

Neural Network Identifiability for a Family of Sigmoidal Nonlinearities

Verner Vlačić and Helmut Bölcskei

Chair for Mathematical Information Science, ETH Zurich, Switzerland

Email: vlasticv@mins.ee.ethz.ch, hboelcskei@ethz.ch

Abstract

This paper addresses the following question of neural network identifiability: Does the input-output map realized by a feed-forward neural network with respect to a given nonlinearity uniquely specify the network architecture, weights, and biases? Existing literature on the subject [1], [2], [3] suggests that the answer should be *yes, up to certain symmetries induced by the nonlinearity, and provided the networks under consideration satisfy certain “genericity conditions”*. The results in [1] and [2] apply to networks with a single hidden layer and in [3] the networks need to be fully connected. In an effort to answer the identifiability question in greater generality, we derive *necessary* genericity conditions for the identifiability of neural networks of arbitrary depth and connectivity with an arbitrary nonlinearity. Moreover, we construct a family of nonlinearities for which these genericity conditions are *minimal*, i.e., both *necessary and sufficient*. This family is large enough to approximate many commonly encountered nonlinearities to within arbitrary precision in the uniform norm.

KEYWORDS

Deep neural networks, identifiability, sigmoidal nonlinearities

MATHEMATICS SUBJECT CLASSIFICATION

68T07 Artificial neural networks and deep learning, 30-XX Functions of a complex variable, 46-XX Functional analysis, 22-XX Topological groups, Lie groups

I. INTRODUCTION

Deep learning has become a highly successful machine learning method employed in a wide range of applications such as optical character recognition [4], image classification [5], and speech recognition [6]. In a typical deep learning scenario one aims to fit a parametric model, realized by a deep neural network, to match a set of training data points. In order to make the ensuing discussion more concrete, we begin with the definition of a neural network and the map it realizes under a nonlinearity.

Definition 1 (Neural network). We call an ordered sequence

$$\mathcal{N} = (D_0, D_1, \dots, D_L; W^1, \theta^1, W^2, \theta^2, \dots, W^L, \theta^L),$$

a neural network, where

– L is a positive integer, referred to as the depth of \mathcal{N} ,

- (D_0, D_1, \dots, D_L) is an $(L + 1)$ -tuple of positive integers, called the layout,
- $W^\ell = (W_{jk}^\ell) \in \mathbb{R}^{D_\ell \times D_{\ell-1}}$, $\ell \in \{1, \dots, L\}$, are matrices whose entries are referred to as the network’s weights, and
- $\theta^\ell = (\theta_j^\ell) \in \mathbb{R}^{D_\ell}$, $\ell \in \{1, \dots, L\}$, are vectors of the so-called biases.

Furthermore, we stipulate that none of the W^ℓ , $\ell \in \{1, \dots, L\}$, have an identically zero row or an identically zero column.

Definition 2. Given a neural network \mathcal{N} and a nonlinear function $\rho : \mathbb{R} \rightarrow \mathbb{R}$, referred to as the nonlinearity, we define the *map realized by \mathcal{N} under ρ* as the function $\langle \mathcal{N} \rangle^\rho : \mathbb{R}^{D_0} \rightarrow \mathbb{R}^{D_L}$ given by

$$\langle \mathcal{N} \rangle^\rho(x) = \rho(W^L(\rho(W^{L-1}(\dots \rho(W^1 x + \theta^1) \dots) + \theta^{L-1})) + \theta^L), \quad x \in \mathbb{R}^{D_0},$$

where ρ acts on real vectors in a componentwise fashion.

The requirement that the matrices W^ℓ in Definition 1 have nonzero rows corresponds to the absence of nodes whose contributions depend on the biases only, and are therefore constant as functions of the input. Similarly, columns that are identically zero correspond to nodes whose contributions do not enter the computation at the next layer. The map of a neural network failing this requirement can be realized by a network obtained by simply removing such spurious nodes. In practical applications, the numbers L, D_0, D_1, \dots, D_L are typically determined through heuristic considerations, whereas the coefficients W^ℓ, θ^ℓ of the affine maps $x \mapsto W^\ell x + \theta^\ell$ are learned based on training data. For an overview of practical techniques for deep learning, see [7]. Neural networks are often studied as mathematical objects in their own right, for instance in approximation theory [8], [9], [10], [11] and in control theory [12], [13]. In this context, a natural question is that of identification: Can a neural network be uniquely identified from the map it is to realize? Specifically, we will be interested in identifiability according to the following definition.

Definition 3 (Identifiability). Given positive integers D_{in} and D_{out} , define $\mathcal{N}^{D_{in}, D_{out}}$ to be the set of all neural networks whose layouts (D_0, \dots, D_L) satisfy $D_0 = D_{in}$ and $D_L = D_{out}$, but are otherwise arbitrary. Let \mathcal{N} be a subset of $\mathcal{N}^{D_{in}, D_{out}}$, ρ a nonlinearity, and \sim an equivalence relation on $\mathcal{N}^{D_{in}, D_{out}}$.

- (i) We say that \sim is *compatible with (\mathcal{N}, ρ)* if, for all $\mathcal{N}_1, \mathcal{N}_2 \in \mathcal{N}$,

$$\mathcal{N}_1 \sim \mathcal{N}_2 \quad \implies \quad \langle \mathcal{N}_1 \rangle^\rho(x) = \langle \mathcal{N}_2 \rangle^\rho(x), \quad \forall x \in \mathbb{R}^{D_{in}}.$$

- (ii) We say that (\mathcal{N}, ρ) is *identifiable up to \sim* if, for all $\mathcal{N}_1, \mathcal{N}_2 \in \mathcal{N}$,

$$\langle \mathcal{N}_1 \rangle^\rho(x) = \langle \mathcal{N}_2 \rangle^\rho(x), \quad \forall x \in \mathbb{R}^{D_{in}} \quad \implies \quad \mathcal{N}_1 \sim \mathcal{N}_2.$$

Thus, by informally saying that a neural network \mathcal{N}_1 in a certain class is identifiable, we mean that any neural network \mathcal{N}_2 in the same class giving rise to the same output map, i.e., $\langle \mathcal{N}_1 \rangle^\rho = \langle \mathcal{N}_2 \rangle^\rho$, is necessarily equivalent to \mathcal{N}_2 . The role of the equivalence relation \sim in the previous definition is thus to “measure the degree of non-uniqueness”, and in particular, to accommodate symmetries within the network that may arise either from symmetries induced by the network weights and biases (such as the presence of clone pairs, to be introduced in Definition 5), symmetries of the nonlinearity (e.g., \tanh is odd), or both simultaneously. These abstract concepts will be incarnated momentarily when

discussing the seminal work by Fefferman [3], and in Section II through Definitions 4 and 5, as well as in the examples leading up to the formulation of the paper’s main results.

In [3], Fefferman showed that neural networks satisfying the following genericity conditions are, indeed, uniquely determined by the map they realize under the nonlinearity $\rho = \tanh$, up to certain obvious isomorphisms of networks:

Assumptions 1 (Fefferman’s genericity conditions).

- (i) $\theta_j^\ell \neq 0$, for all ℓ and j , and $|\theta_j^\ell| \neq |\theta_{j'}^\ell|$, for all ℓ and j, j' with $j \neq j'$.
- (ii) $W_{jk}^\ell \neq 0$, for all ℓ, j , and k , and
- (iii) for all ℓ, k and j, j' with $j \neq j'$,

$$W_{jk}^\ell / W_{j'k}^\ell \notin \{p/q : p, q \in \mathbb{Z}, 1 \leq q \leq 100D_\ell^2\}.$$

More precisely, for fixed positive integers D_{in} and D_{out} , Fefferman showed that $(\mathcal{N}_{A1}^{D_{in}, D_{out}}, \tanh)$ is identifiable up to \sim_\pm , where $\mathcal{N}_{A1}^{D_{in}, D_{out}}$ is defined as the set of all neural networks in $\mathcal{N}^{D_{in}, D_{out}}$ satisfying Assumptions 1, and \sim_\pm is defined by stipulating that $\mathcal{N} \sim_\pm \tilde{\mathcal{N}}$ if and only if

- (i) $L = \tilde{L}$ and $(D_0, D_1, \dots, D_L) = (\tilde{D}_0, \tilde{D}_1, \dots, \tilde{D}_L)$, and
- (ii) there exists a collection of signs $\{\epsilon_j^\ell : 0 \leq \ell \leq L, 1 \leq j \leq D_\ell\}$, $\epsilon_j^\ell \in \{-1, +1\}$, and permutations $\gamma_\ell : \{1, \dots, D_\ell\} \rightarrow \{1, \dots, D_\ell\}$ such that
 - γ_ℓ is the identity permutation and $\epsilon_j^\ell = +1$, $j \in \{1, \dots, D_\ell\}$, whenever $\ell = 0$ or $\ell = L$, and
 - for all $\ell \in \{1, \dots, L\}$, $k \in \{1, \dots, D_{\ell-1}\}$, and $j \in \{1, \dots, D_\ell\}$,

$$\tilde{W}_{jk}^\ell = \epsilon_j^\ell W_{\gamma_\ell(j)\gamma_{\ell-1}(k)}^\ell \epsilon_k^{\ell-1}, \quad \text{and} \quad \tilde{\theta}_j^\ell = \epsilon_j^\ell \theta_{\gamma_\ell(j)}^\ell.$$

It can be verified that \sim_\pm is an equivalence relation on $\mathcal{N}_{A1}^{D_{in}, D_{out}}$. Networks $\mathcal{N}, \tilde{\mathcal{N}}$ such that $\mathcal{N} \sim_\pm \tilde{\mathcal{N}}$ are said to be *isomorphic up to sign changes*. The permutations γ_ℓ reflect the fact that the ordering of the neurons in the hidden layers $1, \dots, L-1$ is not unique, whereas the freedom in choosing the signs ϵ_j^ℓ reflects that \tanh is an odd function. It can be verified that any two networks isomorphic up to sign changes give rise to the same map under the \tanh nonlinearity, so \sim_\pm is compatible with $(\mathcal{N}_{A1}^{D_{in}, D_{out}}, \tanh)$. The crux of Fefferman’s result therefore lies in proving the converse statement, namely that two networks giving rise to the same map with respect to \tanh are necessarily isomorphic up to sign changes. This is effected by the insight that the depth, the layout, and the weights and biases of a network $\mathcal{N} \in \mathcal{N}_{A1}^{D_{in}, D_{out}}$ are encoded in the geometry of the singularities of the analytic continuation of $\langle \mathcal{N} \rangle^{\tanh}$.

We note that Fefferman distilled the precise conditions of Assumptions 1 from his proof technique, in order to define a class of neural networks that is, on the one hand, sufficiently small to guarantee identifiability, and on the other hand, sufficiently large to encompass “generic” networks. Indeed, if we consider the network weights and biases $(W^1, \theta^1, \dots, W^L, \theta^L)$ as elements of the space $\mathbb{R}^{D_1 \times D_0} \times \mathbb{R}^{D_1} \times \dots \times \mathbb{R}^{D_L \times D_{L-1}} \times \mathbb{R}^{D_L}$, then Assumptions 1 rule out only a set of measure zero. In the contemporary practical machine learning literature, however, a network satisfying Assumptions 1 would hardly be considered generic, as Part (i) of Assumptions 1 implies that all biases are nonzero, and Part (ii) imposes full connectivity throughout the network.

Indeed, Fefferman remarks explicitly that it would be interesting to replace Assumptions 1 with minimal hypotheses, and to study nonlinearities other than \tanh . The present paper aims to address

these two issues. Characterizing the fundamental nature of conditions necessary for identifiability with respect to a fixed nonlinearity, even a simple one such as \tanh , is likely a rather formidable task. In fact, the minimal identifiability conditions may generally depend on “fine” properties of the nonlinearity under consideration, and it is hence unclear how much insight can be obtained by having conditions that are specific to a given nonlinearity. We will thus be interested in an identification result with very mild conditions on the weights and biases of the neural networks to be identified, while still accommodating a broad class of nonlinearities.

II. CONTRIBUTIONS

We begin with two motivating examples. These lead up to the statements of our main contributions, whose corresponding proofs are developed in the remainder of the paper. We consider nonlinearities ρ which are not necessarily odd (as \tanh), and thus need an equivalence relation which dispenses with sign changes.

Definition 4 (Neural network isomorphism). We say that the neural networks \mathcal{N} and $\tilde{\mathcal{N}}$ are isomorphic, and write $\mathcal{N} \simeq \tilde{\mathcal{N}}$, if

- (i) $L = \tilde{L}$ and $(D_0, D_1, \dots, D_L) = (\tilde{D}_0, \tilde{D}_1, \dots, \tilde{D}_L)$, and
- (ii) there exist permutations $\gamma_\ell : \{1, \dots, D_\ell\} \rightarrow \{1, \dots, D_\ell\}$ such that
 - γ_ℓ is the identity permutation for $\ell = 0$ and $\ell = L$, and
 - for all $\ell \in \{1, \dots, L\}$, $k \in \{1, \dots, D_{\ell-1}\}$, and $j \in \{1, \dots, D_\ell\}$,

$$\tilde{W}_{jk}^\ell = W_{\gamma_\ell(j)\gamma_{\ell-1}(k)}^\ell, \quad \text{and} \quad \tilde{\theta}_j^\ell = \theta_{\gamma_\ell(j)}^\ell.$$

In the remainder of the paper we will work exclusively with isomorphisms in the sense of Definition 4. Note that any two isomorphic networks give rise to the same map with respect to any nonlinearity ρ , and thus \simeq is an equivalence relation compatible with any pair (\mathcal{N}, ρ) . The requirement that γ_ℓ be the identity map for $\ell \in \{0, L\}$ in the previous definition again corresponds to the fact that the inputs and the outputs of a neural network are not generally interchangeable. Indeed, suppose that $\mathcal{N}^\rho : \mathbb{R}^2 \rightarrow \mathbb{R}^2$, $\mathcal{N}^\rho(x, y) = (x, 2y)$ is the map of a neural network with respect to some nonlinearity ρ . Let $\mathcal{N}_1, \mathcal{N}_2$, and \mathcal{N}_3 be the networks obtained from \mathcal{N} by interchanging the inputs of \mathcal{N} , the outputs of \mathcal{N} , and both inputs and outputs, respectively. Then $\mathcal{N}_1^\rho(x, y) = (y, 2x)$, $\mathcal{N}_2^\rho(x, y) = (2y, x)$, and $\mathcal{N}_3^\rho(x, y) = (2x, y)$ are, indeed, distinct functions. We now give an example that Fefferman uses to motivate the necessity of restricting the class of all neural networks $\mathcal{N}^{D_{in}, D_{out}}$ to a smaller class to be identifiable up to an equivalence relation. In Fefferman’s case, the equivalence relation is \sim_{\pm} , but the example is equally pertinent to the relation \simeq . Suppose that \mathcal{N} is a neural network with $L \geq 2$, and ℓ_0, j_1, j_2 with $1 \leq \ell_0 \leq L - 1$ and $1 \leq j_1 < j_2 \leq D_{\ell_0}$ are such that $\theta_{j_1}^{\ell_0} = \theta_{j_2}^{\ell_0}$ and $W_{j_1 k}^{\ell_0} = W_{j_2 k}^{\ell_0}$ for all k . Then, if $\tilde{\mathcal{N}}$ is obtained from \mathcal{N} by replacing $W_{1j_1}^{\ell_0+1}$ and $W_{1j_2}^{\ell_0+1}$ with an arbitrary pair of numbers $\tilde{W}_{1j_1}^{\ell_0+1}$ and $\tilde{W}_{1j_2}^{\ell_0+1}$ such that $W_{1j_1}^{\ell_0+1} + W_{1j_2}^{\ell_0+1} = \tilde{W}_{1j_1}^{\ell_0+1} + \tilde{W}_{1j_2}^{\ell_0+1}$, then $\langle \tilde{\mathcal{N}} \rangle^\rho = \langle \mathcal{N} \rangle^\rho$, for any ρ . This example motivates the following definition.

Definition 5 (No-clones condition). Let \mathcal{N} be a neural network as in Definition 1. We say that \mathcal{N} has a clone pair if there exist $\ell \in \{1, \dots, L\}$ and $j, j' \in \{1, \dots, D_\ell\}$ with $j \neq j'$ such that

$$(\theta_j^\ell, W_{j1}^\ell, \dots, W_{jD_{\ell-1}}^\ell) = (\theta_{j'}^\ell, W_{j'1}^\ell, \dots, W_{j'D_{\ell-1}}^\ell).$$

If \mathcal{N} does not have a clone pair, we say that \mathcal{N} satisfies the no-clones condition.

As the nonlinearity ρ in the example above is completely arbitrary, the no-clones condition is necessary to have any hope of obtaining identifiability up to \simeq . Hence, with our program in mind, given positive integers D_{in} and D_{out} , we define

$$\mathcal{N}_{nc}^{D_{in}, D_{out}} = \{\mathcal{N} \in \mathcal{N}^{D_{in}, D_{out}} : \mathcal{N} \text{ satisfies the no-clones condition}\},$$

and seek nonlinearities ρ such that $(\mathcal{N}_{nc}^{D_{in}, D_{out}}, \rho)$ is identifiable up to \simeq . As any class strictly containing $\mathcal{N}_{nc}^{D_{in}, D_{out}}$, paired with any nonlinearity, fails identifiability up to \simeq , the no-clones condition furnishes a canonical minimal assumption for identifiability up to \simeq . Similarly to $\mathcal{N}_{AI}^{D_{in}, D_{out}}$, the class $\mathcal{N}_{nc}^{D_{in}, D_{out}}$, paired with any measurable nonlinearity ρ such that $\lim_{x \rightarrow \infty} \rho(x)$ and $\lim_{x \rightarrow -\infty} \rho(x)$ exist and are not equal, satisfies the universal approximation property in the sense of Hornik [14] and Cybenko [15]. The following example demonstrates that insisting on the no-clones condition as the only assumption on the weights, biases, and layout will necessarily come at the cost of restricting the class of nonlinearities that allow for identifiability. Let $\rho(x) = \min\{1, \max\{0, x\}\}$ be the clipped rectified linear unit (ReLU) function. Note that

$$\rho\left(\rho(x) - \frac{1}{2}\rho(2x) - \frac{1}{2}\rho(2x-1) + 0\right) = 0, \quad \text{for all } x \in \mathbb{R}.$$

Now, given an arbitrary neural network $\mathcal{N} = (W^1, \theta^1, W^2, \theta^2, \dots, W^L, \theta^L)$ with $D_L = 1$ satisfying the no-clones condition, the network

$$\mathcal{N}_0 = \left(W^1, \theta^1, W^2, \theta^2, \dots, W^L, \theta^L, \left(\frac{1}{2}\right), \left(\begin{smallmatrix} 0 \\ -1 \end{smallmatrix}\right), \left(1 \quad -\frac{1}{2} \quad -\frac{1}{2}\right), 0\right)$$

also satisfies the no-clones condition, and yields the identically-zero output, i.e., $\mathcal{N}_0^p \equiv 0$. We have thus constructed an infinite collection of distinct networks satisfying the no-clones condition and all yielding the identically-zero map. The class of identically-zero output maps therefore contains networks of different depths and layouts, and thus identifiability up to \simeq fails. This leads to the conclusion that a uniqueness result for neural networks with the clipped ReLU nonlinearity would need to encompass genericity conditions more stringent than the no-clones condition. Nonetheless, we are able to construct a class of real meromorphic nonlinearities σ yielding identifiability without any assumptions on the neural networks beyond the no-clones condition, and which is large enough to uniformly approximate any piecewise C^1 nonlinearity ρ with $\rho' \in BV(\mathbb{R})$, where

$$BV(\mathbb{R}) = \left\{ f \in L^1(\mathbb{R}) : \|f\|_{BV(\mathbb{R})} := \sup_{\substack{\varphi \in C_c^1(\mathbb{R}) \\ \|\varphi\|_{L^\infty(\mathbb{R})} \leq 1}} \int_{\mathbb{R}} f(x)\varphi'(x)dx < \infty \right\}$$

is the space of functions of bounded variation on \mathbb{R} .

Concretely, we have the following main result of this paper.

Theorem 1 (Uniqueness Theorem). *Let D_{in} and D_{out} be arbitrary positive integers. Furthermore, let ρ be a piecewise C^1 function with $\rho' \in BV(\mathbb{R})$ and let $\epsilon > 0$. Then there exists a meromorphic function $\sigma : \mathcal{D} \rightarrow \mathbb{C}$, $\mathcal{D} \supset \mathbb{R}$, $\sigma(\mathbb{R}) \subset \mathbb{R}$ such that $\|\rho - \sigma\|_{L^\infty(\mathbb{R})} < \epsilon$ and $(\mathcal{N}_{nc}^{D_{in}, D_{out}}, \sigma)$ is identifiable up to \simeq .*

We note that, having fixed the input and output dimensions D_{in} and D_{out} , the depths and the layouts of the networks in $\mathcal{N}_{nc}^{D_{in}, D_{out}}$ are completely arbitrary. Examples of nonlinearities $\rho(x)$ covered by

Theorem 1 include many sigmoidal functions such as the aforementioned clipped ReLU, the logistic function $\frac{1}{1+e^{-x}}$, the hyperbolic tangent $\tanh(x)$, the inverse tangent $\arctan(x)$, the softsign function $\frac{x}{1+|x|}$, the inverse square root unit $\frac{x}{\sqrt{1+ax^2}}$, the clipped identity $\frac{x}{\max\{1,|x|/a\}}$, and the soft clipping function $\frac{1}{a} \log \frac{1+e^{ax}}{1+e^{a(x-1)}}$, where $a > 0$ is fixed in the last two cases. Unbounded nonlinearities such as the ReLU are not comprised. The nonlinearities σ for which we have identifiability, unfortunately, need to be constructed, and, at the present time, we do not have an identification result for arbitrary given σ . Furthermore, we remark that the statement of Theorem 1 is “not continuous” in the approximation error ϵ . Indeed, while the clipped ReLU function satisfies the conditions of Theorem 1, as shown in the example above, there exist non-isomorphic networks \mathcal{N}_0 and $\tilde{\mathcal{N}}_0$ satisfying the no-clones condition and $\langle \mathcal{N}_0 \rangle^\rho(x) = 0 = \langle \tilde{\mathcal{N}}_0 \rangle^\rho(x)$, for all $x \in \mathbb{R}^{D_0}$, where ρ is the clipped ReLU function. We will see that Theorem 1 is, in fact, a consequence of the following result, which states that the maps realized by pairwise non-isomorphic networks with $D_L = 1$, under a nonlinearity σ according to Theorem 1, are linearly independent functions $\mathbb{R}^{D_0} \rightarrow \mathbb{R}$.

Theorem 2 (Linear Independence Theorem). *Let D_{in} be an arbitrary positive integer, let ρ be a piecewise C^1 function with $\rho' \in BV(\mathbb{R})$, and let $\epsilon > 0$. Then there exists a meromorphic function $\sigma : \mathcal{D} \rightarrow \mathbb{C}$, $\mathcal{D} \supset \mathbb{R}$, $\sigma(\mathbb{R}) \subset \mathbb{R}$ such that $\|\rho - \sigma\|_{L^\infty(\mathbb{R})} < \epsilon$ with the following property: Suppose that \mathcal{N}_j , $j = 1, 2, \dots, n$, are pairwise non-isomorphic (in the sense of \simeq) neural networks in $\mathcal{N}_{nc}^{D_{in},1}$. Then, $\{\langle \mathcal{N}_j \rangle^\sigma\}_{j=1}^n \cup \{\mathbf{1}\}$ is a linearly independent set of functions $\mathbb{R}^{D_0} \rightarrow \mathbb{R}$, where $\mathbf{1}$ denotes the constant function taking on the value 1.*

Remark. The function $\mathbf{1}$ is included in the linearly independent set both for the sake of greater generality of the statement, and to facilitate the proof of Theorem 2.

Unfortunately, Theorem 2 does not generalize to multiple outputs $D_{out} > 1$, as shown by the following example: Fix an arbitrary network \mathcal{N} according to Definition 1 such that $L \geq 2$, $D_L = 4$, $\theta_L = \mathbf{0}$, and \mathcal{N} satisfies the no-clones condition. Define $U^m \in \mathbb{R}^{2 \times D_{L-1}}$, $m \in \{1, 2, 3, 4\}$, as the submatrices of W^L consisting of the rows 1 and 3, 1 and 4, 2 and 4, and 2 and 3, respectively. Furthermore, define the networks

$$\mathcal{N}_m := (D_0, D_1, \dots, D_{L-1}, 2; W^1, \theta^1, W^2, \theta^2, \dots, W^{L-1}, \theta^{L-1}, U^m, \mathbf{0}),$$

for $m \in \{1, 2, 3, 4\}$. As \mathcal{N} satisfies the no-clones condition, the networks \mathcal{N}_m , $m \in \{1, 2, 3, 4\}$, also satisfy the no-clones condition, and are pairwise non-isomorphic.

Now, let ρ be an arbitrary nonlinearity, and write $\langle \mathcal{N} \rangle^\rho = (f_1, f_2, f_3, f_4)$, where $f_m : \mathbb{R}^{D_0} \rightarrow \mathbb{R}$, $m \in \{1, 2, 3, 4\}$. Then

$$\langle \mathcal{N}_1 \rangle^\rho = (f_1, f_3), \quad \langle \mathcal{N}_2 \rangle^\rho = (f_1, f_4), \quad \langle \mathcal{N}_3 \rangle^\rho = (f_2, f_4), \quad \text{and} \quad \langle \mathcal{N}_4 \rangle^\rho = (f_2, f_3),$$

and so

$$\langle \mathcal{N}_1 \rangle^\rho - \langle \mathcal{N}_2 \rangle^\rho + \langle \mathcal{N}_3 \rangle^\rho - \langle \mathcal{N}_4 \rangle^\rho = \begin{pmatrix} 0 + f_1 - f_1 + f_2 - f_2 \\ 0 + f_3 - f_4 + f_4 - f_3 \end{pmatrix} = \mathbf{0}.$$

The set $\{\langle \mathcal{N}_m \rangle^\rho\}_{m=1}^4$ is hence linearly dependent, showing that Theorem 2 cannot be generalized to multiple outputs by replacing $\mathcal{N}_{nc}^{D_{in},1}$ with $\mathcal{N}_{nc}^{D_{in},D_{out}}$. We now provide a panorama of the proofs of Theorems 1 and 2. The proof of Theorem 1 is by way of contradiction with Theorem 2. Specifically, assume that D_{in} , D_{out} , ρ , and $\epsilon > 0$ are as in the statement of Theorem 1, and let σ be a nonlinearity

satisfying the conclusion of Theorem 2 with these D_{in} , ρ , and ϵ . For a network $\mathcal{N} \in \mathcal{N}_{nc}^{D_{in}, D_{out}}$, we write the map $\langle \mathcal{N} \rangle^\sigma = ((\langle \mathcal{N} \rangle^\sigma)_1, \dots, (\langle \mathcal{N} \rangle^\sigma)_{D_{out}})$ in terms of the coordinate functions $(\langle \mathcal{N} \rangle^\sigma)_j : \mathbb{R}^{D_{in}} \rightarrow \mathbb{R}$, $j \in \{1, \dots, D_{out}\}$. Now, let $\mathcal{N}_1, \mathcal{N}_2 \in \mathcal{N}_{nc}^{D_{in}, D_{out}}$ be networks such that $\langle \mathcal{N}_1 \rangle^\sigma(x) = \langle \mathcal{N}_2 \rangle^\sigma(x)$, for all $x \in \mathbb{R}^{D_{in}}$, and suppose by way of contradiction that they are non-isomorphic. We construct a network \mathcal{M} containing both \mathcal{N}_1 and \mathcal{N}_2 as subnetworks (a precise definition of “subnetwork” is given in Section III, Definition 9). It follows that \mathcal{M} contains subnetworks $\mathcal{M}_{m,j} \in \mathcal{N}_{nc}^{D_{in}, 1}$ with maps satisfying $\langle \mathcal{M}_{m,j} \rangle^\sigma = (\langle \mathcal{N}_m \rangle^\sigma)_j$, for $m \in \{1, 2\}$ and $j \in \{1, \dots, D_{out}\}$. We then show that, as a consequence of \mathcal{N}_1 and \mathcal{N}_2 being non-isomorphic, there exists a $j \in \{1, \dots, D_{out}\}$ such that $\mathcal{M}_{1,j}$ and $\mathcal{M}_{2,j}$ are non-isomorphic. But then

$$0 \cdot \mathbf{1} + \langle \mathcal{M}_{1,j} \rangle^\sigma - \langle \mathcal{M}_{2,j} \rangle^\sigma = (\langle \mathcal{N}_1 \rangle^\sigma)_j - (\langle \mathcal{N}_2 \rangle^\sigma)_j = 0,$$

which stands in contradiction to Theorem 2. This completes the proof of Theorem 1.

