

Variation Spaces for Multi-Output Neural Networks: Insights on Multi-Task Learning and Network Compression

Joseph Shenouda

JSHENOUDA@WISC.EDU

*Department of Electrical and Computer Engineering
University of Wisconsin–Madison
Madison, WI 53706, USA*

Rahul Parhi*

RPARHI@UCSD.EDU

*Department of Electrical and Computer Engineering
University of California, San Diego
La Jolla, CA 92093, USA*

Kangwook Lee

KANGWOOK.LEE@WISC.EDU

Robert D. Nowak

RDNOWAK@WISC.EDU

*Department of Electrical and Computer Engineering
University of Wisconsin–Madison
Madison, WI 53706, USA*

Editor: Aarti Singh

Abstract

This paper introduces a novel theoretical framework for the analysis of vector-valued neural networks through the development of vector-valued variation spaces, a new class of reproducing kernel Banach spaces. These spaces emerge from studying the regularization effect of weight decay in training networks with activation functions like the rectified linear unit (ReLU). This framework offers a deeper understanding of multi-output networks and their function-space characteristics. A key contribution of this work is the development of a representer theorem for the vector-valued variation spaces. This representer theorem establishes that shallow vector-valued neural networks are the solutions to data-fitting problems over these infinite-dimensional spaces, where the network widths are bounded by the square of the number of training data. This observation reveals that the norm associated with these vector-valued variation spaces encourages the learning of features that are useful for multiple tasks, shedding new light on multi-task learning with neural networks. Finally, this paper develops a connection between weight-decay regularization and the multi-task lasso problem. This connection leads to novel bounds for layer widths in deep networks that depend on the intrinsic dimensions of the training data representations. This insight not only deepens the understanding of the deep network architectural requirements, but also yields a simple convex optimization method for deep neural network compression. The performance of this compression procedure is evaluated on various architectures.

Keywords: deep learning, model compression, multi-task lasso, regularization, sparsity, variation spaces, weight decay.

*. Part of this work was done while RP was with the University of Wisconsin–Madison and then with the École polytechnique fédérale de Lausanne. He is now with the University of California, San Diego.

1. Introduction

The investigation of shallow scalar-valued (single-output) neural networks through the lens of function spaces has been extensively developed. This line of work was pioneered by Barron (1993); Kůrková and Sanguineti (2001, 2002); Mhaskar (2004); Barron et al. (2008); Bach (2017), where those authors study the so-called *variation spaces* of shallow networks. On the other hand, the investigation of vector-valued (multi-output) networks is far less developed. To fill that gap, this paper develops a new framework to investigate the characteristics of the functions learned by vector-valued networks. Since each layer of a deep neural network (DNN) is itself a shallow vector-valued neural network, the proposed investigation is crucial to provide insights into learning with *deep* neural networks.

The developed framework is based on a novel class of Banach spaces which are termed *vector-valued variation spaces*. These spaces arise naturally from a precise characterization of the function-space norm that is regularized by weight decay when training neural networks with homogeneous activation functions, such as the rectified linear unit (ReLU). The vector-valued variation spaces reduce to the classical, scalar-valued variation spaces in the single-output setting. In the general vector-valued setting, these spaces exhibit several intriguing properties that shed light on multi-task learning with neural networks as well as the inductive bias of weight-decay regularization. These properties motivate a simple and computationally efficient method to compress pre-trained deep neural networks (DNNs).

1.1 Organization and Main Contributions

This paper is organized as follows. In Section 2 we review the effect of weight-decay regularization in the training of DNNs with homogeneous activation functions. There, it is revealed that weight-decay regularization is equivalent to a constrained form of multi-task lasso regularization. This observation then naturally leads to a variation norm on the space of vector-valued networks and their (appropriately taken) wide limits, which comprises the vector-valued variation space. These vector-valued variation spaces characterize the functions generated by shallow vector-valued networks (or, equivalently, layers of deep networks).

In Sections 3 and 4 we develop the vector-valued variation spaces and investigate some of their properties. We show that the vector-valued variation spaces are “immune” to the curse of dimensionality in the sense that any function in a vector-valued variation space can be ϵ -approximated (in L^2) by a vector-valued neural network whose width scales as $\epsilon^{-1/2}$, which is independent of both the input and output dimensions. We show that these spaces are reproducing kernel Banach spaces and prove a representer theorem. The representer theorem shows that shallow vector-valued neural networks are solutions to data-fitting problems over the infinite-dimensional vector-valued variation spaces. Furthermore, it guarantees that the widths of these neural network solutions are bounded by the square of the number of training data, irrespective of the input and output dimensions.

In Section 5 we show that the norm associated with the vector-valued variation spaces (and hence weight decay in the case of homogeneous activation functions) promotes a remarkable “neuron sharing” property of the solutions. This refers to the fact that the norm encourages solutions in which each neuron contributes to *every* output function, as opposed to networks in which disjoint sets of neurons are used to represent different output func-

tions. This indicates how different outputs may influence each other in the training process. It also provides a new viewpoint of multi-task learning with neural networks: Solutions are encouraged to learn features that are useful for multiple tasks.

Finally, in Section 6 we show that weight-decay regularization in DNNs with homogeneous activation functions is tightly linked to a convex multi-task lasso problem. With this link, we present new bounds on the sufficient widths of layers in DNNs that depend on the intrinsic dimensions of learned data representations at each layer. This result is based on a novel characterization of the sparsity of multi-task lasso solutions, which may be of independent interest. Notably, if the dimensions (ranks of data representation matrices) are low, then there exist layers whose widths are much narrower than the number of data while still being optimal solutions. This leads to a principled approach to DNN compression. This approach is computationally efficient and can dramatically reduce layer widths without the sacrifice of the learned representations, data-fitting error, or optimality (in terms of the weight decay objective and the variation norm). We evaluate the performance of our proposed compression approach on various architectures in Section 7.

2. Weight Decay and the Neural Balance Theorem

Let f_{θ} be a DNN with weights θ and let $\{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^N$ be a dataset where each $\mathbf{x}_i \in \mathbb{R}^d$ and $\mathbf{y}_i \in \mathbb{R}^D$. To fit the data, a common approach is to train the network using gradient descent with *weight decay*. This corresponds to finding a solution to the optimization problem

$$\min_{\theta} \sum_{i=1}^N \mathcal{L}(\mathbf{y}_i, f_{\theta}(\mathbf{x}_i)) + \frac{\lambda}{2} \|\theta\|_2^2, \quad (1)$$

where $\mathcal{L}(\cdot, \cdot)$ is a loss function that is lower semicontinuous in its second argument and $\lambda > 0$ is the regularization parameter (Hanson and Pratt, 1988). Contemporary neural network architectures include a variety of different building blocks, but a commonality among them is neurons with the ReLU activation function. A ReLU neuron is a map of the form $\mathbf{z} \mapsto \mathbf{v}(\mathbf{w}^T \mathbf{z})_+$ with $\mathbf{z}, \mathbf{w} \in \mathbb{R}^{d_{\text{in}}}$, $\mathbf{v} \in \mathbb{R}^{d_{\text{out}}}$, and where $(\cdot)_+ := \max\{0, \cdot\}$ is the ReLU.

More generally, an activation function $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ is said to be *positively homogeneous of degree one* (or simply homogeneous) if, for any $\gamma > 0$, $\sigma(\gamma t) = \gamma \sigma(t)$ for all $t \in \mathbb{R}$. The ReLU, leaky ReLU, absolute value, and linear activation functions satisfy this property. A key observation in both theory (Grandvalet, 1998; Grandvalet and Canu, 1998; Neyshabur et al., 2015; Parhi and Nowak, 2023b) and practice (Kunin et al., 2021, Figure 5) is that in any solution to (1), the 2-norms of the input and output weights of each neuron with a homogeneous activation function must be balanced. This phenomenon is summarized in the *neural balance theorem* (NBT) (Yang et al., 2022; Parhi and Nowak, 2023b).

