x) = \begin{pmatrix} 2^s \\ 2^s \\ 2^s \\ \vdots \\ 2^s \end{pmatrix} x + \begin{pmatrix} 0 \\ -1 \\ -2 \\ \vdots \\ -2^s + 1 \end{pmatrix}, \quad x \in \mathbb{R},$$

$$w_2(x) = \begin{pmatrix} 1 & -2 & 2 & -2 & 2 & \cdots & -2 \end{pmatrix} x, \quad x \in \mathbb{R}^{2^s},$$

and then applying the algorithm. The corresponding CReLU network is given by

$$\psi_g^s = \sigma \circ w_2^* \circ \sigma \circ w_1^*, \quad (20)$$

where

$$w_1^*(x) = \begin{pmatrix} 2^s \\ 2^s \\ 2^s \\ \vdots \\ 2^s \end{pmatrix} x + \begin{pmatrix} 0 \\ -1 \\ -2 \\ \vdots \\ -(2^s - 1) \end{pmatrix}, \quad x \in \mathbb{R},$$

$$w_2(x) = \begin{pmatrix} 1 & -1 & 1 & -1 & \cdots & 1 & -1 \end{pmatrix} x, \quad x \in \mathbb{R}^{2^s}.$$

Application of Lemma 3.4 now allows us to conclude that each σ -neuron in the first layer of ψ_g^s and the single σ -neuron in the second layer have associated MV terms of length $\Theta(2^{2^s})$. Composing MV terms according to (20) hence shows that the MV term corresponding to ψ_g^s has length $\Theta(2^{2^{s+1}})$. This allows us to conclude that starting from a shallow network realization

of g_s leads to an MV term length that grows double-exponentially in s , whereas working off a deep network realization as above yields length that is exponential in s , specifically $4^s = 2^{2s}$. This stark difference is also reflected in the simulation results below. We finally note that ReLU networks realizing multi-dimensional (McNaughton) functions $f : [0, 1]^n \rightarrow [0, 1]$, with $n \geq 2$, do not always have an equivalent shallow counterpart [20].

B. Expressive power of deep networks

The advantage in expressive power of deep networks over shallow ones has been analyzed in the neural network literature from various approximation-theoretic perspectives, see e.g. [17], [21]–[26]. The exponential vs. double-exponential dichotomy just identified establishes the representation benefits of deep networks from a mathematical logic perspective. In particular, our algorithm, when applied to the deep network representation Φ_g^s of the McNaughton function g_s , honors the compositional structure of $g_s = g \circ \dots \circ g$ thereby rendering the algebraic expression of the resulting MV term compositional as well. Functional compositionality is hence turned into algebraic compositionality. In contrast, application of our algorithm to the shallow network representation ϕ_g^s of the McNaughton function g_s ignores the compositional structure of g_s , which, in turn, leads to significantly longer MV terms. This observation can be substantiated through simulation results. Indeed, Fig. 2 displays double-exponential length growth behavior of our algorithm operating on the shallow network representation ϕ_g^s of g_s . For the other three algorithms we see exponential length growth, as predicted by the analysis, with our algorithm working off a deep network resulting in significantly shorter formulae than the Schauder hat and hyperplane methods.

The sawtooth function g_s along with its associated MV term, obtained by iteratively self-composing (19), constitutes an extreme case of compositional structure. We now investigate whether our algorithm retains an advantage over the Schauder hat and the hyperplane methods when the MV term underlying the McNaughton function does not exhibit compositional structure. To this end, we apply the three algorithms to McNaughton functions obtained from randomly generated MV terms of varying lengths without enforcing compositional structure⁴. Specifically, for a randomly selected MV term $\tau(x)$ of length L , the corresponding (piecewise linear) McNaughton function $\tau^{\mathcal{I}} : [0, 1] \rightarrow [0, 1]$ is computed by exploiting the fact that its breakpoints in

⁴The specific details of the approach used to randomly generate MV terms along with code for reproducing all simulation results can be found at https://www.mins.ee.ethz.ch/research/downloads/nm2mv_experiment.html




Fig. 2: Length of the MV terms extracted from g_s using different methods.

the interval $(0, 1)$ are rational numbers with denominators (in irreducible form) no larger than L [27], [28]. It hence suffices to evaluate $\tau^{\mathcal{I}}(x)$ at the points $\{a/b : a, b \in \mathbb{Z}, 1 \leq b \leq L\} \cap [0, 1]$, which is done by interpreting the connectives $\{\oplus, \odot, \neg\}$ in $\tau(x)$ according to Definition 2.1. Each of the resulting samples of $\tau^{\mathcal{I}}(x)$ is then compared with its left and its right neighbor to determine whether it corresponds to a breakpoint. Finally, we perform linear interpolation between the breakpoints. While the Schauder hat and the hyperplane methods are applied to the so-obtained McNaughton function $\tau^{\mathcal{I}}(x)$, our algorithm works off either the (deep) neural network obtained from $\tau(x)$ by following the constructive proof of the forward statement of Theorem 3.2 or the shallow (single-hidden-layer) network which identifies the breakpoints of $\tau^{\mathcal{I}}(x)$ with individual neurons.

