Lederer, Johannes

Statistical guarantees for sparse deep learning

AStA Advances in Statistical Analysis

Provided in Cooperation with:
Springer Nature

Suggested Citation: Lederer, Johannes (2023) : Statistical guarantees for sparse deep learning, AStA Advances in Statistical Analysis, ISSN 1863-818X, Springer, Berlin, Heidelberg, Vol. 108, Iss. 2, pp. 231-258,
https://doi.org/10.1007/s10182-022-00467-3

This Version is available at:
https://hdl.handle.net/10419/307763

AStA Advances in Statistical Analysis (2024) 108:231–258
https://doi.org/10.1007/s10182-022-00467-3

ORIGINAL PAPER

Statistical guarantees for sparse deep learning

Johannes Lederer[1]

Received: 21 April 2022 / Accepted: 17 December 2022 / Published online: 24 January 2023

Abstract
Neural networks are becoming increasingly popular in applications, but our mathematical understanding of their potential and limitations is still limited. In this paper, we further this understanding by developing statistical guarantees for sparse deep learning. In contrast to previous work, we consider different types of sparsity, such as few active connections, few active nodes, and other norm-based types of sparsity. Moreover, our theories cover important aspects that previous theories have neglected, such as multiple outputs, regularization, and ℓ_2-loss. The guarantees have a mild dependence on network widths and depths, which means that they support the application of sparse but wide and deep networks from a statistical perspective. Some of the concepts and tools that we use in our derivations are uncommon in deep learning and, hence, might be of additional interest.

Keywords Sparsity · Regularization · Oracle inequalities · High-dimensionality

1 Introduction

Sparsity reduces network complexities and, consequently, lowers the demands on memory and computation, reduces overfitting, and improves interpretability (Changpinyo et al. 2017; Han et al. 2016; Kim et al. 2016; Liu et al. 2015; Wen et al. 2016). Sparsity is at the heart of many current techniques in deep learning, such as dropouts (Srivastava et al. 2014), lottery tickets (Frankle and Carbin 2019), augmenting small networks (Ash 1989; Bello 1992), pruning large networks (Simonyan and Zisserman 2015; Han et al. 2016), sparsity constraints (Ledent et al. 2019; Neyshabur et al. 2015; Schmidt-Hieber 2020), and sparsity regularization (Taheri et al. 2021).

The many empirical observations of the benefits of sparsity have sparked interest in mathematical support in the form of statistical theories. Two current approaches are based on Rademacher complexities (Bartlett and Mendelson 2002; Neyshabur et al. 2015) and ideas from nonparametric statistics (Schmidt-Hieber 2020),

✉ Johannes Lederer
johannes.lederer@rub.de

[1] Department of Mathematics, Ruhr-University Bochum, Bochum, Germany

respectively. While their results provide important support for sparse deep learning, they still have major limitations: The first approach is restricted to bounded loss functions (which excludes the ℓ_2-loss, for example), is either restricted to a simple form of sparsity (which we will call "connection sparsity" later) or suffers from an exponential dependence on the number of layers (which contradicts the current interest in very deep networks), caters to constraints rather than regularization (which is the predominant implementation in practice), and is limited to a single output node and ReLU activation. The second approach is restricted to ℓ_0-constraints (which are infeasible in practice), assumes bounded weights, and is also limited to a single output node and ReLU activation. In short, while some progress in the statistical understanding of sparse deep learning has been made already, many aspects have not yet been considered.

The goal of this paper is to establish a statistical theory that accounts for these missing aspects. For this, we follow a third, very recent approach introduced in Taheri et al. (2021). This approach is based on ideas from high-dimensional statistics and empirical-process theory (Lederer 2022). The main feature of their results is that they apply to ℓ_2-loss, regularization instead of constraints, and a variety of activation functions. But they still miss some aspects, such as the inclusion of more complex notions of sparsity (we will speak of "node sparsity" later) and the restriction to a single output node. Moreover, their estimator involves an additional, arguably unnatural parameter.

In this paper, we remove these limitations from Taheri et al. (2021). We focus on regression-type settings with layered, feedforward neural networks. The estimators under consideration consist of a standard least-squares estimator with regularizers that induce different types of sparsity—without the need for an additional parameter. We then derive prediction and generalization guarantees by using techniques from high-dimensional statistics (Dalalyan et al. 2017) and empirical-process theory (van de Geer 2000). In the case of sub-Gaussian noise, we find the rates

$$\sqrt{\frac{l\big(\log[mn\overline{p}]\big)^3}{n}} \quad \text{and} \quad \sqrt{\frac{ml\underline{p}\big(\log[mn\overline{p}]\big)^3}{n}}$$

for the connection-sparse and node-sparse estimators (see the following section for the notions of sparsity), respectively, where l is the number of hidden layers, m the number of output nodes, n the number of samples, $\overline{p}$ the total number of parameters, and $\underline{p}$ the maximal width of the network. The rates suggest that sparsity-inducing approaches can provide accurate prediction even in very wide (with connection sparsity) and very deep (with either type of sparsity) networks while, at the same time, ensuring low network complexities. These findings underpin the current trend toward sparse but wide and especially deep networks from a statistical perspective. More generally speaking, our paper complements the existing statistical theories for sparse deep learning with new results, and it refines the techniques that were introduced in (Taheri et al. 2021).

Outline of the paper Section 2 recapitulates the notions of connection and node sparsity and introduces the corresponding deep learning framework and estimators.

Section 3 confirms the empirically observed accuracies of connection- and node-sparse estimation in theory. Section 4 discusses connections of our theoretical results and weight initialization. Section 5 summarizes the key features and limitations of our work. The Appendix contains all proofs.

2 Connection- and node-sparse deep learning

We consider data $(\boldsymbol{y}_1, \boldsymbol{x}_1), \dots, (\boldsymbol{y}_n, \boldsymbol{x}_n) \in \mathbb{R}^m \times \mathbb{R}^d$ that are related via

$$\boldsymbol{y}_i = \boldsymbol{g}_*[\boldsymbol{x}_i] + \boldsymbol{u}_i \qquad \text{for } i \in \{1, \dots, n\} \tag{1}$$

for an unknown data-generating function $\boldsymbol{g}_* : \mathbb{R}^d \to \mathbb{R}^m$ and unknown, random noise $\boldsymbol{u}_1, \dots, \boldsymbol{u}_n \in \mathbb{R}^m$. We allow all aspects, namely $\boldsymbol{y}_i$, $\boldsymbol{g}_*$, $\boldsymbol{x}_i$, and $\boldsymbol{u}_i$, to be unbounded. Our goal is to model the data-generating function with a feedforward neural network of the form

$$\boldsymbol{g}_{\boldsymbol{\Theta}}[\boldsymbol{x}] := \Theta^l \boldsymbol{f}^l \left[\Theta^{l-1} \cdots \boldsymbol{f}^1[\Theta^0 \boldsymbol{x}]\right] \qquad \text{for } \boldsymbol{x} \in \mathbb{R}^d \tag{2}$$

indexed by the parameter space $\mathcal{M} := \{\boldsymbol{\Theta} = (\Theta^l, \dots, \Theta^0) : \Theta^j \in \mathbb{R}^{p^{j+1} \times p^j}\}$. The functions $\boldsymbol{f}^j : \mathbb{R}^{p^j} \to \mathbb{R}^{p^j}$ are called the activation functions (Lederer 2021), and $p^0 := d$ and $p^{l+1} := m$ are called the input and output dimensions, respectively. The depth of the network is l, the maximal width is $\underline{p} := \max_{j \in \{0,\dots,l-1\}} p^{j+1}$, and the total number of parameters is $\overline{p} := \sum_{j=0}^{l} p^{j+1} p^j$.

In practice, the total number of parameters often rivals or exceeds the number of samples: $\overline{p} \approx n$ or $\overline{p} \gg n$. We then speak of *high dimensionality*. A common technique for avoiding overfitting in high-dimensional settings is regularization that induces additional structures, such as *sparsity*. Sparsity has the interesting side-effect of reducing the networks' complexities, which can facilitate interpretations and reduce demands on energy and memory. Three common notions of sparsity are *connection sparsity,* which means that there is only a small number of nonzero connections between nodes, *node sparsity,* which means that there is only a small number of active nodes (Alvarez and Salzmann 2016; Changpinyo et al. 2017; Feng and Simon 2017; Kim et al. 2016; Lee et al. 2008; Liu et al. 2015; Nie et al. 2015; Scardapane et al. 2017; Wen et al. 2016), and *layer sparsity*, which means that there is only a small number of active layers (Hebiri and Lederer 2020).

In the following, we focus on connection- and node sparsity. Our first sparse estimator is

$$\widehat{\boldsymbol{\Theta}}_{\text{con}} \in \arg\min_{\boldsymbol{\Theta} \in \mathcal{M}_1} \left\{ \sum_{i=1}^{n} \left\| \boldsymbol{y}_i - \boldsymbol{g}_{\boldsymbol{\Theta}}[\boldsymbol{x}_i] \right\|_2^2 + r_{\text{con}} |||\Theta^l|||_1 \right\} \tag{3}$$

for a tuning parameter $r_{\text{con}} \in [0, \infty)$, a nonempty set of parameters

$$\mathcal{M}_1 \subset \left\{ \boldsymbol{\Theta} \in \mathcal{M} : \max_{j \in \{0,\dots,l-1\}} |||\Theta^j|||_1 \leq 1 \right\},$$

and the ℓ_1-norm

$$|||\Theta^j|||_1 := \sum_{i=1}^{p^{j+1}} \sum_{k=1}^{p^j} |(\Theta^j)_{ik}| \quad \text{for } j \in \{0, \dots, l\},\ \Theta^j \in \mathbb{R}^{p^{j+1} \times p^j}.$$

This estimator is an analog of the lasso estimator in linear regression (Tibshirani 1996). It induces sparsity on the level of connections: the larger the tuning parameter r_{con}, the fewer connections among the nodes.

Deep learning with ℓ_1-regularization has become common in theory and practice (Kim et al. 2016; Taheri et al. 2021). Our estimator (3) specifies one way to formulate this type of regularization. The estimator is indeed a regularized estimator (rather than a constraint estimator), because the complexity is regulated entirely through the tuning parameter r_{con} in the objective function (rather than through a tuning parameter in the set over which the objective function is optimized). But ℓ_1 -regularization could also be formulated slightly differently. For example, one could consider the estimators

$$\overline{\boldsymbol{\Theta}}_{\text{con}} \in \arg\min_{\boldsymbol{\Theta} \in \mathcal{M}} \left\{ \sum_{i=1}^{n} \|\boldsymbol{y}_i - \boldsymbol{g}_{\boldsymbol{\Theta}}[\boldsymbol{x}_i]\|_2^2 + r_{\text{con}} \prod_{j=0}^{l} |||\Theta^j|||_1 \right\} \tag{4}$$

or

$$\widetilde{\boldsymbol{\Theta}}_{\text{con}} \in \arg\min_{\boldsymbol{\Theta} \in \mathcal{M}} \left\{ \sum_{i=1}^{n} \|\boldsymbol{y}_i - \boldsymbol{g}_{\boldsymbol{\Theta}}[\boldsymbol{x}_i]\|_2^2 + r_{\text{con}} \sum_{j=0}^{l} |||\Theta^j|||_1 \right\}. \tag{5}$$

The differences among the estimators (3)–(5) are small: for example, our theory can be adjusted for (4) with almost no changes of the derivations. The differences among the estimators mainly concern the normalizations of the parameters; we illustrate this in the following proposition.

Proposition 1 (Scaling of Norms) *Assume that the all-zeros parameter* $(\mathbf{0}_{p^{l+1} \times p^l}, \dots, \mathbf{0}_{p^1 \times p^0}) \in \mathcal{M}_1$ *is neither a solution of* (3) *nor of* (5)*, that* $r_{\text{con}} > 0$*, and that the activation functions are nonnegative homogenous*: $\boldsymbol{f}^j[a\boldsymbol{b}] = a\boldsymbol{f}^j[\boldsymbol{b}]$ *for all* $j \in \{1, \dots, l\}$, $a \in [0, \infty)$, *and* $\boldsymbol{b} \in \mathbb{R}^{p^j}$*. Then,* $|||(\widehat{\Theta}_{\text{con}})^0|||_1, \dots, |||(\widehat{\Theta}_{\text{con}})^{l-1}|||_1 = 1$ (*concerns the inner layers*) *for all solutions of* (3)*, while* $|||(\widetilde{\Theta}_{\text{con}})^0|||_1 = \cdots = |||(\widetilde{\Theta}_{\text{con}})^l|||_1$ (*concerns all layers*) *for at least one solution of* (5)*.*

In brief, the goal of our paper is not to promote a new way of implementing sparsity in practice but to reproduce practical implementations as accurately as possible in theory.