Theorem 1 (Neural Balance Theorem) *Let f_{θ} be a DNN of any architecture such that θ minimizes (1). Then, the weights satisfy the following balance constraint: If \mathbf{w} and \mathbf{v} denote the input and output weights of any neuron with a homogeneous activation function, then $\|\mathbf{w}\|_2 = \|\mathbf{v}\|_2$.*

While the NBT is a simple observation, it allows for an alternative perspective on weight-decay regularization. Indeed, first consider the weight-decay regularized problem for

a shallow vector-valued neural network whose activation function $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ is homogeneous

$$\min_{\{\mathbf{w}_k, \mathbf{v}_k\}_{k=1}^K} \sum_{i=1}^N \mathcal{L} \left(\mathbf{y}_i, \sum_{k=1}^K \mathbf{v}_k \sigma(\mathbf{w}_k^T \bar{\mathbf{x}}_i) \right) + \frac{\lambda}{2} \sum_{k=1}^K \|\mathbf{v}_k\|_2^2 + \|\mathbf{w}_k\|_2^2, \quad (2)$$

where $\mathbf{v}_k \in \mathbb{R}^D$, $\mathbf{w}_k \in \mathbb{R}^{d+1}$ and $\bar{\mathbf{x}} := [\mathbf{x} \ 1]^T \in \mathbb{R}^{d+1}$ augments \mathbf{x} to account for a bias term. Observe that, by the NBT, any solution to this weight-decay regularized problem is always a solution to the so-called *path-norm* regularized problem

$$\min_{\{\mathbf{w}_k, \mathbf{v}_k\}_{k=1}^K} \sum_{i=1}^N \mathcal{L} \left(\mathbf{y}_i, \sum_{k=1}^K \mathbf{v}_k \sigma(\mathbf{w}_k^T \bar{\mathbf{x}}_i) \right) + \lambda \sum_{k=1}^K \|\mathbf{v}_k\|_2 \|\mathbf{w}_k\|_2. \quad (3)$$

Furthermore, thanks to the homogeneity of σ , any solution to (3) is always a solution to the constrained problem

$$\min_{\{\mathbf{w}_k, \mathbf{v}_k\}_{k=1}^K} \sum_{i=1}^N \mathcal{L} \left(\mathbf{y}_i, \sum_{k=1}^K \mathbf{v}_k \sigma(\mathbf{w}_k^T \bar{\mathbf{x}}_i) \right) + \lambda \sum_{k=1}^K \|\mathbf{v}_k\|_2 \quad \text{s.t.} \quad \|\mathbf{w}_k\|_2 = 1, \quad k = 1, \dots, K, \quad (4)$$

upon “absorbing” the magnitude of the input weights into the output weights.¹ Clearly, any solution to (4) is always a solution to (3). Moreover, any solution to (3) is always a solution to (2), after balancing the weights. Thus, the problems (2)–(4) should be viewed as equivalent. In particular, the regularizer in (4) is the *multi-task lasso* regularizer which is known to promote a kind of sparsity (Obozinski et al., 2006, 2010; Argyriou et al., 2008). We also remark that the neural balance theorem and the equivalence between (2)–(4) holds in the case of unregularized biases (cf. Parhi and Nowak, 2023b, pp. 65–66). Furthermore, all of the results of this paper also hold in that setting as well. For notational convenience, we focus on the regularized bias scenario (see also Remark 2).

The equivalence between (2)–(4) can then be extended to any layer in a deep neural network with homogeneous activation functions (see Section 6.1). This “secret sparsity of weight decay” has many remarkable implications on the understanding of DNNs trained with weight decay (see Parhi and Nowak, 2023b, for more details). Furthermore, this connection leads to a natural characterization of the function spaces of vector-valued neural networks, which is developed in the next section.

3. Vector-Valued Variation Spaces

The regularization term $\sum_{k=1}^K \|\mathbf{v}_k\|_2$ in the optimization (4) may be viewed as the *representational cost* of a network. We can adopt this as a measure of the cost of any shallow network or layer in a DNN. Consider a shallow neural network or network layer of the form

$$\mathbf{x} \mapsto \sum_{k=1}^K \mathbf{v}_k \sigma(\mathbf{w}_k^T \bar{\mathbf{x}}), \quad \mathbf{x} \in \mathbb{R}^d, \quad (5)$$

1. That is, by reparameterizing the weights of each neuron as $(\mathbf{v}_k, \mathbf{w}_k) \leftarrow (\mathbf{v}_k \|\mathbf{w}_k\|_2, \mathbf{w}_k / \|\mathbf{w}_k\|_2)$, $k = 1, \dots, K$. Observe that this reparameterization does not change the overall function realized by the neural network.

with $\mathbf{w}_k \in \mathbb{R}^{d+1}$ such that $\|\mathbf{w}_k\|_2 = 1$, and $\mathbf{v}_k \in \mathbb{R}^D$, $k = 1, \dots, K$. If a function f can be represented by a finite-width network, then its representational cost is

$$\|f\| := \inf_{\substack{\{\mathbf{w}_k, \mathbf{v}_k\}_{k=1}^K \\ K \in \mathbb{N}}} \sum_{k=1}^K \|\mathbf{v}_k\|_2 \quad \text{s.t.} \quad f = \left(\mathbf{x} \mapsto \sum_{k=1}^K \mathbf{v}_k \sigma(\mathbf{w}_k^T \bar{\mathbf{x}}) \right). \quad (6)$$

The inf arises due to the fact that there are many neural network representations of the same function f . The inf selects the one with the *lowest* representational cost. The reader can quickly verify that this quantity is a *bona fide* norm since it satisfies the following properties.

1. **Triangle Inequality:** $\|f + g\| \leq \|f\| + \|g\|$.
2. **Homogeneity:** $\|\alpha f\| = |\alpha| \|f\|$ for $\alpha \in \mathbb{R}$.
3. **Positive Definiteness:** $\|f\| = 0$ if and only if $f = 0$.

The space of all finite-norm neural network functions of the form (5) and their limits² defines a Banach space that we call a vector-valued variation space. These spaces, which capture all functions that can be represented or approximated by neural networks with finite norms, are developed in the sequel. We begin by reviewing past constructions of variation spaces for scalar (single-output) networks.

3.1 Scalar-Valued Variation Spaces

In this subsection, we review the definition of the classical, scalar-valued variation spaces. The results stated here can be found in the papers of Bengio et al. (2005); Bach (2017); Parhi and Nowak (2021); Siegel and Xu (2023). The main idea is to consider shallow neural networks with possibly continuously many neurons. These neural networks are parameterized by a finite (Radon) measure. The scalar-valued variation space is the space of functions that map $\mathbb{R}^d \rightarrow \mathbb{R}$

$$\mathcal{V}_\sigma(\mathbb{R}^d) := \left\{ f(\mathbf{x}) = \int_{\mathbb{S}^d} \sigma(\mathbf{w}^T \bar{\mathbf{x}}) d\nu(\mathbf{w}) : \mathbf{x} \in \mathbb{R}^d, \nu \in \mathcal{M}(\mathbb{S}^d) \right\}, \quad (7)$$

where $\mathbb{S}^d := \{\mathbf{w} \in \mathbb{R}^{d+1} : \|\mathbf{w}\|_2 = 1\}$ is the unit sphere, $\bar{\mathbf{x}} := [\mathbf{x} \ 1]^T \in \mathbb{R}^{d+1}$ augments \mathbf{x} to account for a bias term, and $\mathcal{M}(\mathbb{S}^d)$ is the space of finite (Radon) measures. The measure ν plays the role of the output weight of each neuron. Here, and in the rest of this paper, the activation function $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ will always be assumed to be continuous.

Since each function $f \in \mathcal{V}_\sigma(\mathbb{R}^d)$ is parameterized by a measure $\nu \in \mathcal{M}(\mathbb{S}^d)$, we introduce the notation

$$f_\nu(\mathbf{x}) := \int_{\mathbb{S}^d} \sigma(\mathbf{w}^T \bar{\mathbf{x}}) d\nu(\mathbf{w}), \quad \mathbf{x} \in \mathbb{R}^d. \quad (8)$$

It is well-known that the space $\mathcal{V}_\sigma(\mathbb{R}^d)$ is a Banach space when equipped with the norm

$$\|f\|_{\mathcal{V}_\sigma(\mathbb{R}^d)} := \inf_{\substack{\nu \in \mathcal{M}(\mathbb{S}^d) \\ f = f_\nu}} \|\nu\|_{\mathcal{M}(\mathbb{S}^d)}, \quad (9)$$

2. These are weak* limits as opposed to norm limits (cf., E et al., 2022).

where $\|\cdot\|_{\mathcal{M}(\mathbb{S}^d)}$ denotes the total variation norm in the sense of measures. If f is a finite-width network, this norm is in fact equal to the norm defined (6), as shown in (12) below.

