



HAL
open science

Geometry-induced Implicit Regularization in Deep ReLU Neural Networks

Joachim Bona-Pellissier, François Malgouyres, François Bachoc

► **To cite this version:**

Joachim Bona-Pellissier, François Malgouyres, François Bachoc. Geometry-induced Implicit Regularization in Deep ReLU Neural Networks. 2024. hal-04452356

HAL Id: hal-04452356

<https://hal.science/hal-04452356v1>

Preprint submitted on 12 Feb 2024

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

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

Geometry-induced Implicit Regularization in Deep ReLU Neural Networks

Joachim Bona-Pellissier ^{*1}, François Malgouyres ^{†1}, and François Bachoc ^{‡1}

¹Institut de Mathématiques de Toulouse ; UMR 5219, Université de Toulouse ; CNRS
, UPS, F-31062 Toulouse Cedex 9, France

February 12, 2024

Abstract

It is well known that neural networks with many more parameters than training examples do not overfit. Implicit regularization phenomena, which are still not well understood, occur during optimization and ‘good’ networks are favored. Thus the number of parameters is not an adequate measure of complexity if we do not consider all possible networks but only the ‘good’ ones.

To better understand which networks are favored during optimization, we study the geometry of the output set as parameters vary. When the inputs are fixed, we prove that the dimension of this set changes and that the local dimension, called *batch functional dimension*, is almost surely determined by the activation patterns in the hidden layers. We prove that the batch functional dimension is invariant to the symmetries of the network parameterization: neuron permutations and positive rescalings. Empirically, we establish that the batch functional dimension decreases during optimization. As a consequence, optimization leads to parameters with low batch functional dimensions. We call this phenomenon *geometry-induced implicit regularization*.

The batch functional dimension depends on both the network parameters and inputs. To understand the impact of the inputs, we study, for fixed parameters, the largest attainable batch functional dimension when the inputs vary. We prove that this quantity, called *computable full functional dimension*, is also invariant to the symmetries of the network’s parameterization, and is determined by the achievable activation patterns. We also provide a sampling theorem, showing a fast convergence of the estimation of the computable full functional dimension for a random input of increasing size. Empirically we find that the computable full functional dimension remains close to the number of parameters, which is related to the notion of local identifiability. This differs from the observed values for the batch functional dimension computed on training inputs and test inputs. The latter are influenced by geometry-induced implicit regularization.

*joachim.bona-pellissier@univ-toulouse.fr

†francois.malgouyres@math.univ-toulouse.fr

‡francois.bachoc@math.univ-toulouse.fr

Keywords— Deep learning, implicit regularization, geometry of neural networks, functional dimension of neural networks, flat minima, identifiability.

1 Introduction

We introduce the context of the present article in Section 1.1, and we give a first glimpse of the objects of study in Section 1.2. We then outline the article in Section 1.3, and we present the related works in Section 1.4.

1.1 On the Importance of Local Complexity Measures for Neural Networks

Learning deep neural networks has a huge impact on many practical aspects of our lives. This requires optimizing a non-convex function, in a large dimensional space. Surprisingly, in many cases, although the number of parameters defining the neural network exceeds by far the amount of training data, the learned neural network generalizes and performs well with unseen data (Zhang u. a., 2021). This is surprising because in this setting the set of global minimizers is large (Cooper, 2021; Li u. a., 2018) and contains elements that generalize poorly (Wu u. a., 2017; Neyshabur u. a., 2017). In accordance with this empirical observation, the good generalization behavior is not explained by the classical statistical learning theory (e.g., Anthony und Bartlett, 2009; Grohs und Kutyniok, 2022) that considers the worst possible parameters in the parameter set. For instance, the Vapnik-Chervonenkis dimension of feedforward neural networks of depth L , with W parameters, with the ReLU activation function is¹ $\widetilde{O}(LW)$ (Bartlett u. a., 2019, 1998; Harvey u. a., 2017; Maass, 1994), leading to an upper bound on the generalization gap of order¹ $\widetilde{O}(\sqrt{\frac{LW}{n}})$, where n is the sample size. This worst-case analysis is not accurate enough to explain the success of deep learning, when $W \gg n$.

This leads to the conclusion that a global analysis, that applies to all global minima and the worst possible neural network that fits the data, will not permit to explain the success of deep learning. A local analysis is needed.

Despite tremendous research efforts in this direction (see, e.g., Grohs und Kutyniok, 2022 and references below) a complete explanation for the good generalization behavior in deep learning is still lacking. The attempts of explanation suggest that optimization algorithms and notably stochastic algorithms discover ‘good minima’. These are minima having special properties that authors would like to model using local complexity measures that are pivotal in the mathematical explanation. Authors aim to establish that stochastic algorithms prioritize outputs (parameterizations at convergence) with low local complexity and to demonstrate that low local complexity explains the good generalization to unseen data (Bartlett u. a., 2020; Chaudhari u. a., 2019; Camuto u. a., 2021; Keskar u. a., 2017). This is sometimes also expressed as some form of implicit regularization (Imaizumi und Schmidt-Hieber, 2023; Belkin, 2021; Neyshabur u. a., 2017).

¹The notation $\widetilde{O}(\cdot)$ ignores logarithmic factors.

In this spirit, many authors contend that the excellent generalization behavior can be attributed to the fulfillment of conditions regarding the flatness of the landscape in the proximity of the algorithm’s output (Keskar u. a., 2017; Foret u. a., 2021; Cha u. a., 2021; Hochreiter und Schmidhuber, 1997). This is known however not to fully capture the good generalization phenomenon (Dinh u. a., 2017). Other studies explain the good generalization performances by constraints involving norms of the neural network parameters (Bartlett u. a., 2020; Neyshabur u. a., 2015b; Golowich u. a., 2018; Bartlett u. a., 2017). Despite being supported by partial arguments, none of the aforementioned local complexity measures fully explain the experimentally observed behaviors.

This is in sharp contrast with linear networks for which implicit regularization is better understood. The consensus is that implicit regularization constrains the rank of the prediction matrix, the matrix obtained when multiplying all the factors of the linear network (Arora u. a., 2019; Razin und Cohen, 2020; Saxe u. a., 2019; Gidel u. a., 2019; Gissin u. a., 2019; Achour u. a., 2022).

1.2 Local Dimensions of the Image and Pre-image Sets

This article investigates properties and computational aspects of local geometrical complexity measures of deep ReLU neural networks, recently introduced by Grigsby u. a. (2022). The considered complexity measures relate to the local geometry of the *image set* as defined by $\{f_\theta(X) \mid \theta \text{ varies}\}$ and the *pre-image set* $\{\theta' \mid f_{\theta'}(X) = f_\theta(X)\}$, where $f_\theta(X)$ is the prediction made by the neural network of parameter θ , for an input sample $X = (x^{(i)})_{i \in \llbracket 1, n \rrbracket} \in \mathbb{R}^{N_0 \times n}$, where $x^{(i)}$ is the i -th column of X and the i -th input of the sample. When the differential $\mathbf{D}f_\theta(X)$ of $\theta \mapsto f_\theta(X)$ is appropriately defined, these concepts of complexity are associated with the local dimension of these sets, see Corollary 3, and related to the rank of the aforementioned differential, denoted $\text{rank}(\mathbf{D}f_\theta(X))$ and called *batch functional dimension* by Grigsby u. a. (2022). Notice that, before Grigsby u. a., the batch functional dimension already appeared in an identifiability condition introduced by Bona-Pellissier u. a. (2022).

1.3 Main Contributions and Organization of the Paper

- In Theorem 1 (Section 3), up to a negligible set, we decompose the parameter space as a finite union of open sets. On each set, the *batch functional dimension*

$$\text{rank}(\mathbf{D}f_\theta(X))$$

is well defined and constant. The construction of the sets shows that almost everywhere, the activation pattern (defined in Section 2) determines the batch functional dimension. We also establish in Proposition 2 (Section 3) that the batch functional dimension is invariant under the symmetries of a ReLU neural network’s parameterization, positive rescaling and neuron permutation, as defined in Section 2. We also provide examples in Sections 3 and 4.

- In Section 4, we illustrate the consequences of the statements of Section 3 when learning a deep ReLU network. In particular, we explain the links between the batch functional dimension, and the local dimensions of the image and pre-image sets, see Corollary 3 and Figure 2. We also illustrate in this figure how the described geometry impacts the iterates trajectory for small learning rates and we describe the geometry-induced implicit regularization.
- In Section 5, we study the *computable full functional dimension*¹ defined by

$$r^*(\theta) = \max_X \text{rank}(\mathbf{D}f_\theta(X)).$$

The first result of the section states that the *achievable activation patterns* for θ determine $r^*(\theta)$, see Theorem 6. It also shows that when more activation patterns can be achieved, r^* increases. As for the batch functional dimension, we establish that the computable full functional dimension is invariant under positive rescalings and neuron permutations. We finish the section with a connection between the computable full functional dimension and the fat-shattering dimension of neural networks.

- In Section 6 we provide the details on the practical computation of $\text{rank}(\mathbf{D}f_\theta(X))$, for given X and θ . We also establish in Theorem 11 that, for a given θ , a random X of sufficient size can be used to compute $r^*(\theta)$. Indeed, we upper bound the probability of not reaching $r^*(\theta)$, as a vanishing function of the number n of columns of X . The upper bound depends on two natural quantities p and $n^*(\theta)$ (see Theorem 11 for details).
- Finally, we provide experiments on the MNIST data set in Section 7. In Section 7.2, we analyze the behavior of the local complexity measures when the width of the network increases. We also describe their behavior during the learning phase in Section 7.3. We also show in Sections 7.4 and 7.5 how they behave when the distribution of (X, Y) is artificially complexified, Y denoting the outputs. The experiments highlight the geometry-induced implicit regularization described in Section 4 both at the learning and test stages. The experiments also highlight the correlation between the batch functional dimension computed using the learning and test samples and the complexity of the learning task. Our experiments also indicate that for corrupted or highly random inputs, the batch functional dimension may be maximal, corresponding almost surely to local identifiability, see Bona-Pellissier u. a. (2022).

All the proofs are in the Appendices and the codes are available at (Bona-Pellissier u. a., 2023b).

¹As its name indicates, the *computable full functional dimension* is a variant of the *full functional dimension* defined in (Grigsby u. a., 2022), that we can compute. The informal definition below does not take into account restrictions guaranteeing that $\theta \mapsto f_\theta(X)$ is differentiable at θ .

1.4 Related Works

To the best of our knowledge, the functional dimensions of deep ReLU neural networks has only been explicitly studied by Grigsby u. a. (2022, 2023). The article Grigsby u. a. (2022) is very rich and it is difficult to summarize it in a few lines². The authors establish sufficient conditions guaranteeing that $\theta \mapsto f_\theta(X)$ is differentiable. The conditions are comparable to but weaker than the one presented here. The benefit of the difference is that our conditions guarantee the value of the batch functional dimension, allowing us to make the connection between the activation patterns and the batch functional dimension. Furthermore, Grigsby u. a. define and provide examples to illustrate that the batch functional dimension and the full functional dimension vary in the parameter space. They also prove that for all narrowing architectures³, the *functional dimension* as defined by $\max_\theta \max_X \text{rank}(\mathbf{D}f_\theta(X))$ reaches its upper-bound $W - W'$ where W' is the number of positive rescalings. They finish their article with several examples showing that the global structure of the *pre-image set* $\{\theta' \mid f_{\theta'}(X) = f_\theta(X)\}$ can vary in several regards. Grigsby u. a. prove that when the input size lower-bounds the other widths there exist parameters for which the batch functional dimension reaches the upper-bound $W - W'$. They also numerically estimate, for several neural network architectures, the size of the sets of parameters that reach this upper bound.

Geometric properties of the pre-image set of a global minimizer have been studied by Cooper (2021). Topological properties of a variant of the image set included in function spaces, $\{f_\theta \mid \theta \text{ varies}\}$, have been established by Petersen u. a. (2021).

There are many articles devoted to the identifiability of neural networks (Petzka u. a., 2020; Carlini u. a., 2020; Rolnick und Kording, 2020; Stock und Gribonval, 2022; Bona-Pellissier u. a., 2022, 2023a). For a given θ , they study conditions guaranteeing that the pre-image set⁴ of $f_\theta(X)$ coincides with the set obtained when considering all the positive rescalings of θ . Of particular interest in our context, Bona-Pellissier u. a. (2022) establish that the condition $\text{rank}(\mathbf{D}f_\theta(X)) = W - W'$ is sufficient to guarantee local identifiability. The same condition is also involved in a necessary condition of local identifiability.

Other local complexity measures, not related to the geometry of neural networks, have been considered. There are complexity measures using the number of achievable activation patterns Montufar u. a. (2014); Raghu u. a. (2017); Hanin und Rolnick (2019). Those based on norms and the flatness are already mentioned in Section 1.1.

The objects studied in this article are also related to the properties of the landscape of the empirical risk, which have been investigated in the literature. Studies of these properties for instance permit to guarantee that first-order algorithms find a global minimizer (Soudry und Carmon, 2016; Nguyen und Hein, 2017; Safran u. a., 2021; Du u. a., 2019), focus on the shape at the bottom of the empirical risk (Ghorbani u. a., 2019; Sagun u. a., 2016; Gur-Ari u. a., 2018) and (again) on flatness.

The local properties studied in the present article also have an impact on the iterates

²A weakness of it is that it considers neural networks whose last layer undergoes a ReLU activation.

³Narrowing architectures are such that widths decrease.

⁴In these articles X sometimes contains infinitely many examples, in which case we let $f_\theta(X)$ denote the function f_θ restricted to X .

trajectory of minimization algorithms and therefore the biases induced by the optimization as studied in (Bartlett u. a., 2020; Chaudhari u. a., 2019; Camuto u. a., 2021; Keskar u. a., 2017).

Finally, Arora u. a. (2018) and Suzuki u. a. (2020) establish generalization bounds of compressed neural networks. This might provide hints for the construction of upper-bounds of the generalization gap based on the local geometric complexity measures considered in this article.

2 ReLU Networks and Notations

This section is devoted to introducing the formalism and notations that we use throughout the article. In Section 2.1, we present the graph formalism that we use for neural networks, and we specify the architectures that we study, and in Section 2.2, we construct the prediction function implemented by a network, and we define the differential $\mathbf{D}f_\theta(X)$ that is central in this work. In Section 2.3, we recall the two classical symmetries of ReLU networks, namely positive rescalings and permutations. Finally, we introduce the activation patterns in Section 2.4 and some additional notations in Section 2.5.

2.1 ReLU Network Architecture

Let us introduce our notations for deep fully-connected ReLU neural networks. In this paper, a network is a graph (E, V) of the following form.

- V is a set of neurons, which is divided into $L + 1$ layers, with $L \geq 2$: $V = \bigcup_{\ell=0}^L V_\ell$. The layer V_0 is the input layer, V_L is the output layer and the layers V_ℓ with $1 \leq \ell \leq L - 1$ are the hidden layers. Using the notation $|C|$ for the cardinality of a finite set C , we denote, for all⁵ $\ell \in \llbracket 0, L \rrbracket$, $N_\ell = |V_\ell|$ the size of the layer V_ℓ .
- E is the set of all oriented edges $v' \rightarrow v$ between neurons in consecutive layers, that is

$$E = \{v' \rightarrow v \mid v' \in V_{\ell-1}, v \in V_\ell, \text{ for } \ell \in \llbracket 1, L \rrbracket\}.$$

A network is parameterized by weights and biases, gathered in its parameterization θ , with

$$\theta = ((w_{v' \rightarrow v})_{v' \rightarrow v \in E}, (b_v)_{v \in B}) \in \mathbb{R}^E \times \mathbb{R}^B,$$

where $B = \bigcup_{\ell=1}^L V_\ell$. We let $W = |E| + |B|$.

The activation function used in the hidden layers, and denoted σ , is always ReLU: for any $p \in \mathbb{N}^*$ and any vector $x = (x_1, \dots, x_p)^T \in \mathbb{R}^p$, we set $\sigma(x) = (\max(x_1, 0), \dots, \max(x_p, 0))^T$. Here and in the sequel, the symbol \mathbb{N}^* denotes the set of natural numbers without 0. We allow the use of a specific activation $\sigma_L : \mathbb{R}^{N_L} \rightarrow \mathbb{R}^{N_L}$ for the output layer, which we only require to be analytic. For instance, σ_L can be the identity, as is generally the case in regression, or the softmax, as is generally the case in classification. The ReLU neural network architectures considered in this article are fully characterized by a triplet (E, V, σ_L) .

⁵Throughout the paper, for $a, b \in \mathbb{N}$, $a \leq b$, $\llbracket a, b \rrbracket$ is the set of consecutive integers $\{a, a + 1, \dots, b\}$.

2.2 ReLU Network Prediction

For a given θ , we define recursively $f_\theta^\ell : \mathbb{R}^{V_0} \rightarrow \mathbb{R}^{V_\ell}$, for $\ell \in \llbracket 0, L \rrbracket$ and $x \in \mathbb{R}^{V_0}$, by

$$\begin{cases} (f_\theta^0(x))_v = x_v & \text{for } v \in V_0, \\ (f_\theta^\ell(x))_v = \sigma \left(\sum_{v' \in V_{\ell-1}} w_{v' \rightarrow v} (f_\theta^{\ell-1}(x))_{v'} + b_v \right) & \text{for } v \in V_\ell, \text{ when } \ell \in \llbracket 1, L-1 \rrbracket, \\ (y_\theta^L(x))_v = \sum_{v' \in V_{L-1}} w_{v' \rightarrow v} (f_\theta^{L-1}(x))_{v'} + b_v & \text{for } v \in V_L, \\ f_\theta^L(x) = \sigma_L(y_\theta^L(x)), \end{cases} \quad (1)$$

where the definition of $f_\theta^L(x)$ takes into account that $\sigma^L : \mathbb{R}^{V_L} \rightarrow \mathbb{R}^{V_L}$ may require the whole pre-activation output. This is for instance the case for the softmax activation function. We define the function $f_\theta : \mathbb{R}^{V_0} \rightarrow \mathbb{R}^{V_L}$ implemented by the network of parameter θ as $f_\theta = f_\theta^L$. We call it the prediction.

For all $n \in \mathbb{N}^*$, we concatenate a set of n inputs in a matrix $X = (x^{(i)})_{i \in \llbracket 1, n \rrbracket} \in \mathbb{R}^{N_0 \times n}$, where $x^{(i)}$ is the i -th column of X and the i -th input of the network. We also allow to write f_θ as operating on an input set X . In this case, we write $f_\theta : \mathbb{R}^{N_0 \times n} \rightarrow \mathbb{R}^{N_L \times n}$ and we define $f_\theta(X)$ as the matrix gathering the outputs $(f_\theta(x^{(i)}))_{i \in \llbracket 1, n \rrbracket}$.

Among other quantities, we study in this article the set

$$\{f_\theta(X) \mid \theta \in \mathbb{R}^E \times \mathbb{R}^B\},$$

for $X \in \mathbb{R}^{N_0 \times n}$ fixed, which we call an *image set*. When it is differentiable at θ , we denote by $\mathbf{D}f_\theta(X)$ the differential, at the point θ , of the mapping

$$\begin{aligned} \mathbb{R}^E \times \mathbb{R}^B &\longrightarrow \mathbb{R}^{N_L \times n} \\ \theta' &\longmapsto f_{\theta'}(X). \end{aligned}$$

We recall that the differential at θ is the linear map

$$\mathbf{D}f_\theta(X) : \mathbb{R}^E \times \mathbb{R}^B \longrightarrow \mathbb{R}^{N_L \times n} \quad (2)$$

such that, for $\theta' \in \mathbb{R}^E \times \mathbb{R}^B$ in a neighborhood of zero,

$$f_{\theta+\theta'}(X) = f_\theta(X) + \mathbf{D}f_\theta(X)(\theta') + o(\|\theta'\|). \quad (3)$$

2.3 Positive rescaling and neuron permutations symmetries

Consider two parameters $\theta, \tilde{\theta} \in \mathbb{R}^{E \times B}$, with $\tilde{\theta} = \left((\tilde{w}_{v' \rightarrow v})_{v' \rightarrow v \in E}, (\tilde{b}_v)_{v \in B} \right)$. We say that θ and $\tilde{\theta}$ are equivalent modulo positive rescaling, and we write $\theta \sim_s \tilde{\theta}$, when the following holds. There are $(\lambda_v)_{v \in V_0 \cup \dots \cup V_L} \in (0, \infty)^{V_0 \cup \dots \cup V_L}$ such that $\lambda_v = 1$ for $v \in V_0 \cup V_L$ and for $\ell \in \llbracket 1, L \rrbracket$, $v' \in V_{\ell-1}$, $v \in V_\ell$,

$$w_{v' \rightarrow v} = \frac{\lambda_v}{\lambda_{v'}} \tilde{w}_{v' \rightarrow v}, \quad (4)$$

$$b_v = \lambda_v \tilde{b}_v. \quad (5)$$

Then it is a well-known property of ReLU networks (Bona-Pellissier u. a., 2023a, 2022; Neyshabur u. a., 2015a; Stock, 2021; Stock und Gribonval, 2022; Yi u. a., 2019) that if $\theta \sim_s \tilde{\theta}$ then $f_\theta = f_{\tilde{\theta}}$, that is, positive rescalings are a symmetry of the network parameterization.

Another classic symmetry consists in swapping neurons, and their corresponding weights, within each hidden layer. If θ stands for the permuted weights, we denote the corresponding equivalence relation $\tilde{\theta} \sim_p \theta$. Again, when $\tilde{\theta} \sim_p \theta$, we have $f_{\tilde{\theta}} = f_\theta$.

We say that $\tilde{\theta} \sim \theta$ if there exists θ' such that $\tilde{\theta} \sim_p \theta'$ and $\theta' \sim_s \theta$. Again, if $\tilde{\theta} \sim \theta$, then $f_\theta = f_{\tilde{\theta}}$.

2.4 Activation Patterns

For any $\ell \in \llbracket 1, L-1 \rrbracket$, $v \in V_\ell$, $\theta \in \mathbb{R}^E \times \mathbb{R}^B$ and $x \in \mathbb{R}^{N_0}$, we define the activation indicator at neuron v by

$$a_v(x, \theta) = \begin{cases} 1 & \text{if } \sum_{v' \in V_{\ell-1}} w_{v' \rightarrow v} (f_\theta^{\ell-1}(x))_{v'} + b_v \geq 0 \\ 0 & \text{otherwise.} \end{cases}$$

Using (1), we have for the ReLU activation function σ , any $\ell \in \llbracket 1, L-1 \rrbracket$ and $v \in V_\ell$,

$$(f_\theta^\ell(x))_v = a_v(x, \theta) \left(\sum_{v' \in V_{\ell-1}} w_{v' \rightarrow v} (f_\theta^{\ell-1}(x))_{v'} + b_v \right). \quad (6)$$

We then define the *activation pattern* as the mapping

$$\begin{aligned} a : \mathbb{R}^{N_0} \times (\mathbb{R}^E \times \mathbb{R}^B) &\longrightarrow \{0, 1\}^{N_1 + \dots + N_{L-1}} \\ (x, \theta) &\longmapsto (a_v(x, \theta))_{v \in V_1 \cup \dots \cup V_{L-1}}. \end{aligned}$$

For $X \in \mathbb{R}^{N_0 \times n}$ as considered above, we let $a(X, \theta) \in \{0, 1\}^{(N_1 + \dots + N_{L-1}) \times n}$ be defined by, for $i \in \llbracket 1, n \rrbracket$ and $v \in V_1 \cup \dots \cup V_{L-1}$, $a_{v,i}(X, \theta) = a_v(x^{(i)}, \theta)$. By extension, we also call *activation patterns* the elements of $\{0, 1\}^{N_1 + \dots + N_{L-1}}$ or $\{0, 1\}^{(N_1 + \dots + N_{L-1}) \times n}$.

2.5 Further Notation

We use the notation $\text{rank}(\cdot)$ for the rank of linear maps and matrices. The determinant of a square matrix M is denoted $\det(M)$. If the matrix $M \in \mathbb{R}^{a \times b}$ for $a, b \in \mathbb{N}^*$, then for $i \in \llbracket 1, a \rrbracket$, we write $M_{i,\cdot}$ for the row i of M .

All considered vector spaces are finite dimensional and they are endowed with the standard Euclidean topology. For a subset $C \subset T$ of a topological space, we denote $\text{Int}(C)$ the topological interior of C , ∂C its boundary and $C^c = T \setminus C$ the complement of C (the ambient topological space T should always be clear from context). For all Euclidean space V , all element $x \in V$, and all real number $r \geq 0$, the open Euclidean ball of radius r centered at x is denoted by $B(x, r)$.

3 Rank Properties

In this section, we give the key technical theorem, namely Theorem 1, on which the remaining of the article relies. We then illustrate the theorem with examples showing the diversity of situations that might occur. The theorem is composed of two parts. In the first one, we study $(X, \theta) \mapsto f_\theta(X)$ over $\mathbb{R}^{N_0 \times n} \times (\mathbb{R}^E \times \mathbb{R}^B)$, and in the second one, we study $\theta \mapsto f_\theta(X)$ over $\mathbb{R}^E \times \mathbb{R}^B$, for X fixed. We must first introduce a few definitions.

For $n \in \mathbb{N}^*$, the function

$$\begin{aligned} \mathbb{R}^{N_0 \times n} \times (\mathbb{R}^E \times \mathbb{R}^B) &\longrightarrow \{0, 1\}^{(N_1 + \dots + N_{L-1}) \times n} \\ (X, \theta) &\longmapsto a(X, \theta) \end{aligned}$$

takes a finite set of values, that we write $\Delta_1^n, \dots, \Delta_q^n$. Let us write, for $j \in \llbracket 1, q \rrbracket$,

$$\tilde{\mathcal{O}}_j^n = \text{Int}\{(X, \theta) \in \mathbb{R}^{N_0 \times n} \times (\mathbb{R}^E \times \mathbb{R}^B) \mid a(X, \theta) = \Delta_j^n\}, \quad (7)$$

and let us only keep the non-empty $\tilde{\mathcal{O}}_j^n$. If $m_n \in \llbracket 1, q \rrbracket$ is the number of such non-empty sets, up to a re-ordering, we can assume that we keep $\tilde{\mathcal{O}}_1^n, \dots, \tilde{\mathcal{O}}_{m_n}^n$. As will be formally established in Lemma 13, (i), third item, for all $j \in \llbracket 1, m_n \rrbracket$, the function $\theta \mapsto f_\theta(X)$ is differentiable at θ when $(X, \theta) \in \tilde{\mathcal{O}}_j^n$. We can therefore define, for $n \in \mathbb{N}^*$ and $j \in \llbracket 1, m_n \rrbracket$,

$$r_j^n = \max_{(X, \theta) \in \tilde{\mathcal{O}}_j^n} \text{rank}(\mathbf{D}f_\theta(X)). \quad (8)$$

We then define the subset of $\tilde{\mathcal{O}}_j^n$ on which the rank is maximal. For $n \in \mathbb{N}^*$ and $j \in \llbracket 1, m_n \rrbracket$,

$$\mathcal{O}_j^n = \{(X, \theta) \in \tilde{\mathcal{O}}_j^n \mid \text{rank}(\mathbf{D}f_\theta(X)) = r_j^n\}. \quad (9)$$

Similarly, for $n \in \mathbb{N}^*$ and $X \in \mathbb{R}^{N_0 \times n}$, the function $\theta \mapsto a(X, \theta)$ takes a finite number of values $\Delta_1^X, \dots, \Delta_{q^X}^X$, and we define, for $j \in \llbracket 1, q^X \rrbracket$,

$$\tilde{\mathcal{U}}_j^X = \text{Int}\{\theta \in \mathbb{R}^E \times \mathbb{R}^B \mid a(X, \theta) = \Delta_j^X\}. \quad (10)$$

Similarly, we keep only the nonempty such sets, and if $p_X \in \llbracket 1, q^X \rrbracket$ is the number of such sets, we can assume up to a re-ordering that we keep $\tilde{\mathcal{U}}_1^X, \dots, \tilde{\mathcal{U}}_{p_X}^X$. Again, as we will establish in Lemma 13, (ii), third item, for all $j \in \llbracket 1, p_X \rrbracket$, the function $\theta \mapsto f_\theta(X)$ is differentiable at θ when $\theta \in \tilde{\mathcal{U}}_j^X$. We can therefore define, for $n \in \mathbb{N}^*$, $X \in \mathbb{R}^{N_0 \times n}$ and $j \in \llbracket 1, p_X \rrbracket$,

$$r_j^X = \max_{\theta \in \tilde{\mathcal{U}}_j^X} \text{rank}(\mathbf{D}f_\theta(X)). \quad (11)$$

We finally define the subset of $\tilde{\mathcal{U}}_j^X$ on which the rank is maximal. For $n \in \mathbb{N}^*$, $X \in \mathbb{R}^{N_0 \times n}$ and $j \in \llbracket 1, p_X \rrbracket$,

$$\mathcal{U}_j^X = \{\theta \in \tilde{\mathcal{U}}_j^X \mid \text{rank}(\mathbf{D}f_\theta(X)) = r_j^X\}. \quad (12)$$

The following theorem is composed of two parts, named (i) and (ii). In (i), we study $(X, \theta) \mapsto f_\theta(X)$ over $\mathbb{R}^{N_0 \times n} \times (\mathbb{R}^E \times \mathbb{R}^B)$, and we provide properties of the sets $\mathcal{O}_1^n, \dots, \mathcal{O}_{m_n}^n$. In (ii), we study $\theta \mapsto f_\theta(X)$ over $\mathbb{R}^E \times \mathbb{R}^B$, for X fixed, and we provide properties of the sets $\mathcal{U}_1^X, \dots, \mathcal{U}_{p_X}^X$. Note that for both parts (i) and (ii), Items 1, 2 and 3 hold trivially by definition, while Items 4, 5 and 6 require detailed proofs.