Another way to formulate ℓ_1-regularization was proposed in Taheri et al. (2021): they reparametrize the networks through a scale parameter and a constraint version of $\mathcal{M}$ and then to focus the regularization on the scale parameter only. Our above-stated estimator (3) is more elegant in that it avoids the reparametrization and the additional parameter.

The factor $|||\Theta^l|||_1$ in the regularization term of (3) measures the complexity of the network over the set $\mathcal{M}_1$, and the factor r_{con} regulates the complexity of the resulting

estimator. This provides a convenient lever for data-adaptive complexity regularization through well-established calibration schemes for the tuning parameter, such as cross-validation. This practical aspect is an advantage of regularized formulations like ours as compared to constraint estimation over sets with a predefined complexity.

The constraints in the set $\mathcal{M}_1$ of the estimator (3) can also retain the expressiveness of the full parameterization that corresponds to the set $\mathcal{M}$: for example, assuming again nonnegative-homogeneous activation, one can check that for every $\boldsymbol{\Gamma} \in \mathcal{M}$, there is a $\boldsymbol{\Gamma}' \in \{\boldsymbol{\Theta} \in \mathcal{M} : \max_{j\in\{0,\dots,l-1\}} |||\Theta^j|||_1 \leq 1\}$ such that $\boldsymbol{g}_{\boldsymbol{\Gamma}} = \boldsymbol{g}_{\boldsymbol{\Gamma}'}$—cf. (Taheri et al. 2021, Proposition 1). In contrast, existing theories on neural networks often require the parameter space to be bounded, which limits the expressiveness of the networks.

Our regularization approach is, therefore, closer to practical setups than constraint approaches. The price is that to develop prediction theories, we have to use different tools than those typically used in theoretical deep learning. For example, we cannot use established risk bounds such as (Bartlett and Mendelson 2002, Theorem 8) (because Rademacher complexities over classes of unbounded functions are unbounded) or (Lederer 2020a, Theorem 1) (because our loss function is not Lipschitz continuous) or established concentration bounds such as McDiarmid's inequality in (McDiarmid 1989, Lemma (3.3)) (because that would require a bounded loss). We instead invoke ideas from high-dimensional statistics, prove Lipschitz properties for neural networks, and use empirical-process theory, specifically concentration inequalities that are based on chaining (see the Appendix).

Our second estimator is

$$\widehat{\boldsymbol{\Theta}}_{\text{node}} \in \arg\min_{\boldsymbol{\Theta}\in\mathcal{M}_{2,1}} \left\{ \sum_{i=1}^{n} \left\| \boldsymbol{y}_i - \boldsymbol{g}_{\boldsymbol{\Theta}}[\boldsymbol{x}_i] \right\|_2^2 + r_{\text{node}} |||\Theta^l|||_{2,1} \right\} \tag{6}$$

for a tuning parameter $r_{\text{node}} \in [0, \infty)$, a nonempty set of parameters

$$\mathcal{M}_{2,1} \subset \left\{ \boldsymbol{\Theta} \in \mathcal{M} : \max_{j\in\{0,\dots,l-1\}} |||\Theta^j|||_{2,1} \leq 1 \right\},$$

and the ℓ_2/ℓ_1-norm

$$|||\Theta^j|||_{2,1} := \sum_{k=1}^{p^j} \sqrt{\sum_{i=1}^{p^{j+1}} |(\Theta^j)_{ik}|^2}$$
$$\text{for } j \in \{0, \dots, l-1\},\ \Theta^j \in \mathbb{R}^{p^{j+1}\times p^j}.$$

This estimator is an analog of the group-lasso estimator in linear regression (Bakin 1999). Again, to avoid ambiguities in the regularization, our formulation is slightly different from the standard formulations in the literature, but the fact that group-lasso regularizers leads to node-sparse networks has been discussed extensively before (Alvarez and Salzmann 2016; Liu et al. 2015; Scardapane et al. 2017): the larger the tuning parameter r_{node}, the fewer active nodes in the network.

The above-stated comments about the specific form of the connection-sparse estimator also apply to the node-sparse estimator.

An illustration of connection and node sparsity is given in Fig. 1. Connection-sparse networks have only a small number of active connections between nodes (left panel of Fig. 1); node-sparse networks have inactive nodes, that is, completely unconnected nodes (right panel of Fig. 1). The two notions of sparsity are connected: for example, connection sparsity can render entire nodes inactive "by accident" (see the layer that follows the input layer in the left panel of the figure). In general, node sparsity is the weaker assumption, because it allows for highly connected nodes; this observation is reflected in the theoretical guarantees in the following section.

The optimal network architecture for given data (such as the optimal width) is hardly known beforehand in a data analysis. A main feature of sparsity-inducing regularization is, therefore, that it adjusts parts of the network architecture to the data. In other words, sparsity-inducing regularization is a data-driven approach to adapting the complexity of the network.

While versions of the estimators (3) and (6) are popular in deep learning, statistical analyses, especially of node-sparse deep learning, are scarce. Such a statistical analysis is, therefore, the goal of the following section.

3 Statistical prediction guarantees

We now develop statistical guarantees for the sparse estimators described above. The guarantees are formulated in terms of the squared *average (in-sample) prediction error*

$$\text{err}[\boldsymbol{\Theta}] := \frac{1}{n}\sum_{i=1}^{n} \left\| g_*[\boldsymbol{x}_i] - g_{\boldsymbol{\Theta}}[\boldsymbol{x}_i] \right\|_2^2 \qquad \text{for } \boldsymbol{\Theta} \in \mathcal{M},$$

which is a measure for how well the network $g_{\boldsymbol{\Theta}}$ fits the unknown function g_* (which does not need to be a neural network) on the data at hand, and in terms of the *prediction risk* (or *generalization error*) for a new sample $(\boldsymbol{y}, \boldsymbol{x})$ that has the same distribution as the original data

$$\text{risk}[\boldsymbol{\Theta}] := E_{\boldsymbol{y},\boldsymbol{x}} \|\boldsymbol{y} - g_{\boldsymbol{\Theta}}[\boldsymbol{x}]\|_2^2 \qquad \text{for } \boldsymbol{\Theta} \in \mathcal{M},$$

which measures how well the network $g_{\boldsymbol{\Theta}}$ can predict a new sample. We first study the prediction error, because it is agnostic to the distribution of the input data; in the end, we then translate the bounds for the prediction error into bounds for the generalization error.

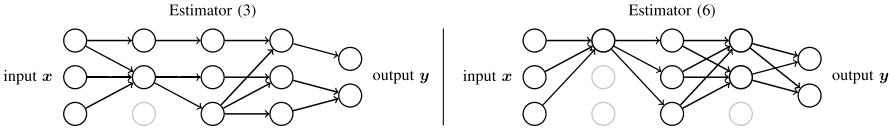

Fig. 1 exemplary networks produced by the connection-sparse estimator (3) and the node-sparse estimator (6)

We first observe that the networks in (2) can be somewhat "linearized:" For every parameter $\boldsymbol{\Theta} \in \mathcal{M}_1$, there is a parameter

$$\bar{\Theta} \in \overline{M} := \{\bar{\Theta} = (\bar{\Theta}^{l-1}, \dots, \bar{\Theta}^0) \ : \ \bar{\Theta}^j \in R^{p^{j+1} \times p^j}, \max_{j \in \{0,\dots,l-1\}} |||\bar{\Theta}^j|||_1 \leq 1\}$$

such that for every $\boldsymbol{x} \in \mathbb{R}^d$

$$\begin{aligned} &g_{\boldsymbol{\Theta}}[\boldsymbol{x}] = \Theta^l \overline{g}_{\overline{\boldsymbol{\Theta}}}[\boldsymbol{x}] \\ &\quad \text{with} \quad \overline{g}_{\overline{\boldsymbol{\Theta}}}[\boldsymbol{x}] := f^l\left[\overline{\Theta}^{l-1} \cdots f^1[\overline{\Theta}^0 \boldsymbol{x}]\right] \in \mathbb{R}^{p^l} . \end{aligned} \tag{7}$$

This additional notation allows us to disentangle the outermost layer (which is regularized directly) from the other layers (which are regularized indirectly). More generally speaking, the additional notation makes a connection to linear regression, where the above holds trivially with $\overline{g}_{\overline{\boldsymbol{\Theta}}}[\boldsymbol{x}] = \boldsymbol{x}$.

We also define

$$\bar{M}_{2,1} := \{\overline{\Theta} = (\overline{\Theta}^{l-1}, \dots, \overline{\Theta}^0) \ : \ \overline{\Theta}^j \in R^{p^{j+1} \times p^j}, \quad \max_{j \in \{0,\dots,l-1\}} |||\overline{\Theta}^j|||_{2,1} \leq 1\}$$

accordingly.

In high-dimensional linear regression, the quantity central to prediction guarantees is the *effective noise* (Lederer and Vogt 2020). The effective noise is in our notation (with $l = 0$ and $m = 1$ to describe linear regression) $2\| \sum_{i=1}^n u_i \boldsymbol{x}_i \|_\infty$. The above linearization allows us to generalize the effective noise to our general deep learning framework:

$$\begin{aligned} r^*_{\text{con}} &:= 2 \sup_{\overline{\Psi} \in \bar{M}_1} |||\sum_{i=1}^n u_i (\overline{g}_{\overline{\Psi}}[x_i])^\top|||_\infty \\ r^*_{\text{node}} &:= 2\sqrt{m} \sup_{\overline{\Psi} \in \bar{M}_{2,1}} |||\sum_{i=1}^n u_i (\overline{g}_{\overline{\Psi}}[x_i])^\top|||_\infty , \end{aligned} \tag{8}$$

where $|||A|||_\infty := \max_{(i,j) \in \{1,\dots,m\} \times \{1,\dots,p^l\}} |A_{ij}|$ for $A \in \mathbb{R}^{m \times p^l}$. The effective noises, as we will see below, are the optimal tuning parameters in our theories; at the same time, the effective noises depend on the noise random variables $\boldsymbol{u}_1, \dots, \boldsymbol{u}_n$, which are unknown in practice. Accordingly, we call the quantities r^*_{con} and r^*_{node} the *oracle tuning parameters*.

We take a moment to compare the effective noises in (8) to Rademacher complexities (Koltchinskii 2001; Koltchinskii and Panchenko 2002). Rademacher complexities are the basis of a line of other statistical theories for deep learning (Bartlett and Mendelson 2002; Golowich et al. 2018; Lederer 2020a; Neyshabur et al. 2015). In our framework, the Rademacher complexities in the case $m = 1$ are (Lederer 2020a, Definition 1)

$$E_{\boldsymbol{x}_1,\dots,\boldsymbol{x}_n,k_1,\dots,k_n}\left[\sup_{\boldsymbol{\Theta}\in\mathcal{M}_1}\left|\frac{1}{n}\sum_{i=1}^{n}k_i\boldsymbol{g}_{\boldsymbol{\Theta}}[\boldsymbol{x}_i]\right|\right]$$

$$\text{and}\quad E_{\boldsymbol{x}_1,\dots,\boldsymbol{x}_n,k_1,\dots,k_n}\left[\sup_{\boldsymbol{\Theta}\in\mathcal{M}_{2,1}}\left|\frac{1}{n}\sum_{i=1}^{n}k_i\boldsymbol{g}_{\boldsymbol{\Theta}}[\boldsymbol{x}_i]\right|\right]$$

for i.i.d. Rademacher random variables $k_1, \dots, k_n$. The effective noises might look like (rescaled) empirical versions of these quantities at first sight, but this is not the case. Two immediate differences are that (8) apply to general m and circumvent the outermost layers of the networks. But more importantly, Rademacher complexities involve external i.i.d. Rademacher random variables that are not connected with the statistical model at hand, while the effective noises involve the noise variables, which are completely specified by the model and, therefore, can have any distribution (see our sub-Gaussian example further below). Hence, there are no general techniques to relate Rademacher complexities and effective noises.