As in (6), the inf arises since the dictionary of neurons $\{\mathbf{x} \mapsto \sigma(\mathbf{w}^T \bar{\mathbf{x}})\}_{\mathbf{w} \in \mathbb{S}^d}$ is highly redundant. Thus, there are many different representations for a given $f \in \mathcal{V}_\sigma(\mathbb{R}^d)$. By choosing the representation with the smallest total variation norm, (9) defines a valid Banach norm on $\mathcal{V}_\sigma(\mathbb{R}^d)$ (see, e.g., Siegel and Xu, 2023, Lemma 3). Here, we use the following definition of $\|\cdot\|_{\mathcal{M}(\mathbb{S}^d)}$

$$\|\nu\|_{\mathcal{M}(\mathbb{S}^d)} := \sup_{\substack{\mathbb{S}^d = \bigcup_{i=1}^n A_i \\ n \in \mathbb{N}}} \sum_{i=1}^n |\nu(A_i)|, \quad (10)$$

where the sup is taken over all partitions of \mathbb{S}^d (i.e., $A_i \cap A_j = \emptyset$ for $i \neq j$). This definition is equal to the more conventional definitions based on the Jordan decomposition of a measure or as a dual norm (Diestel and Uhl, 1977; Bredies and Holler, 2020). We use the definition in (10) as an analogous definition will play an important role in the vector-valued case.

Consider a single neuron $\phi_{v, \mathbf{w}}(\mathbf{x}) := v\sigma(\mathbf{w}^T \bar{\mathbf{x}})$, $\mathbf{x} \in \mathbb{R}^d$, where $v \in \mathbb{R}$ and $\mathbf{w} \in \mathbb{S}^d$. In this scenario, it is clear that the inf in (9) is achieved by the one scaled Dirac measure $v\delta_{\mathbf{w}}$ (since any other combination of dictionary elements that represent a single neuron would have a larger norm). Thus, if we have the shallow neural network

$$\mathbf{x} \mapsto \sum_{k=1}^K v_k \sigma(\mathbf{w}_k^T \bar{\mathbf{x}}), \quad \mathbf{x} \in \mathbb{R}^d, \quad (11)$$

where the input weights $\mathbf{w}_k \in \mathbb{S}^d$ are all unique, it is known (cf., Bach, 2017, p. 6) that the inf in (9) is achieved by $\sum_{k=1}^K v_k \delta_{\mathbf{w}_k}$. Therefore,

$$\left\| \mathbf{x} \mapsto \sum_{k=1}^K v_k \sigma(\mathbf{w}_k^T \bar{\mathbf{x}}) \right\|_{\mathcal{V}_\sigma(\mathbb{R}^d)} = \left\| \sum_{k=1}^K v_k \delta_{\mathbf{w}_k} \right\|_{\mathcal{M}(\mathbb{S}^d)} = \sum_{k=1}^K \|v_k \delta_{\mathbf{w}_k}\|_{\mathcal{M}(\mathbb{S}^d)} = \sum_{k=1}^K |v_k|, \quad (12)$$

where the second equality uses the fact that the Dirac measures have disjoint support and the third equality follows from the property that $\|a\delta_{\mathbf{u}}\|_{\mathcal{M}(\mathbb{S}^d)} = |a|$, where $a \in \mathbb{R}$ and $\mathbf{u} \in \mathbb{S}^d$. The final quantity on the right above is the ℓ^1 -norm of all the output weights, which is precisely the regularization term in (4). Furthermore, if the activation function is homogeneous, then, by the NBT (Theorem 1), the regularization of this quantity is equivalent to weight-decay regularization. Therefore, training a scalar-output shallow neural network with weight decay penalizes the variation norm of the network.

Remark 2 *To consider the unregularized bias scenario, the integral in (8) would instead take the form of an integral combination of neurons of the form $\mathbf{x} \mapsto \sigma(\mathbf{w}^T \mathbf{x} + b)$, $(\mathbf{w}, b) \in \mathbb{S}^{d-1} \times \mathbb{R}$, against the measure $\nu \in \mathcal{M}(\mathbb{S}^{d-1} \times \mathbb{R})$. The precise details are discussed by Ongie et al. (2020, Appendix B) in the case when σ is the ReLU. In that case, the corresponding variation space has an analytic description via the Radon transform (Ongie et al., 2020; Parhi and Nowak, 2021). In this paper, we focus on the regularized bias scenario for notational convenience. As discussed in Section 2, all results presented in this paper hold in the unregularized bias scenario.*

3.2 Vector-Valued Variation Spaces

The vector-valued variation space is the set of functions defined analogously to the scalar-valued variation spaces:

$$\mathcal{V}_\sigma(\mathbb{R}^d; \mathbb{R}^D) := \left\{ f(\mathbf{x}) = \int_{\mathbb{S}^d} \sigma(\mathbf{w}^T \bar{\mathbf{x}}) d\boldsymbol{\nu}(\mathbf{w}) : \mathbf{x} \in \mathbb{R}^d, \boldsymbol{\nu} \in \mathcal{M}(\mathbb{S}^d; \mathbb{R}^D) \right\}, \quad (13)$$

where $\boldsymbol{\nu} = (\nu_1, \dots, \nu_D)$ is now a vector-valued measure (which takes values in \mathbb{R}^D as opposed to \mathbb{R}) and plays the role of the output weight vector of each neuron. Analogous to (10), define the total variation norm of a measure $\boldsymbol{\nu}$ as

$$\|\boldsymbol{\nu}\|_{2,\mathcal{M}} := \sup_{\substack{\mathbb{S}^d = \bigcup_{i=1}^n A_i \\ n \in \mathbb{N}}} \sum_{i=1}^n \|\boldsymbol{\nu}(A_i)\|_2 = \sup_{\substack{\mathbb{S}^d = \bigcup_{i=1}^n A_i \\ n \in \mathbb{N}}} \sum_{i=1}^n \left(\sum_{j=1}^D |\nu_j(A_i)|^2 \right)^{1/2}. \quad (14)$$

The choice of norm in the above display is a common choice for the total variation norm of a vector-valued measure. Furthermore, $(\mathcal{M}(\mathbb{S}^d; \mathbb{R}^D), \|\cdot\|_{2,\mathcal{M}})$ is a Banach space. We refer the reader to the monograph of Diestel and Uhl (1977) for a full treatment of vector-valued measures and the accompanying results. This leads to a norm on functions of the form

$$f_{\boldsymbol{\nu}}(\mathbf{x}) := \int_{\mathbb{S}^d} \sigma(\mathbf{w}^T \bar{\mathbf{x}}) d\boldsymbol{\nu}(\mathbf{w}), \quad \mathbf{x} \in \mathbb{R}^d, \quad (15)$$

as

$$\|f\|_{\mathcal{V}_\sigma(\mathbb{R}^d; \mathbb{R}^D)} := \inf_{\substack{\boldsymbol{\nu} \in \mathcal{M}(\mathbb{S}^d; \mathbb{R}^D) \\ f = f_{\boldsymbol{\nu}}}} \|\boldsymbol{\nu}\|_{2,\mathcal{M}}. \quad (16)$$

To connect back to the optimization (4), consider a single vector-valued neuron $\phi_{\mathbf{v}, \mathbf{w}}(\mathbf{x}) := \mathbf{v} \sigma(\mathbf{w}^T \bar{\mathbf{x}})$, $\mathbf{x} \in \mathbb{R}^d$, where $\mathbf{v} \in \mathbb{R}^D$ and $\mathbf{w} \in \mathbb{S}^d$. As in the scalar-valued scenario, the inf is achieved by the measure $\mathbf{v} \delta_{\mathbf{w}}$. This is a vector (in \mathbb{R}^D) multiplied by a scalar-valued Dirac measure and is therefore a vector-valued measure in $\mathcal{M}(\mathbb{S}^d; \mathbb{R}^D)$. Thus, as in the scalar-valued case, for the shallow vector-valued neural network

$$\mathbf{x} \mapsto \sum_{k=1}^K \mathbf{v}_k \sigma(\mathbf{w}_k^T \bar{\mathbf{x}}), \quad \mathbf{x} \in \mathbb{R}^d, \quad (17)$$

where the input weights $\mathbf{w}_k \in \mathbb{S}^d$ are all unique, the inf in (16) is achieved by

$$\sum_{k=1}^K \mathbf{v}_k \delta_{\mathbf{w}_k}. \quad (18)$$

Writing $\mathbf{v}_k = (v_{k,1}, \dots, v_{k,D})$, a calculation reveals that

$$\left\| \mathbf{x} \mapsto \sum_{k=1}^K \mathbf{v}_k \sigma(\mathbf{w}_k^T \bar{\mathbf{x}}) \right\|_{\mathcal{V}_\sigma(\mathbb{R}^d; \mathbb{R}^D)} = \left\| \sum_{k=1}^K \mathbf{v}_k \delta_{\mathbf{w}_k} \right\|_{2,\mathcal{M}} = \sum_{k=1}^K \|\mathbf{v}_k \delta_{\mathbf{w}_k}\|_{2,\mathcal{M}} = \sum_{k=1}^K \|\mathbf{v}_k\|_2, \quad (19)$$

where we used the property that $\|\mathbf{a}\delta_{\mathbf{u}}\|_{2,\mathcal{M}} = \|\mathbf{a}\|_2$, where $\mathbf{a} \in \mathbb{R}^D$ and $\mathbf{u} \in \mathbb{S}^d$ (cf., Boyer et al., 2019, Section 4.2.3). From (4), we see immediately see that this choice of norm on $\mathcal{M}(\mathbb{S}^d; \mathbb{R}^D)$ corresponds to weight-decay regularization when σ is homogeneous.