For all methods, we verify that the MV terms delivered are, indeed, functionally equivalent to $\tau(x)$. This is done by comparing the breakpoints of the associated McNaughton functions, obtained again through the sampling procedure described above. Fig. 3 shows the average lengths of the extracted MV terms for L ranging from 4 to 14 and 500 Monte Carlo runs for each value of L . We observe that our algorithm applied to deep network representations of $\tau(x)$ consistently achieves the shortest length among all four methods. In fact, it delivers MV terms whose lengths are almost identical to those of $\tau(x)$. This indicates that a deep network representation of the MV term $\tau(x)$ encodes algebraic information in a way not present in the shallow network representation. Our algorithm when applied to trained neural networks can hence honor algebraic




Fig. 3: Lengths of MV terms extracted by different methods, the figure on the right depicts the same results as that on the left, but without the Shallow NN method.

structure present in the training data and expressed through the architecture of the network. The Schauder hat and the hyperplane methods do not have access to this algebraic information as they operate on the functional representation of $\tau^{\mathcal{I}}(x)$. We finally note that the number of breakpoints of $\tau^{\mathcal{I}}(x)$ does not seem to consistently grow with the length of the underlying MV term $\tau(x)$. Fig. 4 shows numerical results substantiating this claim, again for randomly generated MV terms and with 500 Monte Carlo runs for each data point. This observation also explains why the extracted MV term lengths for the Schauder hat, hyperplane, and shallow NN methods in Fig. 3 exhibit fluctuations as a function of L . Concretely, we observed in our simulation studies that the lengths of the MV terms returned by these three methods are governed by the number of breakpoints of the underlying McNaughton function.




Fig. 4: Average number of breakpoints of randomly generated MV terms.

We finally investigate whether our algorithm retains an advantage over the Schauder hat and the

hyperplane methods when $\tau(x)$ exhibits general (as opposed to the extreme case corresponding to $g_s(x)$ analyzed above) compositional structure. To see that this is, indeed, the case we randomly generate s , for $s = 1, \dots, 5$, MV terms $\tau_1(x), \dots, \tau_s(x)$, each of length 2 or 3, and then compute $\tau = \tau_1 \circ \dots \circ \tau_s$. Again, all three methods deliver MV terms that are verified to be functionally equivalent to τ , but have different lengths. Fig. 5 shows the results with 500 Monte Carlo runs for each value of s . We can see that the MV term length increases exponentially in s in all three cases, but again with the neural network algorithm resulting in significantly smaller lengths than the other two methods.




Fig. 5: Lengths of MV terms extracted by different methods.

C. The multi-dimensional case

As mentioned in Section IV, the Schauder hat and the hyperplane methods apply to the one-dimensional case only, whereas the neural network algorithm can handle the general multi-dimensional case. To illustrate this, Fig. 6 shows results obtained by applying our algorithm to neural networks obtained from randomly generated MV terms $\tau(x_1, x_2, x_3)$ of different lengths⁵, with 500 Monte Carlo runs for each data point. We observe that the algorithm returns MV terms of lengths almost identical to those of $\tau(x_1, x_2, x_3)$.

⁵The length of an MV term in d variables is defined as the total number of occurrences of propositional variables x_1, \dots, x_d .




Fig. 6: Length of MV terms extracted from three-dimensional McNaughton functions.

VI. EXTENSIONS TO THE RATIONAL AND REAL CASES

We now describe the extensions of our results to Rational Łukasiewicz logic and $\mathbb{R}\mathcal{L}$ logic, which, as mentioned in the introduction, have truth functions that are continuous piecewise linear, but with rational and, respectively, real coefficients.

A. The Rational Case

Rational Łukasiewicz logic extends MV logic by adding a division (by integers) operation. The algebraic counterpart is given by the so-called divisible many-valued (DMV) algebras [9].

Definition 6.1: Consider the MV algebra \mathcal{I} in Definition 2.1. Define the unary operations $\delta_i x = \frac{1}{i}x, x \in [0, 1]$, for all $i \in \mathbb{N}$. The structure $\mathcal{I}_d = \langle [0, 1], \oplus, \neg, \{\delta_i\}_{i \in \mathbb{N}}, 0 \rangle$ is a DMV algebra [9].

The class of term functions in \mathcal{I}_d is given by the continuous piecewise linear functions à la Theorem 3.1 but with rational coefficients [9], [29], hereafter referred to as rational McNaughton functions. We next extend Theorem 3.2 to the rational case.

