



# Floating-Point Neural Networks are Provably Robust Universal Approximators

Geonho Hwang<sup>1</sup>, Wonyeol Lee<sup>2</sup>, Yeachan Park<sup>3</sup>, Sejun Park<sup>4</sup>, Sejun Park<sup>4</sup>,  
and Feras Saad<sup>5</sup>

<sup>1</sup> GIST, Gwangju, Republic of Korea  
[hgh2134@gist.ac.kr](mailto:ghg2134@gist.ac.kr)

<sup>2</sup> POSTECH, Pohang, Republic of Korea  
[wonyeol.lee@postech.ac.kr](mailto:wonyeol.lee@postech.ac.kr)

<sup>3</sup> Sejong University, Seoul, Republic of Korea  
[ychpark@sejong.ac.kr](mailto:ychpark@sejong.ac.kr)

<sup>4</sup> Korea University, Seoul, Republic of Korea  
[sejun.park000@gmail.com](mailto:sejun.park000@gmail.com)

<sup>5</sup> Carnegie Mellon University, Pittsburgh, PA, USA  
[fsaad@cmu.edu](mailto:fsaad@cmu.edu)

**Abstract.** The classical universal approximation (UA) theorem for neural networks establishes mild conditions under which a feedforward neural network can approximate a continuous function  $f$  with arbitrary accuracy. A recent result shows that neural networks also enjoy a more general *interval* universal approximation (IUA) theorem, in the sense that the abstract interpretation semantics of the network using the interval domain can approximate the direct image map of  $f$  (i.e., the result of applying  $f$  to a set of inputs) with arbitrary accuracy. These theorems, however, rest on the unrealistic assumption that the neural network computes over infinitely precise real numbers, whereas their software implementations in practice compute over finite-precision floating-point numbers. An open question is whether the IUA theorem still holds in the floating-point setting.

This paper introduces the first IUA theorem for *floating-point* neural networks that proves their remarkable ability to *perfectly capture* the direct image map of any rounded target function  $f$ , showing no limits exist on their expressiveness. Our IUA theorem in the floating-point setting exhibits material differences from the real-valued setting, which reflects the fundamental distinctions between these two computational models. This theorem also implies surprising corollaries, which include (i) the existence of *provably robust* floating-point neural networks; and (ii) the *computational completeness* of the class of straight-line programs that use only floating-point additions and multiplications for the class of all floating-point programs that halt.

**Keywords:** Neural networks · Robust machine learning · Floating point · Universal approximation · Abstract interpretation

---

G. Hwang and W. Lee—Equal contribution.

The full version of this article is at <https://doi.org/10.48550/arXiv.2506.16065>.

© The Author(s) 2025

R. Piskac and Z. Rakamarić (Eds.): CAV 2025, LNCS 15932, pp. 301–326, 2025.

[https://doi.org/10.1007/978-3-031-98679-6\\_14](https://doi.org/10.1007/978-3-031-98679-6_14)

## 1 Introduction

**Background.** Despite the remarkable success of neural networks on diverse tasks, these models often lack *robustness* and are subject to adversarial attacks. Slight perturbations to the network inputs can cause the network to produce significantly different outputs [23, 63], raising serious concerns in safety-critical domains such as healthcare [18], cybersecurity [57], and autonomous driving [16].

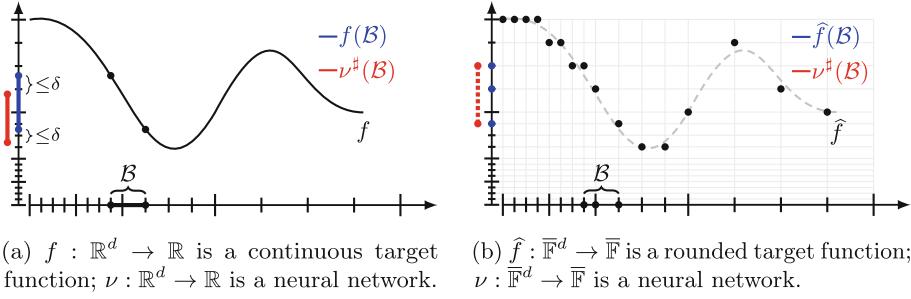
These issues have brought about significant advances in new algorithms for *robustness verification* [2, 34, 42], which prove the robustness of a given network; and *robust training* [24, 48, 56, 68], which train a network to be provably robust. But despite these advances, provably robust networks do not yet achieve state-of-the-art accuracy [38]. For example, on the CIFAR-10 image classification benchmark, non-robust networks achieve over 99% accuracy, whereas the best provably robust networks achieve less than 63% [37]. This performance gap has prompted researchers to explore whether there exists fundamental limits on the *expressiveness* of provably robust networks that restrict their accuracy [3].

Surprisingly, it has been proven that no such fundamental limit exists. Informally, for any continuous function  $f : \mathbb{R}^d \rightarrow \mathbb{R}$  and compact set  $\mathcal{K} \subset \mathbb{R}^d$ , there exists a neural network  $g : \mathbb{R}^d \rightarrow \mathbb{R}$  whose robustness properties are “sufficiently close” to those of  $f$  over  $\mathcal{K}$  and easily provable using abstract interpretation [10] over the interval domain. This result, known as the *interval universal approximation* (IUA) theorem [4, 66], generalizes the classical universal approximation (UA) theorem [12, 27] from pointwise-values to intervals, and confirms that provably robust networks do not suffer from a fundamental loss of expressive power.

**Key Challenges.** The IUA theorem in [4, 66] overlooks a critical aspect of real-world computation, which is the use of *floating-point arithmetic* instead of real arithmetic. It assumes that neural networks and interval analyses operate on arbitrary real numbers with exact operations. In reality, numerical implementations of neural networks use floating-point numbers and operations [22, §4.1], sometimes with extremely low-precision to speed-up performance [14, 29]. This discrepancy means that the existing IUA theorem does not directly apply to neural networks that are implemented in software and actually used in practice.

To our knowledge, no prior work has studied the robustness and expressiveness properties of floating-point neural networks or established an IUA theorem for them. The unique complexities of floating-point arithmetic introduce daunting challenges to any such theoretical study. For example, floating-point numbers are discretized and bounded, and their operations have rounding errors that become infinite in cases of overflow. Whereas the IUA proof over reals requires very large real numbers for network weights or intermediate computations, these values cannot be represented as floats. Naively rounding reals to floats causes approximation errors that invalidate many steps of the IUA proofs in [4, 66].

**This Work.** We formally study the IUA theorem over floating point, as a step toward bridging the theory and practice of provably robust neural networks.



**Fig. 1.** Illustration and comparison of the IUA theorems. (a) In the real-valued setting, the neural network abstract interpretation  $\nu^\#$  forms a  $\delta$ -approximation to the image map of  $f$ . (b) In the floating-point setting,  $\nu^\#$  exactly computes the upper and lower points of the image map of  $f$ :  $\nu^\#(\mathcal{B}) = [\min \hat{f}(\mathcal{B}), \max \hat{f}(\mathcal{B})] \cap \overline{\mathbb{F}}$ .

We first formulate a floating-point analog of the IUA theorem, considering the details of floating point. Let  $f : \mathbb{R}^d \rightarrow \mathbb{R}$  be a target function to approximate. Since all floating-point neural networks are functions between floating-point values, they can at-best approximate the rounded version  $\hat{f} : \overline{\mathbb{F}}^d \rightarrow \overline{\mathbb{F}}$  of  $f$  over floats, where  $\overline{\mathbb{F}}$  denotes the set of all floats. The floating-point version of the IUA theorem asks the following: is there a floating-point neural network  $\nu : \overline{\mathbb{F}}^d \rightarrow \overline{\mathbb{F}}$  whose *interval semantics* is arbitrarily close to the *direct image map* of the rounded target  $\hat{f}$  over  $[-1, 1]^d$ ? More formally, this property means that for any  $\delta > 0$ , there exists a neural network  $\nu$  such that for all boxes  $\mathcal{B} \subseteq [-1, 1]^d \cap \overline{\mathbb{F}}^d$ ,

$$\left| \min \nu^\#(\mathcal{B}) - \min \hat{f}(\mathcal{B}) \right| \leq \delta, \quad \left| \max \nu^\#(\mathcal{B}) - \max \hat{f}(\mathcal{B}) \right| \leq \delta. \quad (1)$$

In Eq. (1),  $\nu^\#(\mathcal{B})$  is the result of abstract interpretation of  $\mathcal{B}$  under  $\nu$  (using the interval domain), and  $\hat{f}(\mathcal{B}) := \left\{ \hat{f}(\mathbf{x}) \mid \mathbf{x} \in \mathcal{B} \right\} \subset \mathbb{R}$  is the image of  $\mathcal{B}$  under  $\hat{f}$ .

We prove that the IUA theorem holds for floating-point networks, despite all their numerical complexities. In particular, we show that for *any* target function  $f$  and a *large* class of activation functions  $\sigma$ , including most practical ones (e.g., ReLU, GELU, sigmoid), it is possible to find a floating-point network  $\nu$  with  $\sigma$  whose interval semantics *exactly* matches the direct image map of the rounded target  $\hat{f}$  over  $[-1, 1]^d \cap \overline{\mathbb{F}}^d$  (Fig. 1). This result implies that no fundamental limit exists on the expressiveness of provably robust floating-point neural networks.

Our result is considerably different from the previous IUA theorem over the reals in three key aspects. The previous theorem considers continuous target functions; requires a restricted class of so-called squashable activation functions; and finds networks that are arbitrarily close to target functions. In contrast, our result considers arbitrary target functions; allows almost all activation functions used in practice; and find networks that are precisely equal to (rounded) target functions. Our IUA theorem even holds for the *identity* activation function, which is not the case for the traditional IUA or UA theorems over real numbers, because any network that uses the identity activation is affine over the reals.

As a corollary of our main theorem, we prove the following existence of provably robust floating-point neural networks: given an ideal floating-point classifier  $\hat{f}$  (not necessarily a neural network) that is robust (not necessarily provably robust), we can find a floating-point neural network  $\nu$  that is *identical* to  $\hat{f}$  and is *provably* robust with interval analysis. We also prove a nontrivial result about “floating-point completeness”, as an unexpected byproduct of the main theorem. Specifically, we show that the class of straight-line floating-point programs that use only floating-point  $+$  and  $\times$  operations is *floating-point interval-complete*: it can simulate *any* terminating floating-point program that takes finite floats as input and returns arbitrary floats as output. The same statement holds under the interval semantics. To our knowledge, no prior work has identified such a small yet powerful class of floating-point programs, suggesting that this corollary is of significant independent interest to the extensive floating-point literature.

**Contributions.** This article makes the following contributions:

- We formalize a *floating-point* analog of the *interval universal approximation* (IUA) theorem, to bridge the theory and practice of *provably robust* neural networks (Sect. 2, Sect. 3). It asks if there is a floating-point network whose interval semantics is close to the direct image map of a given target function.
- We prove the floating-point version of the IUA theorem does hold, for *all* target functions and a *broad* class of activation functions that includes most of the activations used in practice (Sect. 3.1, Sect. 3.2, Sect. 5). This shows no fundamental limit exists on the expressiveness of provably robust networks over floats.
- We rigorously analyze the essential differences between the previous IUA theorem over reals and our IUA theorem over floats (Sect. 3.3). Unlike real-valued networks, floating-point networks can *perfectly* capture the behavior of *any* rounded target function, even with the *identity* activation function.
- We prove that if there exists an ideal robust floating-point classifier, then one can always find a *provably* robust floating-point network that makes *exactly* the same prediction as the classifier (Sect. 4.1).
- We prove that the set of straight-line floating-point programs with only  $(+, \times)$  is *floating-point interval-complete*: it can simulate *any* terminating floating-point programs that take finite inputs and return finite/infinite outputs, under the usual floating-point semantics and interval semantics (Sect. 4.2).

## 2 Preliminaries

This section introduces floating-point arithmetic (Sect. 2.1), neural networks that compute over floating-point numbers (Sect. 2.2), and interval analysis for neural networks (Sect. 2.3). Throughout the paper, we define  $\mathbb{N}$  to be the set of positive integers and let  $[n] := \{1, \dots, n\}$  for each  $n \in \mathbb{N}$ .

## 2.1 Floating Point

**Floating-Point Numbers.** Let  $E, M \in \mathbb{N}$ . The set of *finite* floating-point numbers with  $E$ -bit exponent and  $(M+1)$ -bit significand is typically defined by

$$\mathbb{F}_M^E := \{(-1)^b \times (s_0.s_1 \dots s_M)_2 \times 2^e \mid b, s_i \in \{0, 1\}, e \in \{\epsilon_{\min}, \dots, \epsilon_{\max}\}\}, \quad (2)$$

where  $\epsilon_{\min} := -2^{E-1} + 2$  and  $\epsilon_{\max} := 2^{E-1} - 1$  [52]. The set of *all* floating-point numbers, including non-finite ones, is then defined by  $\bar{\mathbb{F}}_M^E := \mathbb{F}_M^E \cup \{-\infty, +\infty, \perp\}$ , where  $\perp$  denotes NaN (i.e., not-a-number). For brevity, we call a floating-point number simply a *float*, and write  $\mathbb{F}_M^E$  and  $\bar{\mathbb{F}}_M^E$  simply as  $\mathbb{F}$  and  $\bar{\mathbb{F}}$ . In this paper, we assume  $E \geq 5$  and  $2^{E-1} \geq M \geq 3$ , which hold for nearly all practical floating-point formats, including bfloat16 [1] and all the formats defined in the IEEE-754 standard [31] such as float16, float32, and float64.