We remark that several other norms have been previously proposed for vector-valued networks/measures (see Parhi and Nowak, 2022; Korolev, 2022). These prior works essentially treat each output separately. This type of norm is fundamentally different than the one proposed in (16). Furthermore, these other norms do not correspond to weight-decay regularization. These different norms and their relationships are discussed in Appendix A.

3.2.1 THE CURSE OF DIMENSIONALITY

The space $\mathcal{V}_\sigma(\mathbb{R}^d; \mathbb{R}^D)$ has intriguing approximation properties, which carry over from the scalar-valued case (which are known). Let $Q^d := [0, 1]^d$ denote the unit cube in \mathbb{R}^d and define

$$\mathcal{V}_\sigma(Q^d; \mathbb{R}^D) := \{f : Q^d \rightarrow \mathbb{R}^D : \text{there exists } g \in \mathcal{V}_\sigma(\mathbb{R}^d; \mathbb{R}^D) \text{ such that } g|_{Q^d} = f\}. \quad (20)$$

This space is a Banach space when equipped with the norm

$$\|f\|_{\mathcal{V}_\sigma(Q^d; \mathbb{R}^D)} := \inf_{\substack{g \in \mathcal{V}_\sigma(\mathbb{R}^d; \mathbb{R}^D) \\ g|_{Q^d} = f}} \|g\|_{\mathcal{V}_\sigma(\mathbb{R}^d; \mathbb{R}^D)}. \quad (21)$$

By restricting our attention to a bounded domain, we have the continuous embedding $\mathcal{V}_\sigma(Q^d; \mathbb{R}^D) \subset L^2(\mathbb{R}^d; \mathbb{R}^D)$. For each $f = (f_1, \dots, f_D) \in \mathcal{V}_\sigma(Q^d; \mathbb{R}^D)$, we have, for $j = 1, \dots, D$, that $f_j \in \mathcal{V}_\sigma(Q^d)$ (the scalar-valued variation space restricted to Q^d).

In the scalar-valued case, the Maurey–Jones–Barron lemma (Pisier, 1981; Jones, 1992; Barron, 1993) says that, given $f_j \in \mathcal{V}_\sigma(Q^d)$, there exists a K -term approximant

$$f_j^K(\mathbf{x}) = \sum_{k=1}^K v_{k,j} \sigma(\mathbf{w}_{k,j}^T \bar{\mathbf{x}}) \quad (22)$$

with $v_{k,j} \in \mathbb{R}$ and $\mathbf{w}_{k,j} \in \mathbb{S}^d$ such that

$$\|f_j - f_j^K\|_{L^2(Q^d)} \leq C_0 C_\sigma \|f_j\|_{\mathcal{V}_\sigma(Q^d)} K^{-1/2}, \quad (23)$$

where $C_0 > 0$ is an absolute constant independent of d and

$$C_\sigma = \sup_{\mathbf{w} \in \mathbb{S}^d} \|\mathbf{x} \mapsto \sigma(\mathbf{w}^T \bar{\mathbf{x}})\|_{L^2(Q^d)}. \quad (24)$$

This result is remarkable since it establishes that, for any function in $\mathcal{V}_\sigma(Q^d)$, there exists an approximant whose error decays at a rate independent of the input dimension d . Here, we note that the constant C_σ depends on (essentially) the volume the domain Q^d since σ is continuous, and is therefore also independent of the input dimension d (since the volume of Q^d is 1). This result has a straightforward extension to the vector-valued case. This is summarized in the following theorem whose proof can be found in Appendix B, which shows that any function in $\mathcal{V}_\sigma(Q^d; \mathbb{R}^D)$ can be approximated in L^2 by a network of width K with an error that decays at a rate $K^{-1/2}$, independent of the input and output dimensions d and D .

Theorem 3 *Given $f \in \mathcal{V}_\sigma(Q^d; \mathbb{R}^D)$, there exists a K -term approximant of the form*

$$f_K(\mathbf{x}) = \sum_{k=1}^K \mathbf{v}_k \sigma(\mathbf{w}_k^T \bar{\mathbf{x}}), \quad \mathbf{x} \in \mathbb{R}^d, \quad (25)$$

with $\mathbf{v}_k \in \mathbb{R}^D$ and $\mathbf{w}_k \in \mathbb{S}^d$ such that

$$\|f - f_K\|_{L^2(Q^d; \mathbb{R}^D)} \leq C_0 C_\sigma D^{\frac{3}{2}} \|f\|_{\mathcal{V}_\sigma(Q^d; \mathbb{R}^D)} K^{-1/2}, \quad (26)$$

where C_0 and C_σ are as above and the $L^2(Q^d; \mathbb{R}^D)$ -norm is specified by

$$\|f\|_{L^2(Q^d; \mathbb{R}^D)} = \left(\int_{Q^d} \|f(\mathbf{x})\|_2^2 d\mathbf{x} \right)^{1/2}. \quad (27)$$

Remark 4 *Theorem 3 sets the stage for the investigation of dimension-free nonlinear minimax rates of estimation for multi-output neural networks. These results have recently been carried out for single-output neural networks trained with weight decay by Parhi and Nowak (2023a).*

4. Representer Theorem for Vector-Valued Variation Spaces

The discussion above has shown that finite-width neural networks are effective at approximating functions in vector-valued variation spaces. This section considers the problem of fitting data with functions in $\mathcal{V}_\sigma(\mathbb{R}^d; \mathbb{R}^D)$. The main result here is a *representer theorem* that shows that finite-width neural networks are solutions to such problems and bounds the (sufficient) widths of networks in terms of the number of data points. This has an important implication: The infinite-dimensional learning problem can be solved by training a finite-width neural network, and increasing the width beyond the given bound will not yield a smaller objective value.

Theorem 5 *Let $(\mathbf{x}_1, \mathbf{y}_1), \dots, (\mathbf{x}_N, \mathbf{y}_N) \in \mathbb{R}^d \times \mathbb{R}^D$ be a finite dataset. Then, there exists a solution to the variational problem*

$$\inf_{f \in \mathcal{V}_\sigma(\mathbb{R}^d; \mathbb{R}^D)} \sum_{i=1}^N \mathcal{L}(\mathbf{y}_i, f(\mathbf{x}_i)) + \lambda \|f\|_{\mathcal{V}_\sigma(\mathbb{R}^d; \mathbb{R}^D)}, \quad \lambda > 0, \quad (28)$$

where the loss function $\mathcal{L}(\cdot, \cdot)$ is lower semicontinuous in its second argument, which takes the form

$$f^*(\mathbf{x}) = \sum_{k=1}^{K_0} \mathbf{v}_k \sigma(\mathbf{w}_k^T \bar{\mathbf{x}}), \quad \mathbf{x} \in \mathbb{R}^d, \quad (29)$$

where $K_0 \leq \min\{N^2, ND\}$. Here, $\mathbf{v}_k \in \mathbb{R}^D$ and $\mathbf{w}_k \in \mathbb{S}^d$.

The proof of Theorem 5 appears in Appendix C. What is remarkable here is the bound $K_0 \leq N^2$. Indeed, for large D , this bound improves the bound of $ND + 1$ predicted by Carathéodory's theorem (Bredies et al., 2024). Note that Theorem 5 applies to a variation

space based on any continuous activation function σ . Furthermore, we mention again that the result also holds in the unregularized bias scenario upon the appropriate modifications discussed in Remark 2. If σ is homogeneous, then the regularization is equivalent to weight decay. The result of Theorem 5 also applies to the entire solution set to (28). Indeed, the solution set to (28) is nonempty, convex, weak* compact, and its extreme points take the form of sparse vector-valued networks as in (29).

Corollary 6 *Let $\mathcal{L}(\cdot, \cdot)$ be lower semicontinuous in its second argument. Moreover, let $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ be any homogeneous activation function. Then, any solution to the neural network training problem*

$$\min_{\{\mathbf{w}_k, \mathbf{v}_k\}_{k=1}^K} \sum_{i=1}^N \mathcal{L} \left(\mathbf{y}_i, \sum_{k=1}^K \mathbf{v}_k \sigma(\mathbf{w}_k^T \bar{\mathbf{x}}_i) \right) + \frac{\lambda}{2} \sum_{k=1}^K \|\mathbf{v}_k\|_2^2 + \|\mathbf{w}_k\|_2^2, \quad \lambda > 0, \quad (30)$$

is a solution to the variational problem (28), so long as $K \geq \min\{N^2, ND\}$.