Theorem 1. *Consider any deep fully-connected ReLU network architecture (E, V, σ_L) .*

(i) *For all $n \in \mathbb{N}^*$, by definition,*

- *the sets $\mathcal{O}_1^n, \dots, \mathcal{O}_{m_n}^n$ are non-empty and disjoint,*
- *for all $j \in \llbracket 1, m_n \rrbracket$, the function $(X, \theta) \mapsto a(X, \theta)$ is constant on \mathcal{O}_j^n and takes m_n distinct values on $\cup_{j=1}^{m_n} \mathcal{O}_j^n$;*
- *for all $j \in \llbracket 1, m_n \rrbracket$, $(X, \theta) \mapsto \text{rank}(\mathbf{D}f_\theta(X))$ is constant on \mathcal{O}_j^n and equal to r_j^n .*

Furthermore,

- *the sets $\mathcal{O}_1^n, \dots, \mathcal{O}_{m_n}^n$ are open,*
- *$(\cup_{j=1}^{m_n} \mathcal{O}_j^n)^c$ is a closed set with Lebesgue measure zero;*
- *for all $j \in \llbracket 1, m_n \rrbracket$, $(X, \theta) \mapsto f_\theta(X)$ is an analytic function on \mathcal{O}_j^n .*

(ii) *For all $n \in \mathbb{N}^*$, for all $X \in \mathbb{R}^{N_0 \times n}$, by definition,*

- *the sets $\mathcal{U}_1^X, \dots, \mathcal{U}_{p_X}^X$ are non-empty and disjoint,*
- *for all $j \in \llbracket 1, p_X \rrbracket$, the function $\theta \mapsto a(X, \theta)$ is constant on each \mathcal{U}_j^X and takes p_X distinct values on $\cup_{j=1}^{p_X} \mathcal{U}_j^X$;*
- *for all $j \in \llbracket 1, p_X \rrbracket$, $\theta \mapsto \text{rank}(\mathbf{D}f_\theta(X))$ is constant on \mathcal{U}_j^X and equal to r_j^X .*

Furthermore,

- *the sets $\mathcal{U}_1^X, \dots, \mathcal{U}_{p_X}^X$ are open,*
- *$(\cup_{j=1}^{p_X} \mathcal{U}_j^X)^c$ is a closed set with Lebesgue measure zero;*
- *for all $j \in \llbracket 1, p_X \rrbracket$, $\theta \mapsto f_\theta(X)$ is an analytic function on \mathcal{U}_j^X .*

The proof of the theorem is in Appendix A.1.

This theorem formalizes that the sets $(\mathcal{O}_j^n)_{j \in \llbracket 1, m_n \rrbracket}$ (resp. $(\mathcal{U}_j^X)_{j \in \llbracket 1, p_X \rrbracket}$) almost cover the spaces $\mathbb{R}^{N_0 \times n} \times (\mathbb{R}^E \times \mathbb{R}^B)$ (resp. $\mathbb{R}^E \times \mathbb{R}^B$). Moreover, on each set \mathcal{O}_j^n (resp. \mathcal{U}_j^X) the activation pattern is constant, and the function $(X, \theta) \mapsto f_\theta(X)$ (resp. $\theta \mapsto f_\theta(X)$) is analytic. We only state that it is analytic, but when the output activation σ_L is the identity, it is in fact polynomial, and we would like to emphasize here that the structure of the polynomial is very particular. For instance, every variable appears with a degree at most one, and all monomials have the same degree. A more complete description of

the polynomial structure is, for instance, given by Bona-Pellissier u. a. (2022); Stock und Gribonval (2022).

Looking at the definition of $\tilde{\mathcal{O}}_j^n$ (resp $\tilde{\mathcal{U}}_j^X$) and \mathcal{O}_j^n (resp \mathcal{U}_j^X), using that $(\cup_{j=1}^{m_n} \mathcal{O}_j^n)^c$ (resp $(\cup_{j=1}^{p_X} \mathcal{U}_j^X)^c$) is a closed set with Lebesgue measure zero, we find that,

$$\mathcal{O}_j^n \text{ is open and dense in } \tilde{\mathcal{O}}_j^n \quad \text{and} \quad \mathcal{U}_j^X \text{ is open and dense in } \tilde{\mathcal{U}}_j^X.$$

In other words, modulo negligible sets, the activation pattern determines $\text{rank}(\mathbf{D}f_\theta(X))$. Finally, the conclusions concerning $\text{rank}(\mathbf{D}f_\theta(X))$ have direct consequences on the dimensions of the image $\{f_{\theta'}(X) \mid \theta' \in B(\theta, \varepsilon)\}$ and the pre-image $\{\theta' \in B(\theta, \varepsilon) \mid f_{\theta'}(X) = f_\theta(X)\}$, where $\varepsilon > 0$ is small enough. The consequences and their implications in machine learning applications are described in greater detail in the next sections.

When compared to existing similar statements (Stock und Gribonval, 2022; Grigsby u. a., 2022; Bona-Pellissier u. a., 2022; Grigsby und Lindsey, 2022), the particularity of Theorem 1 is that the construction of the sets \mathcal{O}_j^n and \mathcal{U}_j^X permits to include, in the third item of (i) and (ii), a statement on $\text{rank}(\mathbf{D}f_\theta(X))$. To the best of our knowledge, this quantity appears for the first time in conditions of local parameter identifiability introduced by Bona-Pellissier u. a. (2022). It appears independently a few months later, as the core quantity of a study dedicated to the geometric analysis of neural networks carried out by Grigsby u. a. (2022). In the latter article, this quantity is called the ‘batch functional dimension’ and we will use this name in this article.

Because the input space of $\mathbf{D}f_\theta(X)$ is always $\mathbb{R}^E \times \mathbb{R}^B$, the quantity $\text{rank}(\mathbf{D}f_\theta(X))$ is upper bounded by the number of parameters $|E| + |B|$. Furthermore, as formalized by Grigsby u. a. (2022), because of the invariance by positive rescaling, see the definition and discussion of the relation \sim_s in Section 2, we even have $\text{rank}(\mathbf{D}f_\theta(X)) \leq |E| + |B| - N_1 - \dots - N_{L-1}$. In fact, when $\text{rank}(\mathbf{D}f_\theta(X)) = |E| + |B| - N_1 - \dots - N_{L-1}$, under mild conditions on θ , the network function is locally identifiable around θ . That is, $f_\theta(X) = f_{\theta'}(X)$ and $\|\theta - \theta'\|$ small enough imply $\theta \sim_s \theta'$ (see Bona-Pellissier u. a., 2022).

Beyond the case of maximal rank value, $\text{rank}(\mathbf{D}f_\theta(X)) = |E| + |B| - N_1 - \dots - N_{L-1}$, leading to local identifiability, examples of non-identifiable neural networks and rank deficient parameters are in Grigsby u. a. (2022); Bona-Pellissier u. a. (2023a); Grigsby u. a. (2023); Sonoda u. a. (2021). Let us emphasize a simple example illustrating that several rank values can be achieved.

Examples 1. Consider $L \geq 3$, any neuron $v \in V_\ell$, for $\ell \in \{2, \dots, L-1\}$, and $\theta \in \mathbb{R}^E \times \mathbb{R}^B$ such that

$$b_v < 0 \quad \text{and} \quad w_{v' \rightarrow v} < 0, \text{ for all } v' \in V_{\ell-1}. \quad (13)$$

Because of the ReLU activation function, for all $x \in \mathbb{R}^{N_0}$ and all $v' \in V_{\ell-1}$, we have $(f_\theta^{\ell-1}(x))_{v'} \geq 0$, and (1) and (13) guarantee that $(f_\theta^\ell(x))_v = 0$. This holds for all θ in the open set defined by (13). In this set, the parameters $(w_{v' \rightarrow v})_{v' \in V_{\ell-1}}$ and b_v have no impact on $f_\theta(X)$, which leads to a rank deficiency of $\mathbf{D}f_\theta(X)$. Going further, consider any $\theta \in \mathbb{R}^E \times \mathbb{R}^B$. According to the above remark, to diminish $\text{rank}(\mathbf{D}f_\theta(X))$, we can change the weights arriving to a given neuron v , and assign them negative values so that (13)

holds. We can redo this operation to many neurons to diminish the rank further. As a conclusion to the example, many values of $\text{rank}(\mathbf{D}f_\theta(X))$ are reached at different places in the parameter/input space.

Let us conclude the section by showing that the quantity $\text{rank}(\mathbf{D}f_\theta(X))$ is invariant with respect to the positive rescaling and/or neuron permutation symmetries defined in Section 2.

Proposition 2. *Consider any deep fully-connected ReLU network architecture (E, V, σ_L) . Let $\theta, \tilde{\theta} \in \mathbb{R}^E \times \mathbb{R}^B$ such that $\theta \sim \tilde{\theta}$. Then, for any $n \in \mathbb{N}^*$ and $X \in \mathbb{R}^{N_0 \times n}$, $\mathbf{D}f_\theta(X)$ is defined if and only if $\mathbf{D}f_{\tilde{\theta}}(X)$ is defined, and in that case we have*

$$\text{rank}(\mathbf{D}f_{\tilde{\theta}}(X)) = \text{rank}(\mathbf{D}f_\theta(X)).$$

The proof of the proposition is in Appendix A.2. In this appendix, we provide Proposition 14, which includes Proposition 2 as its first statement. Two other statements in Proposition 14 provide invariance properties of the sets $\tilde{\mathcal{O}}_1^n, \dots, \tilde{\mathcal{O}}_{m_n}^n$ and $\mathcal{O}_1^n, \dots, \mathcal{O}_{m_n}^n$ with respect to positive rescalings and neuron permutations.

The invariance in Proposition 2 is a benefit of the complexity measure $\text{rank}(\mathbf{D}f_\theta(X))$. For instance, it does not hold for the local flatness of the empirical risk function studied by Cha u. a. (2021); Foret u. a. (2021); Hochreiter und Schmidhuber (1997); Keskar u. a. (2017). This leads to undesired behaviors (Dinh u. a., 2017). Similarly, complexity measures defined by norms (Bartlett u. a., 2017, 2020; Golowich u. a., 2018; Neyshabur u. a., 2015b) are not invariant to positive rescalings⁶.

4 Geometric Interpretation when X is Fixed

The statement of Theorem 1, (i) is used in Section 5. In this section, we mostly describe the consequences of Theorem 1, (ii). The next corollary is a straightforward consequence of the constant rank theorem and Theorem 1, (ii).

Corollary 3. *Consider any deep fully-connected ReLU network architecture (E, V, σ_L) .*

For any $n \in \mathbb{N}^$, $X \in \mathbb{R}^{N_0 \times n}$, $j \in \llbracket 1, p_X \rrbracket$ and $\theta \in \mathcal{U}_j^X$, there exists $\varepsilon_{X,\theta} > 0$ such that*

- the *local image set*

$$\{f_{\theta'}(X) \in \mathbb{R}^{N_L \times n} \mid \|\theta' - \theta\| < \varepsilon_{X,\theta}\}$$

is a smooth manifold of dimension $\text{rank}(\mathbf{D}f_\theta(X))$;

- the *local pre-image set*

$$\{\theta' \in \mathbb{R}^E \times \mathbb{R}^B \mid f_{\theta'}(X) = f_\theta(X) \text{ and } \|\theta' - \theta\| < \varepsilon_{X,\theta}\}$$

is a smooth manifold of dimension $|E| + |B| - \text{rank}(\mathbf{D}f_\theta(X))$.

⁶For both flatness and norms, it is, of course, possible to consider the minimum of the complexity criterion over the equivalence class of a θ element. However, this is an additional burden that is not necessary for criteria based on the functional dimension.

4.1 Example

We show on Figure 1 the sets $\tilde{\mathcal{U}}_j^X$ (left) and their image $f_{\tilde{\mathcal{U}}_j^X}(X) = \{f_\theta(X) \mid \theta \in \tilde{\mathcal{U}}_j^X\}$ (right), for $j \in \llbracket 1, 6 \rrbracket$, for a one-hidden-layer neural network of widths $(1, 1, 1)$, with the identity activation function on the last layer. To simplify notations, we denote the weights and biases $\theta = (a, b, c, d) \in \mathbb{R}^4$ so that $f_\theta(x) = b\sigma(ax + c) + d$, for all $x \in \mathbb{R}$. We consider $X = (0, 1, 2) \in \mathbb{R}^{1 \times 3}$ and

$$f_\theta(X)^T = \begin{pmatrix} b\sigma(c) + d \\ b\sigma(a + c) + d \\ b\sigma(2a + c) + d \end{pmatrix}.$$

For any $j \in \llbracket 1, 6 \rrbracket$, the sets $\tilde{\mathcal{U}}_j^X$ depend on the activations in the hidden layer. They are separated by the hyperplanes $c = 0$, $a + c = 0$, $2a + c = 0$. The conditions only depend on a and c . We represent the projection of the sets $\tilde{\mathcal{U}}_j^X$ and the lines $c = 0$, $a + c = 0$, $2a + c = 0$ in the plane (a, c) , on the left of Figure 1.

Similarly, for any $j \in \llbracket 1, 6 \rrbracket$, the image set $f_{\tilde{\mathcal{U}}_j^X}(X) \subseteq \mathbb{R}^3$ is invariant to translations by a vector (d, d, d) , for $d \in \mathbb{R}$. On the right of Figure 1, we represent for all j the intersection $\mathcal{V}_j = f_{\tilde{\mathcal{U}}_j^X}(X) \cap \mathcal{P}$ between the image set $f_{\tilde{\mathcal{U}}_j^X}(X)$ and the linear plane \mathcal{P} orthogonal to $(1, 1, 1)$, generated by the vectors $\frac{1}{\sqrt{6}}(1, 1, -2)$ and $\frac{1}{\sqrt{2}}(-1, 1, 0)$. The calculations leading to the construction of the figure are in Appendix B.

4.2 Geometry-induced Implicit Regularization

Corollary 3 is illustrated in Figure 2. There we consider a regression problem with a fixed target data matrix $Y \in \mathbb{R}^{N_L \times n}$ corresponding to the input matrix $X \in \mathbb{R}^{N_0 \times n}$. We consider the square loss $\|Y - f_\theta(X)\|^2$, for $\theta \in \mathbb{R}^E \times \mathbb{R}^B$, where $\|\cdot\|$ is the Euclidean (Frobenius) norm. We also consider θ^* minimizing the square loss.

In Figure 2, we display a (fictive) illustrative case, that can be considered as representative of the practice of deep neural networks, and of our numerical experiments in Section 7.

We consider here that in Corollary 3, $p_X = 7$. Hence, there are 7 sets $\mathcal{U}_1^X, \dots, \mathcal{U}_7^X$ forming a partition of $\mathbb{R}^E \times \mathbb{R}^B$. On Figure 2, for $j \in \llbracket 1, 7 \rrbracket$, the image of \mathcal{U}_j^X , $\{f_\theta(X) \mid \theta \in \mathcal{U}_j^X\}$, is drawn with the same color as \mathcal{U}_j^X . Locally, it has the structure of a smooth manifold of dimension r_j^X . The rank values are $r_1^X = 1$, $r_2^X = 2$, $r_3^X = 1$, $r_4^X = 0$, $r_5^X = 1$, $r_6^X = 1$, $r_7^X = 1$ and thus the full image set $\{f_\theta(X) \mid \theta \in \mathbb{R}^E \times \mathbb{R}^B\}$ is a two-dimensional object. In the figure, this full image set is mainly covered by the two-dimensional image set $\{f_\theta(X) \mid \theta \in \mathcal{U}_2^X\}$, and the six other image sets $\{f_\theta(X) \mid \theta \in \mathcal{U}_j^X\}$, $j \in \llbracket 1, 7 \rrbracket \setminus \{2\}$, of dimension one or zero, are at the boundary of $\{f_\theta(X) \mid \theta \in \mathbb{R}^E \times \mathbb{R}^B\}$. Hence, intuitively they are ‘exposed’, meaning in particular that if Y does not belong to the full image set, then the optimal prediction $f_{\theta^*}(X)$ is in one of the smaller dimensional $\{f_\theta(X) \mid \theta \in \mathcal{U}_j^X\}$, $j \in \llbracket 1, 7 \rrbracket \setminus \{2\}$. This is an illustration of the **geometry-induced implicit regularization** phenomenon put to evidence in this article. In practice, parameters found by minimizing the empirical risk numerically tend to have a small complexity as measured by $\text{rank}(\mathbf{D}f_\theta(X))$, where X is the learning sample.




Figure 1: Representation of the sets $\tilde{\mathcal{U}}_j^X$ in the space (a, c) (left) and restriction to \mathcal{P} of the corresponding image sets $\{f_\theta(X) \mid \theta \in \tilde{\mathcal{U}}_j^X\}$, $j \in \llbracket 1, 6 \rrbracket$ (right). We have $r_1^X = 1$, $r_2^X = 2$, $r_3^X = 3$, $r_4^X = 2$, $r_5^X = 3$, $r_6^X = 2$. The image of $\tilde{\mathcal{U}}_1^X$ such that $r_1^X = 1$ is reduced to $(0, 0)$ (right). The images of the sets $\tilde{\mathcal{U}}_j^X$ with $r_j^X = 2$ (i.e. $j = 2, 4, 6$) are represented with thick lines of their respective colors (right). The images of $\tilde{\mathcal{U}}_3^X$, with $r_3^X = 3$, and $\tilde{\mathcal{U}}_5^X$, with $r_5^X = 3$, are represented by dashed areas, with the corresponding colors (right).




Figure 2: Schematic representation of the sets \mathcal{U}_j^X (left) and the corresponding local image sets $\{f_\theta(X) \mid \theta \in \mathcal{U}_j^X\}$, $j \in \llbracket 1, 7 \rrbracket$ (right). We have $r_1^X = 1$, $r_2^X = 2$, $r_3^X = 1$, $r_4^X = 0$, $r_5^X = 1$, $r_6^X = 1$, $r_7^X = 1$. The image of \mathcal{U}_2^X is the curved diamond-shaped area, hatched in cyan (right). The images of the sets \mathcal{U}_j^X with $r_j^X = 1$ are represented with lines of their respective colors (right). The image of \mathcal{U}_4^X with $r_4^X = 0$ is represented by a magenta bullet point (right). We consider the square loss in $\mathbb{R}^{N_L \times n}$. The target $Y \in \mathbb{R}^{N_L \times n}$ and the global solution $f_{\theta^*}(X)$ of the regression problem are represented (right). The pre-image of $f_{\theta^*}(X)$ is displayed in brown (left). A minimizing sequence is represented by gray stars, in the parameter space (left) and the image space (right).

The illustrative situation in Figure 2 corresponds to the empirical observations made in Section 7.2 for the MNIST classification problem. In this section, we even observe that, consistently in our experiments, a larger optimal loss leads to a smaller batch functional dimension. We will also see empirically in Section 7.2 that for large parameter complexities (W large), the batch functional dimension computed on the learning sample remains moderate. There are two complementary explanations. First, even though the predictions are correct for all training examples, because of the soft-max activation on the last layer, the cross-entropy loss slightly differs from zero. Secondly, although there may exist θ for which the loss is exactly zero, this θ is apparently not in the convergence basin in which the local search algorithm optimizes.

4.3 Influence of the Geometry on the Optimization Trajectory

In Figure 2, we also display a (fictive) minimizing sequence, that is a set of pairs $(\theta_n, f_{\theta_n}(X))_{n \in \mathbb{N}}$ obtained by a numerical gradient-descent-based optimization procedure. This sequence is initialized in \mathcal{U}_2^X , then passes in \mathcal{U}_5^X and then \mathcal{U}_6^X , where the optimal solution lies. This illustrative example is an illustration of the experimental results of Section 7.3. There, during the learning phase, the sequence $(\text{rank}(\mathbf{D}f_{\theta_n}(X)))_{n \in \mathbb{N}}$ typically decreases. According to Corollary 3, this corresponds to an objective landscape that becomes flatter and flatter, in the sense that the local dimension of the pre-image of $f_{\theta_n}(X)$ increases. Locally in the parameter space, the objective function is constant along a smooth manifold of a larger dimension. This new notion of flatness resembles but slightly differs from the notion of ‘flat minima’ usually considered to explain the good generalization properties of deep learning (Keskar u. a., 2017; Dinh u. a., 2017; Foret u. a., 2021; Cha u. a., 2021; Hochreiter und Schmidhuber, 1997).

5 Rank Saturating X , when θ is Fixed

In this section, we define a dense subset of $\mathbb{R}^E \times \mathbb{R}^B$, and for θ in this subset, we analyze the maximum value of $\text{rank}(\mathbf{D}f_{\theta}(X))$, for any X of any size, in a dense set. This is a natural notion of complexity, that we call ‘computable full functional dimension’. In particular, it is independent of X and measures the expressive potential of the neural network defined by θ . It is linked to the full functional dimension defined by Grigsby u. a. (2022), but can be computed (see Section 6.2), thus the name. After giving the main definitions and establishing the first properties of the considered mathematical objects, we give the main result of this section (Theorem 6), which states that the computable full functional dimension depends only on the attainable activation patterns for the considered θ , when X varies. We also show the invariance of the computable full functional dimension with respect to neuron permutations and positive rescalings. We finally establish a simple link with the (local and global) fat-shattering dimensions of the ReLU neural networks of architecture (E, V, σ_L) .

For a fixed $\theta \in \mathbb{R}^E \times \mathbb{R}^B$ and any activation pattern $\delta \in \{0, 1\}^{N_1 + \dots + N_{L-1}}$, we denote

$$D_{\delta}(\theta) = \{x \in \mathbb{R}^{N_0} \mid a(x, \theta) = \delta\}. \quad (14)$$

It is well known (among many others, see Bona-Pellissier u. a., 2023a) that the restriction of the function $x \mapsto f_\theta(x)$ to $D_\delta(\theta)$ is affine. It is therefore smooth in $\text{Int}(D_\delta(\theta))$ but, generically, the function is not differentiable at the boundary of $D_\delta(\theta)$. For any $\theta \in \mathbb{R}^E \times \mathbb{R}^B$, following Stock und Gribonval (2022), we also define the *achievable activation patterns*

$$A(\theta) = \{\delta \in \{0, 1\}^{N_1 + \dots + N_{L-1}} \mid \text{Int}(D_\delta(\theta)) \neq \emptyset\} \quad (15)$$

and

$$\mathcal{X}_\theta = \bigcup_{\delta \in A(\theta)} \text{Int}(D_\delta(\theta)).$$

It is well known that the pieces $D_\delta(\theta)$ are polyhedral (see for instance Bona-Pellissier u. a. 2023a). Hence the complement set \mathcal{X}_θ^c is included in a finite union of hyperplanes. Hence, the set \mathcal{X}_θ is dense (and open) in \mathbb{R}^{N_0} .

We extend this definition to samples and set, for $n \in \mathbb{N}^*$ and $\theta \in \mathbb{R}^E \times \mathbb{R}^B$,

$$\mathcal{X}_\theta^n = \{X \in \mathbb{R}^{N_0 \times n} \mid \forall i \in \llbracket 1, n \rrbracket, x^{(i)} \in \mathcal{X}_\theta\}. \quad (16)$$

The set \mathcal{X}_θ^n is the n th order tensor product of the set \mathcal{X}_θ with itself. By construction, the set \mathcal{X}_θ^n is open and dense in $\mathbb{R}^{N_0 \times n}$, for all n and θ .

The results of this section will apply to all $\theta \in \mathbb{R}^E \times \mathbb{R}^B$ except those in a subset \mathcal{Z} , which will turn out to be of Lebesgue measure zero – see Proposition 4. To define the set \mathcal{Z} , we first define, for all $n \in \mathbb{N}^*$ and all $\theta \in \mathbb{R}^E \times \mathbb{R}^B$,

$$z_n(\theta) = \left\{ X \in \mathbb{R}^{N_0 \times n} \mid (X, \theta) \in \left(\bigcup_{j=1}^{m_n} \mathcal{O}_j^n \right)^c \right\}, \quad (17)$$

where $\mathcal{O}_1^n, \dots, \mathcal{O}_{m_n}^n$ are defined in (9) and described in Theorem 1. The set $z_n(\theta)$ is closed and therefore Lebesgue measurable. We write, for all $n \in \mathbb{N}^*$,

$$\mathcal{Z}_n = \left\{ \theta \in \mathbb{R}^E \times \mathbb{R}^B \mid z_n(\theta) \text{ has positive Lebesgue measure in } \mathbb{R}^{N_0 \times n} \right\} \quad (18)$$

and $\mathcal{Z} = \bigcup_{n \in \mathbb{N}^*} \mathcal{Z}_n$. We state in the following proposition the most important properties of \mathcal{Z} , used in the remaining of the article.

Proposition 4. *Consider any deep fully-connected ReLU network architecture (E, V, σ_L) .*

- (i) *For all $n \in \mathbb{N}^*$, the set \mathcal{Z}_n is Lebesgue measurable and has Lebesgue measure zero on $\mathbb{R}^E \times \mathbb{R}^B$.*
- (ii) *The set \mathcal{Z} is Lebesgue measurable and has Lebesgue measure zero on $\mathbb{R}^E \times \mathbb{R}^B$.*
- (iii) *For all $\theta \in (\mathbb{R}^E \times \mathbb{R}^B) \setminus \mathcal{Z}$, all $n \in \mathbb{N}^*$, and all $X \in \mathcal{X}_\theta^n$, the function $\theta' \mapsto f_{\theta'}(X)$ is analytic in a neighborhood of θ and it is therefore differentiable at the point θ .*

The proof of the proposition is in Appendix C.1.

Using Proposition 4, (iii), we can define, for all $n \in \mathbb{N}^*$ and all $\theta \in (\mathbb{R}^E \times \mathbb{R}^B) \setminus \mathcal{Z}$, the main objects studied in this section

$$r_n^*(\theta) = \max_{X \in \mathcal{X}_\theta^n} \text{rank}(\mathbf{D}f_\theta(X)), \quad (19)$$

and the computable full functional dimension

$$r^*(\theta) = \max_{n \in \mathbb{N}^*} r_n^*(\theta). \quad (20)$$

Notice that, although \mathcal{X}_θ^n is open and dense in $\mathbb{R}^{N_0 \times n}$ and the rank is lower semi-continuous, the existence of $X \in \mathbb{R}^{N_0 \times n} \setminus \mathcal{X}_\theta^n$ such that $\mathbf{D}f_\theta(X)$ is well defined and $\text{rank}(\mathbf{D}f_\theta(X)) > r_n^*(\theta)$ is not excluded. The computable full functional dimension $r^*(\theta)$ therefore may slightly differ from the full functional dimension defined by Grigsby u. a. (2022). It lower bounds the full functional dimension. We will see in Section 6.2 that its advantage is that it can be computed with a random X . Notice finally that in Examples 1 the rank deficiency caused by negative weights is independent of X . Therefore, $r^*(\theta)$ achieves several values, as θ varies.

Notice also that, although we take the maximum over all $n \in \mathbb{N}^*$, we know that since, for all $n \in \mathbb{N}^*$ and all $X \in \mathbb{R}^{N_0 \times n}$, $\theta \mapsto \mathbf{D}f_\theta(X)$ always has the same input dimension $|E| + |B|$, see (2), the maximum is reached for $n \leq |E| + |B|$ (see also Proposition 9 below).

The following proposition states that $r_n^*(\theta)$ equals the largest of all the r_j^n , as defined in (8), that are reachable when X varies, for the given θ .

Proposition 5. *Consider any deep fully-connected ReLU network architecture (E, V, σ_L) . For any $\theta \in (\mathbb{R}^E \times \mathbb{R}^B) \setminus \mathcal{Z}$ and $n \in \mathbb{N}^*$,*

$$r_n^*(\theta) = \max_{j \in I_n(\theta)} r_j^n,$$

where

$$I_n(\theta) = \{j \in \llbracket 1, m_n \rrbracket \mid \exists X \in \mathbb{R}^{N_0 \times n}, (X, \theta) \in \mathcal{O}_j^n\}. \quad (21)$$

The proposition is proved in Appendix C.2.

The following theorem states that the achievable activation patterns $A(\theta)$, as defined in (15), determine $r^*(\theta)$. It also states that when the prediction has more affine areas, that is for a fixed θ , $X \mapsto f_\theta(X)$ is piece-wise affine with more pieces, then this prediction is more complex, in the sense of r^* .

Theorem 6. *Consider any deep fully-connected ReLU network architecture (E, V, σ_L) . For any θ and θ' in $(\mathbb{R}^E \times \mathbb{R}^B) \setminus \mathcal{Z}$,*

$$\text{if } A(\theta) \subseteq A(\theta') \text{ then } r^*(\theta) \leq r^*(\theta');$$

as a consequence,

$$\text{if } A(\theta) = A(\theta') \text{ then } r^*(\theta) = r^*(\theta').$$

The proof of the theorem is in Appendix C.3.

Next, we show that $r_n^*(\theta)$, $r^*(\theta)$ and \mathcal{Z} are invariant by neuron permutation and positive rescaling (recall the relations \sim , \sim_s and \sim_p presented in Section 2).