We introduce several notations and terms related to finite floats. First, we define three key constants: the *smallest* positive float  $\omega := 2^{\epsilon_{\min}-M}$ , the *largest* positive float  $\Omega := 2^{\epsilon_{\max}}(2 - 2^{-M})$ , and the *machine epsilon*  $\varepsilon := 2^{-M-1}$ . Next, consider a finite float  $x \in \mathbb{F}$ . We call  $x$  a *subnormal* number if  $0 < |x| < 2^{\epsilon_{\min}}$ , and a *normal* number otherwise. The *exponent* and *significand* of  $x$  are defined by  $\epsilon_x := \max\{|\log_2|x|\|, \epsilon_{\min}\} \in [\epsilon_{\min}, \epsilon_{\max}]$  and  $\mathfrak{s}_x := |x|/2^{\epsilon_x} \in [0, 2)$ . We use  $\mathfrak{s}_{x,0}, \dots, \mathfrak{s}_{x,M}$  to denote the binary expansion of  $\mathfrak{s}_x$ , i.e.,  $(\mathfrak{s}_{x,0}.\mathfrak{s}_{x,1} \dots \mathfrak{s}_{x,M})_2 = \mathfrak{s}_x$  with  $\mathfrak{s}_{x,i} \in \{0, 1\}$ . The *predecessor* and *successor* of  $x$  in  $\bar{\mathbb{F}}$  are written as  $x^- := \max\{y \in \mathbb{F} \setminus \{\perp\} \mid x > y\}$  and  $x^+ := \min\{y \in \bar{\mathbb{F}} \setminus \{\perp\} \mid x < y\}$ .

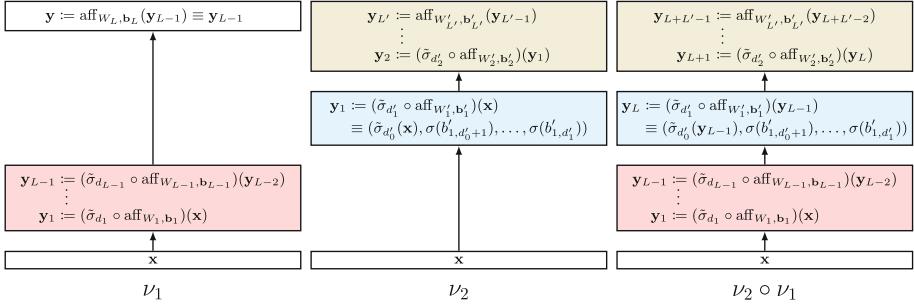
**Floating-Point Operations.** We define the *rounding function*  $\text{rnd} : \mathbb{R} \cup \{-\infty, +\infty\} \rightarrow \bar{\mathbb{F}}$  as follows:  $\text{rnd}(x) := -\infty$  if  $x \in [-\infty, -\Omega - c]$ ,  $\text{rnd}(x) := \arg \min_{y \in \mathbb{F}} |y - x|$  if  $x \in (-\Omega - c, \Omega + c)$ , and  $\text{rnd}(x) := +\infty$  if  $x \in [\Omega + c, +\infty]$ , where  $c := 2^{\epsilon_{\max}}\varepsilon$  and  $\arg \min$  breaks ties by choosing a float  $y$  with  $\mathfrak{s}_{y,M} = 0$ . This function corresponds to the rounding mode “round to nearest (ties to even)”, which is the default rounding mode in the IEEE-754 standard [31].

The floating-point *arithmetic operations*  $\oplus, \ominus, \otimes : \bar{\mathbb{F}} \times \bar{\mathbb{F}} \rightarrow \bar{\mathbb{F}}$  are defined via the rounding function: for finite floats  $x, y \in \mathbb{F}$ ,  $x \oplus y := \text{rnd}(x + y)$ ,  $x \ominus y := \text{rnd}(x - y)$ , and  $x \otimes y := \text{rnd}(x \times y)$ . We omit the definition for non-finite operands because they are unimportant in this paper, except that  $x \oplus 0 = x \ominus 0 = x$  for all  $x \in \{-\infty, +\infty\}$ . For the full definition, refer to the IEEE-754 standard [31].

We introduce two more floating-point operations:  $\text{aff}_{W,\mathbf{b}}$  and  $\text{rnd}(f)$ . First, we define the floating-point *affine transformation*: for a matrix  $W = (w_{i,j})_{i \in [m], j \in [n]} \in \mathbb{F}^{m \times n}$  and a vector  $\mathbf{b} = (b_1, \dots, b_m) \in \mathbb{F}^m$ ,  $\text{aff}_{W,\mathbf{b}} : \bar{\mathbb{F}}^n \rightarrow \bar{\mathbb{F}}^m$  is defined by

$$\text{aff}_{W,\mathbf{b}}(x_1, \dots, x_n) := \left( \left( \sum_{j=1}^n x_j \otimes w_{1,j} \right) \oplus b_1, \dots, \left( \sum_{j=1}^n x_j \otimes w_{m,j} \right) \oplus b_m \right). \quad (3)$$

Here,  $\sum$  denotes the floating-point summation defined in the left-associative way:  $\sum_{i=1}^n y_i := (\dots((y_1 \oplus y_2) \oplus y_3) \dots) \oplus y_n$ , where the order of  $\oplus$  is important



**Fig. 2.** Illustrations of a network  $\nu_1$  without the last affine layer (left), a network  $\nu_2$  without the first affine layer (middle), and their composition  $\nu_2 \circ \nu_1$  (right). Note that  $\text{aff}_{W'_1, b'_1} \circ \text{aff}_{W_L, b_L} = \text{aff}_{W'_1, b'_1}$  is a floating-point affine transformation.

because  $\oplus$  is not associative. Next, we define the *correctly rounded version* of a real-valued function. For  $f : \mathbb{R} \rightarrow \mathbb{R}$ , the function  $\text{rnd}(f) : \overline{\mathbb{F}} \rightarrow \overline{\mathbb{F}}$  is defined by

$$\text{rnd}(f)(x) := \begin{cases} \text{rnd}(f(x)) & \text{if } x \in (-\infty, +\infty) \\ \text{rnd}\left(\lim_{t \rightarrow x} f(t)\right) & \text{if } x \in \{-\infty, +\infty\} \wedge \lim_{t \rightarrow x} f(t) \in \mathbb{R} \cup \{-\infty, +\infty\} \\ \perp & \text{otherwise.} \end{cases} \quad (4)$$

## 2.2 Neural Networks

A neural network typically refers to a composition of affine transformations and activation functions. Formally, for  $L \in \mathbb{N}$  and  $\sigma : \overline{\mathbb{F}} \rightarrow \overline{\mathbb{F}}$ , we call a function  $\nu$  a *depth- $L$   $\sigma$ -neural network* (or a *neural network*) if  $\nu$  is defined by

$$\nu : \overline{\mathbb{F}}^{d_0} \rightarrow \overline{\mathbb{F}}^{d_L}, \quad \nu := \text{aff}_{W_L, b_L} \circ \tilde{\sigma}_{d_{L-1}} \circ \text{aff}_{W_{L-1}, b_{L-1}} \circ \cdots \circ \tilde{\sigma}_{d_1} \circ \text{aff}_{W_1, b_1} \quad (5)$$

for some  $d_\ell \in \mathbb{N}$ ,  $W_\ell \in \mathbb{F}^{d_\ell \times d_{\ell-1}}$ , and  $\mathbf{b}_\ell \in \mathbb{F}^{d_\ell}$ , where  $\tilde{\sigma}_n : \overline{\mathbb{F}}^n \rightarrow \overline{\mathbb{F}}^n$  is the coordinatewise application of  $\sigma$ . Here,  $L$  denotes the number of layers,  $\sigma$  the floating-point activation function,  $d_0$  and  $d_L$  the input and output dimensions,  $d_\ell$  the number of hidden neurons in the  $\ell$ -th layer ( $\ell \in [L-1]$ ), and  $W_\ell$  and  $\mathbf{b}_\ell$  the parameters of the floating-point affine transformation in the  $\ell$ -th layer ( $\ell \in [L]$ ). We emphasize that a neural network in this paper is a function over *floating-point* values, defined in terms of *floating-point* activation function and arithmetic. For instance, a depth-1 neural network is a floating-point affine transformation.

Let  $\nu$  be a neural network defined as Eq. (5). We say  $\nu$  is *without the last affine layer* if  $d_L = d_{L-1}$ ,  $W_L$  is the identity matrix, and  $\mathbf{b}_L = \mathbf{0}$ . Similarly, we say  $\nu$  is *without the first affine layer* if  $d_1 \geq d_0$ ,  $W_1$  is a rectangular diagonal matrix whose diagonal entries are all 1, and  $b_{1,i} = 0$  for all  $i \in [d_0]$ . The two definitions are not perfectly symmetric due to some technical details arising in our proofs. We note that a neural network can be constructed by composing

networks without the first/last affine layer(s) and arbitrary networks (Fig. 2). For example, consider arbitrary networks  $\nu_1 : \bar{\mathbb{F}}^{n_0} \rightarrow \bar{\mathbb{F}}^{n_1}$  and  $\nu_4 : \bar{\mathbb{F}}^{n_3} \rightarrow \bar{\mathbb{F}}^{n_4}$ , a network without the first affine layer  $\nu_2 : \bar{\mathbb{F}}^{n_1} \rightarrow \bar{\mathbb{F}}^{n_2}$ , and a network without the first and last affine layers  $\nu_3 : \bar{\mathbb{F}}^{n_2} \rightarrow \bar{\mathbb{F}}^{n_3}$ . It is easily verified that the function  $\nu : \bar{\mathbb{F}}^{n_0} \rightarrow \bar{\mathbb{F}}^{n_4}$  specified by  $\nu(\mathbf{x}) = (\nu_4 \circ \dots \circ \nu_1)(\mathbf{x})$  denotes a network, whose definition in the form of Eq. (5) can be obtained by “merging” the last layer of  $\nu_1$  and the first layer of  $\nu_2$ , etc.

### 2.3 Interval Semantics

Interval analysis [11, 50] is a technique for analyzing the behavior of numerical programs soundly and efficiently, based on abstract interpretation [10]. It uses intervals to overapproximate the ranges of inputs and expressions, and propagates them through a program to overapproximate the output range. Interval analysis has been used to establish the robustness of practical neural networks [19, 24, 33, 43]. It can overapproximate the output range of a network over perturbed inputs, which is required to prove robustness; and it runs efficiently by performing only simple computations, which is required to analyze large-scale networks.

**Interval Domain and Operations.** We formalize interval analysis for neural networks as follows. We first define the *interval domain*

$$\mathbb{I} := \{\langle a, b \rangle \mid a, b \in \bar{\mathbb{F}} \setminus \{\perp\} \text{ with } a \leq b\} \cup \{\top\}, \quad (6)$$

on which interval analysis operates. Here,  $\langle a, b \rangle$  abstracts the floating-point interval  $[a, b] \cap \bar{\mathbb{F}}$ , and  $\top$  abstracts the entire floating-point set  $\bar{\mathbb{F}}$  including  $\perp$ . The concrete semantics of an *abstract interval*  $\mathcal{I} \in \mathbb{I}$  and an *abstract box*  $\mathcal{B} = (\mathcal{I}_1, \dots, \mathcal{I}_d) \in \mathbb{I}^d$  are defined through the *concretization function*  $\gamma$ , where

$$\gamma : \cup_{d=1}^{\infty} \mathbb{I}^d \rightarrow \cup_{d=1}^{\infty} 2^{\bar{\mathbb{F}}^d}, \quad \gamma(\mathcal{I}) := \begin{cases} [a, b] \cap \bar{\mathbb{F}} & \text{if } \mathcal{I} = \langle a, b \rangle \\ \bar{\mathbb{F}} & \text{if } \mathcal{I} = \top \end{cases}, \quad \gamma(\mathcal{B}) := \prod_{i=1}^d \gamma(\mathcal{I}_i). \quad (7)$$

We say that an abstract box  $\mathcal{B} \in \mathbb{I}^d$  is in a set  $\mathcal{S} \subseteq \mathbb{R}^d$  if  $\gamma(\mathcal{B}) \subseteq \mathcal{S}$ .