Theorem 6.2: For $n \in \mathbb{N}$, let $\tau(x_1, \dots, x_n)$ be a DMV term in n variables and $\tau^{\mathcal{I}_d} : [0, 1]^n \rightarrow [0, 1]$ the associated term function in \mathcal{I}_d . There exists a ReLU network Φ with rational weights, satisfying

$$\Phi(x_1, \dots, x_n) = \tau^{\mathcal{I}_d}(x_1, \dots, x_n),$$

for all $(x_1, \dots, x_n) \in [0, 1]^n$. Conversely, for every ReLU network $\Phi : [0, 1]^n \rightarrow [0, 1]$ with rational weights, there exists a DMV term $\tau(x_1, \dots, x_n)$ whose associated term function in \mathcal{L}_d satisfies

$$\tau^{\mathcal{L}_d}(x_1, \dots, x_n) = \Phi(x_1, \dots, x_n),$$

for all $(x_1, \dots, x_n) \in [0, 1]^n$.

Proof: We already know from the proof of Theorem 3.2 how to realize the operations \oplus, \odot , and \neg with ReLU networks. The division operation $\delta_i : x \rightarrow \frac{1}{i}x$, for $i \in \mathbb{N}$, by virtue of being an affine transformation, is trivially realized by a single-layer ReLU network. ReLU network realizations of formulae in Rational Łukasiewicz logic are obtained by concatenating ReLU networks implementing the operations \oplus, \odot, \neg , and $\{\delta_i\}_{i \in \mathbb{N}}$. Inspection of the proof of the ReLU network composition Lemma B.2 reveals that the resulting ReLU network has rational weights.

Next, we extend the algorithm described in Section III to extract DMV terms from ReLU networks with rational weights. Steps 1 and 3 remain unaltered. We only remark that the additional weights introduced by the replacement procedure (3) in Step 1 are all integer-valued and hence the resulting σ -network is guaranteed to have rational weights. Step 2 needs to be modified as the σ -neurons are now of the form

$$h = \sigma(m_1x_1 + \dots + m_nx_n + b), \quad (21)$$

with $m_1, \dots, m_n, b \in \mathbb{Q}$, rendering Lemma 3.3, which requires $m_1, \dots, m_n, b \in \mathbb{Z}$, inapplicable. We employ an idea from [29] to transform a given σ -neuron with rational coefficients m_1, \dots, m_n, b into multiple σ -neurons with integer coefficients. Concretely, let $s \in \mathbb{N}$ be the least common multiple of the denominators of m_1, \dots, m_n, b . Recognizing that

$$s\sigma(x) = \sigma(sx) + \sigma(sx - 1) + \dots + \sigma(sx - (s - 1)), \quad (22)$$

for $x \in \mathbb{R}$, and setting $h_i = \sigma(s(m_1x_1 + \dots + m_nx_n + b) - i)$, it follows that $h = \sum_{i=0}^{s-1} \frac{1}{s}h_i$. As $h = \sum_{i=0}^{s-1} \frac{1}{s}h_i \leq 1$, the DMV term corresponding to h is given by $\oplus_{i=0}^{s-1} \delta_s \tau_i$, where τ_i denotes the DMV term associated with h_i . ■

The Schauder hat and the hyperplane methods were extended to Rational Łukasiewicz logic in [9], but again only for the univariate case.

B. The Real Case

We finally turn to the case of ReLU networks with real coefficients, which is of particular practical interest as it allows to extract logical formulae from trained ReLU networks. The Riesz many-valued algebra (RMV) [10] extends the MV algebra in Definition 2.1 by adding a multiplication operation.

Definition 6.3: Consider the MV algebra \mathcal{I} in Definition 2.1. Define the unary operations $\Delta_r x = rx, x \in [0, 1]$, for all $r \in [0, 1]$. The structure $\mathcal{I}_r = \langle [0, 1], \oplus, \neg, \{\Delta_r\}_{r \in [0, 1]}, 0 \rangle$ is an RMV algebra [10].

The term functions of the corresponding logic $\mathbb{R}\mathcal{L}$ are continuous piecewise linear functions with real coefficients [10]. Noting that the multiplication operation $\Delta_r : x \rightarrow rx$ is an affine transformation, it follows that ReLU networks realizing RMV term functions are obtained in the same manner as in the integer and rational cases, namely by composing the ReLU networks realizing the individual operations appearing in the RMV term under consideration. Conversely, for every ReLU network $\Phi : [0, 1]^n \rightarrow [0, 1]$ with real weights, there exists an RMV term $\tau(x_1, \dots, x_n)$ whose associated term function in \mathcal{I}_r satisfies

$$\tau^{\mathcal{I}_r}(x_1, \dots, x_n) = \Phi(x_1, \dots, x_n),$$

for all $(x_1, \dots, x_n) \in [0, 1]^n$.

We now generalize our algorithm to extract RMV terms from ReLU networks with real weights. Again, Steps 1 and 3 in Section III remain unaltered. In Step 2, instead of Lemma 3.3, we apply the following result.