Proposition 7. *Let $\theta \in \mathbb{R}^E \times \mathbb{R}^B$ and $\tilde{\theta} \sim \theta$. Then $\theta \in \mathcal{Z} \iff \tilde{\theta} \in \mathcal{Z}$. Also, if $\theta \in (\mathbb{R}^E \times \mathbb{R}^B) \setminus \mathcal{Z}$, $r_n^*(\theta) = r_n^*(\tilde{\theta})$ for all $n \in \mathbb{N}^*$ and $r^*(\theta) = r^*(\tilde{\theta})$.*

The proof of the proposition is in Appendix C.4.

Let us conclude this section by showing that $r^*(\theta)$ provides lower-bounds on various fat-shattering dimensions for neural networks. The fat-shattering dimension of a family of regression functions is a well-known measure of complexity (see for instance Anthony und Bartlett 2009, Chapter 11). In the rest of the section, we let $N_L = 1$ and, for a subset $A \subseteq \mathbb{R}^E \times \mathbb{R}^B$, for $\gamma > 0$, the *fat-shattering dimension* of the family $\{f_\theta \mid \theta \in A\}$, that we write $\text{fS}_{A,\gamma}$, is defined as follows. It is the largest $n \in \mathbb{N}^*$ such that there exist $x^{(1)}, \dots, x^{(n)} \in \mathbb{R}^{N_0}$ and $t_1, \dots, t_n \in \mathbb{R}$ such that for all $I \subseteq \llbracket 1, n \rrbracket$, there is $\theta \in A$ such that for $i \in I$, $f_\theta(x^{(i)}) \geq t_i + \gamma$ and for $i \in \llbracket 1, n \rrbracket \setminus I$, $f_\theta(x^{(i)}) \leq t_i - \gamma$. If this property holds for all n then we let $\text{fS}_{A,\gamma} = \infty$.

The intuition is that $\text{fS}_{A,\gamma}$ is the largest number n of input points for which all 2^n combinations of being above or below the threshold t_i by a margin γ , $i \in \llbracket 1, n \rrbracket$, can be reached by the functions in $\{f_\theta \mid \theta \in A\}$ (see Anthony und Bartlett 2009). When A is a small ball centered at a parameter of interest we shall call $\text{fS}_{A,\gamma}$ a local fat-shattering dimension, and when $A = \mathbb{R}^E \times \mathbb{R}^B$, we shall call $\text{fS}_{A,\gamma}$ a global fat-shattering dimension. We also consider the case where A is the set of parameters yielding a given computable full functional dimension, where we refer to $\text{fS}_{A,\gamma}$ as the *fixed-rank fat-shattering dimension*. The next proposition shows the announced lower bounds.

Proposition 8. *Consider any deep fully-connected ReLU network architecture (E, V, σ_L) such that $N_L = 1$.*

Let $\theta \in (\mathbb{R}^E \times \mathbb{R}^B) \setminus \mathcal{Z}$. Then for any $\varepsilon > 0$, there is $\gamma > 0$ such that we have the following lower bound on the local fat-shattering dimension,

$$\text{fS}_{B(\theta,\varepsilon),\gamma} \geq r^*(\theta). \quad (22)$$

As a first consequence, there is $\gamma' > 0$ such that the global fat-shattering dimension is lower bounded as follows,

$$\text{fS}_{\mathbb{R}^E \times \mathbb{R}^B, \gamma'} \geq \max_{\theta \in (\mathbb{R}^E \times \mathbb{R}^B) \setminus \mathcal{Z}} r^*(\theta).$$

As a second consequence, the fixed-rank fat-shattering dimension is lower bounded as follows. Consider $r \in \llbracket 1, |E| + |B| - N_1 - \dots - N_{L-1} \rrbracket$. Let $A_r = \{\theta \in (\mathbb{R}^E \times \mathbb{R}^B) \setminus \mathcal{Z} \mid r^(\theta) = r\}$. If A_r has non-empty interior, there is $\gamma'' > 0$ such that*

$$\text{fS}_{A_r, \gamma''} \geq r.$$

The proof of the proposition is in Appendix C.5.

It consists first in obtaining local continuous differentiability, with an invertible square Jacobian matrix and second in applying the inverse function theorem. Variations of this second step were already carried out in the literature, in particular by Erlich u. a. (1997).

Remark that the same proof would also apply to other measures of complexity, for instance the Vapnik–Chervonenkis (VC) dimension (Anthony und Bartlett, 2009, Chapter 3) of the binary classifiers indexed by θ and obtained by taking the sign of $f_\theta - f_{\theta_0}$ for any fixed $\theta_0 \in (\mathbb{R}^E \times \mathbb{R}^B) \setminus \mathcal{Z}$ (with σ_L as the identity).

Our motivation, for studying the fat-shattering dimension (or the VC-dimension), is their relationships with notions of generalization errors in machine learning, and with uniform convergence in probability and statistics, (see in particular Alon u. a. 1997; Bartlett und Long 1998; Colomboni u. a. 2023; Vapnik und Chervonenkis 1971 and references therein). In particular, Proposition 8 indicates that the computable full functional dimension $r^*(\theta)$ can be seen as relevant for studying the generalization error of neural networks in machine learning.

Finally, remark that Bartlett u. a. (2019, 1998); Harvey u. a. (2017); Maass (1994) relate the global VC dimension of neural networks to their number of parameters (and their depths) while in Proposition 8 we consider local or global dimensions and relate them to the (smaller) computable full functional dimension.

6 Computational Considerations

In this section, we describe how to compute the quantities of this article in practice. In Section 6.1, we describe how one can efficiently compute $\text{rank}(\mathbf{D}f_\theta(X))$ for a given X , and in Section 6.2 we explain how this allows to compute $r^*(\theta)$ by sampling, for a sufficient number of samples.

6.1 How to Compute $\text{rank}(\mathbf{D}f_\theta(X))$

For a given $X \in \mathbb{R}^{N_0 \times n}$ and a given $\theta \in \mathbb{R}^E \times \mathbb{R}^B$, $\text{rank}(\mathbf{D}f_\theta(X))$ is computed using the backpropagation and numerical linear algebra tools computing the rank of a matrix. To justify the computations, let us first recall the classical backpropagation algorithm for computing the gradients with respect to the parameters of the network, for a given loss $Lo : \mathbb{R}^{N_L} \times \mathbb{R}^{N_L} \rightarrow \mathbb{R}$. We will then describe how to use the backpropagation to compute $\text{rank}(\mathbf{D}f_\theta(X))$. We conclude with implementation recommendations.

For a given input $x \in \mathbb{R}^{N_0}$ and a given output $y \in \mathbb{R}^{N_L}$, backpropagation computes the gradient $\nabla Lo(f_\theta(x), y)$ of the function $\theta \mapsto Lo(f_\theta(x), y)$. To do so, it first computes $f_\theta(x)$ and stores the intermediate pre-activation values $(y_\theta^\ell)_v = \sum_{v' \in V_{\ell-1}} w_{v' \rightarrow v} (f_\theta^{\ell-1}(x))_{v'} + b_v$, for $\ell \in \llbracket 1, L \rrbracket$ and $v \in V_\ell$. This is known as the ‘forward pass’. Then, backpropagation computes the vector of errors η_θ^L defined by

$$\eta_\theta^L = \left(J\sigma_L(y_\theta^L) \right)^T \frac{\partial Lo}{\partial y_1}(f_\theta(x), y),$$

where $\frac{\partial Lo}{\partial y_1}(f_\theta(x), y) \in \mathbb{R}^{N_L}$ is the gradient of $y_1 \mapsto Lo(y_1, y)$, at the point $(f_\theta(x), y)$, and $J\sigma_L(y_\theta^L)$ is the Jacobian matrix of $y^L \mapsto \sigma_L(y^L)$, at y_θ^L . This vector is then backpropagated, from $\ell = L$ to $\ell = 1$ thanks to the equation

$$\forall v' \in V_{\ell-1} \quad \left(\eta_\theta^{\ell-1}\right)_{v'} = \sigma' \left(\left(y_\theta^{\ell-1}\right)_{v'} \right) \sum_{v \in V_\ell} w_{v' \rightarrow v} \left(\eta_\theta^\ell\right)_v \quad (23)$$

where $\sigma'(t) = 1$ if $t > 0$ and $\sigma'(t) = 0$ if $t \leq 0$. This allows to recursively obtain the error vectors $\eta_\theta^\ell \in \mathbb{R}^{N_\ell}$, for all $\ell \in \llbracket 1, L \rrbracket$. We deduce the partial derivatives thanks to the formulas

$$\forall \ell \in \llbracket 1, L \rrbracket, \forall v' \in V_{\ell-1}, \forall v \in V_\ell, \quad \frac{\partial Lo(f_\theta(x), y)}{\partial w_{v' \rightarrow v}} = \sigma \left(\left(y_\theta^{\ell-1}\right)_{v'} \right) \left(\eta_\theta^\ell\right)_v$$

and

$$\forall \ell \in \llbracket 1, L \rrbracket, \forall v \in V_\ell, \quad \frac{\partial Lo(f_\theta(x), y)}{\partial b_v} = \left(\eta_\theta^\ell\right)_v.$$

This allows computing the gradients for one example (x, y) . For a batch, the algorithm is repeated for each example $(x^{(i)}, y^{(i)})$, and the average of the so obtained gradients is computed.

Let us now make the connection between backpropagation and the computation of $\text{rank}(\mathbf{D}f_\theta(X))$. Vectorizing both the input and output spaces of $\theta \mapsto f_\theta(X)$, we first notice that $\text{rank}(\mathbf{D}f_\theta(X)) = \text{rank}(Jf_\theta(X))$, where the Jacobian matrix $Jf_\theta(X) \in \mathbb{R}^{n_{N_L} \times (|E|+|B|)}$ takes the form

$$Jf_\theta(X) = \begin{pmatrix} Jf_\theta(x^{(1)}) \\ \vdots \\ Jf_\theta(x^{(n)}) \end{pmatrix}$$

and, for all $i \in \llbracket 1, n \rrbracket$, $Jf_\theta(x^{(i)}) \in \mathbb{R}^{N_L \times (|E|+|B|)}$ is the Jacobian matrix of $\theta \mapsto f_\theta(x^{(i)})$. We construct the matrix $Jf_\theta(X)$ by successively computing each of its lines, i.e. computing each line of $Jf_\theta(x^{(i)})$ for all $i \in \llbracket 1, n \rrbracket$.

For a given $i \in \llbracket 1, n \rrbracket$ and $v \in V_L$, the line corresponding to v of $Jf_\theta(x^{(i)})$ is indeed simply obtained as the transpose of $\nabla Lo_v(f_\theta(x^{(i)}), y^{(i)})$ for the function $Lo_v : \mathbb{R}^{N_L} \times \mathbb{R}^{N_L} \rightarrow \mathbb{R}$ defined by $Lo_v(y_1, y_2) = (y_1)_v$, for all $(y_1, y_2) \in \mathbb{R}^{N_L} \times \mathbb{R}^{N_L}$. We indeed have $Lo_v(f_{\theta'}(x^{(i)}), y^{(i)}) = f_{\theta'}(x^{(i)})_v$ for all θ' . The gradient $\nabla Lo_v(f_\theta(x^{(i)}), y^{(i)})$ is obtained using the backpropagation algorithm described above. Notice that when σ_L is the identity, for a given $v \in V_L$, using the definition of Lo_v , we always have $(\eta_\theta^L)_v = 1$ and $(\eta_\theta^L)_{v'} = 0$ for all $v' \neq v$. We need however to compute the forward pass in order to compute the vectors y_θ^ℓ , for $\ell \in \llbracket 0, L-1 \rrbracket$. Finally, once $Jf_\theta(X)$ is computed its rank is obtained using standard linear algebra algorithms.

Our implementation uses the existing automatic differentiation of **Tensorflow**. It is possible to call the method `GradientTape.gradients`, which computes $Jf_\theta(x)$ for a single

⁷Neural networks libraries such as **Tensorflow** set $\sigma'(0) = 0$ and we adopt this convention in this calculus. Due to numerical imprecision, we rarely have $y_{\theta v}^{\ell-1} = 0$ in practice. In the theoretical sections of this article, the situation $\sigma'(0)$ never occurs for the cases where $\mathbf{D}f_\theta(X)$ is considered.

example x , and to repeat it for each example $x^{(i)}$. However, it is more efficient to use `GradientTape.jacobian` which allows to compute directly $Jf_\theta(X)$. We do not report the details of the experiments here but we found even more efficient to cut X in sub-batches and repeatedly call `GradientTape.jacobian`, when appropriately choosing the size of the sub-batches.

Once $Jf_\theta(X)$ built, the value of $\text{rank}(Jf_\theta(X))$ can be computed with the `np.linalg.rank` function of `Numpy`, or using the accelerated rank computation of `Pytorch` with a GPU, which improves the speed by some factors. Note that the limiting factor when computing $\text{rank}(Jf_\theta(X))$ for large networks and/or n large is the computation of the rank and not the construction of $Jf_\theta(X)$.

The codes are available at (Bona-Pellissier u. a., 2023b).

6.2 How to Compute $r^*(\theta)$

In this section, our goal is to estimate the maximal rank $r^*(\theta)$, see (20), from $\text{rank}(\mathbf{D}f_\theta(X))$, where $X \in \mathbb{R}^{N_0 \times n}$ is a random data set composed of n i.i.d samples. Such an estimate is already considered by Grigsby u. a. (2023). Intuitively, the bigger n is, the better the estimation. Indeed, we provide an upper bound on the probability that $\text{rank}(\mathbf{D}f_\theta(X)) < r^*(\theta)$ as a function of n , see Theorem 11. This probability also depends on the probability of generating an example in the least probable linear region of $x \mapsto f_\theta(x)$. This result can be compared to the smallest possible sample size $n^*(\theta)$ obtained if an optimal $X \in \mathbb{R}^{N_0 \times n^*(\theta)}$ was provided by an oracle, see Proposition 9.

This proposition proves that this smallest possible sample size $n^*(\theta)$ has the order of magnitude of $r^*(\theta)$. Before stating the proposition, we remind that, since the input space of $\mathbf{D}f_\theta(X)$ is always $\mathbb{R}^E \times \mathbb{R}^B$, we always have⁸ $r^*(\theta) \leq |E| + |B|$.

Proposition 9. *Consider any deep fully-connected ReLU network architecture (E, V, σ_L) .*

Let $\theta \in (\mathbb{R}^E \times \mathbb{R}^B) \setminus \mathcal{Z}$. Consider the sequence $(r_n^(\theta))_{n \in \mathbb{N}^*}$. There exists (a unique) $n^*(\theta) \in \mathbb{N}^*$ such that this sequence is increasing for $1 \leq n \leq n^*(\theta)$ and stationary (constant) for $n^*(\theta) \leq n$. Furthermore, if $r^*(\theta) \neq 0$, we also have*

$$\frac{r^*(\theta)}{N_L} \leq n^*(\theta) \leq r^*(\theta).$$

The proof of the proposition is in Appendix D.1.

As its proof shows, the following proposition is a direct consequence of Proposition 5. It already guarantees that, without any knowledge of the problem, a random X following a sufficiently spread distribution can be used to calculate $r_n^*(\theta)$ and therefore $r^*(\theta)$. Its purpose is to illustrate how the statements in the previous sections can be used to calculate $r^*(\theta)$. A better statement is given in Theorem 11.

Proposition 10. *Consider any deep fully-connected ReLU network architecture (E, V, σ_L) .*

⁸A tighter upper-bound taking into account the positive rescaling invariance of ReLU networks is given by Grigsby u. a. (2022), as discussed in Section 3.

Let $\theta \in (\mathbb{R}^E \times \mathbb{R}^B) \setminus \mathcal{Z}$. Let $n \in \mathbb{N}^*$. The set $\{X \in \mathcal{X}_\theta^n \mid \text{rank}(\mathbf{D}f_\theta(X)) = r_n^*(\theta)\}$ has non-zero Lebesgue measure (on $\mathbb{R}^{N_0 \times n}$).

Proof of Proposition 10. From Proposition 5, there exists $j \in I_n(\theta)$ such that $r_j^n = r_n^*(\theta)$. We then have the inclusion

$$\{X \in \mathcal{X}_\theta^n \mid (X, \theta) \in \mathcal{O}_j^n\} \subseteq \{X \in \mathcal{X}_\theta^n \mid \text{rank}(\mathbf{D}f_\theta(X)) = r_n^*(\theta)\}.$$

Since \mathcal{O}_j^n is open, the left-hand set above is an open set, which is non-empty by definition of $I_n(\theta)$, in (21). Hence, the right-hand set above has a non-zero Lebesgue measure. \square

Theorem 11. Consider any deep fully-connected ReLU network architecture (E, V, σ_L) .

Let us consider a distribution \mathcal{G} over \mathbb{R}^{N_0} , that is absolutely continuous with respect to Lebesgue measure, with (strictly) positive density. Assume we sample randomly and independently the vectors $x^{(i)} \in \mathbb{R}^{N_0}, i \in \llbracket 1, n \rrbracket$, following the distribution \mathcal{G} , for some $n \in \mathbb{N}^*$.

Let $\theta \in (\mathbb{R}^E \times \mathbb{R}^B) \setminus \mathcal{Z}$ and

$$p = \min_{\delta \in A(\theta)} \mathbb{P}(x^{(1)} \in \text{Int } D_\delta(\theta)),$$

where $A(\theta)$ and $D_\delta(\theta)$ are defined in (15) and (14). Note that we have $p > 0$, because $A(\theta)$ is finite and, for all $\delta \in A(\theta)$, $D_\delta(\theta)$ has nonempty interior.

Then the following holds.

1. Consider i.i.d. Bernoulli random variables B_1, \dots, B_n , with $\mathbb{P}(B_1 = 1) = p$ and $\mathbb{P}(B_1 = 0) = 1 - p$. We have

$$\mathbb{P}\left(\text{rank}\left(\mathbf{D}f_\theta\left(\left(x^{(i)}\right)_{1 \leq i \leq n}\right)\right) < r^*(\theta)\right) \leq \mathbb{P}(B_1 + \dots + B_n < n^*(\theta)).$$

2. As a consequence, if $n \geq 2n^*(\theta)/p$,

$$\mathbb{P}\left(\text{rank}\left(\mathbf{D}f_\theta\left(\left(x^{(i)}\right)_{1 \leq i \leq n}\right)\right) < r^*(\theta)\right) \leq \frac{4}{np}.$$

The proof of the theorem is in Appendix D.2.

A first consequence of the theorem is that if one simply adds columns to an input matrix X randomly and independently, the corresponding value of $\text{rank}(\mathbf{D}f_\theta(X))$ will reach the computable full functional dimension, almost surely (this consequence alone could be seen/proved more simply). This can for instance help understand the experimental results of Grigsby u. a. (2023).

The theorem then provides two upper bounds (Items 1 and 2) on the probability of not reaching the computable full functional dimension, as a vanishing function of the number of columns n . A beneficial feature of these upper bounds is that they are based on p , the smallest probability of reaching a given region $\text{Int } D_\delta(\theta)$, relative to a given activation

pattern δ , for a *single* column sample. Importantly, the probabilities of reaching several given activation patterns *simultaneously* over multiple samplings of the columns, which would be typically much smaller than p , are not involved.

Also, the upper bounds are based on $n^*(\theta)$, the smallest number of columns for X such that we can have $\text{rank}(\mathbf{D}f_\theta(X)) = r^*(\theta)$. This is natural since the larger $n^*(\theta)$, the more samples we need to have a non-zero probability of reaching the computable full functional dimension. Typically, for n of the order of magnitude of $n^*(\theta)$ (for instance such that $np \gtrsim 2n^*(\theta)$), since $n^*(\theta)$ is usually large, we already have a high probability that the lower-bound $\text{rank}(\mathbf{D}f_\theta((x^{(i)})_{1 \leq i \leq n}))$ coincides with $r^*(\theta)$.

The first upper bound (Item 1) is the tightest and most general. The second one (Item 2) simply follows from Chebyshev’s inequality and is provided for the sake of obtaining a straightforward compact bound. Other upper bounds could be obtained simply from Item 1, using for instance the Hoeffding inequality.

7 Experiments

The experiments emphasize several aspects of the *geometry-induced implicit regularization* illustrated in Figure 2. The setting of the experiments is described in Section 7.1. In Section 7.2, we describe the results of an experiment in which we compute the functional dimensions as the number of parameters of the network grows. In Section 7.3, we compute functional dimensions throughout the learning phase. In Section 7.4, we investigate the impact of the corruption of the inputs of the learning sample on the functional dimensions. In Section 7.5, we perform the same experiment but corrupt the outputs of the learning sample.

The Python codes implementing the experiments described in this section are available at (Bona-Pellissier u. a., 2023b).

7.1 Experiments Description

In the experiments of Section 7.2, 7.3, 7.4 and 7.5, we evaluate the behavior of different complexity measures for the classification of a subpart of the MNIST data set.

We consider a fully-connected feed-forward ReLU network of depth $L = 4$, of widths $(N_0, N_1, N_2, N_3, N_4) = (784, w, w, w, 10)$, for different values of $w \in \llbracket 1, 60 \rrbracket$. The tested values of w depend on the experiment/section. The hidden layers (1, 2, 3) include a ReLU activation function. The last layer includes a soft-max activation function. We randomly extract a training sample $(X_{\text{train}}, Y_{\text{train}})$, containing 6 000 images and a test sample $(X_{\text{test}}, Y_{\text{test}})$ containing 20 000 images from MNIST.

For given w and $(X_{\text{train}}, Y_{\text{train}})$, we tune the parameters of the network to minimize the cross-entropy. This is achieved using the Glorot uniform initialization for the weights while initializing the biases to 0, and using the stochastic gradient descent ‘sgd’ as optimizer with a learning rate of 0.1 and a batch size of 256. The number of epochs depends on the experiment/section.

In the figures presenting the results of the experiments, we display the following quantities:

- Max rank: the maximal theoretically possible value of $\text{rank}(\mathbf{D}f_\theta(X))$ for any sample X and parameter θ . It is equal to $|E| + |B| - N_1 - \dots - N_{L-1} = N_0N_1 + N_1N_2 + \dots + N_{L-1}N_L + N_L$ (see the bound provided by Grigsby u. a. (2022), Theorem 7.1). With the architecture described above, for a given w , the Max_rank is equal to $2w^2 + 794w + 10$. This is very close to the number of parameters $2w^2 + 797w + 10$. Furthermore, with the values of w considered in the forthcoming experiments, the predominant term is $794w$.
- Rank X_random: an estimation of the computable full functional dimension, according to the statement of Section 6.2, by computing $\text{rank}(\mathbf{D}f_\theta(X_{\text{random}}))$ with a random i.i.d. sample X_{random} , where each example of the sample is a Gaussian random vector. The number of examples is equal to 20 000 or 40 000 depending on the experiment/section.
- Rank X_test: It corresponds to $\text{rank}(\mathbf{D}f_\theta(X_{\text{test}}))$, where X_{test} is the test sample of size 20 000 introduced above. Note that the test set is bigger than the train set, in contrast to classical settings. Indeed, the test set serves two purposes here: it is classically used to compute the classification accuracy, but it is also meant to provide an estimation of the functional dimension over the distribution of the inputs (the MNIST images). The latter estimation differs from the estimation with X_{random} which samples images outside the distribution of the inputs.
- Rank X_train: It corresponds to $\text{rank}(\mathbf{D}f_\theta(X_{\text{train}}))$, where X_{train} is the training sample of size 6 000 mentioned above. This quantity is the batch functional dimension.
- Train loss: the value of the training loss at the end of the training (resp. at the current epoch) in Sections 7.2, 7.4 and 7.5 (resp. in Section 7.3).
- Test error: the proportion of images of X_{test} that are misclassified by the network.
- Train error: the proportion of images of X_{train} that are misclassified by the network.

7.2 Behavior of the Functional Dimensions as the Network Width Increases

In this experiment, we evaluate the functional dimensions when the width w varies between 1 and 85. More precisely, we test all w between 1 and 9, then all w between 10 and 18 with an increment of 2, and then all w between 20 and 85 with an increment of 5. Overall, the number of parameters of the network varies between 809 and 82 205

As described in Section 7.1, we randomly extract a training sample $(X_{\text{train}}, Y_{\text{train}})$ of size 6 000 and a test sample $(X_{\text{test}}, Y_{\text{test}})$ of size 20 000 from MNIST. The size of the i.i.d Gaussian sample X_{random} is 40 000. We optimize the network parameters during 1 000 epochs.

The results of the experiment are in Figure 3. When increasing the number of parameters, the train loss, the train error and the test error decrease. For $w \geq 12$, i.e. when the number of parameters is superior or equal to 9 862, the train error is equal to 0: the network is able




Figure 3: Behavior of different complexity measures as the size of the network increases. The values of the several ranks are to be read on the left axis, titled ‘functional dimension’. The values for the test and train errors and the train loss are to be read on the right axis.

to fit perfectly the training images. However, the test error continues to decrease even after the train error reaches 0: from 0.101 when $w = 12$ to 0.058 when $w = 85$.

As we can see, the quantity $\text{rank}(\mathbf{D}f_\theta(X_{\text{random}}))$ in the case of the 40 000 inputs generated as Gaussian vectors, X_{random} , is nearly always equal to its maximum theoretical value Max_rank . This indicates first that $r^*(\theta)$ is equal to its maximal theoretical value, and second that $\text{rank}(\mathbf{D}f_\theta(X_{\text{random}}))$ provides here a good estimation of $r^*(\theta)$. Furthermore, according to Bona-Pellissier u. a. (2022), the networks parameters are locally identifiable from X_{random} in this case.

The ranks $\text{rank}(\mathbf{D}f_\theta(X_{\text{train}}))$ and $\text{rank}(\mathbf{D}f_\theta(X_{\text{test}}))$ are nearly equal when the number of parameters is smaller than 21 185 ($w = 25$). For these sizes, it seems to indicate that $\text{rank}(\mathbf{D}f_\theta(X_{\text{test}}))$ is indeed equal to the functional dimension over the distribution of inputs, and that $\text{rank}(\mathbf{D}f_\theta(X_{\text{train}}))$ already attains it, which means that adding MNIST images to X_{train} would not increase $\text{rank}(\mathbf{D}f_\theta(X_{\text{train}}))$. It thus also suggests that the functional dimension over the distribution of inputs (the MNIST images) is smaller than $r^*(\theta)$ which consider all X inside but also outside of this distribution, and which, as we have seen in the previous paragraph, is maximal here. In the light of Bona-Pellissier u. a. (2022), this also shows that θ is not locally identifiable from X_{train} nor X_{test} . This suggests that for these networks and the MNIST dataset, using only samples of the input distribution does not allow to identify the parameters of a network, and one needs to add examples outside the input distribution.

Then, for more than 21 185 parameters, a gap appears between the two ranks $\text{rank}(\mathbf{D}f_\theta(X_{\text{train}}))$ and $\text{rank}(\mathbf{D}f_\theta(X_{\text{test}}))$, i.e. $\text{rank}(\mathbf{D}f_\theta(X_{\text{train}}))$ is smaller than the functional dimension over the distribution of the inputs. Furthermore, while both ranks are not far from the maximum rank for small numbers of parameters, the gap increases with the number of parameters, to the point where the shape of the curves seem to diverge: while the maximum rank is nearly

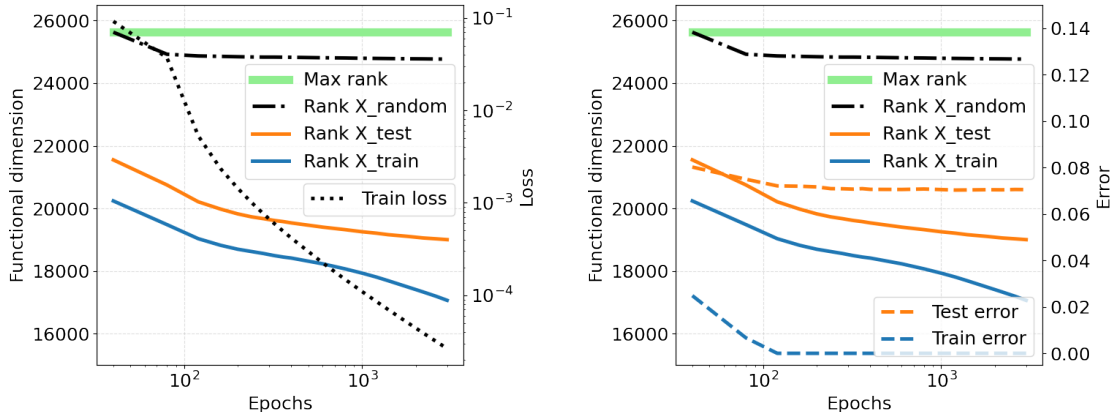


Figure 4: Behavior of different complexity measures during training. The values of the different ranks are to be read on the left axis, titled ‘functional dimension’. The values of the train loss (on the left figure), and the values of the test and train errors (on the right figure) are to be read on the right axis.

proportional to the number of parameters, the ranks $\text{rank}(\mathbf{D}f_\theta(X_{\text{train}}))$ and $\text{rank}(\mathbf{D}f_\theta(X_{\text{test}}))$ seem to increase less and less with the number of parameters. This is the consequence of the geometry-induced implicit regularization described in Section 4.2 and Figure 2. The regularization on the training sample seems to also concern the input distribution in general as the curve $\text{rank}(\mathbf{D}f_\theta(X_{\text{test}}))$ indicates.

