



HAL
open science

Convexity in ReLU Neural Networks: beyond ICNNs?

Anne Gagneux, Mathurin Massias, Emmanuel Soubies, Rémi Gribonval

► **To cite this version:**

Anne Gagneux, Mathurin Massias, Emmanuel Soubies, Rémi Gribonval. Convexity in ReLU Neural Networks: beyond ICNNs?. 2025. hal-04877619

HAL Id: hal-04877619

<https://hal.science/hal-04877619v1>

Preprint submitted on 9 Jan 2025

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

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

Convexity in ReLU Neural Networks: beyond ICNNs?

Anne Gagneux^{1*}, Mathurin Massias², Emmanuel Soubies³, Rémi Gribonval²

^{1*}Ens de Lyon, CNRS, Université Claude Bernard Lyon 1, Inria, LIP, UMR 5668, 69342, Lyon Cedex 07, France.

^{2*}Inria, ENS de Lyon, CNRS, Université Claude Bernard Lyon 1, LIP, UMR 5668, 69342, Lyon Cedex 07, France.

³Université de Toulouse, CNRS, IRIT, UMR 5505, 31000, Toulouse France.

*Corresponding author(s). E-mail(s): anne.gagneux@ens-lyon.fr;

Abstract

Convex functions and their gradients play a critical role in mathematical imaging, from proximal optimization to Optimal Transport. The successes of deep learning has led many to use *learning-based methods*, where fixed functions or operators are replaced by learned neural networks. Regardless of their empirical superiority, establishing rigorous guarantees for these methods often requires to impose structural constraints on neural architectures, in particular convexity. The most popular way to do so is to use so-called Input Convex Neural Networks (ICNNs). In order to explore the expressivity of ICNNs, we provide necessary and sufficient conditions for a ReLU neural network to be convex. Such characterizations are based on product of weights and activations, and write nicely for any architecture in the *path-lifting* framework. As particular applications, we study our characterizations in depth for 1 and 2-hidden-layer neural networks: we show that every convex function implemented by a 1-hidden-layer ReLU network can be also expressed by an ICNN with the same architecture; however this property no longer holds with more layers. Finally, we provide a numerical procedure that allows an exact check of convexity for ReLU neural networks with a large number of affine regions.

1 Introduction

There exists a strong demand for neural networks that implement expressive convex functions. Here are a few use cases. In image processing, Plug-and-Play (PnP) methods [1] replace explicit proximal operators by learned denoisers using neural networks [2–5]. Since proximal operators are (sub)gradients of convex functions [6], an important ingredient in the mathematical analysis of PnP schemes is to enforce the learned denoiser to be the gradient of a convex function – e.g. through automatic differentiation of a convex neural network [7]. More generally, handling convex priors in inverse problems is often beneficial [8, 9]. Another

example is in optimal transport, where Brenier’s theorem [10] states that any Monge map which moves an absolutely continuous distribution to another distribution while minimizing a quadratic cost is the gradient of a convex function. Hence, many applications which build upon the optimal transport framework seek to train convex neural networks to learn Monge maps [11–14].

Though there exist alternatives [7, 15, 16], Input Convex Neural Networks (ICNNs) [17] are by far the most-used solution when it comes to implementing convex functions with neural networks. ICNNs are a simple approach to enforce convexity which only requires slight modifications

to popular architectures (e.g. ResNets, Convolutional Neural Networks, etc). In short, clipping to zero the negative weights of all hidden layers of a ReLU neural network (except the first one) is sufficient to obtain a convex neural network.

Besides providing a very easy-to-implement approach to get convex neural networks and being widespread in practice, ICNNs are also grounded on the theoretical side: Chen et al. [18] show that ICNNs are rich enough to approximate any convex function over a compact domain in the sup norm.

But is the story over? Not quite. First, ICNNs have shown poor scalability properties [19] and there remains a need for more expressiveness, as suggested by the introduction of several variants of the standard ICNN architecture [11, 20]. Second, the architecture that Chen et al. [18] exhibit has one neuron per layer and *as many layers as affine pieces in the implemented function*, which is far from what one would use in practice. These observations lead us to the following:

Given a ReLU neural network architecture with set depth and widths, do ICNNs cover all convex functions that are implementable with this architecture?

In this paper, we first give a positive answer for neural networks with one hidden-layer: we cannot hope to do better than ICNNs. Yet, with 2 hidden layers, we can already exhibit a network which does not abide by the non-negativity constraint of ICNNs but still implements a convex function. Moreover, we show that *no ICNN with the same number of layers and neurons per layer can implement this function*. These results stem from our work to provide *exact* characterizations of ReLU neural network parameters (on general deep ReLU architectures that notably allow modern features such as skip connections: connections – weighted or not – between neurons of non-consecutive layers) which yield convex functions, beyond ICNNs.

Since the functions implemented by such networks are continuous and piecewise affine, the first stage of our analysis is to characterize the convexity of such functions. This characterization is reminiscent of global convexity conditions for piecewise convex functions [21], but is somehow made “minimal” by finely exploiting the affine nature of the pieces.

The paper is organized as follows: Section 2 formalizes the considered neural network architectures, and provides (Proposition 2.1) our simple, two-hidden-layer motivating counter-example of a convex ReLU network that cannot be implemented by an ICNN with the same architecture. Then, we characterize convexity of continuous piecewise affine functions (Section 3) and provide a first translation of these conditions to one-hidden-layer ReLU networks (Section 4), showing (Proposition 4.9) that in this simple case ICNNs “are all you need”.

The characterization of convexity conditions in the one- and two-hidden-layer case (Sections 4 and 5) serve as a gentle entrypoint to handle more general ReLU networks in Sections 6 and 7. The path-lifting formalism of [22] for such DAG (Directed Acyclic Graph) networks is recalled in Section 6 where we translate the convexity conditions of Section 3 to necessary conditions for DAG ReLU networks. *Sufficient* convexity conditions for DAG networks are expressed in Section 7. Finally, we use these theoretical results to derive in Section 8 a practical algorithm that is able to check convexity on small networks.

2 Convex ReLU networks

Our goal throughout this paper is to study how well ReLU networks can express convex functions. Two key observations will guide our study:

1. Functions implemented by ReLU networks are Continuous Piecewise Linear (CPWL) [23].
2. *Convex* CWPL functions can uniformly approximate arbitrarily well any Lipschitz-continuous convex function on a compact domain [18].

In other words, *the ability of a given ReLU neural network architecture to approximate convex functions comes down to its capability to represent convex CPWL functions*. Therefore, we will focus on how effectively ReLU networks can implement convex CPWL functions.

2.1 ReLU MLPs and ICNNs

We consider a ReLU neural network associated to a function $f_\theta : \mathbb{R}^d \rightarrow \mathbb{R}$ which is a function of the input $x \in \mathbb{R}^d$ and some parameters θ .

In the case of a multilayer perceptron (MLP) or its variants with weighted input skip-connections, let $L \geq 2$ denote the number of (linear) layers. The set of neurons is partitioned into $N_{\text{in}} = N_0$ the input layer ($|N_{\text{in}}| = d$), N_1, \dots, N_{L-1} the hidden layers, and $N_{\text{out}} = N_L$ the last (output) layer (to implement convex functions we focus on scalar-valued networks, so that $|N_L| = 1$). We denote $H := \cup_{\ell=1}^{L-1} N_\ell$ the set of all hidden neurons.

The neural network function f_θ is implemented as follows: starting from an input $x \in \mathbb{R}^d$, we iteratively compute the *pre-activation* $z_\ell \in \mathbb{R}^{N_\ell}$ of each layer ℓ to obtain the network output as

$$\begin{aligned} z_1(x, \theta) &:= \text{affine function of } x, \\ z_\ell(x, \theta) &:= \text{affine function of } \text{ReLU}(z_{\ell-1}) \text{ and } x, \\ &\quad \ell \in \{2, \dots, L\}, \\ f_\theta(x) &:= z_L(x, \theta), \end{aligned}$$

where θ represents the collection of weights and biases in the network's linear functions. We can also define the pre-activation of a neuron¹ $\nu \in N_\ell$ as $z_\nu(x, \theta) = (z_\ell(x, \theta))_\nu$.

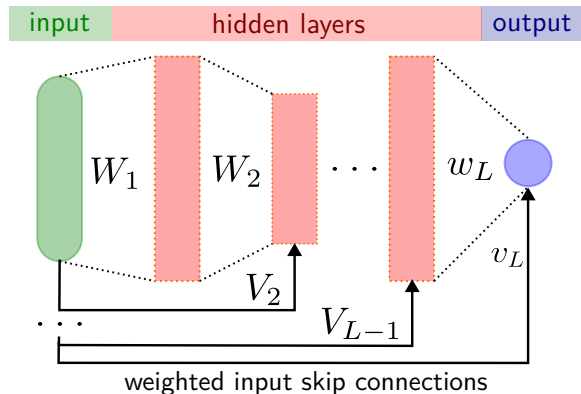


Fig. 1: Architecture of a SkipMLP network. ICNN imposes weights W_2, \dots, W_{L-1}, w_L to be non-negative entry-wise.

ICNNs are a subset of Multi Layers Perceptrons (MLPs) *including linear, weighted, input skip connections* with the additional requirement

¹It will be convenient to consider neurons as distinct elements of the set $N := \cup_{\ell=0}^L N_\ell$ (with $N_0 := N_{\text{in}}$) rather than indices $1 \leq i \leq n_\ell := |N_\ell|$. Hence the notation \mathbb{R}^{N_ℓ} instead of \mathbb{R}^{n_ℓ} . Accordingly $W_\ell \in \mathbb{R}^{N_\ell \times N_{\ell-1}}$ and $b_\ell \in \mathbb{R}^{N_\ell}$.

of non-negativity of certain weights. Based on this difference, we now formalize two classes of functions (see Figure 1): the general class of MLPs with weighted input skip connections (unconstrained) and the constrained subset corresponding to ICNNs. To specify an architecture, we denote $\mathbf{n} = (n_\ell)_{\ell=1}^{L-1}$ the tuple of widths, *excluding the width n_L of the last layer* as we focus on scalar-valued networks ($n_L = 1$). Given \mathbf{n} :

- $\text{SkipMLP}_d(\mathbf{n})$ is the set of all functions $f : \mathbb{R}^d \rightarrow \mathbb{R}$ that can be implemented as a neural network with L layers ($L - 1$ hidden layers) and each hidden layer N_ℓ , $\ell \in \{1, \dots, L - 1\}$ has at most n_ℓ ReLU neurons. In addition are allowed *weighted input skip connections*, linking the input x to each layer $\ell \geq 1$. In this setting, each pre-activation for $\ell \geq 2$ writes as

$$z_\ell(x, \theta) := W_\ell \text{ReLU}(z_{\ell-1}(x, \theta)) + V_\ell x + b_\ell,$$

and $\theta = (W_1, b_1, W_2, V_2, b_2, \dots, w_L, v_L, b_L)$.

- $\text{ICNN}_d(\mathbf{n}) \subseteq \text{SkipMLP}_d(\mathbf{n})$ corresponds to the subset of functions implemented by ICNNs: MLPs with weighted input skip connections *with the additional constraint* that the weight matrices/vectors $(W_2, \dots, W_{L-1}, w_L)$ have non-negative entries. Noticeably, *no non-negativity constraint* is imposed on W_1 nor V_ℓ , $2 \leq \ell \leq L$.

We will usually index weights matrices or bias vectors by neurons: e.g. given two neurons μ and ν in consecutive layers $N_{\ell-1}$ and N_ℓ with $\ell \in \{1, \dots, L\}$, the weight connecting μ to ν is $W_\ell[\nu, \mu]$.

We also refer to *standard* MLPs when we consider MLPs *without input skip connections*: they can be retrieved from the previous classes by imposing zero weights on the input skip connections, *i.e.*, $V_\ell = 0$. Finally, in Section 6, we will consider a general family of ReLU neural networks with Directed Acyclic Graphs (DAGs) architectures.

2.2 Convexity of ICNNs

The convexity of ICNNs is based on rules for composition of convex functions: the composition $f \circ g$ of a convex function g with a non-decreasing convex function f is convex [24, Sec. 3.2.4], and so is any linear combination of convex functions using non-negative coefficients. By recursion over the layers, the entry-wise non-negative weight

constraints $W_\ell \geq 0$ ($\ell \geq 2$) of ICNNs ensure convexity of each neuron’s pre-activation, and of the overall function. Indeed, the first layer implements an affine function, so each component of z_1 is convex. Then, all the subsequent operations are non-decreasing and convex: using convex non-decreasing activation functions (ReLU in our case) and non-negative weight matrices preserves convexity. Besides, this approach also allows for weighted input skip-connections of unconstrained sign between the input and any layer as the sum of a convex function and a linear function remains convex.

2.3 Limits of ICNNs

While Chen et al. [18] have shown that any convex CPWL function can be implemented with an ICNN, the corresponding implementation has a very peculiar architecture, of depth *equal to the number of affine pieces of the function*. This seems potentially suboptimal, and indeed we show in this section that to implement some convex CPWL functions, ICNNs require more neurons or layers than their unconstrained counterpart. In other words, there are convex CPWL functions in $\text{SkipMLP}_d(\mathbf{n}) \setminus \text{ICNN}_d(\mathbf{n})$. To establish this, we simply provide a concrete example.

Proposition 2.1. *Let $f_{\text{EX}} : \mathbb{R}^2 \rightarrow \mathbb{R}$ defined by*

$$f_{\text{EX}} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = (1 \ 1) \text{ReLU} \left(\begin{pmatrix} -1 & 1 \\ 2 & 1 \end{pmatrix} \text{ReLU} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} + \begin{pmatrix} -1 \\ -0.5 \end{pmatrix} \right). \quad (1)$$

This function is convex and CPWL (see Figure F3). By its very definition, it can be implemented by a network in $\text{SkipMLP}((2,2))$, i.e., an MLP with 2 hidden layers, each made of 2 neurons. Yet, it cannot be implemented by any ICNN network with the same architecture, i.e. belonging to $\text{ICNN}_2((2,2))$.

The proof is in Appendix F. This example was constructed from a careful study of necessary and sufficient conditions for convexity of CPWL functions, that we detail in the next section.

3 Convex CPWL functions

Our characterization of the convexity of functions implemented by ReLU networks leverages a general machinery to characterize convexity of CPWL functions, that we establish in this section.

3.1 CPWL functions: definition

Several ways of defining CPWL functions can be found in the literature: either from their linear pieces [25], or from the notion of polyhedral complex and cells [26, 27], or also from the notion of polyhedral partition and regions [28]. We follow the latter approach as it removes some terminology related to polyhedra while allowing a rather simple connection with ReLU networks. A brief complementary reminder is provided in Appendix A.

Definition 3.1 (Polyhedral partition, [28]). A finite family of convex polyhedral sets $(R_k)_{k=1}^K$ is said to form a *polyhedral partition* of \mathbb{R}^d if

- (i) $\mathbb{R}^d = \cup_{k=1}^K R_k$,
- (ii) $\text{int}(R_k) \neq \emptyset$ for every $k \in \{1, \dots, K\}$,
- (iii) $\text{int}(R_k) \cap \text{int}(R_\ell) = \emptyset$ for every $k, \ell \in \{1, \dots, K\}, k \neq \ell$.

The sets $(R_k)_{k=1}^K$ are called *regions* of \mathbb{R}^d .

Definition 3.2 (CPWL function). A function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is said to be continuous and piecewise linear (CPWL) if it is continuous and if there exists a polyhedral partition $(R_k)_{k=1}^K$ of \mathbb{R}^d such that $f|_{R_k}$ is affine for each $k \in \{1, \dots, K\}$.

In Definition 3.2, the polyhedral partition $(R_k)_{k=1}^K$ is not uniquely defined.

Definition 3.3 (Compatible partition). A *compatible partition* for a CPWL function f refers to any polyhedral partition of \mathbb{R}^d for which f is affine on each set of the partition.

In what follows, we introduce a few notions that will come in handy when studying the convexity of CPWL functions. We have not found any equivalent to these definitions in the literature.

Definition 3.4 (Neighboring regions). Two regions R_k and R_ℓ of a polyhedral partition are *neighboring regions* if $k \neq \ell$ and R_k and R_ℓ share a *facet*, i.e. the affine hull of $R_k \cap R_\ell$ is an affine

hyperplane. We denote $R_k \sim R_\ell$ or $k \sim \ell$ in short the neighboring relationship between R_k and R_ℓ .

Definition 3.5 (Frontiers). Given a polyhedral partition $(R_k)_{k=1}^K$, the *frontier* $F_{k,\ell}$ between neighboring regions R_k and R_ℓ is the relative interior of $R_k \cap R_\ell$.

A visualization of neighbouring regions and their frontier is given in Figure C1.

3.2 “Minimal” convexity conditions

We consider in this section a CPWL function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ and investigate convexity conditions. Given a compatible polyhedral partition $(R_k)_{k=1}^K$, we denote $u_k \in \mathbb{R}^d$ and $b_k \in \mathbb{R}$ respectively the slope of $f|_{R_k}$ and its intercept:

$$f|_{R_k} : x \mapsto \langle u_k, x \rangle + b_k. \quad (2)$$

This implies that the gradient of f is well-defined at any $x \in \text{int}(R_k)$, with $\nabla f(x) = u_k$.

Any convex function f must satisfy a well-known monotonicity condition [29, Example 20.3]

$$\langle \nabla f(x) - \nabla f(y), x - y \rangle \geq 0 \quad (3)$$

at every points $x, y \in \mathbb{R}^d$ where f is differentiable. To go further and obtain necessary and sufficient conditions, Bauschke et al. [21] already studied conditions under which a piecewise-defined and piecewise-convex function is globally convex. CPWL functions fall under this framework. More precisely, Bauschke et al. [21] provide sufficient conditions that only need to be checked at boundaries between regions. They also show that a finite number of points can be ignored when checking convexity. Yet, as CPWL functions have *the same slope* on each of their affine components, we are able to derive “minimal” conditions for this specific case: convexity can be checked *at only a finite number of points*. The following proposition ((i) \Leftrightarrow (iv)) asserts that *it suffices to check (3) on a finite number of points* (one pair of points around each frontier $F_{k,\ell}$) to assess convexity.

Proposition 3.6. *Consider a CPWL function $f : \mathbb{R}^d \rightarrow \mathbb{R}$, and any compatible partition $(R_k)_{k=1}^K$. Denote $\mathcal{F} := \bigcup_{k \sim \ell} F_{k,\ell}$ the set of all frontier points. The following are equivalent:*

- (i) *The function f is convex on \mathbb{R}^d .*

- (ii) *For all frontier points $x \in \mathcal{F}$ and $v \in \mathbb{R}^d$, there is $\epsilon > 0$ s.t. $f|_{[x-\epsilon v, x+\epsilon v]}$ is convex.*
- (iii) *For all neighboring regions $R_k \sim R_\ell$, there are $x \in F_{k,\ell}$, $v \in \mathbb{R}^d \setminus \text{span}(R_k \cap R_\ell - x)$ and $\epsilon > 0$ such that $f|_{[x-\epsilon v, x+\epsilon v]}$ is convex.*
- (iv) *(Minimal characterization) For all neighboring regions $R_k \sim R_\ell$, there are $x_k \in \text{int}(R_k)$, $x_\ell \in \text{int}(R_\ell)$ such that² $[x_k, x_\ell] \cap F_{k,\ell} \neq \emptyset$ and*

$$\langle \nabla f(x_k) - \nabla f(x_\ell), x_k - x_\ell \rangle \geq 0. \quad (4)$$

- (v) *(Local characterization) For all frontier points $x \in \mathcal{F}$ there exists a neighborhood \mathcal{N} of x such that for every $x_1, x_2 \in \mathcal{N} \setminus \mathcal{F}$, f is differentiable at x_1 and x_2 and*

$$\langle \nabla f(x_1) - \nabla f(x_2), x_1 - x_2 \rangle \geq 0. \quad (5)$$

The proof in Appendix C uses tools from convex analysis reminded in Appendix B.

4 The one-hidden-layer case

In this section, we use Proposition 3.6 to show that, for one hidden-layer networks, any convex neural network is an ICNN. For this we first introduce a few additional notations and definitions for ReLU networks that are valid *beyond the one-hidden-layer case* and will be re-used throughout the paper.

4.1 Preliminaries

Checking the convexity of a CPWL function using Proposition 3.6 involves checking certain properties near *frontiers between regions*. In the case of ReLU networks, such frontiers are related to *changes in neuron activations*.

We thus start by introducing some notations regarding the activations of neurons of a neural network. Let $f_\theta \in \text{SkipMLP}_d(\mathbf{n})$. It is a CPWL function [23], as the composition and addition of affine and continuous piecewise linear functions. Given an input $x \in \mathbb{R}^d$, if \mathcal{N} is a neighborhood of x where all pre-activations z_ν of all hidden neurons $\nu \in H$ have constant sign, then f_θ is affine on \mathcal{N} . This observation leads us to define the (binary)

² $[x, y]$ denotes a line segment, see (B2) in Appendix B.

activation of a hidden neuron $\nu \in H$:

$$a_\nu(x, \theta) := \mathbf{1}_{z_\nu(x, \theta) > 0} \in \{0, 1\}. \quad (6)$$

We define similarly the layer-wise activation pattern $\mathbf{a}_\ell := (a_\nu)_{\nu \in N_\ell} \in \{0, 1\}^{N_\ell}$ for the ℓ -th hidden layer of ReLU neurons ($1 \leq \ell \leq L - 1$). Note that the function $a_\nu(\cdot, \theta)$ is piecewise constant. To verify convexity of f_θ , we want to look at input points where the slope of f_θ changes, i.e. points around which the activation function of at least one neuron is not constant. To study *necessary* conditions for convexity, we will even focus on points where *only* the activation of one specific neuron changes.

Definition 4.1 (Isolated neuron). Given a parameter θ , for each hidden neuron $\nu \in H := \cup_{\ell=1}^{L-1} N_\ell$, we define³ \mathfrak{X}_ν as the set of input points for which in every small enough neighborhood⁴, *only* the activation of neuron ν switches

$$\mathfrak{X}_\nu := \left\{ x \in \mathbb{R}^d : \exists \epsilon_0 > 0 : \forall 0 < \epsilon \leq \epsilon_0, \right. \\ \left. \begin{cases} a_\mu(\cdot, \theta) \text{ constant on } B(x, \epsilon) & \forall \mu \neq \nu \\ a_\nu(\cdot, \theta) \text{ not constant on } B(x, \epsilon) \end{cases} \right\}.$$

A hidden neuron ν is said to be *isolated* if $\mathfrak{X}_\nu \neq \emptyset$.

4.2 ICNNs are all you need

Consider a one-hidden-layer standard MLP (thus *without* input skip-connection) and denote f_θ the CPWL function it implements, which thus reads:

$$f_\theta : x \in \mathbb{R}^d \mapsto w_2^\top \text{ReLU}(W_1 x + b_1) + b_2 \in \mathbb{R} \quad (7)$$

with $w_2, b_1 \in \mathbb{R}^{N_1}$ and $W_1 \in \mathbb{R}^{N_1 \times N_0}$.