Proof By Theorem 5, there always exists a solution to (28) that takes the form of a shallow vector-valued neural network with less than $\min\{N^2, ND\}$ neurons. Thus, a solution to (28) must exist in the space of all shallow vector-valued neural networks with $K \geq \min\{N^2, ND\}$ neurons. By (19), any solution to

$$\min_{\{\mathbf{w}_k, \mathbf{v}_k\}_{k=1}^K} \sum_{i=1}^N \mathcal{L} \left(\mathbf{y}_i, \sum_{k=1}^K \mathbf{v}_k \sigma(\mathbf{w}_k^T \bar{\mathbf{x}}_i) \right) + \lambda \sum_{k=1}^K \|\mathbf{v}_k\|_2, \quad \text{s.t. } \|\mathbf{w}_k\|_2 = 1, k = 1, \dots, K, \quad (31)$$

is a solution to (28). The result then follows by the equivalence between the problem in the above display with (30) as discussed in Section 2. \blacksquare

Remark 7 *The vector-valued variation space $(\mathcal{V}_\sigma(\mathbb{R}^d; \mathbb{R}^D), \|\cdot\|_{\mathcal{V}_\sigma(\mathbb{R}^d; \mathbb{R}^D)})$ is an example of a reproducing kernel Banach space (RKBS) (Zhang et al., 2009; Lin et al., 2022). Indeed, this can be readily deduced from the fact that the scalar-valued variation spaces are reproducing kernel Banach spaces (Bartolucci et al., 2023; Spek et al., 2022).*

4.1 A Representer Theorem for Deep Neural Networks

Theorem 5 can be extended to DNNs by using the techniques developed by Parhi and Nowak (2022, Theorem 3.2). The extension is summarized in Theorem 8 and the proof can be found in Appendix D.

Theorem 8 *Let $(\mathbf{x}_1, \mathbf{y}_1), \dots, (\mathbf{x}_N, \mathbf{y}_N) \in \mathbb{R}^{d_0} \times \mathbb{R}^{d_L}$ be a finite dataset. Then, there exists a solution to the variational problem*

$$\inf_{\substack{f^{(1)}, \dots, f^{(L)} \\ f^{(\ell)} \in \mathcal{V}_\sigma(\mathbb{R}^{d_{\ell-1}}; \mathbb{R}^{d_\ell})}} \sum_{i=1}^N \mathcal{L}(\mathbf{y}_i, f^{(L)} \circ \dots \circ f^{(1)}(\mathbf{x}_i)) + \lambda \sum_{\ell=1}^L \|f^{(\ell)}\|_{\mathcal{V}_\sigma(\mathbb{R}^{d_{\ell-1}}; \mathbb{R}^{d_\ell})}, \quad \lambda > 0, \quad (32)$$

where the loss function $\mathcal{L}(\cdot, \cdot)$ is lower semicontinuous in its second argument, which takes the form

$$f^*(\mathbf{x}) = \mathbf{A}^{(L)} \circ \boldsymbol{\sigma} \circ \mathbf{A}^{(L-1)} \circ \dots \circ \boldsymbol{\sigma} \circ \mathbf{A}^{(1)}(\mathbf{x}) \quad \mathbf{x} \in \mathbb{R}^d, \quad (33)$$

where, for each layer $\ell = 1, \dots, L$, the function $\mathbf{A}^{(\ell)}(\mathbf{z}) = \mathbf{V}^{(\ell)}\mathbf{z} - \mathbf{b}^{(\ell)}$ is an affine mapping with weight matrix $\mathbf{V}^{(\ell)} \in \mathbb{R}^{d_\ell \times d_{\ell-1}}$ and bias vector $\mathbf{b}^{(\ell)} \in \mathbb{R}^{d_\ell}$, where $\boldsymbol{\sigma}$ applies the activation function $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ component-wise.

5. Neuron Sharing in Neural Network Solutions

This section describes a remarkable ‘‘neuron sharing’’ property of solutions to the weight decay optimization objective (1) and the variational problem (28). This refers to the fact that each neuron in a solution is encouraged to contribute to every output, as opposed to networks in which different neurons are used to represent different output functions. This indicates how different outputs may influence each other in the training process. It also provides a new viewpoint for multi-task learning with neural networks: When the activation function is homogeneous, weight decay encourages the learning of features that are useful for multiple tasks/outputs.

The neuron sharing property arises from the definition of the $\mathcal{V}_\sigma(\mathbb{R}^d; \mathbb{R}^D)$ -norm in (16). In particular, the $\|\cdot\|_{\mathcal{V}_\sigma(\mathbb{R}^d; \mathbb{R}^D)}$ -norm regularized problem

$$\min_{f \in \mathcal{V}_\sigma(\mathbb{R}^d; \mathbb{R}^D)} \left(\mathcal{J}(f) := \sum_{i=1}^N \mathcal{L}(\mathbf{y}_i, f(\mathbf{x}_i)) + \lambda \|f\|_{\mathcal{V}_\sigma(\mathbb{R}^d; \mathbb{R}^D)} \right), \quad \lambda > 0, \quad (34)$$

where $(\mathbf{x}_1, \mathbf{y}_1), \dots, (\mathbf{x}_N, \mathbf{y}_N) \in \mathbb{R}^d \times \mathbb{R}^D$ is any fixed dataset and $\mathcal{L}(\cdot, \cdot)$ is lower semicontinuous in its second argument, favors solutions that share neurons. We quantify this explicitly in Theorem 9.

Theorem 9 *Let f be a finite-width vector-valued neural network with unique input weights of the form*

$$f(\mathbf{x}) = \sum_{k=1}^K \mathbf{v}_k \sigma(\mathbf{w}_k^T \bar{\mathbf{x}}), \quad \mathbf{x} \in \mathbb{R}^d, \quad (35)$$

with $\|\mathbf{w}_k\|_2 = 1$, $k = 1, \dots, K$. Then, there exists $\delta > 0$ (that depends on λ and the data) such that, if $\|\mathbf{w}_1 - \mathbf{w}_2\|_2 < \delta$ and these two neurons contribute exclusively to two disjoint subsets of the outputs, then the neural network

$$\hat{f}(\mathbf{x}) = f(\mathbf{x}) - \mathbf{v}_1 \sigma(\mathbf{w}_1^T \bar{\mathbf{x}}) + \mathbf{v}_1 \sigma(\mathbf{w}_2^T \bar{\mathbf{x}}), \quad (36)$$

which shares one neuron across both sets of outputs has a strictly smaller objective value, i.e., $\mathcal{J}(\hat{f}) < \mathcal{J}(f)$.

Proof Since the input weights are all unit norm, Equation (34) reduces to

$$\mathcal{J}(f) = \sum_{i=1}^N \mathcal{L}(\mathbf{y}_i, f(\mathbf{x}_i)) + \lambda \sum_{k=1}^K \|\mathbf{v}_k\|_2. \quad (37)$$

Without loss of generality, suppose that the two neurons whose input weights are \mathbf{w}_1 and \mathbf{w}_2 contribute to two disjoint subsets of the D outputs. In particular, the first neuron contributes to the outputs in index set $\mathcal{I}_1 \subset \{1, \dots, D\}$, while the second contributes to outputs in the index set $\mathcal{I}_2 \subset \{1, \dots, D\}$, where $\mathcal{I}_1 \cap \mathcal{I}_2 = \emptyset$. That is, \mathbf{v}_1 and \mathbf{v}_2 have disjoint support. Define

$$\epsilon := \lambda(\|\mathbf{v}_1\|_2 + \|\mathbf{v}_2\|_2 - \|\mathbf{v}_1 + \mathbf{v}_2\|_2) > 0. \quad (38)$$

Since the loss \mathcal{L} is lower semicontinuous in its second argument, there exists $\gamma_i > 0$, such that, if $\|\hat{f}(\mathbf{x}_i) - f(\mathbf{x}_i)\|_2 < \gamma_i$, then

$$\mathcal{L}(\mathbf{y}_i, \hat{f}(\mathbf{x}_i)) - \mathcal{L}(\mathbf{y}_i, f(\mathbf{x}_i)) < \epsilon/N, \quad (39)$$

where we note that γ_i depends on ϵ , N , and \mathbf{x}_i . By (36), we have that

$$\begin{aligned} \|\hat{f}(\mathbf{x}_i) - f(\mathbf{x}_i)\|_2 &= \|\mathbf{v}_1 \sigma(\mathbf{w}_1^T \bar{\mathbf{x}}_i) - \mathbf{v}_1 \sigma(\mathbf{w}_2^T \bar{\mathbf{x}}_i)\|_2 \\ &= \|\mathbf{v}_1\|_2 |\sigma(\mathbf{w}_1^T \bar{\mathbf{x}}_i) - \sigma(\mathbf{w}_2^T \bar{\mathbf{x}}_i)|. \end{aligned} \quad (40)$$