The proof of Theorem 2 is significantly more involved, as it requires extensive “fine tuning” of the function σ . Let $\sigma : \mathcal{D} \rightarrow \mathbb{C}$ be as in the statement of Theorem 2. In addition to the properties stated in Theorem 2, the function σ we construct exhibits the following convenient structural properties:

- 1) The domain $\mathcal{D} \subset \mathbb{C}$ of σ is the complement of an (infinite) discrete set of poles,
- 2) σ is i -periodic, i.e., $\sigma(z + i) = \sigma(z)$, for all $z \in \mathcal{D}$, and
- 3) for any network $\mathcal{N} \in \mathcal{N}^{1,1}$, the natural domain $\mathcal{D}_{\langle \mathcal{N} \rangle^\sigma} \subset \mathbb{C}$ of $\langle \mathcal{N} \rangle^\sigma$, viewed as a holomorphic function, is the complement of a closed countable subset of \mathbb{C} , and therefore a connected open set.

These three properties are all satisfied by the function $\tanh(\pi \cdot)$, and are essentially the key insight leading to Fefferman’s identifiability result in [3], which establishes that, under the genericity conditions stated in Assumptions 1, a neural network can be read off from the asymptotic (as the imaginary part of the argument tends to infinity) locations of the singularities of the map it realizes under the \tanh nonlinearity. The properties 1) – 3) will be key to our results as well, but instead of studying the set of singularities of the map in its own right, our proof of Theorem 2 will proceed by contradiction. The proof consists of three steps that we call *amalgamation*, *input splitting*, and *input anchoring*, and involves the use of analytic continuation, graph-theoretic constructions, and Kronecker’s theorem [16], the latter two of which are novel tools in this context and signify a significant departure from Fefferman’s proof technique in [3]. We now briefly describe the proof of Theorem 2 according to the aforementioned program. Suppose that $\mathcal{N}_1, \dots, \mathcal{N}_n$ are pairwise non-isomorphic neural networks satisfying the no-clones condition. For the sake of simplicity of this informal discussion, we assume that $L_1 = L_2 = \dots = L_n$, $D_0^1 = D_0^2 = \dots = D_0^n = 1$, and $D_{L_1}^1 = D_{L_2}^2 = \dots = D_{L_n}^n = 1$. By way of contradiction, we suppose that there exists a nontrivial linear combination such that $\lambda_0 \mathbf{1}(x) + \sum_{j=1}^n \lambda_j \mathcal{N}_j^\sigma(x) = 0$, for all $x \in \mathbb{R}$.

Amalgamation. In Section III we construct a neural network $\mathcal{M} \in \mathcal{N}_{nc}^{1,n}$, called the amalgam of $\{\mathcal{N}_j\}_{j=1}^n$, containing each \mathcal{N}_j as a subnetwork. In particular, we have $(\langle \mathcal{M} \rangle^\sigma)_j = \langle \mathcal{N}_j \rangle^\sigma$, for all $j \in \{1, \dots, n\}$. The linear dependence of $\{\langle \mathcal{N}_j \rangle^\sigma\}_{j=1}^n \cup \{\mathbf{1}\}$ thus translates to

$$\lambda_0 + \sum_{j=1}^n \lambda_j (\langle \mathcal{M} \rangle^\sigma)_j(z) = 0, \tag{1}$$

for all $z \in \mathbb{R}$. By our construction of σ , the natural domains $\mathcal{D}_{\langle \mathcal{N}_j \rangle^\sigma} = \mathcal{D}_{(\langle \mathcal{M} \rangle^\sigma)_j}$ are complements of closed countable sets, and hence, by analytic continuation, (1) is valid for all $z \in \bigcap_{j=1}^n \mathcal{D}_{\langle \mathcal{N}_j \rangle^\sigma}$.

Now define \mathcal{M} to be the set of all neural networks in $\bigcup_{m=1}^n \mathcal{N}_{nc}^{1,m}$ with linear dependency as in (1) between the output functions and the constant function. Note that \mathcal{M} is nonempty, simply as $\mathcal{M} \in \mathcal{M}$. We then fix a network $\mathcal{M}' \in \mathcal{M}$ of minimum size (the precise definition of size will be given in the proof of Theorem 4). Write $(1, D_1^{\mathcal{M}'}, \dots, D_m^{\mathcal{M}'})$ for the layout of \mathcal{M}' , and let $(\omega_1, \dots, \omega_{D_1^{\mathcal{M}'}})$ be the weights of the first layer of \mathcal{M}' (i.e., the entries of W^1 according to Definition 1). At this point the proof splits into two cases, depending on whether there exist $j, j' \in \{1, \dots, D_1^{\mathcal{M}'}\}$, $j \neq j'$, such that $\omega_j/\omega_{j'}$ is irrational.

Input splitting, the easy case. Provided there do exist such j and j' , we use Kronecker's theorem [16] and the properties (i) – (iii) of σ to construct a network $\mathcal{M}'' \in \mathcal{M}$ with layout $(k, D_1^{\mathcal{M}'}, \dots, D_m^{\mathcal{M}'})$, for some $k \in \{2, \dots, D_1^{\mathcal{M}'}\}$, and first-layer weights $\widetilde{W}^1 \in \mathbb{R}^{D_1^{\mathcal{M}'} \times k}$ such that the first k rows of \widetilde{W}^1 form a $k \times k$ identity matrix.

Input anchoring. We then construct a third network $\mathcal{N} \in \mathcal{M}$, obtained by fixing $k-1$ of the k inputs of \mathcal{M}'' to specific real numbers, and “cutting out” all the parts of the network whose contributions to the output map have become constant in the process. The resulting network \mathcal{N} will be a network in \mathcal{M} of size smaller than \mathcal{M}' , which contradicts the minimality of \mathcal{M}' , and thereby completes the proof.

Input splitting, the hard case. If, however, all the ratios $\omega_j/\omega_{j'}$, $j \neq j'$ are rational, the input splitting construction described above cannot be carried out. This problem will be remedied by further refining our initial construction of σ . Specifically, we will ensure that the real parts of the poles of σ form a subset of \mathbb{R} satisfying what we call the *self-avoiding property*, to be introduced in Section V. This will enable an alternative construction of a network \mathcal{M}'' with at least two inputs. The resulting \mathcal{M}'' will, however, not be a neural network in the sense of Definition 1, but rather a generalized network in the sense of Definition 8, to be introduced in Section III.

Input anchoring. Finally, we apply an input anchoring procedure to \mathcal{M}'' similar to the one described above. Even though now \mathcal{M}'' is not a network in the sense of Definition 1, the input anchoring procedure will result in a network $\mathcal{N} \in \mathcal{M}$ which is a network in the sense of Definition 1, and is of smaller size than \mathcal{M}' , again completing the proof by contradiction.

We conclude this section by laying out the organization of the remainder of the paper. In Section III we develop a graph-theoretic framework needed to define amalgams of neural networks and several other technical concepts. In Section IV we state results from complex analysis and Kronecker's theorem needed in arguments involving analytic continuation and input splitting, respectively. The proofs of these results are relegated to the Appendix. In Section V we discuss the fine structural properties of the function σ constructed in the proof of Theorem 2. Finally, Section VI contains the proofs of our two main results.

III. DIRECTED ACYCLIC GRAPHS, GENERAL NEURAL NETWORKS, AND NEURAL NETWORK AMALGAMS

As already mentioned, in the proof of Theorem 2 we will work with a form of neural networks that does not fit in with Definitions 1 and 2. In order to accommodate this notion of neural networks, and to lighten the manipulations needed to formalize the aforementioned techniques of amalgamation and input anchoring, we introduce a graph-theoretic framework.

We start by introducing the concept of a directed acyclic graph (DAG), commonly encountered in the graph theory literature [17].

Definition 6 (Directed acyclic graph).

- A directed graph is an ordered pair $G = (V, E)$ where V is a finite set of nodes, and $E \subset V \times V$ is a set of directed edges.
- A directed cycle of a directed graph G is a set $\{v_1, \dots, v_k\} \subset V$ such that, for every $j \in \{1, \dots, k\}$, $(v_j, v_{j+1}) \in E$, where we set $v_{k+1} := v_1$.
- A directed graph G is said to be a directed acyclic graph (DAG) if it has no directed cycles.

We interpret an edge (v, \tilde{v}) as an arrow connecting the nodes v and \tilde{v} and pointing at \tilde{v} .

Definition 7 (Parent set, input nodes, and node level). Let $G = (V, E)$ be a DAG.

- We define the parent set of a node by $\text{par}(v) = \{\tilde{v} : (\tilde{v}, v) \in E\}$.
- We say that $v \in V$ is an input node if $\text{par}(v) = \emptyset$, and we write $\text{In}(G)$ for the set of input nodes.
- We define the level $\text{lv}(v)$ of a node $v \in V$ recursively as follows. If $\text{par}(v) = \emptyset$, we set $\text{lv}(v) = 0$. If $\text{par}(v) = \{v_1, v_2, \dots, v_k\}$ and $\text{lv}(v_1), \text{lv}(v_2), \dots, \text{lv}(v_k)$ are defined, we set $\text{lv}(v) = \max\{\text{lv}(v_1), \text{lv}(v_2), \dots, \text{lv}(v_k)\} + 1$.

Since the graph G in Definition 7 is assumed to be acyclic, the level is well-defined for all nodes of G . We are now ready to introduce our generalized definition of a neural network.

Definition 8. A general feed-forward neural network (GFNN) is an ordered sextuple $\mathcal{N} = (V, E, V_{in}, V_{out}, \Omega, \Theta)$, where

- $G = (V, E)$ is a DAG, called the architecture of \mathcal{N} ,
- $V_{in} = \text{In}(G)$ is the set of inputs of \mathcal{N} ,
- $V_{out} \subset V \setminus V_{in}$ is the set of outputs of \mathcal{N} ,
- $\Omega = \{\omega_{\tilde{v}v} \in \mathbb{R} \setminus \{0\} : (v, \tilde{v}) \in E\}$ is the set of weights of \mathcal{N} , and
- $\Theta = \{\theta_v \in \mathbb{R} : v \in V \setminus V_{in}\}$ is the set of biases of \mathcal{N} .

The depth of a GFNN is defined as $L(\mathcal{N}) = \max\{\text{lv}(v) : v \in V\}$.

When translating from Definition 1 to Definition 8, we will interpret a zero weight $W_{jk}^\ell = 0$ simply as the absence of a directed edge between the nodes concerned, hence we do not allow the edges of a GFNN to have zero weight. If V^1 and V^2 are the sets of nodes of GFNNs \mathcal{N}_1 and \mathcal{N}_2 , respectively, and $v \in V^1 \cap V^2$, we will say that \mathcal{N}_1 and \mathcal{N}_2 share the node v . When dealing with several networks sharing a node v , we will write $\text{par}_{\mathcal{N}}(v)$ for the parent set of v in the architecture (V, E) of \mathcal{N} , to avoid ambiguity. Note that the set of outputs of a GFNN can be an arbitrary subset of the non-input nodes. In particular, V_{out} can include nodes w with $\text{lv}(w) < L(\mathcal{N})$. Related to the concept of the parent set of a node is the concept of a subnetwork introduced next.

Definition 9 (Subnetwork and ancestor subnetwork). Let $\mathcal{N} = (V, E, V_{in}, V_{out}, \Omega, \Theta)$ be a GFNN. A *subnetwork* of \mathcal{N} is a GFNN $\mathcal{N}' = (V', E', V'_{in}, V'_{out}, \Omega', \Theta')$ such that there exists a set $S \subset V$ so that

- (i) $V' = \{v \in V : v \in \text{par}^r(u) \text{ for some } r \geq 0\}$, where, for a set $W \subset V$, we define $\text{par}^0(W) = W$ and $\text{par}^r(W) = \bigcup_{s \in W} \text{par}^{r-1}(\text{par}(s))$, for $r \geq 1$.

- (ii) $E' = \{(v, \tilde{v}) \in E : v, \tilde{v} \in V'\}$,
- (iii) $V'_{in} = V_{in} \cap V'$,
- (iv) $\Omega' = \{\omega_{\tilde{v}v} : (v, \tilde{v}) \in E'\}$, and
- (v) $\Theta' = \{\theta_v : v \in V'\}$.

If additionally $V'_{out} = S$, then \mathcal{N}' is uniquely specified by S . In this case we say that \mathcal{N}' is the *ancestor subnetwork of S in \mathcal{N}* , and write $\mathcal{N}(S)$ for this network.




Fig. 1: A GFNN of depth 3 with input nodes $\{u_1, u_2, u_3\}$ and output nodes $\{w_1, w_2, w_3\}$. The node levels are indicated by the numbers inside the circles. Note that the output node w_3 is not a “final node”, i.e., it has outgoing edges. As there is an edge (u_3, v) connecting nodes of non-consecutive levels, the network is not layered.

Definition 10. A layered feed-forward neural network (LFNN) is a GFNN satisfying $\text{lv}(\tilde{v}) = \text{lv}(v) + 1$, for all $(v, \tilde{v}) \in E$.

For an example of a GFNN that is not layered, see Figure 1. We notice that LFNNs correspond to neural networks as specified by Definition 1, with the nodes of level ℓ corresponding to the ℓ -th network layer. Specifically, if $\mathcal{N} = (V, E, V_{in}, V_{out}, \Omega, \Theta)$ is a LFNN, we can label the nodes $\{v \in V : \text{lv}(v) = \ell\}$ by v_j^ℓ , $j = 1, \dots, D_\ell$, and let $\theta_j^\ell = \theta_{v_j^\ell}$, $W_{jk}^\ell = \omega_{v_j^\ell v_k^{\ell-1}}$ when $(k, j) \in E$ and $W_{jk}^\ell = 0$ else. Apropos, this correspondence is the reason for the indices of the weight $\omega_{\tilde{v}v}$ associated with the edge (v, \tilde{v}) of a GFNN appearing in “reverse order”. The following definition generalizes Definition 2 to GFNNs.

Definition 11 (Output maps of nodes and networks). Let $\mathcal{N} = (V, E, V_{in}, V_{out}, \Omega, \Theta)$ be a GFNN, and let $\rho : \mathbb{R} \rightarrow \mathbb{R}$ be a nonlinearity. The map realized by a node $v \in V$ under ρ is the function $\langle v \rangle^\rho : \mathbb{R}^{V_{in}} \rightarrow \mathbb{R}$ defined recursively as follows:

- If $v \in V_{in}$, set $\langle v \rangle^\rho(\mathbf{t}) = t_v$, for all $\mathbf{t} = (t_u)_{u \in V_{in}} \in \mathbb{R}^{V_{in}}$.
- Otherwise set $\langle v \rangle^\rho(\mathbf{t}) = \rho\left(\sum_{u \in \text{par}(v)} \omega_{vu} \cdot \langle u \rangle^\rho(\mathbf{t}) + \theta_v\right)$, for all $\mathbf{t} \in \mathbb{R}^{V_{in}}$.

The map realized by \mathcal{N} under ρ is the function $\langle \mathcal{N} \rangle^\rho : \mathbb{R}^{V_{in}} \rightarrow \mathbb{R}^{V_{out}}$ given by $\langle \mathcal{N} \rangle^\rho = (\langle v \rangle^\rho)_{v \in V_{out}}$. When dealing with several networks we will write $\langle v \rangle^{\rho, \mathcal{N}}$ for the map realized by v in \mathcal{N} , to avoid ambiguity.

We will treat nodes $v \in V$ only as “handles”, and never as variables or functions. This is relevant when dealing with several networks with shared nodes, such as depicted in Figure 2. On the other hand, the output map $\langle v \rangle^\rho$ realized by v is a function. In the special case when the nonlinearity is




Fig. 2: The network \mathcal{N}_1 consists of the elements in red and black, and \mathcal{N}_2 consists of the elements in blue and black. Thus \mathcal{N}_1 and \mathcal{N}_2 share the nodes u_1 , u_2 , and w_1 , even though the functions $\langle w_1 \rangle^{\rho, \mathcal{N}_1}$ and $\langle w_1 \rangle^{\rho, \mathcal{N}_2}$ may be “completely unrelated”.

holomorphic on a neighborhood of \mathbb{R} , the output maps realized by the nodes of a network will extend to holomorphic functions on their natural domains, as given by the following definition.

Definition 12 (Natural domain). Let $\mathcal{N} = (V, E, V_{in}, V_{out}, \Omega, \Theta)$ be a GFNN, and let $\sigma : \mathcal{D}_\sigma \rightarrow \mathbb{C}$ be a function holomorphic on an open domain $\mathcal{D}_\sigma \supset \mathbb{R}$ and such that $\sigma(\mathbb{R}) \subset \mathbb{R}$. For a node $v \in V$, we define the natural domain $\mathcal{D}_{\langle v \rangle^\sigma} \subset \mathbb{C}^{V_{in}}$ and extend the definition of the function $\langle v \rangle^\sigma : \mathcal{D}_{\langle v \rangle^\sigma} \rightarrow \mathbb{C}$ recursively as follows:

- For $v \in V_{in}$, let $\mathcal{D}_{\langle v \rangle^\sigma} = \mathbb{C}^{V_{in}}$, and set $\langle v \rangle^\sigma(\mathbf{z}) = z_v$, for all $\mathbf{z} = (z_u)_{u \in V_{in}} \in \mathbb{C}^{V_{in}}$.
- Otherwise, set $\mathcal{D}_{\langle v \rangle^\sigma} = \left\{ \mathbf{z} \in \bigcap_{u \in \text{par}(v)} \mathcal{D}_{\langle u \rangle^\sigma} : \sum_{u \in \text{par}(v)} \omega_{vu} \langle u \rangle^\sigma(\mathbf{z}) + \theta_v \in \mathcal{D}_\sigma \right\}$, and let $\langle v \rangle^\sigma(\mathbf{z}) = \sigma \left(\sum_{u \in \text{par}(v)} \omega_{vu} \cdot \langle u \rangle^\sigma(\mathbf{z}) + \theta_v \right)$, for all $\mathbf{z} \in \mathcal{D}_{\langle v \rangle^\sigma}$.

It follows that the natural domain $\mathcal{D}_{\langle u \rangle^\sigma}$ of a node u is open, as it is the preimage of an open set with respect to a continuous map. Moreover, the output map $\langle u \rangle^\sigma$ realized by u is holomorphic on $\mathcal{D}_{\langle u \rangle^\sigma}$, as it is given explicitly by a concatenation of affine maps and the nonlinearity σ , which are themselves holomorphic functions.

The following definition is a straightforward generalization of Definition 5.

Definition 13 (Clone pairs and the no-clones condition). Let $\mathcal{N} = (V, E, V_{in}, V_{out}, \Omega, \Theta)$ be a GFNN. We say that the nodes $v_1, v_2 \in V$, $v_1 \neq v_2$, are clones if $\text{par}(v_1) = \text{par}(v_2)$, $\theta_{v_1} = \theta_{v_2}$, and $\forall u \in \text{par}(v_1)$, $\omega_{v_1 u} = \omega_{v_2 u}$. We say that \mathcal{N} satisfies the no-clones condition (or briefly, \mathcal{N} is clones-free), if no two nodes $v_1, v_2 \in V$, $v_1 \neq v_2$, are clones.

The following definition generalizes Definition 4 to GFNNs, and introduces two new concepts, termed extensional isomorphism and faithful isomorphism, which will play an important technical role throughout the remainder of the paper.

Definition 14 (Extensional and faithful isomorphisms of GFNNs). Let $\mathcal{N}^1 = (V^1, E^1, V_{in}, V_{out}^1, \Omega^1, \Theta^1)$ and $\mathcal{N}^2 = (V^2, E^2, V_{in}, V_{out}^2, \Omega^2, \Theta^2)$ be GFNNs with the same input nodes V_{in} .

– We say that \mathcal{N}^1 and \mathcal{N}^2 are *extensionally isomorphic*, and write $\mathcal{N}^1 \stackrel{e}{\sim} \mathcal{N}^2$, if there exists a bijection $\pi : V^1 \rightarrow V^2$, called an *extensional isomorphism*, such that the following holds:

- (i) π restricted to V_{in} is the identity map,
- (ii) $\pi(V_{out}^1) = V_{out}^2$,
- (iii) for all $(v, \tilde{v}) \in E^1$, we have $\omega_{\pi(\tilde{v})\pi(v)}^2 = \omega_{\tilde{v}v}^1$, and
- (iv) for all $v \in V^1 \setminus V_{in}$, we have $\theta_{\pi(v)}^2 = \theta_v^1$.

– We say that \mathcal{N}^1 and \mathcal{N}^2 are *faithfully isomorphic*, and write $\mathcal{N}^1 \stackrel{f}{\sim} \mathcal{N}^2$, if they are extensionally isomorphic via $\pi : V^1 \rightarrow V^2$ with the following additional property:

- (v) $V_{out}^1 = V_{out}^2$, and π restricted to V_{out}^1 is the identity map.

In this case we call π a *faithful isomorphism*.

Remark. The concept of faithful isomorphisms in Definition 14 generalizes that of isomorphisms according to Definition 4. It is easily seen that extensional isomorphism is an equivalence relation on the set of all GFNNs with the same input nodes, whereas faithful isomorphism is an equivalence relation on the set of all GFNNs with the same input and output nodes. Furthermore, if $\mathcal{N}^1 \stackrel{e}{\sim} \mathcal{N}^2$ via $\pi : V^1 \rightarrow V^2$, then we have $\langle \pi(v) \rangle^{\rho, \mathcal{N}^2} = \langle v \rangle^{\rho, \mathcal{N}^1}$, for all $v \in V^1$ and any nonlinearity ρ , and if additionally $\mathcal{N}^1 \stackrel{f}{\sim} \mathcal{N}^2$, then $\langle \mathcal{N}^1 \rangle^\rho = \langle \mathcal{N}^2 \rangle^\rho$.

The following definition introduces the non-degeneracy property of a GFNN, which corresponds to the absence of spurious nodes, i.e., nodes that do not contribute to the map realized by the GFNN (with respect to an arbitrary nonlinearity). In the special case of LFNNs considered in the introduction, this property corresponds to the requirement that no matrix W^ℓ in Definition 1 has an identically zero row or column.

Definition 15 (Non-degeneracy). We say that a GFNN $\mathcal{N} = (V, E, V_{in}, V_{out}, \Omega, \Theta)$ is non-degenerate if $V = V^{\mathcal{N}(V_{out})}$, where $V^{\mathcal{N}(V_{out})}$ is the set of nodes of the ancestor subnetwork of V_{out} in \mathcal{N} . Networks that are not non-degenerate are referred to as degenerate.

Informally, a network is non-degenerate if its every node “leads up” to at least one output. This notion is best understood with the help of examples as in Figure 3.

We are now ready to introduce the concept of amalgams of LFNNs.

Definition 16 (Amalgam of two layered neural networks). Let $\mathcal{N}_1 = (V^1, E^1, V_{in}, V_{out}^1, \Omega^1, \Theta^1)$ and $\mathcal{N}_2 = (V^2, E^2, V_{in}, V_{out}^2, \Omega^2, \Theta^2)$ be non-degenerate clones-free LFNNs with the same input set V_{in} .

– Let $\mathcal{A} = (V^{\mathcal{A}}, E^{\mathcal{A}}, V_{in}, V_{out}^{\mathcal{A}}, \Omega^{\mathcal{A}}, \Theta^{\mathcal{A}})$ be a non-degenerate LFNN with the following properties:

- (i) There exist injective maps $\pi_1 : V^1 \rightarrow \pi_1(V^1) \subset V^{\mathcal{A}}$ and $\pi_2 : V^2 \rightarrow \pi_2(V^2) \subset V^{\mathcal{A}}$ such that the networks \mathcal{N}_1 and \mathcal{N}_2 are extensionally isomorphic to the ancestor subnetworks $\mathcal{A}(\pi_1(V_{out}^1))$ and $\mathcal{A}(\pi_2(V_{out}^2))$ via π_1 and π_2 , respectively.
- (ii) $V^{\mathcal{A}} = \pi_1(V^1) \cup \pi_2(V^2)$ and $V_{out}^{\mathcal{A}} = \pi_1(V_{out}^1) \cup \pi_2(V_{out}^2)$.

We then say that \mathcal{A} is a proto-amalgam of \mathcal{N}_1 and \mathcal{N}_2 .

– If \mathcal{A} is a clones-free proto-amalgam of \mathcal{N}_1 and \mathcal{N}_2 , we say that \mathcal{A} is an amalgam of \mathcal{N}_1 and \mathcal{N}_2 .




Fig. 3: These GFNNs are degenerate owing to the presence of spurious nodes (in red) that do not affect the map of the output node w_1 . Such networks obviously need to be excluded from consideration when discussing identifiability from the map realized by the network, as its “spurious parts” cannot be inferred from the map it realizes.

Proposition 1. *Let $\mathcal{N}_1 = (V^1, E^1, V_{in}, V_{out}^1, \Omega^1, \Theta^1)$ and $\mathcal{N}_2 = (V^2, E^2, V_{in}, V_{out}^2, \Omega^2, \Theta^2)$ be non-degenerate clones-free LFNNs with a shared input set V_{in} . Then there exists an amalgam \mathcal{A} of \mathcal{N}_1 and \mathcal{N}_2 . Moreover, the amalgam is unique up to extensional isomorphisms.*

As asserted in Proposition 1 (whose proof is deferred to the Appendix), an amalgam of two given non-degenerate clones-free LFNNs \mathcal{N}_1 and \mathcal{N}_2 always exists and is unique up to extensional isomorphisms. With slight abuse of notation, we will write $\mathcal{N}_1 \vee \mathcal{N}_2$ for an arbitrary element of the equivalence class (induced by $\overset{\sim}{\sim}$) of all the amalgams of \mathcal{N}_1 and \mathcal{N}_2 . A concrete example of an amalgam construction is provided in Figure 4. Having defined the amalgam of two non-degenerate clones-free LFNNs, we define the amalgam of any finite collection $\mathcal{N}_1, \dots, \mathcal{N}_n$ of non-degenerate clones-free LFNNs according to

$$\bigvee_{k=1}^n \mathcal{N}_k = \mathcal{N}_1 \vee \mathcal{N}_2 \vee \dots \vee \mathcal{N}_n := (\dots (\mathcal{N}_1 \vee \mathcal{N}_2) \vee \dots) \vee \mathcal{N}_n.$$

By Definition 16, $\bigvee_{k=1}^n \mathcal{N}_k$ is a non-degenerate clones-free LFNN. Moreover, there exist extensional isomorphisms $\pi_j : \mathcal{N}_j \rightarrow \pi_j(\mathcal{N}_j) \subset \bigvee_{k=1}^n \mathcal{N}_k$, for $j \in \{1, \dots, n\}$, and we have $\langle \pi_j(v) \rangle^{\rho, \bigvee_{k=1}^n \mathcal{N}_k} = \langle v \rangle^{\rho, \mathcal{N}_j}$, for $j \in \{1, \dots, n\}$, $v \in V^{\mathcal{N}_j}$, and any nonlinearity ρ .

We are now in a position to prove two lemmas that form the basis for the proof of Theorem 2. The first lemma formalizes the idea of combining multiple pairwise non-isomorphic single-output networks with linearly dependent output maps into one multiple-output network with linear dependency among the maps of its output nodes.

Lemma 1. *Let $\mathcal{N}_1, \mathcal{N}_2, \dots, \mathcal{N}_n$ be non-degenerate, clones-free LFNNs with a shared input set V_{in} and the same single output node $\{v_{out}\}$. Furthermore, assume that no two networks $\mathcal{N}_{j_1}, \mathcal{N}_{j_2}$, $j_1 \neq j_2$, are extensionally isomorphic. Let ρ be a nonlinearity and suppose that $\mathbf{1}, \langle \mathcal{N}_1 \rangle^\rho, \langle \mathcal{N}_2 \rangle^\rho, \dots, \langle \mathcal{N}_n \rangle^\rho$ are linearly dependent as functions $\mathbb{R}^{V_{in}} \rightarrow \mathbb{R}$. Then there exists a non-degenerate clones-free LFNN $\mathcal{M} = (V^{\mathcal{M}}, E^{\mathcal{M}}, V_{in}^{\mathcal{M}}, V_{out}^{\mathcal{M}}, \Omega^{\mathcal{M}}, \Theta^{\mathcal{M}})$ (obtained by modifying $\bigvee_{k=1}^n \mathcal{N}_k$) with a single input node $V_{in}^{\mathcal{M}} = \{v_{in}\}$, such that $\{\langle w \rangle^\rho : w \in V_{out}^{\mathcal{M}}\} \cup \{\mathbf{1}\}$ is a linearly dependent set of functions from \mathbb{R} to \mathbb{R} .*




Fig. 4: *Top*: LFNNs \mathcal{N}_1 and \mathcal{N}_2 to be amalgamated, with their weights next to the edges and the biases inside the nodes. *Middle*: A proto-amalgam \mathcal{N} of the two LFNNs, obtained by putting \mathcal{N}_1 and \mathcal{N}_2 “side by side”. This network is not an amalgam of \mathcal{N}_1 and \mathcal{N}_2 , as there is a clone pair (c_1, c_2) . *Bottom*: The network \mathcal{N} can be modified by deleting the node c_2 and “grafting” its outgoing edge to c_1 . The resulting network \mathcal{A} is now a clones-free proto-amalgam of \mathcal{N}_1 and \mathcal{N}_2 , and is thus the amalgam $\mathcal{N}_1 \vee \mathcal{N}_2$. For general LFNNs \mathcal{N}_1 and \mathcal{N}_2 , this “deleting and grafting” process can be repeated until there are no clone pairs left.