We prove non-negativity of the last layer w_2 associated to isolated neurons when f_θ is convex.

Lemma 4.2 (Non-negativity of the last layer). *Assume that f_θ expressed as (7) is convex. Then, for any isolated hidden neuron of the (only) hidden layer $\nu \in N_1$, the corresponding weight in the last layer must be non-negative: $w_2[\nu] \geq 0$.*

³this set depends on θ , but for the sake of brevity we omit it in the notation, as θ should always be clear from context

⁴ $B(x, \epsilon)$ denotes the open ball (e.g. with respect to the Euclidean metric) of radius $\epsilon > 0$ centered at x

Proof. Consider an isolated neuron $\nu \in N_1$, $x \in \mathfrak{X}_\nu \neq \emptyset$, and a neighborhood $B(x, \epsilon)$ where only the activation of ν changes, as given by Definition 4.1. Using that $\text{ReLU}(t) = \mathbf{1}_{t > 0} t$ it is standard to rewrite f_θ as

$$f_\theta : x \mapsto w_2^\top \text{diag}(\mathbf{a}_1(x, \theta))(W_1 x + b_1) + b_2. \quad (8)$$

The neuron ν being isolated, the layerwise activation pattern $\mathbf{a}_1(\cdot, \theta) \in \{0, 1\}^{N_1}$ takes two different values on $B(x, \epsilon)$, denoted \mathbf{a}^+ and \mathbf{a}^- with $\mathbf{a}_\mu^+ = \mathbf{a}_\mu^-$ for all hidden neurons $\mu \in N_1 \setminus \{\nu\}$ (because $a_\mu(\cdot, \theta)$ is constant on $B(x, \epsilon)$). Wlog, we assume $\mathbf{a}_\nu^+ = 1$ and $\mathbf{a}_\nu^- = 0$. By construction, there exists $x^+, x^- \in B(x, \epsilon)$ such that $\mathbf{a}_1(x^+, \theta) = \mathbf{a}^+$, $\mathbf{a}_1(x^-, \theta) = \mathbf{a}^-$ and that f_θ is differentiable at both points. Denoting $u^+ = \nabla f_\theta(x^+)$, $u^- = \nabla f_\theta(x^-)$, we have:

$$(u^+ - u^-)^\top = w_2^\top (\text{diag}(\mathbf{a}^+) - \text{diag}(\mathbf{a}^-)) W_1 \\ = w_2[\nu] W_1[\nu, :].$$

Since $\mathbf{a}_\nu^+ = 1$, the pre-activation of ν at x^+ is positive: $z_\nu(x^+, \theta) > 0$. Similarly, because $\mathbf{a}_\nu^- = 0$ we have $z_\nu(x^-, \theta) \leq 0$. This rewrites

$$(W_1 x^+ + b_1)[\nu] > 0, \quad (W_1 x^- + b_1)[\nu] \leq 0,$$

from which we get

$$W_1[\nu, :](x^+ - x^-) > 0. \quad (9)$$

Combining with the above expression of $u^+ - u^-$, we obtain

$$\langle u^+ - u^-, x^+ - x^- \rangle = w_2[\nu] \underbrace{W_1[\nu, :](x^+ - x^-)}_{> 0}.$$

Since f_θ is convex, by monotonicity of the gradient (cf Equation (3)) it holds $\langle u^+ - u^-, x^+ - x^- \rangle \geq 0$ from which we get $w_2[\nu] \geq 0$. \square

Remark 4.3. Even though we focus in this section on the one-hidden-layer case, Lemma 4.2 extends to any standard MLP with L layers with essentially the same proof. In Proposition 6.3 we state the most general result which applies to ReLU networks of any depth, in their general DAG (Directed Acyclic Graph) form.

We have just shown that if f_θ is convex, any weight $w_2[\nu]$ corresponding to an *isolated* hidden

neuron ν (i.e. a neuron which can be activated independently of the other neurons) must be non-negative. Before trivially extending the result to one-hidden-layer ReLU networks *with input skip connections*, we characterize one-hidden-layer networks where *all* hidden neurons $\nu \in N_1$ are isolated.

Assumption 4.4. For *all* neurons ν of the first hidden layer N_1 , we have $W_1[\nu, :] \neq 0$. For all pairwise distinct neurons $\nu_1 \neq \nu_2$ of this layer, the corresponding augmented rows $(W_1[\nu_1, :] \mid b_1[\nu_1]) \in \mathbb{R}^{N_1+1}$ and $(W_1[\nu_2, :] \mid b_1[\nu_2]) \in \mathbb{R}^{N_1+1}$ are not co-linear.

Lemma 4.5. *Consider a one-hidden-layer ReLU network, possibly with weighted input skip connections. Assumption 4.4 holds if and only if every neuron $\nu \in N_1$ is isolated.*

Proof. We prove the forward implication (the other implication is left to the reader). Denote $H_\nu := \{x : W_1[\nu, :]x + b_1[\nu] = 0\}$ the 0-level set of $z_\nu(\cdot, \theta)$ for $\nu \in N_1$. By Assumption 4.4:

- For every neuron $\nu \in N_1$, $z_\nu(\cdot, \theta)$ cannot be of constant sign on \mathbb{R}^d , as otherwise we would have $W_1[\nu, :] = 0$. So, $a_\nu(\cdot, \theta)$ is not constant on \mathbb{R}^d .
- Given two distinct neurons $\nu_1 \neq \nu_2$ in the first layer N_1 , the corresponding level sets H_{ν_1}, H_{ν_2} are hyperplanes, and are distinct.

Fix $\nu \in N_1$. Since all considered hyperplanes are pairwise distinct, there exists $x \in H_\nu$ such that for every $\mu \neq \nu$, we have $x \notin H_\mu$. Since $x \in H_\nu$ we have $z_\nu(x, \theta) = 0$ while since $x \notin H_\mu$ we have $z_\mu(x, \theta) \neq 0$ for every $\mu \neq \nu$. As all functions $z_\mu(\cdot, \theta)$ are continuous, it follows that there exists an open ball $B(x, \epsilon)$ on which $\text{sign}(z_\mu(\cdot, \theta)) \neq 0$ is constant for every $\mu \neq \nu$. This implies that $a_\mu(\cdot, \theta)$ is constant on $B(x, \epsilon)$ for every $\mu \neq \nu$ and establishes $x \in \mathfrak{X}_\nu$, thus proving that $\mathfrak{X}_\nu \neq \emptyset$, i.e. by definition ν is isolated. Since this holds for every $\nu \in N_1$ this proves the forward implication. \square

We are now equipped to characterize the convexity of f_θ of the form (7) when all neurons are isolated.

Proposition 4.6. *Under Assumption 4.4, f_θ of the form (7) is convex if and only if its last layer is non-negative: $\forall \nu \in N_1, w_2[\nu] \geq 0$.*

Remark 4.7. Assumption 4.4 cannot be simply omitted: for example, consider the 1D example $f_\theta : \mathbb{R} \rightarrow \mathbb{R}$ and take $W_1 = (1, -1)^\top, b_1 = (0, 0)^\top$ (2 hidden neurons), and $w_2 = (1, -1)^\top$, then $f_\theta(x) = x$, which is convex although w_2 has a negative entry.

Proof of Proposition 4.6 Assume f_θ convex. By Assumption 4.4 and Lemma 4.5, every neuron $\nu \in N_1$ is isolated. By Lemma 4.2 it follows that $w_2[\nu] \geq 0$ for every $\nu \in N_1$. Conversely, if $w_2[\nu] \geq 0$ for every $\nu \in N_1$ then the network is an ICNN, hence f_θ is convex. \square

The extension to one-hidden-layer ReLU networks with input skip connections is an immediate corollary, and the main result of this section.

Corollary 4.8 (Convex one-hidden-layer ReLU networks are ICNNs). *Under Assumption 4.4, the only convex one hidden-layer ReLU networks with weighted input skip connections are Input Convex Neural Networks, i.e. networks which write as*

$$f_\theta : x \in \mathbb{R}^d \mapsto w_2^\top \text{ReLU}(W_1 x + b_1) + b_2 + v_2^\top x,$$

with $w_2 \geq 0$.

Proof. By Proposition 4.6, under Assumption 4.4, $x \in \mathbb{R}^d \mapsto w_2^\top \text{ReLU}(W_1 x + b_1) + b_2$ is convex if and only if $w_2 \geq 0$. Adding or removing the linear term $x \mapsto v_2^\top x$ (corresponding to a weighted input skip connection) does not change the convexity of the considered function, so $x \in \mathbb{R}^d \mapsto w_2^\top \text{ReLU}(W_1 x + b_1) + b_2 + v_2^\top x$ is convex if and only if $w_2 \geq 0$. The condition $w_2 \geq 0$ is precisely the ICNN requirement. \square

Assumption 4.4 is not restrictive, as networks which do not satisfy it can be rewritten as networks of the same depth and fewer neurons that do satisfy it, leading to the following proposition.

Proposition 4.9. Denote CvxCPwL_d the class of convex CPWL functions from \mathbb{R}^d to \mathbb{R} . Then, it holds for any one-hidden-layer architecture $\mathbf{n} = (n)$ with $n \in \mathbb{N}^*$ and any dimension $d \in \mathbb{N}^*$:

$$\text{SkipMLP}_d(\mathbf{n}) \cap \text{CvxCPwL}_d = \text{ICNN}_d(\mathbf{n}), \quad (10)$$

i.e. a convex function implemented by a one-hidden-layer ReLU network with weighted input skip connections can always be also implemented by a one-hidden-layer ICNN with the same width.

Proof. Consider a one-hidden-layer architecture $\mathbf{n} = (n)$ for some width $n \in \mathbb{N}^*$ and let $d \in \mathbb{N}^*$ be the input dimension. The inclusion $\text{ICNN}_d(\mathbf{n}) \subset \text{SkipMLP}_d(\mathbf{n}) \cap \text{CvxCPwL}_d$ is trivial. Consider a convex CPWL function f implemented by a network $\text{SkipMLP}_d(\mathbf{n})$. If Assumption 4.4 is satisfied, then by Corollary 4.8, f belongs to $\text{ICNN}_d(\mathbf{n})$. Otherwise, we show that f can be implemented by some one-hidden-layer Skip-MLP with width $n' \leq n$ which satisfies Assumption 4.4:

- if $W_1[\nu, \cdot] = 0$ for some $\nu \in N_1$, this neuron can be removed from the network architecture, with no impact on the implemented function f , up to adjusting the bias of the last layer.
- if two neurons have co-linear augmented rows (or equivalently have the same hyperplanes $H_{\nu_1} = H_{\nu_2}$), they can be merged (see the discussion on *twin neurons* in Stock and Gribonval [30]). \square

To summarize, once we exclude trivial degeneracies of network parameterizations corresponding to Assumption 4.4, the *only* convex one-hidden-layer ReLU networks (with or without input skip connection) are Input Convex Neural Networks. Does this extend to deeper ReLU networks? Clearly not: Proposition 2.1 yields a non-ICNN two-hidden-layer ReLU network that *does* implement a convex CPWL function. Moreover this function *cannot* be implemented with *any* ICNN with the same depth and width. Next we investigate extensions and adaptation beyond the one-hidden-layer case of the necessary non-negativity condition of Lemma 4.2.

5 The two-hidden-layer case

Before proving our most general result in Section 6 for any DAG ReLU network, we study a simple two-hidden-layer *standard* MLP. It implements the CPWL function

$$f_\theta : x \mapsto w_3^\top \text{ReLU}(W_2 \text{ReLU}(W_1 x + b_1) + b_2) + b_3. \quad (11)$$

The obtained necessary conditions involve, as in the one-hidden-layer case, non-negativity of the last layer, but also of certain *products of weights*.

Regarding necessary convexity conditions on the weights of the last layer, it is not difficult to adapt the proof of Lemma 4.2 (by replacing $\nu \in N_1$ with $\nu \in N_2$, and identifying proper slope vectors u^+, u^-) to show that if f_θ convex, then all weights of the last layer associated to isolated hidden neurons $\nu \in N_2$ are again non-negative: $w_3[\nu] \geq 0$. This necessary condition will be proved formally in the general case of DAG ReLU networks in the next section (Corollary 6.5).

Concerning other weights, we establish an additional necessary condition associated to isolated neurons of the *first* hidden layer, $\nu \in N_1$, and illustrate on the example of Proposition 2.1, as a sanity check, that this condition indeed holds.

5.1 Additional necessary condition

The additional condition involves the reachable activation patterns of the second hidden-layer, i.e. the image of \mathfrak{X}_ν via $\mathbf{a}_2(\cdot, \theta)$, cf. the notations of Section 4.1.

Lemma 5.1. Consider a convex two-hidden-layer ReLU network $f_\theta : \mathbb{R}^d \rightarrow \mathbb{R}$ of the form (11). For each isolated neuron of the first hidden layer $\nu \in N_1$, it holds

$$\min_{\mathbf{a} \in \mathbf{a}_2(\mathfrak{X}_\nu)} \langle \mathbf{a}, (w_3 \odot W_2[:, \nu]) \rangle \geq 0, \quad (12)$$

where \odot is the entry-wise (Hadamard) product, and $\mathbf{a}_2(\mathfrak{X}_\nu) = \{\mathbf{a}_2(x, \theta) : x \in \mathfrak{X}_\nu\}$.

The proof follows the same approach as the one for non-negativity of the last layer in the one-hidden-layer case (Lemma 4.2). More precisely, for each isolated $\nu \in N_1$, it compares the two possible slopes of f_θ at points within a neighborhood of a point where the activation of ν is non constant while other neurons have constant activations.

Proof of Lemma 5.1. Since $\nu \in N_1$ is isolated, we can pick $x \in \mathfrak{X}_\nu \neq \emptyset$ together with a ball $B(x, \epsilon)$ given by Definition 4.1. The activation pattern of the first layer $\mathbf{a}_1(\cdot, \theta) \in \{0, 1\}^{N_1}$ takes two distinct values on $B(x, \epsilon)$, denoted \mathbf{a}^+ and \mathbf{a}^- , with $(\mathbf{a}^+)_\nu = 1$ and $(\mathbf{a}^-)_\nu = 0$. The patterns \mathbf{a}^+ and \mathbf{a}^- are identical elsewhere. Regarding $\mathbf{a}_2(\cdot, \theta) \in \{0, 1\}^{N_2}$, it takes a constant value \mathbf{a}_2 on $B(x, \epsilon)$ by the very definition of this ball for the considered isolated neuron ν . By definition of $\mathbf{a}^+, \mathbf{a}^-$, there are $x^+, x^- \in B(x, \epsilon)$ such that $\mathbf{a}(x^+, \theta) = \mathbf{a}^+$ and $\mathbf{a}(x^-, \theta) = \mathbf{a}^-$ and f is differentiable at x^+ and x^- . Denoting $u^+ = \nabla f_\theta(x^+)$ and $u^- = \nabla f_\theta(x^-)$ we have

$$\begin{aligned} (u^+ - u^-)^\top &= w_3^\top \text{diag}(\mathbf{a}_2) W_2 \text{diag}(\mathbf{a}^+ - \mathbf{a}^-) W_1 \\ &= (w_3^\top \text{diag}(\mathbf{a}_2) W_2) [\nu] W_1 [\nu, :]. \end{aligned}$$

Since $(\mathbf{a}^+)_\nu = 1$ (resp. $(\mathbf{a}^-)_\nu = 0$), we have $z_\nu(x^+, \theta) > 0$ (resp. $z_\nu(x^-, \theta) \leq 0$), i.e.

$$(W_1 x^+ + b_1) [\nu] > 0, \quad (W_1 x^- + b_1) [\nu] \leq 0,$$

from which we get $W_1 [\nu, :](x^+ - x^-) > 0$ and

$$\begin{aligned} \langle u^+ - u^-, x^+ - x^- \rangle &= \\ &= (w_3^\top \text{diag}(\mathbf{a}_2) W_2) [\nu] \underbrace{W_1 [\nu, :](x^+ - x^-)}_{>0}. \end{aligned}$$

Since f_θ is convex, by monotonicity of its gradient (cf Equation (3)) it holds $\langle u^+ - u^-, x^+ - x^- \rangle \geq 0$, from which we get $(w_3^\top \text{diag}(\mathbf{a}_2) W_2) [\nu] \geq 0$, i.e.

$$(\mathbf{a}_2)^\top (w_3 \odot W_2[:, \nu]) \geq 0. \quad (13)$$

Since (13) must hold for every $x \in \mathfrak{X}_\nu$, i.e. for any $\mathbf{a} \in \mathbf{a}_2(\mathfrak{X}_\nu)$, we recover Equation (12). \square

5.2 Comparison to ICNN condition

We have already seen in Proposition 2.1 a convex CPWL implemented by a non-ICNN two-hidden-layer ReLU network. As a sanity check, and a hands-on exercise to manipulate (12), let us verify that this network *does satisfy* the non-negativity condition of Lemma 5.1.

Example 5.2. Recall the parametrization of f_{EX} , which is a convex CPWL function:

$$w_3 = \begin{pmatrix} 1 \\ 1 \end{pmatrix}, \quad W_2 = \begin{pmatrix} -1 & 1 \\ 2 & 1 \end{pmatrix}, \quad W_1 = \text{Id}_2,$$

$$b_3 = 0, \quad b_2 = \begin{pmatrix} -1 \\ -0.5 \end{pmatrix}, \quad b_1 = \mathbf{0}_2.$$

First, all neurons are isolated, see Section F.2 for a proof. Second, $w_3 \geq 0$ satisfies the non-negativity constraint on the last layer. We denote μ_1 and μ_2 the two neurons of the first layer N_1 with their respective pre-activations given by $z_{\mu_1}(x) = \text{ReLU}(x_1)$ and $z_{\mu_2}(x) = \text{ReLU}(x_2)$ where $x = (x_1, x_2)^\top$. Checking (12) for $\nu = \mu_2$ is trivial since $w_3 \odot W_2[:, \mu_2] = (1 \ 1)^\top$ is non-negative as well as *any* activation $\mathbf{a} \in \{0, 1\}^{N_2}$. Since $w_3 \odot W_2[:, \mu_1] = (-1 \ 2)^\top$, (12) holds for $\nu = \mu_1$ if, and only if,

$$\min_{\mathbf{a} \in \mathbf{a}_2(\mathfrak{X}_{\mu_1})} -\mathbf{a}[1] + 2\mathbf{a}[2] \geq 0. \quad (14)$$

As $\mathbf{a}_2(\mathfrak{X}_{\mu_1}) = \{(0, 0), (0, 1), (1, 1)\}$ (see proof in Section F.2), condition (14) indeed holds.

It is natural to wonder whether the necessary convexity condition of Lemma 5.1 for two-hidden-layer standard MLPs extends to two-hidden-layer SkipMLPs, and whether it is also sufficient. Such questions are directly investigated using the more general path-lifting framework with DAG ReLU networks in the next section.

6 DAG ReLU networks

The two-hidden-layer case has highlighted the appearance of *products of layer weights* and their *inner product with activation patterns* in the exposed necessary conditions for convexity. Such structures also appear naturally in the so-called path-lifting framework introduced by Gonon et al. [31] which encompasses very general deep ReLU architectures including skip connections, max-pooling, or convolutional layers. After recalling the needed definitions, we show that they enable a general characterization of necessary convexity conditions in this framework.

For clarity we restrict the exposition in the main text to DAG ReLU networks *without max-pooling*. We discuss in Appendix E adaptations of the definitions (of path activations, and of isolated neurons) that yield the full extension to the general framework of Gonon et al. [31].

6.1 The DAG ReLU model

Consider a ReLU network with parameters θ which can be described as a Directed Acyclic

Graph (DAG): each vertex represents a neuron with weight θ_ν (the bias), each oriented edge connects neurons ν and μ with a parameter $\theta^{\nu \rightarrow \mu}$. The DAG, denoted G , can be described by the tuple (N, E) made of respectively neurons $\nu \in N$ (i.e. vertices of the graph) and directed edges $(\mu, \nu) \in E \subseteq N \times N$.

To each neuron $\nu \in N$ we associate, with a slight abuse of notation, a function $\nu : \mathbb{R}^d \rightarrow \mathbb{R}$ called its *post-activation*. For input $\nu \in N_{\text{in}}$ we have $\nu(x) := x_\nu$ (there is no nonlinearity); the *pre-activation* of hidden neurons and output neurons $\nu \in N \setminus N_{\text{in}}$ is recursively given by an affine combination of the post-activation of its antecedents in the DAG, $z_\nu(x) = \sum_{\mu: (\mu, \nu) \in E} \theta^{\mu \rightarrow \nu} \mu(x) + \theta_\nu$. The post-activation of hidden neurons is given by applying the ReLU to the corresponding pre-activation $\nu(x) = \text{ReLU}(z_\nu(x))$. For output neurons $\nu \in N_{\text{out}}$, we usually have $\nu(x) = z_\nu(x)$ (as for input neuron, there is no nonlinearity). The formalism below however also covers the case where a ReLU is applied at the output.

We are now ready to define the path-lifting Φ and path-activations A .

Definition 6.1 (Path-lifting and path activations, [31]). Consider a ReLU network described by a DAG G and parametrized by θ , with all hidden neurons equipped by a ReLU activation.

- **Paths** A *path* p is a sequence of neurons (p_0, \dots, p_m) (of arbitrary path length $m \geq 0$) connected in the graph G . We denote it as $p = p_0 \rightarrow \dots \rightarrow p_m$. The set \mathcal{P} denotes the collection of all paths *ending at the output neuron*⁵ of G (note that it includes paths which start at hidden or output neurons).
- **Path-lifting** For each path p we define $\phi_p(\theta)$ as a product of parameters, depending on the nature of the first neuron $p_0 \in p$:

$$\phi_p(\theta) = \begin{cases} \prod_{i=1}^m \theta^{p_{i-1} \rightarrow p_i} & \text{if } p_0 \in N_{\text{in}} \\ b_{p_0} \prod_{i=1}^m \theta^{p_{i-1} \rightarrow p_i} & \text{if } p_0 \in H \\ b_{p_0} & \text{if } p_0 \in N_{\text{out}} \end{cases} \quad (15)$$

The path-lifting is $\Phi(\theta) := (\phi_p(\theta))_{p \in \mathcal{P}} \in \mathbb{R}^{\mathcal{P}}$.

⁵Given our interest in scalar-valued function, we consider only one output neuron for simplicity, but all the definitions generalize easily (see [31]) to several output neurons.

- **Path-activations** The activation of a path p is

$$a_p(x, \theta) = \prod_{i=0}^m a_{p_i}(x, \theta) \in \{0, 1\} \quad (16)$$

where by convention $a_\nu(x, \theta) \equiv 1$ for every linear neuron ν (all input neurons, and usually also the output neuron). In other words, the activation of the path is 1 if all hidden neurons on it are active, and 0 otherwise. We define a path-activation vector $a(x, \theta) := (a_p(x, \theta))_{p \in \mathcal{P}} \in \{0, 1\}^{\mathcal{P}}$ and the *path-activations* matrix $A(x, \theta) \in \mathbb{R}^{\mathcal{P} \times (N_{\text{in}}+1)}$ as the matrix with entries

$$A_{p, \nu}(x, \theta) = \begin{cases} a_p(x, \theta) \mathbf{1}_{p_0 = \nu} & \text{if } \nu \in N_{\text{in}}, \\ a_p(x, \theta) \mathbf{1}_{p_0 \notin N_{\text{in}}} & \text{if } \nu \notin N_{\text{in}}. \end{cases} \quad (17)$$

In the path-activations matrix, the columns indexed by N_{in} collect the activation of paths in \mathcal{P} that start from input neurons, while the last column gathers the activations starting from *hidden* neurons (equipped with biases). Just as the neuron-wise functions $a_\nu(\cdot, \theta)$, the matrix-valued function $A(\cdot, \theta)$ is piecewise constant.