For any function  $\phi : \bar{\mathbb{F}}^d \rightarrow \bar{\mathbb{F}}$  over floats (which is not a neural network or a floating-point affine transformation), the *interval operation*  $\phi^\sharp : \mathbb{I}^d \rightarrow \mathbb{I}$  extends  $\phi$  to the interval domain as follows:

$$\phi^\sharp(\mathcal{B}) := \begin{cases} \langle \min \mathcal{S}, \max \mathcal{S} \rangle & \text{if } \perp \notin \mathcal{S} \\ \top & \text{if } \perp \in \mathcal{S} \end{cases}, \quad \text{where } \mathcal{S} := \phi(\gamma(\mathcal{B})). \quad (8)$$

In the special case that  $\phi = \odot \in \{\oplus, \ominus, \otimes\}$  is a floating-point arithmetic operation, the above definition (using infix notation) is equivalent to the following:

$$\langle a, b \rangle \odot^\sharp \langle c, d \rangle := \begin{cases} \langle \min \mathcal{S}, \max \mathcal{S} \rangle & \text{if } \perp \notin \mathcal{S} \\ \top & \text{if } \perp \in \mathcal{S} \end{cases}, \quad \text{where } \mathcal{S} := \begin{cases} a \odot c, a \odot d, \\ b \odot c, b \odot d \end{cases}, \quad (9)$$

and  $\odot^\#$  returns  $\top$  if at least one of its operands is  $\top$ .<sup>1</sup> We remark that  $\odot^\#$  can be efficiently computed, and so can  $\phi^\#$  when  $\phi : \overline{\mathbb{F}} \rightarrow \overline{\mathbb{F}}$  is piecewise-monotone with finitely many pieces, which holds for the correctly rounded versions of widely-used activation functions (e.g., ReLU, GELU, sigmoid). We then define the *interval affine transformation*  $\text{aff}_{W, \mathbf{b}}^\# : \mathbb{I}^n \rightarrow \mathbb{I}^m$ , which extends its floating-point counterpart  $\text{aff}_{W, \mathbf{b}} : \overline{\mathbb{F}}^n \rightarrow \overline{\mathbb{F}}^m$ :  $\text{aff}_{W, \mathbf{b}}^\#(\mathcal{I}_1, \dots, \mathcal{I}_n) := \left( \left( \overline{\sum}_{j=1}^n \# \mathcal{I}_j \otimes^\# \langle w_{i,j}, w_{i,j} \rangle \right) \oplus^\# \langle b_i, b_i \rangle \right)_{i=1}^m$ , where  $\overline{\sum}^\#$  is the interval summation which uses  $\oplus^\#$  instead of  $\oplus$ .

**Interval Semantics.** The *interval semantics*  $\nu^\# : \mathbb{I}^{d_0} \rightarrow \mathbb{I}^{d_L}$  of a neural network  $\nu : \overline{\mathbb{F}}^{d_0} \rightarrow \overline{\mathbb{F}}^{d_L}$  is defined as the result of interval analysis on  $\nu$ :

$$\nu^\# := \text{aff}_{W_L, \mathbf{b}_L}^\# \circ \tilde{\sigma}_{d_{L-1}}^\# \circ \text{aff}_{W_{L-1}, \mathbf{b}_{L-1}}^\# \circ \dots \circ \tilde{\sigma}_{d_1}^\# \circ \text{aff}_{W_1, \mathbf{b}_1}^\#, \quad (10)$$

where  $\nu$  is assumed to be defined as Eq. (5) and  $\tilde{\sigma}_n^\# : \mathbb{I}^n \rightarrow \mathbb{I}^n$  is the coordinate-wise application of  $\sigma^\# : \mathbb{I} \rightarrow \mathbb{I}$ . It is easily verified that the interval semantics is sound with respect to the floating-point semantics:

$$\nu(\gamma(\mathcal{B})) \subseteq \gamma(\nu^\#(\mathcal{B})) \quad (\mathcal{B} \in \mathbb{I}^{d_0}). \quad (11)$$

That is, the result of interval analysis  $\nu^\#(\mathcal{B}) \in \mathbb{I}^{d_L}$  subsumes the set of all possible outputs of the network  $\nu$  when the input is in the concrete box  $\gamma(\mathcal{B}) \subseteq \overline{\mathbb{F}}^{d_0}$ .

### 3 Interval Universal Approximation Over Floats

This section presents our main result on interval universal approximation (IUA) for floating-point neural networks. We first introduce conditions on activation functions for our result (Sect. 3.1), and then formally describe our result under these conditions (Sect. 3.2). We then compare our IUA theorem over floats with existing IUA theorems over reals, highlighting several nontrivial differences (Sect. 3.3).

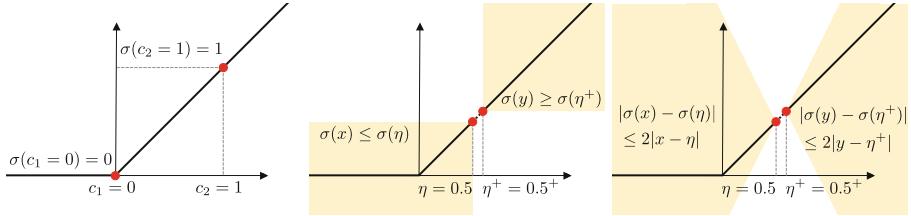
#### 3.1 Conditions on Activation Functions

Our IUA theorem is for floating-point neural networks that use activation functions satisfying the following conditions (Fig. 3).

**Condition 1.** *An activation function  $\sigma : \overline{\mathbb{F}} \rightarrow \overline{\mathbb{F}}$  satisfies the following conditions:*

---

<sup>1</sup> This definition of  $\odot^\#$  differs slightly from the standard definition, as  $\odot^\#$  uses “round to nearest” mode (implicit in  $\odot$ ), whereas the more common mode is “round downward/upward” (e.g.,  $\langle a, b \rangle \oplus^\# \langle c, d \rangle := \langle a \oplus_\downarrow c, b \oplus_\uparrow d \rangle$ ) [26, Section 5]. This choice is due to different goals to achieve: our definition overapproximates floating-point operations (e.g.,  $\oplus$ ), while the usual one overapproximates exact operations (e.g.,  $+$ ).



**Fig. 3.** Illustration of the first (left), second (middle), and third (right) conditions in Condition 1 for the ReLU activation function:  $\sigma(x) := \max\{x, 0\}$  for  $x \in \overline{\mathbb{F}}$ .

- (C1) There exist  $c_1, c_2 \in \mathbb{F}$  such that  $\sigma(c_1) = 0$ ,  $|\sigma(c_2)| \in [\frac{\varepsilon}{2} + 2\varepsilon^2, \frac{5}{4} - 2\varepsilon]$ ,  $\max\{|c_1|, |c_2|\} \geq 2^{\epsilon_{\min}+1}$ , and  $\sigma(x)$  lies between  $\sigma(c_1)$  and  $\sigma(c_2)$  for all  $x$  between  $c_1$  and  $c_2$ , where  $\varepsilon$  is the machine epsilon (see Sect. 2.1).
- (C2) There exists  $\eta \in \mathbb{F}$  with  $|\eta| \in [2^{\epsilon_{\min}+5}, 4-8\varepsilon]$  and  $|\sigma(\eta)|, |\sigma(\eta^+)| \in [2^{\epsilon_{\min}+5}, 2^{\epsilon_{\max}-6} \cdot |\eta|]$  such that for any  $x, y \in \mathbb{F}$  with  $x \leq \eta < \eta^+ \leq y$ ,

$$\sigma(x) \leq \sigma(\eta) < \sigma(\eta^+) \leq \sigma(y) \quad \text{or} \quad \sigma(x) \geq \sigma(\eta) > \sigma(\eta^+) \geq \sigma(y). \quad (12)$$

- (C3) There exists  $\lambda \in [0, 2^{\epsilon_{\max}-7} \cdot \min\{|\sigma(\eta)|, 2^{M+3}\}]$  such that for any  $x, y \in \mathbb{F}$  with  $x \leq \eta < \eta^+ \leq y$ ,

$$|\sigma(x) - \sigma(\eta)| \leq \lambda|x - \eta| \quad \text{and} \quad |\sigma(y) - \sigma(\eta^+)| \leq \lambda|y - \eta^+|. \quad (13)$$

The condition (C1) states that the activation function  $\sigma$  can output the exact zero (i.e.,  $\sigma(c_1)$ ) and some value whose magnitude is approximately in  $[\frac{\varepsilon}{2}, \frac{5}{4}]$  (i.e.,  $\sigma(c_2)$ ); and its output is within  $\sigma(c_1)$  and  $\sigma(c_2)$  for all inputs between  $c_1$  and  $c_2$ . The condition (C2) states that there exists some *threshold*  $\eta$  such that  $\sigma(x)$  is either smaller or greater than  $\sigma(\eta)$  or  $\sigma(\eta^+)$ , depending on whether  $x$  is on the left or right side of  $\eta$ . This condition holds automatically for all monotone activation functions that are non-constant on either  $[2^{\epsilon_{\min}+5}, 4-8\varepsilon] \cap \mathbb{F}$  or  $[-4+8\varepsilon, -2^{\epsilon_{\min}+5}] \cap \mathbb{F}$ . The condition (C3) states that  $\sigma$  does not increase or decrease too rapidly from  $\eta$  and  $\eta^+$ , which implies that  $\sigma(x)$  is finite for all finite floats  $x \in \mathbb{F}$ .

While Condition 1 is mild, verifying whether practical activation functions over floats satisfy Condition 1 can be cumbersome. Floating-point activation functions are typically implemented in complicated ways [7, 44, 51] (e.g., by intermixing floating-point operations with integer/bit-level operations and if-else branches), which makes it challenging to rigorously analyze such implementations [17, 36]. To bypass this issue, we focus on the correctly rounded version  $\sigma : \overline{\mathbb{F}} \rightarrow \overline{\mathbb{F}}$  of a real-valued activation function  $\rho : \mathbb{R} \rightarrow \mathbb{R}$  (i.e.,  $\sigma(x) := \text{rnd}(\rho(x))$ ), when verifying Condition 1. Correctly rounded versions of elementary mathematical functions have been actively developed in several software libraries [13, 39, 40, 58, 72].

Under the correct rounding assumption, we provide an easily verifiable sufficient condition for activation functions on reals that can be used to verify Condition 1 for their rounded versions. The proof of Lemma 1 is in §B.1.

**Lemma 1.** *For any activation function  $\rho : \mathbb{R} \rightarrow \mathbb{R}$ , the correctly rounded activation  $\text{rnd}(\rho) : \overline{\mathbb{F}} \rightarrow \overline{\mathbb{F}}$  satisfies Condition 1 if the following conditions hold:*

(C1') *There exist  $c'_1, c'_2 \in \mathbb{F}$  such that  $|\rho(c'_1)| \leq \frac{\omega}{2}$ ,  $|\rho(c'_2)| \in [\frac{\varepsilon}{2} + 2\varepsilon^2, \frac{5}{4} - 2\varepsilon]$ ,  $\max\{|c'_1|, |c'_2|\} \geq 2^{\varepsilon_{\min}+1}$ , and  $\rho(x)$  lies between  $\rho(c'_1)$  and  $\rho(c'_2)$  for all  $x$  between  $c'_1$  and  $c'_2$ , where  $\omega$  is the smallest positive float (see Sect. 2.1).*

(C2') *There exists  $\delta \in \mathbb{R}$  with  $|\delta| \in [\frac{3}{8}, \frac{7}{8}]$  such that*

- *for all  $x, y \in \mathbb{R}$  satisfying  $x \leq \delta - \frac{1}{8} < \delta + \frac{1}{8} \leq y$ ,*

$$\rho(x) \leq \rho(\delta - \frac{1}{8}) < \rho(\delta + \frac{1}{8}) \leq \rho(y) \quad \text{or} \quad \rho(x) \geq \rho(\delta - \frac{1}{8}) > \rho(\delta + \frac{1}{8}) \geq \rho(y),$$

- *$|\rho(x)| \in [\frac{1}{4}, 1]$  and  $|\rho(x) - \rho(y)| > \frac{1}{8}|x - y|$  for all  $x, y \in [\delta - \frac{1}{8}, \delta + \frac{1}{8}]$ .*

(C3')  *$\rho$  is  $\lambda$ -Lipschitz continuous for some  $\lambda \in [0, \frac{1}{5} \cdot 2^{\varepsilon_{\max}-9}]$ .*

The Conditions (C1')–(C3') in Lemma 1 correspond to the Conditions (C1)–(C3) in Condition 1. The Condition (C1'), corresponding to (C1), can be easily satisfied since modern activation functions are piecewise-monotone and either zero at zero (e.g., ReLU, GELU, softplus, tanh) or close to zero at  $-\Omega$  or  $\Omega$  (e.g., sigmoid). The Condition (C2') roughly states the existence of  $\delta \in \mathbb{R}$  satisfying the following: (i)  $\rho(\delta - \frac{1}{8})$  and  $\rho(\delta + \frac{1}{8})$  are lower/upper bounds of  $\rho$  on  $(-\infty, \delta - \frac{1}{8})$  and  $(\delta + \frac{1}{8}, \infty)$ ; and (ii)  $\rho$  is bounded and strictly monotone on  $[\delta - \frac{1}{8}, \delta + \frac{1}{8}]$ . This condition guarantees the existence of  $\eta \in \mathbb{F}$  in (C2). The Condition (C3'), corresponding to (C3), can also be easily satisfied since  $\lambda < 3$  for most practical activation functions. We note that Lemma 1 gives sufficient but not necessary conditions for a correctly rounded activation function to satisfy Condition 1.

The following corollary uses Lemma 1 to show that many prominent activation functions satisfy Condition 1. Its proof is in §B.2.

**Corollary 1.** *The correctly rounded implementations of the ReLU, LeakyReLU, GELU, ELU, Mish, softplus, sigmoid, and tanh activations satisfy Condition 1.*

### 3.2 Main Result

We are now ready to present our IUA theorem over floating-point arithmetic.