7.3 Behavior of the Functional Dimensions During Training

We consider the setting described in Section 7.1, where we fix the value of w to 30. For this experiment, we keep the original MNIST images and labels in $(X_{\text{train}}, Y_{\text{train}})$ and $(X_{\text{test}}, Y_{\text{test}})$ and we set the size of the random set X_{random} to 20 000. The architecture is $(784, 30, 30, 30, 10)$, which corresponds to a total number of parameters equal to 25 720. The quantities plotted in the previous experiment (see Figure 3) are computed after the training is done. In contrast, here, we fix a total number of epoch to 3 000 and we compute the same quantities during training, throughout the epochs.

More precisely, we study the quantities Max rank, Rank X_random, Rank X_test, Rank X_train, Train loss, Test error and Train error, as described in Section 7.1. They are computed at the epochs $\{40, 80, 120, 160, 200, 240, 280, 320, 360, 400\} \cup \{600, 800, 1000, 1200, 1400\} \cup \{1800, 2200, 2600, 3000\}$. We plot these quantities in Figure 4.

We plot the train loss (on the left), which decreases throughout the epochs, and the train error (on the right), which decreases and reaches 0 at epoch 120, after which all training images are always correctly classified. The test error decreases the most in the first 80 epochs, after which it continues to decrease, although at a slower pace.

At the beginning of the training, the quantity $\text{rank}(\mathbf{D}f_\theta(X_{\text{random}}))$ is equal to the




Figure 5: Effect of the noise on a MNIST image for different amplitudes. From left to right, a MNIST image to which has been added a Gaussian noise, of amplitude 10^{-2} , 10^{-1} , 5×10^{-1} and 10^0 respectively.

maximum rank, here equal to 25 630, and it decreases to reach 24 769 at the end of the training. As can be seen in the experiments of Sections 7.2, 7.4 and 7.5, most of the time, at the end of the training, $\text{rank}(\mathbf{D}f_{\theta}(X_{\text{random}}))$ is equal or very close to the maximum rank, but sometimes, there is a small gap between both quantities. The training studied here corresponds to one of the latter cases. As shown by Bona-Pellissier u. a. (2022), when $\text{rank}(\mathbf{D}f_{\theta}(X_{\text{random}}))$ is maximal, the parameterization θ is locally identifiable from the sample X_{random} . On the contrary, the gap observed here at the end of the training indicates some lack of identifiability from X_{random} .

We observe that the value of $\text{rank}(\mathbf{D}f_{\theta}(X_{\text{train}}))$ consistently decreases during training. Such a behavior is consistent with the geometric interpretation of Section 4.2 and 4.3, and Figure 2. The value of $\text{rank}(\mathbf{D}f_{\theta}(X_{\text{test}}))$ also decreases, with a more gentle slope. This indicates that the geometry-induced implicit regularization occurring on the training sample is ‘communicated’ to the test sample.

7.4 Behavior of the Functional Dimensions when X is Corrupted

We consider the same setting as the experiments of Section 7.1, with $w = 30$, which corresponds to a total number of parameters equal to 25 720.

The network is trained, during 3 000 epochs, repeatedly over different train sets of size 6 000 made of MNIST images. We add to the train images a Gaussian noise, before clipping the values of the pixels between 0 and 1, to stay consistent with black and white images. We do the same for the test images. This blurs the input distribution. For each training, we use a different noise variance, which overall varies between 10^{-3} and 1. We represent visually an image with different levels of noise in Figure 5. The network is trained to the point it is able to interpolate the training examples: for all the settings, the final train error is equal to 0.

Once the training is done, we compute the quantities Max rank, Rank X_{random} , Rank X_{test} , Rank X_{train} , Test error, and Train error described in Section 7.1, for the different noise levels. The size of X_{random} is 20 000. We plot these quantities in Figure 6.

As already said, the expressiveness of the network permits to fit the learning data




Figure 6: Behavior of different complexity measures when noise is added to the input images. The values of $\text{rank}(\mathbf{D}f_\theta(X_{\text{train}}))$ are to be read on the left axis, titled ‘functional dimension’. The values of the test and train errors are to be read on the right axis.

perfectly. The training error is zero for all noise levels. However, the noise has two effects: an effect on the distribution of inputs which becomes more complex, and an effect on the difficulty of the problem. Indeed, as is reflected by the increase in the test error, the problem becomes more difficult. The curve representing $\text{rank}(\mathbf{D}f_\theta(X_{\text{train}}))$ also increases, and gets close to the maximum rank for a noise amplitude getting close to 1. This phenomenon is coherent with the fact that the batch functional dimension is linked to activation patterns, which are linked to the distribution of the inputs, which –as already said– are made more complex by the noise. The quantity $\text{rank}(\mathbf{D}f_\theta(X_{\text{test}}))$ also increases, but for smaller noises than X_{train} . It oscillates close to the maximum rank.

7.5 Behavior of the Functional Dimensions when Y is Corrupted

We consider the setting described in Section 7.1, for three different values of w , equal to 30, 50 and 60 respectively. The first experiment carried out was the one with the width $w = 30$, in harmony with the experiments of Sections 7.3 and 7.4, and we then added the cases $w = 50$ and $w = 60$ to see the impact of the width on this experiment.

Following Neyshabur u. a. (2017); Zhang u. a. (2021), we study what happens when part of the labels of the training set are corrupted with random labels. The network is trained repeatedly over different train sets of size 6 000 made of MNIST images. A varying number of the labels associated with the training images are set to a random value, according to the uniform distribution over $\{0, 1, \dots, 9\}$. The quantity of images with random labels varies from 0 to 4 500. Each time, the network is trained to the point where it interpolates the training set, even with the random labels.

Intuitively, the more corrupted labels there are, the more complex the function interpo-

lating the training data should be. The distribution of the inputs is however the same and it is not clear if the linear regions, defined by the activation patterns, and used to describe the interpolating function need to change much when corruption increases. Another interpretation is that depending on the level of corruption, we select different global minimizers of the empirical risk with respect to the 1500 images that always remain clean. The purpose of this experiment is to see if the functional dimension increases with the proportion of corrupted labels and the test error, when w is such that implicit regularization occurs.

The quantities Max_rank, Rank_X_random, Rank_X_test, Rank_X_train, Train loss, Test error and Train error are described in Section 7.1. We plot these quantities as a function of the number of corrupted labels in Figure 7 for three different values of w . The case $w = 30$, is described in Section 7.5.1, and the case $w = 50$ is described in Section 7.5.2.

7.5.1 Observations when $w = 30$

Consistently with Sections 7.3 and 7.4, we take $w = 30$, which corresponds to a number of parameters equal to 25 720. The size of the test set X_{test} and of the random Gaussian inputs X_{random} are both equal to 20 000. The learning performs 3 000 epochs. The plots associated to this experiment correspond to the left side of Figure 7.

As expected, the test error (the proportion of images in X_{test} that are misclassified) increases with the number of corrupted labels: from 0.08 when no label is corrupted to 0.73 when 4500 images out of 6000 have random labels.

In this experiment, again, we observe that $\text{rank}(\mathbf{D}f_{\theta}(X_{\text{random}}))$ is often equal to the maximum rank, which is here equal to 25630. For numbers of corrupted labels equal to 500 and 1500, we observe that the value is slightly lower.

Although we can observe a clear increase of the ranks $\text{rank}(\mathbf{D}f_{\theta}(X_{\text{train}}))$ and $\text{rank}(\mathbf{D}f_{\theta}(X_{\text{test}}))$ between 0 and 500 corrupted labels, for a greater number of corrupted labels these quantities are relatively stable and do not seem to be much affected by the corruption of the labels. We try to clarify this observation by performing two other experiments with a bigger width w in the following section.

7.5.2 Observations when $w = 50$ and $w = 60$

To better understand the impact of label corruption on the functional dimension, we repeat the previous experiment twice, this time with widths w set to 50 and 60, corresponding to a number of parameters equal to 44860 and 55030 respectively. This corresponds to points on the right of curve on fig. 3, where implicit regularization occurs the most. We keep the sizes of X_{train} and X_{test} to 6000 and 20000 respectively, and we increase the size of X_{random} to 40000. The plots associated with this experiment correspond to the right side of Figure 7.

We observe similar train and test errors as in 7.5.1. We also observe similarly that $\text{rank}(\mathbf{D}f_{\theta}(X_{\text{random}}))$ is often equal to the maximum rank (but note that we increase the size of X_{random} here).

For these two widths, the increase of both quantities $\text{rank}(\mathbf{D}f_{\theta}(X_{\text{train}}))$ and $\text{rank}(\mathbf{D}f_{\theta}(X_{\text{test}}))$ is clearer, in particular in the case of X_{train} .




Figure 7: Behavior of different complexity measures as some corrupted labels are introduced in the training set, tested for three network widths. The values of the different ranks are to be read on the left axis, titled ‘functional dimension’. The values of the test and train errors are to be read on the right axis. Width of the network from left to right: $w = 30$, $w = 50$ and $w = 60$.

Overall, the experiments 7.5.1 and 7.5.2 indicate that there is a positive correlation between the functional dimensions and the complexity of the function learned by the neural network.

8 Conclusion and perspectives

In this article, we describe the local geometry of deep ReLU neural networks. The study shows that the image of a sample X by deep ReLU neural networks of a fixed architecture is a set whose local dimension varies. The local dimension is called the *batch functional dimension* by Grigsby u. a. (2022). We show that the parameter space is divided into pieces where the batch functional dimension is fixed. Empirically, the pieces of small dimensions are on the outside of the ones of large dimensions. They are favored by the optimization. We call this phenomenon *geometry-induced implicit regularization*. When X is allowed to vary, we also study the maximal dimension over all possible X . We call it the *computable full functional dimension*. Both notions of local complexity are determined by the activation patterns. We investigate the practical computation of the functional dimensions and provide experiments emphasizing the *geometry-induced implicit regularization* and the link between functional dimensions and the complexity of the learning task.

This opens up many perspectives in deep learning theory. The formal connection between the notions of local complexity and the generalization gap is still missing. It would permit us to obtain a theory that explains the good performance of deep learning. It would be interesting to study more systematically how the functional dimensions of the learned parameters depend on the distribution of the learned phenomenon. To do so, it would be

interesting to study instances in larger dimensions. Algorithms of a better complexity for computing the functional dimensions are needed. In particular, since we have proved that the batch functional dimension is almost-surely determined by the activation patterns, it would be interesting to compute the former using the activation patterns instead of the gradients.

Acknowledgement

This work has benefited from the AI Interdisciplinary Institute ANITI. ANITI is funded by the French “Investing for the Future – PIA3” program under the Grant agreement n°ANR-19-PI3A-0004.

The authors gratefully acknowledge the support of the DEEL project.⁹

⁹<https://www.deel.ai/>

A Proofs of Section 3

This appendix is devoted to the proofs of Section 3. In Section A.1, we prove Theorem 1, and in Section A.2 we prove Proposition 2.

A.1 Proof of Theorem 1

Let us define, for $\ell \in \llbracket 1, L-1 \rrbracket$ and $v \in V_\ell$ the set

$$\mathcal{T}_v = \left\{ (x, \theta) \in \mathbb{R}^{N_0} \times (\mathbb{R}^E \times \mathbb{R}^B) \mid \sum_{v' \in V_{\ell-1}} w_{v' \rightarrow v} \left(f_\theta^{\ell-1}(x) \right)_{v'} + b_v = 0 \right\}, \quad (24)$$

and let

$$\mathcal{T} = \bigcup_{\ell=1}^{L-1} \bigcup_{v \in V_\ell} \mathcal{T}_v. \quad (25)$$

Similarly, for any $x \in \mathbb{R}^{N_0}$, $\ell \in \llbracket 1, L-1 \rrbracket$ and $v \in V_\ell$, we define the set

$$\mathcal{T}_v^x = \left\{ \theta \in \mathbb{R}^E \times \mathbb{R}^B \mid \sum_{v' \in V_{\ell-1}} w_{v' \rightarrow v} \left(f_\theta^{\ell-1}(x) \right)_{v'} + b_v = 0 \right\},$$

and let

$$\mathcal{T}^x = \bigcup_{\ell=1}^{L-1} \bigcup_{v \in V_\ell} \mathcal{T}_v^x.$$

Lemma 12. (i) – Over $\mathbb{R}^{N_0} \times (\mathbb{R}^E \times \mathbb{R}^B)$, the function $(x, \theta) \mapsto a(x, \theta) \in \{0, 1\}^{N_1 + \dots + N_{L-1}}$ takes exactly $2^{N_1 + \dots + N_{L-1}}$ distinct values.

– For any $\delta \in \{0, 1\}^{N_1 + \dots + N_{L-1}}$, we write

$$A_\delta = \{(x, \theta) \in \mathbb{R}^{N_0} \times (\mathbb{R}^E \times \mathbb{R}^B) \mid a(x, \theta) = \delta\}, \quad (26)$$

which is thus nonempty. Then: On A_δ , the function $(x, \theta) \mapsto f_\theta(x)$ is analytic.

– The set \mathcal{T} has Lebesgue measure zero and $\bigcup_{\delta \in \{0, 1\}^{N_1 + \dots + N_{L-1}}} \partial A_\delta = \mathcal{T}$. Therefore, for any $\delta \in \{0, 1\}^{N_1 + \dots + N_{L-1}}$, ∂A_δ is a closed set of Lebesgue measure zero in $\mathbb{R}^{N_0} \times (\mathbb{R}^E \times \mathbb{R}^B)$.

(ii) – For any fixed $x \in \mathbb{R}^{N_0}$, the function $\theta \mapsto a(x, \theta)$ exactly takes $2^{N_1 + \dots + N_{L-1}}$ distinct values.

– For $x \in \mathbb{R}^{N_0}$ and $\delta \in \{0, 1\}^{N_1 + \dots + N_{L-1}}$, we write

$$A_\delta^x = \{\theta \in \mathbb{R}^E \times \mathbb{R}^B \mid a(x, \theta) = \delta\}.$$

Then: On A_δ^x , the function $\theta \mapsto f_\theta(x)$ is analytic.

– The set \mathcal{T}^x has Lebesgue measure zero and $\bigcup_{\delta \in \{0, 1\}^{N_1 + \dots + N_{L-1}}} \partial A_\delta^x = \mathcal{T}^x$. Therefore, for any $\delta \in \{0, 1\}^{N_1 + \dots + N_{L-1}}$, ∂A_δ^x is a closed set with Lebesgue measure zero in $\mathbb{R}^E \times \mathbb{R}^B$.

Proof of Lemma 12. To avoid repetitions, we only detail the proof of (i). The proof of (ii) is very similar, considering functions of θ only (with x fixed) rather than (x, θ) .

We first prove the first item, i.e. we prove that all activation patterns are reached. The set $\{0, 1\}^{N_1 + \dots + N_{L-1}}$ is finite and its cardinal is $2^{N_1 + \dots + N_{L-1}}$. Observe that for any $\delta \in \{0, 1\}^{N_1 + \dots + N_{L-1}}$, by taking $\theta \in \mathbb{R}^E \times \mathbb{R}^B$ such that $w_{v \rightarrow v'} = 0$ for any $(v \rightarrow v') \in E$, $b_v = 0$ for $v \in V_L$ and $b_v = (-1)^{1 + \delta_v}$ for any $v \in V_1 \cup \dots \cup V_{L-1}$, then, for any $x \in \mathbb{R}^{N_0}$ and any $v \in V_1 \cup \dots \cup V_{L-1}$, we have $a_v(x, \theta) = \delta_v$, i.e. $a(x, \theta) = \delta$.

In order to prove the second item, i.e. that the function $(x, \theta) \mapsto f_\theta(x)$ is analytic, we remind the definition of f_θ^ℓ , in (1), and we define

$$a_{\leq \ell}(x, \theta) = \begin{cases} (a_v(x, \theta))_{v \in V_1 \cup \dots \cup V_\ell} & \text{if } \ell \geq 1, \\ 1 & \text{if } \ell = 0. \end{cases}$$

We prove by induction that the assertion

$$H_\ell : \begin{cases} \forall D \subseteq \mathbb{R}^{N_0} \times (\mathbb{R}^E \times \mathbb{R}^B), \text{ if } (x, \theta) \mapsto a_{\leq \ell}(x, \theta) \text{ is constant on } D, \text{ then} \\ (x, \theta) \mapsto f_\theta^\ell(x) \text{ is polynomial on } D \end{cases}$$

holds, for all $\ell \in \llbracket 0, L-1 \rrbracket$.

The assertion H_0 indeed holds because $f_\theta^0(x) = x$ is polynomial in (x, θ) (of degree 1) on any subset of $\mathbb{R}^{N_0} \times (\mathbb{R}^E \times \mathbb{R}^B)$. Assume now that for some $\ell \in \llbracket 1, L-1 \rrbracket$, $H_{\ell-1}$ holds, and let us prove H_ℓ .

Let $D \subseteq \mathbb{R}^{N_0} \times (\mathbb{R}^E \times \mathbb{R}^B)$ such that $a_{\leq \ell}(x, \theta)$ is constant on D . For $(x, \theta) \in D$ and $v \in V_\ell$, using (6), we have

$$(f_\theta^\ell(x))_v = a_v(x, \theta) \left(\sum_{v' \in V_{\ell-1}} w_{v' \rightarrow v} (f_\theta^{\ell-1}(x))_{v'} + b_v \right).$$

The quantity $a_{\leq \ell-1}(x, \theta)$ is constant on D and thus from $H_{\ell-1}$, for all $v' \in V_{\ell-1}$, $(x, \theta) \mapsto (f_\theta^{\ell-1}(x))_{v'}$ is a polynomial function of (x, θ) on D . Since $a_v(x, \theta)$ is constant on D , $(f_\theta^\ell(x))_v$ is a polynomial function of (x, θ) . This concludes the proof by induction that H_ℓ holds for all $\ell \in \llbracket 0, L-1 \rrbracket$.

If we recall from (1) that $y_\theta^L(x) \in \mathbb{R}^{N_L}$ is the vector satisfying, for all $v \in V_L$,

$$(y_\theta^L(x))_v = \sum_{v' \in V_{L-1}} w_{v' \rightarrow v} (f_\theta^{L-1}(x))_{v'} + b_v,$$

we have

$$f_\theta(x) = \sigma_L(y_\theta^L(x)).$$

We recall the notations $A_\delta, \delta \in \{0, 1\}^{N_1 + \dots + N_{L-1}}$, in (26). For $\delta \in \{0, 1\}^{N_1 + \dots + N_{L-1}}$, $a_{\leq L-1}(x, \theta) = a(x, \theta)$ is constant on A_δ and thus from H_{L-1} , $(x, \theta) \mapsto f_\theta^{L-1}(x)$ is polynomial on A_δ . As a consequence, $(x, \theta) \mapsto y_\theta^L(x)$ is polynomial on A_δ , and since σ_L is analytic, $(x, \theta) \mapsto f_\theta(x)$ is analytic on A_δ . This proves the second item of Lemma 12, (i).

Let us now show the third item, which states that \mathcal{T} has Lebesgue measure zero. For that, let us show that for all $\ell \in \llbracket 1, L-1 \rrbracket$ and $v \in V_\ell$, \mathcal{T}_v has Lebesgue measure zero. To do so, since $\cup_\delta A_\delta = \mathbb{R}^{N_0} \times (\mathbb{R}^E \times \mathbb{R}^B)$, we consider $\ell \in \llbracket 1, L-1 \rrbracket$ and $v \in V_\ell$, and prove that, for all $\delta \in \{0, 1\}^{N_1 + \dots + N_{L-1}}$, $\mathcal{T}_v \cap A_\delta$ has Lebesgue measure zero. For $\delta \in \{0, 1\}^{N_1 + \dots + N_{L-1}}$, $a_{\leq \ell-1}(x, \theta)$ is constant on A_δ and thus from $H_{\ell-1}$, $(x, \theta) \mapsto f_\theta^{\ell-1}(x)$ is a polynomial function of (x, θ) on A_δ and thus $\sum_{v' \in V_{\ell-1}} w_{v' \rightarrow v} \left(f_\theta^{\ell-1}(x) \right)_{v'} + b_v$ also is. Since the variable $(x, \theta) \mapsto b_v$ is not present in the expression of $f_\theta^{\ell-1}(x)$, it only appears in a single monomial of degree and coefficient 1 of $\sum_{v' \in V_{\ell-1}} w_{v' \rightarrow v} \left(f_\theta^{\ell-1}(x) \right)_{v'} + b_v$. The latter polynomial function is therefore non-constant. Hence the set $\mathcal{T}_v \cap A_\delta$, constituted by the zeros of this polynomial function, has Lebesgue measure zero. Since $\cup_\delta A_\delta = \mathbb{R}^{N_0} \times (\mathbb{R}^E \times \mathbb{R}^B)$, we finally conclude that, for any $\ell \in \llbracket 1, L-1 \rrbracket$ and $v \in V_\ell$, \mathcal{T}_v has Lebesgue measure zero.

The set

$$\mathcal{T} = \cup_{\ell=1}^{L-1} \cup_{v \in V_\ell} \mathcal{T}_v$$

is thus also of Lebesgue measure zero.

Let us now prove the set equality:

$$\bigcup_{\delta} \partial A_\delta = \mathcal{T}. \quad (27)$$

We first show the inclusion $\bigcup_{\delta} \partial A_\delta \subseteq \mathcal{T}$. Consider $\delta \in \{0, 1\}^{N_1 + \dots + N_{L-1}}$ and let us now show that $\partial A_\delta \subseteq \mathcal{T}$. To do so, consider $(x, \theta) \in \partial A_\delta$. Since $(x, \theta) \notin \text{Int}(A_\delta)$, for any ε there exists $\delta_\varepsilon \neq \delta$ such that $B((x, \theta), \varepsilon) \cap A_{\delta_\varepsilon} \neq \emptyset$. Since the set of all possible δ_ε is finite, we are sure that there exists $\delta' \neq \delta$ such that $(x, \theta) \in \overline{A_{\delta'}}$. Let $\ell \in \llbracket 1, L-1 \rrbracket$ and $v \in V_\ell$ such that $\delta_v \neq \delta'_v$. We assume without loss of generality that $\delta_v = 0$. The proof is indeed similar when $\delta_v = 1$. There exists $(x_n, \theta_n) \in A_{\delta'}^{\mathbb{N}^*}$ such that $(x_n, \theta_n) \rightarrow (x, \theta)$ as $n \rightarrow \infty$ and there exists $(x'_n, \theta'_n) \in A_{\delta}^{\mathbb{N}^*}$ such that $(x'_n, \theta'_n) \rightarrow (x, \theta)$ as $n \rightarrow \infty$. We have $a_v(x_n, \theta_n) = 0$ and $a_v(x'_n, \theta'_n) = 1$ for all n .

Using that $(x, \theta) \mapsto \sum_{v' \in V_{\ell-1}} w_{v' \rightarrow v} \left(f_\theta^{\ell-1}(x) \right)_{v'} + b_v$ is continuous and taking the limit of this function at (x_n, θ_n) , as n goes to infinity, we obtain that $\sum_{v' \in V_{\ell-1}} w_{v' \rightarrow v} \left(f_\theta^{\ell-1}(x) \right)_{v'} + b_v \leq 0$. Reasoning similarly with the sequence $(x'_n, \theta'_n)_{n \in \mathbb{N}^*}$ we obtain the reverse inequality and conclude that $\sum_{v' \in V_{\ell-1}} w_{v' \rightarrow v} \left(f_\theta^{\ell-1}(x) \right)_{v'} + b_v = 0$. This shows that $(x, \theta) \in \mathcal{T}_v \subseteq \mathcal{T}$. This finishes the proof of $\partial A_\delta \subseteq \mathcal{T}$.

Let us now show the reciprocal inclusion $\mathcal{T} \subseteq \bigcup_{\delta} \partial A_\delta$. Indeed, let $(x, \theta) \in \mathcal{T}$. There exist $\ell \in \llbracket 1, L-1 \rrbracket$ and $v \in V_\ell$ such that $(x, \theta) \in \mathcal{T}_v$. There also exists $\delta \in \{0, 1\}^{N_1 + \dots + N_{L-1}}$ such that $(x, \theta) \in A_\delta$. In particular, since $\sum_{v' \in V_{\ell-1}} w_{v' \rightarrow v} \left(f_\theta^{\ell-1}(x) \right)_{v'} + b_v = 0$, we have $\delta_v = a_v(x, \theta) = 1$. For any $\varepsilon > 0$, by replacing b_v by $b_v - \varepsilon$, we obtain a θ_ε satisfying $\|\theta - \theta_\varepsilon\| \leq \varepsilon$ and $a_v(x, \theta_\varepsilon) = 0 \neq \delta_v$, which shows $(x, \theta_\varepsilon) \notin A_\delta$. This shows $(x, \theta) \in \partial A_\delta \subseteq \bigcup_{\delta} \partial A_\delta$. This shows the desired inclusion, and thus the equality (27).

For all $\delta \in \{0, 1\}^{N_1 + \dots + N_{L-1}}$, ∂A_δ is closed by definition of a boundary. Since \mathcal{T} has been shown to have Lebesgue measure zero, then ∂A_δ has Lebesgue measure zero and thus the proof of the part (i) is concluded. \square

We state and prove another lemma before proving Theorem 1. The lemma resembles Theorem 1 but does not include the statements on $\text{rank}(\mathbf{D}f_\theta(X))$.

For $n \in \mathbb{N}^*$, recalling the definition \mathcal{T} in (25), let us now define

$$\mathcal{T}_n = \cup_{i=1}^n \left\{ (X, \theta) \in \mathbb{R}^{N_0 \times n} \times (\mathbb{R}^E \times \mathbb{R}^B) \mid (x^{(i)}, \theta) \in \mathcal{T} \right\}. \quad (28)$$

Similarly, for $n \in \mathbb{N}^*$ and $X \in \mathbb{R}^{N_0 \times n}$, we define

$$\mathcal{T}^X = \cup_{i=1}^n \left\{ \theta \in \mathbb{R}^E \times \mathbb{R}^B \mid (x^{(i)}, \theta) \in \mathcal{T} \right\}.$$

Lemma 13. (i) For all $n \in \mathbb{N}^*$, the sets $\tilde{\mathcal{O}}_1^n, \dots, \tilde{\mathcal{O}}_{m_n}^n$ defined in (7) are non-empty, open and disjoint, and they satisfy

- $(\cup_{j=1}^{m_n} \tilde{\mathcal{O}}_j^n)^c = \mathcal{T}_n$, and in particular the complement $(\cup_{j=1}^{m_n} \tilde{\mathcal{O}}_j^n)^c$ is a closed set with Lebesgue measure zero;
- for all $j \in \llbracket 1, m_n \rrbracket$, the function $(X, \theta) \mapsto a(X, \theta)$ is constant on each $\tilde{\mathcal{O}}_j^n$ and takes m_n distinct values on $\cup_{j=1}^{m_n} \tilde{\mathcal{O}}_j^n$;
- for all $j \in \llbracket 1, m_n \rrbracket$, $(X, \theta) \mapsto f_\theta(X)$ is an analytic function on $\tilde{\mathcal{O}}_j^n$.

(ii) For all $n \in \mathbb{N}^*$, for all $X \in \mathbb{R}^{N_0 \times n}$, the sets $\tilde{\mathcal{U}}_1^X, \dots, \tilde{\mathcal{U}}_{p_X}^X$ defined in (10) are non-empty, open and disjoint, and they satisfy

- $(\cup_{j=1}^{p_X} \tilde{\mathcal{U}}_j^X)^c = \mathcal{T}^X$, and in particular the complement $(\cup_{j=1}^{p_X} \tilde{\mathcal{U}}_j^X)^c$ is a closed set with Lebesgue measure zero;
- for all $j \in \llbracket 1, p_X \rrbracket$, the function $\theta \mapsto a(X, \theta)$ is constant on each $\tilde{\mathcal{U}}_j^X$ and takes p_X distinct values on $\cup_{j=1}^{p_X} \tilde{\mathcal{U}}_j^X$;
- for all $j \in \llbracket 1, p_X \rrbracket$, $\theta \mapsto f_\theta(X)$ is an analytic function on $\tilde{\mathcal{U}}_j^X$.

Proof of Lemma 13. As in the proof of Lemma 12, the proofs of (i) and (ii) are very similar. To avoid repetitions, we only detail the proof of (i).

By definition, see (7), the sets $\tilde{\mathcal{O}}_1^n, \dots, \tilde{\mathcal{O}}_{m_n}^n$ are non-empty, open and disjoint. Before proving the first point of (i), let us notice that \mathcal{T}_n is of Lebesgue measure zero. Indeed, the third item of Lemma 12 states that \mathcal{T} has Lebesgue measure zero, and therefore, we see in the definition (28) that \mathcal{T}_n is a finite union of zero-measured sets, which shows that it itself has Lebesgue measure zero. Let us also write the following useful characterization: thanks to the characterization of \mathcal{T} in the third item of Lemma 12 (i), we have

$$\mathcal{T}_n = \cup_{i=1}^n \cup_{\delta \in \{0,1\}^{N_1 + \dots + N_{L-1}}} \left\{ (X, \theta) \in \mathbb{R}^{N_0 \times n} \times (\mathbb{R}^E \times \mathbb{R}^B) \mid (x^{(i)}, \theta) \in \partial A_\delta \right\}. \quad (29)$$

Let us now show the first point of (i). Let us prove that $(\cup_{j=1}^{m_n} \tilde{\mathcal{O}}_j^n)^c = \mathcal{T}_n$.