Regarding our work, the main interest of these objects is that they enable to write the network output as a scalar product between $\Phi(\theta)$ and $A(x, \theta)$ [31, Lemma A.1]:

$$f_\theta(x) = \left\langle \Phi(\theta), A(x, \theta) \begin{pmatrix} x \\ 1 \end{pmatrix} \right\rangle. \quad (18)$$

From this, the piecewise affine nature of f_θ is clear: at each point x around which the activation patterns are locally constant, the network output is an affine function whose slope depends on the path-lifting $\Phi(\theta)$ and path activation $A(x, \theta)$.

6.2 Path-lifting factorization

It will also be convenient to manipulate variants of the path-lifting and the path-activation matrix associated to certain sub-graphs of G (and to the associated subset of parameters).

Subgraphs and their paths. Given a hidden neuron ν , as illustrated on Figure 2, we denote $G^{\rightarrow \nu}$ the largest (DAG) sub-graph of G with the same input neurons N_{in} and with ν as the single output neuron. Likewise, $G^{\nu \rightarrow}$ is the largest

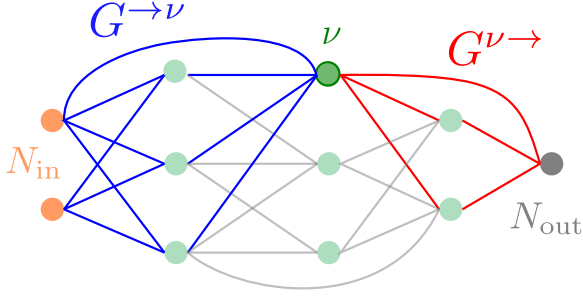


Fig. 2: DAG G and extracted subgraphs $G^{\rightarrow\nu}$ and $G^{\nu\rightarrow}$.

(DAG) sub-graph with ν as its single input neuron and the same output neuron as G .

Path-lifting and activations on $G^{\rightarrow\nu}$. The pre-activation $z_\nu(x, \theta)$ is implemented by the restriction of the initial DAG ReLU network to $G^{\rightarrow\nu}$. Instantiating the above notions yields $\mathcal{P}^{\rightarrow\nu}$, the set of all paths of $G^{\rightarrow\nu}$ ending at (the only) output neuron ν (by convention we equip it with a linear activation instead of its original ReLU), and with a slight abuse of notation (these functions depend on a subset of entries of θ): the path-lifting $\Phi^{\rightarrow\nu}(\theta)$; the path-activation vector $a^{\rightarrow\nu}(x, \theta)$; and the path-activation matrix $A^{\rightarrow\nu}(x, \theta)$. By (18) we have

$$z_\nu(x, \theta) = \left\langle \Phi^{\rightarrow\nu}(\theta), A^{\rightarrow\nu}(x, \theta) \begin{pmatrix} x \\ 1 \end{pmatrix} \right\rangle. \quad (19)$$

Modified definitions on $G^{\nu\rightarrow}$. By analogy with the above notions, and for later notational convenience, we denote $\mathcal{P}^{\nu\rightarrow}$ the set of all paths on $G^{\nu\rightarrow}$ that start at ν and end at the output neuron. The definition of the path-lifting $\Phi^{\nu\rightarrow}$ is:

$$\Phi^{\nu\rightarrow}(\theta) := \left(\prod_{i=1}^m \theta^{p_{i-1} \rightarrow p_i} \right)_{p \in \mathcal{P}^{\nu\rightarrow}} \in \mathbb{R}^{\mathcal{P}^{\nu\rightarrow}}. \quad (20)$$

In contrast with the definitions of Φ and $\Phi^{\rightarrow\nu}$, ν plays here a role analog to that of an input neuron in $G^{\nu\rightarrow}$, and we do not include its bias θ_ν in the path-lifting. The activation of $r = r_0 \rightarrow \dots \rightarrow r_m \in \mathcal{P}^{\nu\rightarrow}$ is defined with a variant of (16)

$$a_r^{\nu\rightarrow}(x, \theta) := \prod_{i=1}^m a_{r_i}(x, \theta), \quad (21)$$

(contrary to a and $a^{\rightarrow\nu}$, the input neuron $r_0 = \nu$ is not part of the product) and we define, similarly to above, the path-activation vector

$$a^{\nu\rightarrow}(x, \theta) := (a_r^{\nu\rightarrow}(x, \theta))_{r \in \mathcal{P}^{\nu\rightarrow}} \in \{0, 1\}^{\mathcal{P}^{\nu\rightarrow}}, \quad (22)$$

We insist that the definition of $\mathcal{P}^{\nu\rightarrow}$ and definitions (20)-(21) introduced on $G^{\nu\rightarrow}$ differ from the standard path-lifting not only because they are restricted to paths that start at ν , but also because the function $a^{\nu\rightarrow}(x, \theta)$ depends on $x \in \mathbb{R}^d$ which is the input vector feeding the initial network defined on G but not the scalar entry z_ν feeding the input neuron ν of $G^{\nu\rightarrow}$.

When x and θ are clear from context, we omit them for brevity. We will exploit the following useful lemma.

Lemma 6.2. *With the above notations, consider an arbitrary hidden neuron ν and any parameter θ . The restriction of $\Phi = \Phi(\theta) \in \mathbb{R}^{\mathcal{P}}$ (resp. of $a = a(x, \theta) \in \{0, 1\}^{\mathcal{P}}$) to paths starting from an input neuron and “containing” ν (i.e. such that $p_i = \nu$ for some i ; this is denoted $p \ni \nu$) satisfies*

$$\Phi_{[p: p_0 \in N_{\text{in}}]}^{p \ni \nu} = \Phi_{[q: q_0 \in N_{\text{in}}]}^{\rightarrow\nu} \otimes \Phi^{\nu\rightarrow}, \quad (23)$$

$$a_{[p: p_0 \in N_{\text{in}}]}^{p \ni \nu} = a_\nu \cdot \left(a_{[q: q_0 \in N_{\text{in}}]}^{\rightarrow\nu} \otimes a^{\nu\rightarrow} \right). \quad (24)$$

Finally,

$$A_{[p: p_0 \in N_{\text{in}}, N_{\text{in}}]}^{p \ni \nu} = a_\nu \cdot \left(A_{[q: q_0 \in N_{\text{in}}, N_{\text{in}}]}^{\rightarrow\nu} \otimes a^{\nu\rightarrow} \right). \quad (25)$$

The proof in Appendix D uses that $\{p \in \mathcal{P} : p \ni \nu, p_0 \in N_{\text{in}}\}$ is in bijection with $\{q \in \mathcal{P}^{\rightarrow\nu} : q_0 \in N_{\text{in}}\} \times \mathcal{P}^{\nu\rightarrow}$. While we choose to focus in this paper on scalar function with a single output neuron as we study convexity, Lemma 6.2 still holds for networks with several output neurons.

6.3 Necessary convexity condition

Thanks to the path-lifting and activation formalism, we can express a necessary convexity condition for generic DAG networks.

Proposition 6.3 (Necessary condition for convex DAG ReLU networks). *Consider a ReLU network described by a DAG and parametrized by θ , which implements a convex CPWL function $f_\theta : \mathbb{R}^d \rightarrow \mathbb{R}$. For every isolated hidden neuron $\nu \in H$, it holds*

$$\min_{x \in \mathfrak{X}_\nu} \langle a^{\nu \rightarrow}(x, \theta), \Phi^{\nu \rightarrow}(\theta) \rangle \geq 0. \quad (26)$$

Remark 6.4. To simplify the exposition we only introduced in the main text the path-lifting and path-activations for DAG ReLU networks in the special case where all hidden neurons are equipped with the ReLU activation. As discussed in Appendix E, these definitions can be extended to also cover max-pooling neurons, again by leveraging the framework of Gonon et al. [31]. With a proper adaptation of the notion of isolated neuron and of the corresponding set \mathfrak{X}_ν (Definition E.3), Proposition 6.3 remains valid.

Introducing the *finite* set of possible activations of the paths $\mathcal{P}^{\nu \rightarrow}$ from ν to the output neuron,

$$\mathbf{a}^{\nu \rightarrow}(\mathfrak{X}_\nu) := \{a^{\nu \rightarrow}(x, \theta) : x \in \mathfrak{X}_\nu\} \subseteq \{0, 1\}^{\mathcal{P}^{\nu \rightarrow}},$$

we get an equivalent expression of (26) which emphasizes its analogy with the special case of (12):

$$\min_{\mathbf{a} \in \mathbf{a}^{\nu \rightarrow}(\mathfrak{X}_\nu)} \langle \mathbf{a}, \Phi^{\nu \rightarrow}(\theta) \rangle,$$

where for two-hidden-layer feedforward networks, $\Phi^{\nu \rightarrow}(\theta)$ is exactly $(w_3 \odot W_2[:, \nu])$.

Before proceeding to the proof of Proposition 6.3, let us highlight a particular consequence: a generalization of Lemma 4.2 on the non-negativity of the last layer. The equivalent of the last layer in a DAG $G = (N, E)$ is the set of hidden neurons $\nu \in H$ that are *antecedents* of the output neurons. Here, since we focus on convex functions, there is a single output neuron ($N_{\text{out}} = \{\mu_{\text{out}}\}$) and $\text{ant}(\mu_{\text{out}}) := \{\nu \in N : (\nu, \mu_{\text{out}}) \in E\}$.

Corollary 6.5 (Necessary condition: non-negativity of the “last layer”). *Isolated hidden neurons in the “last layer”, i.e. $\nu \in H \cap \text{ant}(\mu_{\text{out}})$ such that $\mathfrak{X}_\nu \neq \emptyset$, must have non-negative outgoing weight $\theta^{\nu \rightarrow \mu_{\text{out}}} \geq 0$.*

Remark 6.6. The reader may notice that the above is only expressed for neurons in $\text{ant}(\mu_{\text{out}})$ that are *hidden* neurons: in the presence of input skip connections, there can also be *input neurons* in $\text{ant}(\mu_{\text{out}})$, but there is no non-negativity constraint on the corresponding weight.

Proof. When ν is in the “last layer” $G^{\nu \rightarrow}$ is reduced to a simple graph with two nodes: ν, μ_{out} connected via a single edge $\nu \rightarrow \mu_{\text{out}}$ with weight $\theta^{\nu \rightarrow \mu_{\text{out}}}$. The only path $p \in \mathcal{P}^{\nu \rightarrow}$ starting from ν is $p = \nu \rightarrow \mu_{\text{out}}$, hence $\Phi^{\nu \rightarrow}(\theta) = \phi_p(\theta) = \theta^{\nu \rightarrow \mu_{\text{out}}}$ and $\mathbf{a}^{\nu \rightarrow}(x, \theta) = a_p(x, \theta) = 1$. \square

Remark 6.7. Informally, the non-negativity constraint (26) reads as follows. For every isolated hidden neuron ν of the network, one has to check sign of scalar products between

- the path-lifting for paths which begin at ν and end at the output neuron, and
- the activations of these paths at input points $x \in \mathbb{R}^{N_{\text{in}}}$ at which a_ν is the only activation which switches, i.e. input points x such that $z_\nu(x, \theta) = 0$ and $z_\mu(x, \theta) \neq 0$ for $\mu \neq \nu$.

The proof of Proposition 6.3 relies on the following lemma.

Lemma 6.8 (Local convexity criterion for ReLU networks). *Consider a ReLU network described by a DAG and parametrized by θ which implements a CPWL function $f_\theta : \mathbb{R}^d \rightarrow \mathbb{R}$. Consider an isolated neuron $\nu \in H$ and $x \in \mathfrak{X}_\nu \neq \emptyset$. There exists $\epsilon > 0$ such that*

1. there are $x^+, x^- \in B(x, \epsilon)$ such that f_θ is differentiable at x^+, x^- and $a_\nu(x^+, \theta) = 1$, $a_\nu(x^-, \theta) = 0$;
2. for any such pair x^+, x^- it holds

$$\begin{aligned} & \langle \nabla f_\theta(x^+) - \nabla f_\theta(x^-), x^+ - x^- \rangle \geq 0 \\ \iff & \langle a^{\nu \rightarrow}, \Phi^{\nu \rightarrow}(\theta) \rangle \geq 0. \end{aligned} \quad (27)$$

The proof of this lemma in Appendix E should sound familiar as it is reminiscent of the proof of Lemmas 4.2 and 5.1. Let us highlight that Appendix E actually proves an *extension* of the lemma that covers the case of DAG ReLU networks that may include max-pooling neurons.

Proof of Proposition 6.3 By convexity of f_θ , it holds for any points x^+, x^- where f_θ is differentiable that $\langle u^+ - u^-, x^+ - x^- \rangle \geq 0$ where $u^+ = \nabla f_\theta(x^+), u^- = \nabla f_\theta(x^-)$ denote the slopes of f_θ at x^+, x^- . Combined with Lemma 6.8 this implies that for every isolated neuron we have $\langle a^{\nu \rightarrow}, \Phi^{\nu \rightarrow}(\theta) \rangle \geq 0$. As this non-negativity condition must hold for every $x \in \mathfrak{X}_\nu$, we end up with the claimed result. \square

7 Convex ReLU Networks: sufficient conditions

As in the case of simple feedforward ReLU architectures with few layers, it is natural to wonder whether the necessary convexity condition of Proposition 6.3 is in some sense also sufficient under mild “non-degeneracy” conditions. This is the object of this section.

7.1 Non-degeneracy assumption

Consider a CPWL function $f_\theta : \mathbb{R}^d \rightarrow \mathbb{R}$ implemented by a DAG network, parametrized by θ and fix $(R_k)_{k=1}^K$ any compatible partition. Recall Proposition 3.6(v) on convexity of CPWL functions: it gives sufficient conditions that need to be checked along the set $\mathcal{F} := \bigcup_{k \sim \ell} F_{k,\ell}$ of all frontiers of the CPWL function – where the slope of the CPWL function *may* change. To study convexity, it is in fact sufficient to restrict ourselves to frontiers where the slope of the function f_θ *actually* changes, so we introduce

$$\mathcal{F}_{\text{no-diff}} := \{x \in \mathcal{F} : f_\theta \text{ is not affine on } B(x, \epsilon), \forall \epsilon > 0\} \quad (28)$$

On the other hand, the necessary conditions of Proposition 6.3 involve the set $\bigcup_{\nu \in H} \mathfrak{X}_\nu$, the set of input points where exactly one hidden neuron switches. Note that if the slope of f_θ changes at input point x , then at least one neuron must switch at x . The connection we need to easily establish sufficient conditions is to assume that for any input point on a frontier where the slope changes, *one and only one neuron “switches”*.

Assumption 7.1. Consider a ReLU network described by a DAG with parameters θ which implements a CPWL function $f_\theta : \mathbb{R}^d \rightarrow \mathbb{R}$

and consider a compatible polyhedral partition $(R_k)_{k=1}^K$ associated with this CPWL function. The assumption states

$$\mathfrak{X}_\nu \neq \emptyset \quad \forall \nu \in H \quad (29)$$

and

$$\mathcal{F}_{\text{no-diff}} \subset \bigcup_{\nu \in H} \mathfrak{X}_\nu. \quad (30)$$

Remark 7.2. In words, the assumption requires that every hidden neuron in the network is isolated, and that every change of slope of f_θ at points belonging to a $d - 1$ dimensional face between regions corresponds to a change of activation of a *single* neuron.

For one-hidden-layer networks Assumption 7.1 simply coincides with Assumption 4.4.

Lemma 7.3. *For a one-hidden-layer ReLU network architecture, Assumptions 4.4 and 7.1 are equivalent.*

Proof. Consider a function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ in $\text{SkipMLP}_d(\mathbf{n})$ with $\mathbf{n} = (n)$ for some $n \in \mathbb{N}^*$. In this case, the set of hidden neurons is $H = N_1$. By Lemma 4.5, Assumption 4.4 is equivalent to the fact that each neuron $\nu \in N_1$ is isolated, which by Definition 4.1 is equivalent to (29). Thus, Assumption 7.1 implies Assumption 4.4. Conversely, assume that Assumption 4.4 holds. To recover Assumption 7.1, it suffices to show (30). First, define the 0-level sets of each neuron $\nu \in N_1$ as $H_\nu := \{x \in \mathbb{R}^d : z_\nu(x) = W_1[\nu, :]x + b_1[\nu] = 0\}$. By the non-degeneracy assumption ($W_1[\nu, :] \neq 0$ for each $\nu \in N_1$) of Assumption 4.4, these sets H_ν are affine hyperplanes, and by the non-colinearity assumption of Assumption 4.4, they cannot coincide nor be parallel. Considering $x \in \mathcal{F}_{\text{no-diff}}$ we wish to prove that $x \in \bigcup_{\nu \in H} \mathfrak{X}_\nu$. By contradiction, assume that $x \notin \bigcup_{\nu \in H} \mathfrak{X}_\nu$. Because of the non-differentiability of f at x , there exists $\nu \in N_1$ such that $x \in H_\nu$. Because $x \notin \mathfrak{X}_\nu$, another neuron switches: there exists $\mu \neq \nu \in N_1$ such that $z_\mu(x) = 0$. Hence $x \in H_\nu \cap H_\mu$, so x belongs at least to four different regions which contradicts $x \in \mathcal{F}$. \square

7.2 Characterization of convex ReLU DAG networks

Theorem 7.4 (NSC for convex ReLU networks). *Consider a ReLU network described by a DAG and parametrized by θ which implements a CPWL function $f_\theta : \mathbb{R}^d \rightarrow \mathbb{R}$.*

Under Assumption 7.1, the function f_θ is convex if and only if for every hidden neuron ν of the network, it holds

$$\min_{x \in \mathfrak{X}_\nu} \langle a^{\nu \rightarrow}(x, \theta), \Phi^{\nu \rightarrow}(\theta) \rangle \geq 0. \quad (31)$$

This holds even with max-pooling neurons, with \mathfrak{X}_ν , $a^{\nu \rightarrow}$, $\Phi^{\nu \rightarrow}$ as defined in Appendix E.

Proof. Necessity. Assume f_θ convex. By Assumption 7.1, every hidden neuron ν is isolated (as $\mathfrak{X}_\nu \neq \emptyset$). We conclude with Proposition 6.3.

Sufficiency. Assume that (31) holds for every hidden neuron (notice that, by Assumption 7.1, $\mathfrak{X}_\nu \neq \emptyset$). We will exploit the characterization of Proposition 3.6 (v). For this, consider an arbitrary $x \in \mathcal{F}$. If f_θ is affine for every small enough ball centered at x then we directly get the expected inequality (5). Otherwise we must have $x \in \mathcal{F}_{\text{no-diff}}$. By Assumption 7.1, there exists a hidden neuron ν such that $x \in \mathfrak{X}_\nu$. By (31) we have $\langle a^{\nu \rightarrow}(x, \theta), \Phi^{\nu \rightarrow}(\theta) \rangle \geq 0$. By Lemma 6.8, there exists $\epsilon > 0$ such that for all $x^+, x^- \in B(x, \epsilon) \setminus \mathcal{F}_{\text{no-diff}}$, f_θ is differentiable, and with $u^+ = \nabla f_\theta(x^+)$, $u^- = \nabla f_\theta(x^-)$ we have $\langle u^+ - u^-, x^+ - x^- \rangle \geq 0$. This yields again (5), which therefore holds for any $x \in \mathcal{F}$. We conclude using Proposition 3.6 (v). \square

7.3 Genericity of Assumption 7.1

In this section, we comment on why Assumption 7.1 is not too restrictive.

Back to the one-hidden-layer case, Lemma 7.3 shows that Assumptions 4.4 and 7.1 are equivalent. As already discussed in Section 4, Assumption 4.4 is not restrictive in the sense that every function implemented by a network that does not satisfy it can also be implemented by a network of the same depth and with fewer neurons which *does* satisfy it (resulting in Proposition 4.9).

Still in the one-hidden-layer case, a stronger assumption, which implies Assumption 4.4 and yet is generic, is to assume that for each $\nu \in N_1$

the 0-level set of z_ν is a hyperplane, and that the resulting collection of hyperplanes is in *general position*⁶. This is more restrictive than Assumption 4.4: in \mathbb{R}^2 , it prevents 3 hyperplanes to intersect at the same point. In particular, it does not allow null biases.

When looking at deeper *standard MLPs*⁷ ReLU networks, there exists an extension of this general position property that applies to so-called *bent* hyperplanes, i.e. to the 0-level set of pre-activation of any hidden neuron $\nu \in H$.

Grigsby and Lindsey [26] introduce the *transversality* property, corresponding to every 0-level set of the pre-activation of any neuron $\nu \in N$ either having dimension $d - 1$ or being empty. Masden [27] then introduced the *supertransversality property* ensuring that intersection of k bent hyperplanes has dimension $d - k$. These properties are shown to hold on almost every standard ReLU MLPs. Because of the possibility of having hidden neurons whose activation never changes, these assumptions do not exactly match Assumption 7.1. Yet, since such neurons do not introduce new regions, we believe that the characterization of Theorem 7.4 would still hold under a slight modification of Assumption 7.1 that would be a direct consequence of transversality and supertransversality. Such a variant of Assumption 7.1 would hence be likely to hold for almost every ReLU network.

7.4 Back to the 2D example

We can leverage the sufficient conditions given by Theorem 7.4 to prove convexity of the 2D function f_{EX} of Proposition 2.1. We have already mentioned (see proof in Appendix F) that all its hidden neurons are isolated and that at each frontier point where f_{EX} is non-differentiable, there is a single neuron whose activation changes. Hence, Assumption 7.1 holds. To prove convexity of f_{EX} ,

⁶Let $\nu_1, \dots, \nu_m \in N_1$ and $\mathcal{H} := \{H_{\nu_1}, \dots, H_{\nu_m}\}$ where $H_{\nu_i} := \{x \in \mathbb{R}^d : W_1[\nu_i, :]x + b_1[\nu_i] = 0\}$. If $W_1[\nu_i, :]$ is a non-zero vector for all neurons ν_i , then \mathcal{H} is a *finite (affine) hyperplane arrangement*. Such an arrangement is said to be in *general position* [32] if:

- (i) for $k \leq d$, every k -fold intersection of hyperplanes has dimension $d - k$,
- (ii) for $k > d$, every k -fold intersection of hyperplanes is empty.