Proof: We first create a new node v_{in} and select an arbitrary set $\{\omega_{\tilde{v}v_{in}} : \tilde{v} \in V_{in}\} \subset \mathbb{R} \setminus \{0\}$ of cardinality $\#V_{in}$. Now, we enlarge each \mathcal{N}_j to a new network $\tilde{\mathcal{N}}_j$ by gluing the node v_{in} to the set V_{in} through the edges $\{(v_{in}, \tilde{v}) : \tilde{v} \in V_{in}\}$ along with the corresponding weights $\omega_{\tilde{v}v_{in}}$. The nodes $v \in V_{in}$ are non-input nodes of the $\tilde{\mathcal{N}}_j$, as their parent sets $\text{par}_{\tilde{\mathcal{N}}_j}(v) = \{v_{in}\}$ are non-empty, and we set their biases θ_v to 0. The node v_{in} is now the shared single input of the networks $\tilde{\mathcal{N}}_j$, $j = 1, \dots, n$. Note that, as the networks \mathcal{N}_j are clones-free, and the weights $\omega_{\tilde{v}v_{in}}$ are distinct, the networks $\tilde{\mathcal{N}}_j$ are clones-free by assumption. Further, since \mathcal{N}_j , $j \in \{1, \dots, n\}$, are pairwise non-isomorphic, so are the $\tilde{\mathcal{N}}_j$, $j \in \{1, \dots, n\}$. We now construct a network \mathcal{M} by amalgamating $\tilde{\mathcal{N}}_j$, $j = 1, \dots, n$, according to $\mathcal{M} = (\dots(\tilde{\mathcal{N}}_1 \vee \tilde{\mathcal{N}}_2) \vee \dots) \vee \tilde{\mathcal{N}}_n$. Denote by $\pi_j : V^{\tilde{\mathcal{N}}_j} \rightarrow \pi_j(V^{\tilde{\mathcal{N}}_j}) \subset V^{\mathcal{M}}$ the extensional isomorphism between $\tilde{\mathcal{N}}_j$ and the corresponding subnetwork of \mathcal{M} , and let $w_j = \pi_j(v_{out})$ be the node of \mathcal{M} corresponding to the output node of \mathcal{N}_j . We claim that $w_{j_1} \neq w_{j_2}$, for $j_1 \neq j_2$. To see

this, take j_1, j_2 such that $w_{j_1} = w_{j_2}$, i.e., $\pi_{j_1}(v_{out}) = \pi_{j_2}(v_{out})$. Then, by Property (i) of Definition 16, $\tilde{\mathcal{N}}_{j_1}(v_{out}) \stackrel{e}{\sim} \tilde{\mathcal{N}}_{j_2}(v_{out})$, and therefore $\mathcal{N}_{j_1}(v_{out}) \stackrel{e}{\sim} \mathcal{N}_{j_2}(v_{out})$ as well. But $\mathcal{N}_{j_1}(v_{out}) = \mathcal{N}_{j_1}$ and $\mathcal{N}_{j_2}(v_{out}) = \mathcal{N}_{j_2}$ by the non-degeneracy assumption, and hence $\mathcal{N}_{j_1} \stackrel{e}{\sim} \mathcal{N}_{j_2}$. It follows that $j_1 = j_2$, as $\mathcal{N}_j, j = 1, \dots, n$, are assumed to be pairwise non-isomorphic. Thus the w_j are, indeed, distinct nodes of \mathcal{M} , and we have $V_{out}^{\mathcal{M}} = \{w_1, w_2, \dots, w_n\}$. As $\mathbf{1}, \langle \mathcal{N}_1 \rangle^\rho, \langle \mathcal{N}_2 \rangle^\rho, \dots, \langle \mathcal{N}_n \rangle^\rho$ are linearly dependent by assumption, there exists a nonzero vector $(c, \lambda_1, \lambda_2, \dots, \lambda_n) \in \mathbb{R}^{n+1}$ such that $(c\mathbf{1} + \sum_{j=1}^n \lambda_j \langle \mathcal{N}_j \rangle^\rho)((t_v)_{v \in V_{in}}) = 0$, for all $(t_v)_{v \in V_{in}} \in \mathbb{R}^{V_{in}}$. We then have

$$\begin{aligned} \left(c\mathbf{1} + \sum_{j=1}^n \lambda_j \langle w_j \rangle^{\rho, \mathcal{M}} \right)(t) &= \left(c\mathbf{1} + \sum_{j=1}^n \lambda_j \langle \pi_j(v_{out}) \rangle^{\rho, \mathcal{M}} \right)(t) = \left(c\mathbf{1} + \sum_{j=1}^n \lambda_j \langle v_{out} \rangle^{\rho, \tilde{\mathcal{N}}_j} \right)(t) \\ &= \left(c\mathbf{1} + \sum_{j=1}^n \lambda_j \langle \mathcal{N}_j \rangle^\rho \right)((\omega_{\tilde{v}v_{in}} t)_{\tilde{v} \in V_{in}}) = 0, \end{aligned}$$

for all $t \in \mathbb{R}$. This establishes that $\{\langle w_1 \rangle^{\rho, \mathcal{M}}, \langle w_2 \rangle^{\rho, \mathcal{M}}, \dots, \langle w_n \rangle^{\rho, \mathcal{M}}\} \cup \{\mathbf{1}\}$ is a linearly dependent set, so \mathcal{M} is the desired network. \blacksquare

Before stating the next lemma, we describe the procedure of input anchoring, which is a method for selecting and modifying a subnetwork of a non-degenerate GFNN in a manner that preserves linear dependencies between the maps realized by the output nodes of the original network. Concretely, let $\mathcal{M} = (V^{\mathcal{M}}, E^{\mathcal{M}}, V_{in}^{\mathcal{M}}, V_{out}^{\mathcal{M}}, \Omega^{\mathcal{M}}, \Theta^{\mathcal{M}})$ be a non-degenerate, clones-free GFNN with input nodes $V_{in}^{\mathcal{M}} = \{v_1^0, \dots, v_{D_0}^0\}$, $D_0 \geq 2$. For specificity, let w.l.o.g. $v_{D_0}^0$ be the input node to be anchored, and let $a \in \mathbb{R}$ be the value $v_{D_0}^0$ is anchored to. Furthermore, let ρ be a nonlinearity. We seek to construct a network $\mathcal{M}_a = (V^{\mathcal{M}_a}, E^{\mathcal{M}_a}, V_{in}^{\mathcal{M}_a}, V_{out}^{\mathcal{M}_a}, \Omega^{\mathcal{M}_a}, \Theta^{\mathcal{M}_a})$ with $V_{in}^{\mathcal{M}_a} = \{v_1^0, \dots, v_{D_0-1}^0\}$ and $V_{out}^{\mathcal{M}_a} = V_{out}^{\mathcal{M}} \cap V^{\mathcal{M}_a}$ satisfying the following two properties:

(IA-1) For all $w \in V_{out}^{\mathcal{M}_a}$,

$$\langle w \rangle^{\rho, \mathcal{M}_a}(t_1, t_2, \dots, t_{D_0-1}) = \langle w \rangle^{\rho, \mathcal{M}}(t_1, t_2, \dots, t_{D_0-1}, a),$$

for all $(t_1, t_2, \dots, t_{D_0-1}) \in \mathbb{R}^{D_0-1}$ (after identifying $\mathbb{R}^{V_{in}}$ with \mathbb{R}^{D_0}).

(IA-2) For all $w \in V_{out}^{\mathcal{M}} \setminus V_{out}^{\mathcal{M}_a}$, the function $\mathbb{R}^{D_0-1} \rightarrow \mathbb{R}$ given by

$$(t_1, t_2, \dots, t_{D_0-1}) \mapsto \langle w \rangle^{\rho, \mathcal{M}}(t_1, t_2, \dots, t_{D_0-1}, a)$$

is constant, and we denote its value by $\langle w \rangle^{\rho, \mathcal{M}}(a)$.

As $V^{\mathcal{M}_a} \subset V^{\mathcal{M}} \setminus \{v_{D_0}^0\}$, the network \mathcal{M}_a will, indeed, have fewer nodes than \mathcal{M} . Now suppose that \mathcal{M}_a is such a network, and suppose that $\{\langle w \rangle^{\rho, \mathcal{M}}\}_{w \in V_{out}^{\mathcal{M}}}$ is a linearly dependent set of functions $\mathbb{R}^{D_0} \rightarrow \mathbb{R}$. In particular, let $(\lambda_w)_{w \in V_{out}^{\mathcal{M}}}$ be a nonzero set of scalars such that

$$\sum_{w \in V_{out}^{\mathcal{M}}} \lambda_w \langle w \rangle^{\rho, \mathcal{M}} = 0.$$

We then have

$$\left(\sum_{w \in V_{out}^{\mathcal{M}} \setminus V_{out}^{\mathcal{M}_a}} \lambda_w \langle w \rangle^{\rho, \mathcal{M}}(a) \right) \mathbf{1} + \sum_{w \in V_{out}^{\mathcal{M}_a}} \lambda_w \langle w \rangle^{\rho, \mathcal{M}_a} = \sum_{w \in V_{out}^{\mathcal{M}}} \lambda_w \langle w \rangle^{\rho, \mathcal{M}} = 0,$$

and thus $\{\langle w \rangle^{\rho, \mathcal{M}_a}\}_{w \in V_{out}^{\mathcal{M}_a}} \cup \{\mathbf{1}\}$ is a linearly dependent set of functions $\mathbb{R}^{D_0-1} \rightarrow \mathbb{R}$. Apropos, this derivation illustrates why it is often convenient to include the constant function $\mathbf{1}$ when dealing with

linear dependencies between the outputs of GFNNs. In the following definition we construct a network \mathcal{M}_a with the desired properties, and in Figure 5 we provide an illustration of this construction.

Definition 17. Let $\mathcal{M} = (V^{\mathcal{M}}, E^{\mathcal{M}}, V_{in}^{\mathcal{M}}, V_{out}^{\mathcal{M}}, \Omega^{\mathcal{M}}, \Theta^{\mathcal{M}})$ be a non-degenerate, clones-free GFNN with input nodes $V_{in}^{\mathcal{M}} = \{v_1^0, \dots, v_{D_0}^0\}$, $D_0 \geq 2$. Let $a \in \mathbb{R}$, and let ρ be a nonlinearity. The network obtained from \mathcal{M} by anchoring the input $v_{D_0}^0$ to a is the GFNN $\mathcal{M}_a = (V^{\mathcal{M}_a}, E^{\mathcal{M}_a}, V_{in}^{\mathcal{M}_a}, V_{out}^{\mathcal{M}_a}, \Omega^{\mathcal{M}_a}, \Theta^{\mathcal{M}_a})$ given by the following:

- $V^{\mathcal{M}_a} = \{v \in V^{\mathcal{M}} : \{v_1^0, \dots, v_{D_0-1}^0\} \cap V^{\mathcal{M}(v)} \neq \emptyset\}$, where $\mathcal{M}(v)$ denotes the ancestor network of v ,
- $E^{\mathcal{M}_a} = \{(v, \tilde{v}), v, \tilde{v} \in V^{\mathcal{M}_a}\}$,
- $V_{in}^{\mathcal{M}_a} = \{v_1^0, \dots, v_{D_0-1}^0\}$, $V_{out}^{\mathcal{M}_a} = V_{out}^{\mathcal{M}} \cap V^{\mathcal{M}_a}$, and
- $\Omega^{\mathcal{M}_a} = \{\omega_{\tilde{v}v} : (v, \tilde{v}) \in E^{\mathcal{M}_a}\}$.
- For a node $v \in V^{\mathcal{M}} \setminus V^{\mathcal{M}_a}$ we define recursively

$$a_v = \begin{cases} a, & v = v_{D_0}^0 \\ \rho \left(\sum_{u \in \text{par}_{\mathcal{M}}(v)} \omega_{vu} a_u + \theta_v \right), & v \neq v_{D_0}^0 \end{cases}. \quad (2)$$

(Note that all a_v are well-defined, as $\text{par}_{\mathcal{M}}(v) \subset V^{\mathcal{M}} \setminus V^{\mathcal{M}_a}$ whenever $v \in V^{\mathcal{M}} \setminus V^{\mathcal{M}_a}$.) Now, for $v \in V^{\mathcal{M}_a}$ let

$$\tilde{\theta}_v = \theta_v + \sum_{u \in \text{par}_{\mathcal{M}}(v) \setminus V^{\mathcal{M}_a}} \omega_{vu} a_u, \quad (3)$$

and set $\Theta^{\mathcal{M}_a} = \{\tilde{\theta}_v : v \in V^{\mathcal{M}_a}\}$.




Fig. 5: A concrete example of anchoring the input at u_4 of a network \mathcal{M} with input nodes $\{u_1, u_2, u_3, u_4\}$ and output nodes $\{w_1, w_2\}$ to a real number a . The parts of \mathcal{M} that are connected to u_4 , but not to any of the remaining inputs u_1, u_2, u_3 (dashed lines), are removed, while the rest of \mathcal{M} constitutes \mathcal{M}_a . To ensure that the outputs of \mathcal{M}_a (in this case only the node w_1) obey (IA-1), we need to “propagate” the anchored value through the removed parts of \mathcal{M} . This will manifest itself as a bias modification according to (2) and (3) at some of the nodes of \mathcal{M}_a (the only such node in this example is labeled by v).

The network \mathcal{M}_a satisfies (IA-1) and (IA-2) by construction, and if \mathcal{M} is layered, then so is \mathcal{M}_a . Moreover, \mathcal{M}_a is non-degenerate. To see this, let $v \in V^{\mathcal{M}_a}$ be arbitrary. Then, by non-degeneracy of

\mathcal{M} , there exists a $w \in V_{out}^{\mathcal{M}}$ such that $v \in V^{\mathcal{M}(w)}$. As w is connected directly with a node in $V^{\mathcal{M}_a}$, it follows that $w \in V^{\mathcal{M}_a}$, and so $v \in V_{out}^{\mathcal{M}_a}$.

Therefore $v \in V^{\mathcal{M}_a(w)}$, and, as v was arbitrary, we obtain $V^{\mathcal{M}_a} \subset \bigcup_{w \in V_{out}^{\mathcal{M}_a}} V^{\mathcal{M}_a(w)}$, establishing by Definition 15 that \mathcal{M}_a is non-degenerate. However, \mathcal{M}_a will not, generally, be clones-free. This is unfortunate, as our program for proving Theorem 2 envisages maintaining the no-clones property when constructing networks with linearly dependent outputs. However, not all is lost, as the following lemma says that, for nonlinearities holomorphic on a neighborhood of \mathbb{R} , either there exists some value of $a \in \mathbb{R}$ such that the network \mathcal{M}_a is, indeed, clones-free, or it is possible to modify a subnetwork of \mathcal{M} (different from the subnetwork giving rise to \mathcal{M}_a) to yield a clones-free subnetwork \mathcal{N} of \mathcal{M} with input $\{v_{D_0}^0\}$ and linear dependency among the maps realized by its output nodes. This will be sufficient for our purposes.

Lemma 2 (Input anchoring). *Let $\mathcal{M} = (V^{\mathcal{M}}, E^{\mathcal{M}}, V_{in}^{\mathcal{M}}, V_{out}^{\mathcal{M}}, \Omega^{\mathcal{M}}, \Theta^{\mathcal{M}})$ be a non-degenerate, clones-free GFNN with input nodes $V_{in}^{\mathcal{M}} = \{v_1^0, \dots, v_{D_0}^0\}$, $D_0 \geq 2$. Let $\rho : \mathcal{U} \rightarrow \mathbb{R}$ be holomorphic on an open domain $\mathcal{U} \subset \mathbb{C}$ containing \mathbb{R} , such that $\rho(\mathbb{R}) \subset \mathbb{R}$. Let \mathcal{M}_a denote the network obtained by anchoring the input $v_{D_0}^0$ to some $a \in \mathbb{R}$, according to Definition 17. Then one of the following two statements must be true:*

- (i) *There exists an $a \in \mathbb{R}$ such that \mathcal{M}_a is clones-free.*
- (ii) *There exist a non-degenerate clones-free GFNN $\mathcal{N} = (V^{\mathcal{N}}, E^{\mathcal{N}}, \{v_{D_0}^0\}, V_{out}^{\mathcal{N}}, \Omega^{\mathcal{N}}, \Theta^{\mathcal{N}})$ (obtained by modifying a subnetwork of \mathcal{M}), a real number λ_0 , and nonzero real numbers $(\lambda_w)_{w \in V_{out}^{\mathcal{N}}}$, such that the function $h_{out}^{\mathcal{N}} := \lambda_0 \mathbf{1} + \sum_{w \in V_{out}^{\mathcal{N}}} \lambda_w w^{\rho, \mathcal{N}}$ is identically zero on \mathbb{R} .*

Proof: For a pair of nodes $(c_1, c_2) \in V^{\mathcal{M}} \times V^{\mathcal{M}}$ define

$$E_{(c_1, c_2)} = \{a \in \mathbb{R} : c_1, c_2 \in V^{\mathcal{M}_a}, \text{ and } c_1, c_2 \text{ are clones in } \mathcal{M}_a\}.$$

Suppose that (i) is false, so that, for every $a \in \mathbb{R}$, we have $a \in E_{(c_1, c_2)}$ for some (c_1, c_2) . Then we can write \mathbb{R} as a finite union

$$\mathbb{R} = \bigcup_{(c_1, c_2) \in V^{\mathcal{M}} \times V^{\mathcal{M}}} E_{(c_1, c_2)}.$$

It follows that there exists a pair (c_1, c_2) such that at least one of the sets $E_{(c_1, c_2)}$ is not discrete, i.e., it has a limit point. Fix such a pair (c_1, c_2) . Note that we have $v_{D_0}^0 \in V^{\mathcal{M}(c_j)}$, for at least one of $j = 1$ or $j = 2$, as otherwise we would have $\text{par}_{\mathcal{M}_a}(c_j) = \text{par}_{\mathcal{M}}(c_j)$, for $j \in \{1, 2\}$ and all $a \in E_{(c_1, c_2)}$, and thus c_1, c_2 would be clones in \mathcal{M}_a if and only if they are clones in \mathcal{M} . But, by the no-clones property of \mathcal{M} , this would imply $E_{(c_1, c_2)} = \emptyset$, contradicting the fact that $E_{(c_1, c_2)}$ is not discrete. Thus, we may w.l.o.g. assume that $v_{D_0}^0 \in V^{\mathcal{M}(c_1)}$, which leaves us with the cases $v_{D_0}^0 \in V^{\mathcal{M}(c_2)}$ and $v_{D_0}^0 \notin V^{\mathcal{M}(c_2)}$ that will be treated separately when needed. Define the GFNN $\mathcal{N} = (V^{\mathcal{N}}, E^{\mathcal{N}}, \{v_{D_0}^0\}, V_{out}^{\mathcal{N}}, \Omega^{\mathcal{N}}, \Theta^{\mathcal{N}})$ according to the following:

– Let $S = \{v \in V^{\mathcal{M}(\{c_1, c_2\})} : V_{in}^{\mathcal{M}} \cap V^{\mathcal{M}(v)} = \{v_{D_0}^0\}\}$, and set

$$V^{\mathcal{N}} = \begin{cases} S \cup \{c_1, c_2\}, & \text{if } v_{D_0}^0 \in V^{\mathcal{M}(c_2)} \\ S \cup \{c_1\}, & \text{if } v_{D_0}^0 \notin V^{\mathcal{M}(c_2)}. \end{cases}$$

- $E^{\mathcal{N}} = \{(v, \tilde{v}), v, \tilde{v} \in V^{\mathcal{N}}\}$,
- $V_{out}^{\mathcal{N}} = \{c_1, c_2\} \cap V^{\mathcal{N}}$,

- $\Omega^{\mathcal{N}} = \{\omega_{\tilde{v}v} : (v, \tilde{v}) \in E^{\mathcal{N}}\}$,
- choose a number $r \in \mathbb{R} \setminus (\{\theta_v - \theta_{c_1} : v \in S\} \cup \{\theta_v - \theta_{c_2} : v \in S\})$, and set $\bar{\theta}_{c_1} = \theta_{c_1} + r$, $\bar{\theta}_{c_2} = \theta_{c_2} + r$, and $\bar{\theta}_v = \theta_v$, for $v \in S$. Define $\Theta^{\mathcal{N}} = \{\bar{\theta}_v : v \in V^{\mathcal{N}}\}$.

Informally, the so-constructed network \mathcal{N} consists of the parts of \mathcal{M} propagating the input at $v_{D_0}^0$ to c_1 and c_2 (and it might happen that this input does not reach c_2 , in which case this node is not included in $V^{\mathcal{N}}$), and the biases $\bar{\theta}_{c_1}$ and $\bar{\theta}_{c_2}$ are chosen so as to ensure that \mathcal{N} has no clone pair (v, \tilde{v}) with $v \in \{c_1, c_2\}$ and $\tilde{v} \in S$. Thus, in order to show that \mathcal{N} is clones-free, it suffices to establish that c_1 and c_2 are not clones in \mathcal{N} (note that c_1 and c_2 can be clones in \mathcal{N} only in the case $v_{D_0}^0 \in V^{\mathcal{M}(c_2)}$), as any clone pair (v, \tilde{v}) with $v, \tilde{v} \in S$ would also be a clone pair in \mathcal{M} . By way of contradiction, assume that c_1 and c_2 are clones in \mathcal{N} , i.e.,

$$\begin{aligned} \text{par}_{\mathcal{M}}(c_1) \cap V^{\mathcal{N}} &= \text{par}_{\mathcal{M}}(c_2) \cap V^{\mathcal{N}} \\ \theta_{c_1} + r &= \theta_{c_2} + r, \quad \text{and} \end{aligned} \tag{4}$$

$$(\omega_{c_1 u})_{u \in \text{par}_{\mathcal{M}}(c_1) \cap V^{\mathcal{N}}} = (\omega_{c_2 u})_{u \in \text{par}_{\mathcal{M}}(c_2) \cap V^{\mathcal{N}}}.$$

As the construction of \mathcal{N} does not depend on a , we can fix an arbitrary $a \in E_{(c_1, c_2)}$, and the condition that c_1 and c_2 are clones in \mathcal{M}_a then implies

$$\begin{aligned} \text{par}_{\mathcal{M}}(c_1) \setminus V^{\mathcal{N}} &= \text{par}_{\mathcal{M}}(c_2) \setminus V^{\mathcal{N}}, \\ \theta_{c_1} + \sum_{u \in \text{par}_{\mathcal{M}}(c_1) \cap V^{\mathcal{N}}} \omega_{c_1 u} a_u &= \theta_{c_2} + \sum_{u \in \text{par}_{\mathcal{M}}(c_2) \cap V^{\mathcal{N}}} \omega_{c_2 u} a_u, \quad \text{and} \\ (\omega_{c_1 u})_{u \in \text{par}_{\mathcal{M}}(c_1) \setminus V^{\mathcal{N}}} &= (\omega_{c_2 u})_{u \in \text{par}_{\mathcal{M}}(c_2) \setminus V^{\mathcal{N}}}, \end{aligned} \tag{5}$$

where the real numbers a_u are defined according to (2). This, together with (4), yields

$$\begin{aligned} \text{par}_{\mathcal{M}}(c_1) &= \text{par}_{\mathcal{M}}(c_2), \\ \theta_{c_1} &= \theta_{c_2}, \quad \text{and} \end{aligned} \tag{6}$$

$$(\omega_{c_1 u})_{u \in \text{par}_{\mathcal{M}}(c_1)} = (\omega_{c_2 u})_{u \in \text{par}_{\mathcal{M}}(c_2)},$$

which would say that c_1 and c_2 are clones in \mathcal{M} and hence stands in contradiction to the no-clones property of \mathcal{M} . This establishes the no-clones property of \mathcal{N} . The non-degeneracy of \mathcal{N} follows by its construction. Now, by adding r to both sides of (5) and applying ρ , we find

$$\langle c_1 \rangle^{\rho, \mathcal{N}}(a) = \begin{cases} \langle c_2 \rangle^{\rho, \mathcal{N}}(a), & \text{if } v_{D_0}^0 \in V^{\mathcal{M}(c_2)} \\ \rho(\theta_{c_2} + r) \mathbf{1}(a), & \text{if } v_{D_0}^0 \notin V^{\mathcal{M}(c_2)} \end{cases}, \tag{7}$$

for all $a \in E_{(c_1, c_2)}$ (note that $\text{par}_{\mathcal{M}}(c_2) \cap V^{\mathcal{N}} = \emptyset$ in the case $v_{D_0}^0 \notin V^{\mathcal{M}(c_2)}$), and so the sum on the right-hand side of (5) evaluates to 0 in this case). As ρ is holomorphic on an open neighborhood of \mathbb{R} and $\rho(\mathbb{R}) \subset \mathbb{R}$, we also have that $\langle c_1 \rangle^{\rho, \mathcal{N}}$, $\langle c_2 \rangle^{\rho, \mathcal{N}}$ are holomorphic on a neighborhood of \mathbb{R} . Further, since $E_{(c_1, c_2)}$ has a limit point, it follows by the identity theorem [18, Thm. 10.18] that (7) holds for all $a \in \mathbb{R}$. We have hence shown that Statement (ii) is valid with this \mathcal{N} , and

$$\begin{aligned} \lambda_0 = 0, \lambda_{c_1} = 1, \lambda_{c_2} = -1, & \quad \text{if } v_{D_0}^0 \in V^{\mathcal{M}(c_2)}, \text{ or} \\ \lambda_0 = -\rho(\theta_{c_2} + r), \lambda_{c_1} = 1, & \quad \text{if } v_{D_0}^0 \notin V^{\mathcal{M}(c_2)}. \end{aligned}$$

■

IV. AUXILIARY RESULTS FROM COMPLEX ANALYSIS AND KRONECKER'S THEOREM

We state the remaining auxiliary results needed in the proof of our main statements. Since these results are relatively simple consequences of standard results in complex analysis and of Kronecker's theorem, their proofs are relegated to the appendix.

Recall the definition of the natural domain $\mathcal{D}_{\langle u \rangle}^\sigma$ of the map realized by a GFNN node u with respect to a holomorphic nonlinearity as given in Definition 12.

In the proof of Theorem 2 it will be crucial that $\mathcal{D}_{\langle u \rangle}^\sigma$ be connected for all nodes u of a certain GFNN with a single input. The following lemma establishes this fact.

Lemma 3. *Let $\mathcal{N} = (V, E, \{v_{in}\}, V_{out}, \Omega, \Theta)$ be a GFNN, and let $\sigma : \mathcal{D}_\sigma \rightarrow \mathbb{C}$ be a meromorphic function on \mathbb{C} with its set of poles given by $P \subset \mathbb{C} \setminus \mathbb{R}$. Furthermore, suppose that $\sigma(\mathbb{R}) \subset \mathbb{R}$. Then, for every $u \in V$, we have $\mathcal{D}_{\langle u \rangle}^\sigma = \mathbb{C} \setminus E_u$, where $E_u \subset \mathbb{C}$ is a closed countable subset of $\mathbb{C} \setminus \mathbb{R}$. In particular, we have that $\mathcal{D}_{\langle u \rangle}^\sigma$ is an open connected set with $\mathcal{D}_{\langle u \rangle}^\sigma \supset \mathbb{R}$.*

In the following we write $D_k^\circ(\mathbf{a}, \delta) := \{(z_1, \dots, z_k) \in \mathbb{C}^k : |z_j - a_j| < \delta, \forall j\}$ for the open polydisc of radius $\delta > 0$, centered at $\mathbf{a} = (a_1, \dots, a_k) \in \mathbb{C}^k$. Further, for a set $S \subset \mathbb{C}^k$, we write $\text{cl}(S)$ for the closure of S in \mathbb{C}^k .