To do so, let us first show that $(\cup_{j=1}^{m_n} \tilde{\mathcal{O}}_j^n)^c \subseteq \mathcal{T}_n$. Let $(X, \theta) \in (\cup_{j=1}^{m_n} \tilde{\mathcal{O}}_j^n)^c$. Consider the $\Delta_1^n, \dots, \Delta_q^n$ defined just before (7). There exists $j \in \llbracket 1, q \rrbracket$ such that $a(X, \theta) = \Delta_j^n$. Since

$(X, \theta) \notin \widetilde{\mathcal{O}}_j^n$, there exists a sequence $(X_k, \theta_k)_{k \in \mathbb{N}^*}$ such that $(X_k, \theta_k) \rightarrow (X, \theta)$, as $k \rightarrow \infty$ and $a(X_k, \theta_k) \neq \Delta_j^n$, for all k . Modulo the extraction of a sub-sequence, we can assume that there exist $i \in \llbracket 1, n \rrbracket$ such that for all $k \in \mathbb{N}^*$, $a(x_k^{(i)}, \theta_k) \neq \delta$, where $x_k^{(i)}$ is the i^{th} column of X_k , and where we denote $\delta = (\Delta_j^n)_i$. Thus, we have $(x_k^{(i)}, \theta_k) \notin A_\delta$, for all k , and since $(x^{(i)}, \theta) \in A_\delta$, we conclude $(x^{(i)}, \theta) \in \partial A_\delta$. The characterization (29) thus shows $(X, \theta) \in \mathcal{T}_n$. This shows $(\cup_{j=1}^{m_n} \widetilde{\mathcal{O}}_j^n)^c \subseteq \mathcal{T}_n$.

Let us now show that $\mathcal{T}_n \subseteq (\cup_{j=1}^{m_n} \widetilde{\mathcal{O}}_j^n)^c$. If $(X, \theta) \in \mathcal{T}_n$, there exists $i \in \llbracket 1, n \rrbracket$ and $\delta \in \{0, 1\}^{N_1 + \dots + N_{L-1}}$ such that $(x^{(i)}, \theta) \in \partial A_\delta$. Thus, for any $\varepsilon > 0$, $(x, \theta') \mapsto a(x, \theta')$ is not constant over $B((x^{(i)}, \theta), \varepsilon)$. As a consequence, (X, θ) does not belong to any of the open sets $\widetilde{\mathcal{O}}_j^n$.

This shows $(\cup_{j=1}^{m_n} \widetilde{\mathcal{O}}_j^n)^c = \mathcal{T}_n$, and thus, since \mathcal{T}_n is of Lebesgue measure zero, $(\cup_{j=1}^{m_n} \widetilde{\mathcal{O}}_j^n)^c$ has Lebesgue measure zero. Adding that the complement of an open set is closed, $(\cup_{j=1}^{m_n} \widetilde{\mathcal{O}}_j^n)^c$ is closed, which ends the proof of the first point of (i).

The second point of (i) holds by definition of $\widetilde{\mathcal{O}}_1^n, \dots, \widetilde{\mathcal{O}}_{m_n}^n$.

Let us now show the third point of (i). Let $j \in \llbracket 1, m_n \rrbracket$. The function $(X, \theta) \mapsto a(X, \theta)$ is constant on $\widetilde{\mathcal{O}}_j^n$. The set $\widetilde{\mathcal{O}}_j^n$ is associated to Δ_j^n in (7) and the latter is of the form $(\delta^1, \dots, \delta^n)$ with $\delta^i \in \{0, 1\}^{N_1 + \dots + N_{L-1}}$ for $i \in \llbracket 1, n \rrbracket$. Fix $i' \in \llbracket 1, n \rrbracket$. Then for $X = (x^{(i)})_{i \in \llbracket 1, n \rrbracket}$ with $(X, \theta) \in \widetilde{\mathcal{O}}_j^n$, $(x^{(i')}, \theta) \in A_{\delta^{i'}}$. Hence, Lemma 12 (i) shows that $f_\theta(x^{(i')})$ is an analytic function of $(x^{(i')}, \theta)$ and thus of (X, θ) . The quantity $f_\theta(X)$ is a matrix whose columns are $f_\theta(x^{(i)})$, $i \in \llbracket 1, n \rrbracket$. Hence $(X, \theta) \mapsto f_\theta(X)$ is an analytic function on $\widetilde{\mathcal{O}}_j^n$, which concludes the proof of (i). \square

Proof of Theorem 1. Again, the proofs of (i) and (ii) are very similar and we only detail the proof of (i).

Consider $n \in \mathbb{N}^*$. The sets $\mathcal{O}_1^n, \dots, \mathcal{O}_{m_n}^n$ are non-empty by definition of $r_1^n, \dots, r_{m_n}^n$, and they are disjoint because of the inclusion $\mathcal{O}_j^n \subseteq \widetilde{\mathcal{O}}_j^n$ for all j , and because the sets $\widetilde{\mathcal{O}}_1^n, \dots, \widetilde{\mathcal{O}}_{m_n}^n$ are disjoint as shown in Lemma 13. Hence the first item holds. The second item is a direct consequence of the definition of \mathcal{O}_j^n , in (9). The third item holds by definition.

To see that \mathcal{O}_j^n is open, first recall that $\widetilde{\mathcal{O}}_j^n$ is open, then note that since the function $(X, \theta) \mapsto f_\theta(X)$ is analytic over $\widetilde{\mathcal{O}}_j^n$ (by Lemma 13 (i)), the function $(X, \theta) \mapsto \mathbf{D}f_\theta(X)$ is continuous over $\widetilde{\mathcal{O}}_j^n$. Since the rank is lower semicontinuous, if $\text{rank}(\mathbf{D}f_\theta(X)) = r_j^n$, then there exists $\varepsilon > 0$ such that for any $(X', \theta') \in B((X, \theta), \varepsilon)$, we have $\text{rank}(\mathbf{D}f_{\theta'}(X')) \geq r_j^n$, which by maximality of r_j^n , is equivalent to $\text{rank}(\mathbf{D}f_{\theta'}(X')) = r_j^n$ and to $(X', \theta') \in \mathcal{O}_j^n$. This shows that \mathcal{O}_j^n is open. Hence Item 4 holds.

Item 6 comes directly from Lemma 13 and from the inclusion $\mathcal{O}_j^n \subseteq \widetilde{\mathcal{O}}_j^n$.

To finish the proof, we need to prove Item 5, stating that $(\cup_{j=1}^{m_n} \mathcal{O}_j^n)^c$ is a closed set with Lebesgue measure zero. Let us consider a basis $(e_1, \dots, e_{|E|+|B|})$ of $\mathbb{R}^E \times \mathbb{R}^B$ and a basis $(\varepsilon_1, \dots, \varepsilon_{nN_L})$ of $\mathbb{R}^{N_L \times n}$. For all X , let us write $M_\theta(X)$ for the matrix of the differential $\mathbf{D}f_\theta(X)$ of the function $\theta \mapsto f_\theta(X)$ in these two bases. Then $(X, \theta) \mapsto M_\theta(X)$ is an

analytic function on $\tilde{\mathcal{O}}_j^n$. Recall the notation $r_j^n = \max_{(X,\theta) \in \tilde{\mathcal{O}}_j^n} \text{rank}(\mathbf{D}f_\theta(X))$, and let $(X', \theta') \in \tilde{\mathcal{O}}_j^n$ such that $\text{rank}(\mathbf{D}f_{\theta'}(X')) = r_j^n$. We thus have $\text{rank}(M_{\theta'}(X')) = r_j^n$, and thus there exists a sub-matrix $N_{\theta'}(X')$ of $M_{\theta'}(X')$, of size $r_j^n \times r_j^n$, such that $\det N_{\theta'}(X') \neq 0$. The function $(X, \theta) \mapsto M_\theta(X)$ is an analytic function on $\tilde{\mathcal{O}}_j^n$ and thus $(X, \theta) \mapsto \det(N_\theta(X))$ also is. This latter function is not uniformly zero on $\tilde{\mathcal{O}}_j^n$ and thus the set of its zeros, that we write \mathcal{Y}_j , is a closed set of Lebesgue measure zero (Mityagin, 2020).

For all $(X, \theta) \in \tilde{\mathcal{O}}_j^n \setminus \mathcal{Y}_j$, we have $\det N_\theta(X) \neq 0$ and thus $\text{rank}(N_\theta(X)) = r_j^n$ and thus $\text{rank}(M_\theta(X)) \geq r_j^n$. We also have $\text{rank}(M_\theta(X)) = \text{rank}(\mathbf{D}f_\theta(X)) \leq r_j^n$ by definition of r_j^n . Hence $\text{rank}(\mathbf{D}f_\theta(X)) = r_j^n$. This shows $\tilde{\mathcal{O}}_j^n \setminus \mathcal{Y}_j \subseteq \mathcal{O}_j^n$.

Finally,

$$\begin{aligned} \left(\bigcup_{j=1}^{m_n} \mathcal{O}_j^n\right)^c &= \bigcap_{j=1}^{m_n} (\mathcal{O}_j^n)^c \\ &\subseteq \bigcap_{j=1}^{m_n} (\tilde{\mathcal{O}}_j^n \setminus \mathcal{Y}_j)^c \\ &= \bigcap_{j=1}^{m_n} \left((\tilde{\mathcal{O}}_j^n)^c \cup \mathcal{Y}_j \right) \\ &\subseteq \bigcap_{j=1}^{m_n} \left((\tilde{\mathcal{O}}_j^n)^c \cup \left(\bigcup_{j'=1}^{m_n} \mathcal{Y}_{j'}\right) \right) \\ &= \left(\bigcap_{j=1}^{m_n} (\tilde{\mathcal{O}}_j^n)^c\right) \cup \left(\bigcup_{j=1}^{m_n} \mathcal{Y}_j\right) \\ &= \left(\bigcup_{j=1}^{m_n} \tilde{\mathcal{O}}_j^n\right)^c \cup \left(\bigcup_{j=1}^{m_n} \mathcal{Y}_j\right). \end{aligned}$$

We know from Lemma 13 that $\left(\bigcup_{j=1}^{m_n} \tilde{\mathcal{O}}_j^n\right)^c$ has Lebesgue measure zero. Also each \mathcal{Y}_j has Lebesgue measure zero, thus $\bigcup_{j=1}^{m_n} \mathcal{Y}_j$ has Lebesgue measure zero. Hence, $\left(\bigcup_{j=1}^{m_n} \mathcal{O}_j^n\right)^c$ has Lebesgue measure zero, which concludes the proof of (i) in the theorem. \square

A.2 Proof of Proposition 2

Let us prove more than Proposition 2. We in fact prove the following proposition. Proposition 2 simply corresponds to the first item, but it is convenient to prove the others items at the same time, which will be useful in the proofs of other results.

Proposition 14. *Consider any deep fully-connected ReLU network architecture (E, V, σ_L) . Let $\theta, \tilde{\theta} \in \mathbb{R}^E \times \mathbb{R}^B$. Assume $\theta \sim \tilde{\theta}$. Then for any $n \in \mathbb{N}^*$ and $X \in \mathbb{R}^{N_0 \times n}$:*

1. *if the differential $\mathbf{D}f_\theta(X)$ is well-defined, then the differential $\mathbf{D}f_{\tilde{\theta}}(X)$ is well-defined, and in that case we have*

$$\text{rank}(\mathbf{D}f_{\tilde{\theta}}(X)) = \text{rank}(\mathbf{D}f_\theta(X));$$

2. *if there exists $j \in \llbracket 1, m_n \rrbracket$, such that $(X, \theta) \in \tilde{\mathcal{O}}_j^n$, then there exists $j' \in \llbracket 1, m_n \rrbracket$ such that $(X, \tilde{\theta}) \in \tilde{\mathcal{O}}_{j'}^n$;*

3. if there exists $j \in \llbracket 1, m_n \rrbracket$, such that $(X, \theta) \in \mathcal{O}_j^n$, then there exists $j' \in \llbracket 1, m_n \rrbracket$ such that $(X, \tilde{\theta}) \in \mathcal{O}_{j'}^n$.

Proof. Let $\tilde{\theta} \sim \theta$, let $n \in \mathbb{N}^*$ and let $X \in \mathbb{R}^{N_0 \times n}$. Let us prove the three items of the proposition.

Let us prove the first item. By definition of the relation \sim_j in Section 2, there is an invertible linear map $M : \mathbb{R}^E \times \mathbb{R}^B \rightarrow \mathbb{R}^E \times \mathbb{R}^B$ such that $\tilde{\theta} = M\theta$. Note that when expressed in the canonical basis of $\mathbb{R}^E \times \mathbb{R}^B$, the matrix corresponding to M is the product of a permutation matrix and a diagonal matrix, with strictly positive diagonal components whose values are given by (4) and (5). Notice that since M corresponds to positive rescalings and neuron permutations, as discussed after (5), we have,

$$\text{for any } \theta' \in \mathbb{R}^E \times \mathbb{R}^B, \quad f_{\theta'}(X) = f_{M\theta'}(X). \quad (30)$$

Assume that $\mathbf{D}f_\theta(X)$ is well-defined, i.e. the map $\theta' \mapsto f_{\theta'}(X)$ is differentiable at θ . Then, for all $u \in \mathbb{R}^E \times \mathbb{R}^B$, the following calculation holds, using the fact that M is invertible, using (30) and using (3),

$$\begin{aligned} f_{\tilde{\theta}+u}(X) &= f_{M\theta+u}(X) = f_{M(\theta+M^{-1}u)}(X) \\ &= f_{\theta+M^{-1}u}(X) \\ &= f_\theta(X) + \mathbf{D}f_\theta(X)(M^{-1}u) + o(\|M^{-1}u\|) \\ &= f_\theta(X) + \mathbf{D}f_\theta(X)(M^{-1}u) + o(\|u\|). \end{aligned}$$

Hence, $\theta' \mapsto f_{\theta'}(X)$ is differentiable at $\tilde{\theta}$ and for all $u \in \mathbb{R}^E \times \mathbb{R}^B$,

$$\mathbf{D}f_{\tilde{\theta}}(X)(u) = \mathbf{D}f_\theta(X)(M^{-1}u).$$

Since M^{-1} is invertible, it follows that $\text{rank}(\mathbf{D}f_{\tilde{\theta}}(X)) = \text{rank}(\mathbf{D}f_\theta(X))$. This concludes the proof of the first item.

Let us now prove the second and the third items. Let $j \in \llbracket 1, m_n \rrbracket$ such that $(X, \theta) \in \tilde{\mathcal{O}}_j^n$. Again, let us consider the map M defined above. For any $\theta' \in \mathbb{R}^E \times \mathbb{R}^B$, it is well-known (see for instance Bona-Pellissier u. a. 2023a, Proposition 39, Item 2) that since θ' and $M\theta'$ are equivalent modulo permutation and positive rescaling, for any $x \in \mathbb{R}^{N_0}$, we have $a(x, M\theta') = Pa(x, \theta')$, for some permutation matrix $P \in \{0, 1\}^{(N_1+\dots+N_{L-1}) \times (N_1+\dots+N_{L-1})}$ that only depends on M and not on x .

Since $(X, \theta) \in \tilde{\mathcal{O}}_j^n$, recalling the definition (7), there exists an open neighborhood \tilde{U} of (X, θ) such that for any $(X', \theta') \in \tilde{U}$, we have $a(X', \theta') = \Delta_j^n$. For any $(X', \theta') \in \tilde{U}$, we have $a(X', M\theta') = Pa(X', \theta') = P\Delta_j^n$. Furthermore, since M is invertible, the map \bar{M} defined by $(X', \theta') \mapsto (X', M\theta')$ is invertible, and the image \tilde{V} of \tilde{U} by \bar{M} is thus an open neighborhood of $(X, \tilde{\theta})$. This shows that there exists $j' \in \llbracket 1, m_n \rrbracket$ such that $\Delta_{j'}^n = P\Delta_j^n$ and $(X, \tilde{\theta}) \in \tilde{V} \subseteq \tilde{\mathcal{O}}_{j'}^n$. This proves the second item.

Finally, if we furthermore assume $(X, \theta) \in \mathcal{O}_j^n$, applying Theorem 1, we see that the function $(X', \theta') \mapsto \text{rank}(\mathbf{D}f_{\theta'}(X'))$ is locally constant at (X, θ) : there exists an open

neighborhood of (X, θ) , $U \subseteq \widetilde{U}$, such that $U \subseteq \mathcal{O}_j^n$ and thus, for all $(X', \theta') \in U$, we have $\text{rank}(\mathbf{D}f_{\theta'}(X')) = r_j^n$. Thus, applying the invertible map \overline{M} , which preserves the rank as shown in the first item, this shows that $(X', \theta') \mapsto \text{rank}(\mathbf{D}f_{\theta'}(X'))$ is also constant on $V = \overline{M}(U)$ (note that $V \subseteq \widetilde{V}$), which is an open set and which contains $(X, \widetilde{\theta}) = \overline{M}(X, \theta)$. Since $V \subseteq \widetilde{\mathcal{O}}_j^n$, this shows that $\text{rank}(\mathbf{D}f_{\widetilde{\theta}}(X)) = r_j^n$ because r_j^n is the only value that can be taken on a subset of $\widetilde{\mathcal{O}}_j^n$ of non-zero Lebesgue measure. In other words, we have $(X, \widetilde{\theta}) \in \mathcal{O}_{j'}^n$ (and in particular we have shown $r_j^n = r_{j'}^n$). \square

B Calculations for the Example in Section 4.1

We provide in this appendix, the calculations permitting to construct Figure 1. We consider a one-hidden-layer neural network of widths $(1, 1, 1)$, with the identity activation function on the last layer. To simplify notations, we denote the weights and biases $\theta = (a, b, c, d) \in \mathbb{R}^4$ so that $f_\theta(x) = b\sigma(ax + c) + d$, for all $x \in \mathbb{R}$. We consider $X = (0, 1, 2) \in \mathbb{R}^{1 \times 3}$ and

$$f_\theta(X)^T = \begin{pmatrix} b\sigma(c) + d \\ b\sigma(a + c) + d \\ b\sigma(2a + c) + d \end{pmatrix}.$$

The boundaries of the sets $\widetilde{\mathcal{U}}_j^X$, corresponding to the parameters having the same activation pattern, are defined by the equation $c = 0$, $a + c = 0$ and $2a + c = 0$. There are 6 possible activation patterns corresponding to the zones represented, in the (a, c) plane, on the left of Figure 1.

Since the sets $f_{\widetilde{\mathcal{U}}_j^X}(X) = \{f_\theta(X) \mid \theta \in \widetilde{\mathcal{U}}_j^X\}$, for $j \in \llbracket 1, 6 \rrbracket$, are invariant to translations by vectors (d, d, d) , for $d \in \mathbb{R}$, we consider the plane \mathcal{P} orthogonal to the vector $(1, 1, 1)$ and parameterize its elements using the mapping

$$\begin{aligned} p: \mathbb{R}^2 &\longrightarrow \mathcal{P} \\ (x, y) &\longmapsto \frac{x}{\sqrt{6}}(1, 1, -2) + \frac{y}{\sqrt{2}}(-1, 1, 0). \end{aligned}$$

Instead of representing $f_{\widetilde{\mathcal{U}}_j^X}(X)$, we represent on the right of Figure 1 its intersection with \mathcal{P} , formally defined as the set $\mathcal{V}_j \subseteq \mathbb{R}^2$ such that

$$f_{\widetilde{\mathcal{U}}_j^X}(X) = \{p(x, y) + (z, z, z) \in \mathbb{R}^{1 \times 3} \mid (x, y) \in \mathcal{V}_j \text{ and } z \in \mathbb{R}\}.$$

Below, we construct the sets \mathcal{V}_j , for $j \in \llbracket 1, 6 \rrbracket$.

Case $j = 1$: We have $c \leq 0$, $2a + c \leq 0$ and therefore $a + c \leq 0$. This leads to $f_{(a,b,c,d)}(X) = (d, d, d)$ and $\mathcal{V}_1 = \{(0, 0)\}$.

Case $j = 2$: We have $c \geq 0$, $a + c \leq 0$ and therefore $2a + c \leq 0$. This leads to $f_{(a,b,c,d)}(X) = (bc + d, d, d)$ and

$$\mathcal{V}_2 = \{(x, y) \in \mathbb{R}^2 \mid \exists(a, b, c, d) \in \tilde{\mathcal{U}}_2^X, p(x, y) = (bc + d, d, d)\}.$$

Solving

$$\begin{cases} (1) : & bc + d = \frac{x}{\sqrt{6}} - \frac{y}{\sqrt{2}} \\ (2) : & d = \frac{x}{\sqrt{6}} + \frac{y}{\sqrt{2}} \\ (3) : & d = -2\frac{x}{\sqrt{6}} \end{cases} \iff \begin{cases} (1) - (2) : & -\sqrt{2}y = bc \\ \sqrt{2}((2) - (3)) : & y = -\sqrt{3}x \\ (3) : & d = -2\frac{x}{\sqrt{6}} \end{cases}$$

and we obtain

$$\mathcal{V}_2 = \{(x, y) \in \mathbb{R}^2 \mid \sqrt{3}x + y = 0\}.$$

Case $j = 3$: We have $c \geq 0$, $a + c \geq 0$ and $2a + c \leq 0$. This leads to $f_{(a,b,c,d)}(X) = (bc + d, b(a + c) + d, d)$ and

$$\mathcal{V}_3 = \{(x, y) \in \mathbb{R}^2 \mid \exists(a, b, c, d) \in \tilde{\mathcal{U}}_3^X, p(x, y) = (bc + d, b(a + c) + d, d)\}.$$

We have

$$\begin{cases} (1) : & bc + d = \frac{x}{\sqrt{6}} - \frac{y}{\sqrt{2}} \\ (2) : & b(a + c) + d = \frac{x}{\sqrt{6}} + \frac{y}{\sqrt{2}} \\ (3) : & d = -2\frac{x}{\sqrt{6}} \end{cases} \iff \begin{cases} \sqrt{2}((1) - (3)) : & \sqrt{3}x = y + \sqrt{2}bc \\ (2) - (1) : & \sqrt{2}y = ba \\ (3) : & d = -2\frac{x}{\sqrt{6}} \end{cases}.$$

Using $(a, b, c, d) \in \tilde{\mathcal{U}}_3^X$, we obtain $c \in [-a, -2a]$, where we recall that $a \leq 0$.

- Taking, for simplicity, $b = 1$ and choosing the value of a , the second equation shows that we can reach any $y = \frac{a}{\sqrt{2}} \leq 0$. Moreover, as c goes through $[-a, -2a]$, $\sqrt{2}bc$ goes through $[-\sqrt{2}a, -2\sqrt{2}a] = [-2y, -4y]$. Therefore, we see with the first equation that $\sqrt{3}x$ goes through $[-y, -3y]$, that is x goes through $[-\frac{y}{\sqrt{3}}, -\sqrt{3}y]$. It is not possible to reach other values for other values of $b \geq 0$.
- Similarly, taking $b = -1$ and choosing the value of a , the second equation shows that we can reach any $y = -\frac{a}{\sqrt{2}} \geq 0$. Moreover, as c goes through $[-a, -2a]$, $\sqrt{2}bc$ goes through $[2\sqrt{2}a, \sqrt{2}a] = [-4y, -2y]$. Therefore, we see with the first equation that $\sqrt{3}x$ goes through $[-3y, -y]$, that is x goes through $[-\sqrt{3}y, -\frac{y}{\sqrt{3}}]$. Again, it is not possible to reach other values for other values of $b \leq 0$.

Finally, the set \mathcal{V}_2 is the set in between the two lines $x + \sqrt{3}y = 0$ and $\sqrt{3}x + y = 0$, as on the right of Figure 1.

Case $j = 4$: We have $c \geq 0$, $a + c \geq 0$ and $2a + c \geq 0$. This leads to $f_{(a,b,c,d)}(X) = (bc + d, b(a + c) + d, b(2a + c) + d)$ and

$$\mathcal{V}_4 = \{(x, y) \in \mathbb{R}^2 \mid \exists(a, b, c, d) \in \tilde{\mathcal{U}}_4^X, p(x, y) = (bc + d, b(a + c) + d, b(2a + c) + d)\}.$$

We have

$$\left\{ \begin{array}{l} (1) : \quad bc + d = \frac{x}{\sqrt{6}} - \frac{y}{\sqrt{2}} \\ (2) : \quad b(a + c) + d = \frac{x}{\sqrt{6}} + \frac{y}{\sqrt{2}} \\ (3) : \quad b(2a + c) + d = -2\frac{x}{\sqrt{6}} \end{array} \right. \iff \left\{ \begin{array}{l} (2) - (1) : \quad \sqrt{2}y = ba \\ (3) - (2) : \quad -3\frac{x}{\sqrt{6}} - \frac{y}{\sqrt{2}} = ba \\ (3) : \quad b(2a + c) + d = -2\frac{x}{\sqrt{6}} \end{array} \right.$$

which is equivalent to

$$\left\{ \begin{array}{l} (1) : \quad \sqrt{2}y = ba \\ \sqrt{2}((2) - (1)) : \quad 3y = -\sqrt{3}x \\ (3) : \quad b(2a + c) + d = -2\frac{x}{\sqrt{6}} \end{array} \right.$$

This leads to

$$\mathcal{V}_4 = \{(x, y) \in \mathbb{R}^2 \mid x + \sqrt{3}y = 0\}.$$

Case $j = 5$: We have $c \leq 0$, $a + c \geq 0$ and therefore $2a + c \geq 0$. This leads to $f_{(a,b,c,d)}(X) = (d, b(a + c) + d, b(2a + c) + d)$ and

$$\mathcal{V}_5 = \{(x, y) \in \mathbb{R}^2 \mid \exists(a, b, c, d) \in \tilde{\mathcal{U}}_5^X, p(x, y) = (d, b(a + c) + d, b(2a + c) + d)\}.$$

We have

$$\left\{ \begin{array}{l} (1) : \quad d = \frac{x}{\sqrt{6}} - \frac{y}{\sqrt{2}} \\ (2) : \quad b(a + c) + d = \frac{x}{\sqrt{6}} + \frac{y}{\sqrt{2}} \\ (3) : \quad b(2a + c) + d = -2\frac{x}{\sqrt{6}} \end{array} \right. \iff \left\{ \begin{array}{l} (1) : \quad \frac{x}{\sqrt{6}} - \frac{y}{\sqrt{2}} = d \\ (2) - (1) : \quad \sqrt{2}y = b(a + c) \\ (3) - (1) : \quad -3\frac{x}{\sqrt{6}} + \frac{y}{\sqrt{2}} = b(2a + c) \end{array} \right. .$$

Using b and $a + c \geq 0$, we see with the second equation that y can take any value in \mathbb{R} . Let us consider an arbitrary fixed $y \in \mathbb{R}$. In fact, there are infinitely many choices for b, a and c corresponding to this value. Taking $b = \text{sign}(y)$, we have $a + c = \text{sign}(y)\sqrt{2}y = \sqrt{2}|y|$ and, since $c \leq 0$, a can take any value in $[\sqrt{2}|y|, +\infty)$. Therefore, $2a + c = a + (a + c)$ can take any value in $[2\sqrt{2}|y|, +\infty)$.

- If $y \geq 0$: $-3\frac{x}{\sqrt{6}} + \frac{y}{\sqrt{2}} = 2a + c$ goes through $[2\sqrt{2}y, +\infty)$. Therefore, $-3\frac{x}{\sqrt{6}}$ goes through $[3\frac{y}{\sqrt{2}}, +\infty)$, which means x goes through $(-\infty, -\sqrt{3}y]$.
- If $y \leq 0$: $-3\frac{x}{\sqrt{6}} + \frac{y}{\sqrt{2}} = -(2a + c)$ goes through $(-\infty, -2\sqrt{2}|y|]$. Therefore, $-3\frac{x}{\sqrt{6}}$ goes through $(-\infty, 3\frac{y}{\sqrt{2}}]$, which means x goes through $[-\sqrt{3}y, +\infty)$.

Finally, the set \mathcal{V}_5 is the set in between the two lines $x + \sqrt{3}y = 0$ and $y = 0$, as on the right of Figure 1.

Case $j = 6$: We have $c \leq 0$, $a + c \leq 0$ and $2a + c \geq 0$. This leads to $f_{(a,b,c,d)}(X) = (d, d, b(2a + c) + d)$ and

$$\mathcal{V}_6 = \{(x, y) \in \mathbb{R}^2 \mid \exists (a, b, c, d) \in \widetilde{\mathcal{U}}_6^X, p(x, y) = (d, d, b(2a + c) + d)\}.$$

We have

$$\begin{cases} (1) : & d = \frac{x}{\sqrt{6}} - \frac{y}{\sqrt{2}} \\ (2) : & d = \frac{x}{\sqrt{6}} + \frac{y}{\sqrt{2}} \\ (3) : & b(2a + c) + d = -2\frac{x}{\sqrt{6}} \end{cases} \iff \begin{cases} (1) : & \frac{x}{\sqrt{6}} - \frac{y}{\sqrt{2}} = d \\ ((2) - (1))/\sqrt{2} : & y = 0 \\ (3) - (1) : & -3\frac{x}{\sqrt{6}} + \frac{y}{\sqrt{2}} = b(2a + c) \end{cases}.$$

Using either d or $b(2a + c)$, x can take any value in \mathbb{R} and

$$\mathcal{V}_6 = \{(x, y) \in \mathbb{R}^2 \mid y = 0\}.$$

C Proofs of Section 5

This appendix is devoted to proving the results of Section 5: in Section C.1 we prove Proposition 4, in Section C.2 we prove Proposition 5, in Section C.3 we prove Theorem 6, in Section C.4 we prove Proposition 7, and in Section C.5 we prove Proposition 8.