⁷i.e., with no skip connection

it is sufficient to check that for every neuron ν in $N_1 \cup N_2$, Equation (31) holds. Denoting ν_{out} the output neuron, for each neuron $\nu \in N_2$ of the second hidden layer the set of paths $\mathcal{P}^{\nu \rightarrow} = \{\nu \rightarrow \nu_{\text{out}}\}$ is a singleton and for all $x \in \mathfrak{X}_\nu$, $a^{\nu \rightarrow}(x, \theta) := (a_{\nu_{\text{out}}}(x, \theta)) = (1)$ and $\phi^{\nu \rightarrow} = (\theta^{\nu \rightarrow \nu_{\text{out}}}) = w_3[\nu]$. As $w_3 = (1, 1)^\top$, Equation (31) (which simply reads $w_3[\nu] \geq 0$) is satisfied for all $\nu \in N_2$. Now, consider $\mu \in N_1$ and denote $N_2 := \{\nu_1, \nu_2\}$. We have $\mathcal{P}^{\mu \rightarrow} = \{\mu \rightarrow \nu_1 \rightarrow \nu_{\text{out}}, \mu \rightarrow \nu_2 \rightarrow \nu_{\text{out}}\}$, and for all $x \in \mathfrak{X}_\mu$,

$$a^{\mu \rightarrow}(x, \theta) = \begin{pmatrix} a_{\nu_1}(x, \theta) & \underbrace{a_{\nu_{\text{out}}}(x, \theta)}_{:=1} \\ a_{\nu_2}(x, \theta) & \underbrace{a_{\nu_{\text{out}}}(x, \theta)}_{:=1} \end{pmatrix} = \mathbf{a}_2(x, \theta),$$

$$\phi^{\mu \rightarrow}(\theta) = \begin{pmatrix} \theta^{\mu \rightarrow \nu_1} \theta^{\nu_1 \rightarrow \nu_{\text{out}}} \\ \theta^{\mu \rightarrow \nu_2} \theta^{\nu_2 \rightarrow \nu_{\text{out}}} \end{pmatrix} = (w_3 \odot W_2[:, \mu]),$$

so Equation (31) boils down –as expected– to the same necessary condition given in the two-hidden-layer case

$$\min_{x \in \mathfrak{X}_\mu} \langle \mathbf{a}_2(x, \theta), w_3 \odot W_2[:, \mu] \rangle \geq 0, \quad (32)$$

which we have already verified in Section 5.2.

8 Numerical check of convexity

Given an architecture $G = (N, E)$, the two main computational bottlenecks to check the convexity conditions of Theorem 7.4 for a ReLU network with architecture G and parameters θ are:

1. **Handling the dimension of the vectors** $\phi^{\nu \rightarrow}$ and $a^{\nu \rightarrow}$ involved in the scalar product. Considering *feedforward* networks, the number of paths grows exponentially with the number of layers, and in general with the depth of the network, so this is *a priori* a daunting task.
2. **Computing the set of activation patterns** $a^{\nu \rightarrow}(\mathfrak{X}_\nu)$ for every hidden neuron $\nu \in H$. This requires identifying all input points $x \in \mathfrak{X}_\nu$ where only the activation of ν changes. In practice, it amounts to identify all frontiers and regions of the CPWL function, whose number grows exponentially with the number of neurons.

The first bottleneck can be addressed as follows. Given one hidden neuron $\nu \in H$ and one path-activation vector $\mathbf{a} \in \mathbf{a}^{\nu \rightarrow}(\mathfrak{X}_\nu)$, the constraint to be checked $\langle \mathbf{a}, \Phi^{\nu \rightarrow}(\theta) \rangle \geq 0$ is reminiscent of Equation (18) which gives the network output as a scalar product between the path-activation matrix and the path-lifting. More precisely, $\langle \mathbf{a}, \Phi^{\nu \rightarrow}(\theta) \rangle \geq 0$ is *exactly* the output of a modified network with architecture $G^{\nu \rightarrow}$ (ν is its input neuron), *with a scalar input* equal to 1, *with the biases of all its hidden neurons set to zero*, and with activation values given by \mathbf{a} (i.e., unlike in standard ReLU networks, the activation here is *not* dependent on the modified network’s input; this is further detailed below). Given a hidden neuron $\nu \in H$, suppose we have already identified a finite subset $\mathfrak{X}_\nu^{\text{finite}} \subset \mathfrak{X}_\nu$ such that $\mathbf{a}^{\nu \rightarrow}(\mathfrak{X}_\nu^{\text{finite}}) = \mathbf{a}^{\nu \rightarrow}(\mathfrak{X}_\nu)$ (remember that there is a finite number of activation patterns in $\mathbf{a}^{\nu \rightarrow}(\mathfrak{X}_\nu)$). Then, one can circumvent the explicit computation of the vector $\Phi^{\nu \rightarrow}$ (of combinatorially high dimension) and the enumeration of all paths by proceeding as follows:

- extract the architecture corresponding to the sub-graph $G^{\nu \rightarrow} := (N^{\nu \rightarrow}, E^{\nu \rightarrow})$;
- set all biases of the sub-network to zero, i.e. parametrize the architecture $G^{\nu \rightarrow}$ with $\tilde{\theta}$ such that

$$\tilde{\theta}_\mu := 0 \quad \forall \mu \in N^{\nu \rightarrow} \quad (33)$$

$$\tilde{\theta}^{\mu \rightarrow \mu'} := \theta^{\mu \rightarrow \mu'} \quad \forall (\mu, \mu') \in E^{\nu \rightarrow} \quad (34)$$

- for each identified input point $x \in \mathfrak{X}_\nu^{\text{finite}}$, perform a forward pass on the *initial network* and store the activations neuron-wise, i.e. store $a_\mu := a_\mu(\theta, x)$ for every $\mu \in N$;
- replace the ReLU activation⁸ of each hidden and output neuron $\mu \in N^{\nu \rightarrow}$ by a function that computes a simple product $z \mapsto \mu(x) := a_\mu(x)z$,
- do a forward pass with scalar input equal to 1 to obtain $\langle a^{\nu \rightarrow}(x, \theta), \Phi^{\nu \rightarrow}(\theta) \rangle$.

Regarding the second bottleneck, it requires identifying points in \mathfrak{X}_ν : notice that any input point $x \in \mathfrak{X}_\nu$ is such that $z_\nu(x) = 0$ and $z_\mu(x) \neq 0$ for $\mu \neq \nu$. The problem of identifying bent hyperplanes (i.e. 0-level sets of z_ν) has

⁸The adaptation to max-pooling neurons in the spirit of Appendix E is straightforward.

already been previously studied [33–35]. Notably, Berzins [35] provides an efficient algorithm which –given a *compact* input domain– *exactly* extracts all intersections between bent hyperplanes. This algorithm is able to extract 0–faces and 1–faces of the polyhedral partition of a ReLU network in a few seconds for architectures with up to 10 layers and hundreds of neurons per layer, defining millions of regions.

Remark 8.1. While all the convexity conditions expressed in this paper are for convex functions $f : \mathbb{R}^d \rightarrow \mathbb{R}$, the framework can easily be adapted to study convexity on any non-empty convex domain $\Omega \subseteq \mathbb{R}^d$. This essentially requires replacing \mathfrak{X}_ν from Definition 4.1 (and its variant for max-pooling from Appendix E) and \mathcal{F} by their intersection with Ω everywhere.

We build upon the framework of [35] to provide an algorithm which numerically checks convexity of a given ReLU network on a compact convex domain $\Omega \subseteq \mathbb{R}^d$: it checks if the necessary conditions from Theorem 7.4 are satisfied with \mathfrak{X}_ν replaced by its intersection with Ω . We detail in Appendix G how to proceed to obtain the sets $a^{\nu \rightarrow}(\mathfrak{X}_\nu)$ for each hidden neuron $\nu \in H$. Then, we apply the method described above to compute the scalar product $\langle \mathbf{a}, \Phi^{\nu \rightarrow} \rangle$ and finally we test whether all these inner products are non-negative.

8.1 Experiments

In this section, we look at the probability of randomly sampling a convex ReLU network, for a given architecture, when drawing parameters θ from a Gaussian distribution. Convexity of the implemented function f_θ is numerically checked with the method described previously. We consider standard MLPs with two-hidden-layer networks and input dimension equal to 2. The width of each hidden layer varies from 2 to 7.

For each of these architectures $\mathbf{n} = (n_1, n_2)$ with $n_1, n_2 \in \{2, \dots, 7\}$, we draw 10^4 parameters θ (weights and biases) at random according to a standard Gaussian. We compare the number of sampled convex functions to the number of sampled ICNNs. Note, the probability of drawing a random ICNN can be easily analytically computed: each weight in W_2 and w_3 having probability 1/2 to be non-negative under standard

Gaussian initialization, it follows

$$\mathbb{P}(f_\theta \in \text{ICNN}_2(\mathbf{n})) = \frac{1}{2^{n_2(n_1+1)}}. \quad (35)$$

For a small architecture $\mathbf{n} = (2, 2)$, $\mathbb{P}(f_\theta \in \text{ICNN}_2(\mathbf{n})) \approx 0.016$ which is to be compared with the experimental frequency of convex ReLU networks (including the ICNNs) which is 4 times greater (Figure 3). Increasing the number of neurons in the architecture further increase the ratio between the number of obtained convex ReLU networks and the number of obtained ICNNs: adding a single neuron more than double this ratio. This suggests that for larger architectures, ICNNs are a very small fraction of all the convex ReLU networks implementable.

While our condition for convexity could allow one to analytically derive the probability of having a convex ReLU network (not necessarily ICNN), the combinatorial aspects of the exploitation of these conditions remain currently somewhat challenging. Yet, we anticipate that this should be possible in the $\mathbf{n} = (2, 2)$ case for instance.

9 Conclusion

Through our exploration of convexity conditions for ReLU networks, we confirmed the intuition that there is life beyond ICNNs. Despite being the current standard to implement convex functions, ICNNs have intrinsic expressivity limitations starting from simple two-hidden-layer networks. Thanks to the analysis of necessary and sufficient convexity conditions for the general class of DAG ReLU networks, we have highlighted that there are many parameterizations of such neural networks that yield convex functions, far beyond the restricted scope of ICNNs.

The expressed conditions allow to concretely *test* the convexity of a given network of moderate size, and allowed to concretely observe on small random two-hidden-layer networks the increased probability of drawing a convex function when using Gaussian parameters. Being based on the non-negativity of a finite number of computable criteria, these conditions also naturally give rise to possible regularizers to promote convexity while training a network. One of the main challenges lying ahead is computational, with a particular focus on reducing as much as possible the number

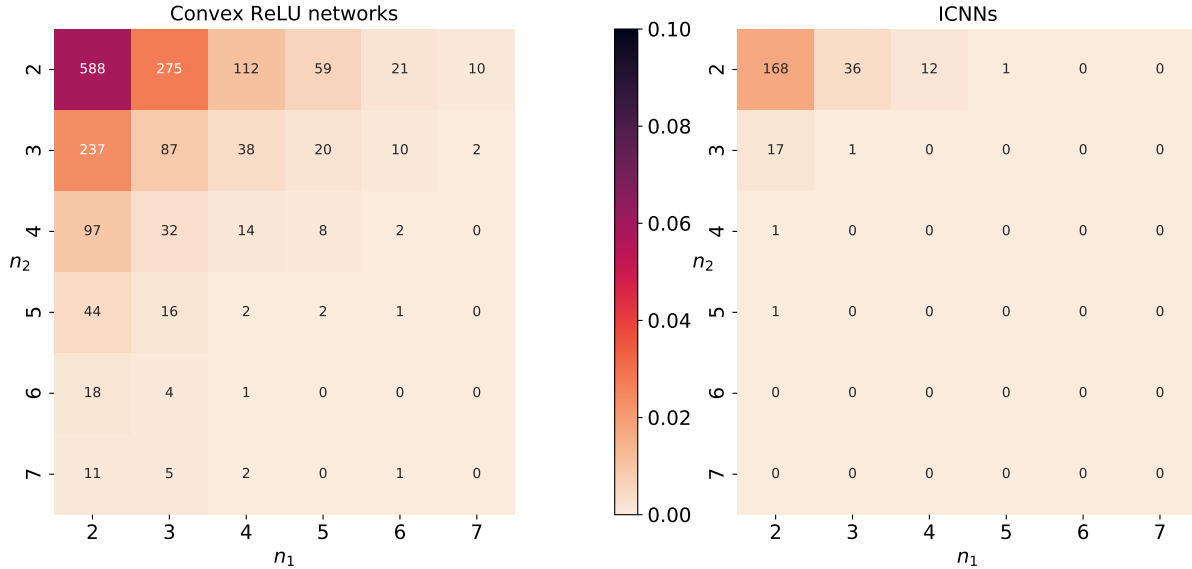


Fig. 3: Number of convex ReLU networks among 10^4 draws. (Left): convex ReLU networks. (Right): ICNNs. Architecture $\mathbf{n} = (n_1, n_2)$.

of convexity conditions to be tested. A possible avenue is to explore possible redundancies between these constraints as well as the graph structure of the network.

10 Acknowledgements

R. Gribonval is grateful to Antonin Chambolle for enlightening discussions on related topics during the SMAI-MODE 2024 conference. This project was supported in part by the AllegroAssai ANR project ANR-19-CHIA0009 and by the SHARP ANR project ANR-23-PEIA-0008 funded in the context of the France 2030 program. A. Gagneux thanks the RT IASIS for supporting her work through the PROSSIMO grant.

References

- [1] Venkatakrishnan, S., Bouman, C., Wohlberg, B.: Plug-and-play priors for model based reconstruction. In: 2013 IEEE Global Conference on Signal and Information Processing, pp. 945–948 (2013). IEEE
- [2] Meinhardt, T., Moller, M., Hazirbas, C., Cremers, D.: Learning proximal operators: Using denoising networks for regularizing inverse imaging problems. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 1781–1790 (2017)
- [3] Zhang, K., Zuo, W., Chen, Y., Meng, D., Zhang, L.: Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising. IEEE transactions on image processing **26**(7), 3142–3155 (2017)
- [4] Zhang, K., Li, Y., Zuo, W., Zhang, L., Van Gool, L., Timofte, R.: Plug-and-play image restoration with deep denoiser prior. IEEE Transactions on Pattern Analysis and Machine Intelligence **44**(10), 6360–6376 (2021)
- [5] Pesquet, J.-C., Repetti, A., Terris, M.,

- Wiaux, Y.: Learning maximally monotone operators for image recovery. *SIAM Journal on Imaging Sciences* **14**(3), 1206–1237 (2021)
- [6] Gribonval, R., Nikolova, M.: A characterization of proximity operators. *Journal of Mathematical Imaging and Vision* **62**(6-7), 773–789 (2020)
- [7] Hurault, S.: Convergent plug-and-play methods for image inverse problems with explicit and nonconvex deep regularization. PhD thesis, Université de Bordeaux (2023)
- [8] Cohen, R., Blau, Y., Freedman, D., Rivlin, E.: It has potential: Gradient-driven denoisers for convergent solutions to inverse problems. *Neurips* (2021)
- [9] Fang, Z., Buchanan, S., Sulam, J.: What’s in a prior? learned proximal networks for inverse problems. In: *ICLR* (2024)
- [10] Brenier, Y.: Polar factorization and monotone rearrangement of vector-valued functions. *Communications on pure and applied mathematics* **44**(4), 375–417 (1991)
- [11] Korotin, A., Egiazarian, V., Asadulaev, A., Safin, A., Burnaev, E.: Wasserstein-2 generative networks. *arXiv preprint arXiv:1909.13082* (2019)
- [12] Makuva, A., Taghvaei, A., Oh, S., Lee, J.: Optimal transport mapping via input convex neural networks. In: *International Conference on Machine Learning*, pp. 6672–6681 (2020). PMLR
- [13] Korotin, A., Li, L., Genevay, A., Solomon, J.M., Filippov, A., Burnaev, E.: Do neural optimal transport solvers work? a continuous wasserstein-2 benchmark. *Neurips* (2021)
- [14] Amos, B.: On amortizing convex conjugates for optimal transport. *arXiv preprint arXiv:2210.12153* (2022)
- [15] Saremi, S.: On approximating ∇f with neural networks. *arXiv preprint arXiv:1910.12744* (2019)
- [16] Chaudhari, S., Pranav, S., Moura, J.M.: Learning gradients of convex functions with monotone gradient networks. In: *ICASSP 2023* (2023). IEEE
- [17] Amos, B., Xu, L., Kolter, J.Z.: Input convex neural networks. In: *International Conference on Machine Learning*, pp. 146–155 (2017). PMLR
- [18] Chen, Y., Shi, Y., Zhang, B.: Optimal control via neural networks: A convex approach. In: *ICLR* (2019)
- [19] Korotin, A., Selikhanovych, D., Burnaev, E.: Neural optimal transport. In: *ICLR* (2023)
- [20] Vesseron, N., Cuturi, M.: On a neural implementation of brenier’s polar factorization. In: *ICML* (2024)
- [21] Bauschke, H.H., Lucet, Y., Phan, H.M.: On the convexity of piecewise-defined functions. *ESAIM: Control, Optimisation and Calculus of Variations* **22**(3), 728–742 (2016)
- [22] Gonon, A., Brisebarre, N., Gribonval, R., Riccietti, E.: Approximation speed of quantized vs. unquantized relu neural networks and beyond. *IEEE Transactions on Information Theory* (2023)
- [23] Arora, R., Basu, A., Mianjy, P., Mukherjee, A.: Understanding Deep Neural Networks with Rectified Linear Units. *ICLR* (2018)
- [24] Boyd, S., Vandenberghe, L.: *Convex Optimization*. Cambridge university press, (2004)
- [25] Goujon, A., Etemadi, A., Unser, M.: On the number of regions of piecewise linear neural networks. *Journal of Computational and Applied Mathematics* **441**, 115667 (2024)
- [26] Grigsby, J.E., Lindsey, K.: On transversality of bent hyperplane arrangements and the topological expressiveness of relu neural networks. *SIAM Journal on Applied Algebra and Geometry* **6**(2), 216–242 (2022)
- [27] Masden, M.: Algorithmic determination of

- the combinatorial structure of the linear regions of relu neural networks. arXiv preprint arXiv:2207.07696 (2022)
- [28] Gorokhovich, V.V., Zorko, O.I., Birkhoff, G.: Piecewise affine functions and polyhedral sets. *Optimization* **31**(3), 209–221 (1994)
- [29] Bauschke, H., Combettes, P.: *Convex Analysis and Monotone Operator Theory in Hilbert Spaces* vol. 408. Springer, (2011)
- [30] Stock, P., Gribonval, R.: *An Embedding of ReLU Networks and an Analysis of their Identifiability*. *Constructive Approximation* (2022). Publisher: Springer Verlag
- [31] Gonon, A., Brisebarre, N., Riccietti, E., Gribonval, R.: A path-norm toolkit for modern networks: consequences, promises and challenges. arXiv preprint arXiv:2310.01225 (2023)
- [32] Stanley, R.P., *et al.*: An introduction to hyperplane arrangements. *Geometric combinatorics* **13**(389-496), 24 (2004)
- [33] Rolnick, D., Kording, K.: Reverse-engineering deep relu networks. In: *International Conference on Machine Learning*, pp. 8178–8187 (2020). PMLR
- [34] Humayun, A.I., Balestrieri, R., Balakrishnan, G., Baraniuk, R.G.: Splinecam: Exact visualization and characterization of deep network geometry and decision boundaries. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3789–3798 (2023)
- [35] Berzins, A.: Polyhedral complex extraction from relu networks using edge subdivision. In: *ICML* (2023)
- [36] Ziegler, G.M.: *Lectures on Polytopes*. Springer, (1995)
- [37] Rockafellar, R.T.: *Convex analysis*:(pms-28) (2015)
- [38] Moreau, J.J.: Fonctionnelles sous-différentiables. *Comptes rendus hebdomadaires des séances de l’Académie des sciences* **257**, 4117–4119 (1963)
- [39] Rockafellar, R.T., Wets, R.J.-B.: *Variational Analysis* vol. 317. Springer, (2009)
- [40] Hanin, B., Rolnick, D.: Deep relu networks have surprisingly few activation patterns. *Neurips* (2019)

Appendix A Background on CPWL functions

Definition A.1 (Affine function). A function $f : \Omega \rightarrow \mathbb{R}$ defined on a convex set $\Omega \subset \mathbb{R}^d$ is *affine* if and only if for every $x, y \in \Omega$, $t \in [0, 1]$, $f((1-t)x + ty) = (1-t)f(x) + tf(y)$.

Proposition A.2 ([24, Section 3.1.1]). A function $f : \Omega \rightarrow \mathbb{R}$ defined on a convex set $\Omega \subset \mathbb{R}^d$ is *affine* if and only if there exists $u \in \mathbb{R}^d$, $b \in \mathbb{R}$ such that, for every $x \in \Omega$,

$$f(x) = \langle u, x \rangle + b. \quad (\text{A1})$$

Definition A.3 (Polyhedron). A polyhedron (also called a polyhedral set) $\Omega \subset \mathbb{R}^d$ is a non-empty intersection of finitely many closed half-spaces. A polytope is a polyhedron that is bounded (in \mathbb{R}^d).

We recall below some useful terminology for polyhedra, see also [36].

Definition A.4. Consider a polyhedral set P .

- A linear inequality $a^\top x \leq b$ is *valid* for P if it is satisfied for all points $x \in P$.
- A set F is called a *face* of P if $F = P \cap \{x \in \mathbb{R}^d : a^\top x = b\}$ where $a^\top x \leq b$ is a valid inequality for P . A face of a polyhedron is also a polyhedron.
- The dimension of a polyhedron P is the dimension of its affine hull $\text{aff}(P)$. Likewise, the dimension of a face is the dimension of its affine hull: $\dim(F) = \dim(\text{aff}(F))$.
- A face whose dimension is $d - 1$ is called a *facet*.
- $F = \emptyset$ and $F = P$ are *improper* faces. All other faces are *proper* faces.

Appendix B Background on convexity

Definition B.1 (Line segment). The line segment between $x \in \mathbb{R}^d$ and $y \in \mathbb{R}^d$ is

$$[x, y] := \{(1-t)x + ty : t \in [0, 1]\}. \quad (\text{B2})$$

Definition B.2 (Convex set). A set $\Omega \subset \mathbb{R}^d$ is convex if and only if it contains all its line segments, i.e. for every $x, y \in \Omega$, $[x, y] \subset \Omega$.

Definition B.3. Let $f : \mathbb{R}^d \rightarrow (-\infty, +\infty]$ be a function. The *domain* of f is $\text{dom } f := \{x \in \mathbb{R}^d : f(x) < +\infty\}$. The function f is *proper* if $\text{dom } f \neq \emptyset$.

Definition B.4 (Subdifferential). The subdifferential of a function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ at x is the (possibly empty) set of slopes of all affine minorants of f that are exact at x :

$$\partial f(x) := \{u \in \mathbb{R}^d : \forall y \in \mathbb{R}^d, f(y) \geq f(x) + \langle u, y - x \rangle\}. \quad (\text{B3})$$

Fact B.5. A proper function $f : \mathbb{R}^d \rightarrow (-\infty, +\infty]$ is convex if and only if $\partial f(x) \neq \emptyset$ for all $x \in \text{ri}(\text{dom}(f))$, where ri denotes the relative interior [29, Def. 6.9].