**Theorem 1.** *Let  $\sigma : \overline{\mathbb{F}} \rightarrow \overline{\mathbb{F}}$  be an activation function satisfying Condition 1.<sup>2</sup> Then, for any target function  $f : \mathbb{R}^d \rightarrow \mathbb{R}$ , there exists a  $\sigma$ -neural network  $\nu : \overline{\mathbb{F}}^d \rightarrow \overline{\mathbb{F}}$  such that*

$$\gamma(\nu^\sharp(\mathcal{B})) = \left[ \min \widehat{f}(\gamma(\mathcal{B})), \max \widehat{f}(\gamma(\mathcal{B})) \right] \cap \overline{\mathbb{F}} \quad (14)$$

<sup>2</sup> Condition 1 is sufficient for Theorem 1 but not necessary. E.g., Theorem 1 still holds under 8-bit floats (both E4M3 and E5M2 formats [46]) for the ReLU activation function; this corresponds to the case where (C1) and (C2) hold but (C3) is violated.

for  $\hat{f} = \text{rnd}(f) : \bar{\mathbb{F}}^d \rightarrow \bar{\mathbb{F}}^d$  and for all abstract boxes  $\mathcal{B}$  in  $[-1, 1]^d$ .<sup>3</sup>

Theorem 1 states that for any activation function  $\sigma : \bar{\mathbb{F}} \rightarrow \bar{\mathbb{F}}$  satisfying Condition 1 and any target function  $f : \mathbb{R}^d \rightarrow \mathbb{R}$ , there exists a  $\sigma$ -network  $\nu$  whose interval semantics *exactly computes* the upper and lower points of the direct image map of the rounded target  $\hat{f} : \bar{\mathbb{F}}^d \rightarrow \bar{\mathbb{F}}$  on  $[-1, 1]^d \cap \mathbb{F}^d$ . A special case of our IUA Theorem 1 is the following universal approximation (UA) theorem for floating-point neural networks:

$$\nu(\mathbf{x}) = \hat{f}(\mathbf{x}) \quad (\mathbf{x} \in [-1, 1]^d \cap \mathbb{F}^d). \quad (15)$$

That is, floating-point neural networks using an activation function satisfying Condition 1 can represent any function  $\hat{f} : [-1, 1]^d \cap \mathbb{F}^d \rightarrow \mathbb{F} \cup \{-\infty, +\infty\}$ ; or the rounded version of any real function  $f : [-1, 1]^d \rightarrow \mathbb{R}$ . Moreover, Theorem 1 easily extends to any target function  $f : \mathbb{R}^d \rightarrow \mathbb{R}^d$  with multiple outputs.

As previous IUA results assume exact operations over reals, they do not extend to our setting of floating-point arithmetic (due to rounding errors, overflow, NaNs, discreteness, boundedness, etc.). As a simple example of these issues, consider the following subnetwork, which is used in the IUA proof of [4]:

$$\mu(x, y) = \frac{1}{2} (\text{ReLU}(x + y) - \text{ReLU}(-x - y) - \text{ReLU}(x - y) - \text{ReLU}(y - x)).$$

This subnetwork returns  $\min\{x, y\}$  if all operations are exact. However, it does not under floating-point arithmetic due to the rounding error: if  $(+, \times)$  is replaced by  $(\oplus, \otimes)$ , then  $\mu(x, y) = 0 \neq \varepsilon = \min\{x, y\}$  for  $x = 1$  and  $y = \varepsilon$ .

In addition, the network construction in [66, Theorem 4.10] requires multiplying a large number  $z$  that depends on the target error and the activation function, to the output of some neuron. However, because  $\mathbb{F}$  is bounded and floating-point operations are subject to overflow, the number  $z$  and the result of the multiplication are not guaranteed to be within  $\mathbb{F}$  when using a small target error (e.g., less than  $\omega$ ) or when using common activation functions (e.g., ReLU, softplus). To bypass these issues, we carefully analyze rounding errors and design a network without infinities in the intermediate layers, when proving Theorem 1.

We present the proof outline of Theorem 1 in Sect. 5, and the full proof in §D–§F. We implemented the proof (i.e., our network construction) in Python and made it available at <https://github.com/yechanp/floating-point-iua-theorem>.

### 3.3 Comparison With Existing Results Over Reals

Theorem 1, which gives an IUA theorem over floats, has notable differences from previous IUA theorems over the reals [4, Theorem 1.1]; [66, Theorem 3.7].

<sup>3</sup> In the literature on universal approximation theorems, it is typically assumed that the inputs are in  $[0, 1]$  or in a compact subset of  $\mathbb{R}$  (e.g., [4, 12, 66, 69]). Since the inputs are often normalized to  $[-1, 1]$ , we focus the theoretical analysis on  $[-1, 1]^d$ .

One difference is the class of target functions and the desired property of networks. Previous IUA theorems find a network that *sufficiently approximates* the direct image map of a *continuous* target function (i.e.,  $\delta > 0$  in Eq. (1)). In contrast, our IUA theorem finds a network that *exactly computes* the direct image map of an *arbitrary* rounded target function (i.e.,  $\delta = 0$  in Eq. (1)). This difference arises from the domains of the functions being approximated: the real-valued setting considers functions  $f$  over  $[-1, 1]^d$  (or a compact  $\mathcal{K} \subset \mathbb{R}^d$ ); the floating-point setting considers functions  $\hat{f}$  over  $[-1, 1]^d \cap \mathbb{F}^d$ .

- Since  $[-1, 1]^d$  is uncountable, exactly computing the direct image map of  $f$  requires a network to fit *uncountably* many input/output pairs and related box/interval pairs. This task is difficult to achieve, and indeed, recent works [3, 5, 47] prove that it is theoretically unachievable even for simple target functions (e.g., continuous piecewise linear functions).
- Since  $[-1, 1]^d \cap \mathbb{F}^d$  is finite, exactly computing the direct image map of  $\hat{f}$  requires a network to fit *finitely* many input/output and box/interval pairs. Our result proves that, despite all the complexities of floating-point computation, this task can be achieved for any rounded target function.

Another key difference is the class of activation functions. There are real-valued activation functions  $\rho, \rho' : \mathbb{R} \rightarrow \mathbb{R}$  such that previous IUA theorems *cannot* hold for  $\rho$  but our IUA theorem *does* hold for  $\text{rnd}(\rho) : \bar{\mathbb{F}} \rightarrow \bar{\mathbb{F}}$ ; and vice versa for  $\rho'$ .

- An example of  $\rho$  is the *identity* function:  $\rho(x) = x$ . No classical IUA or UA theorem can hold for  $\rho$ , since all *real-valued*  $\rho$ -networks  $\mu : \mathbb{R}^d \rightarrow \mathbb{R}$  are *affine over the reals* (i.e., there exists  $A \in \mathbb{R}^{1 \times d}$  and  $b \in \mathbb{R}$  such that  $\mu(\mathbf{x}) = A\mathbf{x} + b$  for all  $\mathbf{x} \in \mathbb{R}^d$ ). In contrast, our IUA theorem does hold for  $\text{rnd}(\rho)$ , because  $\text{rnd}(\rho)$  satisfies all the conditions in Lemma 1 (with constants  $c'_1 = 0$ ,  $c'_2 = 1$ ,  $\delta = 1/2$ , and  $\lambda = 1$ ). This counterintuitive result is made possible because *floating-point*  $\text{rnd}(\rho)$ -networks  $\nu : \bar{\mathbb{F}}^d \rightarrow \bar{\mathbb{F}}$  can be *non-affine over the reals* (i.e., there may not exist  $A \in \mathbb{R}^{1 \times d}$  and  $b \in \mathbb{R}$  such that  $\nu(\mathbf{x}) = A\mathbf{x} + b$  for all  $\mathbf{x} \in \mathbb{F}^d$ ). This non-affineness arises from rounding errors: some floating-point affine transformations  $\text{aff}_{W, b}$  are not actually affine over the reals due to rounding errors. An interesting implication of this result is discussed in Sect. 4.2.
- An example of  $\rho'$  is any function that is non-decreasing on  $\mathbb{R}$ , is constant on  $[-\Omega, \Omega]$ , and satisfies  $\lim_{x \rightarrow -\infty} \rho'(x) < \lim_{x \rightarrow +\infty} \rho'(x)$ , where the two limits exist in  $\mathbb{R}$ . The real-valued IUA theorem holds for  $\rho'$ , because  $\rho'$  satisfies the condition in [66, Definition 2.3]. However, no floating-point IUA or UA theorem can hold for  $\text{rnd}(\rho')$ , because all  $\text{rnd}(\rho')$ -networks  $\nu : \bar{\mathbb{F}} \rightarrow \bar{\mathbb{F}}$  must be monotone if its depth is 1, and must satisfy  $\nu(0) = \nu(\omega)$  otherwise. The monotonicity holds when the depth is 1 since  $\oplus, \otimes$  are monotone when an operand is a constant; and  $\nu(0) = \nu(\omega)$  holds otherwise since  $x \otimes a \oplus b \in [-\Omega, \Omega]$  for all  $x \in \{0, \omega\}$  and  $a, b \in \mathbb{F}$ , and  $\text{rnd}(\rho')$  is constant on  $[-\Omega, \Omega] \cap \mathbb{F}$ .

## 4 Implications of IUA Theorem Over Floats

This section presents two important implications of our IUA theorem, on provable robustness and ‘‘floating-point completeness’’. We first prove the existence of a provably robust floating-point network, given an ideal robust floating-point classifier (Sect. 4.1). We then prove that floating-point  $+$  and  $\times$  are sufficient to simulate all halting programs that return finite/infinite floats when given finite floats (Sect. 4.2).

### 4.1 Provable Robustness of Neural Networks

Consider the task of classifying floating-point inputs  $\mathbf{x} \in \mathcal{X}$  (e.g., images of objects) into  $n \in \mathbb{N}$  classes (e.g., categories of objects), where  $\mathcal{X} := [-1, 1]^d \cap \mathbb{F}^d$  denotes the space of inputs throughout this subsection. For this task, a function  $f : \mathcal{X} \rightarrow \mathbb{F}^n$  is often viewed as a classifier in the following sense:  $f$  predicts  $\mathbf{x}$  to be in the  $i$ -th class ( $i \in [n]$ ), where  $i := \text{class}(f(\mathbf{x}))$  and  $\text{class} : \mathbb{F}^n \rightarrow [n]$  is defined by  $\text{class}(y_1, \dots, y_n) := \arg \max_{i \in [n]} y_i$  with an arbitrary tie-breaking rule.

A typical robustness property of a classifier  $f$  is that  $f$  should predict the same class for all neighboring inputs under the  $\ell_\infty$  distance [38]. We formalize this notion of *robust* classifiers in a way similar to [66, Definition A.4].

**Definition 1.** Let  $\delta > 0$  and  $\mathcal{D} \subseteq \mathcal{X}$ . A classifier  $f : \mathcal{X} \rightarrow \mathbb{F}^n$  is called  $\delta$ -robust on  $\mathcal{D}$  if for all  $\mathbf{x}_0 \in \mathcal{D}$ ,  $\mathbf{y}, \mathbf{y}' \in f(\mathcal{N}_\delta(\mathbf{x}_0))$  implies  $\text{class}(\mathbf{y}) = \text{class}(\mathbf{y}')$ , where  $\mathcal{N}_\delta(\mathbf{x}_0) := \{\mathbf{x} \in \mathcal{X} \mid \|\mathbf{x}_0 - \mathbf{x}\|_\infty \leq \delta\}$  and  $\|\cdot\|_\infty$  denotes the  $\ell_\infty$ -norm.

Neural networks have been widely used as classifiers, but establishing the robustness properties of practical networks as in Definition 1 is intractable due to the enormous number of inputs to be checked (i.e.,  $|\mathcal{N}_\delta(\mathbf{x}_0)| \gg 1$  when  $d \gg 1$ ). Instead, these properties have been proven often by using interval analysis, as mentioned in Sect. 2.3. We formalize the notion of such *provably robust* networks under interval analysis, in a way similar to [66, Definition A.5].

**Definition 2.** Let  $\delta > 0$  and  $\mathcal{D} \subseteq \mathcal{X}$ . A neural network  $\nu : \overline{\mathbb{F}}^d \rightarrow \overline{\mathbb{F}}^n$  is called  $\delta$ -provably robust on  $\mathcal{D}$  if for all  $\mathbf{x}_0 \in \mathcal{D}$ ,  $\mathbf{y}, \mathbf{y}' \in \gamma(\nu^\sharp(\mathcal{B}))$  implies  $\text{class}(\mathbf{y}) = \text{class}(\mathbf{y}')$ , where  $\mathcal{B} \in \mathbb{I}^d$  denotes the abstract box such that  $\gamma(\mathcal{B}) = \mathcal{N}_\delta(\mathbf{x}_0)$ .

Under these definitions, we prove that given an ideal robust classifier  $f$ , we can always find a neural network  $\nu$  (i) whose robustness property is *exactly* the same as that of  $f$  and is easily provable using only interval analysis, and (ii) whose predictions are *precisely* equal to those of  $f$ .