Not only are the two concepts distinct, but also they are used in very different ways. For example, existing theories use Rademacher complexities to measure the size of the function class at hand, while we use effective noises to measure the maximal impact of the stochastic noise on the estimators. (Our proofs also require a measure of the size of the function class, but this measure is entropy—cf. Lemma 1.) In general, our proof techniques are very different from those in the context of Rademacher complexities.

We can now state a general prediction guarantee.

Theorem 1 (General Prediction Guarantees) *If* $r_{\text{con}} \geq r^*_{\text{con}}$, *it holds that*

$$\text{err}[\widehat{\boldsymbol{\Theta}}_{\text{con}}] \leq \inf_{\boldsymbol{\Theta}\in\mathcal{M}_1}\left\{\text{err}[\boldsymbol{\Theta}] + \frac{2r_{\text{con}}}{n}|||\Theta^l|||_1\right\}.$$

Similarly, if $r_{\text{node}} \geq r^*_{\text{node}}$, *it holds that*

$$\text{err}[\widehat{\boldsymbol{\Theta}}_{\text{node}}] \leq \inf_{\boldsymbol{\Theta}\in\mathcal{M}_{2,1}}\left\{\text{err}[\boldsymbol{\Theta}] + \frac{2r_{\text{node}}}{n}|||\Theta^l|||_{2,1}\right\}.$$

Each bound contains an approximation error $\text{err}[\boldsymbol{\Theta}]$ that captures how well the class of networks can approximate the true data-generating function $\boldsymbol{g}_*$ and a statistical error proportional to r_{con}/n and r_{node}/n, respectively, that captures how well the estimator can select within the class of networks at hand. In other words, Theorem 1 ensures that the estimators (3) and (6) predict—up to the statistical error described by r_{con}/n and r_{node}/n, respectively—as well as the best connection- and node-sparse network. This observation can be illustrated further:

Corollary 1 (Parametric Setting) *If additionally* $\boldsymbol{g}_* = \boldsymbol{g}_{\boldsymbol{\Theta}^*}$ *for a* $\boldsymbol{\Theta}^* \in \mathcal{M}_1$, *it holds that*

$$\mathrm{err}[\widehat{\boldsymbol{\Theta}}_{\mathrm{con}}] \leq \frac{2r_{\mathrm{con}}}{n} |||(\boldsymbol{\Theta}^*)^l|||_1 .$$

If instead $\boldsymbol{g}_* = \boldsymbol{g}_{\boldsymbol{\Theta}^*}$ *for a* $\boldsymbol{\Theta}^* \in \mathcal{M}_{2,1}$, *it holds that*

$$\mathrm{err}[\widehat{\boldsymbol{\Theta}}_{\mathrm{node}}] \leq \frac{2r_{\mathrm{node}}}{n} |||(\boldsymbol{\Theta}^*)^l|||_{2,1} .$$

Hence, if the underlying data-generating function is a sparse network itself, the prediction errors of the estimators are essentially bounded by the statistical errors r_{con}/n and r_{node}/n. In high-dimensional statistics, bounds similar to those in Theorem 1 and Corollary 1 are called *oracle inequalities* (Lederer et al. 2019; Lederer 2022).

The above-stated results also identify the oracle tuning parameters r^*_{con} and r^*_{node} as optimal tuning parameters: they give the best prediction guarantees in Theorem 1. But since the oracle tuning parameters are unknown in practice, the guarantees implicitly presume a calibration scheme that satisfies $r_{\mathrm{con}} \approx r^*_{\mathrm{con}}$ in practice. A natural candidate is cross-validation, but there are no guarantees that cross-validation provides such tuning parameters. This is a limitation that our theories share with all other theories in the field.

Rather than dealing with the practical calibration of the tuning parameters, we exemplify the oracle tuning parameters in a specific setting. This analysis will illustrate the rates of convergences that we can expect from Theorem 1, and it will allow us to compare our theories with other theories in the literature. Assume that the activation functions satisfy $\boldsymbol{f}^j[\boldsymbol{0}_{p^j}] = \boldsymbol{0}_{p^j}$ and are 1-Lipschitz continuous with respect to the Euclidean norms on the functions' input and output spaces $\mathbb{R}^{p^j}$. A popular example is ReLU activation, but the conditions are met by many other functions as well. Also, assume that the noise vectors $\boldsymbol{u}_1, \ldots, \boldsymbol{u}_n$ are independent and centered and have uniformly sub-Gaussian entries (van de Geer 2000, Display (8.2) on Page 126). Keep the input vectors fixed and capture their normalizations by

$$\bar{v}_\infty := \sqrt{\frac{1}{n}\sum_{i=1}^n \|\boldsymbol{x}_i\|_\infty^2} \quad \text{and} \quad \bar{v}_2 := \sqrt{\frac{1}{n}\sum_{i=1}^n \|\boldsymbol{x}_i\|_2^2} .$$

Then, we obtain the following bounds for the effective noises.

Proposition 2 (Sub-Gaussian Noise) *There is a constant* $c \in (0, \infty)$ *that depends only on the sub-Gaussian parameters of the noise such that*

$$P\left\{ r^*_{\mathrm{con}} \leq c\bar{v}_\infty \sqrt{nl\big(\log[2mn\bar{p}]\big)^3} \right\} \geq 1 - \frac{1}{n}$$

and

$$P\left\{ r^*_{\mathrm{node}} \leq c\bar{v}_2 \sqrt{mnlp\big(\log[2mn\bar{p}]\big)^3} \right\} \geq 1 - \frac{1}{n} .$$

Broadly speaking, this result combined with Theorem 1 illustrates that accurate prediction with connection- and node-sparse estimators is possible even when using very wide and deep networks. Let us analyze the factors one by one and compare them to the factors in the bounds of Taheri et al. (2021) and Neyshabur et al. (2015), which are the two most related papers. The connection-sparse case compares to the results in Taheri et al. (2021), and it compares to the results in Neyshabur et al. (2015) when setting the parameters in that paper to $p = q = 1$ (which gives a setting that is slightly more restrictive than ours) or $p = 1; q = \infty$ (which gives a setting that is slightly less restrictive than ours), and it compares to (Golowich et al. 2018, Theorem 2). The node-sparse case compares to Neyshabur et al. (2015) with $p = 2; q = \infty$ (which gives a setting that is more restrictive than ours, though). Our setup is also more general than the one in Neyshabur et al. (2015) in the sense that it allows for activation other than ReLU.

The dependence on n is, as usual, $1/\sqrt{n}$ up to logarithmic factors.

In the connection-sparse case, our bounds involve $\bar{v}_\infty = \sqrt{\sum_{i=1}^n \|\boldsymbol{x}_i\|_\infty^2 / n}$ rather than the factor $v_\infty := \max_{i\in\{1,\dots,n\}} \|\boldsymbol{x}_i\|_\infty$ of Golowich et al. (2018) and Neyshabur et al. (2015) or the factor $\bar{v}_2 = \sqrt{\sum_{i=1}^n \|\boldsymbol{x}_i\|_2^2 / n}$ of Taheri et al. (2021). In principle, the improvements of $\bar{v}_\infty$ over v_∞ and $\bar{v}_2$ can be up to a factor $\sqrt{n}$ and up to a factor $\sqrt{d}$, respectively; in practice, the improvements depend on the specifics on the data. For example, on the training data of `MNIST` (LeCun et al. 1998) and `Fashion-MNIST` (Xiao et al. 2017) ($\sqrt{n} \approx 250; \sqrt{d} = 28$ in both data sets), it holds that $\bar{v}_\infty \approx v_\infty \approx \bar{v}_2/9$ and $\bar{v}_\infty \approx v_\infty \approx \bar{v}_2/12$, respectively. In the node-sparse case, our bounds involve $\bar{v}_2$, which is again somewhat smaller than the factor $v_2 := \max_{i\in\{1,\dots,n\}} \|\boldsymbol{x}_i\|_2$ in Neyshabur et al. (2015).

The main difference between the bounds for the connection-sparse and node-sparse estimators is their dependencies on the networks' maximal width p. The bound for the connection-sparse estimator (3) depends on the width p only logarithmically (through $\bar{p}$), while the bound for the node-sparse estimator (6) depends on p sublinearly. The dependence in the connection-sparse case is the same as in Taheri et al. (2021), while Neyshabur et al. (2015) can avoid even that logarithmic dependence (and, therefore, allow for networks with infinite widths). The node-sparse case in Neyshabur et al. (2015) does not involve our linear dependence on the width, but this difference stems from the fact that they use a more restrictive version of the grouping—we take the maximum over each layer, while they take the maximum over each node— and our results can be readily adjusted to their notion of group sparsity. These observations indicate that node sparsity as formulated above is suitable for slim networks ($p \ll n$) but should be strengthened or complemented with other notions of sparsity otherwise. To give a numeric example, the training data in `MNIST` (LeCun et al. 1998) and `Fashion-MNIST` (Xiao et al. 2017) comprise $n = 60\,000$ samples, which means that the width should be considerably smaller than 60 000 when using node sparsity alone. (Note that the input layer does not take part in p, which means that d could be larger.)

For unconstraint estimation, one can expect a linear dependence of the error on the total number of parameters (Anthony and Bartlett 1999). Our bounds for the

sparse estimators, in contrast, only have a $\log[\overline{p}]$ dependence on the total number of parameters. This difference illustrates the virtue of regularization in general, and the virtue of sparsity in particular.

Both of our bounds have a mild $\sqrt{l}$ dependence on the depth. These dependencies align with the results in (Golowich et al. 2018, Theorem 2) but considerably improve on the exponentially increasing dependencies on the depth in Neyshabur et al. (2015) and, therefore, are particularly suited to describe deep network architectures. Replacing the conditions $\max_j |||\Theta^j|||_1 \leq 1$ and $\max_j |||\Theta^j|||_{2,1} \leq 1$ in the definitions of the connection-sparse and node-sparse estimators by the stricter conditions $\sum_j |||\Theta^j|||_1 \leq 1$ and $\sum_j |||\Theta^j|||_{2,1} \leq 1$, respectively (cf. Taheri et al. (2021) and our discussion in Section 2), the dependence on the depth can be improved further from $\sqrt{l}$ to $(2/l)^l\sqrt{l}$ (this only requires a simple adjustment of the last display in the proof of Proposition 4), which is exponentially decreasing in the depth.

Our connection-sparse bounds have a mild $\log[m]$ dependence on the number of output nodes; the node-sparse bound involve an additional factor $\sqrt{m}$. The case of multiple outputs has not been considered in statistical prediction bounds before.

Proposition 2 also highlights another advantage of our regularization approach over theories such as Golowich et al. (2018) and Neyshabur et al. (2015) that apply to constraint estimators. The theories for constraint estimators require bounding the sparsity levels directly, but in practice, suitable values for these bounds are rarely known. In our framework, in contrast, the sparsity is controlled via tuning parameters indirectly, and Proposition 2—although not providing a complete practical calibration scheme—gives insights into how these tuning parameters should scale with n, d, l, and so forth.

We also note that the bounds in Theorem 1 can be generalized readily to every estimator of the form

$$\widehat{\boldsymbol{\Theta}}_{\text{gen}} \in \arg\min_{\boldsymbol{\Theta}\in\mathcal{M}_{\text{gen}}} \left\{ \sum_{i=1}^{n} \|\boldsymbol{y}_i - \boldsymbol{g}_{\boldsymbol{\Theta}}[\boldsymbol{x}_i]\|_2^2 + r_{\text{gen}} |||\Theta^l||| \right\},$$

where $r_{\text{gen}} \in [0, \infty)$ is a tuning parameter, $\mathcal{M}_{\text{gen}}$ any nonempty subset of $\mathcal{M}$, and $|||\cdot|||$ any norm. The bound for such an estimator is then

$$\text{err}[\widehat{\boldsymbol{\Theta}}_{\text{gen}}] \leq \inf_{\boldsymbol{\Theta}\in\mathcal{M}_{\text{gen}}} \left\{ \text{err}[\boldsymbol{\Theta}] + \frac{2r_{\text{gen}}}{n} |||\Theta^l||| \right\}$$

for $r_{\text{gen}} \geq r^*_{\text{gen}}$, where r^*_{gen} is as r^*_{con} but based on the dual norm of $|||\cdot|||$ instead of the dual norm of $|||\cdot|||_1$. For example, one could impose connection sparsity on some layers and node sparsity on others, or one could impose different regularizations altogether. We omit the details to avoid digression.