Proof. The implication \Rightarrow can be found in [37, Thm. 23.4]. The converse property \Leftarrow can be found in [38, Sec. 1]. □

Fact B.6. Let $f : \mathcal{I} := [t_1, t_2] \rightarrow \mathbb{R}$ be a convex function where $t_1, t_2 \in \mathbb{R}$, $t_1 < t_2$. Then, $f'_+(t_1)$ and $f'_-(t_2)$ exist (but may be equal to $-\infty$ or $+\infty$, respectively).

Proof. Recall that the right derivative $f'_+(t_1)$ is defined as:

$$\lim_{\substack{t \rightarrow t_1^+ \\ t \in \mathcal{I}}} \frac{f(t) - f(t_1)}{t - t_1}. \quad (\text{B4})$$

From convexity of f , the function $t \mapsto \frac{f(t) - f(t_1)}{t - t_1}$ is non-decreasing, so it admits a limit from the right at t_1 (potentially equal to $-\infty$). The proof for t_2 is similar. \square

Fact B.7 ([39, Section 2.C]). $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is convex if and only if it is convex along every line segment: for all $x, y \in \mathbb{R}^d$, $t \mapsto f(x + t(y - x))$ is convex on $[0, 1]$.

Lemma B.8. Let $f : \mathcal{I} = [t_{\min}, t_{\max}] \rightarrow \mathbb{R}$ be continuous and piecewise convex, i.e. there exist an integer K and $-\infty \leq t_{\min} = t_0 < t_1 < \dots < t_K = t_{\max} \leq +\infty$ such that $f|_{[t_{i-1}, t_i]}$ is convex, for every $i \in \{1, \dots, K\}$. The function f is convex on \mathcal{I} if and only if, for each breakpoint t_i , $i \in \{1, \dots, K - 1\}$, one has

$$f'_-(t_i) \leq f'_+(t_i). \quad (\text{B5})$$

Proof. We prove the result when $K = 2$. If f has more than two pieces, we can iteratively gather groups of 2 consecutive pieces of 2 and conclude by induction.

Considering f with only two pieces, we denote $I_1 = [t_{\min}, t]$, $I_2 = [t, t_{\max}]$ those pieces, where $t_{\min} < t < t_{\max}$. Define $f_1, f_2 : \mathbb{R} \rightarrow]-\infty, +\infty]$ as:

$$f_1 : s \mapsto \begin{cases} f(s) & s \in I_1 \\ +\infty & \text{otherwise} \end{cases} \quad f_2 : s \mapsto \begin{cases} f(s) & s \in I_2 \\ +\infty & \text{otherwise} \end{cases} \quad (\text{B6})$$

The functions f_1 and f_2 are proper and convex on \mathbb{R} because their epigraphs are convex sets. Hence, from Fact B.5, we get that $\partial f(s) = \partial f_1(s) \neq \emptyset$ for all $s \in]t_{\min}, t[$ and $\partial f(s) = \partial f_2(s) \neq \emptyset$ for all $s \in]t, t_{\max}[$. Then, using again Fact B.5, we get that f is convex on $[t_{\min}, t_{\max}]$ if and only if $\partial f(s) \neq \emptyset$ for all $s \in]t_{\min}, t_{\max}[$. Since $\partial f(s) \neq \emptyset$ for every $s \in]t_{\min}, t[\cup]t, t_{\max}[$, the convexity condition is equivalent to $\partial f(t) \neq \emptyset$. Moreover we have that

$$\partial f(t) = \bigcap_{t' \in I_1 \cup I_2} \{u : f(t') \geq f(t) + u(t' - t)\} \quad (\text{B7})$$

$$= \left(\bigcap_{t' \in I_1} \{u : f(t') \geq f(t) + u(t' - t)\} \right) \cap \left(\bigcap_{t' \in I_2} \{u : f(t') \geq f(t) + u(t' - t)\} \right) \quad (\text{B8})$$

$$= \partial f_1(t) \cap \partial f_2(t) \quad (\text{B9})$$

$$= [f'_-(t), +\infty[\cap]-\infty, f'_+(t)] \quad (\text{B10})$$

This is non-empty if and only if $f'_-(t) \leq f'_+(t)$. \square

Fact B.9 ([29], Proposition 8.16). Let $(f_n)_{n \in \mathbb{N}}$ be a sequence of convex functions from \mathbb{R}^d to \mathbb{R} such that $(f_n)_{n \in \mathbb{N}}$ is pointwise convergent. Then $\lim_n f_n$ is convex.

Appendix C Proof of Proposition 3.6

Scheme of the proof

- (i) \implies (ii) \implies (iii) \implies (iv) \implies (ii) \implies (i)

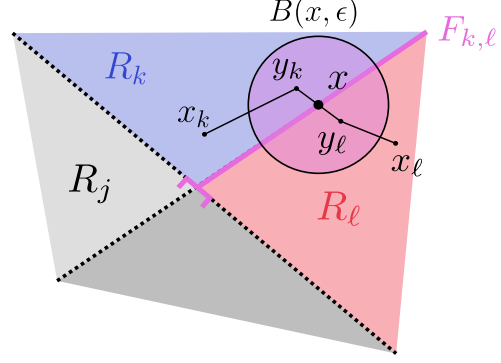


Fig. C1: The regions R_k and R_ℓ are neighbouring, but not R_j and R_ℓ . A frontier $F_{k,\ell}$ is the interior of the facet of dimension $d - 1$ which separates two distinct neighbouring regions. Convexity is a local property that has to be studied on a ball around x .

- (i) \implies (v)
- (v) \implies (iii)

Proof. (i) \implies (ii) Assume f convex on \mathbb{R}^d . From Fact B.7, f is convex on every line, so (ii) follows.

(ii) \implies (iii) is trivial.

(iii) \implies (iv) Consider two neighboring regions $R_k \sim R_\ell$. Choose $x \in F_{k,\ell}$, $v \in \mathbb{R}^d \setminus \text{span}(R_k \cap R_\ell - x)$ and $\epsilon > 0$ from (iii). Since $v \notin \text{span}(R_k \cap R_\ell - x)$ we have that $[x - \epsilon v, x + \epsilon v] \cap \text{int}(R_k) \neq \emptyset$ and $[x - \epsilon v, x + \epsilon v] \cap \text{int}(R_\ell) \neq \emptyset$ so we can pick $x_k \in [x - \epsilon v, x + \epsilon v] \cap \text{int}(R_k)$ and $x_\ell \in [x - \epsilon v, x + \epsilon v] \cap \text{int}(R_\ell)$. Then, we rewrite $f(x_k)$ and $f(x_\ell)$ with their respective affine expression

$$f(x_k) = \langle \nabla f(x_k), x_k \rangle + b_k \quad (\text{C11})$$

$$f(x_\ell) = \langle \nabla f(x_\ell), x_\ell \rangle + b_\ell \quad (\text{C12})$$

We define $\bar{t} \in (0, 1)$ the scalar for which $x = (1 - \bar{t})x_k + \bar{t}x_\ell$. Then, by convexity of $f|_{[x_k, x_\ell]}$, we have

- for $t \in (0, \bar{t}]$,

$$\langle \nabla f(x_k), (1 - t)x_k + tx_\ell \rangle + b_k \leq (1 - t) (\langle \nabla f(x_k), x_k \rangle + b_k) + t (\langle \nabla f(x_\ell), x_\ell \rangle + b_\ell) , \quad (\text{C13})$$

hence,

$$b_k - b_\ell \leq \langle \nabla f(x_\ell) - \nabla f(x_k), x_\ell \rangle . \quad (\text{C14})$$

- for $t \in [\bar{t}, 1)$,

$$\langle \nabla f(x_\ell), (1 - t)x_k + tx_\ell \rangle + b_\ell \leq (1 - t) (\langle \nabla f(x_k), x_k \rangle + b_k) + t (\langle \nabla f(x_\ell), x_\ell \rangle + b_\ell) . \quad (\text{C15})$$

hence,

$$b_\ell - b_k \leq -\langle \nabla f(x_\ell) - \nabla f(x_k), x_k \rangle . \quad (\text{C16})$$

By summing (C14) and (C16), we recover $\langle \nabla f(x_\ell) - \nabla f(x_k), x_\ell - x_k \rangle \geq 0$.

(iv) \implies (ii): Take $x \in \mathcal{F}$ and $v \in \mathbb{R}^d$. By the very definition of \mathcal{F} there exist neighboring regions $R_k \sim R_\ell$ such that $x \in F_{k,\ell}$. Moreover, there is $\epsilon > 0$ such that $B(x, \epsilon) \subset R_k \cup R_\ell$. The existence of such a ball comes from the definition of a frontier (Definition 3.5) and is illustrated in Figure C1. We now distinguish two cases. In the case where $v \in \text{span}(R_k \cap R_\ell - x)$, then, for ϵ small enough, $[x - \epsilon v, x + \epsilon v] \subset F_{k,\ell}$ so $f|_{[x - \epsilon v, x + \epsilon v]}$ is affine, thus convex, as claimed. We now treat the case where $v \in \mathbb{R}^d \setminus \text{span}(R_k \cap R_\ell - x)$. This implies that the tuple $(y_k, y_\ell) := (x - \epsilon v, x + \epsilon v)$ either belongs to $\text{int}(R_k) \times \text{int}(R_\ell)$ or to $\text{int}(R_\ell) \times \text{int}(R_k)$. Assume, without loss of generality, that $y_k \in \text{int}(R_k)$ and $y_\ell \in \text{int}(R_\ell)$. From (iv), there exists $x_k \in \text{int}(R_k)$ and $x_\ell \in \text{int}(R_\ell)$ such that $\langle \nabla f(x_k) - \nabla f(x_\ell), x_k - x_\ell \rangle \geq 0$. See Figure C1. Let η be a normal vector to $R_k \cap R_\ell$ which points from R_k to R_ℓ . As $\dim(\text{span}(R_k \cap R_\ell - x)) = \dim(\text{aff}(R_k \cap R_\ell)) = d - 1$, we can complete η with $\zeta_1, \dots, \zeta_{d-1} \in \text{span}(R_k \cap R_\ell - x)$ to form a basis of \mathbb{R}^d . We write $x_k - x_\ell$ and $y_k - y_\ell$

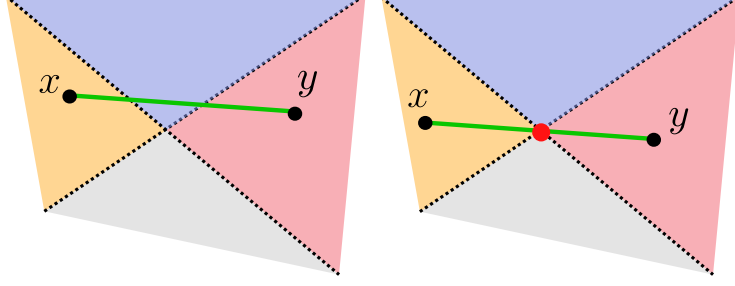


Fig. C2: Left: the segment does not cross any pathological points. Right: the red point is pathological: it belongs to more than two neighboring regions.

in this basis:

$$x_k - x_\ell = |\alpha^x| \eta + \sum_{i=1}^{d-1} \beta_i^x \zeta_i, \quad (\text{C17})$$

$$y_k - y_\ell = |\alpha^y| \eta + \sum_{i=1}^{d-1} \beta_i^y \zeta_i. \quad (\text{C18})$$

where $|\alpha^x| > 0$, $|\alpha^y| > 0$ since η points from R_k to R_ℓ and the considered points are interior to the regions. By the affine expressions (C11) of f on the regions R_k, R_ℓ and the continuity of f , we have

$$x \in R_k \cap R_\ell \subseteq \{y : \langle \nabla f(x_k), y \rangle + b_k = \langle \nabla f(x_\ell), y \rangle + b_\ell\} = \{y : \langle \nabla f(x_k) - \nabla f(x_\ell), y \rangle = b_\ell - b_k\}$$

hence $\text{span}(R_k \cap R_\ell - x) \subseteq \{y : \langle \nabla f(x_k) - \nabla f(x_\ell), y - x \rangle = 0\}$. As a result $\langle \nabla f(x_k) - \nabla f(x_\ell), \zeta_i \rangle = 0$ for every $i \in \{1, \dots, d-1\}$ (the same holds for $\nabla f(y_k) - \nabla f(y_\ell)$ since $y_k \in \text{int}(R_k)$ and $y_\ell \in \text{int}(R_\ell)$). Then,

$$\langle \nabla f(x_k) - \nabla f(x_\ell), y_k - y_\ell \rangle = \frac{|\alpha^y|}{|\alpha^x|} \langle \nabla f(x_k) - \nabla f(x_\ell), x_k - x_\ell \rangle \stackrel{(iv)}{\geq} 0. \quad (\text{C19})$$

Consider $t \in [0, 1] \mapsto g(t) := f((1-t)y_k + ty_\ell)$. Showing convexity of $f|_{[y_k, y_\ell]}$ amounts to showing that g is convex. We conclude using straightforward calculus and Lemma B.8 (with $K = 2$ pieces).

(ii) \implies (i): Consider any line segment in \mathbb{R}^d parametrized as $[x, y]$ for some $x, y \in \mathbb{R}^d$. To show (i), we need to prove that $f|_{[x, y]}$ is convex (Fact B.7). We distinguish the case where this line segment $[x, y]$ only contains non-pathological points, i.e. which are either in the interior of a region (where f is affine hence convex) or on a frontier (where (ii) ensures convexity of $f|_{[z-\epsilon(y-x), z+\epsilon(y-x)]}$), from the case where it contains pathological points that belong to more than two neighboring regions (see Figure C2). These pathological points belong to affine subspaces which are faces of the convex regions with dimension at most $d-2$ (otherwise, $d-1$ faces of the convex regions are the so-called frontiers).

- *Case 1: $[x, y]$ only has non-pathological points.* More formally, we assume here that $[x, y]$ is such that, for every $z \in [x, y]$, either $z \in \text{int}(R)$ for some region R from the considered compatible partition, either $z \in F$ for some frontier F between two neighboring regions. According to Lemma B.8, it suffices to study convexity at breakpoints $(z_i)_{i \in I}$ belonging to $(F_i)_{i \in I}$, frontiers between neighboring regions, using left and right derivatives. Then, (ii) gives convexity of $f|_{[z_i \pm \epsilon(y-x)]}$ for $\epsilon > 0$ small enough. So, $f|_{[x, y]}$ is locally convex around each breakpoint implying that it is convex overall.
- *Case 2: $[x, y]$ has pathological points.* First, pathological points belong to faces of the convex regions with dimension at most $d-2$. The number of such subspaces is finite due to the finite number of convex regions. We assume there exist pathological points $(z_i)_{i \in I}$ in $[x, y]$ which belong to affine subspaces $(A_i)_{i \in I}$ of dimension at most $d-2$. Note that either $[x, y]$ belongs to some face, making $I = [x, y]$ (with a slight abuse of notation), either it crosses a finite number of faces, with I being finite. We rewrite each affine subspace $A_i = z_i + V_i$ where V_i is the corresponding vectorial subspace and $z_i \in [x, y]$.

The strategy is to find a direction δ so as to perturb the segment $[x, y]$ along δ and avoid pathological subspaces A_i . To do so, we consider $\tilde{V}_i = V_i + \text{Vect}(y - x)$ which is of dimension at most $d - 1$. Then, we take $\delta \in \mathbb{R}^d \setminus \cup_{i \in I} \tilde{V}_i$ ($\delta \neq 0$). Now, we built a new segment $s_t := [x + t\delta, y + t\delta]$ for $t > 0$. We show that s_t does not contain any of the initial pathological points of $[x, y]$. Consider $z \in s_t$. The following equivalences hold: $z \in A_i \iff \exists v_i \in V_i : z = z_i + v_i \iff \exists \epsilon, \epsilon' > 0 : (1 - \epsilon)x + \epsilon y + t\delta = (1 - \epsilon')x + \epsilon'y + v_i \iff t\delta = v_i + (\epsilon' - \epsilon)(y - x)$. As $t\delta \notin \tilde{V}_i$ for $t > 0$ and $v_i + (\epsilon' - \epsilon)(y - x) \in \tilde{V}_i$, if $z \in s_t$ with $t > 0$, then by contradiction $z \notin A_i$. So, for every $t > 0$, for every $i \in I$, $s_t \cap A_i = \emptyset$. Since the number of pathological regions is finite, there exists $\tau > 0$ such that for every $0 < t < \tau$, $f_{|[x+t\delta, y+t\delta]}$ does not contain any pathological points: according to Case 1, $f_{|[x+t\delta, y+t\delta]}$ is convex. Then $f_{|[x, y]}$ is the pointwise limit, when $t \rightarrow 0$ of the convex functions $f_{|[x+t\delta, y+t\delta]}$ so, by Fact B.9 it is convex.

(i) \implies (v). Assume the function f is convex and consider $x \in \mathcal{F}$, that is $x \in F_{k, \ell}$ for some neighboring regions $R_k \sim R_\ell$. By definition of a frontier, there exists a neighborhood $\mathcal{N} \subset R_k \cup R_\ell$. Take $x_1, x_2 \in \mathcal{N} \setminus \mathcal{F}$. Then, $x_1 \in \text{int}(R_k) \cup \text{int}(R_\ell)$ (resp. x_2) so f is differentiable at x_1 (resp. x_2): the inequality follows from the characterization of convexity [29, Example 20.3] (see also Equation (3)).

(v) \implies (iii) Let $R_k \sim R_\ell$, $x \in F_{k, \ell}$, and $v \in \mathbb{R}^d \setminus \text{span}(R_k \cap R_\ell - x)$. Consider the neighborhood \mathcal{N} given by (v). Then there exists $\epsilon > 0$ such that points $x \pm \epsilon v$ are in \mathcal{N} . Besides, $F_{k, \ell} \cap \mathcal{N} \subset x + \text{span}(R_k \cap R_\ell - x)$ so $x \pm \epsilon v \notin F_{k, \ell}$. Then, the inequality in (v) implies increasing partial derivatives of $f_{|[x - \epsilon v, x + \epsilon v]}$ which proves its convexity (Lemma B.8). \square

Appendix D Proof of path-lifting factorization

Proof of Lemma 6.2 Consider a path $p = p_0 \rightarrow \dots \rightarrow p_m \in \mathcal{P}$ on the initial graph G such that $p_0 \in N_{\text{in}}$ and $p_i = \nu$ for some i . Considering the paths $q := p_0 \rightarrow \dots \rightarrow p_i$ and $r := p_i \rightarrow \dots \rightarrow p_m$, one can easily check that $q \in \mathcal{P}^{\rightarrow \nu}$, $r \in \mathcal{P}^{\nu \rightarrow}$. Moreover, by definition of the path-lifting $\phi_p(\theta)$, we have

$$\Phi_p = \phi_p(\theta) = \phi_q(\theta)\phi_r(\theta) = \Phi_q^{\rightarrow \nu}\Phi_r^{\nu \rightarrow}.$$

To prove a similar property for the activations, first observe that since ν is an output neuron of $G^{\rightarrow \nu}$, its activation on this DAG ReLU network (but not necessarily on the original one with G) is $a_\nu^{\rightarrow \nu}(x, \theta) = 1$ for every x, θ , by convention, while for every other neuron μ in $G^{\rightarrow \nu}$, $a_\mu^{\rightarrow \nu}(x, \theta) = a_\mu(x, \theta)$. As a consequence

$$a_q^{\rightarrow \nu} \stackrel{(16)}{=} \prod_{j=0}^i a_{q_j}^{\rightarrow \nu} = \prod_{j=0}^{i-1} a_{q_j} \tag{D20}$$

$$\begin{aligned} \text{so that } a_p &\stackrel{(16)}{=} \prod_{j=0}^m a_{p_j} \\ &= \left(\prod_{j=0}^{i-1} a_{p_j} \right) a_\nu \left(\prod_{j=i+1}^m a_{p_j} \right) \\ &\stackrel{(21) \& (D20)}{=} a_q^{\rightarrow \nu} a_\nu a_r^{\nu \rightarrow} \end{aligned} \tag{D21}$$

Since $\{p \in \mathcal{P} : p \ni \nu, p_0 \in N_{\text{in}}\}$ is in bijection with $\{q \in \mathcal{P}^{\rightarrow \nu} : q_0 \in N_{\text{in}}\} \times \mathcal{P}^{\nu \rightarrow}$, (23)–(24) follow using the definition of a Kronecker product (adapted to arbitrary cartesian product of index sets).