**Theorem 2.** Let  $f : \mathcal{X} \rightarrow \mathbb{F}^n$  be a classifier that is  $\delta$ -robust on  $\mathcal{D}$ , and  $\sigma : \overline{\mathbb{F}} \rightarrow \overline{\mathbb{F}}$  be an activation function satisfying Condition 1. Then, there exists a  $\sigma$ -neural network  $\nu : \overline{\mathbb{F}}^d \rightarrow \overline{\mathbb{F}}^n$  that is  $\delta$ -provably robust on  $\mathcal{D}$  and makes the same prediction as  $f$  on  $\mathcal{D}$  (i.e.,  $\text{class}(\nu(\mathbf{x})) = \text{class}(f(\mathbf{x}))$  for all  $\mathbf{x} \in \mathcal{D}$ ).

*Proof Sketch.* We show this (i) by applying Theorem 1 to  $n$  target functions that are constructed from  $f$ , and (ii) by using the following observation: the network constructed in the proof of Theorem 1 has depth not depending on a target function (when  $d$  is fixed). The full proof is in §C.1.  $\square$

## 4.2 Floating-Point Interval-Completeness

To motivate our result, we recall the notion of Turing completeness. A computation model is called *Turing-complete* if for every Turing machine  $T$ , there exists a program in the model that can simulate the machine [6, 35, 49]. Extensive research has established the Turing completeness of numerous computation models: from untyped  $\lambda$ -calculus [8, 64] and  $\mu$ -recursive functions [9, 20], to type systems (e.g., Haskell [67], Java [25]) and neural networks over the rationals (e.g., RNNs [59], Transformers [54]). These results identify simpler computation models as powerful as Turing machines, and shed light on the computational power of new models.

We ask an analogous question for *floating-point* computations instead of *binary* computations, where the former is captured by floating-point programs and the latter by Turing machines. That is, which small class of floating-point programs can simulate all (or nearly all) floating-point programs?

Formally, let  $\mathcal{F}$  be the set of all terminating programs that take finite floats and return finite or infinite floats, where these programs can use any floating-point constants/operations (e.g.,  $-\infty$ ,  $\otimes$ ) and language constructs (e.g., if-else, while). Then,  $\mathcal{F}$  semantically denotes the set of all functions from  $\mathbb{F}^n$  to  $(\mathbb{F} \cup \{-\infty, +\infty\})^m$  for all  $n, m \in \mathbb{N}$ , because each such function can be expressed with if-else branches and floating-point constants. For this class of programs, we define the notion of (*interval-*)*simulation* and *floating-point (interval-)completeness* as follows.

**Definition 3.** Let  $P, Q \in \mathcal{F}$  be programs with arity  $n$ . We say  $Q$  simulates  $P$  if  $Q(\mathbf{x}) = P(\mathbf{x})$  for all  $\mathbf{x} \in \mathbb{F}^n$ , where  $P(\mathbf{x})$  denotes the concrete semantics of  $P$  on  $\mathbf{x}$ . We say  $Q$  interval-simulates  $P$  if  $\gamma(Q^\sharp(\mathcal{B})) = [\min P(\gamma(\mathcal{B})), \max P(\gamma(\mathcal{B}))] \cap \overline{\mathbb{F}}$  for all abstract boxes  $\mathcal{B}$  in  $\mathbb{F}^n$ , where  $Q^\sharp(\mathcal{B})$  denotes the interval semantics of  $Q$  on  $\mathcal{B}$ .

**Definition 4.** We say a class of programs  $\mathcal{G} \subseteq \mathcal{F}$  is floating-point (interval-)complete if for every  $P \in \mathcal{F}$ , there exists  $Q \in \mathcal{G}$  such that  $Q$  (interval-)simulates  $P$ .

We prove that a surprisingly small class of programs is floating-point interval-complete (so floating-point complete). In particular, we show that only floating-point addition, multiplication, and constants are sufficient to interval-simulate *all* halting programs that output finite/infinite floats when given finite floats.

**Theorem 3.**  $\mathcal{F}_{\oplus, \otimes} \subset \mathcal{F}$  is floating-point interval-complete, where  $\mathcal{F}_{\oplus, \otimes}$  denotes the class of straight-line programs that use only  $\oplus$ ,  $\otimes$ , and floating-point constants.

*Proof Sketch.* We show this by extending the key lemma used in the proof of Theorem 1: there exist  $\sigma$ -networks that capture the direct image maps of indicator functions over  $[-1, 1]^n \cap \mathbb{F}^n$  (Lemma 2). In particular, we prove that  $[-1, 1]^n \cap \mathbb{F}^n$  can be extended to  $\mathbb{F}^n$  if  $\sigma$  is the identity function. The full proof is in §C.2.  $\square$

To our knowledge, this is the first non-trivial result on floating-point (interval-) completeness. This result is an extension of our IUA theorem (Theorem 1) for the identity activation function  $\sigma_{\text{id}}$ , in that floating-point interval-completeness considers the input domain  $\mathbb{F}^n$  (not  $[-1, 1]^n \cap \mathbb{F}^n$ ) and  $\mathcal{F}_{\oplus, \otimes}$  includes all  $\sigma_{\text{id}}$ -networks (but no other  $\sigma$ -networks). Theorem 3, however, *cannot* be extended to the input domain  $(\mathbb{F} \cup \{-\infty, +\infty\})^n$  (instead of  $\mathbb{F}^n$ ), since no program in  $\mathcal{F}_{\oplus, \otimes}$  can represent a non-constant function that maps an infinite float to a finite float—this is because  $\oplus$  and  $\otimes$  do not return finite floats when applied to  $\pm\infty$ .

## 5 Proof of IUA Theorem Over Floats

We now prove Theorem 1 by constructing a  $\sigma$ -neural network that computes the upper and lower points of the direct image map of a rounded target function  $\hat{f}$ . For  $a, b \in \mathbb{R}$ , we let  $[a, b]_{\mathbb{F}} := [a, b] \cap \mathbb{F}$  and  $\mathbb{I}_{[a, b]} := \{\mathcal{I} \in \mathbb{I} \mid \gamma(\mathcal{I}) \subseteq [a, b]\}$ . With this notation,  $(\mathbb{I}_{[a, b]})^d$  is the set of all abstract boxes in  $[a, b]^d$ .

We start with defining indicator functions for a set of floating-point values and for an abstract box, which play a key role in our proof.

**Definition 5.** Let  $d \in \mathbb{N}$ . For  $\mathcal{S} \subseteq \overline{\mathbb{F}}^d$ , we define  $\iota_{\mathcal{S}} : \overline{\mathbb{F}}^d \rightarrow \overline{\mathbb{F}}$  as  $\iota_{\mathcal{S}}(\mathbf{x}) := 1$  if  $\mathbf{x} \in \mathcal{S}$ , and  $\iota_{\mathcal{S}}(\mathbf{x}) := 0$  otherwise. For  $a \in \mathbb{F}$ , we define  $\iota_{>a} : \overline{\mathbb{F}} \rightarrow \overline{\mathbb{F}}$  by  $\iota_{\{x > a \mid x \in \mathbb{F}\}}$ , and define  $\iota_{\geq a}$ ,  $\iota_{< a}$ ,  $\iota_{\leq a}$  analogously. For  $\mathcal{C} \in \mathbb{I}^d$ , we define  $\iota_{\mathcal{C}} : \overline{\mathbb{F}}^d \rightarrow \overline{\mathbb{F}}$  by  $\iota_{\gamma(\mathcal{C})}$ .

Our proof of Theorem 1 consists of two parts. We first show the existence of  $\sigma$ -networks that precisely compute indicator functions under the interval semantics. We then construct a  $\sigma$ -network stated in Theorem 1 by composing the  $\sigma$ -networks for indicator functions and using the properties of indicator functions.

Both parts of our proof are centered around a new property of activation functions, which we call “ $([a, b]_{\mathbb{F}}, \eta, K, L_{\phi}, L_{\psi})$ -separability” and define as follows.

**Definition 6.** We say that  $\sigma : \overline{\mathbb{F}} \rightarrow \overline{\mathbb{F}}$  is  $([a, b]_{\mathbb{F}}, \eta, K, L_{\phi}, L_{\psi})$ -separable for  $a, b, \eta, K \in \mathbb{F}$  and  $L_{\phi}, L_{\psi} \in \mathbb{N}$  if the following hold:

- For every  $z \in [a, b]_{\mathbb{F}}$ , there exist depth- $L_{\phi}$   $\sigma$ -networks  $\phi_{<z}, \phi_{\geq z} : \overline{\mathbb{F}} \rightarrow \overline{\mathbb{F}}$  without the last affine layer such that  $\phi_{\leq z}^{\sharp} = (K\iota_{\leq z})^{\sharp}$  and  $\phi_{\geq z}^{\sharp} = (K\iota_{\geq z})^{\sharp}$  on  $\mathbb{I}_{[a, b]}$ .
- There exists a depth- $L_{\psi}$   $\sigma$ -network  $\psi_{>\eta} : \overline{\mathbb{F}} \rightarrow \overline{\mathbb{F}}$  without the first and last affine layers such that  $\psi_{>\eta}^{\sharp} = (K\iota_{>\eta})^{\sharp}$  on  $\mathbb{I}_{[a, b]}$ .

The first condition in Definition 6 ensures the existence of  $\sigma$ -networks that perfectly implement scaled indicator functions  $K\iota_{\leq z}$  and  $K\iota_{\geq z}$  under the interval semantics, for all  $z \in [a, b]_{\mathbb{F}}$ . Since these networks should have the same depth  $L_\phi$  without the last affine layer, a function  $\nu : \overline{\mathbb{F}}^n \rightarrow \overline{\mathbb{F}}$  defined, e.g., by

$$\nu(x_1, \dots, x_n) = \left( \sum_{i=1}^n \alpha \otimes \phi_{\leq z_i}(x_i) \right) \oplus \beta \quad (16)$$

denotes a depth- $L_\phi$   $\sigma$ -network for any  $z_i \in [a, b]_{\mathbb{F}}$  and  $\alpha, \beta \in \mathbb{F}$ . The second condition in Definition 6 guarantees that another scaled indicator function  $K\iota_{>\eta}$  can be precisely implemented by a depth- $L_\psi$   $\sigma$ -network  $\psi_{>\eta}$  without the first and last affine layers. This implies, e.g., that  $\psi_{>\eta} \circ \nu$  denotes a depth- $(L_\phi + L_\psi - 1)$   $\sigma$ -network, where  $\nu$  denotes the network presented in Eq. (16).

Using the separability property, we can formally state the two parts of our proof as Lemmas 2 and 3. Theorem 1 is a direct corollary of the two lemmas. We present the proofs of Lemmas 2 and 3 in the next subsections (Sect. 5.1 and 5.2).

**Lemma 2.** *Suppose that  $\sigma : \overline{\mathbb{F}} \rightarrow \overline{\mathbb{F}}$  satisfies Condition 1 with constants  $c_2, \eta \in \mathbb{F}$ . Then,  $\sigma$  is  $([-1, 1]_{\mathbb{F}}, \eta, K, L_\phi, L_\psi)$ -separable for some  $L_\phi, L_\psi \in \mathbb{N}$ , where  $\eta$  and  $K := \sigma(c_2)$  satisfy  $|\eta| \in [2^{\epsilon_{\min}+5}, 4 - 8\epsilon]$  and  $|K| \in [\frac{\epsilon}{2} + 2\epsilon^2, \frac{5}{4} - 2\epsilon]$ .*

**Lemma 3.** *Suppose that  $\sigma : \overline{\mathbb{F}} \rightarrow \overline{\mathbb{F}}$  is  $([a, b]_{\mathbb{F}}, \eta, K, L_\phi, L_\psi)$ -separable for some  $a, b, \eta, K \in \mathbb{F}$  and  $L_\phi, L_\psi \in \mathbb{N}$  with  $|\eta| \in [2^{\epsilon_{\min}+5}, 4 - 8\epsilon]$  and  $|K| \in [\frac{\epsilon}{2} + 2\epsilon^2, \frac{5}{4} - 2\epsilon]$ . Then, for every  $d \in \mathbb{N}$  and function  $h : \overline{\mathbb{F}}^d \rightarrow \overline{\mathbb{F}} \setminus \{\perp\}$ , there exists a  $\sigma$ -neural network  $\nu : \overline{\mathbb{F}}^d \rightarrow \overline{\mathbb{F}}$  such that  $\nu^\sharp(\mathcal{B}) = h^\sharp(\mathcal{B})$  for all abstract boxes  $\mathcal{B}$  in  $[a, b]^d$ .*

To prove Lemma 2, we construct a  $\sigma$ -network for the scaled indicator function  $K\iota_{\geq z}$  in two steps. We first construct a  $\sigma$ -network that maps all inputs smaller than  $z$  to some point  $x_1$ , and all other inputs to another point  $x_2 \neq x_1$  (Lemma 4 and 6), where we exploit round-off errors to obtain such “contraction” (Lemma 17 in §F). We then map  $x_1$  to  $c_1$  and  $x_2$  to  $c_2$ , and apply  $\sigma$  to the result so that the final network maps all inputs smaller than  $z$  to  $\sigma(c_1) = 0$  and all other inputs to  $\sigma(c_2) = K$  (Lemma 5). We construct  $\sigma$ -networks for  $K\iota_{\leq z}$  and  $K\iota_{>\eta}$  analogously.

To prove Lemma 3, we construct  $\sigma$ -networks for the scaled indicator functions of every box in  $([a, b]_{\mathbb{F}})^d$  (Lemma 7) and every subset of  $([a, b]_{\mathbb{F}})^d$  (Lemma 8), using the indicator functions constructed in Lemma 2. We construct the final  $\sigma$ -network (i.e., universal interval approximator) as a floating-point linear combination of the  $\sigma$ -networks that represent the scaled indicator functions of the level sets of the target function (Lemma 9).