The above oracle inequalities bound the prediction error, a standard measure of accuracy in statistics. Broadly speaking, this measure captures "how well the estimator describes the data-generating process." So our comparison with Neyshabur et al. (2015) and Golowich et al. (2018) might seem questionable, because they instead bound the generalization error, a measure that is more common in

machine learning and captures "how well the estimator describes new samples." But we can derive such bounds as well. For simplicity, we consider a parametric setting and sub-Gaussian noise again. We then find the following bounds:

Proposition 3 (Generalization Guarantees) *Assume that the inputs* $\boldsymbol{x}, \boldsymbol{x}_1, \dots, \boldsymbol{x}_n$ *are i.i.d. random vectors, that the noise vectors* $\boldsymbol{u}_1, \dots, \boldsymbol{u}_n$ *are independent and centered and have uniformly sub-Gaussian entries, and that* $r_{\text{con}} = r^*_{\text{con}}, r_{\text{node}} = r^*_{\text{node}} \to 0$ *as* $n \to \infty$. *Consider an arbitrary positive constant* $b \in (0, \infty)$. *If* $\boldsymbol{g}_* = \boldsymbol{g}_{\boldsymbol{\Theta}^*}$ *for a* $\boldsymbol{\Theta}^* \in \mathcal{M}_1$, *it holds with probability at least* $1 - 1/n$ *that*

$$\text{risk}[\widehat{\boldsymbol{\Theta}}_{\text{con}}] \leq (1+b)\text{risk}[\boldsymbol{\Theta}^*] + c\bar{v}_\infty \sqrt{\frac{l\big(\log[2mn\bar{p}]\big)^3}{n}} \,|||(\boldsymbol{\Theta}^*)^l|||_1$$

for a constant $c \in (0, \infty)$ *that depends only on b and the sub-Gaussian parameters of the noise. Similarly, if* $\boldsymbol{g}_* = \boldsymbol{g}_{\boldsymbol{\Theta}^*}$ *for a* $\boldsymbol{\Theta}^* \in \mathcal{M}_{2,1}$, *it holds with probability at least* $1 - 1/n$ *that*

$$\text{risk}[\widehat{\boldsymbol{\Theta}}_{\text{con}}] \leq (1+b)\text{risk}[\boldsymbol{\Theta}^*] + c\bar{v}_2 \sqrt{\frac{ml\underline{p}\big(\log[2mn\bar{p}]\big)^3}{n}} \,|||(\boldsymbol{\Theta}^*)^l|||_{2,1}$$

for a constant $c \in (0, \infty)$ *that depends only on b and the sub-Gaussian parameters of the noise.*

Hence, the generalization errors are bounded by the same terms as the prediction errors.

4 Outlook: Initialization

Our theoretical results also suggest further research on a practical problem in deep learning: weight initialization (Glorot and Bengio 2010; He et al. 2015; Mishkin and Matas 2015). To highlight the connection between our work and weight initialization, we consider once more our guarantees' dependence on the depth l. Proposition 3, for example, comprises a sublinear dependence through the factor $\sqrt{l}$ and a logarithmic dependence through the total number of parameters $\bar{p}$ inside the logarithm—we have discussed these dependencies in detail. But there is another potential source of dependence on l: the factor $|||(\boldsymbol{\Theta}^*)^l|||_1$. Naively thinking, one could suspect that this factor scales exponentially in l: the argument would be that the weight matrices of each of the $l-1$ inner layers needs to be rescaled to fit into $\mathcal{M}_1$ or $\mathcal{M}_{2,1}$, which means that the weight matrix of the outer layer needs to be rescaled by a product of these $l-1$ factors.

The argument is intuitive, but it is wrong: the problem with it is that the optimal weight matrices $(\boldsymbol{\Theta}^*)^l$ change with the depth of the network, while the

data-generating process remains unaffected by what function we use to approximate it. In other words, we cannot expect a simple relationship between $(\boldsymbol{\Theta}^*)^l$ and $(\boldsymbol{\Theta}^*)^{l-1}$, but we can expect the overall "scales" of the corresponding networks to be similar, that is, $|||(\boldsymbol{\Theta}^*)^l|||_1 \approx |||(\boldsymbol{\Theta}^*)^{l-1}|||_1$. Hence, we can assume that the factor $|||(\boldsymbol{\Theta}^*)^l|||_1$ in our bounds to be approximately independent of l.

One can also argue that the recent results on approximation properties of sparse neural networks, such as Beknazaryan (2021); Schmidt-Hieber (2020), suggest that sparse networks with parameters in $\mathcal{M}_1$ or $\mathcal{M}_{2,1}$ and fixed $|||(\boldsymbol{\Theta}^*)^l|||_1$ or $|||(\boldsymbol{\Theta}^*)^l|||_{2,1}$ norms, respectively, can indeed approximate large classes of functions.

In any case, we can draw two conclusions: First, our bounds indeed depend on the network depth as advertised. Second, our results hint at the fact that initialization schemes should take network depths into account, and it might be favorable to use sparse initialization schemes rather than distributing weights "uniformly" across the entire network. More generally, we conclude that the connection between sparse networks and weight initializations might be an interesting topic for further research.

5 Discussion

We have developed guarantees for sparse deep learning both in terms of the prediction error (Theorems 1 and Corollary 1 together with Proposition 2), a standard measure of accuracy in statistics, and in terms of the generalization error (Proposition 3), a standard measure of accuracy in machine learning. These results extend and complement existing guarantees in the literature—see Table 1 below.

Even though many deep learning applications fall into the framework of classification, we have focussed on regression with least-squares loss. The reason is that the regression setting is much more challenging: since the loss is unbounded, many of the techniques regularly used in classification (like McDiarmid's inequality (McDiarmid 1989, Lemma (3.3))) are not applicable. In this sense, our derivations are more general, and we expect that our approach will provide very similar classifications bounds in the future as well (see Appendix 1 for possible extensions more generally).

Evidence for the benefits of deep networks has been established in practice (LeCun et al. 2015; Schmidhuber 2015), approximation theory (Liang and Srikant 2016; Telgarsky 2016; Yarotsky 2017), and statistics (Taheri et al. 2021; Kohler et al. 2019). Since our guarantees scale at most sublinearly in the number of layers (or even improve with increasing depth—see our comment on Page 5), our paper complements these lines of research and shows that sparsity-inducing regularization is an effective approach to coping with the complexity of deep and very deep networks.

While previous theories mostly considered connection sparsity (small number of active connections between nodes), we also include node sparsity (small number of active nodes). Moreover, as discussed on Page 5, Theorem 1 can be readily extended to any norm-based regularization. Hence, it is straightforward to adjust our results to granularities between connection and node sparsity—cf. Mao et al. (2017). On the

Table 1 Presence (✓) or absence ( ) of certain features in previous statistical theories for sparse deep learning

Approach	Mathematical techniques	References
FatShat	Fat-shattering dimension	Bartlett (1998)
RadCon	Rademacher complexity	Bartlett and Mendelson (2002); Golowich et al. (2018) Neyshabur et al. (2015); Lederer (2020a)
RadNode	Rademacher complexity	Bartlett and Mendelson (2002); Neyshabur et al. (2015) Lederer (2020a)
NonPar	Nonparametric statistics	Schmidt-Hieber (2020)
HighDim	High-dimensional statistics & concentration inequalities	Taheri et al. (2021)

Approach	Feasible	ℓ_2-loss	Prediction error	Regularized	(sub)linear in l	Node sparsity	Multiple outputs
FatShat	✓						
RadCon	✓					✓	
RadNode	✓			✓			
NonPar		✓			✓		
HighDim	✓	✓	✓	✓	✓		

We extend the HighDim results to node sparsity and multiple outputs. Moreover, we improve the dependence of the HighDim bounds on the data, and we avoid their auxiliary parameters

other hand, our techniques do not seem appropriate for "hard-coded" types of sparsity, such as 2:4 ("two-to-four") sparsity (Mishra et al. 2021).

Connection sparsity limits the number of nonzero entries in each parameter matrix, while node sparsity only limits the total number of nonzero rows. Hence, the number of columns in a parameter matrix, that is, the width of the preceding layer, is regularized only in the case of connection sparsity. Our theoretical results reflect this insight in that the bounds for the connection- and node-sparse estimators depend on the networks' width logarithmically and sublinearly, respectively. Practically speaking, our results indicate that connection sparsity is suitable to handle wide networks, but node sparsity is suitable for wide networks only when complemented by connection sparsity or other strategies.

The mild logarithmic dependence of our connection-sparse bounds on the number of output nodes illustrates that networks with many outputs can be learned in practice. Our prediction theory is the first one to consider multiple output nodes; a classification theory with a logarithmic dependence on the output nodes has been established very recently in Ledent et al. (2019).

The mathematical underpinnings of our theory are very different from those of most other papers in theoretical deep learning. The proof of the main theorem shares similarities with proofs in high-dimensional statistics, such as the concept of the effective noise (Lederer 2022). The treatments of the relevant empirical-processes use metric entropy, chaining, and Lipschitz properties of neural networks. These concepts and tools are not standard in deep learning and, therefore, might be of more general interest (see again Appendix 1 for further ideas).

Our theory has three limitations: First, the bounds apply only to global optima of the optimization landscapes rather than local optima or other points in which certain algorithms might be trapped. However, there is evidence that global optimization can be feasible at least in wide and deep networks (Lederer 2020b). Second, the theory does not entail a practical scheme for the calibration of the tuning parameters. However, the inclusion of regularization (rather than constraints) is already a step forward, because it reveals how the tuning parameters should scale with the problem dimensions (see our Proposition 2). Third, the network architecture is limited to fully connected feedforward layers, which excludes some aspects of modern pipelines (such as convolutions, dropout, and so forth). In any case, all three limitations are open problems in the literature; in particular, the mentioned limitations are shared by most theories on the topic.

We can summarize what this paper contributes—and what it does not—as follows: From a practical perspective, it is well established that sparsity can benefit deep learning, and there are several methods to generate sparsity in practice. Thus, this paper does not provide new practical insights or methods. Instead, our paper (i) backs up these practical observations with statistical theories that are more general and closer to practice than previous theories, and it (ii) establishes refined concepts and techniques for the statistical analysis of deep learning more generally.

Appendix

The Appendix consists of two auxiliary results and the proofs of Theorem 1 and Propositions 1 and 2. Our approach combines techniques from high-dimensional statistics and empirical-process theory that are very different from the techniques used in most other approaches in the literature.

A Lipschitz property

In this section, we prove a Lipschitz property that we use in the proof of Proposition 2.

Proposition 4 (Lipschitz Property) *In the framework of Sections 2 and 3, it holds for all* $\overline{\boldsymbol{\Theta}}, \overline{\boldsymbol{\Gamma}} \in \overline{\mathcal{M}}_1$ *that*

$$\left\| \overline{g}_{\overline{\boldsymbol{\Theta}}}[\boldsymbol{x}] - \overline{g}_{\overline{\boldsymbol{\Gamma}}}[\boldsymbol{x}] \right\|_\infty \leq \sqrt{l} \|\boldsymbol{x}\|_\infty |||\overline{\boldsymbol{\Theta}} - \overline{\boldsymbol{\Gamma}}|||_{\mathrm{F}}$$

and for all $\overline{\boldsymbol{\Theta}}, \overline{\boldsymbol{\Gamma}} \in \overline{\mathcal{M}}_{2,1}$ *that*

$$\left\| \overline{g}_{\overline{\boldsymbol{\Theta}}}[\boldsymbol{x}] - \overline{g}_{\overline{\boldsymbol{\Gamma}}}[\boldsymbol{x}] \right\|_2 \leq \sqrt{l} \|\boldsymbol{x}\|_2 |||\overline{\boldsymbol{\Theta}} - \overline{\boldsymbol{\Gamma}}|||_{\mathrm{F}} .$$

The Frobenius norm is defined as

$$|||\overline{\boldsymbol{\Theta}}|||_{\mathrm{F}} := \sqrt{\sum_{j=0}^{l-1} |||\overline{\Theta}^j|||_{\mathrm{F}}^2} := \sqrt{\sum_{j=0}^{l-1} \sum_{i=1}^{p^{j+1}} \sum_{k=1}^{p^j} |(\overline{\Theta}^j)_{ik}|^2} \qquad \text{for } \overline{\boldsymbol{\Theta}} \in \overline{\mathcal{M}}_{2,1} = \overline{\mathcal{M}}_1 \cup \overline{\mathcal{M}}_{2,1} .$$

Proposition 4 generalizes (Taheri et al. 2021, Proposition 2) to vector-valued network outputs and to node sparsity, and it replaces their $\|\boldsymbol{x}\|_2$ with the smaller $\|\boldsymbol{x}\|_\infty$ in the connection-sparse case.