Finally, to establish (25), explicit each entry of $A_{[p: p_0 \in N_{\text{in}}, \nu \in N_{\text{in}}]}$ from (17)

$$A_{p, \mu} = \begin{cases} 0 & \text{if } p_0 \neq \mu \\ a_p = a_q^{\rightarrow \nu} a_\nu a_r^{\nu \rightarrow} & \text{otherwise} \end{cases} \tag{D22}$$

and similarly for $A_{\{q:q_0 \in N_{\text{in}}\}, N_{\text{in}}}^{\rightarrow \nu}$, again from (17)

$$A_{q, \mu}^{\rightarrow \nu} = \begin{cases} 0 & \text{if } q_0 \neq \mu \\ a_q^{\rightarrow \nu} & \text{otherwise} \end{cases} \quad (\text{D23})$$

Then, denoting the input neurons $N_{\text{in}} = \{\mu_1, \dots, \mu_d\}$ and enumerating $\{q : q_0 \in N_{\text{in}}\} = \{q^1, \dots, q^Q\}, \{r : r \in \mathcal{P}^{\nu \rightarrow}\} = \{r^1, \dots, r^R\}$ and $\{p : p \ni \nu, p_0 \in N_{\text{in}}\} = \{p^1, \dots, p^{QR}\}$:

$$A_{\left[\begin{smallmatrix} p: p \ni \nu \\ p_0 \in N_{\text{in}}, N_{\text{in}} \end{smallmatrix} \right]} \stackrel{(\text{D22})}{=} \begin{pmatrix} \delta_{p_0^1 = \mu_1} a_{q^1}^{\rightarrow \nu} a_\nu a_{r^1}^{\nu \rightarrow} & \dots & \delta_{p_0^1 = \mu_d} a_{q^1}^{\rightarrow \nu} a_\nu a_{r^1}^{\nu \rightarrow} \\ \vdots & \ddots & \vdots \\ \delta_{p_0^R = \mu_1} a_{q^1}^{\rightarrow \nu} a_\nu a_{r^R}^{\nu \rightarrow} & \dots & \delta_{p_0^R = \mu_d} a_{q^1}^{\rightarrow \nu} a_\nu a_{r^R}^{\nu \rightarrow} \\ \delta_{p_0^{R+1} = \mu_1} a_{q^2}^{\rightarrow \nu} a_\nu a_{r^1}^{\nu \rightarrow} & \dots & \delta_{p_0^{R+1} = \mu_d} a_{q^2}^{\rightarrow \nu} a_\nu a_{r^1}^{\nu \rightarrow} \\ \vdots & \ddots & \vdots \\ \delta_{p_0^{QR} = \mu_1} a_{q^Q}^{\rightarrow \nu} a_\nu a_{r^R}^{\nu \rightarrow} & \dots & \delta_{p_0^{QR} = \mu_d} a_{q^Q}^{\rightarrow \nu} a_\nu a_{r^R}^{\nu \rightarrow} \end{pmatrix} \quad (\text{D24})$$

$$= \begin{pmatrix} \delta_{q_0^1 = \mu_1} a_{q^1}^{\rightarrow \nu} a_\nu a_{r^1}^{\nu \rightarrow} & \dots & \delta_{q_0^1 = \mu_d} a_{q^1}^{\rightarrow \nu} a_\nu a_{r^1}^{\nu \rightarrow} \\ \vdots & \ddots & \vdots \\ \delta_{q_0^1 = \mu_1} a_{q^1}^{\rightarrow \nu} a_\nu a_{r^R}^{\nu \rightarrow} & \dots & \delta_{q_0^1 = \mu_d} a_{q^1}^{\rightarrow \nu} a_\nu a_{r^R}^{\nu \rightarrow} \\ \delta_{q_0^2 = \mu_1} a_{q^2}^{\rightarrow \nu} a_\nu a_{r^1}^{\nu \rightarrow} & \dots & \delta_{q_0^2 = \mu_d} a_{q^2}^{\rightarrow \nu} a_\nu a_{r^1}^{\nu \rightarrow} \\ \vdots & \ddots & \vdots \\ \delta_{q_0^Q = \mu_1} a_{q^Q}^{\rightarrow \nu} a_\nu a_{r^R}^{\nu \rightarrow} & \dots & \delta_{q_0^Q = \mu_d} a_{q^Q}^{\rightarrow \nu} a_\nu a_{r^R}^{\nu \rightarrow} \end{pmatrix} \quad (\text{D25})$$

$$= \begin{pmatrix} \delta_{q_0^1 = \mu_1} a_{q^1}^{\rightarrow \nu} a_\nu a^{\nu \rightarrow} & \dots & \delta_{q_0^1 = \mu_d} a_{q^1}^{\rightarrow \nu} a_\nu a^{\nu \rightarrow} \\ \vdots & \ddots & \vdots \\ \delta_{q_0^Q = \mu_1} a_{q^Q}^{\rightarrow \nu} a_\nu a^{\nu \rightarrow} & \dots & \delta_{q_0^Q = \mu_d} a_{q^Q}^{\rightarrow \nu} a_\nu a^{\nu \rightarrow} \end{pmatrix} \quad (\text{D26})$$

$$\stackrel{(\text{D23})}{=} a_\nu \begin{pmatrix} A_{q^1, \mu_1}^{\rightarrow \nu} a^{\nu \rightarrow} & \dots & A_{q^1, \mu_d}^{\rightarrow \nu} a^{\nu \rightarrow} \\ \vdots & \ddots & \vdots \\ A_{q^Q, \mu_1}^{\rightarrow \nu} a^{\nu \rightarrow} & \dots & A_{q^Q, \mu_d}^{\rightarrow \nu} a^{\nu \rightarrow} \end{pmatrix} \quad (\text{D27})$$

$$= a_\nu \left(A_{[q:q_0 \in N_{\text{in}}, N_{\text{in}}]}^{\rightarrow \nu} \otimes a^{\nu \rightarrow} \right) \quad (\text{D28})$$

□

Appendix E Proof of Lemma 6.8

We begin with the proof for DAG ReLU networks *where all hidden neurons are equipped with the ReLU activation function*. Then we prove that Lemma 6.8 (and therefore Proposition 6.3, Theorem 7.4 and their consequences) extend to the full DAG framework of Gonon et al. [31] which also contains pooling neurons in addition to ReLU neurons. We however restrict such pooling to *max-pooling* neurons, instead of the general “*k*-max-pooling” considered by Gonon et al. [31], since those also include *min-pooling* neurons, whose relation to convexity is somewhat opposite to max-pooling neurons.

E.1 The case of ReLU hidden neurons

Since $\nu \in H$ is isolated we can pick $x \in \mathfrak{X}_\nu \neq \emptyset$ and a ball $B(x, \epsilon)$ as in Definition 4.1. By definition of $B(x, \epsilon)$, the functions $A^{\rightarrow \nu}(\cdot, \theta)$ and $a^{\nu \rightarrow}(\cdot, \theta)$ are constant on this ball (they only involve products with neuron activations $a_\mu(x, \theta)$, $\mu \neq \nu$), and we denote $A^{\rightarrow \nu}$, $a^{\nu \rightarrow}$ their respective values. For the same

reason, $z_\nu(\cdot, \theta)$ is affine on $B(x, \epsilon)$ and $a_\nu(\cdot, \theta)$ takes two values. As a consequence, on this ball, the path-activation matrix $A(\cdot, \theta)$ takes exactly two values. This implies the existence of x^+, x^- in $B(x, \epsilon)$ such that $a_\nu(x^+, \theta) = 1$, $a_\nu(x^-, \theta) = 0$, and $A(\cdot, \theta)$ is locally constant around x^+ and x^- . As a consequence, f_θ is differentiable at x^+ and x^- establishing the first point. Now consider *any* two such points $x^+, x^- \in B(x, \epsilon)$ and denote $A^+ := A(x^+, \theta)$, $A^- := A(x^-, \theta)$. Since $a_\nu(x^+, \theta) = 1$ and $a_\nu(x^-, \theta) = 0$ we have $z_\nu(x^+, \theta) > 0$ and $z_\nu(x^-, \theta) \leq 0$. By (19) we have (we omit the dependency of $\Phi^{\rightarrow\nu}(\theta)$ in θ for brevity)

$$\left\langle \Phi^{\rightarrow\nu}, A^{\rightarrow\nu} \begin{pmatrix} x^+ \\ 1 \end{pmatrix} \right\rangle = z_\nu(x^+, \theta) > 0 \quad (\text{E29})$$

$$\left\langle \Phi^{\rightarrow\nu}, A^{\rightarrow\nu} \begin{pmatrix} x^- \\ 1 \end{pmatrix} \right\rangle = z_\nu(x^-, \theta) \leq 0. \quad (\text{E30})$$

Denote $u^+ = \nabla f_\theta(x^+)$ (resp. $u^- = \nabla f_\theta(x^-)$) the two slopes of f_θ given by A^+ (resp. A^-). To get the slope of f_θ from the expression $\langle \Phi, A(x, 1)^\top \rangle$ in (18), we need to extract the block matrix corresponding to the paths which start with inputs neurons N_{in} (recall that the remaining rows/column of A collect the biases of the network)

$$(u^+ - u^-)^\top = \Phi_{[p:p_0 \in N_{\text{in}}]}^\top (A^+ - A^-)_{[p:p_0 \in N_{\text{in}}, N_{\text{in}}]}.$$

Because the only entries that change between A^+ and A^- are the ones corresponding to paths $p \in \mathcal{P}$ which contain ν , it holds

$$(u^+ - u^-)^\top = \Phi_{[p:p_0 \in N_{\text{in}}]}^\top \begin{matrix} p \ni \nu \\ p_0 \in N_{\text{in}} \end{matrix} (A^+ - A^-)_{[p:p_0 \in N_{\text{in}}, N_{\text{in}}]}.$$

Moreover, $a_p(x', \theta) = 0$ for all paths p which contain ν and all x' such that $z_\nu(x', \theta) \leq 0$, so $A_{[p:p \ni \nu, N_{\text{in}}]}^- = \mathbf{0}$, leading to

$$(u^+ - u^-)^\top = \Phi_{[p:p_0 \in N_{\text{in}}]}^\top A_{[p:p_0 \in N_{\text{in}}, N_{\text{in}}]}^+ \quad (\text{E31})$$

Subtracting (E30) to (E29) similarly yields:

$$[\Phi_{[p:p_0 \in N_{\text{in}}]}^{\rightarrow\nu}]^\top A_{[p:p_0 \in N_{\text{in}}, N_{\text{in}}]}^{\rightarrow\nu} (x^+ - x^-) > 0. \quad (\text{E32})$$

By Lemma 6.2, using that $A^{\rightarrow\nu}(\cdot, \theta) = A^{\rightarrow\nu}$ and $a^{\nu\rightarrow}(\cdot, \theta) = a^{\nu\rightarrow}$ on $B(x, \epsilon)$, and $a_\nu(x^+, \theta) = 1$ (cf (E29)) the quantity $\Delta u := (u^+ - u^-)^\top$ from (E31) rewrites as

$$\begin{aligned} \Delta u &= \left(\Phi_{[q:q_0 \in N_{\text{in}}]}^{\rightarrow\nu} \otimes \Phi^{\nu\rightarrow} \right)^\top \\ &\quad \times \left(A_{[q:q_0 \in N_{\text{in}}, N_{\text{in}}]}^{\rightarrow\nu} \otimes a^{\nu\rightarrow} \right) \\ &= \left([\Phi_{[q:q_0 \in N_{\text{in}}]}^{\rightarrow\nu}]^\top A_{[q:q_0 \in N_{\text{in}}, N_{\text{in}}]}^{\rightarrow\nu} \right) \\ &\quad \otimes \underbrace{([\Phi^{\nu\rightarrow}]^\top a^{\nu\rightarrow})}_{\text{scalar}} \\ &= ([\Phi^{\nu\rightarrow}]^\top a^{\nu\rightarrow}) \\ &\quad \times \left([\Phi_{[q:q_0 \in N_{\text{in}}]}^{\rightarrow\nu}]^\top A_{[q:q_0 \in N_{\text{in}}, N_{\text{in}}]}^{\rightarrow\nu} \right). \end{aligned}$$

It follows

$$\begin{aligned} \langle u^+ - u^-, x^+ - x^- \rangle &= ([\Phi^{\nu \rightarrow}]^\top a^{\nu \rightarrow}) \\ &\times \underbrace{\left([\Phi_{[q:q_0 \in N_{\text{in}}]}^{\rightarrow \nu}]^\top A_{[q:q_0 \in N_{\text{in}}, N_{\text{in}}]}^{\rightarrow \nu} \right)}_{>0 \text{ by (E32)}} (x^+ - x^-). \end{aligned}$$

As a result

$$\begin{aligned} \langle u^+ - u^-, x^+ - x^- \rangle &\geq 0 \\ \iff \langle a^{\nu \rightarrow}, \Phi^{\nu \rightarrow}(\theta) \rangle &\geq 0. \end{aligned} \quad (\text{E33})$$

E.2 DAG ReLU networks with max-pooling activation function

The max-pooling function

$$\text{maxpool}(x) := \max_{i \in \{1, \dots, d\}} x_i \quad (\text{E34})$$

returns the largest coordinate of $x \in \mathbb{R}^d$. This section explains how to incorporate max-pooling activations in the DAG neural network model considered in this paper and recover the necessary conditions for convexity of the network given in Proposition 6.3. Consider a network with architecture G described by the tuple (N, E) and parameters θ . The key is to judiciously modify the definition of neuron activation and preactivation, and to introduce *edge* activations [31]. Consider a neuron ν with max-pooling activation and denote $\text{ant}(\nu) := \{\mu \in N : (\mu, \nu) \in E\}$. The post-activation of the neuron ν for a given input point $x \in \mathbb{R}^d$ writes as (there is no bias for such a neuron)

$$\nu(x) := \max_{\mu \in \text{ant}(\nu)} \theta^{\mu \rightarrow \nu} \mu(x) \quad (\text{E35})$$

$$= \langle (\theta^{\mu \rightarrow \nu} \mu(x))_{\mu \in \text{ant}(\nu)}, e_{\mu^*} \rangle \quad \text{with } \mu^* := \operatorname{argmax}_{\mu \in \text{ant}(\nu)} \theta^{\mu \rightarrow \nu} \mu(x), \quad (\text{E36})$$

where e_{μ^*} in Equation (E36) is the binary vector with value 1 for index $\mu = \mu^*$ and 0 otherwise⁹. It is then natural to define the activation of a max-pooling neuron ν as the binary *vector* $\mathbf{a}_\nu(x, \theta) \in \{0, 1\}^{\text{ant}(\nu)}$ with

$$[\mathbf{a}_\nu(x, \theta)]_\mu := \begin{cases} 1 & \text{if } \mu = \operatorname{argmax}_{\mu' \in \text{ant}(\nu)} \theta^{\mu' \rightarrow \nu} \mu'(x), \\ 0 & \text{otherwise,} \end{cases} \quad (\text{E37})$$

and the pre-activation of ν as the *vector* $z_\nu(x) := (\theta^{\mu \rightarrow \nu} \mu(x))_{\mu \in \text{ant}(\nu)} \in \mathbb{R}^{\text{ant}(\nu)}$. The post-activation of ν rewrites

$$\nu(x) = \langle z_\nu(x), \mathbf{a}_\nu(x, \theta) \rangle \in \mathbb{R}. \quad (\text{E38})$$

The same holds true with ReLU neurons with the convention $\mathbf{a}_\nu(\theta, x) := a_\nu(\theta, x) \in \{0, 1\}$ for such neurons.

Following Gonon et al. [31], we also define the activation of an edge $(\mu, \nu) \in E$ as

$$a_{\mu \rightarrow \nu}(x, \theta) := \begin{cases} a_\nu(x, \theta) & \text{if } \nu \text{ is a ReLU neuron} \\ [a_\nu(x, \theta)]_\mu & \text{if } \nu \text{ is a maxpool neuron } \in \{0, 1\} \\ 1 & \text{if } \nu \text{ is a linear neuron} \end{cases} \quad (\text{E39})$$

These extended definitions enable us to re-define the activation of a path $p := p_0 \rightarrow p_1 \rightarrow \dots \rightarrow p_m$ as $a_p(x, \theta) := a_{p_0}(x, \theta) \prod_{i=1}^m a_{p_{i-1} \rightarrow p_i}(x, \theta)$ (with the convention that $a_{p_0}(\theta, x) = 1$ when p_0 is a max-pooling neuron, an input neuron, or a linear neuron, see [31, Definition A.3]) : we describe the activation of a path using the edge-wise activations instead of the neuron-wise activations. On the extracted sub-graph

⁹we arbitrarily number the neurons in $\text{ant}(\nu)$ so that if the argmax contains more than one neuron, we systematically pick the one indexed by the smallest number

$G^{\nu \rightarrow}$, the path-activation of $r := r_0 \rightarrow \dots \rightarrow r_m \in \mathcal{P}^{\nu \rightarrow}$ (where $r_0 = \nu$ by definition of $\mathcal{P}^{\nu \rightarrow}$) is defined as $a_r^{\nu \rightarrow}(x, \theta) := \prod_{i=1}^m a_{r_{i-1} \rightarrow r_i}(x, \theta)$ (without multiplication by $a_{r_0}(\theta, x)$). Note that when only considering ReLU neurons, the modified definitions provided here yield exactly the same path-activations as the ones given in the core of the paper.

Remark E.1. Remember that when ν is equipped with the ReLU, considering the subnetwork with architecture $G^{\rightarrow \nu}$ we replace the activation of ν by a linear activation since it is the output of this subnetwork. When ν is a maxpool neuron, we preserve its activation function. Whether the proof for the ReLU case can be adapted to fit this framework is left to future work.

Besides, since we follow the framework from [31] which already includes max-pooling activations, all relevant results from their work – e.g. Equation (18) – still hold in the framework of this section.

We provide below a (slightly) modified path-activations factorization which aligns with the edge-wise description of path-activations instead of the neuron-wise one.

Lemma E.2 (Path-activations factorization – variant of Lemma 6.2). *Within the framework of this section, consider an arbitrary hidden neuron ν and any parameter θ . The restrictions of $\Phi = \Phi(\theta) \in \mathbb{R}^{\mathcal{P}}$ and $a = a(x, \theta) \in \{0, 1\}^{\mathcal{P}}$ to paths starting from an input neuron and “containing” an edge $\mu \rightarrow \nu$ (i.e. such that $p_i = \mu$ and $p_{i+1} = \nu$ for some i ; this is denoted $p \ni (\mu, \nu)$) satisfy*

$$\Phi \left[\begin{array}{l} p: p \ni (\mu, \nu) \\ p_0 \in N_{\text{in}} \end{array} \right] = \theta^{\mu \rightarrow \nu} \cdot \left(\Phi_{[q: q_0 \in N_{\text{in}}]}^{\rightarrow \mu} \otimes \Phi^{\nu \rightarrow} \right) \quad (\text{E40})$$

$$a \left[\begin{array}{l} p: p \ni (\mu, \nu) \\ p_0 \in N_{\text{in}} \end{array} \right] = a_{\mu \rightarrow \nu} \cdot \left(a_{[q: q_0 \in N_{\text{in}}]}^{\rightarrow \mu} \otimes a^{\nu \rightarrow} \right). \quad (\text{E41})$$

For the path-activation matrix, one has

$$A \left[\begin{array}{l} p: p \ni (\mu, \nu), N_{\text{in}} \\ p_0 \in N_{\text{in}} \end{array} \right] = a_{\mu \rightarrow \nu} \cdot \left(A_{[q: q_0 \in N_{\text{in}}, N_{\text{in}}]}^{\rightarrow \mu} \otimes a^{\nu \rightarrow} \right). \quad (\text{E42})$$

Proof. The proof is exactly the same as the one done in Appendix D. We only detail below the factorization of a_p for a path $p = p_0 \rightarrow \dots \rightarrow p_m \in \mathcal{P}$ such that $p_0 \in N_{\text{in}}$ and $p_i = \mu$, $p_{i+1} = \nu$ for some i . We denote $q := p_0 \rightarrow \dots \rightarrow p_i \in \mathcal{P}^{\rightarrow \mu}$ and $r := p_{i+1} \rightarrow \dots \rightarrow p_m \in \mathcal{P}^{\nu \rightarrow}$.

$$a_q^{\rightarrow \mu} = \underbrace{a_{p_0}^{\rightarrow \mu}}_{=1 \text{ by convention since } p_0 \in N_{\text{in}}} \prod_{j=1}^i a_{p_{j-1} \rightarrow p_j}^{\rightarrow \mu} = \prod_{j=1}^i a_{p_{j-1} \rightarrow p_j}, \quad (\text{E43})$$

$$a_r^{\nu \rightarrow} = a_{p_{i+1} \rightarrow p_{i+2}} \cdots a_{p_{m-1} \rightarrow p_m}, \quad (\text{E44})$$

from which we have $a_p = a_q^{\rightarrow \mu} a_{\mu \rightarrow \nu} a_r^{\nu \rightarrow}$. \square

Now, we are almost equipped to prove a variant of Proposition 6.3 for networks which include max-pooling activations. First, we need to redefine what an isolated neuron when considering max-pooling activations. Recall that for a ReLU activation, the set $\mathcal{X}_\nu \subset \mathbb{R}^d$ is defined as input points for which there exists a neighborhood \mathcal{N} over which only the activation of ν changes. In other words, this means that over \mathcal{N} , a_ν takes *exactly two values*: 0 and 1, while a_μ for $\mu \neq \nu$ is constant.

This motivates the following definition of isolated neurons, which handles the maxpool case while being equivalent to Definition 4.1 in the case of ReLU neurons.

Definition E.3 (Isolated neurons – variant of Definition 4.1). Given a parameter θ , for each hidden neuron $\nu \in H := \cup_{\ell=1}^{L-1} N_\ell$, we define $\mathfrak{X}_\nu \subset \mathbb{R}^d$ as the set of input points for which in every small enough neighborhood, *only* the activation of neuron ν changes, and takes exactly two distinct values:

$$\mathfrak{X}_\nu := \left\{ x \in \mathbb{R}^d : \exists \epsilon_0 > 0 : \forall 0 < \epsilon \leq \epsilon_0, \begin{cases} \mathbf{a}_\mu(\cdot, \theta) \text{ is constant on } B(x, \epsilon) \\ \mathbf{a}_\nu(\cdot, \theta) \text{ takes exactly 2 values on } B(x, \epsilon) \end{cases} \forall \mu \neq \nu \right\}. \quad (\text{E45})$$

A hidden neuron ν is said to be *isolated* if $\mathfrak{X}_\nu \neq \emptyset$.

E.3 Extension Lemma 6.8 to cover maxpool activations.

We now prove an extension of Lemma 6.8 that also covers DAG ReLU networks including maxpool activations. We highlight in blue the subtle difference with the statement Lemma 6.8.

Lemma E.4 (Local convexity criterion for ReLU networks). *Consider a ReLU network described by a DAG and parametrized by θ which implements a CPWL function $f_\theta : \mathbb{R}^d \rightarrow \mathbb{R}$. Consider an isolated neuron $\nu \in H$ and $x \in \mathfrak{X}_\nu \neq \emptyset$. Denote $\mathbf{a}^+, \mathbf{a}^-$ the two distinct values of $\mathbf{a}_\nu(\cdot, \theta)$ from Definition E.3. There exists $\epsilon > 0$ such that*

1. there are $x^+, x^- \in B(x, \epsilon)$ such that f_θ is differentiable at x^+, x^- and $\mathbf{a}_\nu(x^+, \theta) = \mathbf{a}^+, \mathbf{a}_\nu(x^-, \theta) = \mathbf{a}^-$;
2. for any such pair x^+, x^- it holds

$$\begin{aligned} & \langle \nabla f_\theta(x^+) - \nabla f_\theta(x^-), x^+ - x^- \rangle \geq 0 \\ \iff & \langle \mathbf{a}^{\nu \rightarrow}, \Phi^{\nu \rightarrow}(\theta) \rangle \geq 0. \end{aligned} \quad (\text{E46})$$

Proof. Consider $\nu \in H$ an isolated neuron. In case of a ReLU neuron, $\mathbf{a}_\nu(\cdot, \theta) = a_\nu(\cdot, \theta)$, and without loss of generality $\mathbf{a}^+ = 1, \mathbf{a}^-$, so the result is a consequence of Lemma 6.8. Assume now that ν an isolated neuron with maxpool activation. We highlight the parts of the proof identical in spirit to the corresponding proof of Lemma 6.8.

[Unchanged] Pick $x \in \mathfrak{X}_\nu \neq \emptyset$ and a ball $B(x, \epsilon)$ as in Definition E.3. By definition of $B(x, \epsilon)$, $\mathbf{a}^{\nu \rightarrow}(\cdot, \theta)$ is constant on this ball (it only involves products with edge activations $\mathbf{a}^{\mu \rightarrow \mu'}, \mu' \neq \nu$).

[New] By definition of \mathfrak{X}_ν , \mathbf{a}_ν takes exactly two different values on $B(x, \epsilon)$; denote μ^+ and μ^- the two antecedents of ν which realize the max-pool on this ball. The only path-activations which change are for paths containing the edges $\mu^+ \rightarrow \nu$ and $\mu^- \rightarrow \nu$. Without loss of generality we assume that the arbitrary numbering of the antecedents of ν chosen to uniquely break ties in the arg max is such that μ^- is associated to the smallest number. This implies that as soon as $\mathbf{a}_\nu(x, \theta) = e_{\mu^+}$ the strict inequality

$$[z_\nu(x)]_{\mu^+} > [z_\nu(x)]_{\mu^-} \quad (\text{E47})$$

holds, while when $\mathbf{a}_\nu(x, \theta) = e_{\mu^-}$ we have

$$[z_\nu(x)]_{\mu^-} \geq [z_\nu(x)]_{\mu^+}. \quad (\text{E48})$$

[Unchanged] The path-activation matrix $A(\cdot, \theta)$ takes exactly two values on $B(x, \epsilon)$.