Lemma 4. *Let $F : \mathcal{U} \rightarrow \mathbb{C}$ be holomorphic on a connected open domain $\mathcal{U} \subset \mathbb{C}^k$ containing \mathbb{R}^k . Let $\mathbf{a} = (a_1, \dots, a_k) \in \mathbb{R}^k$ and $\delta > 0$ be given, and let*

$$T = \{(a_1 + iz_1, \dots, a_k + iz_k) : z_j \in (-\delta, \delta), j = 1, \dots, k\}.$$

Suppose that $D_k^\circ(\mathbf{a}, \delta) \subset \mathcal{U}$, and $F(z) = 0$, for all $z \in T$. Then $F = 0$ identically on \mathcal{U} .

Lemma 5. *Let $t^* \in \mathbb{C}$, $\mathbf{a} = (a_1, \dots, a_k) \in \mathbb{R}^k$, and $\delta > 0$, and let $F : \mathcal{U} \rightarrow \mathbb{C}$ be holomorphic on a connected open domain $\mathcal{U} \subset \mathbb{C}^{1+k}$ containing $\{t^*\} \times \mathbb{R}^k$. Define the set*

$$T = \{(t^*, a_1 + iz_1, \dots, a_k + iz_k) : z_j \in (-\delta, \delta), j = 1, \dots, k\},$$

and suppose that $D_{1+k}^\circ(\mathbf{a}, \delta) \subset \mathcal{U}$. If there exists a set $\tilde{T} \subset \mathbb{C}^{1+k}$ such that $\tilde{T} \subset (\mathbb{C} \setminus \{t^\}) \times \mathbb{C}^k$, $\text{cl}(\tilde{T}) \supset T$, and $F|_{\tilde{T}} \equiv 0$, then $F|_{\mathcal{U}} \equiv 0$.*

We will now elaborate on the tools needed in the proof of Theorem 2. The material touches upon the theory of Lie groups and representation theory, and will be presented in a self-contained fashion, only assuming familiarity with finitely-generated abelian groups and basic point-set topology. We write $T^d = \mathbb{R}^d / \mathbb{Z}^d$ for the d -dimensional torus considered as a compact abelian topological group. For a finite set of real numbers $\{\alpha_j\}_{j=1}^d$ we let $\langle \alpha_1, \dots, \alpha_d \rangle_{\mathbb{Q}}$ denote the span of $\{\alpha_j\}_{j=1}^d$ in the vector space \mathbb{R} over the scalar field \mathbb{Q} , and we write $\dim \langle \alpha_1, \dots, \alpha_d \rangle_{\mathbb{Q}}$ for its dimension. We will need the following lemma, which is an easy consequence of Kronecker's theorem [16]. For the sake of completeness, we provide an elementary proof from first principles.

Lemma 6 ([16] Kronecker). *Let $d \in \mathbb{N}$ and let $\{\alpha_j\}_{j=1}^d$ be an arbitrary set of nonzero real numbers with $k = \dim \langle \alpha_1, \dots, \alpha_d \rangle_{\mathbb{Q}}$. Define the following subset of T^d :*

$$M = \text{cl}\{(\alpha_1 t, \alpha_2 t, \dots, \alpha_d t) + \mathbb{Z}^d : t \in \mathbb{R}\},$$




Fig. 6: The line $\ell : t \mapsto (\alpha_1 t, \alpha_2 t) + \mathbb{Z}^2$, $t \in \mathbb{R}$, depicted in the fundamental cell $[-\frac{1}{2}, \frac{1}{2}] \times [-\frac{1}{2}, \frac{1}{2}]$ of the torus $T^2 = \{(x_1, x_2) + \mathbb{Z}^2 : (x_1, x_2) \in \mathbb{R}^2\}$, with $(\alpha_1, \alpha_2) = (1, 1.4)$ (left), and $(\alpha_1, \alpha_2) = (1, \sqrt{2})$ (right).

where cl denotes the closure in T^d . Then M is isomorphic to a k -dimensional torus as a Lie group, i.e., there exists a $\Psi : M \rightarrow \mathbb{R}^k / \mathbb{Z}^k$ that is both a homeomorphism (between M and $\mathbb{R}^k / \mathbb{Z}^k$ as topological spaces) and a homomorphism (between M and $\mathbb{R}^k / \mathbb{Z}^k$ as abelian groups).

When $d = 2$, Lemma 6 simply says that the line $\ell : t \mapsto (\alpha_1 t, \alpha_2 t) + \mathbb{Z}^2$, $t \in \mathbb{R}$, either exhibits discrete periodic behavior and is thus homeomorphic to a 1-dimensional torus, which is the case if $k = 1$, i.e., α_1 / α_2 is rational, or otherwise, if $k = 2$, i.e., when α_1 / α_2 is irrational, ℓ is dense in the whole square, and so its closure is a 2-dimensional torus, namely $\mathbb{R}^2 / \mathbb{Z}^2$ itself. This is illustrated in Figure 6. When $d \geq 3$, the situation can be more complicated, as illustrated in Figure 7. Specifically, the torus M obtained as the closure of the line $\ell : t \mapsto (\alpha_1 t, \dots, \alpha_d t) + \mathbb{Z}^d$, $t \in \mathbb{R}$, may not occupy the entirety of $\mathbb{R}^d / \mathbb{Z}^d$. In this case, Lemma 6 provides the precise dimension of M , namely $k = \dim \langle \alpha_1, \dots, \alpha_d \rangle_{\mathbb{Q}}$. For the purpose of proving Theorem 2, it will suffice to consider the behavior of ℓ in a neighborhood of the point $\mathbf{0} + \mathbb{Z}^d \in T^d$. Concretely, if $Q \in \mathbb{Q}^{d \times k}$ is the matrix representing $\alpha_1, \dots, \alpha_d$ in the basis $\{\alpha_1, \dots, \alpha_k\}$, the following lemma states that, in a neighborhood of $\mathbf{0}$, ℓ visits points arbitrarily close to the k -dimensional subspace of \mathbb{R}^d spanned by the columns of Q .

Lemma 7. *Suppose that $\{\alpha_j\}_{j=1}^d$ are nonzero real numbers, and let $k = \dim \langle \alpha_1, \dots, \alpha_d \rangle_{\mathbb{Q}}$. Furthermore, assume that $\{\alpha_j\}_{j=1}^k$ is a basis for $\langle \alpha_1, \dots, \alpha_d \rangle_{\mathbb{Q}}$ over \mathbb{Q} , and let $Q = (Q_{pj}) \in \mathbb{Q}^{d \times k}$ be the matrix such that $(\alpha_1, \dots, \alpha_d) = Q \cdot (\alpha_1, \dots, \alpha_k)$. Then there exists an open set $C \subset \mathbb{R}^k$ with $\mathbf{0} \in C$, such that, for every $\mathbf{s} = (s_1, \dots, s_k) \in C$, there are sequences $(t^{n,\mathbf{s}})_{n \in \mathbb{N}} \subset \mathbb{R}$ and $(\mathbf{r}^{n,\mathbf{s}})_{n \in \mathbb{N}} = (r_1^{n,\mathbf{s}}, \dots, r_k^{n,\mathbf{s}})_{n \in \mathbb{N}} \subset C$ with the following properties:*

- (i) $(\alpha_1 t^{n,\mathbf{s}}, \alpha_2 t^{n,\mathbf{s}}, \dots, \alpha_d t^{n,\mathbf{s}}) + \mathbb{Z}^d = Q \cdot (\alpha_1 r_1^{n,\mathbf{s}}, \dots, \alpha_k r_k^{n,\mathbf{s}}) + \mathbb{Z}^d$, for all $n \in \mathbb{N}$,
- (ii) $|t^{n,\mathbf{s}}| \rightarrow \infty$ as $n \rightarrow \infty$,
- (iii) $\mathbf{r}^{n,\mathbf{s}} \rightarrow \mathbf{s}$ in \mathbb{R}^k , as $n \rightarrow \infty$.

V. IMAGINARY PERIOD AND THE SELF-AVOIDING PROPERTY

We say that a holomorphic function $f : \mathcal{D} \rightarrow \mathbb{C}$ is i -periodic if $f(z + i) = f(z)$, for all $z \in \mathcal{D}$. An example of such a function is the scaled hyperbolic tangent function $\tanh(\pi \cdot)$. More generally,




Fig. 7: The line $t \mapsto (\alpha_1 t, \alpha_2 t, \alpha_3 t)$, $t \in \mathbb{R}$, depicted in the fundamental cell $[-\frac{1}{2}, \frac{1}{2}]^3$ of the torus $T^3 = \{(x_1, x_2, x_3) + \mathbb{Z}^3 : (x_1, x_2, x_3) \in \mathbb{R}^3\}$, with $(\alpha_1, \alpha_2, \alpha_3) = (\frac{2}{5}, -\frac{4}{5}, \frac{3}{2})$ (left), and $(\alpha_1, \alpha_2, \alpha_3) = (1, \sqrt{2}, \frac{1}{2} + \sqrt{2})$ (right). Note that in a neighborhood of $\mathbf{0}$ (marked in red), ℓ is dense in a k -dimensional subspace of \mathbb{R}^3 , with $k = \dim \langle \frac{2}{5}, -\frac{4}{5}, \frac{3}{2} \rangle_{\mathbb{Q}} = 1$ (left), and $k = \dim \langle 1, \sqrt{2}, \frac{1}{2} + \sqrt{2} \rangle_{\mathbb{Q}} = 2$ (right).

for an arbitrary discrete set $S \subset \mathbb{R}$, and arbitrary $C \in \mathbb{R}$ and real sequence $\{c_s\}_{s \in S} \in \ell^1(S)$, the function $\sigma = C + \sum_{s \in S} c_s \tanh(\pi(\cdot - s))$ is also i -periodic, and in particular, the set of its poles P has the structure $P = \bigcup_{n \in \mathbb{Z}} (S + (n + \frac{1}{2})i)$. We now introduce a property defined for discrete subsets of \mathbb{R} , which will, when applied to the set S , be the final technical ingredient in the proof of our main results.

Definition 18 (Self-avoiding set). Let $S \subset \mathbb{R}$ be a discrete set. We say that S is self-avoiding if, for every finite collection of distinct pairs $\{(\omega_j, \theta_j)\}_{j=1}^m \subset (2\mathbb{Z} + 1) \times \mathbb{R}$, there exist a $j^* \in \{1, \dots, m\}$ and a t^* such that

$$t^* \in \frac{S - \theta_{j^*}}{\omega_{j^*}} \setminus \bigcup_{j \neq j^*} \frac{S - \theta_j}{\omega_j}.$$

Remark. In other words, a set S is self-avoiding if the union of a finite number of distinct copies of S obtained by translating and scaling by an odd integer contains a real number which is an element of exactly one of the copies.

Proposition 2. Let $S = \{s_k : k \in \mathbb{Z}\}$, $s_k - s_{k-1} > 0$, $\forall k \in \mathbb{Z}$, be an infinite discrete set such that $\{s_k - s_{k-1} : k \in \mathbb{Z}\}$ is rationally independent. Then S is self-avoiding.

Proof: We use the shorthand notation $S_{\omega, \theta} = \frac{S - \theta}{\omega}$. Suppose by way of contradiction that $A \subset (2\mathbb{Z} + 1) \times \mathbb{R}$, $\#A \geq 2$, is a set of pairs such that, for every $(\omega, \theta) \in A$ and every $t \in S_{\omega, \theta}$, there exists a pair $(\omega', \theta') \in A \setminus \{(\omega, \theta)\}$ such that $t \in S_{\omega', \theta'}$. Fix a pair $(\omega_1, \theta_1) \in A$. We then have, by assumption,

$$S_{\omega_1, \theta_1} = \bigcup_{(\omega', \theta') \in A \setminus \{(\omega_1, \theta_1)\}} S_{\omega_1, \theta_1} \cap S_{\omega', \theta'}.$$

Since S is infinite, there exists a $(\omega_2, \theta_2) \in A \setminus \{(\omega_1, \theta_1)\}$ such that $\#(S_{\omega_1, \theta_1} \cap S_{\omega_2, \theta_2}) \geq 3$. Pick an arbitrary subset $\{t_1 < t_2 < t_3\} \subset S_{\omega_1, \theta_1} \cap S_{\omega_2, \theta_2}$ and note that there exist $k_1^1, k_2^1, k_3^1 \in \mathbb{Z}$ and

$k_1^2, k_2^2, k_3^2 \in \mathbb{Z}$ such that

$$t_j = \frac{s_{k_j^1} - \theta_1}{\omega_1} = \frac{s_{k_j^2} - \theta_2}{\omega_2}, \quad \text{for } j = 1, 2, 3. \quad (8)$$

Moreover, for $r = 1, 2$, we have $k_1^r < k_2^r < k_3^r$ if $\omega_r > 0$ and $k_1^r > k_2^r > k_3^r$ if $\omega_r < 0$. Define the index sets

$$K_j^r = \begin{cases} \{k_j^r + 1, k_j^r + 2, \dots, k_{j+1}^r\}, & \text{if } \omega_r > 0 \\ \{k_{j+1}^r + 1, k_{j+1}^r + 2, \dots, k_j^r\}, & \text{if } \omega_r < 0 \end{cases}, \quad \text{for } j = 1, 2, r = 1, 2.$$

For brevity write $a_k = s_k - s_{k-1}$, $\forall k \in \mathbb{Z}$. We then have

$$(t_2 - t_1, t_3 - t_2) = \left(\frac{1}{|\omega_1|} \sum_{k \in K_1^1} a_k, \frac{1}{|\omega_1|} \sum_{k \in K_2^1} a_k \right) = \left(\frac{1}{|\omega_2|} \sum_{k \in K_1^2} a_k, \frac{1}{|\omega_2|} \sum_{k \in K_2^2} a_k \right). \quad (9)$$

Now, since $\{a_k : k \in \mathbb{Z}\}$ is rationally independent and $|\omega_1|, |\omega_2| \in \mathbb{Z}$, (9) implies $|\omega_1| = |\omega_2|$ and $K_j^1 = K_j^2$, for $j = 1, 2$. In particular, $K_j^1 = K_j^2$, for $j = 1, 2$, implies $\text{sgn}(\omega_1) = \text{sgn}(\omega_2)$, so we have $\omega_1 = \omega_2$. Then, from the definition of K_j^r , it follows that $k_j^1 = k_j^2$, for $j = 1, 2, 3$. We thus obtain from (8) that $\theta_1 = \theta_2$, contradicting $(\omega_1, \theta_1) \neq (\omega_2, \theta_2)$. Therefore, our initial assumption was false, so we deduce that S must be self-avoiding. \blacksquare

The following proposition formalizes the notion that nonlinearities σ of the form considered at the beginning of the chapter are dense in the set of sigmoidal nonlinearities, even after imposing the additional constraint that S be self-avoiding.

Proposition 3. *Let ρ be a piecewise C^1 nonlinearity with $\rho' \in BV(\mathbb{R}) \cap L^1(\mathbb{R})$. Then, for every $\epsilon > 0$, there exist a discrete self-avoiding set $S \subset \mathbb{R}$, a sequence $\{c_s\}_{s \in S} \in \ell^1(S)$ with $c_s \neq 0$, for all $s \in S$, and real numbers $\alpha > 0$ and C , such that the function σ given by*

$$\sigma = C + \sum_{s \in S} c_s \tanh(\alpha(\cdot - s))$$

satisfies $\|\sigma - \rho\|_{L^\infty(\mathbb{R})} < \epsilon$.

Proof: First note that

$$\rho(-\infty) = \lim_{x \rightarrow -\infty} \rho(x) = \rho(0) - \int_{-\infty}^0 \rho'(y) dy$$

is a well-defined real number, as $\rho' \in L^1(\mathbb{R})$. Let H denote the Heaviside step function. We now have, for all $x \in \mathbb{R}$,

$$\rho(x) = \rho(-\infty) + \int_{\mathbb{R}} \rho'(y) H(x - y) dy.$$

Denote $h_\alpha = \frac{1}{2}(1 + \tanh(\alpha \cdot))$ and consider the function ρ_α defined by

$$\rho_\alpha(x) = \rho(-\infty) + \int_{\mathbb{R}} \rho'(y) h_\alpha(x - y) dy, \quad x \in \mathbb{R}. \quad (10)$$

We then have

$$\begin{aligned} \sup_{x \in \mathbb{R}} |\rho(x) - \rho_\alpha(x)| &= \sup_{x \in \mathbb{R}} \left| \int_{\mathbb{R}} \rho'(y) [H(x - y) - h_\alpha(x - y)] dy \right| \\ &= \sup_{x \in \mathbb{R}} \left| \int_{\mathbb{R}} \rho'(x - y) [H(y) - h_\alpha(y)] dy \right| \\ &\leq \|\rho'\|_{L^\infty(\mathbb{R})} \|H - h_\alpha\|_{L^1(\mathbb{R})}. \end{aligned}$$

Now note that $\|\rho'\|_{L^\infty(\mathbb{R})} < \infty$ as $\rho' \in BV(\mathbb{R})$, and $\|H - h_\alpha\|_{L^1(\mathbb{R})} \rightarrow 0$ as $\alpha \rightarrow \infty$ by dominated convergence, so there exists $\alpha > 0$ such that $\|\rho - \rho_\alpha\|_{L^\infty(\mathbb{R})} < \frac{\epsilon}{3}$. Let $b : \mathbb{Z} \rightarrow \mathbb{N}$ be a bijection, and $\beta \in (0, 1)$ a parameter to be specified. Define the infinite discrete set $S_\beta = \{s_k^\beta := \beta(k + \pi^{-b(k)}) : k \in \mathbb{Z}\} \subset \mathbb{R}$. Then, since π is transcendental, Proposition 2 implies that S_β is self-avoiding. Now, since ρ' is integrable on \mathbb{R} and piecewise continuous, and h_α is bounded and continuous, we have that $\rho' \cdot h_\alpha(x - \cdot)$ is integrable on \mathbb{R} and piecewise continuous. Hence, as $\text{mesh}(S_\beta) := \sup_{k \in \mathbb{Z}} |s_k^\beta - s_{k-1}^\beta| \rightarrow 0$ for $\beta \rightarrow 0$, we have the following convergence of Riemann sums

$$\sum_{k \in \mathbb{Z}} (s_k^\beta - s_{k-1}^\beta) \rho'(s_k^\beta) h_\alpha(x - s_k^\beta) \rightarrow \int_{\mathbb{R}} \rho'(y) h_\alpha(x - y) dy \quad \text{as } \beta \rightarrow 0, \quad \text{for all } x \in \mathbb{R}.$$

Therefore $\rho(-\infty) + \sum_{k \in \mathbb{Z}} (s_k^\beta - s_{k-1}^\beta) \rho'(s_k^\beta) h_\alpha(\cdot - s_k^\beta) \rightarrow \rho_\alpha$ pointwise. To upgrade this to convergence in $\|\cdot\|_{L^\infty(\mathbb{R})}$, we proceed as follows. By the mean value theorem, for any $x \in \mathbb{R}$ and $\beta > 0$, there exist $y_k^{\beta, x} \in [s_{k-1}^\beta, s_k^\beta]$ such that

$$\int_{s_{k-1}^\beta}^{s_k^\beta} \rho'(y) h_\alpha(x - y) dy = (s_k^\beta - s_{k-1}^\beta) \rho'(y_k^{\beta, x}) h_\alpha(x - y_k^{\beta, x}).$$

We can therefore write

$$\begin{aligned} & \sup_{x \in \mathbb{R}} \left| \sum_{k \in \mathbb{Z}} (s_k^\beta - s_{k-1}^\beta) \rho'(s_k^\beta) h_\alpha(x - s_k^\beta) - \int_{\mathbb{R}} \rho'(y) h_\alpha(x - y) dy \right| \tag{11} \\ &= \sup_{x \in \mathbb{R}} \left| \sum_{k \in \mathbb{Z}} (s_k^\beta - s_{k-1}^\beta) \left[\rho'(s_k^\beta) h_\alpha(x - s_k^\beta) - \rho'(y_k^{\beta, x}) h_\alpha(x - y_k^{\beta, x}) \right] \right| \\ &\leq \text{mesh}(S_\beta) \cdot \sup_{x \in \mathbb{R}} \sum_{k \in \mathbb{Z}} \left| \rho'(s_k^\beta) h_\alpha(x - s_k^\beta) - \rho'(y_k^{\beta, x}) h_\alpha(x - y_k^{\beta, x}) \right| \\ &\leq \text{mesh}(S_\beta) \cdot \sup_{x \in \mathbb{R}} \|\rho' \cdot h_\alpha(x - \cdot)\|_{BV(\mathbb{R})} \\ &\leq \text{mesh}(S_\beta) (\|\rho'\|_{L^\infty(\mathbb{R})} \|h_\alpha\|_{BV(\mathbb{R})} + \|h_\alpha\|_{L^\infty(\mathbb{R})} \|\rho'\|_{BV(\mathbb{R})}). \end{aligned}$$

Since $\rho' \in BV(\mathbb{R})$ by assumption, and $h_\alpha \in BV(\mathbb{R})$ by definition, the quantities in the parentheses are all finite. As they are moreover independent of β , and $\text{mesh}(S_\beta) \rightarrow 0$ for $\beta \rightarrow 0$, we can pick a $\beta > 0$ such that

$$\left\| \sum_{k \in \mathbb{Z}} (s_k^\beta - s_{k-1}^\beta) \rho'(s_k^\beta) h_\alpha(\cdot - s_k^\beta) - \rho_\alpha + \rho(-\infty) \right\|_{L^\infty(\mathbb{R})} < \frac{\epsilon}{3}, \tag{12}$$

where we used (10) to replace $\int_{\mathbb{R}} \rho'(y) h_\alpha(x - y) dy$ in (11) with $\rho_\alpha - \rho(-\infty)$. Finally, let $\{d_s\}_{s \in S_\beta}$ be an arbitrary sequence of real numbers such that $\text{mesh}(S_\beta) \sum_{k \in \mathbb{Z}} |d_{s_k^\beta}| < \frac{\epsilon}{3}$ and, for each $s \in S_\beta$, $d_s = 0$ if and only if $\rho'(s) \neq 0$. We then have

$$\left\| \sum_{k \in \mathbb{Z}} (s_k^\beta - s_{k-1}^\beta) d_{s_k^\beta} h_\alpha(\cdot - s_k^\beta) \right\|_{L^\infty(\mathbb{R})} \leq \text{mesh}(S_\beta) \sum_{k \in \mathbb{Z}} |d_{s_k^\beta}| \cdot \|h_\alpha\|_{L^\infty(\mathbb{R})} < \frac{\epsilon}{3}. \tag{13}$$

Now, combining the estimates (12), (13), and $\|\rho - \rho_\alpha\|_{L^\infty(\mathbb{R})} < \frac{\epsilon}{3}$ yields

$$\left\| \rho(-\infty) + \sum_{k \in \mathbb{Z}} (s_k^\beta - s_{k-1}^\beta) (\rho'(s_k^\beta) + d_{s_k^\beta}) h_\alpha(\cdot - s_k^\beta) - \rho \right\|_{L^\infty(\mathbb{R})} < \epsilon,$$

so the claim of the proposition holds with $S = S_\beta$, $c_{s_k^\beta} = \frac{1}{2}(s_k^\beta - s_{k-1}^\beta)(\rho'(s_k^\beta) + d_{s_k^\beta})$, and $C = \rho(-\infty) + \sum_{k \in \mathbb{Z}} c_{s_k^\beta}$. \blacksquare

VI. THE MAIN THEOREMS

Theorem 3. Let \mathcal{N}_1 and \mathcal{N}_2 be non-degenerate clones-free LFNNs with the same input and output sets V_{in} and V_{out} . Let

$$\sigma = C + \sum_{s \in S} c_s \tanh(\pi(\cdot - s)),$$

where $C \in \mathbb{R}$, S is a discrete self-avoiding set, and $\{c_s\}_{s \in S} \in \ell^1(S)$ are all nonzero and real. Suppose that $\langle \mathcal{N}_1 \rangle^\sigma(\mathbf{t}) = \langle \mathcal{N}_2 \rangle^\sigma(\mathbf{t})$, for all $\mathbf{t} \in \mathbb{R}^{V_{in}}$. Then \mathcal{N}_1 and \mathcal{N}_2 are faithfully isomorphic.

Theorem 4. Let \mathcal{N}_j , $j \in \{1, 2, \dots, n\}$, be non-degenerate clones-free LFNNs with the same input set V_{in} and the same single output node $\{v_{out}\}$. Furthermore, suppose that no two networks \mathcal{N}_{j_1} , \mathcal{N}_{j_2} , $j_1 \neq j_2$, are extensionally isomorphic. Consider the nonlinearity

$$\sigma = C + \sum_{s \in S} c_s \tanh(\pi(\cdot - s)),$$

with $C \in \mathbb{R}$, S a discrete self-avoiding set, and $\{c_s\}_{s \in S} \in \ell^1(S)$, where each c_s is nonzero and real. Then $\{\langle \mathcal{N}_j \rangle^\sigma\}_{j=1}^n \cup \{\mathbf{1}\}$ is a linearly independent set of functions from $\mathbb{R}^{V_{in}}$ to \mathbb{R} .

Before embarking on the proofs of Theorems 3 and 4, we show how Theorems 1 and 2 follow from these two results together with Proposition 3.