C.1 Proof of Proposition 4

Before proving Proposition 4, we state and prove a lemma connecting the sets $\widetilde{\mathcal{O}}_j^n$, defined in (7), and \mathcal{X}_θ^n , defined in (16).

Lemma 15. Let $\theta \in (\mathbb{R}^E \times \mathbb{R}^B) \setminus \mathcal{Z}$, and let $n \in \mathbb{N}^*$. We have

$$\mathcal{X}_\theta^n = \{X \in \mathbb{R}^{N_0 \times n} \mid (X, \theta) \in \bigcup_{j=1}^{m_n} \widetilde{\mathcal{O}}_j^n\}.$$

Proof. Consider $\theta \in (\mathbb{R}^E \times \mathbb{R}^B) \setminus \mathcal{Z}$ and $n \in \mathbb{N}^*$ and let us first prove that $\{X \in \mathbb{R}^{N_0 \times n} \mid (X, \theta) \in \bigcup_{j=1}^{m_n} \widetilde{\mathcal{O}}_j^n\} \subseteq \mathcal{X}_\theta^n$.

Let $X = (x^{(i)})_{i \in \llbracket 1, n \rrbracket} \in \mathbb{R}^{N_0 \times n}$ and let $j \in \llbracket 1, m_n \rrbracket$ such that $(X, \theta) \in \widetilde{\mathcal{O}}_j^n$. Denote $\delta^1, \dots, \delta^n \in \{0, 1\}^{N_1 + \dots + N_{L-1}}$ such that $X \in \prod_{i=1}^n D_{\delta^i}(\theta)$. Using Lemma 13, (i), Item 2, we have

$$X \in \{X' \in \mathbb{R}^{N_0 \times n} \mid (X', \theta) \in \widetilde{\mathcal{O}}_j^n\} \subseteq \prod_{i=1}^n D_{\delta^i}(\theta).$$

Moreover, $\{X' \in \mathbb{R}^{N_0 \times n} \mid (X', \theta) \in \widetilde{\mathcal{O}}_j^n\}$ is open, since $\widetilde{\mathcal{O}}_j^n$ is open, and therefore

$$X \in \text{Int} \left(\prod_{i=1}^n D_{\delta^i}(\theta) \right) = \prod_{i=1}^n \text{Int}(D_{\delta^i}(\theta)).$$

Using the definition of \mathcal{X}_θ^n , in (16), we conclude that $X \in \mathcal{X}_\theta^n$. This concludes the proof of the first inclusion.

Before proving the converse inclusion, let us establish that $\mathcal{T}_n \subseteq \left(\bigcup_{j=1}^{m_n} \mathcal{O}_j^n\right)^c$, where the open sets $\mathcal{O}_1^n, \dots, \mathcal{O}_{m_n}^n$ are as in Theorem 1 and \mathcal{T}_n is defined in (28). By definition, we have $\mathcal{O}_j^n \subseteq \widetilde{\mathcal{O}}_j^n$ for $j \in \llbracket 1, m_n \rrbracket$. Thus $\bigcup_{j=1}^{m_n} \mathcal{O}_j^n \subseteq \bigcup_{j=1}^{m_n} \widetilde{\mathcal{O}}_j^n$. Also from Lemma 13, (i), Item 1, $\mathcal{T}_n = \left(\bigcup_{j=1}^{m_n} \widetilde{\mathcal{O}}_j^n\right)^c$. Hence

$$\mathcal{T}_n = \left(\bigcup_{j=1}^{m_n} \widetilde{\mathcal{O}}_j^n\right)^c \subseteq \left(\bigcup_{j=1}^{m_n} \mathcal{O}_j^n\right)^c. \quad (31)$$

Let us now prove the inclusion $\mathcal{X}_\theta^n \subseteq \{X \in \mathbb{R}^{N_0 \times n} \mid (X, \theta) \in \bigcup_{j=1}^{m_n} \widetilde{\mathcal{O}}_j^n\}$.

Let $X = (x^{(i)})_{i \in \llbracket 1, n \rrbracket} \in \mathcal{X}_\theta^n$. Since, as written above, $\left(\bigcup_{j=1}^{m_n} \widetilde{\mathcal{O}}_j^n\right)^c = \mathcal{T}_n$, proving that $(X, \theta) \in \bigcup_{j=1}^{m_n} \widetilde{\mathcal{O}}_j^n$ is equivalent to proving that $(X, \theta) \notin \mathcal{T}_n$.

Assume by contradiction that $(X, \theta) \in \mathcal{T}_n$. There exist $i \in \llbracket 1, n \rrbracket$, $\ell \in \llbracket 1, L-1 \rrbracket$ and $v \in V_\ell$ such that $(x^{(i)}, \theta) \in \mathcal{T}_v$, which means that

$$\sum_{v' \in V_{\ell-1}} w_{v' \rightarrow v} \left(f_\theta^{\ell-1}(x^{(i)})\right)_{v'} + b_v = 0.$$

Since $X \in \mathcal{X}_\theta^n$, we have $x^{(i)} \in \mathcal{X}_\theta$, and there exists $\delta \in A(\theta)$ such that $x^{(i)} \in \text{Int}(D_\delta(\theta))$. Let us show that this implies $\text{Int}(D_\delta(\theta)) \times \{\theta\} \subseteq \mathcal{T}_v$. Indeed, the function

$$x \mapsto \sum_{v' \in V_{\ell-1}} w_{v' \rightarrow v} \left(f_\theta^{\ell-1}(x)\right)_{v'} + b_v$$

is affine over the open set $\text{Int}(D_\delta(\theta))$. If it is not constantly equal to zero over this set, then since its value at $x^{(i)}$ is zero, it takes both positive and negative values over $\text{Int}(D_\delta(\theta))$, and thus $a_v(x, \theta)$ is not constant over $\text{Int}(D_\delta(\theta))$. This contradicts the definition of $D_\delta(\theta)$. Thus, the function is constantly equal to zero on $\text{Int}(D_\delta(\theta))$ and, using the definition of \mathcal{T}_v , in (24), $\text{Int}(D_\delta(\theta)) \times \{\theta\} \subseteq \mathcal{T}_v$. Therefore, $\text{Int}(D_\delta(\theta)) \times \{\theta\} \subseteq \mathcal{T}$ and, using the definition of \mathcal{T}_n in (28) and (31)

$$\left(\{X' \in \mathbb{R}^{N_0 \times n} \mid (x')^{(i)} \in \text{Int}(D_\delta(\theta))\} \times \{\theta\}\right) \subseteq \mathcal{T}_n \subseteq \left(\bigcup_{j=1}^{m_n} \mathcal{O}_j^n\right)^c.$$

Therefore, $\{X' \in \mathbb{R}^{N_0 \times n} \mid (x')^{(i)} \in \text{Int}(D_\delta(\theta))\} \subseteq z_n(\theta)$, where $z_n(\theta)$ is defined in (17). Since $\text{Int}(D_\delta(\theta))$ is non-empty, the Lebesgue measure of $\{X' \in \mathbb{R}^{N_0 \times n} \mid (x')^{(i)} \in \text{Int}(D_\delta(\theta))\}$ is not zero and therefore $\theta \in \mathcal{Z}_n$, as defined in (18). This contradicts the hypothesis on θ and finishes the proof of the statement $(X, \theta) \notin \mathcal{T}_n$.

This concludes the proof of the inclusion $\mathcal{X}_\theta^n \subseteq \{X \in \mathbb{R}^{N_0 \times n} \mid (X, \theta) \in \bigcup_{j=1}^{m_n} \widetilde{\mathcal{O}}_j^n\}$ and finishes the proof of Lemma 15. \square

Proof of Proposition 4. We remind that for all $n \in \mathbb{N}^*$ and all $\theta \in \mathbb{R}^E \times \mathbb{R}^B$,

$$z_n(\theta) = \left\{ X \in \mathbb{R}^{N_0 \times n} \mid (X, \theta) \in \left(\bigcup_{j=1}^{m_n} \mathcal{O}_j^n \right)^c \right\},$$

$$\mathcal{Z}_n = \left\{ \theta \in \mathbb{R}^E \times \mathbb{R}^B \mid z_n(\theta) \text{ has strictly positive Lebesgue measure} \right\}$$

and $\mathcal{Z} = \bigcup_{n \in \mathbb{N}^*} \mathcal{Z}_n$.

Let us first prove (i).

Let us write λ for Lebesgue measure on $\mathbb{R}^{N_0 \times n} \times (\mathbb{R}^E \times \mathbb{R}^B)$, and λ^1 for Lebesgue measure on $\mathbb{R}^{N_0 \times n}$. Note that $\left(\bigcup_{j=1}^{m_n} \mathcal{O}_j^n \right)^c$ is measurable in $\mathbb{R}^{N_0 \times n} \times (\mathbb{R}^E \times \mathbb{R}^B)$, as a closed set, and thus for all $\theta \in \mathbb{R}^E \times \mathbb{R}^B$, $z_n(\theta)$ is measurable in $\mathbb{R}^{N_0 \times n}$. With similar arguments, the function $\theta \mapsto \lambda^1(z_n(\theta))$ is measurable on $\mathbb{R}^E \times \mathbb{R}^B$ as the integral with respect to X of a measurable function of X and θ . Hence, \mathcal{Z}_n , as the set where this function is strictly positive, is indeed measurable.

Let $n \in \mathbb{N}^*$. We first prove that \mathcal{Z}_n has Lebesgue measure zero on $\mathbb{R}^E \times \mathbb{R}^B$. Let us assume by contradiction that \mathcal{Z}_n has strictly positive Lebesgue measure.

Let us write $\mathcal{C} = \bigcup_{\theta \in \mathcal{Z}_n} (z_n(\theta) \times \{\theta\})$. For all θ , for all $X \in z_n(\theta)$, $(X, \theta) \in \left(\bigcup_{j=1}^{m_n} \mathcal{O}_j^n \right)^c$. Hence, we have, for all θ , $z_n(\theta) \times \{\theta\} \subseteq \left(\bigcup_{j=1}^{m_n} \mathcal{O}_j^n \right)^c$, and thus $\mathcal{C} \subseteq \left(\bigcup_{j=1}^{m_n} \mathcal{O}_j^n \right)^c$. We have

$$\begin{aligned} \lambda(\mathcal{C}) &= \int_{\mathbb{R}^E \times \mathbb{R}^B} \int_{\mathbb{R}^{N_0 \times n}} 1_{\mathcal{C}}(X, \theta) dX d\theta \\ &= \int_{\mathcal{Z}_n} \int_{z_n(\theta)} 1 dX d\theta \\ &= \int_{\mathcal{Z}_n} \lambda^1(z_n(\theta)) d\theta \\ &> 0, \end{aligned}$$

as the integral of the strictly positive function $\lambda^1(z_n(\theta))$, on the set \mathcal{Z}_n with non-zero Lebesgue measure, is strictly positive (note that since all the functions in the above display are non-negative and measurable, their integrability is guaranteed). This is in contradiction with the fact that $\mathcal{C} \subseteq \left(\bigcup_{j=1}^{m_n} \mathcal{O}_j^n \right)^c$, since $\left(\bigcup_{j=1}^{m_n} \mathcal{O}_j^n \right)^c$ has Lebesgue measure zero from Theorem 1 (i).

This concludes the proof of the statement that \mathcal{Z}_n has Lebesgue measure zero on $\mathbb{R}^E \times \mathbb{R}^B$ and concludes the proof of (i).

The item (ii) is an immediate consequence of (i), since \mathcal{Z} is a countable union of measurable sets of measure 0.

Let us now prove (iii).

For any $\theta \in (\mathbb{R}^E \times \mathbb{R}^B) \setminus \mathcal{Z}$, $n \in \mathbb{N}^*$, and $X \in \mathcal{X}_\theta^n$, Lemma 15 shows that there exists $j \in \llbracket 1, m_n \rrbracket$ such that $(X, \theta) \in \tilde{\mathcal{O}}_j^n$. Lemma 13 (i) Item 3 shows that $(X', \theta') \mapsto f_{\theta'}(X')$ is an analytic function on the open set $\tilde{\mathcal{O}}_j^n$. Its restriction to the open neighborhood $\{\theta' \in \mathbb{R}^E \times \mathbb{R}^B \mid (X, \theta') \in \tilde{\mathcal{O}}_j^n\}$ of θ is also an analytic function. In particular, $\theta' \mapsto f_{\theta'}(X)$ is differentiable at θ . \square

C.2 Proof of Proposition 5

Consider $\theta \in (\mathbb{R}^E \times \mathbb{R}^B) \setminus \mathcal{Z}$ and $n \in \mathbb{N}^*$ and let us first prove that $r_n^*(\theta) \geq \max_{j \in I_n(\theta)} r_j^n$.

Consider $j \in I_n(\theta)$ such that $r_j^n = \max_{j' \in I_n(\theta)} r_{j'}^n$. Because of the definition of $I_n(\theta)$, in (21), there exists $X \in \mathbb{R}^{N_0 \times n}$ such that $(X, \theta) \in \mathcal{O}_j^n$. As a first consequence, using Theorem 1, (i) third item, we know that $r_j^n = \text{rank}(\mathbf{D}f_\theta(X))$. As a second consequence, since $\theta \in (\mathbb{R}^E \times \mathbb{R}^B) \setminus \mathcal{Z}$ and $\mathcal{O}_j^n \subseteq \widetilde{\mathcal{O}}_j^n$ we can use Lemma 15, at the beginning of Section C.1, and we have $X \in \mathcal{X}_\theta^n$.

This leads to the conclusion that

$$\max_{j' \in I_n(\theta)} r_{j'}^n = r_j^n = \text{rank}(\mathbf{D}f_\theta(X)) \leq \max_{X' \in \mathcal{X}_\theta^n} \text{rank}(\mathbf{D}f_\theta(X')) = r_n^*(\theta).$$

Let us now prove that $r_n^*(\theta) \leq \max_{j \in I_n(\theta)} r_j^n$.

Because the rank can only take a finite number of values and considering the definition of $r_n^*(\theta)$ in (19), we know there exists $X = (x^{(i)})_{i \in \llbracket 1, n \rrbracket} \in \mathcal{X}_\theta^n$ such that

$$\text{rank}(\mathbf{D}f_\theta(X)) = r_n^*(\theta).$$

Since $\theta \notin \mathcal{Z}$ and $X \in \mathcal{X}_\theta^n$, we know thanks to Lemma 15 that there exists $j \in \llbracket 1, m_n \rrbracket$ such that $(X, \theta) \in \widetilde{\mathcal{O}}_j^n$. Thus, by definition of r_j^n , see (8), we have

$$r_j^n \geq \mathbf{D}f_\theta(X) = r_n^*(\theta).$$

To conclude the proof, we only need to show that $j \in I_n(\theta)$.

Since by definition we have $\mathcal{O}_j^n \subseteq \widetilde{\mathcal{O}}_j^n$, we have the inclusion

$$\{X' \in \mathbb{R}^{N_0 \times n} \mid (X', \theta) \in \mathcal{O}_j^n\} \subseteq \{X' \in \mathbb{R}^{N_0 \times n} \mid (X', \theta) \in \widetilde{\mathcal{O}}_j^n\}. \quad (32)$$

Let us show that the set on the left of this inclusion is non-empty, which, by definition of $I_n(\theta)$, will show that $j \in I_n(\theta)$. First, both sets are open since \mathcal{O}_j^n and $\widetilde{\mathcal{O}}_j^n$ are open, and the set on the right is non-empty since it contains X . Assume by contradiction that the set on the left is empty. Since for all $j' \in \llbracket 1, m_n \rrbracket \setminus \{j\}$, we have $\mathcal{O}_{j'}^n \cap \widetilde{\mathcal{O}}_j^n = \emptyset$, this means, recalling the definition of $z_n(\theta)$ in (17), that we have

$$\{X' \in \mathbb{R}^{N_0 \times n} \mid (X', \theta) \in \widetilde{\mathcal{O}}_j^n\} \subseteq \left\{X' \in \mathbb{R}^{N_0 \times n} \mid (X', \theta) \in \left(\bigcup_{j=1}^{m_n} \mathcal{O}_j^n\right)^c\right\} = z_n(\theta).$$

Since, as we have just shown, the set $\{X' \in \mathbb{R}^{N_0 \times n} \mid (X', \theta) \in \widetilde{\mathcal{O}}_j^n\}$ is open and non-empty, this means that $z_n(\theta)$ has positive Lebesgue measure, and thus, using (18), that $\theta \in \mathcal{Z}_n \subseteq \mathcal{Z}$. This contradicts the hypothesis $\theta \in (\mathbb{R}^E \times \mathbb{R}^B) \setminus \mathcal{Z}$. This concludes the proof establishing that the open set $\{X' \in \mathbb{R}^{N_0 \times n} \mid (X', \theta) \in \mathcal{O}_j^n\}$ is non-empty, and thus that $j \in I_n(\theta)$. This proves the desired inequality $r_n^*(\theta) \leq \max_{j \in I_n(\theta)} r_j^n$.

C.3 Proof of Theorem 6

We begin the proof with a lemma. To state the lemma, let us recall first that for any $n \in \mathbb{N}^*$ and any $j \in \llbracket 1, m_n \rrbracket$, the mapping $(X, \theta) \mapsto a(X, \theta)$, is constant on $\widetilde{\mathcal{O}}_j^n$ and that we denote by $\Delta_j^n \in \{0, 1\}^{(N_1 + \dots + N_{L-1}) \times n}$ its constant value, see (7). Since by definition we have $\mathcal{O}_j^n \subseteq \widetilde{\mathcal{O}}_j^n$, in particular the mapping $(X, \theta) \mapsto a(X, \theta)$, is also constant equal to Δ_j^n over \mathcal{O}_j^n . Let us write $\Delta_j^n = (\delta_1^{j,n}, \dots, \delta_n^{j,n})$. We therefore have, for all $(X, \theta) = ((x^{(i)})_{i \in \llbracket 1, n \rrbracket}, \theta) \in \mathcal{O}_j^n$,

$$a(x^{(i)}, \theta) = \delta_i^{j,n} \in \{0, 1\}^{N_1 + \dots + N_{L-1}}, \quad \text{for all } i \in \llbracket 1, n \rrbracket. \quad (33)$$

For the next lemma, note that $I_n(\theta)$ in (21) can in fact be defined for any $\theta \in \mathbb{R}^E \times \mathbb{R}^B$.

Lemma 16. For any $n \in \mathbb{N}^*$, $\theta \in \mathbb{R}^E \times \mathbb{R}^B$, and $j \in \llbracket 1, m_n \rrbracket$, we have

$$j \in I_n(\theta) \quad \implies \quad \forall i \in \llbracket 1, n \rrbracket, \quad \delta_i^{j,n} \in A(\theta).$$

Conversely, for any $n \in \mathbb{N}^*$, $\theta \in (\mathbb{R}^E \times \mathbb{R}^B) \setminus \mathcal{Z}$, and $j \in \llbracket 1, m_n \rrbracket$, we have

$$j \in I_n(\theta) \quad \iff \quad \forall i \in \llbracket 1, n \rrbracket, \quad \delta_i^{j,n} \in A(\theta).$$

Proof. Consider $n \in \mathbb{N}^*$, $\theta \in \mathbb{R}^E \times \mathbb{R}^B$ and $j \in \llbracket 1, m_n \rrbracket$.

Let us first prove that if $j \in I_n(\theta)$ then, for all $i \in \llbracket 1, n \rrbracket$, $\delta_i^{j,n} \in A(\theta)$. Indeed, let $j \in I_n(\theta)$. For any $i \in \llbracket 1, n \rrbracket$, given the definition of $A(\theta)$ in (15), in order to prove that $\delta_i^{j,n} \in A(\theta)$, it is sufficient to prove $\text{Int}(D_{\delta_i^{j,n}}(\theta)) \neq \emptyset$, where the notation $D_\delta(\theta)$ is defined in (14).

Since $j \in I_n(\theta)$, there exists X such that $(X, \theta) = ((x^{(i)})_{i \in \llbracket 1, n \rrbracket}, \theta) \in \mathcal{O}_j^n$. Moreover, since \mathcal{O}_j^n is open

$$\{X' \in \mathbb{R}^{N_0 \times n} \mid (X', \theta) \in \mathcal{O}_j^n\}$$

is open and, since it contains X , non-empty. We also have using (33)

$$\{X' \in \mathbb{R}^{N_0 \times n} \mid (X', \theta) \in \mathcal{O}_j^n\} \subseteq \prod_{i=1}^n D_{\delta_i^{j,n}}(\theta)$$

and thus

$$\{X' \in \mathbb{R}^{N_0 \times n} \mid (X', \theta) \in \mathcal{O}_j^n\} \subseteq \prod_{i=1}^n \text{Int}(D_{\delta_i^{j,n}}(\theta)),$$

which shows that, for $i \in \llbracket 1, n \rrbracket$, $\text{Int}(D_{\delta_i^{j,n}}(\theta)) \neq \emptyset$ and therefore $\delta_i^{j,n} \in A(\theta)$. This finishes the proof of the first implication.

Consider $n \in \mathbb{N}^*$, $\theta \in (\mathbb{R}^E \times \mathbb{R}^B) \setminus \mathcal{Z}$ and $j \in \llbracket 1, m_n \rrbracket$.

Let us now prove that if, for all $i \in \llbracket 1, n \rrbracket$, $\delta_i^{j,n} \in A(\theta)$, then $j \in I_n(\theta)$. To do so, assume that for all $i \in \llbracket 1, n \rrbracket$, $\delta_i^{j,n} \in A(\theta)$. Considering the definition of $I_n(\theta)$, in (21), it suffices to put to evidence $X \in \mathbb{R}^{N_0 \times n}$ such that $(X, \theta) \in \mathcal{O}_j^n$ to prove that $j \in I_n(\theta)$.

Since, for all $i \in \llbracket 1, n \rrbracket$, $\delta_i^{j,n} \in A(\theta)$, using the definition of $A(\theta)$ in (15), $\text{Int}(D_{\delta_i^{j,n}}(\theta)) \neq \emptyset$. Therefore, $\Pi_{i=1}^n \text{Int}(D_{\delta_i^{j,n}}(\theta)) \subseteq \mathbb{R}^{N_0 \times n}$ is a non-empty open set. Moreover, since $\theta \notin \mathcal{Z}$,

$$\{X \in \mathbb{R}^{N_0 \times n} \mid (X, \theta) \in (\cup_{j=1}^{m_n} \mathcal{O}_j^n)^c\}$$

has Lebesgue measure zero in $\mathbb{R}^{N_0 \times n}$. Therefore, since $\Pi_{i=1}^n \text{Int}(D_{\delta_i^{j,n}}(\theta))$ is a non-empty open set, there exists $j' \in \llbracket 1, m_n \rrbracket$ such that

$$\Pi_{i=1}^n \text{Int}(D_{\delta_i^{j',n}}(\theta)) \cap \{X \in \mathbb{R}^{N_0 \times n} \mid (X, \theta) \in \mathcal{O}_{j'}^n\} \neq \emptyset.$$

Consider X in this set, we have

$$a(X, \theta) = \delta^{j',n} = \delta^{j',n}.$$

Using Theorem 1, (i), second item, we conclude that $j = j'$. Finally, $(X, \theta) \in \mathcal{O}_j^n$ and $j \in I_n(\theta)$. This concludes the proof. \square

Proof of Theorem 6. Consider θ and θ' in $(\mathbb{R}^E \times \mathbb{R}^B) \setminus \mathcal{Z}$ and assume $A(\theta) \subseteq A(\theta')$.

Let $n \in \mathbb{N}^*$ and $j \in I_n(\theta)$. Using Lemma 16, we know that for all $i \in \llbracket 1, n \rrbracket$, $\delta_i^{j,n} \in A(\theta)$, and therefore for all $i \in \llbracket 1, n \rrbracket$, $\delta_i^{j,n} \in A(\theta')$. Using Lemma 16 again, we obtain that $j \in I_n(\theta')$. As a conclusion, for all $n \in \mathbb{N}^*$, $I_n(\theta) \subseteq I_n(\theta')$.

Using Proposition 5, we obtain that for all $n \in \mathbb{N}^*$, $r_n^*(\theta) \leq r_n^*(\theta')$. We conclude, using the definition of r^* , in (19), that $r^*(\theta) \leq r^*(\theta')$.

This concludes the proof of the first statement of Theorem 6.

The second statement follows by applying the first statement twice. Once with the hypothesis $A(\theta) \subseteq A(\theta')$ and a second time with the hypothesis $A(\theta') \subseteq A(\theta)$. This concludes the proof of the theorem. \square

C.4 Proof of Proposition 7

Let $\theta \in \mathbb{R}^E \times \mathbb{R}^B$ and let $\tilde{\theta} \sim \theta$. The first statement is immediate: for $n \in \mathbb{N}^*$, the third item of Proposition 14 shows that

$$(X, \theta) \in \bigcup_{j=1}^{m_n} \mathcal{O}_j^n \iff (X, \tilde{\theta}) \in \bigcup_{j=1}^{m_n} \mathcal{O}_j^n,$$

i.e. $z_n(\theta) = z_n(\tilde{\theta})$. This immediately yields $\theta \in \mathcal{Z} \iff \tilde{\theta} \in \mathcal{Z}$.

Let us now show the second statement. Consider $\theta \in (\mathbb{R}^E \times \mathbb{R}^B) \setminus \mathcal{Z}$ and $\tilde{\theta} \sim \theta$. The first statement shows that $\tilde{\theta} \in (\mathbb{R}^E \times \mathbb{R}^B) \setminus \mathcal{Z}$. Let us fix $n \in \mathbb{N}^*$ and show that $r_n^*(\theta) = r_n^*(\tilde{\theta})$. Since, as discussed in Section 3, the rank is bounded by $|E| + |B|$, there is a finite number of possible values, and thus there exists $X \in \mathcal{X}_\theta^n$ that realizes the maximum in the definition (19), i.e. satisfying $\text{rank}(\mathbf{D}f_\theta(X)) = r_n^*(\theta)$. Since $\theta \in (\mathbb{R}^E \times \mathbb{R}^B) \setminus \mathcal{Z}$, by Lemma 15 there exists $j \in \llbracket 1, m_n \rrbracket$ such that $(X, \theta) \in \mathcal{O}_j^n$. The first item of Proposition 14 shows that

$\text{rank}(\mathbf{D}f_{\tilde{\theta}}(X)) = \text{rank}(\mathbf{D}f_{\theta}(X))$, and the second item shows that there exists $j' \in \llbracket 1, m_n \rrbracket$ such that $(X, \tilde{\theta}) \in \tilde{\mathcal{O}}_{j'}^n$. Using again Lemma 15 and the fact that $\tilde{\theta} \in (\mathbb{R}^E \times \mathbb{R}^B) \setminus \mathcal{Z}$, we find that $X \in \mathcal{X}_{\tilde{\theta}}^n$, and we thus have, using the definition (19):

$$r_n^*(\tilde{\theta}) = \max_{X' \in \mathcal{X}_{\tilde{\theta}}^n} \text{rank}(\mathbf{D}f_{\tilde{\theta}}(X')) \geq \text{rank}(\mathbf{D}f_{\tilde{\theta}}(X)) = \text{rank}(\mathbf{D}f_{\theta}(X)) = r_n^*(\theta).$$

We have shown $r_n^*(\theta) \leq r_n^*(\tilde{\theta})$, and by swapping the roles of θ and $\tilde{\theta}$, we obtain $r_n^*(\theta) = r_n^*(\tilde{\theta})$.

Finally, since $r_n^*(\theta) = r_n^*(\tilde{\theta})$ for all $n \in \mathbb{N}^*$, then (20) guarantees that $r^*(\theta) = r^*(\tilde{\theta})$. This concludes the proof of Proposition 7.

C.5 Proof of Proposition 8

Similarly to what was done in the proof of Theorem 1, all along the proof of Proposition 8, we consider the canonical basis $(e_1, \dots, e_{|E|+|B|})$ of $\mathbb{R}^E \times \mathbb{R}^B$ and the canonical basis $(\varepsilon_1, \dots, \varepsilon_n)$ of \mathbb{R}^n , where this time we have $N_L = 1$. For all $\theta' \in \mathbb{R}^E \times \mathbb{R}^B$ and all $X' \in \mathbb{R}^n$ for which $\theta'' \mapsto f_{\theta''}(X')$ is differentiable at θ' , we write $Jf_{\theta'}(X')$ for the $n \times (|E| + |B|)$ matrix of $\mathbf{D}f_{\theta'}(X')$ in these two bases.

Consider $\theta \in (\mathbb{R}^E \times \mathbb{R}^B) \setminus \mathcal{Z}$ and $\varepsilon > 0$ as in the first statement of the proposition. From Proposition 5, there is $n \in \mathbb{N}^*$, $j \in \llbracket 1, m_n \rrbracket$ and $\bar{X} = (x^{(1)}, \dots, x^{(n)})$ such that $(\bar{X}, \theta) \in \mathcal{O}_j^n$ and $\text{rank}(\mathbf{D}f_{\theta}(\bar{X})) = r^*(\theta)$. Also, from Theorem 1 (i), there is $\varepsilon' > 0$ such that $\theta' \mapsto f_{\theta'}(\bar{X})$ is continuously differentiable on $B(\theta, \varepsilon')$. We consider such an ε' satisfying $\varepsilon' \leq \varepsilon$.