[New] This implies the existence of x^+ (resp. x^-) in $B(x, \epsilon)$ such that $[\mathbf{a}_\nu(x^+, \theta)]_{\mu^+} = 1, [\mathbf{a}_\nu(x^+, \theta)]_{\mu^-} = 0$ (resp. $[\mathbf{a}_\nu(x^-, \theta)]_{\mu^+} = 0, [\mathbf{a}_\nu(x^-, \theta)]_{\mu^-} = 1$).

[Unchanged] As $A(\cdot, \theta)$ is locally constant around x^+ and x^- , f_θ is differentiable at x^+ and x^- establishing the first point. Now, consider any two such points $x^+, x^- \in B(x, \epsilon)$ and denote $A^+ := A(x^+, \theta)$, $A^- := A(x^-, \theta)$. Denote $u^+ = \nabla f_\theta(x^+)$ (resp. $u^- = \nabla f_\theta(x^-)$) the two slopes of f_θ given by A^+ (resp. A^-). To get the slope of f_θ from the expression $\langle \Phi, A \begin{pmatrix} x \\ 1 \end{pmatrix} \rangle$ in Equation (18), we need to extract

the block matrix corresponding to the paths which start with input neurons N_{in} (recall that the remaining rows/columns of A collect the biases of the network)

$$(u^+ - u^-)^\top = \Phi_{[p:p_0 \in N_{\text{in}}]}^\top (A^+ - A^-)_{[p:p_0 \in N_{\text{in}}, N_{\text{in}}]}. \quad (\text{E49})$$

[New] Because the only entries that change between A^+ and A^- are the ones corresponding to paths $p \in \mathcal{P}$ which contain the edge $\mu^+ \rightarrow \nu$ or $\mu^- \rightarrow \nu$, it holds

$$\Delta u := (u^+ - u^-)^\top = \left(\Phi_{[p:p \ni (\mu^+, \nu), p_0 \in N_{\text{in}}]}^\top \quad \Phi_{[p:p \ni (\mu^-, \nu), p_0 \in N_{\text{in}}]}^\top \right) \begin{pmatrix} (A^+ - A^-)_{[p:p \ni (\mu^+, \nu), N_{\text{in}}]} \\ (A^+ - A^-)_{[p:p \ni (\mu^-, \nu), N_{\text{in}}]} \end{pmatrix}. \quad (\text{E50})$$

Using the factorization of Lemma E.2

$$A_{p:p \ni (\mu^+, \nu), p_0 \in N_{\text{in}}}^+ = \underbrace{a_{\mu^+ \rightarrow \nu}(x^+, \theta)}_{=1} A_{[q:q_0 \in N_{\text{in}}, N_{\text{in}}]}^{\rightarrow \mu^+} \otimes a^{\nu \rightarrow}, \quad (\text{E51})$$

$$A_{p:p \ni (\mu^-, \nu), p_0 \in N_{\text{in}}}^+ = \underbrace{a_{\mu^- \rightarrow \nu}(x^+, \theta)}_{=0} A_{[q:q_0 \in N_{\text{in}}, N_{\text{in}}]}^{\rightarrow \mu^-} \otimes a^{\nu \rightarrow} = \mathbf{0}, \quad (\text{E52})$$

$$A_{p:p \ni (\mu^+, \nu), p_0 \in N_{\text{in}}}^- = \underbrace{a_{\mu^+ \rightarrow \nu}(x^-, \theta)}_{=0} A_{[q:q_0 \in N_{\text{in}}, N_{\text{in}}]}^{\rightarrow \mu^+} \otimes a^{\nu \rightarrow} = \mathbf{0}, \quad (\text{E53})$$

$$A_{p:p \ni (\mu^-, \nu), p_0 \in N_{\text{in}}}^- = \underbrace{a_{\mu^- \rightarrow \nu}(x^-, \theta)}_{=1} A_{[q:q_0 \in N_{\text{in}}, N_{\text{in}}]}^{\rightarrow \mu^-} \otimes a^{\nu \rightarrow}, \quad (\text{E54})$$

from which Δu rewrites as

$$\Delta u = \left(\Phi_{[p:p \ni (\mu^+, \nu), p_0 \in N_{\text{in}}]}^\top \quad \Phi_{[p:p \ni (\mu^-, \nu), p_0 \in N_{\text{in}}]}^\top \right) \begin{pmatrix} A_{[q:q_0 \in N_{\text{in}}, N_{\text{in}}]}^{\rightarrow \mu^+} \otimes a^{\nu \rightarrow} \\ -A_{[q:q_0 \in N_{\text{in}}, N_{\text{in}}]}^{\rightarrow \mu^-} \otimes a^{\nu \rightarrow} \end{pmatrix} \quad (\text{E55})$$

$$= \left(\Phi_{[p:p \ni (\mu^+, \nu), p_0 \in N_{\text{in}}]}^\top (A_{[q:q_0 \in N_{\text{in}}, N_{\text{in}}]}^{\rightarrow \mu^+} \otimes a^{\nu \rightarrow}) \right) - \left(\Phi_{[p:p \ni (\mu^-, \nu), p_0 \in N_{\text{in}}]}^\top (A_{[q:q_0 \in N_{\text{in}}, N_{\text{in}}]}^{\rightarrow \mu^-} \otimes a^{\nu \rightarrow}) \right) \quad (\text{E56})$$

The first term factorizes as

$$\Phi_{[p:p \ni (\mu^+, \nu), p_0 \in N_{\text{in}}]}^\top (A_{[q:q_0 \in N_{\text{in}}, N_{\text{in}}]}^{\rightarrow \mu^+} \otimes a^{\nu \rightarrow}) = (\theta^{\mu^+ \rightarrow \nu} \Phi_{[q:q_0 \in N_{\text{in}}]}^{\rightarrow \mu^+} \otimes \Phi^{\nu \rightarrow})^\top (A_{[q:q_0 \in N_{\text{in}}, N_{\text{in}}]}^{\rightarrow \mu^+} \otimes a^{\nu \rightarrow}) \quad (\text{E57})$$

$$= \theta^{\mu^+ \rightarrow \nu} ((\Phi_{[q:q_0 \in N_{\text{in}}]}^{\rightarrow \mu^+})^\top A_{[q:q_0 \in N_{\text{in}}, N_{\text{in}}]}^{\rightarrow \mu^+}) \otimes \underbrace{((\Phi^{\nu \rightarrow})^\top a^{\nu \rightarrow})}_{\text{scalar}} \quad (\text{E58})$$

$$= ((\Phi^{\nu \rightarrow})^\top a^{\nu \rightarrow}) \left(\theta^{\mu^+ \rightarrow \nu} ((\Phi_{[q:q_0 \in N_{\text{in}}]}^{\rightarrow \mu^+})^\top A_{[q:q_0 \in N_{\text{in}}, N_{\text{in}}]}^{\rightarrow \mu^+}) \right), \quad (\text{E59})$$

while similarly the second term rewrites

$$\Phi_{[p:p \ni (\mu^-, \nu), p_0 \in N_{\text{in}}]}^\top (A_{[q:q_0 \in N_{\text{in}}, N_{\text{in}}]}^{\rightarrow \mu^-} \otimes a^{\nu \rightarrow}) = (\theta^{\mu^- \rightarrow \nu} \Phi_{[q:q_0 \in N_{\text{in}}]}^{\rightarrow \mu^-} \otimes \Phi^{\nu \rightarrow})^\top (A_{[q:q_0 \in N_{\text{in}}, N_{\text{in}}]}^{\rightarrow \mu^-} \otimes a^{\nu \rightarrow}) \quad (\text{E60})$$

$$= \theta^{\mu^- \rightarrow \nu} ((\Phi_{[q:q_0 \in N_{\text{in}}]}^{\rightarrow \mu^-})^\top A_{[q:q_0 \in N_{\text{in}}, N_{\text{in}}]}^{\rightarrow \mu^-}) \otimes \underbrace{((\Phi^{\nu \rightarrow})^\top a^{\nu \rightarrow})}_{\text{scalar}} \quad (\text{E61})$$

$$= ((\Phi^{\nu \rightarrow})^\top a^{\nu \rightarrow}) \left(\theta^{\mu^- \rightarrow \nu} ((\Phi_{[q:q_0 \in N_{\text{in}}]}^{\rightarrow \mu^-})^\top A_{[q:q_0 \in N_{\text{in}}, N_{\text{in}}]}^{\rightarrow \mu^-}) \right), \quad (\text{E62})$$

leading to

$$\Delta u = \langle a^{\nu \rightarrow}, \Phi^{\nu \rightarrow}(\theta) \rangle \left[\left(\theta^{\mu^+ \rightarrow \nu} (\Phi_{[q:q_0 \in N_{\text{in}}]}^{\rightarrow \mu^+})^\top A_{[q:q_0 \in N_{\text{in}}, N_{\text{in}}]}^{\rightarrow \mu^+} \right) - \left(\theta^{\mu^- \rightarrow \nu} (\Phi_{[q:q_0 \in N_{\text{in}}]}^{\rightarrow \mu^-})^\top A_{[q:q_0 \in N_{\text{in}}, N_{\text{in}}]}^{\rightarrow \mu^-} \right) \right] \quad (\text{E63})$$

Then we have

$$\langle u^+ - u^-, x^+ - x^- \rangle = \langle a^{\nu \rightarrow}, \Phi^{\nu \rightarrow}(\theta) \rangle \left[\left(\theta^{\mu^+ \rightarrow \nu} (\Phi_{[q:q_0 \in N_{\text{in}}]}^{\rightarrow \mu^+})^\top A_{[q:q_0 \in N_{\text{in}}, N_{\text{in}}]}^{\rightarrow \mu^+} (x^+ - x^-) \right) - \left(\theta^{\mu^- \rightarrow \nu} (\Phi_{[q:q_0 \in N_{\text{in}}]}^{\rightarrow \mu^-})^\top A_{[q:q_0 \in N_{\text{in}}, N_{\text{in}}]}^{\rightarrow \mu^-} (x^+ - x^-) \right) \right] \quad (\text{E64})$$

To prove the equivalence (E46) it is thus enough to show that

$$\left[\left(\theta^{\mu^+ \rightarrow \nu} (\Phi_{[q:q_0 \in N_{\text{in}}]}^{\rightarrow \mu^+})^\top A_{[q:q_0 \in N_{\text{in}}, N_{\text{in}}]}^{\rightarrow \mu^+} (x^+ - x^-) \right) - \left(\theta^{\mu^- \rightarrow \nu} (\Phi_{[q:q_0 \in N_{\text{in}}]}^{\rightarrow \mu^-})^\top A_{[q:q_0 \in N_{\text{in}}, N_{\text{in}}]}^{\rightarrow \mu^-} (x^+ - x^-) \right) \right] > 0. \quad (\text{E65})$$

By definition of x^+ , x^- and (E47)-(E48), we have

$$[z_\nu(x^+, \theta)]_{\mu^+} > [z_\nu(x^+, \theta)]_{\mu^-}, \quad \text{and} \quad [z_\nu(x^-, \theta)]_{\mu^-} \geq [z_\nu(x^-, \theta)]_{\mu^+}$$

In other words,

$$\theta^{\mu^+ \rightarrow \nu} \mu^+(x^+) > \theta^{\mu^- \rightarrow \nu} \mu^-(x^+) \quad \text{and} \quad \theta^{\mu^- \rightarrow \nu} \mu^-(x^-) \geq \theta^{\mu^+ \rightarrow \nu} \mu^+(x^-) \quad (\text{E66})$$

which implies

$$\theta^{\mu^+ \rightarrow \nu} (\mu^+(x^+) - \mu^+(x^-)) - \theta^{\mu^- \rightarrow \nu} (\mu^-(x^+) - \mu^-(x^-)) > 0. \quad (\text{E67})$$

By Equation (18), here used on two neural networks which output neuron μ^+ (resp. μ^-) is *not* linear, the post-activation μ^+ and μ^- can be explicitated for any input x using the path-lifting as

$$\mu^+(x) = \left\langle \Phi^{\rightarrow \mu^+}, A_{[q:q_0 \in N_{\text{in}}, N_{\text{in}}]}^{\rightarrow \mu^+}(x, \theta) \begin{pmatrix} x \\ 1 \end{pmatrix} \right\rangle, \quad (\text{E68})$$

$$\mu^-(x) = \left\langle \Phi^{\rightarrow \mu^-}, A_{[q:q_0 \in N_{\text{in}}, N_{\text{in}}]}^{\rightarrow \mu^-}(x, \theta) \begin{pmatrix} x \\ 1 \end{pmatrix} \right\rangle, \quad (\text{E69})$$

so that, as $A_{[q:q_0 \in N_{\text{in}}, N_{\text{in}}]}^{\rightarrow \mu^+}(x^+, \theta) = A_{[q:q_0 \in N_{\text{in}}, N_{\text{in}}]}^{\rightarrow \mu^+}(x^-, \theta) = A_{[q:q_0 \in N_{\text{in}}, N_{\text{in}}]}^{\rightarrow \mu^+}$

$$\mu^+(x^+) - \mu^+(x^-) = \left\langle \Phi^{\rightarrow \mu^+}, A_{[q:q_0 \in N_{\text{in}}, N_{\text{in}}]}^{\rightarrow \mu^+} \begin{pmatrix} x^+ - x^- \\ 0 \end{pmatrix} \right\rangle \quad (\text{E70})$$

$$= (\Phi_{[q:q_0 \in N_{\text{in}}]}^{\rightarrow \mu^+})^\top A_{[q:q_0 \in N_{\text{in}}, N_{\text{in}}]}^{\rightarrow \mu^+} (x^+ - x^-). \quad (\text{E71})$$

Doing the same computation for μ^- , Equation (E65) is equivalent to

$$\theta^{\mu^+ \rightarrow \nu} (\mu^+(x^+) - \mu^+(x^-)) - \theta^{\mu^- \rightarrow \nu} (\mu^-(x^+) - \mu^-(x^-)) > 0.$$

which indeed holds, cf (E67). \square

Appendix F Extensive study of the 2D counter example and proof of Proposition 2.1

F.1 Proof of Proposition 2.1

Recall the CPWL function $f_{\text{EX}} : \mathbb{R}^2 \rightarrow \mathbb{R}$ is implemented by a ReLU network with 2 hidden layers and 2 neurons per layer (i.e. belonging to $\text{SkipMLP}((2, 2))$):

$$f_{\text{EX}} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = (1 \ 1) \text{ReLU} \left(\begin{pmatrix} -1 & 1 \\ 2 & 1 \end{pmatrix} \text{ReLU} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} + \begin{pmatrix} -1 \\ -0.5 \end{pmatrix} \right). \quad (\text{F72})$$

The network written above is not an ICNN because of the negative weight -1 in W_2 of the second hidden layer. Still, f_{θ} is convex: it can be easily checked using the sufficient conditions of Proposition 3.6. Plots of f_{EX} are displayed in Figure F3.

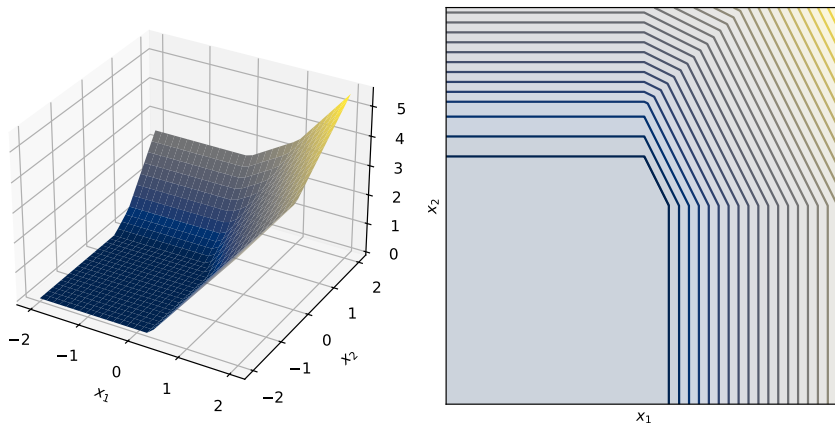


Fig. F3: Visual representations of f_{EX} . Left: 3D plot of f_{EX} . Right: Level lines of f_{EX} .

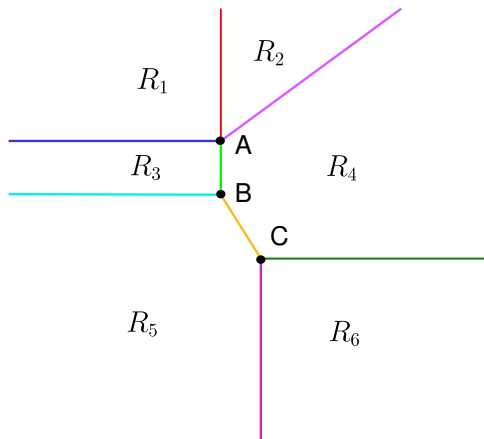


Fig. F4: Frontiers and regions of the convex 2D CPWL function

To prove Proposition 2.1, we adopt a reverse-engineering approach: we only use the knowledge of the landscape of the convex CPWL function $f_{\text{EX}} \in \text{CvxCPWL}$. The landscape displays 8 pieces of

hyperplanes (represented in color in Figure F4), which we name frontiers, and 6 corresponding regions $R_k, k \in \{1, \dots, 6\}$. Frontiers are denoted $F_{k,l}$ with $k, l \in \{1, \dots, 6\}$ and correspond to points at which f_{EX} is non-differentiable. By contradiction, we assume that f_{EX} can be implemented as an ICNN:

$$f_{\text{EX}}(x) = w_3^\top \text{ReLU}(W_2 \text{ReLU}(W_1 x + b_1) + V_2 x + b_2) + v_3^\top x + b_3, \quad (\text{F73})$$

with w_3, W_2 having non-negative entries.

Now, we aim to identify the parameters $(W_1, W_2, w_3, V_2, v_3, b_1, b_2, b_3)$ of this ICNN. Frontiers, i.e. points of non-differentiability of f_{EX} , arise from the switch of -at least- one neuron (otherwise, the function is locally affine, hence differentiable). A priori, each frontier (i.e. each colored segment in Figure F4) could be matched to one or more neurons which switch, so there would be an exponential number of cases to consider. Fortunately, not all arrangements of pieces of hyperplanes generated by neurons of ReLU networks are possible. The sketch of the proof is as follows: we first state general considerations on ReLU networks which enable us to restrict the possible matching between frontiers and neurons to a few cases. Then, writing down the equations which describe the frontiers, we can identify the parameters of the first layer and second. We end up showing there must be a negative weight in W_2 , which contradicts the ICNN hypothesis.

Notation

We denote $\mu_1, \mu_2 \in N_1$ the neurons of the first layer and $\nu_1, \nu_2 \in N_2$ the neurons of the second layer. We use the following notation: $z_\eta : \mathbb{R}^2 \rightarrow \mathbb{R}$ is the pre-activation of the neuron $\eta \in \{\mu_1, \mu_2, \nu_1, \nu_2\}$ and $y_\eta := \text{ReLU} \circ z_\eta : \mathbb{R}^2 \rightarrow \mathbb{R}$ is its post-activation. For a layer $i \in \{1, 2\}$, we denote $z_i : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ its pre-activation and $y_i : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ its post-activation. We denote $H_\eta := \{x \in \mathbb{R}^2 : z_\eta(x) = 0\}$ the 0-level set of the pre-activation of neuron η . For $k, l \in \{1, \dots, 6\}$, we denote $F_{k,l}$ the frontier between neighboring regions R_k and R_l following the numbering given Figure F4.

Frontiers from the first layer There are two cases to consider separately regarding W_1 : the full rank case and degenerated configurations. If W_1 is full rank (i.e. rank 2), then we get from $H_{\mu_i} = \{x \in \mathbb{R}^2 : (W_1)_i x + (b_1)_i = 0\}$ that $\dim(H_{\mu_i}) = 1$ for every $i \in \{1, 2\}$ and $H_{\mu_1} \cap H_{\mu_2}$ is a singleton. If W_1 is not full rank, then different (degenerated) configurations can occur: one H_{μ_i} is of dimension 0 or 2 or H_{μ_1} and H_{μ_2} are parallel (or coincident). In any case, all level lines of z_{μ_1} and z_{μ_2} are parallel.

THE FULL RANK CASE.

Next, we assume W_1 is full rank.

Frontiers from the second layer

In this case, the two hyperplanes generated by the first layer neurons are crossing lines, it means that the first layer divides the input space \mathbb{R}^2 into 4 regions, that we denote $A_k, k \in \{1, \dots, 4\}$. We now describe the possible shapes that the 0-level set of pre-activations of second layer neurons (so called bent hyperplanes [40]) can take. This result is illustrated in Figure F5.

On each of the A_k , z_{ν_i} is an affine function of the input x : y_1 and $V_2 x$ are both affine in x on each region. As in the reasoning for the first layer, $H_{\nu_i} \cap \text{int}(A_k)$ can only be: the empty set, the full region A_k , a segment which joins the two boundaries of A_k or a ray which intersects one boundary of A_k . Besides, if H_{ν_i} intersects two neighboring regions A_k and A_l with rays/segment, then these rays/segments must have a common point at the frontier $A_k \cap A_l$ (to see it, consider $x \in H_{\nu_i} \cap A_k \cap A_l$ and any point $y \in H_{\nu_i} \cap \text{int} A_l$, then $[x, y] \subset H_{\nu_i} \cap A_l$). Last, if $H_{\nu_i} \cap A_k$ is of dimension one (ray or segment) and H_{ν_i} intersects the frontier $A_k \cap A_l$, then $H_{\nu_i} \cap A_l$ is also of dimension 1. In other words, if a segment or ray of H_{ν_i} hits a frontier, it must continue on the other side of the frontier.

Following the previous remarks, we provide in Figure F5 all admissible shapes for the bent hyperplane of a neuron of the second layer.

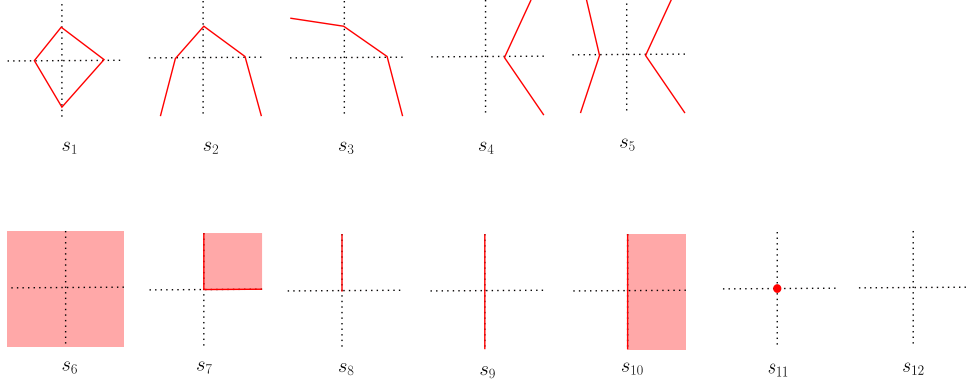


Fig. F5: Admissible shapes of the bent hyperplane of the second layer. Dotted lines correspond to crossing hyperplanes given by neurons of the first layer. Top: non degenerated cases. Bottom: degenerated cases.