Lemma 6.4 ([10]): Consider the function $f(x_1, \dots, x_n) = m_1 x_1 + \dots + m_n x_n + b, (x_1, \dots, x_n) \in [0, 1]^n$, with $m_1, \dots, m_n, b \in \mathbb{R}$. For all $m \in (0, 1]$ and $i \in \{1, \dots, n\}$, with $f_o(x_1, \dots, x_n) = m_1 x_1 + \dots + m_{i-1} x_{i-1} + (m_i - m)x_i + m_{i+1} x_{i+1} + \dots + m_n x_n + b$, it holds that

$$\sigma(f) = (\sigma(f_o) \oplus (m x_i)) \odot \sigma(f_o + 1). \quad (23)$$

For completeness, the proof of Lemma 6.4 is provided in Appendix D. We demonstrate the application of Lemma 6.4 through a simple example. Consider the σ -neuron $\sigma\left(\frac{1}{\sqrt{2}}x_1 - 2x_2\right)$. First, apply Lemma 6.4 with $m = \frac{1}{\sqrt{2}}$ and $i = 1$ to get

$$\sigma\left(\frac{1}{\sqrt{2}}x_1 - 2x_2\right) = \left(\sigma(-2x_2) \oplus \left(\frac{1}{\sqrt{2}}x_1\right)\right) \odot \sigma(-2x_2 + 1). \quad (24)$$

As $\sigma(-2x_2) = 0$ and $0 \oplus x = x$, (24) reduces to $(\frac{1}{\sqrt{2}}x_1) \odot \sigma(-2x_2 + 1)$. Next, using $\sigma(x) = 1 - \sigma(1 - x)$, $x \in \mathbb{R}$, we obtain $\sigma(-2x_2 + 1) = 1 - \sigma(2x_2) = \neg(x_2 \oplus x_2)$. The overall RMV term associated with $\sigma\left(\frac{1}{\sqrt{2}}x_1 - 2x_2\right)$ is therefore given by $\Delta_{\frac{1}{\sqrt{2}}}x_1 \odot \neg(x_2 \oplus x_2)$.

The hyperplane method was extended to the real case in [10], again only for the univariate case. An extension of the Schauder hat method to the real case does not seem to be available in the literature, but can readily be devised.

ACKNOWLEDGMENT

The authors are grateful to Prof. Olivia Caramello for drawing their attention to the McNaughton theorem and, more generally, to MV logic.

APPENDIX A

MV ALGEBRA

Definition A.1 ([5]): A many-valued algebra is a structure $\mathcal{A} = \langle A, \oplus, \neg, 0 \rangle$ consisting of a nonempty set A , a constant $0 \in A$, a binary operation \oplus , and a unary operation \neg satisfying the following axioms:

$$x \oplus (y \oplus z) = (x \oplus y) \oplus z \quad (25.1)$$

$$x \oplus y = y \oplus x \quad (25.2)$$

$$x \oplus 0 = x \quad (25.3)$$

$$\neg\neg x = x \quad (25.4)$$

$$x \oplus \neg 0 = \neg 0 \quad (25.5)$$

$$\neg(\neg x \oplus y) \oplus y = \neg(\neg y \oplus x) \oplus x. \quad (25.6)$$

An MV algebra $\langle A, \oplus, \neg, 0 \rangle$ is said to be nontrivial iff A contains more than one element. In each MV algebra we can define a constant 1 and a binary operation \odot as follows:

$$1 := \neg 0 \quad (26)$$

$$x \odot y := \neg(\neg x \oplus \neg y). \quad (27)$$

The ensuing identities are then direct consequences of Definition A.1:

$$x \odot (y \odot z) = (x \odot y) \odot z \quad (28.1)$$

$$x \odot y = y \odot x \quad (28.2)$$

$$x \odot 1 = x \quad (28.3)$$

$$x \odot 0 = 0. \quad (28.4)$$

We will frequently use the notions of MV terms and term functions formalized as follows.

Definition A.2 ([5]): Let $n \in \mathbb{N}$ and $S_n = \{(\ , \), 0, \neg, \oplus, x_1, \dots, x_n\}$. An MV term in the variables x_1, \dots, x_n is a finite string over S_n arising from a finite number of applications of the operations \neg and \oplus as follows. The elements 0 and x_i , for $i = 1, \dots, n$, considered as one-element strings, are MV terms.

- 1) If the string τ is an MV term, then $\neg\tau$ is also an MV term.
- 2) If the strings τ and γ are MV terms, then $(\tau \oplus \gamma)$ is also an MV term.

We write $\tau(x_1, \dots, x_n)$ to emphasize that τ is an MV term in the variables x_1, \dots, x_n .

For instance, the following strings over $S_2 = \{(\ , \), 0, \neg, \oplus, x_1, x_2\}$ are MV terms:

$$0, x_1, x_2, \neg 0, \neg x_2, (x_1 \oplus \neg x_2).$$

We shall always omit the outermost pair of brackets for conciseness, e.g., we write $x_1 \oplus \neg x_2$ instead of $(x_1 \oplus \neg x_2)$. Besides, for brevity we use the symbols 1 and \odot according to (26) and (27), respectively.