Proof of Theorem 1: Let ρ be as in the statement of Theorem 1, and let $\epsilon > 0$ be arbitrary. Proposition 3 guarantees the existence of a discrete self-avoiding set $S \subset \mathbb{R}$, a sequence $\{c_s\}_{s \in S} \in \ell^1(S)$ with $c_s \neq 0$, for all $s \in S$, and real numbers $\alpha > 0$ and C , such that the function σ defined by

$$\sigma = C + \sum_{s \in S} c_s \tanh(\alpha(\cdot - s))$$

satisfies $\|\sigma - \rho\|_{L^\infty(\mathbb{R})} < \epsilon$. Now suppose that $\mathcal{N} = (V, E, V_{in}, V_{out}, \Omega, \Theta)$ and $\tilde{\mathcal{N}} = (\tilde{V}, \tilde{E}, V_{in}, V_{out}, \tilde{\Omega}, \tilde{\Theta})$ are clones-free non-degenerate LFNNs with the same input set V_{in} and such that $\langle \mathcal{N} \rangle^\sigma(x) = \langle \tilde{\mathcal{N}} \rangle^\sigma(x)$, for all $x \in \mathbb{R}^{V_{in}}$. Consider the scaled objects $\sigma_\alpha := \sigma\left(\frac{\pi}{\alpha}\cdot\right)$, $S_\alpha = \frac{\alpha}{\pi}S$, $\mathcal{N}^\alpha = (V, E, V_{in}, V_{out}, \frac{\alpha}{\pi}\Omega, \frac{\alpha}{\pi}\Theta)$, and $\tilde{\mathcal{N}}^\alpha = (\tilde{V}, \tilde{E}, V_{in}, V_{out}, \frac{\alpha}{\pi}\tilde{\Omega}, \frac{\alpha}{\pi}\tilde{\Theta})$, where $\frac{\alpha}{\pi}\Omega = \{\frac{\alpha}{\pi}\omega : \omega \in \Omega\}$, and $\frac{\alpha}{\pi}\Theta, \frac{\alpha}{\pi}\tilde{\Omega}, \frac{\alpha}{\pi}\tilde{\Theta}$ are defined analogously. Then $\langle \mathcal{N}^\alpha \rangle^{\sigma_\alpha}(x) = \langle \mathcal{N} \rangle^\sigma(x) = \langle \tilde{\mathcal{N}} \rangle^\sigma(x) = \langle \tilde{\mathcal{N}}^\alpha \rangle^{\sigma_\alpha}(x)$, for all $x \in \mathbb{R}^{V_{in}}$. Moreover,

$$\sigma_\alpha = C + \sum_{s \in S_\alpha} c_s \tanh(\pi(\cdot - s)),$$

and S_α is a discrete self-avoiding set (as the self-avoiding property is preserved under scaling by a nonzero real number), so by Theorem 3 we obtain $\mathcal{N}^\alpha \stackrel{f}{\sim} \tilde{\mathcal{N}}^\alpha$, which implies $\mathcal{N} \simeq \tilde{\mathcal{N}}$. ■

Proof of Theorem 2: Let ρ be as in the statement of Theorem 2, and let $\epsilon > 0$ be arbitrary. Proposition 3 guarantees the existence of a discrete self-avoiding set $S \subset \mathbb{R}$, a sequence $\{c_s\}_{s \in S} \in \ell^1(S)$ with $c_s \neq 0$, for all $s \in S$, and real numbers $\alpha > 0$ and C , such that the function σ defined by

$$\sigma = C + \sum_{s \in S} c_s \tanh(\alpha(\cdot - s))$$

satisfies $\|\sigma - \rho\|_{L^\infty(\mathbb{R})} < \epsilon$. Now suppose that $\mathcal{N}_j = (V^j, E^j, V_{in}, \{v_{out}\}, \Omega^j, \Theta^j)$, $j \in \{1, \dots, n\}$, are non-degenerate clones-free LFNNs such that no two \mathcal{N}_{j_1} , \mathcal{N}_{j_2} , $j_1 \neq j_2$, are faithfully isomorphic. As $\{v_{out}\}$ is a singleton, it follows that no two \mathcal{N}_{j_1} , \mathcal{N}_{j_2} , $j_1 \neq j_2$, are extensionally isomorphic either. Now, define the scaled objects $\sigma_\alpha := \sigma\left(\frac{\pi}{\alpha}\cdot\right)$, $S_\alpha = \frac{\alpha}{\pi}S$, and $\mathcal{N}_j^\alpha = (V^j, E^j, V_{in}, \{v_{out}\}, \frac{\alpha}{\pi}\Omega^j, \frac{\alpha}{\pi}\Theta^j)$,

for $j \in \{1, \dots, n\}$, where $\frac{\alpha}{\pi}\Omega^j = \{\frac{\alpha}{\pi}\omega : \omega \in \Omega_j\}$ and $\frac{\alpha}{\pi}\Theta^j = \{\frac{\alpha}{\pi}\theta : \theta \in \Theta^j\}$. Then the \mathcal{N}_j^α are non-degenerate and clones-free, and no two $\mathcal{N}_{j_1}^\alpha, \mathcal{N}_{j_2}^\alpha, j_1 \neq j_2$, are extensionally isomorphic. Moreover,

$$\sigma_\alpha = C + \sum_{s \in S_\alpha} c_s \tanh(\pi(\cdot - s)),$$

and S_α is a discrete self-avoiding set, so by Theorem 4 we obtain that $\{\langle \mathcal{N}_j^\alpha \rangle^{\sigma_\alpha}\}_{j=1}^n \cup \{\mathbf{1}\}$ is linearly independent. Now, suppose by way of contradiction that there is linear dependency $\lambda_0 + \sum_{j=1}^n \lambda_j \langle \mathcal{N}_j \rangle^\sigma = 0$ among $\{\langle \mathcal{N}_j \rangle^\sigma\}_{j=1}^n \cup \{\mathbf{1}\}$. But then

$$\lambda_0 + \sum_{j=1}^n \lambda_j \langle \mathcal{N}_j^\alpha \rangle^{\sigma_\alpha} = \lambda_0 + \sum_{j=1}^n \lambda_j \langle \mathcal{N}_j \rangle^\sigma = 0,$$

which contradicts the linear independence of $\{\langle \mathcal{N}_j^\alpha \rangle^{\sigma_\alpha}\}_{j=1}^n \cup \{\mathbf{1}\}$. We deduce that $\{\langle \mathcal{N}_j \rangle^\sigma\}_{j=1}^n \cup \{\mathbf{1}\}$ must be linearly independent, as desired. \blacksquare

Proof of Theorem 4: We argue by contradiction, so suppose that the statement is false. Specifically, let $\mathcal{N}_j, j \in \{1, 2, \dots, n\}$, be LFNNs and σ a nonlinearity as in the statement of the theorem, and suppose that $\{\langle \mathcal{N}_j \rangle^\sigma\}_{j=1}^n \cup \{\mathbf{1}\}$ is linearly dependent. Then, by Lemma 1, there exists a non-degenerate clones-free LFNN $\mathcal{M} = (V^\mathcal{M}, E^\mathcal{M}, V_{in}^\mathcal{M}, V_{out}^\mathcal{M}, \Omega^\mathcal{M}, \Theta^\mathcal{M})$ with a single input node $V_{in}^\mathcal{M} = \{v_{in}\}$, such that $\{\langle w \rangle^\sigma : w \in V_{out}^\mathcal{M}\} \cup \{\mathbf{1}\}$ is a linearly dependent set of functions from \mathbb{R} to \mathbb{R} . Let \mathcal{M} denote the set of all non-degenerate clones-free LFNNs $\widetilde{\mathcal{M}} = (V^{\widetilde{\mathcal{M}}}, E^{\widetilde{\mathcal{M}}}, \{v_{in}\}, V_{out}^{\widetilde{\mathcal{M}}}, \Omega^{\widetilde{\mathcal{M}}}, \Theta^{\widetilde{\mathcal{M}}})$ such that $\{\langle w \rangle^\sigma : w \in V_{out}^{\widetilde{\mathcal{M}}}\} \cup \{\mathbf{1}\}$ is linearly dependent. We then have $\mathcal{M} \neq \emptyset$, simply as $\mathcal{M} \in \mathcal{M}$. Denote by \mathcal{M}_{min} the set of all networks in \mathcal{M} of minimum depth, and fix a network $\mathcal{M}' \in \mathcal{M}_{min}$ with the minimal number of nodes among all the networks in \mathcal{M}_{min} . The proof proceeds by constructing a network $\mathcal{N} \in \mathcal{M}_{min}$ with a strictly smaller number of nodes than \mathcal{M}' , thereby deriving a contradiction and concluding the proof. First note that linear dependence of $\{\langle w \rangle^\sigma : w \in V_{out}^{\mathcal{M}'}\} \cup \{\mathbf{1}\}$ is equivalent to the existence of a nonzero set of real numbers $\{\lambda_w\}_{w \in V_{out}^{\mathcal{M}'}}$ and a real number $c \in \mathbb{R}$ such that $h_{out} : \mathbb{R} \rightarrow \mathbb{R}$, given by

$$h_{out} := \sum_{w \in V_{out}^{\mathcal{M}'}} \lambda_w \langle w \rangle^\sigma,$$

is constant-valued, i.e., $h_{out}(t) = c$, for all $t \in \mathbb{R}$. Note that $\lambda_w \neq 0$, for all $w \in V_{out}^{\mathcal{M}'}$, for otherwise the ancestor subnetwork $\mathcal{M}'(\{w \in V_{out}^{\mathcal{M}'}, \lambda_w \neq 0\})$ would be an element of \mathcal{M}_{min} with strictly fewer nodes than \mathcal{M}' , contradicting the minimality of \mathcal{M}' .

Next, note that σ is a real meromorphic function whose set of poles is

$$P = \bigcup_{n \in \mathbb{Z}} \left(S + \left(n + \frac{1}{2} \right) i \right), \quad (14)$$

and in particular, \mathcal{M}' and σ satisfy the assumptions of Lemma 3, and so the sets $\mathbb{C} \setminus \mathcal{D}_{\langle w \rangle^\sigma}$ are closed and countable, where $\mathcal{D}_{\langle w \rangle^\sigma}$ denotes the natural domain of $\langle w \rangle^\sigma$, for $w \in V_{out}^{\mathcal{M}'}$. Therefore, as a linear combination of holomorphic functions, h_{out} is a holomorphic function on $\mathcal{D}_{h_{out}} := \bigcap_{w \in V_{out}^{\mathcal{M}'}} \mathcal{D}_{\langle w \rangle^\sigma}$. As the sets $\mathbb{C} \setminus \mathcal{D}_{\langle w \rangle^\sigma}$ are closed and countable, $\mathbb{C} \setminus \mathcal{D}_{h_{out}}$ is also closed and countable, and therefore $\mathcal{D}_{h_{out}}$ is a connected open set. It follows by the identity theorem [18, Thm. 10.18] that h_{out} continues in a unique fashion to a holomorphic function on $\mathcal{D}_{h_{out}}$ with $h_{out}(t) = c$, for all $t \in \mathcal{D}_{h_{out}}$.

Set $V_\ell = \{v \in V^{\mathcal{M}'} : \text{lv}(v) = \ell\}$, for $\ell \geq 1$. Let $k = \dim \langle \{\omega_{uv_{in}} : u \in V_1\} \rangle_{\mathbb{Q}}$ and enumerate the nodes $V_1 = \{v_1^1, \dots, v_{D_1}^1\}$ so that $\{\omega_{v_1^1 v_{in}}, \dots, \omega_{v_k^1 v_{in}}\}$ is a basis for $\langle \omega_{v_1^1 v_{in}}, \dots, \omega_{v_{D_1}^1 v_{in}} \rangle_{\mathbb{Q}}$. In the remainder of the proof, we distinguish between the cases $k \geq 2$ and $k = 1$.

The case $k \geq 2$. Fix a real number

$$A \in [0, 1] \setminus \bigcup_{p=1}^{D_1} \frac{S - \theta_{v_p^1}}{\omega_{v_p^1 v_{in}}}, \quad (15)$$

chosen so that none of $\langle v_p^1 \rangle^\sigma(z) = \sigma(\omega_{v_p^1 v_{in}} z + \theta_{v_p^1})$, $p \in \{1, \dots, D_1\}$, has singularities along $A + i\mathbb{R}$. Such a number always exists, as $\bigcup_{p=1}^{D_1} (S - \theta_{v_p^1})/\omega_{v_p^1 v_{in}}$ is a discrete set. Now, write $(\omega_{v_p^1 v_{in}})_{p=1}^{D_1} = Q \cdot (\omega_{v_p^1 v_{in}})_{p=1}^k$, where $Q = (q_{pj}) \in \mathbb{Q}^{D_1 \times k}$ is a rational matrix whose first k rows form a $k \times k$ identity matrix. Let $C \subset \mathbb{R}^k$ be a set satisfying the conclusion of Lemma 7 applied with $\alpha_p = \omega_{v_p^1 v_{in}}$, $p \in \{1, \dots, D_1\}$. Given an arbitrary $\mathbf{s} = (s_1, s_2, \dots, s_k) \in C$, Lemma 7 yields sequences $(t^{n, \mathbf{s}})_{n \in \mathbb{N}} \subset \mathbb{R}$ and $(\mathbf{r}^{n, \mathbf{s}})_{n \in \mathbb{N}} \subset C$ such that

$$(\omega_{v_1^1 v_{in}} t^{n, \mathbf{s}}, \dots, \omega_{v_{D_1}^1 v_{in}} t^{n, \mathbf{s}}) + \mathbb{Z}^{D_1} = Q \cdot (\omega_{v_1^1 v_{in}} r_1^{n, \mathbf{s}}, \dots, \omega_{v_k^1 v_{in}} r_k^{n, \mathbf{s}}) + \mathbb{Z}^{D_1}, \quad (16)$$

$$|t^{n, \mathbf{s}}| \rightarrow \infty \quad \text{as } n \rightarrow \infty, \quad (17)$$

$$\mathbf{r}^{n, \mathbf{s}} \rightarrow \mathbf{s} \quad \text{in } \mathbb{R}^k, \text{ as } n \rightarrow \infty. \quad (18)$$

We now perform a calculation that will enable us to interpret the single input variable of \mathcal{M}' as a rational linear combination of k input variables of another LFNN \mathcal{M}'' , to be specified below. The argument will then proceed by anchoring at all but one of the inputs of \mathcal{M}'' . It is this last step that uses $k \geq 2$ as a key assumption, as anchoring requires at least two input nodes to be meaningful. We thus have

$$\begin{aligned} & \sigma\left(\omega_{v_p^1 v_{in}}(A + i t^{n, \mathbf{s}}) + \theta_{v_p^1}\right) \\ &= \sigma\left(\omega_{v_p^1 v_{in}} A + i(\omega_{v_p^1 v_{in}} t^{n, \mathbf{s}} + \mathbb{Z}) + \theta_{v_p^1}\right) \end{aligned} \quad (19)$$

$$= \sigma\left(\omega_{v_p^1 v_{in}} A + i \cdot \left(\sum_{j=1}^k q_{pj} \omega_{v_j^1 v_{in}} r_j^{n, \mathbf{s}} + \mathbb{Z}\right) + \theta_{v_p^1}\right) \quad (20)$$

$$= \sigma\left(\sum_{j=1}^k q_{pj} \omega_{v_j^1 v_{in}} (A + i r_j^{n, \mathbf{s}}) + \theta_{v_p^1}\right), \quad (21)$$

for $p \in \{1, \dots, D_1\}$, where in (19) we used the i -periodicity of σ , in (20) we used (16), and in (21) we used $\omega_{v_p^1 v_{in}} = \sum_{j=1}^k q_{pj} \omega_{v_j^1 v_{in}}$ and the i -periodicity of σ again. Owing to (15), none of $\langle v_p^1 \rangle^\sigma$, $p \in \{1, \dots, D_1\}$, has singularities along $A + i\mathbb{R}$, and thus all the quantities in (19) – (21) are well-defined. The calculation just presented suggests constructing a new LFNN by “splitting” the input node v_{in} of \mathcal{M}' into k new input nodes. Formally, we define an LFNN $\mathcal{M}'' = (V^{\mathcal{M}''}, E^{\mathcal{M}''}, V_{in}^{\mathcal{M}''}, V_{out}^{\mathcal{M}''}, \Omega^{\mathcal{M}''}, \Theta^{\mathcal{M}''})$ as follows:

- $V_{in}^{\mathcal{M}''} = \{u_1, \dots, u_k\}$ is a set of k newly-created input nodes (disjoint from $V^{\mathcal{M}'}$),
- $V^{\mathcal{M}''} := V_{in}^{\mathcal{M}''} \cup \bigcup_{\ell \geq 1} V_\ell$,
- $E^{\mathcal{M}''} := \{(v, \tilde{v}) \in E^{\mathcal{M}'} : \text{lv}(v) \geq 1\} \cup \{(u_j, v_p^1) : 1 \leq p \leq D_1, 1 \leq j \leq k, q_{pj} \neq 0\}$,
- $V_{out}^{\mathcal{M}''} := V_{out}^{\mathcal{M}'}$,
- define $\omega_{v_p^1 u_j} := q_{pj} \omega_{v_j^1 v_{in}}$, for $p \in \{1, \dots, D_1\}$, $j \in \{1, \dots, k\}$, and let

$$\Omega^{\mathcal{M}''} := \{\omega_{\tilde{v}v} \in \Omega^{\mathcal{M}'} : \text{lv}(v) \geq 1\} \cup \{(\omega_{v_p^1 u_j} : 1 \leq p \leq D_1, 1 \leq j \leq k, q_{pj} \neq 0\},$$

- $\Theta^{\mathcal{M}''} := \Theta^{\mathcal{M}'}$.

The procedure for constructing \mathcal{M}'' for a given \mathcal{M}' is illustrated in Figure 8.

Owing to (19) – (21) and the construction of \mathcal{M}'' , we have the following “input splitting” relationship

$$\langle v_p^1 \rangle^{\sigma, \mathcal{M}'}(A + i t^{n, \mathbf{s}}) = \langle v_p^1 \rangle^{\sigma, \mathcal{M}''}(A + i r_1^{n, \mathbf{s}}, \dots, A + i r_k^{n, \mathbf{s}}), \quad (22)$$

for $p \in \{1, \dots, D_1\}$. We now show that \mathcal{M}'' is non-degenerate and clones-free. To this end, first note




Fig. 8: *Input splitting, case $k \geq 2$. Left:* A neural network \mathcal{M}' , assumed to be a minimal element of \mathcal{M}_{min} , with $D_1 = 5$. *Right:* The corresponding neural network \mathcal{M}'' , assuming $k = 3$. Note that, as the first k rows of Q form a $k \times k$ identity matrix, we have $(u_j, v_p^1) \in E^{\mathcal{M}''} \iff p = j$, for all $j, p \in \{1, \dots, k\}$. The function F in (24) and (25) corresponds to the map realized by the shared part (in red) of \mathcal{M}' and \mathcal{M}'' .

that, for every $j \in \{1, \dots, k\}$, there exists a $w \in V_{out}^{\mathcal{M}'}$ such that $v_j^1 \in V^{\mathcal{M}'(w)}$, by non-degeneracy of \mathcal{M}' , and as $u_j \in \text{par}(v_j^1)$, we have $u_j \in V^{\mathcal{M}''(w)}$. This establishes non-degeneracy. Next, we observe that a clone pair in \mathcal{M}'' would have to consist of nodes in $\{v_1^1, v_2^1, \dots, v_{D_1}^1\}$, as a clone pair in \mathcal{M}'' consisting only of nodes in $\bigcup_{\ell \geq 2} V_\ell$ would also be a clone pair in \mathcal{M}' . Thus, by way of contradiction, suppose that $(v_{p_1}^1, v_{p_2}^1)$, $1 \leq p_1 < p_2 \leq D_1$, is a clone pair in \mathcal{M}'' . Then $\theta_{p_1}^1 = \theta_{p_2}^1$ and $\omega_{v_{p_1}^1 v_{in}} = \sum_{j=1}^k q_{p_1 j} \omega_{v_j^1 v_{in}} = \sum_{j=1}^k q_{p_2 j} \omega_{v_j^1 v_{in}} = \omega_{v_{p_2}^1 v_{in}}$, so $(v_{p_1}^1, v_{p_2}^1)$ is a clone pair in \mathcal{M}' , which stands in contradiction to the no-clones property of \mathcal{M}' , and hence establishes that \mathcal{M}'' is clones-free. We now revisit the constant-valued function $h_{out}(t) = \sum_{w \in V_{out}^{\mathcal{M}'}} \lambda_w \langle w \rangle^{\sigma, \mathcal{M}'}(t) = c$, for all $t \in \mathcal{D}_{h_{out}}$. Examining the structure of \mathcal{M}' , we see that, for each $w \in V_{out}^{\mathcal{M}'}$, we can write

$$\langle w \rangle^{\sigma, \mathcal{M}'}(z) = F_w \left(\left(\langle v_p^1 \rangle^{\sigma, \mathcal{M}'}(z) \right)_{p=1}^{D_1} \right), \quad \text{for all } z \in \mathcal{D}_{\langle w \rangle^{\sigma, \mathcal{M}'}}$$

where F_w corresponds to the map realized by the LFNN with nodes

$$V^{\mathcal{M}'} \setminus \{v_{in}\}, \quad (23)$$

inputs $\{v_1^1, \dots, v_{D_1}^1\}$, output $\{w\}$, and edges, weights, and biases inherited from \mathcal{M}' . As F_w is the map realized by a node of a GFNN according to Definition 12, it is holomorphic on its natural domain $\mathcal{D}_{F_w} \subset \mathbb{C}^{D_1}$ containing \mathbb{R}^{D_1} . We can therefore write

$$h_{out}(z) = F \left(\left(\langle v_p^1 \rangle^{\sigma, \mathcal{M}'}(z) \right)_{p=1}^{D_1} \right), \quad \text{for all } z \in \mathcal{D}_{h_{out}}, \quad (24)$$

where $F : \mathcal{D}_F \rightarrow \mathbb{C}$, $F = \sum_{w \in V_{out}^{\mathcal{M}'}} \lambda_w F_w$, is holomorphic on $\mathcal{D}_F := \bigcap_{w \in V_{out}^{\mathcal{M}'}} \mathcal{D}_{F_w} \supset \mathbb{R}^{D_1}$. Now, by definition of natural domain, for each $w \in V_{out}^{\mathcal{M}'}$, we have

$$\mathcal{D}_{\langle w \rangle^{\sigma, \mathcal{M}''}} = \left\{ (z_1, \dots, z_k) \in \bigcap_{p=1}^{D_1} \mathcal{D}_{\langle v_p^1 \rangle^{\sigma, \mathcal{M}''}} : \left(\langle v_p^1 \rangle^{\sigma, \mathcal{M}''}(z_1, \dots, z_k) \right)_{p=1}^{D_1} \in \mathcal{D}_{F_w} \right\},$$

where the variables z_1, \dots, z_k correspond to the input nodes u_1, \dots, u_k , respectively. Therefore, for (z_1, \dots, z_k) in the open domain $\mathcal{D}_{\tilde{h}_{out}} := \bigcap_{w \in V_{out}^{\mathcal{M}''}} \mathcal{D}_{\langle w \rangle^{\sigma, \mathcal{M}''}}$, we can define the function $\tilde{h}_{out} : \mathcal{D}_{\tilde{h}_{out}} \rightarrow \mathbb{C}$ according to

$$\tilde{h}_{out}(z_1, \dots, z_k) = F \left(\left(\langle v_p^1 \rangle^{\sigma, \mathcal{M}''}(z_1, \dots, z_k) \right)_{p=1}^{D_1} \right). \quad (25)$$

Moreover, as \mathcal{M}' and \mathcal{M}'' share the nodes in (23), as well as the associated edges, weights, and biases, we have

$$\langle w \rangle^{\sigma, \mathcal{M}''}(z_1, \dots, z_k) = F_w \left(\left(\langle v_p^1 \rangle^{\sigma, \mathcal{M}''}(z_1, \dots, z_k) \right)_{p=1}^{D_1} \right),$$

for all $w \in V_{out}^{\mathcal{M}''}$, and thus

$$\tilde{h}_{out} = \sum_{w \in V_{out}^{\mathcal{M}''}} \lambda_w \langle w \rangle^{\sigma, \mathcal{M}''}.$$

We are now in a position to show that, like h_{out} , the function \tilde{h}_{out} is constant valued. As this will be effected by an analytic continuation argument through Lemma 4, we first need to ensure that the relevant quantities lie in $\mathcal{D}_{\tilde{h}_{out}}$. To this end, as $\langle v_p^1 \rangle^{\sigma, \mathcal{M}''}(z_1, \dots, z_k) \in \mathbb{R}$, for all $(z_1, \dots, z_k) \in \mathbb{R}^k$, $p \in \{1, \dots, D_1\}$, and \mathcal{D}_F is an open set containing \mathbb{R}^{D_1} , we can choose a small enough $\delta > 0$ so that $\mathcal{D}_{\tilde{h}_{out}} \supset D_k^\circ((A, \dots, A), \delta)$. Now, fix an arbitrary $\mathbf{s} = (s_1, \dots, s_k)$ in the smaller open set $C \cap D_k^\circ(\mathbf{0}, \delta)$. We then have

$$(A + is_1, \dots, A + is_k) \in D_k^\circ((A, \dots, A), \delta) \subset \mathcal{D}_{\tilde{h}_{out}},$$

and since

$$(A + ir_1^{n, \mathbf{s}}, \dots, A + ir_k^{n, \mathbf{s}}) \rightarrow (A + is_1, \dots, A + is_k),$$

as $n \rightarrow \infty$, we obtain

$$(A + ir_1^{n, \mathbf{s}}, \dots, A + ir_k^{n, \mathbf{s}}) \in \mathcal{D}_{\tilde{h}_{out}},$$

for large enough $n \in \mathbb{N}$. We may assume w.l.o.g. that this is true for all $n \in \mathbb{N}$ by discarding finitely many elements of the sequence $(\mathbf{r}^{n, \mathbf{s}})_{n \in \mathbb{N}}$. Now, we use (22), (24), and (25) to get

$$\begin{aligned} \tilde{h}_{out}(A + ir_1^{n, \mathbf{s}}, \dots, A + ir_k^{n, \mathbf{s}}) &= F \left(\left(\langle v_p^1 \rangle^{\sigma, \mathcal{M}''}(A + ir_1^{n, \mathbf{s}}, \dots, A + ir_k^{n, \mathbf{s}}) \right)_{p=1}^{D_1} \right) \\ &= F \left(\left(\langle v_p^1 \rangle^{\sigma, \mathcal{M}'}(A + it^{n, \mathbf{s}}) \right)_{p=1}^{D_1} \right) \\ &= h_{out}(A + it^{n, \mathbf{s}}) = c, \quad \text{for all } n \in \mathbb{N}. \end{aligned}$$

Define the set

$$T = \{(A + ir_1^{n, \mathbf{s}}, \dots, A + ir_k^{n, \mathbf{s}}) : \mathbf{s} \in C \cap D_k^\circ(\mathbf{0}, \delta), n \in \mathbb{N}\}$$

and note that $\text{cl}(T) \supset ((A, \dots, A) + (iC) \cap D_k^\circ(\mathbf{0}, \delta))$, so it follows by Lemma 4 that $\tilde{h}_{out} - c \equiv 0$ everywhere in a neighborhood of \mathbb{R}^k , and thus, in particular, $\tilde{h}_{out}|_{\mathbb{R}^k} \equiv c$. We now repeatedly apply Lemma 2 to \mathcal{M}'' , anchoring successively each of the inputs u_1, \dots, u_{k-1} . Observe that we will never find ourselves in the circumstance (ii) of Lemma 2, as this would mean that we have obtained a network $\mathcal{N} \in \mathcal{M}_{min}$ with a strictly smaller number of nodes than \mathcal{M}' . Moreover, as the first k rows of Q form an identity matrix, we have

$$(v_p^1, u_j) \in E^{\mathcal{M}''} \iff q_{pj} \neq 0 \iff p = j,$$

for all $p, j \in \{1, \dots, k\}$. Therefore, for each $j \in \{1, \dots, k\}$, the node v_j^1 will be removed when anchoring the input u_j . A concrete example of this input anchoring procedure in the case $k \geq 2$ is shown schematically in Figure 9. Thus, having anchored the nodes u_1, u_2, \dots, u_{k-1} to appro-




Fig. 9: *Input anchoring. Left:* The neural network \mathcal{M}'' as in Figure 8. *Right:* Anchoring the inputs of \mathcal{M}'' at the nodes u_1, u_2, \dots, u_{k-1} results in the removal of the nodes $v_1^1, v_2^1, \dots, v_{k-1}^1$ (and possibly some other nodes deeper in the network). As $k \geq 2$, the resulting network \mathcal{N} has fewer nodes than the original network \mathcal{M}' .

ropriate real numbers a_1, \dots, a_{k-1} , we will be left with a non-degenerate clones-free LFNN $\mathcal{N} = (V^{\mathcal{N}}, E^{\mathcal{N}}, \{u_k\}, V_{out}^{\mathcal{N}}, \Omega^{\mathcal{N}}, \Theta^{\mathcal{N}})$ such that the function $h_{out}^{\mathcal{N}} := \sum_{w \in V_{out}^{\mathcal{N}}} \lambda_w \langle w \rangle^{\sigma, \mathcal{N}}$ satisfies

$$h_{out}^{\mathcal{N}}(t) = \tilde{h}_{out}(a_1, \dots, a_{k-1}, t) - \sum_{w \in V_{out}^{\mathcal{M}''} \setminus V_{out}^{\mathcal{N}}} \lambda_w \langle w \rangle^{\sigma, \mathcal{M}''}(a_1, \dots, a_{k-1}, t), \quad \forall t \in \mathbb{R}. \quad (26)$$

We have shown that the first term on the right-hand side of (26) evaluates identically to c . Moreover, as input anchoring yields networks satisfying (IA-2), the values $\langle w \rangle^{\sigma, \mathcal{M}''}$, for $w \in V_{out}^{\mathcal{M}''} \setminus V_{out}^{\mathcal{N}}$, are constant with respect to the input at u_k . Therefore the value of the sum on the right-hand side of (26) is independent of t , that is, $h_{out}^{\mathcal{N}} \equiv c_{\mathcal{N}}$, for some $c_{\mathcal{N}} \in \mathbb{R}$. As $\lambda_w \neq 0$, for $w \in V_{out}^{\mathcal{M}''}$, it follows that $\{\langle w \rangle^{\sigma, \mathcal{N}} : w \in V_{out}^{\mathcal{N}}\} \cup \{1\}$ is linearly dependent. We have thus shown that the network \mathcal{N} is in \mathcal{M}_{min} . As \mathcal{N} has strictly fewer nodes than \mathcal{M}' , we have established the desired contradiction and proved the theorem for $k \geq 2$.