## 5.1 Proof of Lemma 2

To prove Lemma 2, we assume that the activation function  $\sigma : \overline{\mathbb{F}} \rightarrow \overline{\mathbb{F}}$  satisfies Condition 1 with some constants  $c_1, c_2, \eta \in \mathbb{F}$ . By Condition 1, the constants

$\eta$  and  $K := \sigma(c_2)$  clearly satisfy the range condition in Lemma 2. Hence, it remains to show the  $([-1, 1]_{\mathbb{F}}, \eta, K, L_{\phi}, L_{\psi})$ -separability of  $\sigma$  for some  $L_{\phi}, L_{\psi} \in \mathbb{N}$ . This requires us to construct  $\sigma$ -networks  $\psi_{>\eta}$  and  $\phi_{\leq z}, \phi_{\geq z}$  for every  $z \in [-1, 1]_{\mathbb{F}}$  such that  $\psi_{>\eta}^{\sharp} = (K\iota_{>\eta})^{\sharp}$ ,  $\phi_{\leq z}^{\sharp} = (K\iota_{\leq z})^{\sharp}$ , and  $\phi_{\geq z}^{\sharp} = (K\iota_{\geq z})^{\sharp}$  on  $\mathbb{I}_{[-1, 1]}$  (Definition 6).

We first construct  $\psi_{>\eta}$  using Lemmas 4 and 5 (Fig. 4). The proofs of these lemmas, presented in §D.1 and §D.2, rely heavily on (C1)–(C3) of Condition 1.

**Lemma 4.** *There exists a  $\sigma$ -network  $\mu : \overline{\mathbb{F}} \rightarrow \overline{\mathbb{F}}$  without the first affine layer such that  $\mu^{\sharp}(\langle -\Omega, \eta \rangle) = \langle \eta, \eta \rangle$ ,  $\mu^{\sharp}(\langle \eta^+, \Omega \rangle) = \langle \eta^+, \eta^+ \rangle$ , and  $\mu^{\sharp}(\langle -\Omega, \Omega \rangle) = \langle \eta, \eta^+ \rangle$ .*

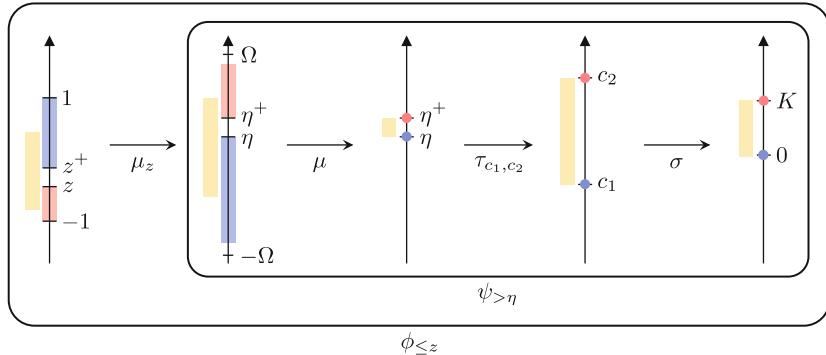
**Lemma 5.** *Let  $(\theta, \theta')$  be either  $(c_1, c_2)$  or  $(c_2, c_1)$ . Then, there exists a depth-2  $\sigma$ -network  $\tau_{\theta, \theta'} : \overline{\mathbb{F}} \rightarrow \overline{\mathbb{F}}$  without the first affine layer such that  $\tau_{\theta, \theta'}^{\sharp}(\langle \eta, \eta \rangle) = \langle \theta, \theta \rangle$ ,  $\tau_{\theta, \theta'}^{\sharp}(\langle \eta^+, \eta^+ \rangle) = \langle \theta', \theta' \rangle$ , and  $\tau_{\theta, \theta'}^{\sharp}(\langle \eta, \eta^+ \rangle) = \langle \min\{\theta, \theta'\}, \max\{\theta, \theta'\} \rangle$ .*

Lemma 4 states that we can construct a  $\sigma$ -network  $\mu$  without the first affine layer, whose interval semantics maps all finite (abstract) intervals left of  $\eta$  to the singleton interval  $\langle \eta, \eta \rangle$ , all finite intervals right of  $\eta^+$  to  $\langle \eta^+, \eta^+ \rangle$ , and all the remaining finite intervals to  $\langle \eta, \eta^+ \rangle$ . Similarly, Lemma 5 shows that there exists a  $\sigma$ -network  $\tau_{\theta, \theta'}$  without the first affine layer, whose interval semantics maps  $\langle \eta, \eta \rangle$  to  $\langle \theta, \theta \rangle$ ,  $\langle \eta^+, \eta^+ \rangle$  to  $\langle \theta', \theta' \rangle$ , and  $\langle \eta, \eta^+ \rangle$  to the interval between  $\theta$  and  $\theta'$ . By composing these networks with  $\sigma$ , we construct  $\psi_{>\eta}$  as

$$\psi_{>\eta} := \sigma \circ \tau_{c_1, c_2} \circ \mu. \quad (17)$$

This function  $\psi_{>\eta}$  is a  $\sigma$ -network without the first and last affine layers, since  $\tau_{c_1, c_2}$  and  $\mu$  are without the first affine layer. Moreover,  $\psi_{>\eta}^{\sharp} = (K\iota_{>\eta})^{\sharp}$  on  $\mathbb{I}_{[-1, 1]}$  by the aforementioned properties of  $\tau_{c_1, c_2}$  and  $\mu$ , and by the next properties of  $\sigma$  from Condition 1 of Condition 1:  $\sigma(c_1) = 0$ ,  $\sigma(c_2) = K$ , and  $\sigma(x)$  lies between them for all  $x$  between  $c_1$  and  $c_2$ . Lastly, we choose  $L_{\psi}$  as the depth of  $\psi_{>\eta}$ .

We next construct  $\phi_{\leq z}$  and  $\phi_{\geq z}$  using Lemma 6 (Fig. 4). The proof of this lemma is provided in §D.3.



**Fig. 4.** Illustration of networks  $\mu$ ,  $\tau_{c_1, c_2}$ ,  $\mu_z$  (Lemmas 4–6) and  $\psi_{>\eta}$ ,  $\phi_{\leq z}$  (Eqs. (17) and (18)), assuming (b) in Lemma 6. A box/dot denotes an abstract interval.

**Lemma 6.** Let  $z \in \mathbb{F}$  with  $|z| \leq 1^+$ . Then, there exists a depth-1  $\sigma$ -network  $\mu_z : \overline{\mathbb{F}} \rightarrow \overline{\mathbb{F}}$  such that one of the following holds.

- (a)  $\gamma(\mu_z^\sharp(\langle -1, z \rangle)) \subset [-\Omega, \eta]$  and  $\gamma(\mu_z^\sharp(\langle z^+, 1 \rangle)) \subset [\eta^+, \Omega]$ .
- (b)  $\gamma(\mu_z^\sharp(\langle -1, z \rangle)) \subset [\eta^+, \Omega]$  and  $\gamma(\mu_z^\sharp(\langle z^+, 1 \rangle)) \subset [-\Omega, \eta]$ .

Lemma 6 ensures the existence of a depth-1  $\sigma$ -network  $\mu_z$ , whose interval semantics maps  $\langle -1, z \rangle$  and  $\langle z^+, 1 \rangle$  to an interval left of  $\eta$  and an interval right of  $\eta^+$ . By composing  $\mu_z$  with the previous networks  $\tau_{\theta, \theta'}$  and  $\mu$ , we construct  $\phi_{\leq z}$  as

$$\phi_{\leq z} := \begin{cases} \sigma \circ \tau_{c_2, c_1} \circ \mu \circ \mu_z & \text{if (a) holds in Lemma 6} \\ \sigma \circ \tau_{c_1, c_2} \circ \mu \circ \mu_z & \text{if (b) holds in Lemma 6.} \end{cases} \quad (18)$$

By a similar argument used above, the function  $\phi_{\leq z}$  is a  $\sigma$ -network without the last affine layer, and it satisfies the desired equation:  $\phi_{\leq z}^\sharp = (K\iota_{\leq z})^\sharp$  on  $\mathbb{I}_{[-1, 1]}$ . We construct  $\phi_{\geq z}$  analogously, but using  $\mu_{z^-}$  instead of  $\mu_z$ . Since the depths of  $\phi_{\leq z}$  and  $\phi_{\geq z}$  are identical for all  $z$ , we denote this depth by  $L_\phi$ . This completes the construction of  $\psi_{>\eta}$ ,  $\phi_{\leq z}$ , and  $\phi_{\geq z}$ , finishing the proof of Lemma 2.

## 5.2 Proof of Lemma 3

To prove Lemma 3, we assume that the activation function  $\sigma$  is  $([a, b]_{\mathbb{F}}, \eta, K, L_\phi, L_\psi)$ -separable for some  $\eta, K \in \mathbb{F}$  with  $|\eta| \in [2^{\epsilon_{\min}+5}, 4 - 8\epsilon]$  and  $|K| \in [\frac{\epsilon}{2} + 2\epsilon^2, \frac{5}{4} - 2\epsilon]$ . Given this, we construct a  $\sigma$ -network whose interval semantics exactly computes that of the target function  $h : \overline{\mathbb{F}}^d \rightarrow \overline{\mathbb{F}} \setminus \{\perp\}$  for all abstract boxes in  $[a, b]^d$ . In our construction, we progressively implement the following functions using  $\sigma$ -networks: (i) scaled indicator functions of arbitrary boxes, (ii) scaled indicator functions of arbitrary sets, and (iii) the target function.

We first construct a  $\sigma$ -network  $\tilde{\nu}_{\mathcal{B}}$ , for any abstract box  $\mathcal{B}$  in  $[a, b]^d$ , that implements the scaled indicator function  $K\iota_{\mathcal{B}}$  under the interval semantics.

**Lemma 7.** For any  $\mathcal{B} \in (\mathbb{I}_{[a, b]})^d$ , there exists a depth- $L$   $\sigma$ -network  $\tilde{\nu}_{\mathcal{B}} : \overline{\mathbb{F}}^d \rightarrow \overline{\mathbb{F}}$  without the last affine layer such that  $\tilde{\nu}_{\mathcal{B}}^\sharp = (K\iota_{\mathcal{B}})^\sharp$  on  $(\mathbb{I}_{[a, b]})^d$ , where  $L := L_\phi + (L_\psi - 1)(\lceil \log_{2^M} d \rceil + 1)$ .

In the proof of Lemma 7, we design  $\tilde{\nu}_{\mathcal{B}}$  using the networks  $\psi_{>\eta}$ ,  $\phi_{\leq z}$ , and  $\phi_{\geq z}$  constructed in Sect. 5.1. Specifically, for an abstract box  $\mathcal{B} = (\langle a_1, b_1 \rangle, \dots, \langle a_d, b_d \rangle)$ , we define a  $\sigma$ -network  $\tilde{\nu}_i : \overline{\mathbb{F}} \rightarrow \overline{\mathbb{F}}$  as

$$\tilde{\nu}_i(x) := \psi_{>\eta}((\alpha \otimes \phi_{\geq a_i}(x)) \oplus (\alpha \otimes \phi_{\leq b_i}(x)) \oplus \beta), \quad (19)$$

where  $\alpha, \beta \in \mathbb{F}$  are constants such that  $\beta \leq \eta$ ,  $(\alpha \otimes K) \oplus \beta \leq \eta$ , and  $(\alpha \otimes K) \oplus (\alpha \otimes K) \oplus \beta > \eta$ . Then, we can show that  $\tilde{\nu}_i^\sharp = (K\iota_{\langle a_i, b_i \rangle})^\sharp$  on  $\mathbb{I}_{[a, b]}$ . When  $d$  is small (e.g.,  $d \leq 2^{M+1}$ ), we construct  $\tilde{\nu}_{\mathcal{B}}$  using  $\tilde{\nu}_i$  and  $\psi_{>\eta}$ , as follows:

$$\tilde{\nu}_{\mathcal{B}}(x_1, \dots, x_d) := \psi_{>\eta} \left( \left( \sum_{i=1}^d \alpha' \otimes \tilde{\nu}_i(x_i) \right) \oplus \beta' \right), \quad (20)$$

where  $\alpha', \beta' \in \mathbb{F}$  are suitably chosen so that  $\tilde{\nu}_{\mathcal{B}}^\sharp = (K\iota_{\mathcal{B}})^\sharp$  on  $(\mathbb{I}_{[a,b]})^d$ . When  $d$  is large (e.g.,  $d > 2^{M+1}$ ), this construction does not work since  $\sum_{i=1}^d \alpha' \otimes \tilde{\nu}_i(x_i)$  may not be computed as we want due to rounding errors (e.g.,  $\sum_{i=1}^n 1 = 2^{M+1} < n$  for all  $n > 2^{M+1}$ ). In such a case, we construct  $\tilde{\nu}_{\mathcal{B}}$  hierarchically using more layers, but based on a similar idea. A rigorous proof of Lemma 7, including the proof that appropriate  $\alpha, \alpha', \beta, \beta' \in \mathbb{F}$  exist, is presented in §E.1

Using  $\tilde{\nu}_{\mathcal{B}}$ , we next construct a  $\sigma$ -network  $\tilde{\nu}_{\mathcal{S}}$ , for any set  $\mathcal{S}$  in  $([a, b]_{\mathbb{F}})^d$ , whose interval semantics computes that of the scaled indicator function  $K\iota_{\mathcal{S}}$ .