MV terms are logical formulae of purely syntactic nature. To endow them with semantics, an MV algebra must be specified. The associated truth functions, a.k.a. term functions which we define presently, are then obtained by interpreting the operations \oplus and \neg according to how they are specified in the MV algebra.

Definition A.3 ([5]): Let $\tau(x_1, \dots, x_n)$ be an MV term in the variables x_1, \dots, x_n . For the MV algebra $\mathcal{A} = \langle A, \oplus, \neg, 0 \rangle$, the term function $\tau^{\mathcal{A}}(a_1, \dots, a_n) : A^n \rightarrow A$ associated with τ in \mathcal{A} is obtained by substituting a_i for all occurrences of the variable x_i and interpreting the symbol 0 and the operations \oplus, \neg as the corresponding symbol 0 and operations \oplus, \neg in \mathcal{A} .

APPENDIX B
RELU NETWORKS

Definition B.1 ([17]): Let $L \in \mathbb{N}$ and $N_0, N_1, \dots, N_L \in \mathbb{N}$. A ReLU network is a map $\Phi : \mathbb{R}^{N_0} \rightarrow \mathbb{R}^{N_L}$ given by

$$\Phi = \begin{cases} W_1, & L = 1 \\ W_2 \circ \rho \circ W_1, & L = 2, \\ W_L \circ \rho \circ W_{L-1} \circ \rho \circ \dots \circ \rho \circ W_1, & L \geq 3 \end{cases}$$

where, for $\ell \in \{1, 2, \dots, L\}$, $W_\ell : \mathbb{R}^{N_{\ell-1}} \rightarrow \mathbb{R}^{N_\ell}$, $W_\ell(x) := A_\ell x + b_\ell$ are affine transformations with weight matrices $A_\ell = \mathbb{R}^{N_\ell \times N_{\ell-1}}$ and bias vectors $b_\ell \in \mathbb{R}^{N_\ell}$, and the ReLU activation function $\rho : \mathbb{R} \rightarrow \mathbb{R}$, $\rho(x) := \max\{0, x\}$ acts component-wise. The number of layers of the network Φ , denoted by $\mathcal{L}(\Phi)$, is given by L . We denote by $\mathcal{N}_{d,d'}$ the set of ReLU networks with input dimension $N_0 = d$ and output dimension $N_L = d'$.

The next result formalizes properties of ReLU network compositions.

Lemma B.2 ([17]): Let $d_1, d_2, d_3 \in \mathbb{N}$, $\Phi_1 \in \mathcal{N}_{d_1, d_2}$, and $\Phi_2 \in \mathcal{N}_{d_2, d_3}$. There exists a network $\Psi \in \mathcal{N}_{d_1, d_3}$ with $\mathcal{L}(\Psi) = \mathcal{L}(\Phi_1) + \mathcal{L}(\Phi_2)$, and satisfying

$$\Psi(x) = (\Phi_2 \circ \Phi_1)(x), \quad \text{for all } x \in \mathbb{R}^{d_1}.$$

Proof: The proof is based on the identity $x = \rho(x) - \rho(-x)$. First, note that by Definition B.1, setting $L_1 = \mathcal{L}(\Phi_1)$ and $L_2 = \mathcal{L}(\Phi_2)$, we can write

$$\Phi_1 = W_{L_1}^1 \circ \rho \circ W_{L_1-1}^1 \circ \dots \circ \rho \circ W_1^1$$

and

$$\Phi_2 = W_{L_2}^2 \circ \rho \circ W_{L_2-1}^2 \circ \dots \circ \rho \circ W_1^2,$$

with the appropriate modifications when either L_1 or L_2 or both are equal to 1 or 2. Next, let $N_{L_1-1}^1$ denote the input dimension of $W_{L_1}^1$ and N_1^2 the output dimension of W_1^2 . We define the affine transformations $\widetilde{W}_{L_1}^1 : \mathbb{R}^{N_{L_1-1}^1} \rightarrow \mathbb{R}^{2d_2}$ and $\widetilde{W}_1^2 : \mathbb{R}^{2d_2} \rightarrow \mathbb{R}^{N_1^2}$ according to

$$\begin{aligned} \widetilde{W}_{L_1}^1(x) &:= \begin{pmatrix} \mathbb{I}_{d_2} \\ -\mathbb{I}_{d_2} \end{pmatrix} W_{L_1}^1(x) \\ \widetilde{W}_1^2(x) &:= W_1^2 \left(\begin{pmatrix} \mathbb{I}_{d_2} & -\mathbb{I}_{d_2} \end{pmatrix} x \right). \end{aligned}$$

The proof is completed upon noting that the network

$$\Psi := W_{L_2}^2 \circ \cdots \circ W_2^2 \circ \rho \circ \widetilde{W}_1^2 \circ \rho \circ \widetilde{W}_{L_1}^1 \circ \rho \circ W_{L_1-1}^1 \circ \cdots \circ W_1^1$$

satisfies the claimed properties. ■

Lemma B.3: There exist ReLU networks $\Phi^\oplus \in \mathcal{N}_{2,1}$ and $\Phi^\ominus \in \mathcal{N}_{2,1}$ satisfying

$$\Phi^\oplus(x, y) = \min\{1, x + y\}$$

$$\Phi^\ominus(x, y) = \max\{0, x + y - 1\},$$

for all $x, y \in [0, 1]$.