Proof of Proposition 4 This proof generalizes and sharpens the proof of Taheri et al. (2021), and it simplifies some arguments of that proof. We define the "inner subnetworks" of a network $\overline{g}_{\overline{\boldsymbol{\Theta}}}$ with $\overline{\boldsymbol{\Theta}} \in \overline{\mathcal{M}}_{2,1}$ as the vector-valued functions

$$\begin{aligned} S_0 \overline{g}_{\overline{\boldsymbol{\Theta}}} \; : \; \mathbb{R}^d &\to \mathbb{R}^{p^1} \\ \boldsymbol{x} &\mapsto S_0 \overline{g}_{\overline{\boldsymbol{\Theta}}}[\boldsymbol{x}] := \overline{\Theta}^0 \boldsymbol{x} \end{aligned}$$

and

$$\begin{aligned} S_j \overline{g}_{\overline{\boldsymbol{\Theta}}} \; : \; \mathbb{R}^d &\to \mathbb{R}^{p^{j+1}} \\ \boldsymbol{x} &\mapsto S_j \overline{g}_{\overline{\boldsymbol{\Theta}}}[\boldsymbol{x}] := \overline{\Theta}^j f^j \left[\cdots f^1 [\overline{\Theta}^0 \boldsymbol{x}] \right] \end{aligned}$$

for $j \in \{1, \dots, l-1\}$. Similarly, we define the "outer subnetworks" of $\overline{g}_{\overline{\Theta}}$ as the real-valued functions

$$\begin{aligned} S^j\overline{g}_{\overline{\Theta}} \;:\; \mathbb{R}^{p^j} &\to \mathbb{R}^{p^l} \\ z &\mapsto S^j\overline{g}_{\overline{\Theta}}[z] := \boldsymbol{f}^l\big[\overline{\Theta}^{l-1} \cdots \boldsymbol{f}^j[z]\big] \end{aligned}$$

for $j \in \{1, \dots, l-1\}$ and

$$\begin{aligned} S^l\overline{g}_{\overline{\Theta}} \;:\; \mathbb{R}^{p^l} &\to \mathbb{R}^{p^l} \\ z &\mapsto S^l\overline{g}_{\overline{\Theta}}[z] := \boldsymbol{f}^l[z]\,. \end{aligned}$$

The initial network can be split into an inner and an outer network along every layer $j \in \{1, \dots, l\}$:

$$\overline{g}_{\overline{\Theta}}[\boldsymbol{x}] = S^j\overline{g}_{\overline{\Theta}}\big[S_{j-1}\overline{g}_{\overline{\Theta}}[\boldsymbol{x}]\big] \qquad \text{for } \boldsymbol{x} \in \mathbb{R}^d\,.$$

We call this our *splitting argument*.

To exploit the splitting argument, we derive a contraction result for the inner subnetworks and a Lipschitz result for the outer subnetworks. We denote the ℓ_2-operator norm of a matrix A, that is, the largest singular value of A, by $|||A|||_{\mathrm{op}}$. Using then the assumptions that the activation functions are 1-Lipschitz and $\boldsymbol{f}^j[\mathbf{0}_{p^j}] = \mathbf{0}_{p^j}$, we get for every $\overline{\boldsymbol{\Theta}} = (\overline{\Theta}^{l-1}, \dots, \overline{\Theta}^{0}) \in \overline{\mathcal{M}}_{2,1}$ and $\boldsymbol{x} \in \mathbb{R}^d$ that

$$\begin{aligned} \big\|S_{j-2}\overline{g}_{\overline{\Theta}}[\boldsymbol{x}]\big\|_2 &= \big\|\overline{\Theta}^{j-2}\boldsymbol{f}^{j-2}\big[S_{j-3}\overline{g}_{\overline{\Theta}}[\boldsymbol{x}]\big]\big\|_2 \\ &\le |||\overline{\Theta}^{j-2}|||_{\mathrm{op}}\big\|\boldsymbol{f}^{j-2}\big[S_{j-3}\overline{g}_{\overline{\Theta}}[\boldsymbol{x}]\big]\big\|_2 \\ &\le |||\overline{\Theta}^{j-2}|||_{\mathrm{op}}\big\|S_{j-3}\overline{g}_{\overline{\Theta}}[\boldsymbol{x}]\big\|_2 \\ &\le \cdots \\ &\le \left(\prod_{k=1}^{j-2} |||\overline{\Theta}^{k}|||_{\mathrm{op}}\right)\|\overline{\Theta}^{0}\boldsymbol{x}\|_2 \\ &\le \left(\prod_{k=0}^{j-2} |||\overline{\Theta}^{k}|||_{\mathrm{op}}\right)\|\boldsymbol{x}\|_2 \end{aligned}$$

for all $j \in \{2, \dots, l\}$. Now, since $|||\overline{\Theta}^{k}|||_{\mathrm{op}} \le |||\overline{\Theta}^{k}|||_{\mathrm{F}} \le |||\overline{\Theta}^{k}|||_{2,1}$ and $\overline{\boldsymbol{\Theta}} \in \overline{\mathcal{M}}_{2,1}$, we can deduce from the display that

$$\big\|S_{j-2}\overline{g}_{\overline{\Theta}}[\boldsymbol{x}]\big\|_2 \le \left(\prod_{k=0}^{j-2} |||\overline{\Theta}^{k}|||_{2,1}\right)\|\boldsymbol{x}\|_2\,.$$

This inequality is our *contraction property*.

By similar arguments, we get for every $z_1, z_2 \in \mathbb{R}^{p^j}$ that

$$
\begin{aligned}
&\left\|S^j\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Theta}}}[z_1] - S^j\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Theta}}}[z_2]\right\|_2 \\
&= \left\|\boldsymbol{f}^l\left[\overline{\Theta}^{l-1}\cdots\boldsymbol{f}^j[z_1]\right] - \boldsymbol{f}^l\left[\overline{\Theta}^{l-1}\cdots\boldsymbol{f}^j[z_2]\right]\right\|_2 \\
&\le \left\|\overline{\Theta}^{l-1}\left[\boldsymbol{f}^{l-1}\cdots\boldsymbol{f}^j[z_1]\right] - \overline{\Theta}^{l-1}\left[\boldsymbol{f}^{l-1}\cdots\boldsymbol{f}^j[z_2]\right]\right\|_2 \\
&\le \left|\!\left|\!\left|\overline{\Theta}^{l-1}\right|\!\right|\!\right|_{\mathrm{op}}\left\|\boldsymbol{f}^{l-1}\left[\cdots\boldsymbol{f}^j[z_1]\right] - \boldsymbol{f}^{l-1}\left[\cdots\boldsymbol{f}^j[z_2]\right]\right\|_2 \\
&\le \cdots \\
&\le \left(\prod_{k=j}^{l-1}\left|\!\left|\!\left|\overline{\Theta}^{k}\right|\!\right|\!\right|_{\mathrm{op}}\right)\|z_1 - z_2\|_2
\end{aligned}
$$

for $j \in \{1, \dots, l\}$, where $\prod_{k=l}^{l-1}\left|\!\left|\!\left|\overline{\Theta}^{k}\right|\!\right|\!\right|_{\mathrm{op}} := 1$. Hence, similarly as above,

$$
\left\|S^j\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Theta}}}[z_1] - S^j\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Theta}}}[z_2]\right\|_2 \le \left(\prod_{k=j}^{l-1}\left|\!\left|\!\left|\overline{\Theta}^{k}\right|\!\right|\!\right|_{2,1}\right)\|z_1 - z_2\|_2 .
$$

This inequality is our *Lipschitz property*.

We now use the contraction and Lipschitz properties of the subnetworks to derive a Lipschitz result for the entire network. We consider two networks $\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Theta}}}$ and $\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Gamma}}}$ with parameters $\overline{\boldsymbol{\Theta}} = (\overline{\Theta}^{l-1}, \dots, \overline{\Theta}^{0}) \in \overline{\mathcal{M}}_{2,1}$ and $\overline{\boldsymbol{\Gamma}} = (\overline{\Gamma}^{l-1}, \dots, \overline{\Gamma}^{0}) \in \overline{\mathcal{M}}_{2,1}$, respectively. Our above-derived splitting argument applied with $j = 1$ and $j = l$, respectively, yields

$$
\left\|\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Theta}}}[\boldsymbol{x}] - \overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Gamma}}}[\boldsymbol{x}]\right\|_2 = \left\|S^1\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Theta}}}\left[S_0\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Theta}}}[\boldsymbol{x}]\right] - S^l\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Gamma}}}\left[S_{l-1}\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Gamma}}}[\boldsymbol{x}]\right]\right\|_2 .
$$

Elementary algebra and the fact that $S^{j-1}\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Theta}}}[S_{j-2}\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Gamma}}}[\boldsymbol{x}]] = S^j\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Theta}}}[\overline{\Theta}^{j-1}\boldsymbol{f}^{j-1}[S_{j-2}\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Gamma}}}[\boldsymbol{x}]]$ for $j \in \{2, \dots, l\}$ then allow us to derive

$$
\begin{aligned}
&\left\|\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Theta}}}[\boldsymbol{x}]-\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Gamma}}}[\boldsymbol{x}]\right\|_2 \\
&=\left\|S^1\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Theta}}}\left[S_0\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Theta}}}[\boldsymbol{x}]\right]-\sum_{j=1}^{l}\left(S^j\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Theta}}}\left[S_{j-1}\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Gamma}}}[\boldsymbol{x}]\right]-S^j\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Theta}}}\left[S_{j-1}\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Gamma}}}[\boldsymbol{x}]\right]\right)-S^l\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Gamma}}}\left[S_{l-1}\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Gamma}}}[\boldsymbol{x}]\right]\right\|_2 \\
&=\Bigg\|S^1\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Theta}}}\left[S_0\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Theta}}}[\boldsymbol{x}]\right]-S^1\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Theta}}}\left[S_0\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Gamma}}}[\boldsymbol{x}]\right] \\
&\quad-\sum_{j=2}^{l}\left(S^j\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Theta}}}\left[S_{j-1}\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Gamma}}}[\boldsymbol{x}]\right]-S^{j-1}\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Theta}}}\left[S_{j-2}\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Gamma}}}[\boldsymbol{x}]\right]\right) \\
&\quad+S^l\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Theta}}}\left[S_{l-1}\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Gamma}}}[\boldsymbol{x}]\right]-S^l\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Gamma}}}\left[S_{l-1}\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Gamma}}}[\boldsymbol{x}]\right]\Bigg\|_2 \\
&=\Bigg\|S^1\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Theta}}}\left[S_0\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Theta}}}[\boldsymbol{x}]\right]-S^1\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Theta}}}\left[S_0\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Gamma}}}[\boldsymbol{x}]\right] \\
&\quad-\sum_{j=2}^{l}\left(S^j\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Theta}}}\left[S_{j-1}\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Gamma}}}[\boldsymbol{x}]\right]-S^j\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Theta}}}\left[\overline{\Theta}^{j-1}\boldsymbol{f}^{j-1}\left[S_{j-2}\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Gamma}}}[\boldsymbol{x}]\right]\right]\right) \\
&\quad+S^l\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Theta}}}\left[S_{l-1}\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Gamma}}}[\boldsymbol{x}]\right]-S^l\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Gamma}}}\left[S_{l-1}\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Gamma}}}[\boldsymbol{x}]\right]\Bigg\|_2 \\
&\leq\left\|S^1\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Theta}}}\left[S_0\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Theta}}}[\boldsymbol{x}]\right]-S^1\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Theta}}}\left[S_0\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Gamma}}}[\boldsymbol{x}]\right]\right\|_2 \\
&\quad+\sum_{j=2}^{l}\left\|S^j\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Theta}}}\left[S_{j-1}\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Gamma}}}[\boldsymbol{x}]\right]-S^j\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Theta}}}\left[\overline{\Theta}^{j-1}\boldsymbol{f}^{j-1}\left[S_{j-2}\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Gamma}}}[\boldsymbol{x}]\right]\right]\right\|_2 \\
&\quad+\left\|S^l\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Theta}}}\left[S_{l-1}\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Gamma}}}[\boldsymbol{x}]\right]-S^l\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Gamma}}}\left[S_{l-1}\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Gamma}}}[\boldsymbol{x}]\right]\right\|_2.
\end{aligned}
$$