The continuity of the activation function guarantees that there exists a $\delta_i > 0$ (that depends on $\gamma_i/\|\mathbf{v}_1\|_2$) such that

$$\|\mathbf{w}_1 - \mathbf{w}_2\|_2 < \delta_i \implies |\sigma(\mathbf{w}_1^T \bar{\mathbf{x}}_i) - \sigma(\mathbf{w}_2^T \bar{\mathbf{x}}_i)| < \gamma_i/\|\mathbf{v}_1\|_2. \quad (41)$$

Therefore, if $\|\mathbf{w}_1 - \mathbf{w}_2\|_2 < \delta := \min_{i=1, \dots, N} \delta_i$, then $|\sigma(\mathbf{w}_1^T \bar{\mathbf{x}}_i) - \sigma(\mathbf{w}_2^T \bar{\mathbf{x}}_i)| < \gamma_i/\|\mathbf{v}_1\|_2$, for any $i = 1, \dots, N$. Consequently, this implies that

$$\left(\sum_{i=1}^N \mathcal{L}(\mathbf{y}_i, \hat{f}(\mathbf{x}_i)) - \mathcal{L}(\mathbf{y}_i, f(\mathbf{x}_i)) \right) < \epsilon \quad (42)$$

whenever $\|\mathbf{w}_1 - \mathbf{w}_2\|_2 < \delta$. To complete the proof, observe that whenever $\|\mathbf{w}_1 - \mathbf{w}_2\|_2 < \delta$, we have that

$$\begin{aligned} \mathcal{J}(\hat{f}) - \mathcal{J}(f) &= \left(\sum_{i=1}^N \mathcal{L}(\mathbf{y}_i, \hat{f}(\mathbf{x}_i)) - \mathcal{L}(\mathbf{y}_i, f(\mathbf{x}_i)) \right) + \lambda \left(\|\hat{f}\|_{\mathcal{V}_\sigma(\mathbb{R}^d; \mathbb{R}^D)} - \|f\|_{\mathcal{V}_\sigma(\mathbb{R}^d; \mathbb{R}^D)} \right) \\ &= \left(\sum_{i=1}^N \mathcal{L}(\mathbf{y}_i, \hat{f}(\mathbf{x}_i)) - \mathcal{L}(\mathbf{y}_i, f(\mathbf{x}_i)) \right) + \lambda (\|\mathbf{v}_1 + \mathbf{v}_2\|_2 - (\|\mathbf{v}_1\|_2 + \|\mathbf{v}_2\|_2)) \\ &< \epsilon + \lambda (\|\mathbf{v}_1 + \mathbf{v}_2\|_2 - (\|\mathbf{v}_1\|_2 + \|\mathbf{v}_2\|_2)) \\ &= 0. \end{aligned} \quad (43)$$

Thus, if two neurons have sufficiently close input weights, removing one and having the other one be shared strictly decreases the objective in (34). \blacksquare

Note that by sharing neurons \hat{f} always has a strictly smaller $\mathcal{V}_\sigma(\mathbb{R}^d; \mathbb{R}^D)$ -norm than f . Indeed, we have that

$$\|\hat{f}\|_{\mathcal{V}_\sigma(\mathbb{R}^d; \mathbb{R}^D)} - \|f\|_{\mathcal{V}_\sigma(\mathbb{R}^d; \mathbb{R}^D)} = \|\mathbf{v}_1 + \mathbf{v}_2\|_2 - (\|\mathbf{v}_1\|_2 + \|\mathbf{v}_2\|_2) < 0, \quad (44)$$

where the equality follows from the fact that the only neurons that are different between the two functions are the ones with input weights \mathbf{w}_1 and \mathbf{w}_2 and the inequality follows from the triangle inequality (which is strict since \mathbf{v}_1 and \mathbf{v}_2 have disjoint support).

Theorem 9 along with the discussion in Section 2 shows that, when the activation function σ is homogeneous, vector-valued neural networks trained with weight decay are encouraged to share neurons. Trained networks that exhibit neuron sharing are important in multi-task learning problems, e.g., multi-class classification, where components of the labels could have some relationships or correlations. Therefore, the neuron sharing phenomenon exhibited by solutions to weight decay regularized problems provides some explanation towards its efficacy when used for classification tasks. We illustrate the types of architectures that are favored by weight decay regularization in Figure 1. We also verify this numerically in Section 7.1.

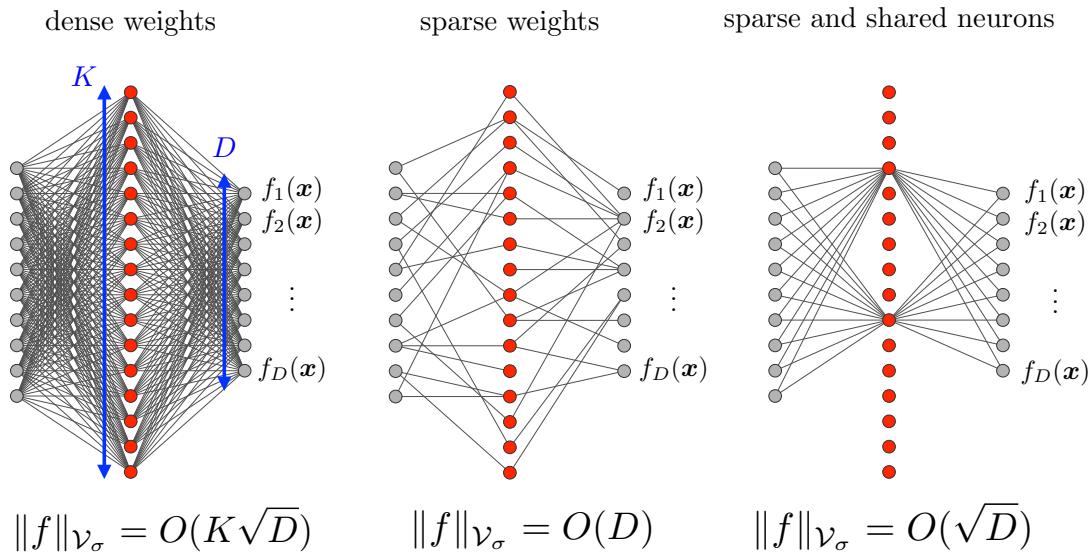


Figure 1: Three neural networks with different weight-sparsity patterns. The input weights are normalized to lie on the sphere and the components of the output weights are all $O(1)$. In the case of homogeneous activation functions, weight decay minimizes the $\mathcal{V}_\sigma(\mathbb{R}^d; \mathbb{R}^D)$ norm and therefore favors the right-most architecture. This architecture exhibits both neuron sparsity and neuron sharing. Each output depends on the same few neurons. This observation also gives insight into the regularity of the optimal functions: They favor functions that only vary in a few directions across all outputs. This is in contrast with the middle network where each output has variation in a small number of directions, but this set of directions can be different for each output.

6. Data-Dependent Width Bounds and DNN Compression

In this section, we complement our previous analysis of the vector-valued variation spaces. We use the NBT (Theorem 1) to refine the number of neurons predicted by the representer theorems (Theorems 5 and 8). In particular, these new bounds are *data dependent* and improve previous results that have been reported in the literature (e.g., Jacot et al., 2022, Proposition 7). Furthermore, these bounds are applicable to all layers with homogeneous activation functions within *any* DNN.

We use the NBT, to relate the weights that minimize the weight decay objective with minimizers of a *convex* multi-task lasso problem. This convex problem is data dependent and is based on the learned representations of the training data. We use this reduction to derive bounds on the widths of DNNs trained with weight decay (Theorem 10). A by-product of our analysis are, to the best of our knowledge, the first sparsity bounds for the (convex) multitask lasso problem (Theorem 11). This result may be of independent interest. This investigation also motivates the proposal of a new principled and computationally efficient procedure to compress pre-trained DNNs (Section 7.3).