Proof: First, to realize the operation $x \oplus y = \min\{1, x + y\}$, we note that addition can be implemented by a single-layer ReLU network according to

$$x + y = \begin{pmatrix} 1 & 1 \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix}.$$

It remains to implement the “min” operation by a ReLU network. To this end, we observe that

$$\min\{1, x\} = 1 - \rho(1 - x) = (W_2 \circ \rho \circ W_1)(x),$$

for $x \in [0, 1]$, where

$$W_1(x) = -x + 1, \quad W_2(x) = -x + 1.$$

Now, applying Lemma B.2 to concatenate the networks $\Phi_1(x, y) = \begin{pmatrix} 1 & 1 \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix}$ and $\Phi_2(x) = (W_2 \circ \rho \circ W_1)(x)$ yields the desired ReLU network realization of $x \oplus y$ according to

$$x \oplus y = (W_2^\oplus \circ \rho \circ W_1^\oplus)(x, y),$$

for $x, y \in [0, 1]$, where

$$W_1^\oplus(x, y) = \begin{pmatrix} -1 & -1 \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} + 1, \quad W_2^\oplus(x) = -x + 1.$$

For the operation $x \ominus y = \max\{0, x + y - 1\}$, we simply note that

$$\max\{0, x + y - 1\} = \rho \left(\begin{pmatrix} 1 & 1 \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} - 1 \right) = (W_2^\ominus \circ \rho \circ W_1^\ominus)(x, y),$$

for $x, y \in [0, 1]$, where

$$W_1^\ominus(x, y) = \begin{pmatrix} 1 & 1 \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} - 1, \quad W_2^\ominus(x) = x.$$

■

APPENDIX C
PROOF OF LEMMA 3.3

Proof: We follow the line of arguments in [7] and consider four different cases.

Case 1: $f_{\circ}(x) \geq 1$, for all $x \in [0, 1]^n$. In this case, the LHS of (5) is

$$\sigma(f) = 1$$

and the RHS evaluates to

$$(\sigma(f_{\circ}) \oplus x_1) \odot \sigma(f_{\circ} + 1) = (1 \oplus x_1) \odot 1 = 1.$$

Case 2: $f_{\circ}(x) \leq -1$, for all $x \in [0, 1]^n$. In this case, the LHS of (5) is

$$\sigma(f) = 0$$

and the RHS satisfies

$$(\sigma(f_{\circ}) \oplus x_1) \odot \sigma(f_{\circ} + 1) = (0 \oplus x_1) \odot 0 = 0.$$

Case 3: $-1 < f_{\circ}(x) \leq 0$, for all $x \in [0, 1]^n$. In this case, $f \in (-1, 1]$ as $x_1 \in [0, 1]$. The RHS of (5) becomes

$$\begin{aligned} & (\sigma(f_{\circ}) \oplus x_1) \odot \sigma(f_{\circ} + 1) \\ &= (0 \oplus x_1) \odot (f_{\circ} + 1) \\ &= x_1 \odot (f_{\circ} + 1) \\ &= \max\{0, x_1 + f_{\circ} + 1 - 1\} \\ &= \max\{0, f\} \\ &= \sigma(f). \end{aligned}$$

Case 4: $0 < f_{\circ}(x) < 1$, for all $x \in [0, 1]^n$. In this case, $f \in (0, 2)$. The RHS of (5) becomes

$$\begin{aligned} & (\sigma(f_{\circ}) \oplus x_1) \odot \sigma(f_{\circ} + 1) \\ &= (f_{\circ} \oplus x_1) \odot 1 \\ &= f_{\circ} \oplus x_1 \\ &= \min\{1, f_{\circ} + x_1\} \\ &= \min\{1, f\} \\ &= \sigma(f). \end{aligned}$$

■

APPENDIX D
PROOF OF LEMMA 6.4

We can follow the line of arguments in the proof of Lemma 3.3 and consequently consider four different cases.