**Lemma 8.** *Suppose that for any  $\mathcal{B} \in (\mathbb{I}_{[a,b]})^d$ , there exists a depth- $L$   $\sigma$ -network  $\tilde{\nu}_{\mathcal{B}}$  without the last affine layer such that  $\tilde{\nu}_{\mathcal{B}}^\sharp = (K\iota_{\mathcal{B}})^\sharp$  on  $(\mathbb{I}_{[a,b]})^d$ . Then, for any  $\mathcal{S} \subseteq ([a, b]_{\mathbb{F}})^d$ , there exists a depth- $(L + L_{\psi} - 1)$   $\sigma$ -network  $\tilde{\nu}_{\mathcal{S}} : \mathbb{F}^d \rightarrow \mathbb{F}$  without the last affine layer such that  $\tilde{\nu}_{\mathcal{S}}^\sharp = (K\iota_{\mathcal{S}})^\sharp$  on  $(\mathbb{I}_{[a,b]})^d$ .*

In the proof of Lemma 8, we construct  $\tilde{\nu}_{\mathcal{S}}$  using  $\tilde{\nu}_{\mathcal{B}}$  and  $\psi_{>\eta}$ , as follows:

$$\tilde{\nu}_{\mathcal{S}}(\mathbf{x}) := \psi_{>\eta} \left( \left( \sum_{\mathcal{B} \in \mathcal{T}} \alpha'' \otimes \tilde{\nu}_{\mathcal{B}}(\mathbf{x}) \right) \oplus \eta \right), \quad (21)$$

where  $\mathcal{T}$  denotes the collection of all abstract boxes in  $\mathcal{S}$ , and  $\alpha'' \in \mathbb{F}$  is a constant such that  $\eta < (\sum_{i=1}^n \alpha'' \otimes K) \oplus \eta < \infty$  for all  $n \geq 1$ . We remark that it is possible to make the summation not overflow even for a large  $n$ , by cleverly exploiting the rounding errors from  $\oplus$ . With a proper choice of  $\alpha''$ , we can further show that  $\tilde{\nu}_{\mathcal{S}}^\sharp = (K\iota_{\mathcal{S}})^\sharp$  on  $(\mathbb{I}_{[a,b]})^d$ . A formal proof of Lemma 8 is given in §E.2.

Using  $\tilde{\nu}_{\mathcal{S}}$ , we finally construct a  $\sigma$ -network that coincides, under the interval semantics, with the target function  $h$  over  $([a, b]_{\mathbb{F}})^d$ . This result (Lemma 9) and the above results (Lemma 7 and 8) directly imply Lemma 3.

**Lemma 9.** *Assume that for any  $\mathcal{S} \subseteq ([a, b]_{\mathbb{F}})^d$ , there exists a depth- $L'$   $\sigma$ -network  $\tilde{\nu}_{\mathcal{S}}$  without the last affine layer such that  $\tilde{\nu}_{\mathcal{S}}^\sharp = (K\iota_{\mathcal{S}})^\sharp$  on  $(\mathbb{I}_{[a,b]})^d$ . Then, for any  $h : \mathbb{F}^d \rightarrow \mathbb{F} \setminus \{\perp\}$ , there exists a  $\sigma$ -network  $\nu : \mathbb{F}^d \rightarrow \mathbb{F}$  such that  $\nu^\sharp = h^\sharp$  on  $(\mathbb{I}_{[a,b]})^d$ .*

We now illustrate the main idea of the proof of Lemma 9. For a simpler argument, we assume that  $h$  is non-negative; the proof for the general case is similar (see §E.3). Let  $0 = z_0 < z_1 < \dots < z_n = +\infty$  be all non-negative floats (except  $\perp$ ) in increasing order, and let  $\mathcal{S}_i := \{\mathbf{x} \in ([a, b]_{\mathbb{F}})^d \mid h(\mathbf{x}) \geq z_i\}$  be the level set of  $h$  for  $z_i$ . Under this setup, we construct  $\nu$  using  $\tilde{\nu}_{\mathcal{S}_i}$ , as follows:

$$\nu(\mathbf{x}) := \sum_{i=1}^m \alpha_i \otimes \tilde{\nu}_{\mathcal{S}_i}(\mathbf{x}), \quad (22)$$

where  $m \in \mathbb{N} \cup \{0\}$  and  $\alpha_i \in \mathbb{F}$  are chosen so that  $z_m = \max\{h(\mathbf{x}) \mid \mathbf{x} \in ([a, b]_{\mathbb{F}})^d\}$  and  $\alpha_i \otimes K \approx z_i - z_{i-1}$  for all  $i \in [m]$ . If  $\alpha_i \otimes K$  is close enough to  $z_i - z_{i-1}$ , then the floating-point summation  $\sum_{i=1}^k \alpha_i \otimes K$  is exactly equal to the exact summation  $\sum_{i=1}^k z_i - z_{i-1} = z_k$  for all  $k \in [m]$ , by the rounding errors of  $\oplus$ . Using this observation, we can show that  $\nu(\mathbf{x}) = h(\mathbf{x})$  for all  $\mathbf{x} \in ([a, b]_{\mathbb{F}})^d$ , and more importantly,  $\nu^\sharp = h^\sharp$  on  $(\mathbb{I}_{[a,b]})^d$ . The full proof of Lemma 9 is in §E.3.

## 6 Related Work

**Universal Approximation.** Universal approximation theorems for neural networks are widely studied in the literature, which include results for feedforward networks [12, 27, 28, 55], convolutional networks [71], residual networks [41], and transformers [70]. With the advent of low-precision computing for neural networks (e.g., 8-bit E5M2, 8-bit E4M3 [46, 65]; float16 [45]; bfloat16 [1]), there has been growing interest among researchers in characterizing their expressiveness power in this setting. New UA theorems for “quantized” neural networks, which use finite-precision network parameters with *exact* real arithmetic, have been studied in [15, 21]. These networks differ from the floating-point networks considered in this work, because our networks use *inexact* floating-point arithmetic.

To the best of current knowledge, [30, 53] are the only works that study UA theorems for floating-point neural networks. [53] proves UA theorems for ReLU and step activation functions. Our IUA Theorem 1, by virtue of Eq. (15), is a strict generalization of [53] in two senses: (i) it applies to a much broader class of activations that satisfy Condition 1, which subsumes ReLU and step functions; and (ii) it provides a result for abstract interpretation via interval analysis, of which the pointwise approximation considered in [53] is a special case. Concurrent with this article, [30] generalizes [53] to support a wider range of activation functions and larger input domains. Our Theorem 1 partially subsumes [30] in that it is a result for interval approximation, whereas [30] considers only pointwise approximation. Conversely, a special case of our Theorem 1 for pointwise approximation (i.e., Eq. (15)) is subsumed by [30] in that it applies to smaller classes of activation functions and input domains.

**Interval Universal Approximation.** The first work to establish an IUA theorem for neural networks used interval analysis with the ReLU activation [4], which was later extended to the more general class of so-called “squashable” activation functions [66]. Whereas these previous IUA theorems assume the neural network can compute over arbitrary real numbers with infinitely precise real arithmetic, the IUA result (Theorem 1) in this work applies to “machine-implementable” neural networks that use floating-point numbers and operations. To the best of our knowledge, no previous work has established an IUA theorem for floating-point neural networks. These different computational models lead to substantial differences in both the proof methods (cf. Sect. 3.2 and 5) and the specific technical results—Sect. 3.3 gives a detailed discussion of how Theorem 1 differs from previous IUA and robustness results [4, Theorem 1.1]; [66, Theorem 3.7].

**Provable Robustness.** There is an extensive literature on robustness verification and robust training for neural networks, which is surveyed in, e.g., [4, Chapter 1]; [38, 60]. Notable methods among these works are [61, 62], which verify the robustness of a neural network using abstract interpretation with the zonotope and polyhedra domains for a restricted class of activations, and are sound

with respect to floating-point arithmetic. Compared to these methods, our contribution is a theoretical result on the inherent expressiveness of provably robust floating-point networks under the interval domain for a broad class of activation functions, rather than new verification algorithms or abstract domains. Indeed, our existence result directly applies to the zonotope and polyhedra domains, as they are more precise than the interval domain. More specific IUA theorems tailored to these domains may yield more compact constructions that witness the existence of a provably robust floating-point neural network. Recently, [32] shows that even if a neural network is provably robust over real arithmetic, it can be non-robust over floating-point arithmetic and remain vulnerable to adversarial attacks. This highlights the importance of establishing robustness in the floating-point setting.

**Acknowledgments.** G. Hwang and Y. Park were supported by Korea Institute for Advanced Study (KIAS) Individual Grants AP092801 and AP090301, via the Center for AI and Natural Sciences at KIAS. G. Hwang was also supported by National Research Foundation of Korea (NRF) Grants RS-2025-00515264 and RS-2024-00406127, funded by the Korea Ministry of Science and ICT (MSIT); and the Gwangju Institute of Science and Technology (GIST) Global University Project in 2025. Y. Park was also supported by the Sejong University faculty research fund in 2025. S. Park was supported by the Korea Institute of Information & Communications Technology Planning & Evaluation (IITP) Grant RS-2019-II190079, funded by the Korea MSIT; the Information Technology Research Center (IITP-ITRC) Grant IITP-2025-RS-2024-00436857, funded by the Korea MSIT; and the Culture, Sports, and Tourism R&D Program through the Korea Creative Content Agency (KOCCA) Grants RS-2024-00348469 and RS-2024-00345025, funded by the Korea Ministry of Culture, Sports and Tourism (MCST) in 2024. W. Lee and F. Saad were supported by the United States National Science Foundation (NSF) under Grant No. 2311983 and funds from the Computer Science Department at Carnegie Mellon University. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the funding agencies.

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

## References

1. Abadi, M., et al.: TensorFlow: large-scale machine learning on heterogeneous distributed systems. arXiv 1603.04467 (2016). <https://doi.org/10.48550/arXiv.1603.04467>
2. Albargouthi, A.: Introduction to neural network verification (2021). <https://doi.org/10.48550/arXiv.2109.10317>
3. Baader, M.: Expressivity of certified neural networks. Ph.D. thesis, ETH Zurich, Zurich, Switzerland (2024). <https://doi.org/10.3929/ETHZ-B-000677199>

4. Baader, M., Mirman, M., Vechev, M.T.: Universal approximation with certified networks. In: Proceedings of the International Conference on Learning Representations (2020). <https://doi.org/10.48550/arXiv.1909.13846>
5. Baader, M., Müller, M.N., Mao, Y., Vechev, M.T.: Expressivity of ReLU-networks under convex relaxations. In: Proceedings of the International Conference on Learning Representations (2024). <https://openreview.net/forum?id=awHTL3Hpto>
6. Barak, B.: Introduction to Theoretical Computer Science (2023). <https://introtcs.org>
7. Beebe, N.: The Mathematical-Function Computation Handbook: Programming Using the MathCW Portable Software Library. Springer (2017). <https://doi.org/10.1007/978-3-319-64110-2>
8. Church, A.: A set of postulates for the foundation of logic. *Ann. Math.* **34**(4), 839–864 (1933). <https://doi.org/10.2307/1968702>
9. Church, A.: An unsolvable problem of elementary number theory. *Am. J. Math.* **58**(2), 345–363 (1936). <https://doi.org/10.2307/2371045>
10. Cousot, P., Cousot, R.: Abstract interpretation: a unified lattice model for static analysis of programs by construction or approximation of fixpoints. In: Proceedings of the ACM Symposium on Principles of Programming Languages, pp. 238–252 (1977). <https://doi.org/10.1145/512950.512973>
11. Cousot, P., Cousot, R.: Static determination of dynamic properties of generalized type unions. In: Proceedings of an ACM Conference on Language Design for Reliable Software, pp. 77–94 (1977). <https://doi.org/10.1145/800022.808314>
12. Cybenko, G.: Approximation by superpositions of a sigmoidal function. *Math. Control Sig. Syst.* **2**(4), 303–314 (1989). <https://doi.org/10.1007/BF02551274>
13. Daramy-Loirat, C., et al.: CR-LIBM: a library of correctly rounded elementary functions in double-precision. Research Report ENSL-01529804, Laboratoire de l’Informatique du Parallelisme, December 2006. <https://ens-lyon.hal.science/ensl-01529804>
14. Dettmers, T., Pagnoni, A., Holtzman, A., Zettlemoyer, L.: QLoRA: efficient fine-tuning of quantized LLMs. In: Proceedings of the International Conference on Neural Information Processing Systems (2023). <https://doi.org/10.5555/3666122.3666563>
15. Ding, Y., Liu, J., Xiong, J., Shi, Y.: On the universal approximability and complexity bounds of quantized ReLU neural networks. In: International Conference on Learning Representations (2019). <https://doi.org/10.48550/arXiv.1802.03646>
16. Eykholt, K., et al.: Robust physical-world attacks on deep learning visual classification. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 1625–1634 (2018). <https://doi.org/10.1109/CVPR.2018.00175>
17. Faïssole, F., de Lamarlière, P.G., Melquiond, G.: End-to-end formal verification of a fast and accurate floating-point approximation. In: Proceedings of the International Conference on Interactive Theorem Proving, pp. 14:1–14:18 (2024). <https://doi.org/10.4230/LIPIcs.ITP.2024.14>
18. Finlayson, S.G., Bowers, J.D., Ito, J., Zittrain, J.L., Beam, A.L., Kohane, I.S.: Adversarial attacks on medical machine learning. *Science* **363**(6433), 1287–1289 (2019). <https://doi.org/10.1126/science.aaw4399>
19. Gehr, T., Mirman, M., Drachsler-Cohen, D., Tsankov, P., Chaudhuri, S., Vechev, M.T.: AI2: safety and robustness certification of neural networks with abstract interpretation. In: Proceedings of the IEEE Symposium on Security and Privacy, pp. 3–18 (2018). <https://doi.org/10.1109/SP.2018.00058>