The case $k = 1$. We have $\dim \langle \omega_{v_1^1 v_{in}}, \dots, \omega_{v_{D_1}^1 v_{in}} \rangle_{\mathbb{Q}} = 1$, so we can write $\omega_{v_j^1 v_{in}} = N_j a$, where $a \in \mathbb{R}$ and $N_j \in \mathbb{Z}$, for $j = 1, \dots, D_1$. Moreover, by replacing a with $2^\ell a$ and all N_j with $N_j/2^\ell$ for an appropriate integer ℓ , we may assume w.l.o.g. that at least one of the N_j is odd. We make the following crucial observation. For all $j = 1, \dots, D_1$ and $t \in \mathbb{R}$, we have

$$\begin{aligned} \langle v_j^1 \rangle^{\sigma, \mathcal{M}'} \left(t + \frac{i}{2a} \right) &= \sigma \left(N_j a t + \theta_{v_j^1} + \frac{N_j i}{2} \right) \\ &= \begin{cases} C + \sum_{s \in S} c_s \tanh(\pi(N_j a t + \theta_{v_j^1} - s)), & N_j \text{ even} \\ C + \sum_{s \in S} c_s \coth(\pi(N_j a t + \theta_{v_j^1} - s)), & N_j \text{ odd} \end{cases}. \end{aligned} \quad (27)$$

We see that, along the line $\mathbb{R} + \frac{i}{2a}$, the functions $\langle v_j^1 \rangle^{\sigma, \mathcal{M}'}$ are real-valued, for all $j = 1, \dots, D_1$, and, provided that N_j is odd, they have poles at the points $\frac{1}{a} \left[\frac{S - \theta_{v_j^1}}{N_j} + \frac{i}{2} \right]$. As S is self-avoiding, and at least one of the N_j is odd, there exist a $j^* \in \{1, \dots, D_1\}$ and a $t^* \in \mathbb{R} + \frac{i}{2a}$ such that $\langle v_{j^*}^1 \rangle^{\sigma, \mathcal{M}'}$ has a pole at t^* , and all the other $\langle v_j^1 \rangle^{\sigma, \mathcal{M}'}$, $j \in \{1, \dots, D_1\} \setminus \{j^*\}$, are analytic and real-valued at t^* . Let $\epsilon > 0$ be such that $\langle v_j^1 \rangle^{\sigma, \mathcal{M}'}$, $j \in \{1, \dots, D_1\} \setminus \{j^*\}$, are analytic on an open set containing the

closed disk $D(t^*, \epsilon)$, and such that $\langle v_{j^*}^1 \rangle^{\sigma, \mathcal{M}'}$ is analytic on the punctured disk $D(t^*, \epsilon) \setminus \{t^*\}$. Before embarking on the construction of \mathcal{N} in the case $k = 1$, we verify the following auxiliary statement:

Claim 1: We have $L(\mathcal{M}') \geq 2$ and $\{\tilde{v} \in V_2 : (v_{j^}^1, \tilde{v}) \in E^{\mathcal{M}'}\} \neq \emptyset$.*

Proof of Claim 1. We first show that $L(\mathcal{M}') \geq 2$. To this end, suppose by way of contradiction that $L(\mathcal{M}') = 1$. Then $V_{out}^{\mathcal{M}'} = V_1$ by non-degeneracy, so the function $h_{out} = \sum_{w \in V_{out}^{\mathcal{M}'}} \lambda_w \langle w \rangle^{\sigma, \mathcal{M}'}$ can be written as

$$h_{out}(t) = \lambda_{v_{j^*}^1} \langle v_{j^*}^1 \rangle^{\sigma, \mathcal{M}'}(t) + g(t), \quad (28)$$

where g is analytic in an open neighborhood of t^* . But $\langle v_{j^*}^1 \rangle^{\sigma, \mathcal{M}'}$ has a pole at t^* , and so h_{out} has a pole at t^* , which stands in contradiction to $h_{out} \equiv c$, and thus establishes $L(\mathcal{M}') \geq 2$.

Next, by way of contradiction assume that $\{\tilde{v} \in V_2 : (v_{j^*}^1, \tilde{v}) \in E^{\mathcal{M}'}\} = \emptyset$. Then, by non-degeneracy of \mathcal{M}' , we have $v_{j^*}^1 \in V_{out}^{\mathcal{M}'}$, and $\langle w \rangle^{\sigma, \mathcal{M}'}$, for $w \in V_{out}^{\mathcal{M}'} \setminus \{v_{j^*}^1\}$, are real holomorphic functions of $(\langle v_j^1 \rangle^{\sigma, \mathcal{M}'})_{j \in \{1, \dots, D_1\} \setminus \{j^*\}}$. Now, as $\langle v_j^1 \rangle^{\sigma, \mathcal{M}'}$, $j \in \{1, \dots, D_1\} \setminus \{j^*\}$, are analytic and real-valued at t^* , the function h_{out} can again be written in the form (28) with g analytic in an open neighborhood of t^* . This again contradicts $h_{out} \equiv c$, and thus $\{\tilde{v} \in V_2 : (v_{j^*}^1, \tilde{v}) \in E^{\mathcal{M}'}\} \neq \emptyset$, establishing the claim. We can therefore enumerate the nodes $V_2 = \{v_1^2, \dots, v_d^2, v_{d+1}^2, \dots, v_{D_2}^2\}$ so that

- $v_{j^*}^1 \in \bigcap_{p \leq d} \text{par}(\{v_p^2\}) \setminus \bigcup_{p > d} \text{par}(\{v_p^2\})$, and
- $\{\omega_{v_1^2 v_{j^*}^1}, \dots, \omega_{v_d^2 v_{j^*}^1}\}$ is a basis for $\langle \omega_{v_1^2 v_{j^*}^1}, \dots, \omega_{v_d^2 v_{j^*}^1} \rangle_{\mathbb{Q}}$.

In particular, we have $\bar{k} = \dim \langle \omega_{v_1^2 v_{j^*}^1}, \dots, \omega_{v_d^2 v_{j^*}^1} \rangle_{\mathbb{Q}}$. We will apply a similar input splitting procedure as in the case $k \geq 2$, but this time with the nodes $v_{j^*}^1$ and v_1^2, \dots, v_d^2 taking on the roles of v_{in} and $v_1^1, \dots, v_{D_1}^1$. Specifically, we will use the pole of $\langle v_{j^*}^1 \rangle^{\sigma, \mathcal{M}'}$ at t^* to obtain sequences $(t^{n, \mathbf{s}})_{n \in \mathbb{N}}$ and $(r^{n, \mathbf{s}})_{n \in \mathbb{N}}$ according to Lemma 7, that is to say, we will “split the non-input node” $v_{j^*}^1$ of \mathcal{M}' into input nodes of the new network \mathcal{M}'' to be constructed. We remark that the outputs of v_1^2, \dots, v_d^2 depend on $\langle v_{j^*}^1 \rangle^{\sigma, \mathcal{M}'}$, which, in turn, is a function of the input variables. This “extra level of separation” will cause the construction of \mathcal{M}'' to be more involved in the case $k = 1$ than it was in the case $k \geq 2$.

In order to motivate the construction of \mathcal{M}'' in the case $k = 1$, we will carry out a calculation analogous to (19)–(21). We begin by determining a $B \in \mathbb{R}$ such that none of the functions

$$\langle v_p^2 \rangle^{\sigma, \mathcal{M}'}(z) = \sigma \left(\omega_{v_p^2 v_{j^*}^1} \langle v_{j^*}^1 \rangle^{\sigma, \mathcal{M}'}(z) + f_p(z) + \theta_{v_p^2} \right), \quad z \in \mathcal{D}_{v_p^2}, \quad (29)$$

for $p \in \{1, \dots, d\}$, have singularities in the set $\mathcal{L}_B := \{z \in D(t^*, \epsilon) : \langle v_{j^*}^1 \rangle^{\sigma, \mathcal{M}'}(z) \in B + i\mathbb{R}\}$, where the functions $f_p : \mathcal{D}_{f_p} \rightarrow \mathbb{C}$, for $p \in \{1, \dots, d\}$, are defined according to

$$f_p(z) = \sum_{j \in \{1, \dots, D_1\} \setminus \{j^*\}} \omega_{v_p^2 v_j^1} \langle v_j^1 \rangle^{\sigma, \mathcal{M}'}(z), \quad z \in \mathcal{D}_{f_p}. \quad (30)$$

When $D_1 = 1$, the functions f_p are all identically zero. For given $p \in \{1, \dots, d\}$, $z \in \mathcal{L}_B$ is a singularity of $\langle v_p^2 \rangle^{\sigma, \mathcal{M}'}$ if and only if z is an element of $D(t^*, \epsilon)$ such that

$$\langle v_{j^*}^1 \rangle^{\sigma, \mathcal{M}'}(z) \in (B + i\mathbb{R}) \cap \frac{P - f_p(z) - \theta_{v_p^2}}{\omega_{v_p^2 v_{j^*}^1}},$$

where P is the set of poles of σ , expressed in terms of S by (14). But

$$\frac{P - f_p(z) - \theta_{v_p^2}}{\omega_{v_p^2 v_{j^*}^1}} \subset \frac{S - \text{Re}[f_p(D(t^*, \epsilon))] - \theta_{v_p^2}}{\omega_{v_p^2 v_{j^*}^1}} + i\mathbb{R},$$

for all $z \in D(t^*, \epsilon)$, so it suffices to ensure that

$$B \notin \bigcup_{p=1}^d \frac{S - \operatorname{Re}[f_p(D(t^*, \epsilon))] - \theta_{v_p^2}^2}{\omega_{v_p^2 v_{j^*}^1}}. \quad (31)$$

Next, let

$$\eta(\epsilon) = \sup_{1 \leq p \leq d} \sup_{z \in D(t^*, \epsilon)} |f_p(z) - f_p(t^*)|$$

and note that, as f_p , $p = 1, \dots, d$, are continuous in a neighborhood of t^* , we have $\eta(\epsilon) \rightarrow 0$ as $\epsilon \rightarrow 0$. Let Leb denote the Lebesgue measure on \mathbb{R} . We then have

$$\begin{aligned} & \operatorname{Leb} \left\{ [0, 1] \cap \bigcup_{p=1}^d \frac{S - \operatorname{Re}[f_p(D(t^*, \epsilon))] - \theta_{v_p^2}^2}{\omega_{v_p^2 v_{j^*}^1}} \right\} \\ & \leq \sum_{p=1}^d \frac{2\eta(\epsilon)}{|\omega_{v_p^2 v_{j^*}^1}|} \cdot \# \left\{ [0, 1] \cap \frac{S - \theta_{v_p^2}^2}{\omega_{v_p^2 v_{j^*}^1}} \right\} < 1 \end{aligned}$$

for small enough values of ϵ . Therefore, by choosing a sufficiently small ϵ , we can ensure that there exists a $B \in [0, 1]$ such that (31) holds, as desired. Now, write $(\omega_{v_p^2 v_{j^*}^1})_{p=1}^d = \bar{Q} \cdot (\omega_{v_p^2 v_{j^*}^1})_{p=1}^{\bar{k}}$, where $\bar{Q} = (\bar{q}_{pj})_{p,j} \in \mathbb{Q}^{d \times \bar{k}}$ is a rational matrix whose first \bar{k} rows form a $\bar{k} \times \bar{k}$ identity matrix. Let $C \subset \mathbb{R}^{\bar{k}}$ be a set satisfying the conclusion of Lemma 7 applied with $\alpha_p = \omega_{v_p^2 v_{j^*}^1}$, $p = 1, \dots, \bar{k}$.

Given an arbitrary $\mathbf{s} = (s_1, s_2, \dots, s_{\bar{k}}) \in C$, Lemma 7 yields sequences $(t^{n,\mathbf{s}})_{n \in \mathbb{N}} \subset \mathbb{R}$, $(\mathbf{r}^{n,\mathbf{s}})_{n \in \mathbb{N}} \subset C$ such that

$$(\omega_{v_1^2 v_{j^*}^1} t^{n,\mathbf{s}}, \dots, \omega_{v_d^2 v_{j^*}^1} t^{n,\mathbf{s}}) + \mathbb{Z}^d = \bar{Q} \cdot (\omega_{v_1^2 v_{j^*}^1} r_1^{n,\mathbf{s}}, \dots, \omega_{v_{\bar{k}}^2 v_{j^*}^1} r_{\bar{k}}^{n,\mathbf{s}}) + \mathbb{Z}^d, \quad \forall n \in \mathbb{N}, \quad (32)$$

$$|t^{n,\mathbf{s}}| \rightarrow \infty \quad \text{as } n \rightarrow \infty, \quad (33)$$

$$\mathbf{r}^{n,\mathbf{s}} \rightarrow \mathbf{s} \quad \text{as } n \rightarrow \infty. \quad (34)$$

As $\langle v_{j^*}^1 \rangle^{\sigma, \mathcal{M}'}$ is analytic on the punctured disk $D(t^*, \epsilon) \setminus \{t^*\}$ and its singularity at t^* is a pole, it follows that the reciprocal $1/\langle v_{j^*}^1 \rangle^{\sigma, \mathcal{M}'}$ is holomorphic on $D(t^*, \epsilon)$ with a zero at t^* . Thus, by the complex open mapping theorem [18, Thm. 10.32] applied to $1/\langle v_{j^*}^1 \rangle^{\sigma, \mathcal{M}'}$, there exists a $\delta > 0$ such that, for every $y \in D(0, \delta)$, there is a $z_y \in D(t^*, \epsilon)$ with $1/\langle v_{j^*}^1 \rangle^{\sigma, \mathcal{M}'}(z_y) = y$. Now, since $|t^{n,\mathbf{s}}| \rightarrow \infty$, we also have $|B + i t^{n,\mathbf{s}}| \rightarrow \infty$, so it follows that there exists a sequence $(z^{n,\mathbf{s}})_{n \in \mathbb{N}}$ in $D(t^*, \epsilon) \setminus \{t^*\}$ with $z^{n,\mathbf{s}} \rightarrow t^*$, such that $\langle v_{j^*}^1 \rangle^{\sigma, \mathcal{M}'}(z^{n,\mathbf{s}}) = B + i t^{n,\mathbf{s}}$ (a finite number of elements of the sequence $(t^{n,\mathbf{s}})_{n \in \mathbb{N}}$ may need to be discarded to ensure that $(z^{n,\mathbf{s}})_{n \in \mathbb{N}}$ is, indeed, contained in $D(t^*, \epsilon) \setminus \{t^*\}$). Now, for $p \in \{1, \dots, d\}$, compute

$$\begin{aligned} & \sigma \left(\omega_{v_p^2 v_{j^*}^1} \langle v_{j^*}^1 \rangle^{\sigma, \mathcal{M}'}(z^{n,\mathbf{s}}) + f_p(z^{n,\mathbf{s}}) + \theta_{v_p^2} \right) \\ & = \sigma \left(\omega_{v_p^2 v_{j^*}^1} (B + i t^{n,\mathbf{s}}) + f_p(z^{n,\mathbf{s}}) + \theta_{v_p^2} \right) \end{aligned} \quad (35)$$

$$= \sigma \left(\omega_{v_p^2 v_{j^*}^1} B + i(\omega_{v_p^2 v_{j^*}^1} t^{n,\mathbf{s}} + \mathbb{Z}) + f_p(z^{n,\mathbf{s}}) + \theta_{v_p^2} \right) \quad (36)$$

$$= \sigma \left(\omega_{v_p^2 v_{j^*}^1} B + i \cdot \left(\sum_{j=1}^{\bar{k}} \bar{q}_{pj} \omega_{v_j^2 v_{j^*}^1} r_j^{n,\mathbf{s}} + \mathbb{Z} \right) + f_p(z^{n,\mathbf{s}}) + \theta_{v_p^2} \right) \quad (37)$$

$$= \sigma \left(\sum_{j=1}^{\bar{k}} \bar{q}_{pj} \omega_{v_j^2 v_{j^*}^1} (B + i r_j^{n,\mathbf{s}}) + f_p(z^{n,\mathbf{s}}) + \theta_{v_p^2} \right), \quad (38)$$

where in (35) we used the definition of $z^{n,s}$, in (36) we used the i -periodicity of σ , in (37) we used (32), and in (38) we used $\omega_{v_p^2 v_{j^*}^1} = \sum_{j=1}^{\bar{k}} \bar{q}_{pj} \omega_{v_j^2 v_{j^*}^1}$ and the i -periodicity of σ again. As B was chosen so that the functions (29) do not have singularities in \mathcal{L}_B , all the quantities in the calculation (35)–(38) are well-defined.

Motivated by (35)–(38), we construct a GFNN $\mathcal{M}'' = (V^{\mathcal{M}''}, E^{\mathcal{M}''}, V_{in}^{\mathcal{M}''}, V_{out}^{\mathcal{M}''}, \Omega^{\mathcal{M}''}, \Theta^{\mathcal{M}''})$ as follows

- First, \bar{k} new nodes are created and enumerated as $\{u_1, \dots, u_{\bar{k}}\}$. Now, if $D_1 > 1$, then let $V_{in}^{\mathcal{M}''} = \{v_{in}, u_1, \dots, u_{\bar{k}}\}$, and if $D_1 = 1$, set $V_{in}^{\mathcal{M}''} = \{u_1, \dots, u_{\bar{k}}\}$.
- $V^{\mathcal{M}''} := V_{in}^{\mathcal{M}''} \cup (V_1 \setminus \{v_{j^*}^1\}) \cup \bigcup_{\ell \geq 2} V_\ell$.
- $E^{\mathcal{M}''} := \{(v, \tilde{v}) \in E^{\mathcal{M}'} : \text{lv}(v) \geq 2\} \cup \{(v_j^1, v_p^2) : j \in \{1, \dots, D_1\} \setminus \{j^*\}, p \in \{1, \dots, D_2\}\} \cup \{(u_j, v_p^2) : p \in \{1, \dots, d\}, j \in \{1, \dots, \bar{k}\}, \bar{q}_{pj} \neq 0\}$,
- $V_{out}^{\mathcal{M}''} := V_{out}^{\mathcal{M}'} \setminus \{v_{j^*}^1\}$,
- define $\omega_{v_p^2 u_j} := \bar{q}_{pj} \omega_{v_j^2 v_{j^*}^1}$, for $p = 1, \dots, d, j = 1, \dots, \bar{k}$, and let

$$\Omega^{\mathcal{M}''} := \{\omega_{\tilde{v}v} \in \Omega^{\mathcal{M}'} : \text{lv}(v) \geq 2\} \cup \{\omega_{v_p^2 v_j^1} : j \in \{1, \dots, D_1\} \setminus \{j^*\}, p \in \{1, \dots, D_2\}\} \cup \{\omega_{v_p^2 u_j} : p \in \{1, \dots, d\}, j \in \{1, \dots, \bar{k}\}, \bar{q}_{pj} \neq 0\},$$

– let

$$\Theta^{\mathcal{M}''} := \{\theta_v \in \Theta^{\mathcal{M}'} : \text{lv}(v) \geq 2\} \cup \{\theta_{v_j^1} : j \in \{1, \dots, D_1\} \setminus \{j^*\}\}.$$

The construction of \mathcal{M}'' for a concrete \mathcal{M}' is illustrated in Figure 10. Note that \mathcal{M}'' is not layered in the case $D_1 > 1$, due to the presence of the node v_{in} . Owing to (35)–(38) and the construction of \mathcal{M}'' , we have the following “input splitting” relationship:

$$\langle v_p^2 \rangle^{\sigma, \mathcal{M}''} (z^{n,s}) = \begin{cases} \langle v_p^2 \rangle^{\sigma, \mathcal{M}''} (z^{n,s}, B + i r_1^{n,s}, \dots, B + i r_{\bar{k}}^{n,s}), & \text{if } D_1 > 1 \\ \langle v_p^2 \rangle^{\sigma, \mathcal{M}''} (B + i r_1^{n,s}, \dots, B + i r_{\bar{k}}^{n,s}), & \text{if } D_1 = 1 \end{cases}, \quad (39)$$

for $p \in \{1, \dots, d\}$. We next show that \mathcal{M}'' is non-degenerate and clones-free. To establish non-

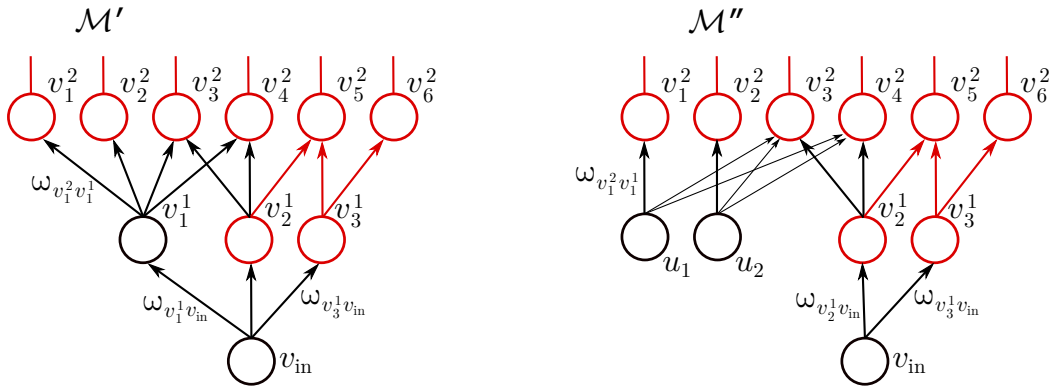


Fig. 10: *Input splitting, case $k = 1$. Left:* A neural network \mathcal{M}' , assumed to be a minimal element of \mathcal{M} , with $D_1 = 3$, $D_2 = 6$, and $d = 4$. *Right:* The corresponding network \mathcal{M}'' , assuming $j^* = 1$ and $\bar{k} = 2$. The function H in (42) and (43) corresponds to the map realized by the shared part (in red) of \mathcal{M}' and \mathcal{M}'' .

degeneracy, it suffices to show $V_{in}^{\mathcal{M}''} \subset \bigcup_{w \in V_{out}^{\mathcal{M}''}} V^{\mathcal{M}''(w)}$. First note that, in both cases $D_1 = 1$ and $D_1 > 1$, for a given $j \in \{1, \dots, \bar{k}\}$, there exists a $w \in V_{out}^{\mathcal{M}'} \setminus \{v_{j^*}^1\}$ such that $v_j^2 \in V^{\mathcal{M}'(w)}$, by non-degeneracy of \mathcal{M}' . It follows that $v_j^2 \in V^{\mathcal{M}''(w)}$ and thus $u_j \in V^{\mathcal{M}''(w)}$. As j was arbitrary, we have

$\{u_1, \dots, u_{\bar{k}}\} \subset \bigcup_{w \in V_{out}^{\mathcal{M}''}} V^{\mathcal{M}''(w)}$, which establishes non-degeneracy of \mathcal{M}'' in the case $D_1 = 1$. For $D_1 > 1$ we need to additionally show that $v_{in} \in V^{\mathcal{M}''(w)}$. To this end, note that there exist an $m^* \in \{1, \dots, D_1\} \setminus \{j^*\}$ and a $w \in V_{out}^{\mathcal{M}'} \setminus \{v_{j^*}^1\}$ such that $v_{m^*}^1 \in V^{\mathcal{M}'(w)}$, and so $v_{in} \in V^{\mathcal{M}''(w)}$, as desired. The clones-free property of \mathcal{M}'' follows by the same argument as in the case $k \geq 2$.

Once again, we revisit the function $h_{out}(t) = \sum_{w \in V_{out}^{\mathcal{M}'}} \lambda_w \langle w \rangle^{\sigma, \mathcal{M}'}(t) = c$, for all $t \in \mathcal{D}_{h_{out}}$, and proceed in a similar fashion as in the case $k \geq 2$. This time, however, the output sets $V_{out}^{\mathcal{M}'}$ and $V_{out}^{\mathcal{M}''}$ may differ by the node $v_{j^*}^1$. This is a nuisance that will be dealt with below in Claim 2, but in the meantime, it is convenient to introduce the ‘‘truncated’’ linear dependency function

$$h_{tr} := \sum_{w \in V_{out}^{\mathcal{M}'} \setminus \{v_{j^*}^1\}} \lambda_w \langle w \rangle^{\sigma, \mathcal{M}'}, \quad (40)$$

and proceed exactly as in the case $k \geq 2$. By examining the structure of \mathcal{M}' , we see that, for each $w \in V_{out}^{\mathcal{M}'} \setminus \{v_{j^*}^1\}$, we can write

$$\langle w \rangle^{\sigma, \mathcal{M}'}(z) = H_w \left(\left(\langle v_p^2 \rangle^{\sigma, \mathcal{M}'}(z) \right)_{p=1}^d, \left(\langle v_j^1 \rangle^{\sigma, \mathcal{M}'}(z) \right)_{j \in \{1, \dots, D_1\} \setminus \{j^*\}} \right), \quad \forall z \in \mathcal{D}_{\langle w \rangle^{\sigma, \mathcal{M}'}}$$

where $H_w : \mathcal{D}_{H_w} \rightarrow \mathbb{C}$ corresponds to the map realized by the GFNN with nodes

$$V^{\mathcal{M}} \setminus \{v_{in}, v_{j^*}^1\}, \quad (41)$$

inputs $\{v_p^2\}_{p=1}^d \cup \{v_j^1\}_{j \in \{1, \dots, D_1\} \setminus \{j^*\}}$, single output $\{w\}$, and edges, weights, and biases inherited from \mathcal{M}' . The function $H_w : \mathcal{D}_{H_w} \rightarrow \mathbb{C}$ is holomorphic on its natural domain $\mathcal{D}_{H_w} \subset \mathbb{C}^{d+(D_1-1)}$ containing $\mathbb{R}^{d+(D_1-1)}$. We can therefore write

$$h_{tr}(z) = H \left(\left(\langle v_p^2 \rangle^{\sigma, \mathcal{M}'}(z) \right)_{p=1}^d, \left(\langle v_j^1 \rangle^{\sigma, \mathcal{M}'}(z) \right)_{j \in \{1, \dots, D_1\} \setminus \{j^*\}} \right), \quad \forall z \in \mathcal{D}_{h_{out}}, \quad (42)$$

where $H : \mathcal{D}_H \rightarrow \mathbb{C}$, $H = \sum_{w \in V_{out}^{\mathcal{M}'} \setminus \{v_{j^*}^1\}} \lambda_w H_w$, is holomorphic on $\mathcal{D}_H = \bigcap_{w \in V_{out}^{\mathcal{M}'} \setminus \{v_{j^*}^1\}} \mathcal{D}_{H_w} \supset \mathbb{R}^{d+(D_1-1)}$.

Now, by definition of natural domain, for each $w \in V_{out}^{\mathcal{M}''}$, the natural domain $\mathcal{D}_{\langle w \rangle^{\sigma, \mathcal{M}''}}$ is the set of all $z \in \bigcap_{p=1}^d \mathcal{D}_{\langle v_p^2 \rangle^{\sigma, \mathcal{M}''}} \cap \bigcap_{j \neq j^*} \mathcal{D}_{\langle v_j^1 \rangle^{\sigma, \mathcal{M}''}}$ such that

$$\left(\left(\langle v_p^2 \rangle^{\sigma, \mathcal{M}''}(z) \right)_{p=1}^d, \left(\langle v_j^1 \rangle^{\sigma, \mathcal{M}''}(z) \right)_{j \in \{1, \dots, D_1\} \setminus \{j^*\}} \right) \in \mathcal{D}_{H_w},$$

where the variable $z = (z_0, z_1, \dots, z_{\bar{k}})$ corresponds to the input nodes $v_{in}, u_1, \dots, u_{\bar{k}}$, in the case $D_1 > 1$, and $z = (z_1, \dots, z_{\bar{k}})$ corresponds to the input nodes $u_1, \dots, u_{\bar{k}}$, in the case $D_1 = 1$. Therefore, for z in the open domain $\mathcal{D}_{\tilde{h}_{out}} := \bigcap_{w \in V_{out}^{\mathcal{M}''}} \mathcal{D}_{\langle w \rangle^{\sigma, \mathcal{M}''}}$, we can define the function $\tilde{h}_{out} : \mathcal{D}_{\tilde{h}_{out}} \rightarrow \mathbb{C}$ according to

$$\tilde{h}_{out}(z) = H \left(\left(\langle v_p^2 \rangle^{\sigma, \mathcal{M}''}(z) \right)_{p=1}^d, \left(\langle v_j^1 \rangle^{\sigma, \mathcal{M}''}(z) \right)_{j \in \{1, \dots, D_1\} \setminus \{j^*\}} \right). \quad (43)$$

Moreover, as \mathcal{M}' and \mathcal{M}'' share the nodes in (41), as well as the associated edges, weights, and biases, we have

$$\langle w \rangle^{\sigma, \mathcal{M}''}(z) = H_w \left(\left(\langle v_p^2 \rangle^{\sigma, \mathcal{M}''}(z) \right)_{p=1}^d, \left(\langle v_j^1 \rangle^{\sigma, \mathcal{M}''}(z) \right)_{j \in \{1, \dots, D_1\} \setminus \{j^*\}} \right)$$

for all $w \in V_{out}^{\mathcal{M}''}$, and thus

$$\tilde{h}_{out} = \sum_{w \in V_{out}^{\mathcal{M}''}} \lambda_w \langle w \rangle^{\sigma, \mathcal{M}''}.$$

At this point we verify another auxiliary claim, which states that h_{tr} and h_{out} are always, in fact, the same function, and therefore $\tilde{h}_{out} \equiv c$ follows by a similar argument as in the case $k \geq 2$.