Case 1: $f_{\circ}(x) \geq 1$, for all $x \in [0, 1]^n$. In this case, the LHS of (23) is

$$\sigma(f) = 1.$$

and the RHS evaluates to

$$(\sigma(f_{\circ}) \oplus (mx_i)) \odot \sigma(f_{\circ} + 1) = (1 \oplus (mx_i)) \odot 1 = 1.$$

Case 2: $f_{\circ}(x) \leq -1$, for all $x \in [0, 1]^n$. In this case, the LHS of (23) is

$$\sigma(f) = 0$$

and the RHS is given by

$$(\sigma(f_{\circ}) \oplus (mx_i)) \odot \sigma(f_{\circ} + 1) = (0 \oplus (mx_i)) \odot 0 = 0.$$

Case 3: $-1 < f_{\circ}(x) \leq 0$, for all $x \in [0, 1]^n$. In this case, $f \in (-1, 1]$ as $mx_i \in [0, 1]$, for all $i \in \{1, \dots, n\}$. The RHS of (23) becomes

$$\begin{aligned} & (\sigma(f_{\circ}) \oplus (mx_i)) \odot \sigma(f_{\circ} + 1) \\ &= (0 \oplus (mx_i)) \odot (f_{\circ} + 1) \\ &= (mx_i) \odot (f_{\circ} + 1) \\ &= \max\{0, mx_i + f_{\circ} + 1 - 1\} \\ &= \max\{0, f\} \\ &= \sigma(f). \end{aligned}$$

Case 4: $0 < f_{\circ}(x) < 1$, for all $x \in [0, 1]^n$. In this case, $f \in (0, 2)$. The RHS of (23) becomes

$$\begin{aligned}
& (\sigma(f_{\circ}) \oplus (mx_i)) \odot \sigma(f_{\circ} + 1) \\
&= (f_{\circ} \oplus (mx_i)) \odot 1 \\
&= f_{\circ} \oplus (mx_i) \\
&= \min\{1, f_{\circ} + mx_i\} \\
&= \min\{1, f\} \\
&= \sigma(f).
\end{aligned}$$

APPENDIX E

PROOF OF LEMMA 3.4

We prove the statement by induction on m . For the base case $m = 1$, $\sigma(f)$ is of the form $\sigma(x_i + b)$ or $\sigma(-x_i + b)$, for some $i \in \{1, \dots, n\}$ and $b \in \mathbb{Z}$. As $\sigma(x_i + b) \equiv 0$ for $b \leq -1$, $\sigma(x_i + b) \equiv 1$ for $b \geq 1$, $\sigma(-x_i + b) \equiv 0$ for $b \leq 0$, and $\sigma(-x_i + b) \equiv 1$ for $b \geq 2$, the only two cases where $\sigma(f)$ is not a constant are $\sigma(f) = \sigma(x_i)$, with corresponding MV term x_i , and $\sigma(f) = \sigma(1 - x_i)$, with corresponding MV term $\neg x_i$. Recalling that the length of an MV term is given by the total number of occurrences of propositional variables, this establishes the statement in the base case. For the induction step, we conclude from (5) that the length of the MV term corresponding to $\sigma(f)$ equals the sum of the lengths of the MV terms corresponding to $\sigma(f_{\circ})$ and $\sigma(f_{\circ} + 1)$ plus 1. By the induction hypothesis, $\sigma(f_{\circ})$ has an MV term of length at most $2^{m-1} - 1$. Further, the length of the MV term corresponding to $\sigma(f_{\circ} + 1)$ is upper-bounded by that of the MV term associated with $\sigma(f_{\circ})$. We can therefore conclude that the length of the MV term corresponding to $\sigma(f)$ is no larger than $2^m - 1$.

APPENDIX F

CONSTRUCTION OF THE NETWORK Φ^{τ} IN SECTION III

By Lemma B.3, we have the ReLU network associated with $x \oplus x$ according to

$$W_2^{\oplus} \circ \rho \circ \left(\begin{pmatrix} -1 & -1 \end{pmatrix} \begin{pmatrix} x \\ x \end{pmatrix} + 1 \right),$$

which can be rewritten as

$$W_2^{\oplus} \circ \rho \circ (-2x + 1). \tag{29}$$

As $y = \rho(y) - \rho(-y)$, for $y \in \mathbb{R}$, and $\neg y = 1 - y$, for $y \in [0, 1]$, we obtain the 2-layer ReLU network associated with $\neg y$ according to

$$-\rho(y) + \rho(-y) + 1. \quad (30)$$

Finally, the ReLU network realizing the term function associated with $\tau = (x \oplus x) \odot \neg y$ is obtained by composing the network Φ^\odot in Lemma B.3 with (29) and (30) according to

$$W_2^\odot \circ \rho \circ W_1^\odot \circ \left(W_2^\oplus \circ \rho \circ (-2x + 1) \right),$$

which, through simple algebraic manipulations, is found to be equivalent to the network Φ^τ in (2).