6.1 Width Bounds

Let f_{θ} be a DNN that minimizes (1). Consider any layer of this DNN that has homogeneous activation functions and suppose that it has a width of K neurons. Let φ_i denote the learned feature representation of the i th training example \mathbf{x}_i that is input to this layer. Furthermore, let $\{\mathbf{w}_k\}_{k=1}^K$ and $\{\mathbf{v}_k\}_{k=1}^K$ denote the input and output weights of this layer, respectively. Then, the *output features* learned by this layer are given by

$$\psi_i := \sum_{k=1}^K \mathbf{v}_k \sigma(\mathbf{w}_k^T \overline{\varphi_i}), \quad i = 1, \dots, N, \quad (45)$$

where we recall that $\overline{\varphi_i}$ augments a 1 to φ_i to account for a bias term. Thanks to the homogeneity of σ , we can write

$$\psi_i := \sum_{k=1}^K \tilde{\mathbf{v}}_k \sigma(\tilde{\mathbf{w}}_k^T \overline{\varphi_i}), \quad i = 1, \dots, N, \quad (46)$$

where $\tilde{\mathbf{w}}_k := \mathbf{w}_k / \|\mathbf{w}_k\|_2$ and $\tilde{\mathbf{v}}_k := \mathbf{v}_k / \|\mathbf{w}_k\|_2$. Let $\phi_{k,i} := \sigma(\tilde{\mathbf{w}}_k^T \overline{\varphi_i})$. Consequently, let

$$\phi_i = (\phi_{1,i}, \dots, \phi_{K,i}) \in \mathbb{R}^K, \quad i = 1, \dots, N, \quad (47)$$

denote the *post-activation* features.

A corollary to the representer theorem (Theorem 5) is that there exists an optimal representation of this layer with $\leq N^2$ neurons. This is often a loose upper bound, especially when the learned features are highly structured. In Theorem 10 we present *data-dependent* bounds on the widths of DNN layers based on the intrinsic dimensions (ranks) of the learned data representations.

Theorem 10 *Let $(\mathbf{x}_1, \mathbf{y}_1), \dots, (\mathbf{x}_N, \mathbf{y}_N)$ be any finite dataset and let $\mathcal{L}(\cdot, \cdot)$ be any loss function that is lower semicontinuous in its second argument. Furthermore, let f_{θ} be a*

DNN that solves

$$\min_{\boldsymbol{\theta}} \sum_{i=1}^N \mathcal{L}(\mathbf{y}_i, f_{\boldsymbol{\theta}}(\mathbf{x}_i)) + \frac{\lambda}{2} \|\boldsymbol{\theta}\|_2^2. \quad (48)$$

Consider any layer of $f_{\boldsymbol{\theta}}$ with homogeneous activation functions and let $\{\phi_i\}_{i=1}^N$ and $\{\psi_i\}_{i=1}^N$ denote the learned post-activation and output features defined according to (47) and (46), respectively. If

$$\begin{aligned} r_{\Phi} &= \dim \text{span}\{\phi_i\}_{i=1}^N \\ r_{\Psi} &= \dim \text{span}\{\psi_i\}_{i=1}^N \end{aligned} \quad (49)$$

denotes the dimensions of the subspaces spanned by the learned features, then there exists an equivalent representation of the DNN (in the sense that it still minimizes (48)) where this layer has at most $r_{\Phi} r_{\Psi} \leq N^2$ neurons. Furthermore, the equivalent representation can be found by solving a convex multi-task lasso problem.

Proof Given the learned post-activation and output features $\{\phi_i\}_{i=1}^N$ and $\{\psi_i\}_{i=1}^N$ of the training examples $\{\mathbf{x}_i\}_{i=1}^N$ from one layer of the DNN $f_{\boldsymbol{\theta}}$ with homogeneous activation functions, observe that, by the NBT (Theorem 1), $\{\tilde{\mathbf{v}}_k\}_{k=1}^K$ (as defined in (46)) must minimize

$$\min_{\{\mathbf{v}_k\}_{k=1}^K} \sum_{k=1}^K \|\mathbf{v}_k\|_2 \quad \text{s.t.} \quad \psi_i = \sum_{k=1}^K \mathbf{v}_k \phi_{k,i}, \quad i = 1, \dots, N. \quad (50)$$

Therefore, if we replace $\{\tilde{\mathbf{v}}_k\}_{k=1}^K$ with any solution to (50) and rebalance the weights, the new DNN would still minimize (48).

To complete the proof, we observe that (50) is a convex multi-task lasso problem. As we shall prove in Theorem 11 in Section 6.2 below, there exists a solution to that multi-task lasso problem with at most $r_{\Phi} r_{\Psi}$ nonzero vectors. By a dimension-counting argument we always have the bound $r_{\Phi} r_{\Psi} \leq N^2$. Finally, observe that we can always find the compressed representation by solving the convex multi-task lasso problem. Thanks to the NBT, upon rebalancing the weights into and out of this layer, the compressed DNN still minimizes (48). ■

Note that this result is not restricted to only one layer in a DNN. Indeed, it can be applied in a layer-wise manner. Thus, this result provides bounds on the widths of any layer in a DNN with homogeneous activation functions. This result also improved previous results that have appeared in the literature since it is data dependent. Furthermore, our generic bound of N^2 also improves the recent bound of $N(N + 1)$ of Jacot et al. (2022, Proposition 7).

Our data-dependent bounds are particularly relevant since it has been observed empirically that the subspaces spanned by learned features are often low-dimensional (Nar et al., 2019; Waleffe and Rekatsinas, 2022; Feng et al., 2022; Huh et al., 2023; Kwon et al., 2024; Yaras et al., 2024). These empirical observations have also been backed by theoretical arguments (Papyan et al., 2020; Le and Jegelka, 2022). Therefore, we see that Theorem 10 motivates the design of a principled and computationally efficient procedure to compress pre-trained DNNs. This procedure proceeds in a layer-by-layer manner and solves the convex multi-task lasso problem (50). We empirically evaluate the performance of this compression procedure in Section 7.3 on various architectures.

6.2 Sparsity of Solutions to the Multi-Task Lasso Problem

The main ingredients of the proof of Theorem 10 were (i) a reduction to the convex multi-task lasso problem (ii) the invocation of the sparsity bounds on the multi-task lasso problem, which we derive in this section. We consider the multi-task lasso problem as formulated by Obozinski et al. (2006, 2010); Argyriou et al. (2008); Obozinski et al. (2011). This problem has also been studied in the context of inverse problems with multiple measurement vectors (Cotter et al., 2005; Chen and Huo, 2006; Sun et al., 2009; Bajwa et al., 2015). Our result on the sparsity of multi-task lasso minimizers appears in Theorem 11. This result is new, to the best of our knowledge, and may be of independent interest. The proof can be found in Appendix E.

Theorem 11 *Consider the multi-task lasso problem*

$$\min_{\mathbf{V}=[\mathbf{v}_1, \dots, \mathbf{v}_K]} \sum_{k=1}^K \|\mathbf{v}_k\|_2 \quad \text{s.t.} \quad \mathbf{\Psi} = \mathbf{V}\mathbf{\Phi}, \quad (51)$$

where $\mathbf{\Psi} \in \mathbb{R}^{D \times N}$ and $\mathbf{\Phi} \in \mathbb{R}^{K \times N}$ are matrices with ranks r_{Φ} and r_{Ψ} , respectively. Assume that the row space of $\mathbf{\Psi}$ is contained in the row space of $\mathbf{\Phi}$, which implies a solution exists. For any $K > 0$ there exists a solution with at most $r_{\Phi}r_{\Psi}$ nonzero columns. Furthermore, there exists $\mathbf{\Psi}$ such that no solution has fewer than r_{Φ} nonzero columns.

Unlike the traditional lasso problem, which seeks a sparse solution in an unstructured manner, the multi-task lasso seeks a structured sparsity. In this setting, each column of \mathbf{V} is associated with a block. This block will either be entirely zero or (typically) entirely nonzero. Since each column of \mathbf{V} corresponds to a feature, this implies that each feature is used in either all prediction tasks or none. The proof utilizes a generalized version of Carathéodory's theorem along with the fact that for any solution \mathbf{V} , the dimension of the column space must be r_{Ψ} .

The intuition behind the upper and lower bounds in Theorem 11 is as follows. If every row of $\mathbf{\Psi}$ can be synthesized using the same r_{Φ} rows in $\mathbf{\Phi}$, while still minimizing the objective, then the lower bound is achieved. On the other hand, if the minimum can be achieved by having r_{Ψ} linearly independent rows in $\mathbf{\Psi}$ synthesized by different subsets of r_{Φ} rows in $\mathbf{\Phi}$, then the upper bound is met. Our numerical experiments in Section 7.2 demonstrate that the minimum number of nonzero columns in any solution may range between the upper and lower bounds depending on precise structure of $\mathbf{\Psi}$ and $\mathbf{\Phi}$.