We bound this further by using the above-derived Lipschitz property of the outer networks and the observation that $S^l\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Theta}}}[S_{l-1}\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Gamma}}}[\boldsymbol{x}]]=S^l\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Gamma}}}[S_{l-1}\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Gamma}}}[\boldsymbol{x}]]$:

$$
\begin{aligned}
\left\|\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Theta}}}[\boldsymbol{x}]-\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Gamma}}}[\boldsymbol{x}]\right\|_2 &\leq\left(\prod_{k=1}^{l-1}|||\overline{\Theta}^k|||_{2,1}\right)\left\|S_0\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Theta}}}[\boldsymbol{x}]-S_0\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Gamma}}}[\boldsymbol{x}]\right\|_2 \\
&\quad+\sum_{j=2}^{l}\left(\prod_{k=j}^{l-1}|||\overline{\Theta}^k|||_{2,1}\right)\left\|S_{j-1}\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Gamma}}}[\boldsymbol{x}]-\overline{\Theta}^{j-1}\boldsymbol{f}^{j-1}\left[S_{j-2}\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Gamma}}}[\boldsymbol{x}]\right]\right\|_2,
\end{aligned}
$$

which is by the definition of the inner networks equivalent to

$$
\begin{aligned}
\left\|\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Theta}}}[\boldsymbol{x}]-\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Gamma}}}[\boldsymbol{x}]\right\|_2 &\leq\left(\prod_{k=1}^{l-1}|||\overline{\Theta}^k|||_{2,1}\right)\|\overline{\Theta}^0\boldsymbol{x}-\overline{\Gamma}^0\boldsymbol{x}\|_2 \\
&\quad+\sum_{j=2}^{l}\left(\prod_{k=j}^{l-1}|||\overline{\Theta}^k|||_{2,1}\right)\left\|\overline{\Gamma}^{j-1}\boldsymbol{f}^{j-1}\left[S_{j-2}\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Gamma}}}[\boldsymbol{x}]\right]-\overline{\Theta}^{j-1}\boldsymbol{f}^{j-1}\left[S_{j-2}\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Gamma}}}[\boldsymbol{x}]\right]\right\|_2.
\end{aligned}
$$

Using the properties of the operator norm, we can deduce from this inequality that

$$\left\|\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Theta}}}[\boldsymbol{x}] - \overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Gamma}}}[\boldsymbol{x}]\right\|_2 \leq \left(\prod_{k=1}^{l-1} |||\overline{\Theta}^k|||_{2,1}\right) |||\overline{\Theta}^0 - \overline{\Gamma}^0|||_{\mathrm{op}} \|\boldsymbol{x}\|_2$$
$$+ \sum_{j=2}^{l} \left(\prod_{k=j}^{l-1} |||\overline{\Theta}^k|||_{2,1}\right) |||\overline{\Gamma}^{j-1} - \overline{\Theta}^{j-1}|||_{\mathrm{op}} \left\|\boldsymbol{f}^{j-1}\left[S_{j-2}\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Gamma}}}[\boldsymbol{x}]\right]\right\|_2 .$$

Invoking the mentioned conditions on theactivation functions and the contraction property for the inner subnetworks then yields

$$\left\|\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Theta}}}[\boldsymbol{x}] - \overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Gamma}}}[\boldsymbol{x}]\right\|_2$$
$$\leq \left(\max_{v \in \{0,\dots,l-1\}} \prod_{\substack{k \in \{0,\dots,l-1\} \\ k \neq v}} \max\left\{|||\overline{\Theta}^k|||_{2,1}, |||\overline{\Gamma}^k|||_{2,1}\right\}\right)$$
$$\left(\sum_{j=0}^{l-1} |||\overline{\Gamma}^j - \overline{\Theta}^j|||_{\mathrm{op}}\right) \|\boldsymbol{x}\|_2$$
$$\leq \sqrt{l}\|\boldsymbol{x}\|_2 |||\overline{\boldsymbol{\Theta}} - \overline{\boldsymbol{\Gamma}}|||_{\mathrm{F}} .$$

The proof for the connection-sparse case is almost the same. The main difference is that one needs to use the $\|\cdot\|_\infty$- and $|||\cdot|||_1$-norms (rather than the $\|\cdot\|_2$- and $|||\cdot|||_{\mathrm{op}}$-norms) and the inequality $\|A\boldsymbol{b}\|_\infty \leq |||A|||_1 \|\boldsymbol{b}\|_\infty$ (rather than the inequality $\|A\boldsymbol{b}\|_2 \leq |||A|||_{\mathrm{op}} \|\boldsymbol{b}\|_2$) to establish suitable contraction and Lipschitz properties. □

B Entropy bound

In this section, we establish bounds for the entropies of $\overline{\mathcal{M}}_1$ and $\overline{\mathcal{M}}_{2,1}$. The distance between two networks $\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Theta}}}$ and $\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Gamma}}}$ is defined as $\mathrm{dist}[\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Theta}}}, \overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Gamma}}}] := \sqrt{\sum_{i=1}^n \|\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Theta}}}[\boldsymbol{x}_i] - \overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Gamma}}}[\boldsymbol{x}_i]\|_\infty^2 / n}$. Given this distance function and a radius $t \in (0, \infty)$, the metric entropy of a nonempty set $\mathcal{A} \subset \{\overline{\boldsymbol{\Theta}} = (\overline{\Theta}^{l-1}, \dots, \overline{\Theta}^0) : \overline{\Theta}^j \in \mathbb{R}^{p^{j+1} \times p^j}\}$ is denoted by $H[t, \mathcal{A}]$. We then get the following entropy bounds.

Lemma 1 (Entropy Bounds) *In the framework of Sections 2 and 3, it holds for a constant $c_H \in (0, \infty)$ and every $t \in (0, \infty)$ that*

$$H[t, \overline{\mathcal{M}}_1] \leq c_H \left\lceil \frac{(\overline{v}_\infty)^2 l}{t^2} \right\rceil \log\left[\frac{\overline{p} t^2}{(\overline{v}_\infty)^2 l} + 2\right]$$

and

$$H[t, \overline{\mathcal{M}}_{2,1}] \leq c_H \left\lceil \frac{(\overline{v}_\infty)^2 l \underline{p}}{t^2} \right\rceil \log\left[\frac{\overline{p} t^2}{(\overline{v}_\infty)^2 l} + 2\right] .$$

Proof of Lemma 1 The first bound can be derived by combining established deterministic and randomization arguments (Carl 1985);(Lederer 2010, Proof of Theorem 1.1);(Taheri et al. 2021, Proposition 3).

For the second bound, observe that

$$|||\Theta^j|||_1 = \sum_{i=1}^{p^{j+1}} \sum_{k=1}^{p^j} |(\Theta^j)_{ik}| \leq \sqrt{p^{j+1}} \sum_{k=1}^{p^j} \sqrt{\sum_{i=1}^{p^{j+1}} |(\Theta^j)_{ik}|^2} = \sqrt{p^{j+1}} |||\Theta^j|||_{2,1} = \sqrt{\underline{p}} |||\Theta^j|||_{2,1}$$

for all $j \in \{0, \dots, l-1\}$ and $\Theta^j \in \mathbb{R}^{p^{j+1} \times p^j}$. We used in turn 1. the definition of the $|||\cdot|||_1$-norm on Page 2, 2. the linearity and interchangeability of finite sums and the inequality $\|\boldsymbol{a}\|_1 \leq \sqrt{b}\|\boldsymbol{a}\|_2$ for all $\boldsymbol{a} \in \mathbb{R}^b$, 3. the definition of the $|||\cdot|||_{2,1}$-norm on Page ??, and 4. the definition of the width $\underline{p}$ on Page 2. Hence, $\overline{\mathcal{M}_{2,1}} \subset \sqrt{\underline{p}}\overline{\mathcal{M}_1}$. A bound for the entropies of $\overline{\mathcal{M}_{2,1}}$ can, therefore, be derived from the first bound by replacing the radii t on the right-hand side by $t/\sqrt{\underline{p}}$.

C Proof of Theorem 1

In this section, we state a proof for Theorem 1. The proof is inspired by derivations in high-dimensional statistics—see, for example, (Zhuang and Lederer 2018; Lederer 2022) and references therein.

Proof of Theorem 1 The main idea of the proof is to contrast the estimators' objective functions evaluated at their minima with the estimators' objective functions at other points. Our first step is to derive what we call a *basic inequality*. By the definition of the estimator in (6), it holds for every $\boldsymbol{\Theta} \in \mathcal{M}_{2,1}$ that

$$\sum_{i=1}^{n} \|\boldsymbol{y}_i - \boldsymbol{g}_{\widehat{\boldsymbol{\Theta}}}[\boldsymbol{x}_i]\|_2^2 + r_{\text{node}} |||\widehat{\Theta}^l|||_{2,1} \leq \sum_{i=1}^{n} \|\boldsymbol{y}_i - \boldsymbol{g}_{\boldsymbol{\Theta}}[\boldsymbol{x}_i]\|_2^2 + r_{\text{node}} |||\Theta^l|||_{2,1} ,$$

where we use the shorthand $\widehat{\boldsymbol{\Theta}} := \widehat{\boldsymbol{\Theta}}_{\text{node}}$. We then invoke the model in (1) to rewrite this inequality as

$$\sum_{i=1}^{n} \|\boldsymbol{g}_*[\boldsymbol{x}_i] + \boldsymbol{u}_i - \boldsymbol{g}_{\widehat{\boldsymbol{\Theta}}}[\boldsymbol{x}_i]\|_2^2 + r_{\text{node}} |||\widehat{\Theta}^l|||_{2,1} \leq \sum_{i=1}^{n} \|\boldsymbol{g}_*[\boldsymbol{x}_i] + \boldsymbol{u}_i - \boldsymbol{g}_{\boldsymbol{\Theta}}[\boldsymbol{x}_i]\|_2^2 + r_{\text{node}} |||\Theta^l|||_{2,1} .$$

Expanding the squared terms and rearranging the inequality then yields

$$\begin{aligned} \sum_{i=1}^{n} \|\boldsymbol{g}_*[\boldsymbol{x}_i] - \boldsymbol{g}_{\widehat{\boldsymbol{\Theta}}}[\boldsymbol{x}_i]\|_2^2 &\leq \sum_{i=1}^{n} \|\boldsymbol{g}_*[\boldsymbol{x}_i] - \boldsymbol{g}_{\boldsymbol{\Theta}}[\boldsymbol{x}_i]\|_2^2 \\ &+ 2 \sum_{i=1}^{n} \left(\boldsymbol{g}_{\widehat{\boldsymbol{\Theta}}}[\boldsymbol{x}_i]\right)^\top \boldsymbol{u}_i - 2 \sum_{i=1}^{n} \left(\boldsymbol{g}_{\boldsymbol{\Theta}}[\boldsymbol{x}_i]\right)^\top \boldsymbol{u}_i + r_{\text{node}} |||\Theta^l|||_{2,1} - r_{\text{node}} |||\widehat{\Theta}^l|||_{2,1} . \end{aligned}$$

This is our basic inequality.