Claim 2: Recall that $t^ \in \mathbb{R} + \frac{i}{2a}$ is such that $\langle v_{j^*}^1 \rangle^{\sigma, \mathcal{M}'}$ has a pole at t^* , and all the other $\langle v_j^1 \rangle^{\sigma, \mathcal{M}'}$, $j \in \{1, \dots, D_1\} \setminus \{j^*\}$, are analytic and real-valued at t^* . Further recall the open set $C \subset \mathbb{R}^{\bar{k}}$ containing $\mathbf{0}$. We have $\{t^*\} \times \mathbb{R}^{\bar{k}} \subset \mathcal{D}_{\tilde{h}_{out}}$ and $\tilde{h}_{out}|_{\mathbb{R}^{\bar{k}+1}} \equiv c$, in the case $D_1 > 1$, and $\mathbb{R}^{\bar{k}} \subset \mathcal{D}_{\tilde{h}_{out}}$ and $\tilde{h}_{out}|_{\mathbb{R}^{\bar{k}}} \equiv c$, in the case $D_1 = 1$. Moreover, in both cases we have $v_{j^*}^1 \notin V_{out}^{\mathcal{M}'}$.*

Proof of Claim 2. First assume that $D_1 > 1$. To show that $\{t^*\} \times \mathbb{R}^{\bar{k}} \subset \mathcal{D}_{\tilde{h}_{out}}$, first observe that, for $j \in \{1, \dots, D_1\} \setminus \{j^*\}$ and $(z_1, \dots, z_{\bar{k}}) \in \mathbb{R}^{\bar{k}}$, we have $\langle v_j^1 \rangle^{\sigma, \mathcal{M}''}(t^*, z_1, \dots, z_{\bar{k}}) = \langle v_j^1 \rangle^{\sigma, \mathcal{M}'}(t^*)$, which, by (27), is a real number. By (30), this further implies $f_p(t^*) \in \mathbb{R}$, for $p = 1, \dots, d$. Therefore

$$\langle v_p^2 \rangle^{\sigma, \mathcal{M}''}(t^*, z_1, \dots, z_{\bar{k}}) = \sigma \left(\sum_{j=1}^{\bar{k}} \bar{q}_{pj} \omega_{v_j^2 v_{j^*}^1} z_j + f_p(t^*) + \theta_{v_p^2} \right) \in \mathbb{R},$$

for $p \in \{1, \dots, d\}$ and $(z_1, \dots, z_{\bar{k}}) \in \mathbb{R}^{\bar{k}}$. As $\mathbb{R}^{d+(D_1-1)} \subset \mathcal{D}_H$, we deduce that

$$\left(\left(\langle v_p^2 \rangle^{\sigma, \mathcal{M}''}(z) \right)_{p=1}^d, \left(\langle v_j^1 \rangle^{\sigma, \mathcal{M}''}(z) \right)_{j \in \{1, \dots, D_1\} \setminus \{j^*\}} \right) \in \mathcal{D}_H, \text{ for all } z \in \{t^*\} \times \mathbb{R}^{\bar{k}}.$$

This establishes $\{t^*\} \times \mathbb{R}^{\bar{k}} \subset \mathcal{D}_{\tilde{h}_{out}}$. We proceed to showing $\tilde{h}_{out}|_{\mathbb{R}^{\bar{k}+1}} \equiv c$. As $\mathcal{D}_{\tilde{h}_{out}}$ is open, it follows that $\mathcal{D}_{\tilde{h}_{out}} \supset \mathcal{U}$, for some connected open $\mathcal{U} \subset \mathbb{C}^{1+\bar{k}}$ containing $\{t^*\} \times \mathbb{R}^{\bar{k}}$. Choose a small enough $\delta > 0$ so that $\mathcal{U} \supset D_1^\circ(t^*, \delta) \times D_{\bar{k}}^\circ((B, \dots, B), \delta)$. Now, fix an arbitrary $\mathbf{s} = (s_1, \dots, s_{\bar{k}})$ in the smaller open set $C \cap D_{\bar{k}}^\circ(\mathbf{0}, \delta)$. We then have

$$(t^*, B + is_1, \dots, B + is_{\bar{k}}) \in D_1^\circ(t^*, \delta) \times D_{\bar{k}}^\circ((B, \dots, B), \delta) \subset \mathcal{U} \subset \mathcal{D}_{\tilde{h}_{out}},$$

and since

$$(z^{n, \mathbf{s}}, B + ir_1^{n, \mathbf{s}}, \dots, B + ir_{\bar{k}}^{n, \mathbf{s}}) \rightarrow (t^*, B + is_1, \dots, B + is_{\bar{k}}),$$

as $n \rightarrow \infty$, we obtain

$$(z^{n, \mathbf{s}}, B + ir_1^{n, \mathbf{s}}, \dots, B + ir_{\bar{k}}^{n, \mathbf{s}}) \in \mathcal{D}_{\tilde{h}_{out}},$$

for large enough $n \in \mathbb{N}$. We may again assume w.l.o.g. that this is true for all $n \in \mathbb{N}$ by discarding finitely many elements of the sequences $(z^{n, \mathbf{s}})_{n \in \mathbb{N}}$ and $(r^{n, \mathbf{s}})_{n \in \mathbb{N}}$. Now, we use (39), (42), and (43) to get

$$\tilde{h}_{out}(z^{n, \mathbf{s}}, B + ir_1^{n, \mathbf{s}}, \dots, B + ir_{\bar{k}}^{n, \mathbf{s}}) = h_{tr}(z^{n, \mathbf{s}}), \quad \forall n \in \mathbb{N}, \quad (44)$$

for all $\mathbf{s} \in C \cap D_{\bar{k}}^\circ(\mathbf{0}, \delta)$. We are now ready to show that $v_{j^*}^1 \notin V_{out}^{\mathcal{M}'}$ (still in the case $D_1 > 1$). To this end, suppose by way of contradiction that $v_{j^*}^1 \in V_{out}^{\mathcal{M}'}$ and set $\mathbf{s} = \mathbf{0}$. Note that $\tilde{h}_{out}(t^*, B, \dots, B)$ is a well-defined (finite) complex number, simply as $(t^*, B, \dots, B) \in \{t^*\} \times \mathbb{R}^{\bar{k}} \subset \mathcal{D}_{\tilde{h}_{out}}$. Thus, by (40) and (44), we have

$$\begin{aligned} \langle v_{j^*}^1 \rangle^{\sigma, \mathcal{M}'}(z^{n, \mathbf{0}}) &= c - h_{out}(z^{n, \mathbf{0}}) + \langle v_{j^*}^1 \rangle^{\sigma, \mathcal{M}'}(z^{n, \mathbf{0}}) \\ &= c - h_{tr}(z^{n, \mathbf{0}}) \\ &= c - \tilde{h}_{out}(z^{n, \mathbf{0}}, B + ir_1^{n, \mathbf{0}}, \dots, B + ir_{\bar{k}}^{n, \mathbf{0}}) \\ &\rightarrow c - \tilde{h}_{out}(t^*, B, \dots, B) \end{aligned}$$

as $n \rightarrow \infty$, which contradicts the fact that $\langle v_{j^*}^1 \rangle^{\sigma, \mathcal{M}'}$ has a pole at t^* . This establishes $v_{j^*}^1 \notin V_{out}^{\mathcal{M}'}$. As a consequence we further have $h_{tr} = h_{out}$, and so (44) reads

$$\tilde{h}_{out}(z^{n, \mathbf{s}}, B + ir_1^{n, \mathbf{s}}, \dots, B + ir_{\bar{k}}^{n, \mathbf{s}}) = h_{out}(z^{n, \mathbf{s}}) = c, \quad \forall n \in \mathbb{N},$$

for all $\mathbf{s} \in C \cap D_k^\circ(0, \delta)$. Now, define the set

$$\tilde{T} = \{(z^{n,\mathbf{s}}, B + i r_1^{n,\mathbf{s}}, \dots, B + i r_k^{n,\mathbf{s}}) : \mathbf{s} \in C \cap D_k^\circ(0, \delta), n \in \mathbb{N}\}.$$

Note that \tilde{T} satisfies

$$\begin{aligned} \tilde{T} &\subset (D(t^*, \epsilon) \setminus \{t^*\}) \times \mathbb{C}^{\bar{k}} \quad \text{and} \\ \text{cl}(\tilde{T}) &\supset \{t^*\} \times ((B, \dots, B) + (iC) \cap D_k^\circ(0, \delta)), \end{aligned}$$

so by Lemma 5, it follows that $\tilde{h}_{out} - c \equiv 0$ everywhere in an open neighborhood of $\mathbb{R}^{\bar{k}+1}$, and thus $\tilde{h}_{out}|_{\mathbb{R}^{\bar{k}+1}} \equiv c$ in particular. This establishes Claim 2 in the case $D_1 > 1$. It remains to prove the claim for $D_1 = 1$. Showing that $\mathbb{R}^{\bar{k}} \subset \mathcal{D}_{\tilde{h}_{out}}$ is fully analogous to showing $\{t^*\} \times \mathbb{R}^{\bar{k}} \subset \mathcal{D}_{\tilde{h}_{out}}$ in the case $D_1 > 1$. We can hence proceed to establishing $\tilde{h}_{out}|_{\mathbb{R}^{\bar{k}}} \equiv c$. To this end, we first note that there is a connected open set \mathcal{U} and a $\delta > 0$ such that $\mathbb{R}^{\bar{k}} \subset \mathcal{U} \subset \mathcal{D}_{\tilde{h}_{out}}$ and $D_k^\circ((B, \dots, B), \delta) \subset \mathcal{U}$, and we similarly obtain

$$(B + i r_1^{n,\mathbf{s}}, \dots, B + i r_k^{n,\mathbf{s}}) \in \mathcal{D}_{\tilde{h}_{out}},$$

for all $n \in \mathbb{N}$ and $\mathbf{s} \in C \cap D_k^\circ(0, \delta)$. Again, showing $v_{j^*}^1 \notin V_{out}^{\mathcal{M}'}$ now proceeds in a manner entirely analogous to the case $D_1 > 1$, as does obtaining the identity

$$\tilde{h}_{out}(B + i r_1^{n,\mathbf{s}}, \dots, B + i r_k^{n,\mathbf{s}}) = h_{out}(z^{n,\mathbf{s}}) \equiv c, \quad \forall n \in \mathbb{N},$$

for all $\mathbf{s} \in C \cap D_k^\circ(0, \delta)$. Now, define the set

$$T = \{(B + i r_1^{n,\mathbf{s}}, \dots, B + i r_k^{n,\mathbf{s}}) : \mathbf{s} \in C \cap D_k^\circ(0, \delta), n \in \mathbb{N}\}.$$

Note that T satisfies $\text{cl}(T) \supset ((B, \dots, B) + (iC) \cap D_k^\circ(0, \delta))$, so, by Lemma 4, we have $\tilde{h}_{out} \equiv c$ everywhere in an open neighborhood of $\mathbb{R}^{\bar{k}}$, which concludes the proof of Claim 2.




Fig. 11: *Input anchoring.* Left: The neural network \mathcal{M}'' as in Figure 10. Note that \mathcal{M}'' is not layered, but every network obtained from \mathcal{M}'' by anchoring all but one of its input nodes is layered. Right: Anchoring the inputs of \mathcal{M}'' at the nodes $v_{in}, u_1, u_2, \dots, u_{\bar{k}-1}$ yields a layered neural network \mathcal{N} with $L(\mathcal{N}) = L(\mathcal{M}) - 1$.

Finally, it remains to apply an input anchoring procedure to \mathcal{M}'' , which will conclude the proof in a manner similar to the case $k \geq 2$. Specifically, we use Lemma 2 to successively eliminate inputs of \mathcal{M}'' , starting with v_{in} (if present), and proceeding with $u_1, \dots, u_{\bar{k}-1}$. If $D_1 > 1$, the network \mathcal{M}'' is not layered (unlike in the case $k \geq 2$ and the case $k = 1, D_1 = 1$). However, every network obtained from \mathcal{M}'' by anchoring all but one of the input nodes $\{v_{in}, u_1, \dots, u_{\bar{k}}\}$ is layered. This means that,

when anchoring v_{in} , we do not find ourselves in the circumstance (ii) of Lemma 2, as this would mean we have obtained a network $\mathcal{N} \in \mathcal{M}_{min}$ with strictly fewer nodes than \mathcal{M} . Thus, after having anchored v_{in} , we are left with a layered network with inputs $u_1, \dots, u_{\bar{k}}$. At this point we proceed completely analogously to the case $k \geq 2$ by successively eliminating the inputs $u_1, \dots, u_{\bar{k}-1}$. We are left with a non-degenerate clones-free LFNN $\mathcal{N} = (V^{\mathcal{N}}, E^{\mathcal{N}}, \{u_{\bar{k}}\}, V_{out}^{\mathcal{N}}, \Omega^{\mathcal{N}}, \Theta^{\mathcal{N}})$ and a vector of real constants \mathbf{a} (specifically, $\mathbf{a} \in \mathbb{R}^{\bar{k}}$ in the case $D_1 > 1$, and $\mathbf{a} \in \mathbb{R}^{\bar{k}-1}$ in the case $D_1 = 1$), such that the function $h_{out}^{\mathcal{N}} := \sum_{w \in V_{out}^{\mathcal{N}}} \lambda_w \langle w \rangle^{\sigma, \mathcal{N}}$ satisfies

$$h_{out}^{\mathcal{N}}(t) = \tilde{h}_{out}(\mathbf{a}, t) - \sum_{w \in V_{out}^{\mathcal{M}''} \setminus V_{out}^{\mathcal{N}}} \lambda_w \langle w \rangle^{\sigma, \mathcal{M}''}(\mathbf{a}, t), \quad \forall t \in \mathbb{R}. \quad (45)$$

A concrete example of this input anchoring procedure in the case $k \geq 2$ is shown schematically in Figure 11. By Claim 2, the first term on the right-hand side of (45) evaluates identically to c . Moreover, as input anchoring yields networks satisfying (IA-2), the values of the functions $\langle w \rangle^{\sigma, \mathcal{M}''}$, for $w \in V_{out}^{\mathcal{M}''} \setminus V_{out}^{\mathcal{N}}$, do not depend on the input at $u_{\bar{k}}$. Therefore $h_{out}^{\mathcal{N}} \equiv c_{\mathcal{N}}$, for some $c_{\mathcal{N}} \in \mathbb{R}$. We have thus shown that the network \mathcal{N} is in \mathcal{M} . But $L(\mathcal{N}) = L(\mathcal{M}) - 1$, which stands in contradiction to the minimality of depth of the elements of \mathcal{M}_{min} , and therefore completes the proof of the theorem. ■

Proof of Theorem 3: Let $\mathcal{N}_j = (V^j, E^j, V_{in}, V_{out}, \Omega^j, \Theta^j)$, $j \in \{1, 2\}$, be networks as in the theorem statement. Let $\mathcal{N} = \mathcal{N}_1 \vee \mathcal{N}_2$ be their amalgam and $\pi_j : V^{\mathcal{N}_j} \rightarrow \pi_j(V^{\mathcal{N}_j}) \subset V^{\mathcal{N}}$ the extensional isomorphisms between \mathcal{N}_j and the corresponding subnetworks of \mathcal{N} , for $j \in \{1, 2\}$. We start by claiming that $\pi_1(w) = \pi_2(w)$, for all $w \in V_{out}$. Indeed, suppose to the contrary that we have $\pi_1(w') \neq \pi_2(w')$, for some $w' \in V_{out}$, and denote $w_j = \pi_j(w')$, $j \in \{1, 2\}$. Since $w_1 \neq w_2$, it follows that $\mathcal{N}(w_1)$ and $\mathcal{N}(w_2)$ are not extensionally isomorphic, for otherwise w_1 and w_2 would be clones, contradicting the no-clones condition for \mathcal{N} . Now,

$$\langle \mathcal{N}(w_1) \rangle^{\sigma}(\mathbf{t}) - \langle \mathcal{N}(w_2) \rangle^{\sigma}(\mathbf{t}) = \langle w' \rangle^{\sigma, \mathcal{N}_1}(\mathbf{t}) - \langle w' \rangle^{\sigma, \mathcal{N}_2}(\mathbf{t}) = 0, \quad \text{for all } \mathbf{t} \in \mathbb{R}^{V_{in}},$$

by assumption. But this contradicts the conclusion of Theorem 4, and thus establishes $\pi_1(w) = \pi_2(w)$, for all $w \in V_{out}$. By non-degeneracy of \mathcal{N}_1 , for every $v \in V^1$, there exists a $w \in V_{out}$ such that $v \in V^{\mathcal{N}_1(w)}$. Then $\pi_1(v) \in V^{\mathcal{N}(\pi_1(w))} = V^{\mathcal{N}(\pi_2(w))} = \pi_2(V^{\mathcal{N}_2(w)}) \subset \pi_2(V^2)$. Similarly, for every $v \in V^2$, we have $\pi_2(v) \in \pi_1(V^1)$. Thus, the function $\psi : V^1 \rightarrow V^2$ given by $\psi = \pi_2^{-1} \circ \pi_1$ is well-defined. This function is invertible with inverse $\pi_1^{-1} \circ \pi_2$, so it is a bijection. Therefore ψ is an extensional isomorphism between \mathcal{N}_1 and \mathcal{N}_2 , by virtue of being a composition of two extensional isomorphisms. Moreover, we have $\psi(w) = \pi_2^{-1}(\pi_1(w)) = w$, for all $w \in V_{out}$, so ψ restricted to V_{out} is the identity map, and thus ψ is a faithful isomorphism. ■

ACKNOWLEDGMENT

The authors would like to thank Thomas Allard for useful suggestions regarding the proof of Proposition 3 and an anonymous reviewer for proposing a clearer exposition of Lemma 6.

REFERENCES

- [1] H. J. Sussman, "Uniqueness of the weights for minimal feedforward nets with a given input-output map," *Neural Networks*, vol. 5, no. 4, pp. 589–593, July 1992.

- [2] F. Albertini, E. D. Sontag, and V. Maillot, “Uniqueness of weights for neural networks,” *Artificial Neural Networks for Speech and Vision*, pp. 113–125, 1993.
- [3] C. Fefferman, “Reconstructing a neural net from its output,” *Revista Matemática Iberoamericana*, vol. 10, no. 3, pp. 507–555, 1994.
- [4] Y. LeCun, L. D. Jackel, L. Bottou, A. Brunot, C. Cortes, J. S. Denker, H. Drucker, I. Guyon, U. A. Müller, E. Säckinger, P. Simard, and V. Vapnik, “Comparison of learning algorithms for handwritten digit recognition,” *International Conference on Artificial Neural Networks*, pp. 53–60, 1995.
- [5] A. Krizhevsky, I. Sutskever, and G. E. Hinton, “Imagenet classification with deep convolutional neural networks,” in *Advances in Neural Information Processing Systems 25*. Curran Associates, Inc., 2012, pp. 1097–1105. [Online]. Available: <http://papers.nips.cc/paper/4824-imagenet-classification-with-deep-convolutional-neural-networks.pdf>
- [6] G. Hinton, L. Deng, D. Yu, G. E. Dahl, A. R. Mohamed, N. Jaitly, A. Senior, V. Vanhoucke, P. Nguyen, T. N. Sainath, and B. Kingsbury, “Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups,” *IEEE Signal Process. Mag.*, vol. 29, no. 6, pp. 82–97, 2012.
- [7] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. MIT Press, 2016.
- [8] H. Bölcskei, P. Grohs, G. Kutyniok, and P. Petersen, “Optimal approximation with sparsely connected deep neural networks,” *SIAM Journal on Mathematics of Data Science*, vol. 1, no. 1, pp. 8–45, 2019.
- [9] S. Mallat, “Group invariant scattering,” *Comm. Pure and Appl. Math.*, vol. 65, no. 10, pp. 1331–1398, 2012.
- [10] P. Petersen and F. Voigtländer, “Optimal approximation of piecewise smooth functions using deep ReLU neural networks,” *Neural Networks*, no. 108, pp. 296–330, 2018.
- [11] T. Wiatowski and H. Bölcskei, “A mathematical theory of deep convolutional neural networks for feature extraction,” *IEEE Transactions on Information Theory*, vol. 64, no. 3, pp. 1845–1866, Mar. 2018.
- [12] F. Albertini and E. D. Sontag, “For neural networks, function determines form,” *Neural Networks*, vol. 6, pp. 975–990, 1993.
- [13] —, “Uniqueness of weights for recurrent networks,” vol. 2. Akademie Verlag, Regensburg, 1993, pp. 599–602.
- [14] K. Hornik, M. Stinchcombe, and H. White, “Multilayer feedforward networks are universal approximators,” *Neural Networks*, vol. 2, no. 5, pp. 359–366, 1989.
- [15] G. Cybenko, “Approximation by superpositions of a sigmoidal function,” *Mathematics of Control, Signals, and Systems*, vol. 2, no. 4, pp. 303–314, Dec. 1989.
- [16] L. Kronecker, *Werke (Näherungsweise ganzzahlige Auflösung linearer Gleichungen)*. reprint, Chelsea, 1979, vol. 3.
- [17] A. Bondy and M. R. Murty, *Graph Theory*, ser. Graduate Texts in Mathematics. Springer, 2008.
- [18] W. Rudin, *Real and Complex Analysis*, 3rd ed., ser. Higher Mathematics. McGraw-Hill, 1987.
- [19] V. Scheidemann, *Introduction to Complex Analysis in Several Variables*. Birkhäuser, 2005.
- [20] B. C. Hall, *Lie Groups, Lie Algebras, and Representations: An Elementary Introduction*, 2nd ed., ser. Graduate Texts in Mathematics. Springer, 2015.
- [21] D. Garling, *A Course in Mathematical Analysis*. Cambridge University Press, 2013, vol. 2.

APPENDIX: PROOFS OF AUXILIARY RESULTS

Proof of Proposition 1: Fix \mathcal{N}_1 and \mathcal{N}_2 as in the statement of the proposition. We begin by establishing the existence of a corresponding amalgam \mathcal{A} . Let \mathcal{A} denote the set of all proto-amalgams of \mathcal{N}_1 and \mathcal{N}_2 . To see that \mathcal{A} is non-empty, consider the LFNN $\mathcal{N} = (V^{\mathcal{N}}, E^{\mathcal{N}}, V_{in}, V_{out}^{\mathcal{N}}, \Omega^{\mathcal{N}}, \Theta^{\mathcal{N}})$ specified as follows:

- Let S be a set of cardinality $\#(V^1 \setminus V_{in}) + \#(V^2 \setminus V_{in})$ disjoint from V_{in} , and set $V^{\mathcal{N}} := V_{in} \cup S$. Furthermore, let $\pi_j^{\mathcal{N}} : V^j \rightarrow \pi_j^{\mathcal{N}}(V^j) \subset V^{\mathcal{N}}$ be injective functions such that $\pi_j^{\mathcal{N}}(v) = v$, for $v \in V_{in}$, $j \in \{1, 2\}$, and $\pi_1^{\mathcal{N}}(V^1 \setminus V_{in}) \cap \pi_2^{\mathcal{N}}(V^2 \setminus V_{in}) = \emptyset$, but otherwise arbitrary.
- $E^{\mathcal{N}} := \bigcup_{j=1,2} \{(\pi_j^{\mathcal{N}}(v), \pi_j^{\mathcal{N}}(\tilde{v})) : v, \tilde{v} \in V^j, (v, \tilde{v}) \in E^j\}$.
- $V_{out}^{\mathcal{N}} := \pi_1^{\mathcal{N}}(V_{out}^1) \cup \pi_2^{\mathcal{N}}(V_{out}^2)$.
- For $j \in \{1, 2\}$ and $v, \tilde{v} \in V^j$ such that $(v, \tilde{v}) \in E^j$, let $\omega_{\pi_j^{\mathcal{N}}(\tilde{v})\pi_j^{\mathcal{N}}(v)} = \omega_{\tilde{v}v}$, and set $\Omega^{\mathcal{N}} := \{\omega_{vu} : (u, v) \in E^{\mathcal{N}}\}$.
- For $j = 1, 2$ and $v \in V^j \setminus V_{in}$, let $\theta_{\pi_j^{\mathcal{N}}(v)} = \theta_v$, and set $\Theta^{\mathcal{N}} := \{\theta_u : u \in V^{\mathcal{N}} \setminus V_{in}\}$.

Informally, the network \mathcal{N} is obtained by putting \mathcal{N}_1 and \mathcal{N}_2 “side by side”, sharing only the input nodes V_{in} . As \mathcal{N}_1 and \mathcal{N}_2 are non-degenerate, so is \mathcal{N} . Moreover, Properties (i) and (ii) of Definition 16 hold for \mathcal{N} with $\pi_j^{\mathcal{N}} : V^j \rightarrow \pi_j(V^j) \subset V^{\mathcal{N}}$, for $j = 1, 2$.

Thus \mathcal{N} is a proto-amalgam of \mathcal{N}_1 and \mathcal{N}_2 , and so $\mathcal{A} \neq \emptyset$. Now, let $\mathcal{A} = (V^{\mathcal{A}}, E^{\mathcal{A}}, V_{in}^{\mathcal{A}}, V_{out}^{\mathcal{A}}, \Omega^{\mathcal{A}}, \Theta^{\mathcal{A}}) \in \mathcal{A}$ be a network with the least possible number of nodes among all the networks in \mathcal{A} , and let $\pi_j : V^j \rightarrow \pi_j(V^j) \subset V^{\mathcal{A}}$, for $j \in \{1, 2\}$, be extensional isomorphisms between \mathcal{N}_j and the appropriate subnetworks of \mathcal{A} . We now show that \mathcal{A} is clones-free. To this end, suppose by way of contradiction that $c_1, c_2 \in V^{\mathcal{A}}$ are clones. As \mathcal{N}_1 is clones-free, c_1, c_2 cannot both be in $\pi_1(V^1)$, for otherwise $\pi_1^{-1}(c_1)$ and $\pi_1^{-1}(c_2)$ would be clones in \mathcal{N}_1 . By the same token, c_1, c_2 cannot both be in $\pi_2(V^2)$. Thus, we may write w.l.o.g. $c_1 = \pi_1(v_1)$ and $c_2 = \pi_2(v_2)$, for some $v_1 \in V^1$ and $v_2 \in V^2$. Now, let $\tilde{\mathcal{A}}$ be the network obtained from \mathcal{A} by making the following alterations:

- For every edge $(c_2, v) \in E^{\mathcal{A}}$, where $v \in V^{\mathcal{A}}$, introduce a new edge (c_1, v) together with the associated weight ω_{vc_2} , and delete the edge (c_2, v) .
- Delete the edges $(v, c_2) \in E^{\mathcal{A}}$, as well as the node c_2 .
- If c_2 was a node in $\pi_2(V_{out}^2)$, then add c_1 to the set $V_{out}^{\tilde{\mathcal{A}}}$.

The network $\tilde{\mathcal{A}}$ is a proto-amalgam of \mathcal{N}_1 and \mathcal{N}_2 via the extensional isomorphisms $\tilde{\pi}_1 = \pi_1$ and

$$\tilde{\pi}_2(v) = \begin{cases} \pi_2(v), & v \in V^2 \setminus \{\pi_2^{-1}(c_2)\} \\ c_1, & v = \pi_2^{-1}(c_2) \end{cases}, \quad \text{for } v \in V^{\mathcal{N}_2}.$$

But $\tilde{\mathcal{A}}$ has strictly fewer nodes than \mathcal{A} , which contradicts the minimality of \mathcal{A} , and thereby establishes that \mathcal{A} is clones-free, and hence \mathcal{A} is an amalgam of \mathcal{N}_1 and \mathcal{N}_2 , completing the proof of existence. To establish uniqueness—up to extensional isomorphisms—of the amalgam, suppose that \mathcal{A} and \mathcal{A}' are both amalgams of \mathcal{N}_1 and \mathcal{N}_2 via extensional isomorphisms $\pi_j : V^j \rightarrow \pi_j(V^j) \subset V^{\mathcal{A}}$, $\pi'_j : V^j \rightarrow \pi'_j(V^j) \subset V^{\mathcal{A}'}$, for $j \in \{1, 2\}$. We first show that

$$(\pi'_1 \circ \pi_1^{-1})(v) = (\pi'_2 \circ \pi_2^{-1})(v), \quad \text{for all } v \in \pi_1(V_1) \cap \pi_2(V_2), \quad (46)$$

by induction on $\text{lv}_{\mathcal{A}}(v)$. If $v \in V_{in}$, then (46) holds trivially as the restrictions of the maps π_j, π'_j , for $j \in \{1, 2\}$, to the set V_{in} , both equal the identity map $\text{id}_{V_{in}}$. Now, let $L \geq 1$ and suppose that (46) holds for all $u \in \pi_1(V_1) \cap \pi_2(V_2)$ with $\text{lv}_{\mathcal{A}}(u) < L$. Let $v \in \pi_1(V_1) \cap \pi_2(V_2)$ with $\text{lv}_{\mathcal{A}}(v) = L$, but otherwise arbitrary, and write $w_j = (\pi'_j \circ \pi_j^{-1})(v)$, for $j = 1, 2$. By Property (i) of Definition 16 for the amalgam \mathcal{A} we have $\mathcal{N}_1 \stackrel{e}{\sim} \mathcal{A}(\pi_1(V_{out}^1))$ and $\mathcal{N}_2 \stackrel{e}{\sim} \mathcal{A}(\pi_2(V_{out}^2))$, and so $\mathcal{N}_1(\pi_1^{-1}(v)) \stackrel{e}{\sim} \mathcal{A}(v)$ and $\mathcal{N}_2(\pi_2^{-1}(v)) \stackrel{e}{\sim} \mathcal{A}(v)$ by appropriately restricting π_1 and π_2 . Similarly, $\mathcal{N}_1((\pi'_1)^{-1}(w_1)) \stackrel{e}{\sim} \mathcal{A}'(w_1)$ and $\mathcal{N}_2((\pi'_2)^{-1}(w_2)) \stackrel{e}{\sim} \mathcal{A}'(w_2)$. But $(\pi'_j)^{-1}(w_j) = \pi_j^{-1}(v)$, and so $\mathcal{N}_j((\pi'_j)^{-1}(w_j)) = \mathcal{N}_j(\pi_j^{-1}(v))$, for $j \in \{1, 2\}$. Therefore $\mathcal{A}'(w_1) \stackrel{e}{\sim} \mathcal{A}(v)$ and $\mathcal{A}'(w_2) \stackrel{e}{\sim} \mathcal{A}(v)$ via $\pi_1 \circ (\pi'_1)^{-1}$ and $\pi_2 \circ (\pi'_2)^{-1}$, respectively. Now, as \mathcal{A}' is an amalgam, it is clones-free, and thus we deduce that $w_1 = w_2$, for otherwise w_1 and w_2 would be clones in \mathcal{A}' . This establishes (46).