REFERENCES

- [1] G. Lample and F. Charton, “Deep learning for symbolic mathematics,” in *International Conference on Learning Representations*, 2020.
- [2] S. Bubeck, V. Chandrasekaran, R. Eldan *et al.*, “Sparks of artificial general intelligence: Early experiments with GPT-4,” *arXiv preprint arXiv:2303.12712*, 2023.
- [3] G. Weiss, Y. Goldberg, and E. Yahav, “Thinking like transformers,” in *International Conference on Machine Learning*, 2021, pp. 11 080–11 090.
- [4] R. O’Donnell, *Analysis of Boolean Functions*. Cambridge, UK: Cambridge University Press, 2014.
- [5] R. Cignoli, I. d’Ottaviano, and D. Mundici, *Algebraic Foundations of Many-Valued Reasoning*. Springer, 2000.
- [6] R. McNaughton, “A theorem about infinite-valued sentential logic,” *The Journal of Symbolic Logic*, vol. 16, no. 1, pp. 1–13, 1951.
- [7] D. Mundici, “A constructive proof of McNaughton’s theorem in infinite-valued logic,” *The Journal of Symbolic Logic*, vol. 59, no. 2, pp. 596–602, 1994.
- [8] S. Aguzzoli, “Geometrical and Proof Theoretical Issues in Łukasiewicz Propositional Logics,” PhD thesis, University of Siena, Italy, 1998.
- [9] B. Gerla, “Rational Łukasiewicz logic and DMV-algebras,” *Neural Network World*, vol. 6, no. 6, 2001.
- [10] A. Di Nola and I. Leuştean, “Łukasiewicz logic and Riesz spaces,” *Soft Computing*, vol. 18, no. 12, pp. 2349–2363, 2014.
- [11] P. Amato, A. Di Nola, and B. Gerla, “Neural networks and rational Łukasiewicz logic,” in *Annual Meeting of the North American Fuzzy Information Processing Society Proceedings*, 2002, pp. 506–510.
- [12] —, “Neural networks and rational McNaughton Functions,” *Journal of Multiple-Valued Logic and Soft Computing*, vol. 11, no. 1-2, pp. 95–110, 2005.
- [13] A. Di Nola, B. Gerla, and I. Leuştean, “Adding real coefficients to Łukasiewicz logic: An application to neural networks,” *Fuzzy Logic and Applications: 10th International Workshop*, pp. 77–85, 2013.
- [14] A. Tarski, *Logic, semantics, metamathematics: Papers from 1923 to 1938*. Hackett Publishing, 1983.
- [15] C. C. Chang, “Algebraic analysis of many-valued logics,” *Transactions of the American Mathematical Society*, vol. 88, no. 2, pp. 467–490, 1958.

- [16] —, “A new proof of the completeness of the Łukasiewicz axioms,” *Transactions of the American Mathematical Society*, vol. 93, no. 1, pp. 74–80, 1959.
- [17] D. Elbrächter, D. Perekrestenko, P. Grohs, and H. Bölskei, “Deep neural network approximation theory,” *IEEE Transactions on Information Theory*, vol. 67, no. 5, pp. 2581–2623, 2021.
- [18] A. Rose and J. B. Rosser, “Fragments of many-valued statement calculi,” *Transactions of the American Mathematical Society*, vol. 87, no. 1, pp. 1–53, 1958.
- [19] S. Aguzzoli, “The complexity of McNaughton functions of one variable,” *Advances in Applied Mathematics*, vol. 21, no. 1, pp. 58–77, 1998.
- [20] C. A. Haase, C. Hertrich, and G. Loho, “Lower bounds on the depth of integral ReLU neural networks via lattice polytopes,” in *International Conference on Learning Representations*, 2023.
- [21] J. Schmidt-Hieber, “Nonparametric regression using deep neural networks with ReLU activation function,” *The Annals of Statistics*, vol. 48, no. 4, pp. 1875–1897, 2020.
- [22] M. Bianchini and F. Scarselli, “On the complexity of neural network classifiers: A comparison between shallow and deep architectures,” *IEEE Transactions on Neural Networks and Learning Systems*, vol. 25, no. 8, pp. 1553–1565, 2014.
- [23] D. Yarotsky, “Error bounds for approximations with deep ReLU networks,” *Neural Networks*, vol. 94, pp. 103–114, 2017.
- [24] M. Raghu, B. Poole, J. Kleinberg, S. Ganguli, and J. Sohl-Dickstein, “On the expressive power of deep neural networks,” in *International Conference on Machine Learning*. PMLR, 2017, pp. 2847–2854.
- [25] H. W. Lin, M. Tegmark, and D. Rolnick, “Why does deep and cheap learning work so well?” *Journal of Statistical Physics*, vol. 168, pp. 1223–1247, 2017.
- [26] M. Telgarsky, “Representation benefits of deep feedforward networks,” *arXiv preprint arXiv:1509.08101*, 2015.
- [27] S. Aguzzoli, “An asymptotically tight bound on countermodels for Łukasiewicz logic,” *International Journal of Approximate Reasoning*, vol. 43, no. 1, pp. 76–89, 2006.
- [28] S. Aguzzoli and A. Ciabattoni, “Finiteness in infinite-valued Łukasiewicz logic,” *Journal of Logic, Language and Information*, vol. 9, no. 1, pp. 5–29, 2000.
- [29] M. Baaz and H. Veith, “Interpolation in fuzzy logic,” *Archive for Mathematical Logic*, vol. 38, no. 7, pp. 461–489, 1999.