In the remainder of the proof, we need to bound the first two terms in the last line of the basic inequality. We call these terms the *empirical-process terms*. Using the reformulation of the networks in (7), we can write the empirical-process term of a general parameter $\mathbf{\Gamma} \in \mathcal{M}_{2,1}$ according to

$$2\sum_{i=1}^{n}\big(\boldsymbol{g}_{\mathbf{\Gamma}}[\boldsymbol{x}_i]\big)^\top \boldsymbol{u}_i = 2\sum_{i=1}^{n}\big(\Gamma^l \overline{\boldsymbol{g}}_{\overline{\mathbf{\Gamma}}}[\boldsymbol{x}_i]\big)^\top \boldsymbol{u}_i$$

with $\overline{\mathbf{\Gamma}} \in \overline{\mathcal{M}}_{2,1}$. Using the 1. the properties of transpositions, 2. the definition of the trace function, 3. the cyclic property of the trace function, and 4. the linearity of the trace function yields further

$$\begin{aligned}
2\sum_{i=1}^{n}\big(\boldsymbol{g}_{\mathbf{\Gamma}}[\boldsymbol{x}_i]\big)^\top \boldsymbol{u}_i &= 2\sum_{i=1}^{n}\big(\overline{\boldsymbol{g}}_{\overline{\mathbf{\Gamma}}}[\boldsymbol{x}_i]\big)^\top (\Gamma^l)^\top \boldsymbol{u}_i \\
&= 2\sum_{i=1}^{n}\mathrm{trace}\Big[\big(\overline{\boldsymbol{g}}_{\overline{\mathbf{\Gamma}}}[\boldsymbol{x}_i]\big)^\top (\Gamma^l)^\top \boldsymbol{u}_i\Big] \\
&= 2\sum_{i=1}^{n}\mathrm{trace}\Big[\boldsymbol{u}_i\big(\overline{\boldsymbol{g}}_{\overline{\mathbf{\Gamma}}}[\boldsymbol{x}_i]\big)^\top (\Gamma^l)^\top\Big] \\
&= 2\mathrm{trace}\Big[\Big(\sum_{i=1}^{n}\boldsymbol{u}_i\big(\overline{\boldsymbol{g}}_{\overline{\mathbf{\Gamma}}}[\boldsymbol{x}_i]\big)^\top\Big)(\Gamma^l)^\top\Big].
\end{aligned}$$

Now, 1. denoting the column-vector that corresponds to the kth column of a matrix A by $A_{\bullet k}$, 2. using Hölder's inequality, 3. using Hölder's inequality again, and 4. again Hölder's inequality and our definitions of the elementwise ℓ_∞-and ℓ_1-norms, we find

$$\begin{aligned}
2\sum_{i=1}^{n}\big(\boldsymbol{g}_{\mathbf{\Gamma}}[\boldsymbol{x}_i]\big)^\top \boldsymbol{u}_i &= 2\sum_{k=1}^{p^l}\Big\langle\Big(\sum_{i=1}^{n}\boldsymbol{u}_i\big(\overline{\boldsymbol{g}}_{\overline{\mathbf{\Gamma}}}[\boldsymbol{x}_i]\big)^\top\Big)_{\bullet k}, (\Gamma^l)_{\bullet k}\Big\rangle \\
&\le 2\sum_{k=1}^{p^l}\Big\|\Big(\sum_{i=1}^{n}\boldsymbol{u}_i\big(\overline{\boldsymbol{g}}_{\overline{\mathbf{\Gamma}}}[\boldsymbol{x}_i]\big)^\top\Big)_{\bullet k}\Big\|_2\big\|(\Gamma^l)_{\bullet k}\big\|_2 \\
&\le 2\max_{k\in\{1,\dots,p^l\}}\Big\|\Big(\sum_{i=1}^{n}\boldsymbol{u}_i\big(\overline{\boldsymbol{g}}_{\overline{\mathbf{\Gamma}}}[\boldsymbol{x}_i]\big)^\top\Big)_{\bullet k}\Big\|_2\sum_{k=1}^{p^l}\big\|(\Gamma^l)_{\bullet k}\big\|_2 \\
&\le 2\sqrt{m}\Big|\Big|\Big|\sum_{i=1}^{n}\boldsymbol{u}_i\big(\overline{\boldsymbol{g}}_{\overline{\mathbf{\Gamma}}}[\boldsymbol{x}_i]\big)^\top\Big|\Big|\Big|_\infty |||\Gamma^l|||_{2,1},
\end{aligned}$$

which implies in view of the definition of the effective noise in (8)

$$2\sum_{i=1}^{n}\big(\boldsymbol{g}_{\mathbf{\Gamma}}[\boldsymbol{x}_i]\big)^\top \boldsymbol{u}_i \le r^*_{\mathrm{node}}|||\Gamma^l|||_{2,1}.$$

This inequality is our bound on the empirical-process terms.

We can combine the bound on the empirical-process term and the basic inequality to find

$$\begin{aligned}&\sum_{i=1}^{n} \left\|\boldsymbol{g}_*[\boldsymbol{x}_i] - \boldsymbol{g}_{\widehat{\boldsymbol{\Theta}}}[\boldsymbol{x}_i]\right\|_2^2 \\ &\leq \sum_{i=1}^{n} \left\|\boldsymbol{g}_*[\boldsymbol{x}_i] - \boldsymbol{g}_{\boldsymbol{\Theta}}[\boldsymbol{x}_i]\right\|_2^2 \\ &+ r_{\text{node}}^* |||\widehat{\Theta}^l|||_{2,1} + r_{\text{node}}^* |||\Theta^l|||_{2,1} + r_{\text{node}} |||\Theta^l|||_{2,1} - r_{\text{node}} |||\widehat{\Theta}^l|||_{2,1} .\end{aligned}$$

Using then the assumption $r_{\text{node}} \geq r_{\text{node}}^*$ yields

$$\sum_{i=1}^{n} \left\|\boldsymbol{g}_*[\boldsymbol{x}_i] - \boldsymbol{g}_{\widehat{\boldsymbol{\Theta}}}[\boldsymbol{x}_i]\right\|_2^2 \leq \sum_{i=1}^{n} \left\|\boldsymbol{g}_*[\boldsymbol{x}_i] - \boldsymbol{g}_{\boldsymbol{\Theta}}[\boldsymbol{x}_i]\right\|_2^2 + 2r_{\text{node}} |||\Theta^l|||_{2,1} .$$

Multiplying both sides by $1/n$ and taking the infimum over $\boldsymbol{\Theta} \in \mathcal{M}_{2,1}$ on the right-hand side then gives

$$\frac{1}{n} \sum_{i=1}^{n} \left\|\boldsymbol{g}_*[\boldsymbol{x}_i] - \boldsymbol{g}_{\widehat{\boldsymbol{\Theta}}}[\boldsymbol{x}_i]\right\|_2^2 \leq \inf_{\boldsymbol{\Theta} \in \mathcal{M}_{2,1}} \left\{ \frac{1}{n} \sum_{i=1}^{n} \left\|\boldsymbol{g}_*[\boldsymbol{x}_i] - \boldsymbol{g}_{\boldsymbol{\Theta}}[\boldsymbol{x}_i]\right\|_2^2 + \frac{2r_{\text{node}}}{n} |||\Theta^l|||_{2,1} \right\} .$$

Invoking the definition of the prediction error on Page 3 gives the desired result.

The proof for the connection-sparse estimator is virtually the same. □

D Proof of Proposition 1

In this section, we give a short proof of Proposition 1.

Proof of Proposition 1 Verify the fact that if the all-zeros parameter is neither a solution of (3) nor of (5), all solutions $\widehat{\boldsymbol{\Theta}}_{\text{con}}$ and $\widetilde{\boldsymbol{\Theta}}_{\text{con}}$ of (3) and (5), respectively, satisfy $(\widehat{\Theta}_{\text{con}})^j, (\widetilde{\Theta}_{\text{con}})^j \neq \boldsymbol{0}_{p^{j+1} \times p^j}$ for all $j \in \{0, \ldots, l\}$.

It then follows from the assumed nonnegative homogeneity, $r_{\text{con}} > 0$, and the definition of the estimator in (3) that $|||(\widehat{\Theta}_{\text{con}})^0|||_1, \ldots, |||(\widehat{\Theta}_{\text{con}})^{l-1}|||_1 = 1$ for all solutions $\widehat{\boldsymbol{\Theta}}_{\text{con}}$.

Given a solution $\widetilde{\boldsymbol{\Theta}}_{\text{con}}$ of (5), define $a := |||(\widetilde{\Theta}_{\text{con}})^0|||_1/(l+1) + \cdots + |||(\widetilde{\Theta}_{\text{con}})^l|||_1/(l+1)$ $a := |||(\widetilde{\Theta}_{\text{con}})^0|||_1/(l+1) + \cdots + |||(\widetilde{\Theta}_{\text{con}})^l|||_1/(l+1)$ and verify the fact that $\boldsymbol{\Gamma} \in \mathcal{M}$ with $\Gamma^0 := a(\widetilde{\Theta}_{\text{con}})^0/|||(\widetilde{\Theta}_{\text{con}})^0|||_1, \Gamma^1 := a(\widetilde{\Theta}_{\text{con}})^1/|||(\widetilde{\Theta}_{\text{con}})^1|||_1, \ldots$ has the same value in the objective function as $\widetilde{\boldsymbol{\Theta}}_{\text{con}}$. □

E Proof of Proposition 2

In this section, we establish a proof of Proposition 2. The key tools are the Lipschitz property of Proposition 4 and the entropy bounds of Lemma 1.

Proof of Proposition 2 The main idea is to rewrite the event under consideration in a form that is amenable to known tail bounds for suprema of empirical-processes with sub-Gaussian random variables.

The connection-sparse bound follows from

$$
\begin{aligned}
&P\left\{ r^*_{\text{con}} \geq c\bar{v}_\infty \sqrt{nl\big(\log[2mn\bar{p}]\big)^3} \right\} \\
&\quad = P\left\{ 2 \sup_{\overline{\Psi} \in \overline{\mathcal{M}}_1} \left\| \sum_{i=1}^n \boldsymbol{u}_i \big(\overline{\boldsymbol{g}}_{\overline{\Psi}}[\boldsymbol{x}_i]\big)^\top \right\|_\infty \geq c\bar{v}_\infty \sqrt{nl\big(\log[2mn\bar{p}]\big)^3} \right\} \\
&\quad \leq mp^l \max_{\substack{j \in \{1,\dots,m\} \\ k \in \{1,\dots,p^l\}}} P\left\{ 2 \sup_{\overline{\Psi} \in \overline{\mathcal{M}}_1} \left| \left(\sum_{i=1}^n \boldsymbol{u}_i \big(\overline{\boldsymbol{g}}_{\overline{\Psi}}[\boldsymbol{x}_i]\big)^\top \right)_{jk} \right| \geq c\bar{v}_\infty \sqrt{nl\big(\log[2mn\bar{p}]\big)^3} \right\} \\
&\quad \leq mp^l \cdot \frac{1}{mn\bar{p}} \\
&\quad \leq \frac{1}{n},
\end{aligned}
$$

where we use in turn 1. the definition of r^*_{con} in (8), 2. the union bound, 3. (van de Geer 2000, Corollary 8.3) and our Proposition 4 and Lemma 1, and 4. the inequality $p^l \leq \bar{p} = \sum_{j=0}^l p^{j+1} p^j$ and consolidating the factors. The key concept underlying (van de Geer 2000, Corollary 8.3 on Page 128) is chaining (van der Vaart and Wellner 1996, Page 90).

The same considerations also apply to the node-sparse case, but we get an additional factor $\sqrt{m}$ from the definition of the effective noise in (8) and a factor $\sqrt{\bar{p}}$ from the entropy bound in Lemma 1. The differences between the bounds for the connection- and node-sparse cases in terms of $\bar{v}_\infty$ vs. $\bar{v}_2$ stem from the different Lipschitz constants in Proposition 4. □

F Proof of Proposition 3

Proof of Proposition 3 The proof is based on standard empirical-process theory, including contraction and symmetrization arguments.

Using basic algebra and measure theory, one can easily show that

$$
\begin{aligned}
&\text{risk}[\widehat{\boldsymbol{\Theta}}_{\text{con}}] \leq (1+b)\text{risk}[\boldsymbol{\Theta}^*] + c_b \text{err}[\widehat{\boldsymbol{\Theta}}_{\text{con}}] \\
&+ c_b \left| \frac{1}{n} \sum_{i=1}^n \Big(\big\| \boldsymbol{g}_*[\boldsymbol{x}_i] \\
&- \boldsymbol{g}_{\widehat{\boldsymbol{\Theta}}_{\text{con}}}[\boldsymbol{x}_i] \big\|_2^2 - E\big\| \boldsymbol{g}_*[\boldsymbol{x}_i] \\
&- \boldsymbol{g}_{\widehat{\boldsymbol{\Theta}}_{\text{con}}}[\boldsymbol{x}_i] \big\|_2^2 \Big) \right|
\end{aligned}
$$

for a constant $c_b \in (0, \infty)$ that depends only on b. The first term in this bound is the minimal risk as stated in the proposition, and the second term can be bounded by Corollary 1 and Proposition 2. Hence, it remains to bound the third term.