As f_{EX} has strictly more than 4 regions, there exists at least one neuron of the second layer, wlog $\nu_1 \in N_2$, such that $H_{\nu_1} \cap \text{int } A_k$ has dimension 1 for some region A_k (other cases correspond to degenerated cases where ν_i does not introduce new regions in f_{EX}). Using the notations in Figure F5, we end-up with $H_{\nu_1} \in \mathcal{S} := \{s_1, s_2, s_3, s_4, s_5\}$ and $H_{\nu_2} \in \mathcal{S} \cup \{s_6, s_7, s_8, s_9, s_{10}\}$.

Frontiers of the full ICNN The last layer is linear with respect to y_2 so it can not create additional frontiers (i.e. can not add points of non differentiability of f_{EX}). So, the landscape we observe in Figure F4 must correspond to the superposition of the bent hyperplanes of the second layer. On the other hand, hyperplanes corresponding to the first layer may be partially hidden by the effect of the second layer.

Identification of the neurons At points A, B and C in Figure F4, pieces of hyperplanes cross, and there is some bend. Hence, they correspond to at least one neuron of the first layer and one neuron of the second layer. In particular, the points A, B, C all belong to $H_{\mu_1} \cup H_{\mu_2}$. As they are not aligned, they can not all belong to a single H_{μ_i} , $i \in \{1, 2\}$: wlog, either $[A, B] \subset H_{\mu_1}$, either $[B, C] \subset H_{\mu_1}$.

- *Case $[B, C] \subset H_{\mu_1}$* (Figure F6a). In that case $B, C \in H_{\mu_1}$ and $A \in H_{\mu_2}$. Then, $F_{5,6} \cup F_{6,4}$ corresponds to a bent hyperplane associated with a neuron of the second layer, wlog H_{ν_1} (shapes s_4 or s_5 in Figure F5). The frontiers $F_{3,5} \cup F_{3,4}$ cannot also belong to H_{ν_1} (shape s_5): indeed, it would imply that H_{μ_1} and H_{μ_2} cross on $[B, C]$ and A would not belong to H_{μ_2} . So, the only admissible pattern for H_{ν_1} is then s_4 . So, $F_{3,5} \cup F_{3,4} \subset H_{\nu_2}$. We end up on a contradiction: at A , this bent hyperplane should split in two pieces, impossible.

- *Case $[A, B] \subset H_{\mu_1}$* (Figure F6b). In that case $A, B \in H_{\mu_1}$ and $C \in H_{\mu_2}$. We have already excluded $[B, C] \subset H_{\mu_2}$, so either $F_{4,6} \subset H_{\mu_2}$, either $F_{5,6} \subset H_{\mu_2}$. The latter would correspond to a degenerate configuration (parallel hyperplanes). Since we are focusing on the full-rank case for now, we do not study it at this point. We thus end-up with the identification of Figure F6c (up to a permutation of neurons in the same layer). Note that it corresponds to the initial configuration given by the non-ICNN network implementing f_{EX} .

Identification of the parameters Up to a scaling and to the sign, identifying the hyperplanes of the first layer gives the weights and bias of the first layer:

$$W_1 = \begin{pmatrix} \pm 1 & 0 \\ 0 & \pm 1 \end{pmatrix}, \quad b_1 = \begin{pmatrix} 0 \\ 0 \end{pmatrix} \quad (\text{F74})$$

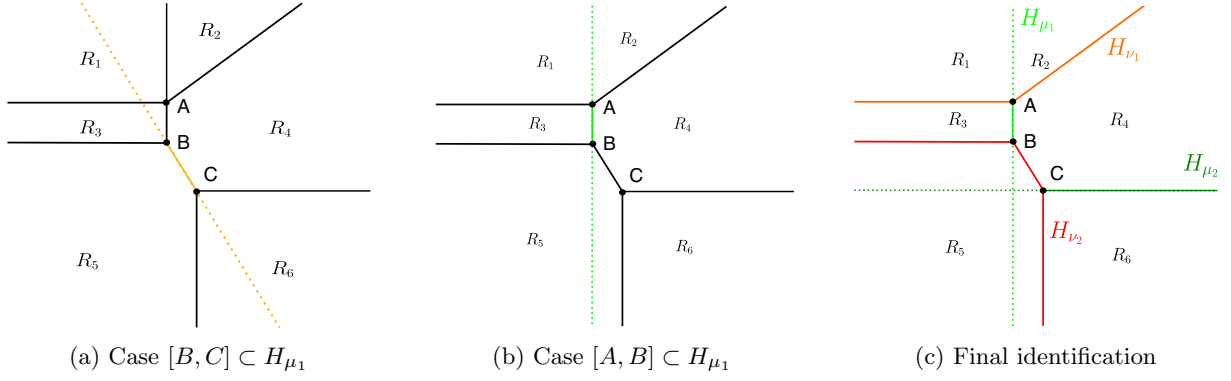


Fig. F6: Identification of the neurons.

We can fix the output layer $w_3 = (1, 1)$ as ReLU networks are invariant modulo rescaling of parameters [30]. The ICNN we consider is thus parametrized as follows:

$$(1 \ 1) \operatorname{ReLU} \left(\begin{pmatrix} a & b \\ c & d \end{pmatrix} \operatorname{ReLU} \left(\begin{pmatrix} \epsilon_1 x_1 \\ \epsilon_2 x_2 \end{pmatrix} + \begin{pmatrix} e & f \\ g & h \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} + \begin{pmatrix} i \\ j \end{pmatrix} \right) + V_3 \mathbf{x} + b_3. \quad (\text{F75})$$

where, for $k = 1, 2$, $\epsilon_k \in \{-1, 1\}$ denotes the weight $W_1[\mu_k, \mu_k]$. Define $\zeta_k = \frac{1+\epsilon_k}{2} \in \{0, 1\}$. Let us list all the equations the ICNN must satisfy to have the same frontiers as our target function.

- $F_{1:3}$. The equation of $F_{1:3}$ is $\alpha(x_2 - 1) = 0$ with $\alpha \neq 0$. The zero-level set of z_{ν_1} on $x_1 \in \mathbb{R}_-, x_2 \in \mathbb{R}_+$ rewrites

$$-(1 - \zeta_1)ax_1 + \zeta_2bx_2 + ex_1 + fx_2 + i = 0,$$

and we identify the equality constraints

$$-(1 - \zeta_1)a + e = 0 \quad (\text{F76})$$

$$\zeta_2b + f = \alpha \quad (\text{F77})$$

$$i = -\alpha \quad (\text{F78})$$

- $F_{2:4}$. The equation of $F_{2:4}$ is $\beta(x_2 - x_1 - 1) = 0$ with $\beta \neq 0$. The zero level of z_{ν_1} on $x_1 \in \mathbb{R}_+, x_2 \in \mathbb{R}_+$ is

$$\zeta_1ax_1 + \zeta_2bx_2 + ex_1 + fx_2 + i = 0$$

and we identify the equality constraints

$$\zeta_1a + e = -\beta \quad (\text{F79})$$

$$\zeta_2b + f = \beta \quad (\text{F80})$$

$$i = -\beta \quad (\text{F81})$$

We get from Equations (F78) and (F81) that $\alpha = \beta$. We study two cases: $\alpha > 0$ and $\alpha < 0$. In the first case, wlog consider $\alpha = 1$. Then, a first set of equations the ICNN weights must satisfy is:

$$-(1 - \zeta_1)a + e = 0 \quad (\text{F82})$$

$$\zeta_1a + e = -1 \quad (\text{F83})$$

$$\zeta_1, \zeta_2 \in \{0, 1\} \quad (\text{activations of } \mu_1, \mu_2 \in N_1) \quad (\text{F84})$$

$$a, b, c, d \geq 0 \quad (\text{ICNN architecture}) \quad (\text{F85})$$

By injecting (F83) into (F82), we get $a = -1$ which contradicts the non-negativity constraint on a . Now, we study the case where $\alpha < 0$, wlog $\alpha = -1$:

$$-(1 - \zeta_1)a + e = 0 \quad (\text{F86})$$

$$\zeta_1 a + e = 1 \quad (\text{F87})$$

$$\zeta_1, \zeta_2 \in \{0, 1\} \quad (\text{activations of } \mu_1, \mu_2 \in N_1) \quad (\text{F88})$$

$$a, b, c, d \geq 0 \quad (\text{ICNN architecture}) \quad (\text{F89})$$

In this case, we find $a = 1$ and there is no immediate contradiction at this stage. By (F78) $i = 1$ and $z_{\nu_1}((0, 0)) = i = 1 > 0$ so ν_1 is active on R_5 (it is active at one point in the region and the activations are constant on region).

Let us now identify weights associated to the second neuron $\nu_2 \in N_2$.

- $F_{3:5}$. The equation of $F_{3:5}$ is $\gamma(x_2 - 0.5) = 0$ with $\gamma \neq 0$. The 0-level set of z_{ν_2} on $x_1 \in \mathbb{R}_-, x_2 \in \mathbb{R}_+$ rewrites

$$-(1 - \zeta_1)cx_1 + d\zeta_2x_2 + gx_1 + hx_2 + j = 0 \quad (\text{F90})$$

and we identify the equality constraints

$$-(1 - \zeta_1)c + g = 0 \quad (\text{F91})$$

$$\zeta_2d + h = \gamma \quad (\text{F92})$$

$$j = -0.5\gamma \quad (\text{F93})$$

- $F_{4:5}$. The equation of $F_{4:5}$ is $\delta(x_2 + 2x_1 - 0.5) = 0$ with $\delta \neq 0$. The 0-level set of z_{ν_2} on $x_1 \in \mathbb{R}_+, x_2 \in \mathbb{R}_+$ is

$$\zeta_1cx_1 + \zeta_2dx_2 + gx_1 + hx_2 + j = 0 \quad (\text{F94})$$

and we identify the equality constraints

$$\zeta_1c + g = 2\delta \quad (\text{F95})$$

$$\zeta_2d + h = \delta \quad (\text{F96})$$

$$j = -0.5\delta \quad (\text{F97})$$

- $F_{5:6}$. The equation of $F_{5:6}$ is $\eta(2x_1 - 0.5) = 0$ with $\eta \neq 0$. The 0-level set of z_{ν_2} on $x_1 \in \mathbb{R}_+, x_2 \in \mathbb{R}_-$ is

$$\zeta_1cx_1 - (1 - \zeta_2)dx_2 + gx_1 + hx_2 + j = 0 \quad (\text{F98})$$

and we identify the equality constraints

$$\zeta_1c + g = 2\eta \quad (\text{F99})$$

$$-(1 - \zeta_2)d + h = 0 \quad (\text{F100})$$

$$j = -0.5\eta \quad (\text{F101})$$

So, it follows $\gamma = \delta = \eta$. We study two cases $\gamma < 0$ and $\gamma > 0$. Assume $\gamma < 0$, wlog consider $\gamma = -1$. It gives the following set of constraints

$$-(1 - \zeta_1)c + g = 0 \quad (\text{F102})$$

$$\zeta_1c + g = -2 \quad (\text{F103})$$

$$\zeta_1, \zeta_2 \in \{0, 1\} \quad (\text{F104})$$

$$c, d \geq 0, \quad (\text{F105})$$

which is not feasible as $c = -2$ contradicts the ICNN constraint $c \geq 0$. So, γ has to be positive.

Rewriting the equation of $F_{4:5}$ (second bullet) as $2\gamma x_1 + \gamma x_2 - 0.5\gamma = 0$, we see that for $x_1, x_2 < 0$, $z_{\nu_2} < 0$ so ν_2 is then inactive on R_5 . Recall that on this region ν_1 is active and the pre-activation z_{ν_1} has 4 different affine pieces (depending on which side of H_{μ_1} and H_{μ_2} the input x is). Yet, we observe on the landscape that f_{EX} is 0 on this region. The skip connection term, $v_3^\top x + b_3$, cannot put at 0 a piecewise affine function with 4 distinct linear pieces.

We end up with a non-feasibility argument: the frontiers of f_{EX} can not be implemented by a 2 hidden-layers, 2 neurons per layer ICNN.

THE DEGENERATED CASE.

Now, assume W_1 is not full rank. Then, the first layer can either generate two parallel but distinct frontiers, either a single frontier. But these frontiers must pass through the 3 points A, B and C . This requirement eliminates all cases, apart from H_{μ_1} and H_{μ_2} corresponding to $F_{1:2} \cup F_{3:4}$ and $F_{5:6}$. Up to a scaling and to the sign, identifying the hyperplanes of the first layer gives the weights and bias of the first layer:

$$W_1 = \begin{pmatrix} \pm 1 & 0 \\ \pm 1 & 0 \end{pmatrix}, \quad b_1 = \begin{pmatrix} 0 \\ \pm \frac{1}{4} \end{pmatrix} \quad (\text{F106})$$

As done previously, we can fix the output layer $w_3 = (1, 1)$. The ICNN we consider is thus parametrized as follows:

$$(1 \ 1) \text{ReLU} \left(\begin{pmatrix} a & b \\ c & d \end{pmatrix} \text{ReLU} \left(\begin{pmatrix} \epsilon_1 x_1 \\ \epsilon_2 (x_1 - \frac{1}{4}) \end{pmatrix} + \begin{pmatrix} e & f \\ g & h \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} + \begin{pmatrix} i \\ j \end{pmatrix} \right) + V_3 \mathbf{x} + b_3. \quad (\text{F107})$$

where, for $k = 1, 2$, $\epsilon_k \in \{-1, 1\}$ denotes the weight $W_1[\mu_k, \mu_k]$. In this case, the two hyperplanes generated by the first layer neurons are parallel and divide the input space \mathbb{R}^2 into 3 regions. On each of these regions, z_{ν_1} and z_{ν_2} are affine function of the input. Applying the same reasoning as before, the 0-level sets of the pre-activation of neurons $\nu_1, \nu_2 \in N_2$ on neighboring regions must be connected. We thus identify H_{ν_1} to $F_{1:3} \cup F_{2:4}$ and H_{ν_2} to $F_{3:5} \cup F_{4:5} \cup F_{4:6}$. We now study the constraints associated to H_{ν_2} .

- $F_{3:5}$. The equation of $F_{3:5}$ is $\gamma(x_2 - 0.5) = 0$ with $\gamma \neq 0$. The 0-level set of z_{ν_2} on $x_1 \in \mathbb{R}_-, x_2 \in \mathbb{R}_+$ rewrites

$$-(1 - \zeta_1)cx_1 - d(1 - \zeta_2)(x_1 - 0.25) + gx_1 + hx_2 + j = 0 \quad (\text{F108})$$

so we identify the equality constraints

$$-(1 - \zeta_1)c - d(1 - \zeta_2) + g = 0 \quad (\text{F109})$$

$$h = \gamma \quad (\text{F110})$$

$$0.25d(1 - \zeta_2) + j = -0.5\gamma \quad (\text{F111})$$

- $F_{4:5}$. The equation of $F_{4:5}$ is $\delta(x_2 + 2x_1 - 0.5) = 0$ with $\delta \neq 0$. The 0-level set of z_{ν_2} on $x_1 \in \mathbb{R}_+, x_2 \in \mathbb{R}_+$ is

$$\zeta_1 cx_1 - d(1 - \zeta_2)(x_1 - 0.25) + gx_1 + hx_2 + j = 0 \quad (\text{F112})$$

$$\zeta_1 c - d(1 - \zeta_2) + g = 2\delta \quad (\text{F113})$$

$$h = \delta \quad (\text{F114})$$

$$0.25d(1 - \zeta_2) + j = -0.5\delta \quad (\text{F115})$$

- $F_{4.6}$. The equation of $F_{4.6}$ is $\eta x_2 = 0$ with $\eta \neq 0$. The 0-level set of z_{ν_2} on $x_1 \in \mathbb{R}_+, x_2 \in \mathbb{R}_-$ is

$$\zeta_1 c x_1 + \zeta_2 d (x_1 - 0.25) + g x_1 + h x_2 + j = 0 \quad (\text{F116})$$

$$\zeta_1 c + \zeta_2 d = 0 \quad (\text{F117})$$

$$h = \eta \quad (\text{F118})$$

$$-0.25 \zeta_2 d + j = 0 \quad (\text{F119})$$

From which we identify:

- by Equations (F110), (F114) and (F118), $\gamma = \delta = \eta$,
- by Equation (F109)

$$-c - d + \underbrace{\zeta_1 c + \zeta_2 d}_{=0 \text{ by (F117)}} + g = 0 \quad (\text{F120})$$

and by Equation (F113)

$$-d + \underbrace{\zeta_1 c + \zeta_2 d}_{=0 \text{ by (F117)}} + g = 2\gamma \quad (\text{F121})$$

from which it follows

$$c = 2\gamma, \quad (\text{F122})$$

together with $c \geq 0$ (ICNN condition), it follows γ (which is $\neq 0$) must be positive.

- by Equations (F111) and (F119)

$$0.25d = -0.5\gamma - \underbrace{(j - 0.25d\zeta_2)}_{=0 \text{ by (F119)}} \quad (\text{F123})$$

leading to $0.25d = -0.5\gamma < 0$ which contradicts the non-negativity constraint on d .

We can conclude that the degenerate case is non-feasible with respect to the ICNNs constraints.

To wrap up, we cannot find an ICNN with 2 layers and 2 neurons per layer such that it implements f_{EX} , which concludes the proof.

F.2 Sanity check of the necessary convexity condition given in Lemma 5.1

For any neuron $\eta \in (N_1 = \{\mu_1, \mu_2\}) \cup (N_2 = \{\nu_1, \nu_2\})$, note that \mathfrak{X}_η is by definition the set of points such that $z_\eta(x_1, x_2) = 0$ and $z_{\eta'}(x_1, x_2) \neq 0$ for $\eta' \neq \eta$. Each pre-activation is a function from \mathbb{R}^2 to \mathbb{R} given by:

$$z_{\mu_1} : (x_1, x_2) \mapsto x_1 \quad (\text{F124})$$

$$z_{\mu_2} : (x_1, x_2) \mapsto x_2 \quad (\text{F125})$$

$$z_{\nu_1} : (x_1, x_2) \mapsto -\text{ReLU}(x_1) + \text{ReLU}(x_2) - 1 \quad (\text{F126})$$

$$z_{\nu_2} : (x_1, x_2) \mapsto 2\text{ReLU}(x_1) + \text{ReLU}(x_2) - 0.5 \quad (\text{F127})$$

It follows

$$\mathfrak{X}_{\mu_1} = \{x \in \mathbb{R}^2 : x_1 = 0\} \setminus \{(0, 0), (0, 1), (0, 0.5)\} = \{0\} \times (-\infty, 0) \cup (0, 0.5) \cup (0.5, 1) \cup (1, +\infty) \quad (\text{F128})$$

$$\mathfrak{X}_{\mu_2} = \{x \in \mathbb{R}^2 : x_2 = 0\} \setminus \{(0, 0), (0, 0.25)\} = (-\infty, 0) \cup (0, 0.25) \cup (0.25, +\infty) \times \{0\} \quad (\text{F129})$$

$$\mathfrak{X}_{\nu_1} = \{x \in \mathbb{R}_+^2 : -x_1 + x_2 - 1 = 0\} \setminus \{(0, 1)\} \quad (\text{F130})$$

$$\mathfrak{X}_{\nu_2} = \{x \in \mathbb{R}_+^2 : 2x_1 + x_2 - 0.5 = 0\} \setminus \{(0, 0.25), (0.25, 0)\} \quad (\text{F131})$$

All neurons are isolated

By definition, a neuron $\eta \in (N_1 = \{\mu_1, \mu_2\}) \cup (N_2 = \{\nu_1, \nu_2\})$ is isolated if $\mathfrak{X}_\eta \neq \emptyset$, which is clearly satisfied.

Determine $\mathbf{a}_2(\mathfrak{X}_{\mu_1})$

- Consider $x = (x_1 = 0, x_2)$ with $x_2 \in (1, +\infty)$:

$$z_{\nu_1}(x) > 0, \quad z_{\nu_2}(x) > 0, \quad (\text{F132})$$

so $(1, 1) \in \mathbf{a}_2(\mathfrak{X}_{\mu_1})$.

- Consider $x = (x_1 = 0, x_2)$ with $x_2 \in (0.5, 1)$:

$$z_{\nu_1}(x) < 0, \quad z_{\nu_2}(x) > 0, \quad (\text{F133})$$

so $(0, 1) \in \mathbf{a}_2(\mathfrak{X}_{\mu_1})$.

- Consider $x = (x_1 = 0, x_2)$ with $x_2 \in (-\infty, 0) \cup (0, 0.5)$:

$$z_{\nu_1}(x) < 0, \quad z_{\nu_2}(x) < 0, \quad (\text{F134})$$

so $(0, 0) \in \mathbf{a}_2(\mathfrak{X}_{\mu_1})$.

We end up with $\mathbf{a}_2(\mathfrak{X}_{\mu_1}) = \{(0, 0), (0, 1), (1, 1)\}$.

Appendix G Algorithm

In the case of ReLU neurons, any input point x in \mathbb{R}^d can be assigned a *sign-vector* $\sigma(x) := (\text{sign}(z_\nu(x)))_{\nu \in H}$ – where $\text{sign}(u) = u/|u|$ except for $\text{sign}(0)$ whose value is 0. The sign-vector of a point is simply related to its activations by $(a_\nu(x))_{\nu \in H} = \text{ReLU}(\sigma(x))$.

The algorithm of Berzins [35] extracts the 1-faces of the polyhedral complex, i.e. its edges, together with their sign-vectors. We use this information to reconstruct the set of reachable activations $\mathbf{a}(\mathfrak{X}_\nu)$: first, we use sign-vectors of points on 1-faces to derive sign-vectors of points on $d - 1$ -faces. Then, taking the ReLU of this vector yields the activations needed to check convexity.

The sign-vector of any point belonging to a 1-face contains $d - 1$ zero entries: it lies on the 0-level set of $d - 1$ different preactivations. On the other hand, sign-vectors of points on $d - 1$ -faces have exactly one zero, because they correspond to points where only one preactivation is 0. Sign-vectors on such point can be recovered by following the so-called *perturbation* process described in Berzins [35, Section 4.1] which enables to recover all faces, in particular the $d - 1$ faces. For a neuron ν :

- we find all possible sign-vectors of points on 1-faces where the coordinate ν is 0
- for each vector, for every subset of $d - 2$ vanishing entries, we replace those 0 values by every possible combination of 1 and -1 (*perturbation* process). This gives all the sign-vectors of points on $d - 1$ -faces.
- taking the ReLU of those sign-vectors gives the corresponding activations, that we use in the same fashion as for the first bottleneck in Section 8 to compute $\langle \mathbf{a}, \Phi^{\nu \rightarrow} \rangle$.