20. Gödel, K.: On undecidable propositions of formal mathematics systems (Notes by Kleene, S.C., Rosser, J.B.) Institute for Advanced Study (1934). <https://albert.ias.edu/20.500.12111/7996>
21. Gonon, A., Brisebarre, N., Gribonval, R., Riccietti, E.: Approximation speed of quantized versus unquantized ReLU neural networks and beyond. *IEEE Trans. Inf. Theory* **69**(6), 3960–3977 (2023). <https://doi.org/10.1109/TIT.2023.3240360>
22. Goodfellow, I.J., Bengio, Y., Courville, A.C.: Deep Learning. MIT Press (2016). <http://www.deeplearningbook.org/>
23. Goodfellow, I.J., Shlens, J., Szegedy, C.: Explaining and harnessing adversarial examples. In: Proceedings of the International Conference on Learning Representations (2015). <https://doi.org/10.48550/arXiv.1412.6572>
24. Gowal, S., et al.: Scalable verified training for provably robust image classification. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 4841–4850 (2019). <https://doi.org/10.1109/ICCV.2019.00494>
25. Grigore, R.: Java generics are Turing complete. In: Proceedings of the ACM Symposium on Principles of Programming Languages, pp. 73–85 (2017). <https://doi.org/10.1145/3093333.3009871>
26. Hickey, T.J., Ju, Q., van Emden, M.H.: Interval arithmetic: from principles to implementation. *J. ACM* **48**(5), 1038–1068 (2001). <https://doi.org/10.1145/502102.502106>
27. Hornik, K., Stinchcombe, M., White, H.: Multilayer feedforward networks are universal approximators. *Neural Netw.* **2**(5), 359–366 (1989). [https://doi.org/10.1016/0893-6080\(89\)90020-8](https://doi.org/10.1016/0893-6080(89)90020-8)
28. Hornik, K., Stinchcombe, M., White, H.: Universal approximation of an unknown mapping and its derivatives using multilayer feedforward networks. *Neural Netw.* **3**(5), 551–560 (1990). [https://doi.org/10.1016/0893-6080\(90\)90005-6](https://doi.org/10.1016/0893-6080(90)90005-6)
29. Hubara, I., Courbariaux, M., Soudry, D., El-Yaniv, R., Bengio, Y.: Quantized neural networks: training neural networks with low precision weights and activations. *J. Mach. Learn. Res.* **18**(1), 6869–6898 (2017). <https://doi.org/10.5555/3122009.3242044>
30. Hwang, G., Park, Y., Lee, W., Park, S.: Floating-point neural networks can represent almost all floating-point functions. In: Proceedings of the International Conference on Machine Learning (2025)
31. Institute of Electrical and Electronics Engineers: IEEE Standard for Floating-Point Arithmetic (IEEE Std 754-2019). IEEE, Piscataway, NJ, USA (2019). <https://doi.org/10.1109/IEEEESTD.2019.8766229>
32. Jin, J., Ohrimenko, O., Rubinstein, B.I.P.: Getting a-round guarantees: floating-point attacks on certified robustness. In: Proceedings of the Workshop on Artificial Intelligence and Security, pp. 53–64 (2024). <https://doi.org/10.1145/3689932.3694761>
33. Jovanovic, N., Balunovic, M., Baader, M., Vechev, M.T.: On the paradox of certified training. *Trans. Mach. Learn. Res.* (2022). <https://doi.org/10.48550/arXiv.2102.06700>
34. Katz, G., Barrett, C., Dill, D.L., Julian, K., Kochenderfer, M.J.: Reluplex: an efficient SMT solver for verifying deep neural networks. In: Majumdar, R., Kunčak, V. (eds.) CAV 2017. LNCS, vol. 10426, pp. 97–117. Springer, Cham (2017). [https://doi.org/10.1007/978-3-319-63387-9\\_5](https://doi.org/10.1007/978-3-319-63387-9_5)
35. Kozen, D.: Automata and Computability. Springer (1997). <https://doi.org/10.1007/978-1-4612-1844-9>

36. Lee, W., Sharma, R., Aiken, A.: On automatically proving the correctness of math.h implementations. *Proc. ACM Programm. Lang.* **2**(POPL), 47:1–47:32 (2018). <https://doi.org/10.1145/3158135>
37. Li, L., Xie, T., Li, B.: Leaderboard for “SoK: certified robustness for deep neural networks” (2023). <https://sokcertifiedrobustness.github.io/leaderboard/>. Accessed 19 May 2025
38. Li, L., Xie, T., Li, B.: SoK: certified robustness for deep neural networks. In: *Proceedings of the IEEE Symposium on Security and Privacy*, pp. 1289–1310. IEEE Press (2023). <https://doi.org/10.1109/SP46215.2023.10179303>
39. Lim, J.P., Nagarakatte, S.: High performance correctly rounded math libraries for 32-bit floating point representations. In: *Proceedings of the ACM Conference on Programming Languages Design and Implementation*, pp. 359–374 (2021). <https://doi.org/10.1145/3453483.3454049>
40. Lim, J.P., Nagarakatte, S.: One polynomial approximation to produce correctly rounded results of an elementary function for multiple representations and rounding modes. *Proc. ACM Programm. Lang.* **6**(POPL), 1–28 (2022). <https://doi.org/10.1145/3498664>
41. Lin, H., Jegelka, S.: ResNet with one-neuron hidden layers is a universal approximator. In: *Proceedings of the International Conference on Neural Information Processing Systems* (2018). <https://dl.acm.org/doi/10.5555/3327345.3327515>
42. Liu, C., Arnon, T., Lazarus, C., Strong, C.A., Barrett, C.W., Kochenderfer, M.J.: Algorithms for verifying deep neural networks. *Found. Trends Optim.* **4**(3–4), 244–404 (2021). <https://doi.org/10.1561/2400000035>
43. Mao, Y., Müller, M.N., Fischer, M., Vechev, M.T.: Understanding certified training with interval bound propagation. In: *Proceedings of the International Conference on Learning Representations* (2024). <https://doi.org/10.48550/arXiv.2306.10426>
44. Markstein, P.: IA-64 and Elementary Functions: Speed and Precision. Hewlett-Packard Professional Books, Prentice Hall (2000)
45. Micikevicius, P., et al.: Mixed precision training. In: *International Conference on Learning Representations* (2018). <https://doi.org/10.48550/arXiv.1710.03740>
46. Micikevicius, P., et al.: FP8 formats for deep learning. *arXiv* 2209.05433 (2022). <https://doi.org/10.48550/arXiv.2209.05433>
47. Mirman, M., Baader, M., Vechev, M.T.: The fundamental limits of neural networks for interval certified robustness. *Trans. Mach. Learn. Res.* (2022). <https://openreview.net/forum?id=fsacLLU35V>
48. Mirman, M., Gehr, T., Vechev, M.T.: Differentiable abstract interpretation for provably robust neural networks. In: *Proceedings of the International Conference on Machine Learning* (2018). <http://proceedings.mlr.press/v80/mirman18b.html>
49. Moore, C., Mertens, S.: *The Nature of Computation*. Oxford University Press (2011). <https://doi.org/10.1093/acprof:oso/9780199233212.001.0001>
50. Moore, R.E., Kearfott, R.B., Cloud, M.J.: *Introduction to interval analysis*. Soc. Ind. Appl. Math. (2009). <https://doi.org/10.1137/1.9780898717716>
51. Muller, J.: *Elementary Functions: Algorithms and Implementation*, 3rd edn. Birkhäuser, Boston, MA (2016). <https://doi.org/10.1007/978-1-4899-7983-4>
52. Muller, J.M., et al.: *Handbook of Floating-Point Arithmetic*. Springer (2018). <https://doi.org/10.1007/978-3-319-76526-6>
53. Park, Y., Hwang, G., Lee, W., Park, S.: Expressive power of ReLU and step networks under floating-point operations. *Neural Netw.* (2024). <https://doi.org/10.1016/j.neunet.2024.106297>
54. Pérez, J., Barceló, P., Marinkovic, J.: Attention is Turing-complete. *J. Mach. Learn. Res.* **22**, 75:1–75:35 (2021). <https://dl.acm.org/doi/10.5555/3546258.3546333>

55. Pinkus, A.: Approximation theory of the MLP model in neural networks. *Acta Numerica* **8**, 143–195 (1999). <https://doi.org/10.1017/S0962492900002919>
56. Raghunathan, A., Steinhardt, J., Liang, P.: Semidefinite relaxations for certifying robustness to adversarial examples. In: Proceedings of the International Conference on Neural Information Processing Systems (2018). <https://dl.acm.org/doi/10.5555/3327546.3327746>
57. Rosenberg, I., Shabtai, A., Elovici, Y., Rokach, L.: Adversarial machine learning attacks and defense methods in the cyber security domain. *ACM Comput. Surv.* **54**(5) (2021). <https://doi.org/10.1145/3453158>
58. Sibidanov, A., Zimmermann, P., Glondu, S.: The CORE-MATH project. In: Proceedings of the IEEE Symposium on Computer Arithmetic, pp. 26–34 (2022). <https://doi.org/10.1109/ARITH54963.2022.00014>
59. Siegelmann, H., Sontag, E.: On the computational power of neural nets. *J. Comput. Syst. Sci.* **50**(1), 132–150 (1995). <https://doi.org/10.1006/jcss.1995.1013>
60. Singh, G.: Building trust and safety in artificial intelligence with abstract interpretation. In: Hermenegildo, M.V., Morales, J.F. (eds.) *Static Analysis*, pp. 28–38. Springer, Cham (2023). [https://doi.org/10.1007/978-3-031-44245-2\\_3](https://doi.org/10.1007/978-3-031-44245-2_3)
61. Singh, G., Gehr, T., Mirman, M., Püschel, M., Vechev, M.T.: Fast and effective robustness certification. In: Proceedings of the International Conference on Neural Information Processing Systems (2018). <https://dl.acm.org/doi/10.5555/3327546.3327739>
62. Singh, G., Gehr, T., Püschel, M., Vechev, M.T.: An abstract domain for certifying neural networks. *Proc. ACM Programm. Lang.* **3**(POPL), 41:1–41:30 (2019). <https://doi.org/10.1145/3290354>
63. Szegedy, C., et al.: Intriguing properties of neural networks. In: Proceedings of the International Conference on Learning Representations (2014). <https://doi.org/10.48550/arXiv.1312.6199>
64. Turing, A.M.: On computable numbers, with an application to the Entscheidungsproblem. *Proc. London Math. Soc.* **s2-42**(1), 230–265 (1937). <https://doi.org/10.1112/plms/s2-42.1.230>
65. Wang, N., Choi, J., Brand, D., Chen, C.Y., Gopalakrishnan, K.: Training deep neural networks with 8-bit floating point numbers. In: Proceedings of the International Conference on Neural Information Processing Systems (2018). <https://doi.org/10.5555/3327757.3327866>
66. Wang, Z., Albargouthi, A., Prakriya, G., Jha, S.: Interval universal approximation for neural networks. *Proc. ACM Programm. Lang.* **6**(POPL), 14.1–14.29 (2022). <https://doi.org/10.1145/3498675>
67. Wansbrough, K.: Instance declarations are universal (1998). <http://www.lochan.org/keith/publications/undec.html>
68. Wong, E., Kolter, J.Z.: Provable defenses against adversarial examples via the convex outer adversarial polytope. In: Proceedings of the International Conference on Machine Learning (2018). <https://doi.org/10.48550/arXiv.1711.00851>
69. Yarotsky, D.: Optimal approximation of continuous functions by very deep ReLU networks. In: Proceedings of the Conference on Learning Theory (2018). <https://doi.org/10.48550/arXiv.1802.03620>
70. Yun, C., Bhojanapalli, S., Rawat, A.S., Reddi, S., Kumar, S.: Are transformers universal approximators of sequence-to-sequence functions? In: Proceedings of the International Conference on Learning Representations (2020). <https://doi.org/10.48550/arXiv.1912.10077>

71. Zhou, D.X.: Universality of deep convolutional neural networks. *Appl. Comput. Harmon. Anal.* **48**(2), 787–794 (2020). <https://doi.org/10.1016/j.acha.2019.06.004>
72. Ziv, A., Olshansky, M., Henis, E., Retiman, A.: IBM accurate portable Mathlib (2001). <https://github.com/dreal-deps/mathlib>

**Open Access** This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