In view of the law of large numbers, it is reasonable to hope for the third term to be small. But to make this precise, we have to keep in mind that the estimator itself depends on the input vectors. We, therefore, need to prepare the third term for the application of a uniform version of the law of large numbers. Using standard contraction arguments—see (Boucheron et al. 2013, Chapter 11.3), for example—and Hölder's inequality, we can bound the third term by bounding

$$\max\left\{|||(\boldsymbol{\Theta}^*)^l|||_1, |||(\widehat{\boldsymbol{\Theta}}_{\text{con}})^l|||_1\right\} \sup_{\overline{\boldsymbol{\Theta}} \in \overline{\mathcal{M}}_1} \left\| \sum_{i=1}^n \left(\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Theta}}^*}[\boldsymbol{x}_i] - \overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Theta}}}[\boldsymbol{x}_i] - E\left[\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Theta}}^*}[\boldsymbol{x}_i] - \overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Theta}}}[\boldsymbol{x}_i]\right] \right) \right\|_\infty^2,$$

which removes the dependence on the estimator $\widehat{\boldsymbol{\Theta}}_{\text{con}}$ up to the leading factor. To see that we can also neglect that factor, verify (see Proposition 2 and the proof of Theorem 1) that $|||(\widehat{\boldsymbol{\Theta}}_{\text{con}})^l|||_1 \leq 2|||(\boldsymbol{\Theta}^*)^l|||_1$ with high probability as long as $r^*_{\text{con}} \geq c\overline{v}_\infty \sqrt{nl(\log[2mn\overline{p}])^3}$ with c large enough. Consequently, we just need to consider the quantity

$$\sup_{\overline{\boldsymbol{\Theta}} \in \overline{\mathcal{M}}_1} \left\| \sum_{i=1}^n \left(\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Theta}}^*}[\boldsymbol{x}_i] - \overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Theta}}}[\boldsymbol{x}_i] - E\left[\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Theta}}^*}[\boldsymbol{x}_i] - \overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Theta}}}[\boldsymbol{x}_i]\right] \right) \right\|_\infty^2$$

in the following.

The last step is to bring this term in a form that is amenable to our earlier proofs. Using standard symmetrization arguments—see (van der Vaart and Wellner 1996, Chapter 2.3), for example)—we can bound this quantity by bounding

$$\sup_{\overline{\boldsymbol{\Theta}} \in \overline{\mathcal{M}}_1} \left\| \sum_{i=1}^n k_i \left(\overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Theta}}^*}[\boldsymbol{x}_i] - \overline{\boldsymbol{g}}_{\overline{\boldsymbol{\Theta}}}[\boldsymbol{x}_i] \right) \right\|_\infty^2,$$

where $k_1, \ldots, k_n$ are i.i.d. Rademacher random variables. But even though $k_1, \ldots, k_n$ are i.i.d. Rademacher random variables, we do not resort to Rademacher complexities; instead, we use that Rademacher random variables are sub-Gaussian, so that we can then proceed similarly as in the proof of Proposition 2.

The node-sparse case can be treated along the same lines. □

G Extensions

Our proof approach disentangles the specifics of the objective function (proof of Theorem 1), of the network structure (proof of Proposition 4), and of the stochastic terms (proofs of Lemma 1 and Proposition 2). This feature allows one to generalize and extend the results of this paper in straightforward ways. For example, extensions to different noise distributions only need a corresponding version of Proposition 2—with everything else unchanged. One could envision, for example, using concentration inequalities for heavy-tailed distributions such as in Lederer and van de Geer

(2014). Extensions to different loss functions, to give another example, can be established by adjusting Theorem 1 accordingly. This can be done, for example, by invoking ideas from specialized literature on high-dimensional logistic regression such as Li and Lederer (2019). We avoid going into further details to avoid digression; the key message is that the flexibility of the proofs is yet another advantage of our approach.

Acknowledgements I thank Shih-Ting Huang, Mahsa Taheri, Fang Xie, and the anonymous referees for their insightful comments on a draft version of this paper.

Funding Open Access funding enabled and organized by Projekt DEAL.

Declarations

Conflict of interest The author declares that he has no conflict of interest.

References

Alvarez, J., Salzmann, M.: Learning the number of neurons in deep networks. In: Proceedings of the NIPS, pp/ 2270–2278 (2016)

Anthony, M., Bartlett, P.: Neural network learning: theoretical foundations. Cambridge University Press, Cambridge (1999)

Ash, T.: Dynamic node creation in backpropagation networks. Connect. Sci. **1**(4), 365–375 (1989)

Bakin, S.: Adaptive regression and model selection in data mining problems. In: PhD thesis, The Australian National University (1999)

Bartlett, P.: The sample complexity of pattern classification with neural networks: the size of the weights is more important than the size of the network. IEEE Trans. Inform. Theory **44**(2), 525–536 (1998)

Bartlett, P., Mendelson, S.: Rademacher and Gaussian complexities: risk bounds and structural results. J. Mach. Learn. Res. **3**, 463–482 (2002)

Beknazaryan, A.: Function approximation by deep neural networks with parameters $\{0, \pm\frac{1}{2}, \pm 1, 2\}$. arXiv: 2103.08659, 2021

Bello, M.: Enhanced training algorithms, and integrated training/architecture selection for multilayer perceptron networks. IEEE Trans. Neural Netw. **3**(6), 864–875 (1992)

Boucheron, S., Lugosi, G., Massart, P.: Concentration inequalities: a nonasymptotic theory of independence. Oxford University Press, Oxford (2013)

Carl, B.: Inequalities of Bernstein–Jackson-type and the degree of compactness of operators in Banach spaces **35**(3):79–118 (1985)

Changpinyo, S., Sandler, M., Zhmoginov,A.: The power of sparsity in convolutional neural networks (2017). arXiv:1702.06257

Dalalyan, A., Hebiri, M., Lederer, J.: On the prediction performance of the lasso. Bernoulli **23**(1), 552–581 (2017)

Feng, J., Simon, N.: Sparse-input neural networks for high-dimensional nonparametric regression and classification (2017). arxiv:1711.07592

Frankle, J., Carbin, M.: The Lottery Ticket Hypothesis: Finding Sparse, Trainable Neural Networks. In Proc, ICLR (2019)
Glorot , X., Bengio, Y.: Understanding the difficulty of training deep feedforward neural networks. In: Proceedings of the AISTATS, pp. 249–256 (2010)
Golowich, N., Rakhlin, A., Shamir, O.: Size-Independent Sample Complexity of Neural Networks. In Proc, COLT (2018)
Han, S., Mao, H., Dally, W.: Deep Compression: Compressing Deep Neural Networks with Pruning, Trained Quantization and Huffman Coding. In Proc, ICML (2016)
He, K., Zhang, X., Ren, S., Sun, J.: Delving deep into rectifiers: surpassing human-level performance on imagenet classification. In: Proceedngs of the ICCV, pp. 1026–1034 (2015)
Hebiri, M., Lederer, J.: Layer sparsity in neural networks. arXiv:2006.15604 (2020)
Kim, J., Calhoun, V., Shim, E., Lee, J.-H.: Deep neural network with weight sparsity control and pre-training extracts hierarchical features and enhances classification performance: evidence from whole-brain resting-state functional connectivity patterns of schizophrenia. Neuroimage **124**, 127–146 (2016)
Kohler, M., Krzyzak, A., Langer, S.: Estimation of a function of low local dimensionality by deep neural networks. arXiv:1908.11140 (2019)
Koltchinskii, V.: Rademacher penalties and structural risk minimization. IEEE Trans. Inform. Theory **47**(5), 1902–1914 (2001)
Koltchinskii, V., Panchenko, D.: Empirical margin distributions and bounding the generalization error of combined classifiers. Ann. Stat. **30**(1), 1–50 (2002)
LeCun, Y., Bottou, L., Bengio, Y., Haffner, P.: Gradient-based learning applied to document recognition. Proc. IEEE **86**(11), 2278–2324 (1998)
LeCun, Y., Bengio, Y., Hinton, G.: Deep learning. Nature **521**, 436–444 (2015)
Ledent, A., Lei, Y., Kloft, M.: Norm-based generalisation bounds for multi-class convolutional neural networks. arXiv:1905.12430 (2019)
Lederer, J.: Bounds for Rademacher processes via chaining. arXiv:1010.5626 (2010)
Lederer, J.: Risk bounds for robust deep learning. arXiv:2009.06202 (2020a)
Lederer, J.: Optimization landscapes of wide deep neural networks are benign. arXiv:2010.00885 (2020b)
Lederer, J.: Activation functions in artificial neural networks: a systematic overview. arXiv:2101.09957 (2021)
Lederer, J.: Fundamentals of High-Dimensional Statistics. Springer Texts in Statistics (2022)
Lederer, J., van de Geer, S.: New concentration inequalities for suprema of empirical processes. Bernoulli, 20(4) (2014)
Lederer, J., Vogt, M.: Estimating the lasso's effective noise. arXiv:2004.11554 (2020)
Lederer, J., Yu, L., Gaynanova, I.: Oracle inequalities for high-dimensional prediction. Bernoulli **25**(2), 1225–1255 (2019)
Lee, H., Ekanadham, C., Ng,A.: Sparse deep belief net model for visual area V2. In: Proceedings of the NIPS, pp. 873–880 (2008)
Li, W., Lederer, J.: Tuning parameter calibration for ℓ_1-regularized logistic regression. J. Statist. Plann. Inference **202**, 80–98 (2019)
Liang, S., Srikant, R.: Why deep neural networks for function approximation? arXiv:1610.04161 (2016)
Liu, B., Wang, M., Foroosh, H., Tappen, M., Pensky, M.: Sparse convolutional neural networks. In: Proceedingas of the CVPR, pp. 806–814 (2015)
Mao, H., Han, S., Pool, J., Li, W., Liu, X., Wang, Y., Dally, W.: Exploring the granularity of sparsity in convolutional neural networks. In: Proceeding of the CVPR Workshops, pp. 13–20 (2017)
McDiarmid, C.: On the method of bounded differences. Surv. Comb. **141**(1), 148–188 (1989)
Mishkin, D., Matas, J.: All you need is a good init. arXiv preprint arXiv:1511.06422 (2015)
Mishra, A., Latorre, J., Pool, J., Stosic, D., Stosic, D., Venkatesh, G., Yu, C., Micikevicius, P.: Accelerating sparse deep neural networks. arXiv:2104.08378 (2021)
Neyshabur, B., Tomioka, R., Srebro, N.: Norm-based capacity control in neural networks. In Conf. Learn. Theory, pp. 1376–1401 (2015)
Nie, L., Wang, M., Zhang, L., Yan, S., Zhang, B., Chua, T.-S.: Disease inference from health-related questions via sparse deep learning. IEEE Trans. Knowl. Data Eng. **27**(8), 2107–2119 (2015)
Scardapane, S., Comminiello, D., Hussain, A., Uncini, A.: Group sparse regularization for deep neural networks. Neurocomputing **241**, 81–89 (2017)
Schmidhuber, J.: Deep learning in neural networks: an overview. Neural Networks **61**, 85–117 (2015)
Schmidt-Hieber, J.: Nonparametric regression using deep neural networks with ReLU activation function. Ann. Statist **48**(4), 1875–1897 (2020)

Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. In Proc, ICML (2015)
Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R.: Dropout: a simple way to prevent neural networks from overfitting. J. Mach. Learn. Res. **15**(1), 1929–1958 (2014)
Taheri, M., Xie, F., Lederer, J.: Statistical guarantees for regularized neural networks. Neural Networks **142**, 148–161 (2021)
Telgarsky, M.: Benefits of depth in neural networks. In Proc. COLT **49**, 1517–1539 (2016)
Tibshirani, R.: Regression shrinkage and selection via the lasso. J. R. Stat. Soc. Ser. B. Stat. Methodol. **58**(1), 267–288 (1996)
van de Geer, S.: Empirical processes in M-estimation. Cambridge Univ, Press (2000)
van der Vaart, A., Wellner, J.: Weak convergence and empirical processes. Springer (1996)
Wen, W., Wu, C., Wang, Y., Chen, Y., Li, H.: Learning structured sparsity in deep neural networks. In: Proceedings of the NIPS, pp. 2082–2090 (2016)
Xiao, H., Rasul, K., Vollgraf, R.: Fashion-MNIST: a novel image dataset for benchmarking machine learning algorithms. arXiv:1708.07747 (2017)
Yarotsky, D.: Error bounds for approximations with deep ReLU networks. Neural Networks **94**, 103–114 (2017)
Zhuang, R., Lederer, J.: Maximum regularized likelihood estimators: a general prediction theory and applications. Stat **7**(1), e186 (2018)