Then, $r^*(\theta) = \text{rank}(\mathbf{D}f_{\theta}(\bar{X})) = \text{rank}(Jf_{\theta}(\bar{X}))$. Hence we can extract $r^*(\theta)$ rows from $Jf_{\theta}(\bar{X})$ such that, up to reordering, we have, with $X = (x^{(1)}, \dots, x^{(r^*(\theta))})$, $\text{rank}(\mathbf{D}f_{\theta}(X)) = r^*(\theta)$. Furthermore, still, $\theta' \mapsto f_{\theta'}(X)$ is continuously differentiable on $B(\theta, \varepsilon')$.

Then, we can extract $r^*(\theta)$ columns from $Jf_{\theta}(X)$ for which the resulting $r^*(\theta) \times r^*(\theta)$ matrix is invertible. These $r^*(\theta)$ columns are associated to a subset of $E \cup B$, that we write S . For $\phi \in S$, we define $\theta(\phi) \in \mathbb{R}^E \times \mathbb{R}^B$ by $\theta(\phi)_w = \theta_w$ for $w \in (E \cup B) \setminus S$ and by $\theta(\phi)_w = \phi_w$ for $w \in S$. We also let $\theta_{|S} = (\theta_w)_{w \in S}$. Then we consider the function $g : \mathbb{R}^S \rightarrow \mathbb{R}^{r^*(\theta)}$ defined by the row vector $g(\phi) = f_{\theta(\phi)}(X)$.

The differential of g at $\theta_{|S}$ is defined by the invertible $r^*(\theta) \times r^*(\theta)$ matrix extracted from $Jf_{\theta}(X)$, discussed above, in the canonical bases of \mathbb{R}^S and $\mathbb{R}^{r^*(\theta)}$. In addition, g is continuously differentiable on $B(\theta_{|S}, \varepsilon')$.

Hence we can apply the inverse function theorem. There is an open set $U \subseteq B(\theta_{|S}, \varepsilon')$ containing $\theta_{|S}$ and an open set V containing $g(\theta_{|S}) = f_{\theta(\theta_{|S})}(X) = f_{\theta}(X)$ such that g is bijective from U to V . We let g^{-1} be the inverse of g . Then, there is $\gamma > 0$ small enough such that $f_{\theta}(X) + \gamma\{-1, 1\}^{r^*(\theta)} \subseteq V$, where “+” denotes the Minkowski sum.

Let $t = (t_1, \dots, t_{r^*(\theta)}) = f_{\theta}(X)$. Then, for each $I \subseteq \llbracket 1, r^*(\theta) \rrbracket$, define $t' \in \mathbb{R}^{r^*(\theta)}$ by $t'_i = t_i + \gamma$ if $i \in I$ and $t'_i = t_i - \gamma$ if $i \notin I$. Since $t' \in t + \gamma\{-1, 1\}^{r^*(\theta)} \subseteq V$, we can define $\theta' = \theta(g^{-1}(t'))$ with $\theta' \in \theta(U) \subseteq B(\theta, \varepsilon') \subseteq B(\theta, \varepsilon)$. This yields, $f_{\theta'}(X) = f_{\theta(g^{-1}(t'))}(X) = g(g^{-1}(t')) = t'$. Hence, for $i \in I$, $f_{\theta'}(x^{(i)}) = t'_i = t_i + \gamma$ and for $i \notin I$, $f_{\theta'}(x^{(i)}) = t'_i = t_i - \gamma$. By definition, this implies that $\text{fs}_{B(\theta, \varepsilon), \gamma} \geq r^*(\theta)$. This shows the first part of the proposition.

The second part of the proposition is a consequence of the first part. Indeed, let us fix some $\varepsilon > 0$, let us take $\theta \in (\mathbb{R}^E \times \mathbb{R}^B) \setminus \mathcal{Z}$ such that $r^*(\theta) = \max_{\tilde{\theta} \in (\mathbb{R}^E \times \mathbb{R}^B) \setminus \mathcal{Z}} r^*(\tilde{\theta})$, and let us take $\gamma' > 0$ such that (22) holds with γ there replaced by γ' here. Then we have

$$\text{fS}_{\mathbb{R}^E \times \mathbb{R}^B, \gamma'} \geq \text{fS}_{B(\theta, \varepsilon), \gamma'} \geq r^*(\theta).$$

Finally, the third part of the proposition is also a consequence of the first part. Consider r such that $A_r = \{\theta \in (\mathbb{R}^E \times \mathbb{R}^B) \setminus \mathcal{Z} \mid r^*(\theta) = r\}$ has non-empty interior. Consider θ in the interior of A_r . There is $\varepsilon > 0$ such that $B(\theta, \varepsilon) \subseteq A_r$. Hence for any $\gamma'' > 0$ such that (22) holds, with γ replaced by γ'' ,

$$\text{fS}_{A_r, \gamma''} \geq \text{fS}_{B(\theta, \varepsilon), \gamma''} \geq r^*(\theta).$$

Since $r^*(\theta) = r$, we obtain $\text{fS}_{A_r, \gamma''} \geq r$ for some $\gamma'' > 0$. Hence the third part is proved.

This concludes the proof.

D Proofs of Section 6

In this appendix, we prove the two results of Section 6.2: in Section D.1 we prove Proposition 9 and in Section D.2 we prove Theorem 11.

D.1 Proof of Proposition 9

Before proving Proposition 9, we state and prove two lemmas.

Lemma 17. Let $\theta \in (\mathbb{R}^E \times \mathbb{R}^B) \setminus \mathcal{Z}$.

Let $n, n' \in \mathbb{N}^*$, let $X \in \mathcal{X}_\theta^n$ and let $X' \in \mathcal{X}_\theta^{n'}$. Let us write $X'' \in \mathcal{X}_\theta^{n+n'}$ the matrix obtained by concatenating the columns of X and X' . We then have the following inequalities

$$\begin{aligned} \text{rank}(\mathbf{D}f_\theta(X)) &\leq \text{rank}(\mathbf{D}f_\theta(X'')), \\ \text{rank}(\mathbf{D}f_\theta(X')) &\leq \text{rank}(\mathbf{D}f_\theta(X'')), \\ \text{rank}(\mathbf{D}f_\theta(X'')) &\leq \text{rank}(\mathbf{D}f_\theta(X)) + \text{rank}(\mathbf{D}f_\theta(X')), \end{aligned}$$

where using Proposition 4, (iii), all the differentials are well defined.

Proof of Lemma 17. Let us consider the canonical basis \mathcal{B} of $\mathbb{R}^E \times \mathbb{R}^B$. Also, for any $m \in \mathbb{N}^*$, let us consider the canonical basis \mathcal{B}_m of $\mathbb{R}^{N_L \times m}$. With the bases \mathcal{B} and \mathcal{B}_n , the linear operator $\mathbf{D}f_\theta(X) : \mathbb{R}^E \times \mathbb{R}^B \rightarrow \mathbb{R}^{N_L \times n}$ corresponds to a Jacobian matrix that we write $Jf_\theta(X)$. We define $Jf_\theta(X')$ similarly with the bases \mathcal{B} and $\mathcal{B}_{n'}$ and $Jf_\theta(X'')$ similarly with the bases \mathcal{B} and $\mathcal{B}_{n+n'}$. Writing ∇_θ the gradient of a scalar quantity depending on $\theta \in \mathbb{R}^E \times \mathbb{R}^B$, then the matrix $Jf_\theta(X)$ is composed of the nN_L rows

$$Jf_\theta(X)_{(i-1)N_L+j,:} = \left(\nabla_\theta \left[f_\theta(x^{(i)})_j \right] \right)^\top, \quad \text{for all } i \in \llbracket 1, n \rrbracket \text{ and } j \in \llbracket 1, N_L \rrbracket,$$

the matrix $Jf_\theta(X')$ is composed of the $n'N_L$ rows

$$Jf_\theta(X')_{(i-1)N_L+j,:} = \left(\nabla_\theta \left[f_\theta(x'^{(i)})_j \right] \right)^\top, \text{ for all } i \in \llbracket 1, n' \rrbracket \text{ and } j \in \llbracket 1, N_L \rrbracket,$$

and the matrix $Jf_\theta(X'')$ is obtained by concatenating the rows of $Jf_\theta(X)$ and $Jf_\theta(X')$. Well-known properties of the matrix rank yield

$$\begin{aligned} \text{rank}(Jf_\theta(X)) &\leq \text{rank}(Jf_\theta(X'')), \\ \text{rank}(Jf_\theta(X')) &\leq \text{rank}(Jf_\theta(X'')), \\ \text{rank}(Jf_\theta(X'')) &\leq \text{rank}(Jf_\theta(X)) + \text{rank}(Jf_\theta(X')). \end{aligned}$$

These three inequalities are equivalent to the three inequalities of the lemma. \square

Lemma 18. Let $\theta \in (\mathbb{R}^E \times \mathbb{R}^B) \setminus \mathcal{Z}$. Let $n \in \mathbb{N}^*$ and let $X \in \mathcal{X}_\theta^n$. If $\text{rank}(\mathbf{D}f_\theta(X)) < r^*(\theta)$, then there exists $x \in \mathcal{X}_\theta$ such that, writing $X_x \in \mathcal{X}_\theta^{n+1}$ for the matrix obtained by adding x as an additional last column to X , we have

$$\text{rank}(\mathbf{D}f_\theta(X_x)) \geq \text{rank}(\mathbf{D}f_\theta(X)) + 1.$$

Furthermore, the set of such x 's has non-zero Lebesgue measure on \mathbb{R}^{N_0} .

Proof of Lemma 18. Let $\theta \in (\mathbb{R}^E \times \mathbb{R}^B) \setminus \mathcal{Z}$, $n \in \mathbb{N}^*$ and $X \in \mathcal{X}_\theta^n$ such that $\text{rank}(\mathbf{D}f_\theta(X)) < r^*(\theta)$. Let $n' \in \mathbb{N}^*$ and $X' \in \mathcal{X}_\theta^{n'}$ such that $\text{rank}(\mathbf{D}f_\theta(X')) = r^*(\theta)$ (the existence being guaranteed by definition of $r^*(\theta)$). Let $X'' \in \mathcal{X}_\theta^{n+n'}$ be obtained by concatenating the columns of X and X' . From Lemma 17, we have

$$\text{rank}(\mathbf{D}f_\theta(X'')) \geq \text{rank}(\mathbf{D}f_\theta(X')) = r^*(\theta) \geq \text{rank}(\mathbf{D}f_\theta(X)) + 1.$$

Using the matrix notations defined in the proof of Lemma 17, there exists a row of $Jf_\theta(X')$, written $\left(\nabla_\theta \left[f_\theta(x'^{(i)})_j \right] \right)^\top$ for $i \in \llbracket 1, n' \rrbracket$ and $j \in \llbracket 1, N_L \rrbracket$, with $x'^{(i)} \in \mathcal{X}_\theta$, that does not belong to the range of $Jf_\theta(X)$. Recall that for $x \in \mathcal{X}_\theta$, we write X_x for the matrix of \mathcal{X}_θ^{n+1} obtained by concatenating X and x . From the above, we have $\text{rank}(\mathbf{D}f_\theta(X_{x'^{(i)}})) \geq \text{rank}(\mathbf{D}f_\theta(X)) + 1$.

This proves the first statement of the lemma and we still need to prove that the set of such x 's has non-zero Lebesgue measure on \mathbb{R}^{N_0} .

Consider $x'^{(i)} \in \mathcal{X}_\theta$ as constructed in the first part of the proof. By continuity of $x \mapsto \mathbf{D}f_\theta(X_x)$ at $x'^{(i)}$, and since the rank is a lower semi-continuous function and \mathcal{X}_θ is open, then there exists $\varepsilon > 0$ such that $B(x'^{(i)}, \varepsilon) \subseteq \mathcal{X}_\theta$ and for all $x \in B(x'^{(i)}, \varepsilon)$,

$$\text{rank}(\mathbf{D}f_\theta(X_x)) \geq \text{rank}(\mathbf{D}f_\theta(X_{x'^{(i)}})) \geq \text{rank}(\mathbf{D}f_\theta(X)) + 1.$$

The ball $B(x'^{(i)}, \varepsilon)$ has a non-zero Lebesgue measure, which concludes the proof. \square

We can now prove Proposition 9.

Proof of Proposition 9. Let us first show that the sequence $(r_n^*(\theta))_{n \in \mathbb{N}^*}$ is non-decreasing. Let $n \in \mathbb{N}^*$ and $X \in \mathcal{X}_\theta^n$ such that $\text{rank}(\mathbf{D}f_\theta(X)) = r_n^*(\theta)$. For any $x \in \mathcal{X}_\theta$, using the notation X_x of Lemma 18, Lemma 17 shows that $r_n^*(\theta) = \text{rank}(\mathbf{D}f_\theta(X)) \leq \text{rank}(\mathbf{D}f_\theta(X_x))$, and thus $r_n^*(\theta) \leq r_{n+1}^*(\theta)$. Hence, the sequence $(r_n^*(\theta))_{n \in \mathbb{N}^*}$ is non-decreasing.

Since the input space of $\mathbf{D}f_\theta(X)$ is $\mathbb{R}^E \times \mathbb{R}^B$ for all X , the sequence is also upper bounded by $r^*(\theta) \leq |E| + |B|$. Therefore, since it only takes integer values, there exists n such that $r_n^*(\theta) = r^*(\theta)$ and we can write $n^*(\theta) \in \mathbb{N}^*$ for the smallest of these n . Let $n \in \llbracket 1, n^*(\theta) - 1 \rrbracket$ and $X \in \mathcal{X}_\theta^n$ such that $\text{rank}(\mathbf{D}f_\theta(X)) = r_n^*(\theta)$. We have $r_n^*(\theta) < r^*(\theta)$ and using Lemma 18, there exists $x \in \mathcal{X}_\theta$ such that $\text{rank}(\mathbf{D}f_\theta(X_x)) \geq \text{rank}(\mathbf{D}f_\theta(X)) + 1$. Hence we have

$$r_{n+1}^*(\theta) \geq r_n^*(\theta) + 1. \quad (34)$$

The sequence is therefore increasing for $1 \leq n \leq n^*(\theta)$. Because of the definition $n^*(\theta)$, it is also stationary (constant) for $n^*(\theta) \leq n$.

Let us now prove the upper and lower bounds on $n^*(\theta)$.

Now, consider $n \in \llbracket 1, n^*(\theta) - 1 \rrbracket$ and $X \in \mathcal{X}_\theta^{n+1}$ such that $\text{rank}(\mathbf{D}f_\theta(X)) = r_{n+1}^*(\theta)$. Let us write $X_n \in \mathcal{X}_\theta^n$ for the matrix obtained by removing the last column $x^{(n+1)}$ from X . Lemma 17 shows that

$$\text{rank}(\mathbf{D}f_\theta(X)) \leq \text{rank}(\mathbf{D}f_\theta(X_n)) + \text{rank}(\mathbf{D}f_\theta(x^{(n+1)})).$$

Furthermore, for $x \in \mathcal{X}_\theta$, because of the size of $\mathbf{D}f_\theta(x)$, we have

$$\text{rank}(\mathbf{D}f_\theta(x)) \leq N_L. \quad (35)$$

Hence

$$\begin{aligned} r_{n+1}^*(\theta) = \text{rank}(\mathbf{D}f_\theta(X)) &\leq \text{rank}(\mathbf{D}f_\theta(X_n)) + N_L \\ &\leq r_n^*(\theta) + N_L. \end{aligned} \quad (36)$$

Grouping (34) and (36), we have

$$r_n^*(\theta) + 1 \leq r_{n+1}^*(\theta) \leq r_n^*(\theta) + N_L. \quad (37)$$

This allows an induction argument. Recall that for this part of the proposition we assume here $r^*(\theta) \geq 1$ (in the case $r^*(\theta) = 0$, the sequence $r_n^*(\theta)$ is constantly zero and $n^*(\theta) = 1$). Then, there exists $n \in \mathbb{N}^*$ and $X \in \mathcal{X}_\theta^n$ such that $\text{rank}(\mathbf{D}f_\theta(X)) \geq 1$, and by extracting $x^{(k)}$ from X such that $\text{rank}(\mathbf{D}f_\theta(x^{(k)})) \geq 1$, we find that $1 \leq r_1^*(\theta)$. Adding the inequality (35), we have $1 \leq r_1^*(\theta) \leq N_L$. Using (37) we can then show that for all $n \in \llbracket 1, n^*(\theta) \rrbracket$, we have $n \leq r_n^*(\theta) \leq nN_L$. Applying these latter inequalities to $n = n^*(\theta)$ yields

$$\frac{r^*(\theta)}{N_L} \leq n^*(\theta) \leq r^*(\theta)$$

and the proof is concluded. \square

D.2 Proof of Theorem 11

Before proving Theorem 11, we state and prove a lemma.

Lemma 19. Fix $p \in (0, 1)$ and $k \in \mathbb{N}^*$. Consider random variables X_1, \dots, X_k valued in $\{0, 1\}$ such that $\mathbb{P}(X_1 = 1) \geq p$ and for any $i \in \llbracket 2, k \rrbracket$, for any $x_1, \dots, x_{i-1} \in \{0, 1\}$, $\mathbb{P}(X_i = 1 | X_1 = x_1, \dots, X_{i-1} = x_{i-1}) \geq p$. Consider B_1, \dots, B_k independent Bernoulli random variables such that for $i \in \llbracket 1, k \rrbracket$, $\mathbb{P}(B_i = 1) = p$ and $\mathbb{P}(B_i = 0) = 1 - p$.

Then there exists a finite probability space (Ω_k, \mathbb{P}_k) and random variables $Y_1, \dots, Y_k, C_1, \dots, C_k$ from Ω_k to $\{0, 1\}$ such that

- (1) For $i \in \llbracket 1, k \rrbracket$ and $\omega \in \Omega_k$, $Y_i(\omega) \geq C_i(\omega)$;
- (2) (X_1, \dots, X_k) and (Y_1, \dots, Y_k) have the same distribution;
- (3) (B_1, \dots, B_k) and (C_1, \dots, C_k) have the same distribution.

As a consequence, for each $t \geq 0$,

$$\mathbb{P}(X_1 + \dots + X_k \leq t) \leq \mathbb{P}(B_1 + \dots + B_k \leq t). \quad (38)$$

Proof of Lemma 19. We prove the first part of the lemma by induction. Let \mathcal{L}_k correspond to the statements (1) to (3) for a given $k \in \mathbb{N}^*$.

Let us first show that \mathcal{L}_1 is true. Let $\Omega_1 = \{1, 2, 3\}$ and $\mathbb{P}_1(1) = p$, $\mathbb{P}_1(2) = \mathbb{P}(X_1 = 1) - p$ and $\mathbb{P}_1(3) = 1 - \mathbb{P}(X_1 = 1)$. Let then $Y_1(1) = Y_1(2) = 1$, $Y_1(3) = 0$, $C_1(1) = 1$, $C_1(2) = C_1(3) = 0$. With this choice of Ω_1 , \mathbb{P}_1 , Y_1 and C_1 , \mathcal{L}_1 indeed holds.

Assume now that \mathcal{L}_k holds for some $k \in \mathbb{N}^*$ and let us show that \mathcal{L}_{k+1} holds. We thus consider $\Omega_k, \mathbb{P}_k, Y_1, \dots, Y_k$ and C_1, \dots, C_k as in the statements (1) to (3). We define $\Omega_{k+1} = \Omega_k \times \{1, 2, 3\}$. Since Ω_k is finite by \mathcal{L}_k , Ω_{k+1} is also finite. For $\omega \in \Omega_k$, $i \in \llbracket 1, k \rrbracket$ and $j \in \{1, 2, 3\}$, let us define $Y_i(\omega, j) = Y_i(\omega)$ and $C_i(\omega, j) = C_i(\omega)$. Note that we use the convenient abuse of notation of defining Y_1, \dots, Y_k and C_1, \dots, C_k as both functions on Ω_k and Ω_{k+1} . For $\omega \in \Omega_k$, let us define $Y_{k+1}(\omega, 1) = Y_{k+1}(\omega, 2) = 1$, $Y_{k+1}(\omega, 3) = 0$, $C_{k+1}(\omega, 1) = 1$ and $C_{k+1}(\omega, 2) = C_{k+1}(\omega, 3) = 0$. Then Item (1) of \mathcal{L}_{k+1} is satisfied.

In order to define \mathbb{P}_{k+1} , we denote, for $\omega \in \Omega_k$,

$$\mathbb{P}_\omega = \mathbb{P}(X_{k+1} = 1 | X_1 = Y_1(\omega), \dots, X_k = Y_k(\omega)).$$

Then, for $\omega \in \Omega_k$, we define $\mathbb{P}_{k+1}(\omega, 1) = \mathbb{P}_k(\omega)p$, $\mathbb{P}_{k+1}(\omega, 2) = \mathbb{P}_k(\omega)(\mathbb{P}_\omega - p)$ and $\mathbb{P}_{k+1}(\omega, 3) = \mathbb{P}_k(\omega)(1 - \mathbb{P}_\omega)$. Using that \mathbb{P}_k is a probability measure on Ω_k , it is simple to see that \mathbb{P}_{k+1} is a probability measure on Ω_{k+1} .

Consider now $x_1, \dots, x_k \in \{0, 1\}$. If $\mathbb{P}(Y_1 = x_1, \dots, Y_k = x_k) = 0$ then, from Item (2) of \mathcal{L}_k , $\mathbb{P}(X_1 = x_1, \dots, X_k = x_k) = 0$. In this case, for all $x_{k+1} \in \{0, 1\}$, we have

$$0 = \mathbb{P}(Y_1 = x_1, \dots, Y_k = x_k, Y_{k+1} = x_{k+1}) = \mathbb{P}(X_1 = x_1, \dots, X_k = x_k, X_{k+1} = x_{k+1}).$$

Consider then the case where $\mathbb{P}(Y_1 = x_1, \dots, Y_k = x_k) > 0$. We have

$$\mathbb{P}(Y_1 = x_1, \dots, Y_k = x_k, Y_{k+1} = 1) = \sum_{j=1}^2 \sum_{\substack{\omega \in \Omega_k \text{ s.t.} \\ Y_i(\omega) = x_i \\ \text{for } i \in \llbracket 1, k \rrbracket}} \mathbb{P}_{k+1}(\omega, j) = \sum_{\substack{\omega \in \Omega_k \text{ s.t.} \\ Y_i(\omega) = x_i \\ \text{for } i \in \llbracket 1, k \rrbracket}} \mathbb{P}_k(\omega) \mathbb{P}_\omega.$$

Using now Item (2) of \mathcal{L}_k , and then using the definition of \mathbb{P}_ω , we have

$$\begin{aligned} & \mathbb{P}(Y_1 = x_1, \dots, Y_k = x_k, Y_{k+1} = 1) \\ &= \sum_{\substack{\omega \in \Omega_k \text{ s.t.} \\ Y_i(\omega) = x_i \\ \text{for } i \in \llbracket 1, k \rrbracket}} \mathbb{P}_k(\omega) \mathbb{P}_\omega \\ &= \mathbb{P}(X_1 = x_1, \dots, X_k = x_k) \frac{1}{\mathbb{P}(Y_1 = x_1, \dots, Y_k = x_k)} \sum_{\substack{\omega \in \Omega_k \text{ s.t.} \\ Y_i(\omega) = x_i \\ \text{for } i \in \llbracket 1, k \rrbracket}} \mathbb{P}_k(\omega) \mathbb{P}_\omega \\ &= \mathbb{P}(X_1 = x_1, \dots, X_k = x_k) \frac{1}{\mathbb{P}(Y_1 = x_1, \dots, Y_k = x_k)} \sum_{\substack{\omega \in \Omega_k \text{ s.t.} \\ Y_i(\omega) = x_i \\ \text{for } i \in \llbracket 1, k \rrbracket}} \mathbb{P}_k(\omega) \\ & \quad \mathbb{P}(X_{k+1} = 1 | X_1 = x_1, \dots, X_k = x_k) \\ &= \mathbb{P}(X_1 = x_1, \dots, X_k = x_k) \mathbb{P}(X_{k+1} = 1 | X_1 = x_1, \dots, X_k = x_k) \\ &= \mathbb{P}(X_1 = x_1, \dots, X_k = x_k, X_{k+1} = 1). \end{aligned}$$

We treat the case $Y_{k+1} = 0$ similarly, writing the details for the sake of completeness. We have

$$\begin{aligned} & \mathbb{P}(Y_1 = x_1, \dots, Y_k = x_k, Y_{k+1} = 0) \\ &= \mathbb{P}(X_1 = x_1, \dots, X_k = x_k) \frac{1}{\mathbb{P}(Y_1 = x_1, \dots, Y_k = x_k)} \sum_{\substack{\omega \in \Omega_k \text{ s.t.} \\ Y_i(\omega) = x_i \\ \text{for } i \in \llbracket 1, k \rrbracket}} \mathbb{P}_{k+1}(\omega, 3) \\ &= \mathbb{P}(X_1 = x_1, \dots, X_k = x_k) \frac{1}{\mathbb{P}(Y_1 = x_1, \dots, Y_k = x_k)} \sum_{\substack{\omega \in \Omega_k \text{ s.t.} \\ Y_i(\omega) = x_i \\ \text{for } i \in \llbracket 1, k \rrbracket}} \mathbb{P}_k(\omega) (1 - \mathbb{P}_\omega) \\ &= \mathbb{P}(X_1 = x_1, \dots, X_k = x_k) \frac{1}{\mathbb{P}(Y_1 = x_1, \dots, Y_k = x_k)} \sum_{\substack{\omega \in \Omega_k \text{ s.t.} \\ Y_i(\omega) = x_i \\ \text{for } i \in \llbracket 1, k \rrbracket}} \mathbb{P}_k(\omega) \\ & \quad \mathbb{P}(X_{k+1} = 0 | X_1 = x_1, \dots, X_k = x_k) \\ &= \mathbb{P}(X_1 = x_1, \dots, X_k = x_k) \mathbb{P}(X_{k+1} = 0 | X_1 = x_1, \dots, X_k = x_k) \\ &= \mathbb{P}(X_1 = x_1, \dots, X_k = x_k, X_{k+1} = 0). \end{aligned}$$

Hence Item (2) of \mathcal{L}_{k+1} holds.

Let us now show Item (3) of \mathcal{L}_{k+1} . The method is similar as above, but we give the details for completeness. Consider $c_1, \dots, c_k, c_{k+1} \in \{0, 1\}$. Using the definition of C_1, \dots, C_{k+1} , we have

$$\mathbb{P}(C_1 = c_1, \dots, C_k = c_k, C_{k+1} = c_{k+1}) = \mathbb{P}_{k+1}(A_{c_1} \times \dots \times A_{c_{k+1}}),$$

where for $i \in \llbracket 1, k+1 \rrbracket$, $A_{c_i} = \{1\}$ if $c_i = 1$ and $A_{c_i} = \{2, 3\}$ if $c_i = 0$. Note that this holds by constructing $\mathbb{P}_1, \dots, \mathbb{P}_{k+1}$, $\Omega_1, \dots, \Omega_{k+1}$ and C_1, \dots, C_{k+1} as above, by induction, with in particular $\Omega_{k+1} = \{1, 2, 3\}^{k+1}$. If $c_{k+1} = 1$, then

$$\begin{aligned} \mathbb{P}_{k+1}(A_{c_1} \times \dots \times A_{c_{k+1}}) &= \mathbb{P}_{k+1}(A_{c_1} \times \dots \times A_{c_k} \times \{1\}) \\ &= \sum_{\substack{i_1 \in A_{c_1} \\ \dots \\ i_k \in A_{c_k}}} \mathbb{P}_{k+1}((i_1, \dots, i_k, 1)) \\ &= \sum_{\substack{i_1 \in A_{c_1} \\ \dots \\ i_k \in A_{c_k}}} \mathbb{P}_k((i_1, \dots, i_k)) p \\ &= \mathbb{P}_k(C_1 = c_1, \dots, C_k = c_k) p \\ &= \mathbb{P}_k(B_1 = c_1, \dots, B_k = c_k) p \\ &= \mathbb{P}_k(B_1 = c_1, \dots, B_k = c_k, B_{k+1} = c_{k+1}), \end{aligned}$$

where we have used Item 3 of \mathcal{L}_k for the second equality before last above and the definition of B_1, \dots, B_{k+1} , in the lemma statement, for the last equality. Similarly, we have, if $c_{k+1} = 0$,

$$\begin{aligned} \mathbb{P}_{k+1}(A_{c_1} \times \dots \times A_{c_{k+1}}) &= \mathbb{P}_{k+1}(A_{c_1} \times \dots \times A_{c_k} \times \{2, 3\}) \\ &= \sum_{\substack{i_1 \in A_{c_1} \\ \dots \\ i_k \in A_{c_k}}} (\mathbb{P}_{k+1}((i_1, \dots, i_k, 2)) + \mathbb{P}_{k+1}((i_1, \dots, i_k, 3))) \\ &= \sum_{\substack{i_1 \in A_{c_1} \\ \dots \\ i_k \in A_{c_k}}} \mathbb{P}_k((i_1, \dots, i_k)) (1 - p) \\ &= \mathbb{P}_k(C_1 = c_1, \dots, C_k = c_k) (1 - p) \\ &= \mathbb{P}_k(B_1 = c_1, \dots, B_k = c_k) (1 - p) \\ &= \mathbb{P}_k(B_1 = c_1, \dots, B_k = c_k, B_{k+1} = c_{k+1}). \end{aligned}$$

Hence in all cases,

$$\mathbb{P}(C_1 = c_1, \dots, C_k = c_k, C_{k+1} = c_{k+1}) = \mathbb{P}(B_1 = c_1, \dots, B_k = c_k, B_{k+1} = c_{k+1})$$

and \mathcal{L}_{k+1} is proved. This finishes the proof by induction that \mathcal{L}_k holds for all $k \in \mathbb{N}^*$.