Now define $\psi : V^{\mathcal{A}} \rightarrow V^{\mathcal{A}'}$ according to

$$\psi(v) = \begin{cases} (\pi'_1 \circ \pi_1^{-1})(v), & v \in \pi_1(V_1) \\ (\pi'_2 \circ \pi_2^{-1})(v), & v \in \pi_2(V_2) \end{cases}. \quad (47)$$

It follows by (46) that this definition is consistent, in the sense that the two cases in (47) yield the same value for $\psi(v)$ when $v \in \pi_1(V_1) \cap \pi_2(V_2)$. Now, Properties (i) and (ii) of Definition 14 for ψ follow, so ψ is an extensional isomorphism between \mathcal{A} and \mathcal{A}' , finishing the proof. \blacksquare

Proof of Lemma 3: Denote by $\mathcal{D}_\sigma = \mathbb{C} \setminus P$ the domain of holomorphy of σ . We proceed by induction on $\text{lv}(u)$. In the base case $\text{lv}(u) = 0$, i.e., $u = v_{in}$, the claim is trivially true with $E_u = \emptyset$. Now suppose that $\text{lv}(u) \geq 1$, and assume the statement holds for all $v \in V$ with $\text{lv}(v) < \text{lv}(u)$, i.e., $\mathcal{D}_{\langle v \rangle^\sigma} = \mathbb{C} \setminus E_v$, where E_v are closed countable subsets of $\mathbb{C} \setminus \mathbb{R}$. Set $E_u = \mathbb{C} \setminus \mathcal{D}_{\langle u \rangle^\sigma}$. We will show that E_u is a closed countable subset of $\mathbb{C} \setminus \mathbb{R}$. To this end, first note that $S := \bigcup_{v \in \text{par}(u)} E_v$ is a closed countable subset of $\mathbb{C} \setminus \mathbb{R}$, and thus $\mathbb{C} \setminus S$ is an open connected set containing \mathbb{R} . We claim that if z^* is a limit point of $E_u \setminus S$, then $z^* \in S$. Suppose otherwise, i.e., there exist a sequence $(z_n)_{n \in \mathbb{N}}$ of distinct elements of $E_u \setminus S$, and a point $z^* \in \mathbb{C} \setminus S$, such that $z_n \rightarrow z^*$. Define the function $f : \mathbb{C} \setminus S \rightarrow \mathbb{C}$, $f(z) = \sum_{v \in \text{par}(u)} \omega_{uv} \langle v \rangle^\sigma(z) + \theta_u$. As the functions $\langle v \rangle^\sigma$ are holomorphic on $\mathcal{D}_{\langle v \rangle^\sigma}$, they are, in particular, continuous, and so f is continuous. Therefore $f(z_n) \rightarrow f(z^*)$ as $n \rightarrow \infty$. As

$$z_n \in E_u \setminus S = \bigcap_{v \in \text{par}(u)} \mathcal{D}_{\langle v \rangle^\sigma} \setminus \mathcal{D}_{\langle u \rangle^\sigma},$$

it follows by definition of natural domain that $f(z_n) \in P$, for all $n \in \mathbb{N}$. Moreover, since P is discrete, we deduce that there exists a point $p^* \in P$ such that $f(z_n) = p^*$, for all sufficiently large $n \in \mathbb{N}$. Now, since $\mathbb{C} \setminus S$ is connected and f is holomorphic, it follows that $f(z) = p^*$, for all $z \in \mathbb{C} \setminus S$. But $0 \in \mathbb{R} \subset \mathbb{C} \setminus S$, which thus implies $p^* = f(0) = \sum_{v \in \text{par}(u)} \omega_{uv} \langle v \rangle^\sigma(0) + \theta_u \in \mathbb{R}$, contradicting $P \subset \mathbb{C} \setminus \mathbb{R}$. This completes the proof that any limit point of $E_u \setminus S$ is contained in S . Now define the sets $E_u^N := \{z \in E_u : |z| \leq N, d(z, S) \geq 1/N\}$, for $N \in \mathbb{N}$, where d denotes the Euclidean distance in \mathbb{C} . We see that E_u^N is finite, for each $N \in \mathbb{N}$, for otherwise there would exist a sequence $(z_n)_{n \in \mathbb{N}}$ of distinct elements of E_u^N converging to a point $z^* \in \mathbb{C}$. But then, by the claim above, we have $z^* \in S$, which contradicts $d(z_n, S) \geq 1/N$, for all $n \in \mathbb{N}$. We deduce that $E_u = S \cup \bigcup_{N \in \mathbb{N}} E_u^N$ is a closed countable set, and therefore $\mathcal{D}_{\langle u \rangle^\sigma} = \mathbb{C} \setminus E_u$ is an open connected set. To see that $\mathcal{D}_{\langle u \rangle^\sigma} \supset \mathbb{R}$, note that, for $z \in \mathbb{R}$, we have $z \in \mathbb{C} \setminus S = \bigcap_{v \in \text{par}(u)} \mathcal{D}_{\langle v \rangle^\sigma}$, and $f(z) \in \mathbb{R} \subset \mathcal{D}_\sigma$, so $z \in \mathcal{D}_{\langle u \rangle^\sigma}$. ■

Proof of Lemma 4: Let \mathbf{a} , δ , and T be as in the statement of the lemma, such that $D_k^\circ(\mathbf{a}, \delta) \subset \mathcal{U}$ and $F|_T \equiv 0$. Then the function $F_{\mathbf{a}} := F(\cdot + \mathbf{a})$ is holomorphic on $\mathcal{U} - \mathbf{a}$, and $F_{\mathbf{a}}|_{T-\mathbf{a}} \equiv 0$. Thus, as $F|_{\mathcal{U}} \equiv 0$ if and only if $F_{\mathbf{a}}|_{\mathcal{U}-\mathbf{a}} \equiv 0$, it suffices to prove the result for $\mathbf{a} = \mathbf{0}$. Let $T_0 := T$, $T_k := D_k^\circ(\mathbf{0}, \delta)$, and, for $r = 1, \dots, k-1$, define the sets

$$T_r = \{(iz_1, \dots, iz_{k-r}, s_{k-r+1}, \dots, s_k) : z_j \in (-\delta, \delta), \forall j; s_j \in D_1^\circ(0, \delta), \forall j\}.$$

Note that $T_r \subset D_k^\circ(\mathbf{0}, \delta) \subset \mathcal{U}$, for $r \in \{0, \dots, k\}$. We establish by induction over r that $F|_{T_r} \equiv 0$, $r \in \{0, \dots, k\}$. The base case $F|_{T_0} \equiv 0$ holds by assumption. So suppose that $F|_{T_r} \equiv 0$, for some $r \in \{0, \dots, k-1\}$. If $0 \leq r < k-1$, fix arbitrary $z_j \in (-\delta, \delta)$, for $j \in \{1, \dots, k-r-1\}$. Similarly, if $0 < r \leq k-1$, fix arbitrary $s_j \in D_1^\circ(0, \delta)$, for $j \in \{k-r+1, \dots, k\}$. Consider the function $G : D_1^\circ(0, \delta) \rightarrow \mathbb{C}$ defined by

$$G(z) = \begin{cases} F(iz_1, \dots, iz_{k-1}, iz), & \text{if } r = 0 \\ F(iz_1, \dots, iz_{k-r-1}, iz, s_{k-r+1}, \dots, s_k), & \text{if } 1 \leq r < k-1 \\ F(iz, s_2, \dots, s_k), & \text{if } r = k-1 \end{cases}$$

Note that G is holomorphic, and $G|_{(-\delta, \delta)} \equiv 0$ by the induction hypothesis. Since the zero set of a nonzero holomorphic function in one variable does not have a limit point in the domain, we deduce that $G|_{D_1^\circ(0, \delta)} \equiv 0$. But z_j and s_j were arbitrary, so we have $F|_{T_{r+1}} \equiv 0$. We have thus shown that

F is identically zero on an open subset $T_k = D_k^\circ(\mathbf{0}, \delta)$ of its connected domain \mathcal{U} , and so, by the multivariate identity theorem [19, 1.2.12], it must be identically zero on \mathcal{U} . \blacksquare

Proof of Lemma 5: Let t^* , \mathbf{a} , δ , T , and \tilde{T} be as in the statement of the lemma, such that $D_k^\circ(\mathbf{a}, \delta) \subset \mathcal{U}$, $\tilde{T} \subset (\mathbb{C} \setminus \{t^*\}) \times \mathbb{C}^k$, $\text{cl}(\tilde{T}) \supset T$, and $F|_{\tilde{T}} \equiv 0$, and denote $\mathcal{V} := D_{1+k}^\circ(\mathbf{a}, \delta)$. The function $F_{(t^*, \mathbf{a})} = F(\cdot + (t^*, \mathbf{a}))$ is holomorphic on $\mathcal{U} - (t^*, \mathbf{a})$, and the sets

$$T_{(t^*, \mathbf{a})} := T - (t^*, \mathbf{a}) = \{(0, iz_1, \dots, iz_k) : z_j \in (-\delta, \delta), j = 1, \dots, k\}$$

and $\tilde{T}_{(t^*, \mathbf{a})} := \tilde{T} - (t^*, \mathbf{a})$ satisfy $\tilde{T}_{(t^*, \mathbf{a})} \subset (\mathbb{C} \setminus \{0\}) \times \mathbb{C}^k$, $\text{cl}(\tilde{T}_{(t^*, \mathbf{a})}) \supset T_{(t^*, \mathbf{a})}$, and $F_{(t^*, \mathbf{a})}|_{\tilde{T}_{(t^*, \mathbf{a})}} \equiv 0$. Therefore, as $F|_{\mathcal{U}} \equiv 0$ if and only if $F_{(t^*, \mathbf{a})}|_{\mathcal{U} - (t^*, \mathbf{a})} \equiv 0$, and (t^*, \mathbf{a}) was arbitrary, it suffices to prove the result for $(t^*, \mathbf{a}) = (0, \mathbf{0})$. Assume by way of contradiction that $F|_{\mathcal{V}}$ is not identically 0. Then, by inspection of the power series expansion of F in the open neighborhood \mathcal{V} of $(0, \mathbf{0})$, we obtain that there exists a maximal $p \in \mathbb{N}_0$ such that $z_0^{-p}F(z_0, z_1, \dots, z_k)$ is holomorphic in \mathcal{V} . Write $G(z_0, z_1, \dots, z_k) = z_0^{-p}F(z_0, z_1, \dots, z_k)$, with $G : \mathcal{V} \rightarrow \mathbb{C}$ holomorphic and not identically 0. Now, due to $\tilde{T} \subset (\mathbb{C} \setminus \{0\}) \times \mathbb{C}^k$, we have $z_0 \neq 0$, for every $(z_0, z_1, \dots, z_k) \in \tilde{T}$. Moreover, as $F|_{\tilde{T}} \equiv 0$, we have $G(z_0, z_1, \dots, z_k) = z_0^{-p} \cdot 0 = 0$, for all $(z_0, z_1, \dots, z_k) \in \tilde{T}$. Now, since G is continuous and $\text{cl}(\tilde{T}) \supset T$ by assumption, it follows that $G(0, z_1, \dots, z_k) = 0$, for all $(0, z_1, \dots, z_k) \in T$. The mapping $(z_1, \dots, z_k) \mapsto G(0, z_1, \dots, z_k)$ is holomorphic on $D_k^\circ(\mathbf{0}, \delta)$ and identically zero on the set

$$\{(iz_1, \dots, iz_k) : z_j \in (-\delta, \delta), j = 1, \dots, k\},$$

and so, by Lemma 4, we obtain $G(0, z_1, \dots, z_k) = 0$, for all $(0, z_1, \dots, z_k) \in \mathcal{V}$. By inspection of the power series expansion of G in \mathcal{V} , we find that G must have the form $G(z_0, z_1, \dots, z_k) = z_0 \frac{\partial G}{\partial z_0}(z_0, z_1, \dots, z_k)$. As the function $\frac{\partial G}{\partial z_0}$ is holomorphic in \mathcal{V} , we have that $z_0^{-(p+1)}F(z_0, \dots, z_k) = \frac{\partial G}{\partial z_0}(z_0, \dots, z_k)$ is holomorphic in \mathcal{V} , contradicting the maximality of p . Our hypothesis that $F|_{\mathcal{V}}$ is not identically zero must hence be false, i.e., we have $F|_{\mathcal{V}} \equiv 0$. Finally, by the multivariate identity theorem [19, 1.2.12], we deduce that $F|_{\mathcal{U}} \equiv 0$. \blacksquare

Proof of Lemma 6: First note that M is the closure of a one-parameter subgroup of $T^d = \mathbb{R}^d / \mathbb{Z}^d$. Since T^d is compact and abelian, so is M . Moreover, M is connected (as the closure of a connected set), and so, by [20, Theorem 11.2], it is itself isomorphic to a torus. It remains to determine its dimension. A character on a compact abelian group G is a continuous group homomorphism $\chi : G \rightarrow S^1$, where $S^1 = \{z \in \mathbb{C} : |z| = 1\}$ is the multiplicative circle group, and we denote by \widehat{G} the set of all characters on G . We claim that

$$M = \bigcap_{\substack{\chi \in \widehat{T^d} \\ M \subset \ker(\chi)}} \ker(\chi). \quad (48)$$

The inclusion of M in the right-hand side is clear, so we only need to show the reverse inclusion. Note that, since M is closed, T^d/M is a Lie group. We will rewrite the right-hand side of (48) by establishing a bijective correspondence between the characters $\chi : T^d \rightarrow S^1$ such that $M \subset \ker(\chi)$, and the characters $f : T^d/M \rightarrow S^1$. To this end, let $\pi : T^d \rightarrow T^d/M$ be the projection map, and suppose that $\chi : T^d \rightarrow S^1$ is a character such that $M \subset \ker(\chi)$. Then χ factors according to $\chi = f \circ \pi$, for some continuous homomorphism $f : T^d/M \rightarrow S^1$, in other words, f is a character on T^d/M . Conversely, for any such f we have that $f \circ \pi$ is a character χ on T^d with $M \subset \ker(\chi)$.

Therefore it suffices to show that

$$\bigcap_{f \in \widehat{T^d/M}} \ker(f) = \{0\}. \quad (49)$$

Indeed, if this is the case, then

$$M = \pi^{-1}(\{0\}) = \bigcap_{f \in \widehat{T^d/M}} \pi^{-1}(\ker(f)) \supset \bigcap_{f \in \widehat{T^d/M}} \ker(f \circ \pi) = \bigcap_{\substack{\chi \in \widehat{T^d} \\ M \subset \ker(\chi)}} \ker(\chi),$$

as desired. We thus proceed to establishing (49). First note that, as T^d is compact, connected, and abelian, then so is T^d/M , and thus by [20, Theorem 11.2] we have that T^d/M is isomorphic (as a Lie group) to the torus T^r of some dimension $r \geq 0$. Now suppose that $(u_1, u_2, \dots, u_r) \in T^r$ is such that $f(u_1, u_2, \dots, u_r) = 1$, for all characters $f : T^r \rightarrow S^1$. Our goal is to show that $u_j = 0 \pmod{\mathbb{Z}}$, for all $j = 1, \dots, r$. For a given $j \in \{1, \dots, r\}$ let $f_j(t_1, t_2, \dots, t_r) = e^{2\pi i t_j}$. Since $f_j : T^r \rightarrow S^1$ is a character, we have $1 = f_j(u_1, \dots, u_r) = e^{2\pi i u_j}$, and thus $u_j = 0 \pmod{\mathbb{Z}}$. Since this holds for all j , we have (49), and therefore also (48). Note that any character on T^d has the form

$$\chi_{\mathbf{m}}(t_1, t_2, \dots, t_d) = e^{2\pi i(m_1 t_1 + m_2 t_2 + \dots + m_d t_d)}, \quad \text{for } (t_1, \dots, t_d) \in T^d, \quad (50)$$

where $\mathbf{m} = (m_1, m_2, \dots, m_d) \in \mathbb{Z}^d$ (this is easily seen for $d = 1$, and follows by induction for other values of d). Now, for any character $\chi_{\mathbf{m}} : T^d \rightarrow S^1$ such that $M \subset \ker(\chi_{\mathbf{m}})$, we have

$$1 = \chi_{\mathbf{m}}(\alpha_1 t, \alpha_2 t, \dots, \alpha_d t) = e^{2\pi i(m_1 \alpha_1 + m_2 \alpha_2 + \dots + m_d \alpha_d)t}, \quad \text{for all } t \in \mathbb{R},$$

by definition of M , which is equivalent to

$$m_1 \alpha_1 + m_2 \alpha_2 + \dots + m_d \alpha_d = 0.$$

It follows immediately that $Z = \{\mathbf{m} \in \mathbb{Z}^d : \chi_{\mathbf{m}} \in \widehat{T^d}, M \subset \ker(\chi)\}$ is a free abelian group of dimension $r = n - k$, where $k = \dim\langle \alpha_1, \dots, \alpha_d \rangle_{\mathbb{Q}}$. We can thus pick a basis $\{\mathbf{m}^1, \dots, \mathbf{m}^r\}$ for Z , and then, for any character $\chi_{\mathbf{m}}$ with $\mathbf{m} \in Z$, we have $\chi_{\mathbf{m}} = \chi_{\mathbf{m}^1}^{n_1} \dots \chi_{\mathbf{m}^r}^{n_r}$, for some $n_1, \dots, n_r \in \mathbb{Z}$. Therefore M is the kernel of the continuous surjective homomorphism $\Phi : T^n \rightarrow S^r$ given by $\Phi = (\chi_{\mathbf{m}^1}, \dots, \chi_{\mathbf{m}^r})$, and hence its dimension is $n - r = k$, as desired. \blacksquare

Proof of Lemma 7: Define the following subsets of T^d :

$$M = \{(\alpha_1 t, \alpha_2 t, \dots, \alpha_d t) + \mathbb{Z}^d : t \in \mathbb{R}\},$$

$$M_R = \{(\alpha_1 t, \alpha_2 t, \dots, \alpha_d t) + \mathbb{Z}^d : t \in \mathbb{R} \setminus [-R, R]\}, \text{ for } R > 0, \text{ and}$$

$$M' = \{Q \cdot (u_1, \dots, u_k) + \mathbb{Z}^d : u_1, \dots, u_k \in \mathbb{R}\},$$

as well as the map $\Phi : \mathbb{R}^k \rightarrow T^d$

$$\begin{aligned} \Phi(u_1, \dots, u_k) &= Q \cdot (u_1, \dots, u_k) + \mathbb{Z}^d \\ &= \left(u_1, \dots, u_k, \sum_{j=1}^k q_{k+1,j} u_j, \dots, \sum_{j=1}^k q_{d,j} u_j \right) + \mathbb{Z}^d. \end{aligned}$$

Let $K = \ker \Phi$, and note that M' is the image of Φ . Further, note that K is an abelian group, and a subgroup of \mathbb{Z}^k . For $j = 1, \dots, k$, let $N_j \in \mathbb{Z}$ be such that $q_{pj} N_j \in \mathbb{Z}$, for all $p = 1, \dots, d$. Let $\mathbf{e}_j \in \mathbb{R}^k$ be the vector with N_j in the j -th entry, and 0 in all the other entries. Then $\Phi(\mathbf{e}_j) = \mathbf{0} + \mathbb{Z}^d$,

for all $j = 1, \dots, k$, so $E := \{e_1, \dots, e_k\} \subset K$. Moreover, E is a basis for \mathbb{R}^k , so K is a lattice of rank k . Therefore M' and \mathbb{R}^k/K are isomorphic as groups via the induced map

$$\tilde{\Phi} : \mathbb{R}^k/K \rightarrow M', \quad \mathbf{u} + K \mapsto Q \cdot \mathbf{u}.$$

Since $\tilde{\Phi}$ is a continuous bijection, \mathbb{R}^k/K is compact, and T^d is Hausdorff, it follows that the map $\tilde{\Phi}$ is, in fact, a Lie group isomorphism (when M' is equipped with the subspace topology inherited from T^d). In particular, M' is a torus of dimension k . Let $\{\mathbf{b}_1, \dots, \mathbf{b}_k\}$ be a basis for K , and let

$$B = \left\{ c_1 \mathbf{b}_1 + \dots + c_k \mathbf{b}_k : c_1, \dots, c_k \in \left[-\frac{1}{2}, \frac{1}{2} \right) \right\} \subset \mathbb{R}^k$$

be a fundamental domain of the lattice K . Then, for any $\mathbf{u} \in \mathbb{R}^k$ we can write $\mathbf{u} = \mathbf{b} + \mathbf{k}$ with $\mathbf{b} \in B$ and $\mathbf{k} \in K$. We will prove the lemma with

$$C = \{(u_1/\alpha_1, \dots, u_k/\alpha_k) : (u_1, \dots, u_k) \in \text{int}(B)\},$$

where $\text{int}(B)$ denotes the interior of B . Note that C is open and $0 \in C$. For $t \in \mathbb{R}$ we have

$$\begin{aligned} (\alpha_1 t, \alpha_2 t, \dots, \alpha_d t) + \mathbb{Z}^d &= \left(\alpha_1 t, \dots, \alpha_k t, \sum_{j=1}^k q_{k+1,j} \alpha_j t, \dots, \sum_{j=1}^k q_{d,j} \alpha_j t \right) + \mathbb{Z}^d \\ &= Q \cdot (\alpha_1 t, \alpha_2 t, \dots, \alpha_k t) + \mathbb{Z}^d \in M', \end{aligned} \quad (51)$$

and so $M \subset M'$. Moreover, by Lemma 6 we have that $\text{cl}(M)$ is a torus of dimension k , so we deduce $\text{cl}(M) = M'$. We next establish that $\text{cl}(M_R) = M'$, for every $R > 0$. To this end, we distinguish between the cases $k = 1$ and $k \geq 2$.

The case $k = 1$. Let $(\alpha_1 t, \alpha_2 t, \dots, \alpha_d t) + \mathbb{Z}^d$, $t \in \mathbb{R}$, be an arbitrary element of M . As $\dim_{\mathbb{Q}} \langle \alpha_1, \dots, \alpha_d \rangle_{\mathbb{Q}} = k = 1$, there exist $a \in \mathbb{R} \setminus \{0\}$ and $m_1, \dots, m_d \in \mathbb{Z}$ such that $(\alpha_1, \alpha_2, \dots, \alpha_d) = (a m_1, a m_2, \dots, a m_d)$. Now let $n \in \mathbb{Z}$ be an integer such that $t + n/a \notin [-R, R]$. Then

$$\begin{aligned} (\alpha_1 t, \alpha_2 t, \dots, \alpha_d t) + \mathbb{Z}^d &= (\alpha_1 t, \alpha_2 t, \dots, \alpha_d t) + (n m_1, n m_2, \dots, n m_d) + \mathbb{Z}^d \\ &= \left(\alpha_1 \left(t + \frac{n}{a} \right), \alpha_2 \left(t + \frac{n}{a} \right), \dots, \alpha_d \left(t + \frac{n}{a} \right) \right) + \mathbb{Z}^d \in M_R. \end{aligned}$$

Therefore $M_R = M$, and so $\text{cl}(M_R) = \text{cl}(M) = M'$.

The case $k \geq 2$. First note that

$$L_R := M \setminus M_R = \{(\alpha_1 t, \alpha_2 t, \dots, \alpha_d t) + \mathbb{Z}^d : t \in [-R, R]\}$$

is the image of $[-R, R] \subset \mathbb{R}$ under a continuous bijective map from \mathbb{R} to T^d . Since $[-R, R] \subset \mathbb{R}$ is compact and T^d is Hausdorff, it follows by [21, Cor. 15.1.7] that L_R is homeomorphic to $[-R, R]$. In particular, L_R is a 1-dimensional submanifold of M with boundary. Now, by general properties of the closure, we have $\text{cl}(M_R) = \text{cl}(M \setminus L_R) \supset \text{cl}(M) \setminus \text{cl}(L_R) = M' \setminus L_R$. Therefore, as M' has dimension $k > 1$ and L_R has dimension 1, we have $\text{cl}(M_R) = \text{cl}(\text{cl}(M_R)) \supset \text{cl}(M' \setminus L_R) = M'$. On the other hand, $\text{cl}(M_R) \subset \text{cl}(M) = M'$, and thus $\text{cl}(M_R) = M'$, as desired. Now fix some $\mathbf{s} = (u_1/\alpha_1, \dots, u_k/\alpha_k) \in C$, where $\mathbf{u} = (u_1, \dots, u_k) \in \text{int}(B)$. Since M_R is dense in M' , for every $R > 0$, there exists a sequence $(t^{n,\mathbf{s}})_{n \in \mathbb{N}}$ in \mathbb{R} with $|t^{n,\mathbf{s}}| \rightarrow \infty$ such that

$$(\alpha_1 t^{n,\mathbf{s}}, \alpha_2 t^{n,\mathbf{s}}, \dots, \alpha_d t^{n,\mathbf{s}}) + \mathbb{Z}^d \rightarrow Q \cdot \mathbf{u} + \mathbb{Z}^d. \quad (52)$$

As $M \subset M'$, there exists a sequence $(\tilde{\mathbf{u}}^{n,s})_{n \in \mathbb{N}}$ such that

$$(\alpha_1 t^{n,s}, \alpha_2 t^{n,s}, \dots, \alpha_d t^{n,s}) + \mathbb{Z}^d = Q \cdot \tilde{\mathbf{u}}^{n,s} + \mathbb{Z}^d, \quad (53)$$

for all $n \in \mathbb{N}$. With this, (52) reads

$$Q \cdot \tilde{\mathbf{u}}^{n,s} + \mathbb{Z}^d \rightarrow Q \cdot \mathbf{u} + \mathbb{Z}^d,$$

and after applying the isomorphism $\tilde{\Phi}^{-1}$, we obtain $\tilde{\mathbf{u}}^{n,s} + K \rightarrow \mathbf{u} + K$ as $n \rightarrow \infty$. Now, for each $n \in \mathbb{N}$, let $\mathbf{u}^{n,s} = (u_1^{n,s}, \dots, u_k^{n,s}) \in B$ be such that $\mathbf{u}^{n,s} - \tilde{\mathbf{u}}^{n,s} \in K$. Then we have $\mathbf{u}^{n,s} + K \rightarrow \mathbf{u} + K$ as $n \rightarrow \infty$. Since $\mathbf{u} \in \text{int}(B)$, there exists an $n_0 \in \mathbb{N}$ such that $\mathbf{u}^{n,s} \in \text{int}(B)$, for $n \geq n_0$. By discarding the first n_0 terms of the sequences $(t^{n,s})_{n \in \mathbb{N}}$ and $(\tilde{\mathbf{u}}^{n,s})_{n \in \mathbb{N}}$, we may assume w.l.o.g. that $n_0 = 0$. It follows that $\mathbf{u}^{n,s} \rightarrow \mathbf{u}$ as $n \rightarrow \infty$. Now define $\mathbf{r}^{n,s} = (u_1^{n,s}/\alpha_1, \dots, u_k^{n,s}/\alpha_k)$. We then have $\mathbf{r}^{n,s} \in C$, $\mathbf{r}^{n,s} \rightarrow \mathbf{s}$, and (53) yields

$$(\alpha_1 t^{n,s}, \alpha_2 t^{n,s}, \dots, \alpha_d t^{n,s}) + \mathbb{Z}^d = \Phi(\tilde{\mathbf{u}}^{n,s}) = \Phi(\mathbf{u}^{n,s}) = Q \cdot (\alpha_1 r_1^{n,s}, \dots, \alpha_k r_k^{n,s}) + \mathbb{Z}^d,$$

as desired. ■