Let us now show (38). We have from Items (2) and (3) of \mathcal{L}_k ,

$$\begin{aligned} & \mathbb{P}(X_1 + \dots + X_k \leq t) - \mathbb{P}(B_1 + \dots + B_k \leq t) \\ &= \mathbb{P}(Y_1 + \dots + Y_k \leq t) - \mathbb{P}(C_1 + \dots + C_k \leq t) \\ &= \mathbb{E}(1_{\{Y_1 + \dots + Y_k \leq t\}} - 1_{\{C_1 + \dots + C_k \leq t\}}) \\ &\leq 0, \end{aligned}$$

because from Item (1) of \mathcal{L}_k , the random variable

$$1_{\{Y_1 + \dots + Y_k \leq t\}} - 1_{\{C_1 + \dots + C_k \leq t\}}$$

takes the values -1 or 0 . □

Proof of Theorem 11. Let us consider random vectors $x^{(1)}, \dots, x^{(n)} \in \mathbb{R}^{N_0}$, sampled as described in the statement of the theorem. Let us first show that, almost surely, for all $\ell \in \llbracket 1, n \rrbracket$, there exists $j \in \llbracket 1, m_\ell \rrbracket$ such that $((x^{(i)})_{1 \leq i \leq \ell}, \theta) \in \mathcal{O}_j^\ell$. Indeed, for any $\ell \in \llbracket 1, n \rrbracket$, since $\theta \notin \mathcal{Z}$, the set $\{X \in \mathbb{R}^{N_0 \times \ell} \mid (X, \theta) \in (\cup_{j=1}^{m_\ell} \mathcal{O}_j^\ell)^c\}$ has Lebesgue measure zero. Since the vectors $x^{(i)}$ are independent, the matrix $(x^{(i)})_{1 \leq i \leq \ell} \in \mathbb{R}^{N_0 \times \ell}$ follows the product distribution \mathcal{G}^ℓ , which is absolutely continuous with respect to Lebesgue measure of $\mathbb{R}^{N_0 \times \ell}$. Therefore, we have

$$\mathbb{P}\left(\left((x^{(i)})_{1 \leq i \leq \ell}, \theta\right) \in \left(\cup_{j=1}^{m_\ell} \mathcal{O}_j^\ell\right)^c\right) = 0,$$

and since this is true for all $\ell \in \llbracket 1, n \rrbracket$, we thus have

$$\mathbb{P}\left(\exists \ell \in \llbracket 1, n \rrbracket \text{ such that } \left((x^{(i)})_{1 \leq i \leq \ell}, \theta\right) \in \left(\cup_{j=1}^{m_\ell} \mathcal{O}_j^\ell\right)^c\right) = 0.$$

As a consequence, for the rest of the proof, up to intersecting with an event of probability one, we will assume that for all $\ell \in \llbracket 1, n \rrbracket$, there exists $j \in \llbracket 1, m_\ell \rrbracket$ such that $((x^{(i)})_{1 \leq i \leq \ell}, \theta) \in \mathcal{O}_j^\ell$.

To ease the reading, let us denote $N = n^*(\theta)$ in this proof. By definition of N , we have $r_N^*(\theta) = r^*(\theta)$, and Proposition 5 shows that there exists $j \in I_N(\theta)$ such that

$$r_N^*(\theta) = r_j^N.$$

Consequently, there exist deterministic $\tilde{x}^{(1)}, \dots, \tilde{x}^{(N)} \in \mathbb{R}^{N_0}$ such that if $\tilde{X} = (\tilde{x}^{(i)})_{1 \leq i \leq N}$, we have $(\tilde{X}, \theta) \in \mathcal{O}_j^N$ and

$$\text{rank}(\mathbf{D}f_\theta(\tilde{X})) = r^*(\theta).$$

For $\delta \in A(\theta)$, let us define the deterministic integer

$$c_\delta^* = \text{Card} \{i \in \llbracket 1, N \rrbracket \mid \tilde{x}^{(i)} \in \text{Int } D_\delta(\theta)\}.$$

We have

$$\sum_{\delta \in A(\theta)} c_\delta^* = N,$$

also up to intersecting with an event of probability one, since $(\cup_{\delta \in A(\theta)} \text{Int } D_\delta(\theta))^c = \mathcal{X}_\theta^c$ is included in a finite union of hyperplanes, as observed after (15), and thus has Lebesgue measure zero. Let us also define, for $\ell \in \llbracket 1, n \rrbracket$ and $\delta \in A(\theta)$, the random integers

$$c_\delta(\ell) = \min \left(\text{Card} \{i \in \llbracket 1, \ell \rrbracket \mid x^{(i)} \in \text{Int } D_\delta(\theta)\} , c_\delta^* \right),$$

and

$$c(\ell) = \sum_{\delta \in A(\theta)} c_\delta(\ell).$$

We have $c_\delta(\ell) \leq c_\delta^*$, for all $\delta \in A(\theta)$, thus $c(\ell) \leq N$. The sequence $c(\ell)$ is nondecreasing, and at each step, the increment $c(\ell + 1) - c(\ell)$ is either 0 or 1.

Let us first show that, almost surely, for $\ell \in \llbracket 1, n \rrbracket$,

$$\{c(\ell) = N\} \implies \left\{ \text{rank} \left(\mathbf{D}f_\theta \left((x^{(i)})_{1 \leq i \leq \ell} \right) \right) = r^*(\theta) \right\}. \quad (39)$$

Suppose indeed that, for some $\ell \in \llbracket 1, n \rrbracket$, $c(\ell) = N$. Then for all $\delta \in A(\theta)$, we have $c_\delta(\ell) = c_\delta^*$. Up to a re-ordering, we can assume that for all $i \in \llbracket 1, N \rrbracket$,

$$a(x^{(i)}, \theta) = a(\tilde{x}^{(i)}, \theta). \quad (40)$$

As assumed earlier, there exists $j' \in \llbracket 1, m_N \rrbracket$ such that $((x^{(i)})_{1 \leq i \leq N}, \theta) \in \mathcal{O}_{j'}^N$. The equality (40) and Item 2 of Theorem 1 (i) show that $j' = j$. Item 3 of Theorem 1 (i) shows that the rank is constant over \mathcal{O}_j^N , and thus $\text{rank} \left(\mathbf{D}f_\theta \left((x^{(i)})_{1 \leq i \leq N} \right) \right) = \text{rank} \left(\mathbf{D}f_\theta \left((\tilde{x}^{(i)})_{1 \leq i \leq N} \right) \right) = r^*(\theta)$. This shows (39) as desired.

Define now $\bar{c}(\ell)$ by $\bar{c}(\ell) = c(\ell)$ if $c(\ell) < N$ and by $\bar{c}(\ell) = N + (\ell - M)$ if $c(\ell) = N$, where M is the smallest index i such that $c(i) = N$. Then, for all $\ell \in \llbracket 1, n \rrbracket$, $c(\ell) = N \iff \bar{c}(\ell) \geq N$. Hence, we have

$$\begin{aligned} \mathbb{P} \left(\text{rank} \left(\mathbf{D}f_\theta \left((x^{(i)})_{1 \leq i \leq n} \right) \right) = r^*(\theta) \right) &\geq \mathbb{P}(c(n) = N) \\ &= \mathbb{P}(\bar{c}(n) \geq N). \end{aligned}$$

Thus

$$\mathbb{P} \left(\text{rank} \left(\mathbf{D}f_\theta \left((x^{(i)})_{1 \leq i \leq n} \right) \right) < r^*(\theta) \right) \leq \mathbb{P}(\bar{c}(n) < N). \quad (41)$$

Define $X_1 = \bar{c}(1)$ and $X_k = \bar{c}(k) - \bar{c}(k-1)$ for $k \in \llbracket 2, n \rrbracket$. Notice that $X_1 + \dots + X_k = \bar{c}(k)$ for all $k \geq 2$. Consider also the i.i.d. Bernoulli variables B_1, \dots, B_n from the first item of the theorem statement. We will apply Lemma 19 to p , as defined in Theorem 11, (X_1, \dots, X_n) and (B_1, \dots, B_n) . To do so we need to prove that (X_1, \dots, X_n) satisfies the hypotheses of Lemma 19. First, $\mathbb{P}(X_1 = 1)$ is the probability that $x^{(1)}$ falls into one of the sets $\text{Int } D_\delta(\theta)$ for any δ such that $c_\delta^* \geq 1$. Since there exists at least one such δ , and since the probability of falling into $\text{Int } D_\delta(\theta)$ is lower bounded by p , by definition of p , we thus have $\mathbb{P}(X_1 = 1) \geq p$. Now let us show that, for $k \in \llbracket 1, n-1 \rrbracket$, $x_1, \dots, x_k \in \{0, 1\}$,

$$\mathbb{P}(X_{k+1} = 1 \mid X_1 = x_1, \dots, X_k = x_k) \geq p.$$

Let us work conditionnally to $X_1 = x_1, \dots, X_k = x_k$. If $x_1 + \dots + x_k \geq N$, then $\bar{c}(k) = x_1 + \dots + x_k \geq N$ and thus, by definition of $\bar{c}(k+1)$ and X_{k+1} , $X_{k+1} = 1$. So $\mathbb{P}(X_{k+1} = 1 | X_1 = x_1, \dots, X_k = x_k) = 1 \geq p$.

Consider now the case $x_1 + \dots + x_k < N$, i.e. $\bar{c}(k) < N$. Conditionally to $x^{(1)}, \dots, x^{(k)}$ for which $X_1 = x_1, \dots, X_k = x_k$ with $x_1 + \dots + x_k < N$, there is at least one $\delta \in A(\theta)$ such that $c_\delta(k) < c_\delta^*$, and if $x^{(k+1)}$ falls into $\text{Int } D_\delta(\theta)$, then $X_{k+1} = 1$ because $c(k+1) = c(k) + 1$ and thus $\bar{c}(k+1) = \bar{c}(k) + 1$. Hence $\mathbb{P}(X_{k+1} = 1 | x^{(1)}, \dots, x^{(k)}) \geq p$ and thus

$$\begin{aligned} & \mathbb{P}(X_{k+1} = 1 | X_1 = x_1, \dots, X_k = x_k) \\ &= \frac{1}{\mathbb{P}(X_1 = x_1, \dots, X_k = x_k)} \mathbb{P}(X_1 = x_1, \dots, X_k = x_k, X_{k+1} = 1) \\ &= \frac{1}{\mathbb{P}(X_1 = x_1, \dots, X_k = x_k)} \mathbb{E} \left[\mathbb{E} \left[1_{\{X_1=x_1, \dots, X_k=x_k\}} 1_{\{X_{k+1}=1\}} \middle| x^{(1)}, \dots, x^{(k)} \right] \right] \\ &= \frac{1}{\mathbb{P}(X_1 = x_1, \dots, X_k = x_k)} \mathbb{E} \left[1_{\{X_1=x_1, \dots, X_k=x_k\}} \mathbb{P}(X_{k+1} = 1 | x^{(1)}, \dots, x^{(k)}) \right] \\ &\geq p. \end{aligned}$$

Hence we can apply Lemma 19. From this lemma and (41), for all $n \in \mathbb{N}$,

$$\begin{aligned} \mathbb{P}\left(\text{rank}\left(\mathbf{D}f_\theta\left((x^{(i)})_{1 \leq i \leq n}\right)\right) < r^*(\theta)\right) &\leq \mathbb{P}(\bar{c}(n) < N) \\ &= \mathbb{P}(X_1 + \dots + X_n < N) \\ &\leq \mathbb{P}(B_1 + \dots + B_n < N). \end{aligned}$$

Hence Item 1 of Theorem 11 holds. Let us now consider Item 2. The expectation of $B_1 + \dots + B_n$ is np and the variance is $np(1-p) \leq np$. Hence using Chebyshev's inequality, for $np \geq 2N$

$$\begin{aligned} \mathbb{P}\left(\text{rank}\left(\mathbf{D}f_\theta\left((x^{(i)})_{1 \leq i \leq n}\right)\right) < r^*(\theta)\right) &\leq \mathbb{P}(B_1 + \dots + B_n < N) \\ &= \mathbb{P}(B_1 + \dots + B_n - np < N - np) \\ &\leq \frac{np}{(np - N)^2} \\ &\leq \frac{4}{np}. \end{aligned}$$

This concludes the proof. □

References

[Achour u. a. 2022] ACHOUR, El M. ; MALGOUYRES, François ; GERCHINOVITZ, Sébastien: The loss landscape of deep linear neural networks: a second-order analysis. In: *arXiv preprint arXiv:2107.13289* (2022)

- [Alon u. a. 1997] ALON, Noga ; BEN-DAVID, Shai ; CESA-BIANCHI, Nicolo ; HAUSSLER, David: Scale-sensitive dimensions, uniform convergence, and learnability. In: *Journal of the ACM (JACM)* 44 (1997), Nr. 4, S. 615–631
- [Anthony und Bartlett 2009] ANTHONY, Martin ; BARTLETT, Peter L.: *Neural Network Learning: Theoretical Foundations*. Cambridge University Press, 2009
- [Arora u. a. 2019] ARORA, Sanjeev ; COHEN, Nadav ; HU, Wei ; LUO, Yuping: Implicit regularization in deep matrix factorization. In: *Advances in Neural Information Processing Systems*, 2019, S. 7413–7424
- [Arora u. a. 2018] ARORA, Sanjeev ; GE, Rong ; NEYSHABUR, Behnam ; ZHANG, Yi: Stronger Generalization Bounds for Deep Nets via a Compression Approach. In: *Proceedings of the 35th International Conference on Machine Learning* Bd. 80, 2018, S. 254–263
- [Bartlett u. a. 2017] BARTLETT, Peter L. ; FOSTER, Dylan J. ; TELGARSKY, Matus J.: Spectrally-normalized margin bounds for neural networks. In: *Advances in neural information processing systems* 30 (2017)
- [Bartlett u. a. 2019] BARTLETT, Peter L. ; HARVEY, Nick ; LIAW, Christopher ; MEHRABIAN, Abbas: Nearly-tight VC-dimension and Pseudodimension Bounds for Piecewise Linear Neural Networks. In: *Journal of Machine Learning Research* 20 (2019), Nr. 63, S. 1–17
- [Bartlett und Long 1998] BARTLETT, Peter L. ; LONG, Philip M.: Prediction, learning, uniform convergence, and scale-sensitive dimensions. In: *Journal of Computer and System Sciences* 56 (1998), Nr. 2, S. 174–190
- [Bartlett u. a. 2020] BARTLETT, Peter L. ; LONG, Philip M. ; LUGOSI, Gábor ; TSIGLER, Alexander: Benign overfitting in linear regression. In: *Proceedings of the National Academy of Sciences* 117 (2020), Nr. 48, S. 30063–30070
- [Bartlett u. a. 1998] BARTLETT, Peter L. ; MAIOROV, Vitaly ; MEIR, Ron: Almost Linear VC-Dimension Bounds for Piecewise Polynomial Networks. In: *Neural Computation* 10 (1998), Nr. 8, S. 2159–2173
- [Belkin 2021] BELKIN, Mikhail: Fit without fear: remarkable mathematical phenomena of deep learning through the prism of interpolation. In: *Acta Numerica* 30 (2021), S. 203–248
- [Bona-Pellissier u. a. 2023a] BONA-PELLISSIER, Joachim ; BACHOC, François ; MALGOUYRES, François: Parameter identifiability of a deep feedforward ReLU neural network. In: *Machine Learning* 112 (2023), Nr. 11, S. 4431–4493
- [Bona-Pellissier u. a. 2022] BONA-PELLISSIER, Joachim ; MALGOUYRES, François ; BACHOC, François: Local Identifiability of Deep ReLU Neural Networks: the Theory. In: *Advances in Neural Information Processing Systems*, 2022

- [Bona-Pellissier u. a. 2023b] BONA-PELLISSIER, Joachim ; MALGOUYRES, François ; BACHOC, François: <https://github.com/JoachimBP/Functional-dimension>. 2023.
– Code of the experiments of this article
- [Camuto u. a. 2021] CAMUTO, Alexander ; DELIGIANNIDIS, George ; ERDOGDU, Murat A. ; GURBUZBALABAN, Mert ; SIMSEKLI, Umut ; ZHU, Lingjiong: Fractal structure and generalization properties of stochastic optimization algorithms. In: *Advances in Neural Information Processing Systems* 34 (2021)
- [Carlini u. a. 2020] CARLINI, Nicholas ; JAGIELSKI, Matthew ; MIRONOV, Ilya: Cryptanalytic extraction of neural network models. In: *Annual International Cryptology Conference* Springer (Veranst.), 2020, S. 189–218
- [Cha u. a. 2021] CHA, Junbum ; CHUN, Sanghyuk ; LEE, Kyungjae ; CHO, Han-Cheol ; PARK, Seunghyun ; LEE, Yunsung ; PARK, Sungrae: SWAD: Domain generalization by seeking flat minima. In: *Advances in Neural Information Processing Systems* 34 (2021), S. 22405–22418
- [Chaudhari u. a. 2019] CHAUDHARI, Pratik ; CHOROMANSKA, Anna ; SOATTO, Stefano ; LECUN, Yann ; BALDASSI, Carlo ; BORGS, Christian ; CHAYES, Jennifer ; SAGUN, Levent ; ZECCHINA, Riccardo: Entropy-SGD: Biasing gradient descent into wide valleys. In: *Journal of Statistical Mechanics: Theory and Experiment* 2019 (2019), Nr. 12, S. 124018
- [Colomboni u. a. 2023] COLOMBONI, Roberto ; ESPOSITO, Emmanuel ; PAUDICE, Andrea: An Improved Uniform Convergence Bound with Fat-Shattering Dimension. In: *arXiv preprint arXiv:2307.06644* (2023)
- [Cooper 2021] COOPER, Yaim: Global minima of overparameterized neural networks. In: *SIAM Journal on Mathematics of Data Science* 3 (2021), Nr. 2, S. 676–691
- [Dinh u. a. 2017] DINH, Laurent ; PASCANU, Razvan ; BENGIO, Samy ; BENGIO, Yoshua: Sharp minima can generalize for deep nets. In: *International Conference on Machine Learning*, 2017, S. 1019–1028
- [Du u. a. 2019] DU, Simon ; LEE, Jason ; LI, Haochuan ; WANG, Liwei ; ZHAI, Xiyu: Gradient descent finds global minima of deep neural networks. In: *International conference on machine learning*, 2019, S. 1675–1685
- [Erlich u. a. 1997] ERLICH, Yossi ; CHAZAN, Dan ; PETRACK, Scott ; LEVY, Avraham: Lower Bound on VC-Dimension by Local Shattering. In: *Neural Computation* 9 (1997), Nr. 4, S. 771–776
- [Foret u. a. 2021] FORET, Pierre ; KLEINER, Ariel ; MOBAHI, Hossein ; NEYSHABUR, Behnam: Sharpness-aware Minimization for Efficiently Improving Generalization. In: *International Conference on Learning Representations*, 2021

- [Ghorbani u. a. 2019] GHORBANI, Behrooz ; KRISHNAN, Shankar ; XIAO, Ying: An investigation into neural net optimization via Hessian eigenvalue density. In: *International Conference on Machine Learning*, 2019, S. 2232–2241
- [Gidel u. a. 2019] GIDEL, Gauthier ; BACH, Francis ; LACOSTE-JULIEN, Simon: Implicit regularization of discrete gradient dynamics in linear neural networks. In: *Advances in Neural Information Processing Systems*, 2019, S. 3202–3211
- [Gissin u. a. 2019] GISSIN, Daniel ; SHALEV-SHWARTZ, Shai ; DANIELY, Amit: The Implicit Bias of Depth: How Incremental Learning Drives Generalization. In: *International Conference on Learning Representations*, 2019
- [Golowich u. a. 2018] GOLOWICH, Noah ; RAKHLIN, Alexander ; SHAMIR, Ohad: Size-independent sample complexity of neural networks. In: *Conference On Learning Theory*, 2018, S. 297–299
- [Grigsby u. a. 2023] GRIGSBY, Elisenda ; LINDSEY, Kathryn ; ROLNICK, David: Hidden symmetries of ReLU networks. In: *International Conference on Machine Learning*, 2023, S. 11734–11760
- [Grigsby und Lindsey 2022] GRIGSBY, J E. ; LINDSEY, Kathryn: On transversality of bent hyperplane arrangements and the topological expressiveness of ReLU neural networks. In: *SIAM Journal on Applied Algebra and Geometry* 6 (2022), Nr. 2, S. 216–242
- [Grigsby u. a. 2022] GRIGSBY, J E. ; LINDSEY, Kathryn ; MEYERHOFF, Robert ; WU, Chenxi: Functional dimension of feedforward ReLU neural networks. In: *arXiv preprint arXiv:2209.04036* (2022)
- [Grohs und Kutyniok 2022] GROHS, Philipp (Hrsg.) ; KUTYNIOK, Gitta (Hrsg.): *Mathematical Aspects of Deep Learning*. Cambridge University Press, 2022
- [Gur-Ari u. a. 2018] GUR-ARI, Guy ; ROBERTS, Daniel A. ; DYER, Ethan: Gradient descent happens in a tiny subspace. In: *arXiv preprint arXiv:1812.04754* (2018)
- [Hanin und Rolnick 2019] HANIN, Boris ; ROLNICK, David: Complexity of linear regions in deep networks. In: *International Conference on Machine Learning*, 2019, S. 2596–2604
- [Harvey u. a. 2017] HARVEY, Nick ; LIAW, Christopher ; MEHRABIAN, Abbas: Nearly-tight VC-dimension bounds for piecewise linear neural networks. In: *Conference on learning theory*, 2017, S. 1064–1068
- [Hochreiter und Schmidhuber 1997] HOCHREITER, Sepp ; SCHMIDHUBER, Jürgen: Flat minima. In: *Neural computation* 9 (1997), Nr. 1, S. 1–42
- [Imaizumi und Schmidt-Hieber 2023] IMAIZUMI, Masaaki ; SCHMIDT-HIEBER, Johannes: On generalization bounds for deep networks based on loss surface implicit regularization. In: *IEEE Transactions on Information Theory* 69 (2023), Nr. 2

- [Keskar u. a. 2017] KESKAR, Nitish S. ; MUDIGERE, Dheevatsa ; NOCEDAL, Jorge ; SMELYANSKIY, Mikhail ; TANG, Ping Tak P.: On Large-Batch Training for Deep Learning: Generalization Gap and Sharp Minima. In: *International Conference on Learning Representations*, 2017
- [Li u. a. 2018] LI, Chunyuan ; FARKHOOR, Heerad ; LIU, Rosanne ; YOSINSKI, Jason: Measuring the Intrinsic Dimension of Objective Landscapes. In: *International Conference on Learning Representations*, 2018
- [Maass 1994] MAASS, Wolfgang: Neural nets with superlinear VC-dimension. In: *Neural Computation* 6 (1994), Nr. 5, S. 877–884
- [Mityagin 2020] MITYAGIN, Boris: The zero set of a real analytic function. In: *Math Notes* 107 (2020), S. 529–530
- [Montufar u. a. 2014] MONTUFAR, Guido F. ; PASCANU, Razvan ; CHO, Kyunghyun ; BENGIO, Yoshua: On the number of linear regions of deep neural networks. In: *Advances in neural information processing systems* 27 (2014)
- [Neyshabur u. a. 2017] NEYSHABUR, Behnam ; BHOJANAPALLI, Srinadh ; MCALLESTER, David ; SREBRO, Nati: Exploring generalization in deep learning. In: *Advances in neural information processing systems* 30 (2017)
- [Neyshabur u. a. 2015a] NEYSHABUR, Behnam ; SALAKHUTDINOV, Russ R. ; SREBRO, Nati: Path-SGD: Path-normalized optimization in deep neural networks. In: *Advances in neural information processing systems* 28 (2015)
- [Neyshabur u. a. 2015b] NEYSHABUR, Behnam ; TOMIOKA, Ryota ; SREBRO, Nathan: Norm-based capacity control in neural networks. In: *Conference on learning theory*, 2015, S. 1376–1401
- [Nguyen und Hein 2017] NGUYEN, Quynh ; HEIN, Matthias: The loss surface of deep and wide neural networks. In: *International conference on machine learning*, 2017, S. 2603–2612
- [Petersen u. a. 2021] PETERSEN, Philipp ; RASLAN, Mones ; VOIGTLAENDER, Felix: Topological properties of the set of functions generated by neural networks of fixed size. In: *Foundations of computational mathematics* 21 (2021), Nr. 2, S. 375–444
- [Petzka u. a. 2020] PETZKA, Henning ; TRIMMEL, Martin ; SMINCHISESCU, Cristian: Notes on the Symmetries of 2-Layer ReLU-Networks. In: *Proceedings of the Northern Lights Deep Learning Workshop* Bd. 1, 2020, S. 6–6
- [Raghu u. a. 2017] RAGHU, Maithra ; POOLE, Ben ; KLEINBERG, Jon ; GANGULI, Surya ; SOHL-DICKSTEIN, Jascha: On the expressive power of deep neural networks. In: *International Conference on Machine Learning*, 2017, S. 2847–2854

- [Razin und Cohen 2020] RAZIN, Noam ; COHEN, Nadav: Implicit Regularization in Deep Learning May Not Be Explainable by Norms. In: *Advances in Neural Information Processing Systems* 33 (2020), S. 21174–21187
- [Rolnick und Kording 2020] ROLNICK, David ; KORDING, Konrad: Reverse-engineering deep ReLU networks. In: *Proceedings of the 37th International Conference on Machine Learning* Bd. 119, 2020, S. 8178–8187
- [Safran u. a. 2021] SAFRAN, Itay M. ; YEHUDAI, Gilad ; SHAMIR, Ohad: The effects of mild over-parameterization on the optimization landscape of shallow ReLU neural networks. In: *Conference on Learning Theory*, 2021, S. 3889–3934
- [Sagun u. a. 2016] SAGUN, Levent ; BOTTOU, Leon ; LECUN, Yann: Eigenvalues of the Hessian in deep learning: Singularity and beyond. In: *arXiv preprint arXiv:1611.07476* (2016)
- [Saxe u. a. 2019] SAXE, Andrew M. ; MCCLELLAND, James L. ; GANGULI, Surya: A mathematical theory of semantic development in deep neural networks. In: *Proceedings of the National Academy of Sciences* 116 (2019), Nr. 23, S. 11537–11546
- [Sonoda u. a. 2021] SONODA, Sho ; ISHIKAWA, Isao ; IKEDA, Masahiro: Ghosts in Neural Networks: Existence, Structure and Role of Infinite-Dimensional Null Space. In: *arXiv preprint arXiv:2106.04770* (2021)
- [Soudry und Carmon 2016] SOUDRY, Daniel ; CARMON, Yair: No bad local minima: Data independent training error guarantees for multilayer neural networks. In: *arXiv preprint arXiv:1605.08361* (2016)
- [Stock 2021] STOCK, Pierre: *Efficiency and Redundancy in Deep Learning Models : Theoretical Considerations and Practical Applications*, Université de Lyon, PhD thesis, April 2021. – URL <https://tel.archives-ouvertes.fr/tel-03208517>
- [Stock und Gribonval 2022] STOCK, Pierre ; GRIBONVAL, Rémi: An embedding of ReLU networks and an analysis of their identifiability. In: *Constructive Approximation* (2022)
- [Suzuki u. a. 2020] SUZUKI, Taiji ; ABE, Hiroshi ; NISHIMURA, Tomoaki: Compression based bound for non-compressed network: unified generalization error analysis of large compressible deep neural network. In: *International Conference on Learning Representations*, 2020
- [Vapnik und Chervonenkis 1971] VAPNIK, Vladimir N. ; CHERVONENKIS, Alexey Y.: On the Uniform Convergence of Relative Frequencies of Events to Their Probabilities. In: *Theory of Probability and its Applications* 16 (1971), Nr. 2, S. 264
- [Wu u. a. 2017] WU, Lei ; ZHU, Zhanxing ; E, Weinan: Towards understanding generalization of deep learning: Perspective of loss landscapes. In: *Workshop 'Principled Approaches to Deep Learning', ICML* (2017)

[Yi u. a. 2019] YI, Mingyang ; MENG, Qi ; CHEN, Wei ; MA, Zhi-ming ; LIU, Tie-Yan: Positively scale-invariant flatness of ReLU neural networks. In: *arXiv preprint arXiv:1903.02237* (2019)

[Zhang u. a. 2021] ZHANG, Chiyuan ; BENGIO, Samy ; HARDT, Moritz ; RECHT, Benjamin ; VINYALS, Oriol: Understanding deep learning (still) requires rethinking generalization. In: *Communications of the ACM* 64 (2021), Nr. 3, S. 107–115