Remark 12 *The lower bound follows from the following observation. If a solution exists with fewer than r_{Φ} nonzero columns, then all the rows of $\mathbf{\Psi}$ belong one of the (finitely many) lower dimensional subspaces in the row space of $\mathbf{\Phi}$ spanned by certain subsets of its rows. But collectively, these subspaces do not contain all points in the row space of $\mathbf{\Phi}$, and so there are infinitely many points that cannot be represented with fewer than r_{Φ} rows of $\mathbf{\Phi}$.*

Remark 13 *Observe that when $D = 1$, (51) reduces to the classical lasso problem. In that case, the sparsity of solutions to the classical lasso has been investigated by Rosset et al. (2004); Tibshirani (2013).*

7. Experiments

In this section we present three numerical experiments that validate our theory and demonstrate its utility in practice. Our first experiment demonstrates that weight decay encourages neuron sharing. Our second experiment validates the bound presented for the multi-task lasso problem in Theorem 11. Our third experiment compresses layers for pre-trained VGG-19 (Simonyan and Zisserman, 2014; Wang et al., 2021) and AlexNet (Krizhevsky et al., 2017) models via the multi-task lasso convex optimization problem as in Theorem 10. In particular, we show that this principled compression approach preserves the training loss, accuracy, and weight decay objective of the model.³

7.1 Neuron Sharing Simulation

To demonstrate that weight decay encourages neuron sharing we train three vector-valued shallow ReLU neural networks to fit a synthetic two-dimensional dataset. The dataset consists of 50 samples where the features are two-dimensional vectors drawn i.i.d. from a multivariate normal distribution. The labels are three-dimensional and generated by passing the feature vectors through a randomly initialized ReLU neural network with five neurons.

All networks were initialized with one hundred and fifty neurons. The first network was trained with weight decay regularization. The second network was trained with ℓ^1 regularization. The third network was trained with no regularization. We used full batch gradient descent with the Adam optimizer. We used a learning rate of 2×10^{-3} for two million iterations. The regularization parameter was $\lambda = 10^{-6}$ for weight decay and $\lambda = 10^{-9}$ for ℓ^1 -regularization. We chose λ to be as large as possible such that the networks interpolate the data.

In Figure 2 we plot the locations of the neurons contributing to each output of the trained network. Thanks to the homogeneity of the ReLU, after training, we normalize the input weights to be unit norm by absorbing the magnitude into the output weight. This reparameterization does not change the overall function mapping. This allows us to express each input weight as $\mathbf{w}_k = (\cos \theta_k, \sin \theta_k)$ with a bias $b_k \in \mathbb{R}$. In other words, we can plot each neuron $\mathbf{x} \mapsto \sigma(\mathbf{w}_k^T \mathbf{x} + b_k)$ with the two-dimensional coordinate (θ_k, b_k) .

In Figure 2, we show the (θ_k, b_k) pairs for every active neuron in the trained network. Our results show that, when regularizing with weight decay, the learned network is not only sparser (in terms of the number of active neurons), but also exhibits strong neuron sharing. In contrast with ℓ^1 -regularization there is not much neuron sparsity or neuron sharing. Finally, we see that no regularization results in a very dense network where all neurons are active. To generate the plots, we deem a neuron *active* if the ℓ^1 -norm of the output weight is greater than 10^{-3} . At the end of training we had 5 neurons active for weight decay regularization, 85 neurons active for ℓ^1 -regularization and 130 neurons active with no regularization.

3. The code to reproduce our experiments can be found at <https://github.com/joeshenouda/vv-spaces-nn-width>. We follow the reproducibility guidance of Shenouda and Bajwa (2023).

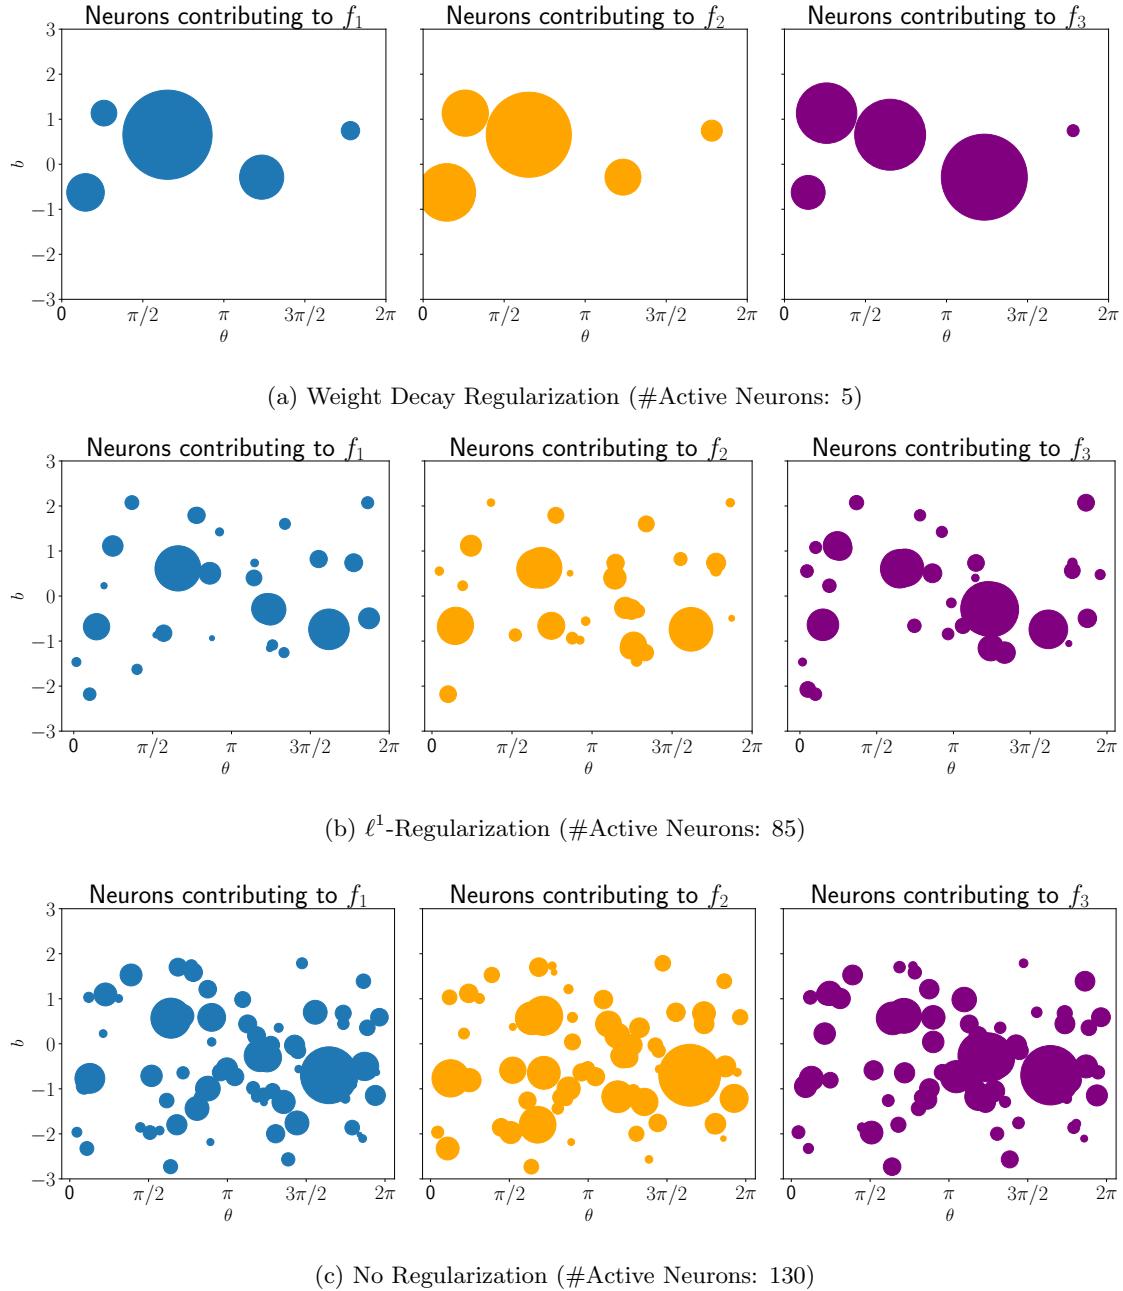


Figure 2: We trained a three output two-dimensional ReLU neural network of the form $\mathbf{f}(x) = \sum_{k=1}^K \mathbf{v}_k \sigma(\mathbf{w}_k^T \mathbf{x} + b_k)$ with weight decay, ℓ^1 -regularization, and no regularization. Let f_1 , f_2 and f_3 denote the first, second, and third components of the outputs. We plot the locations of each active neurons under the (θ_k, b_k) -parameterization. The size of the circles indicate the magnitude of the corresponding output weight vector. We see that in the case of weight decay, we have very few active neurons. Furthermore, those neurons that remain are shared across all outputs.

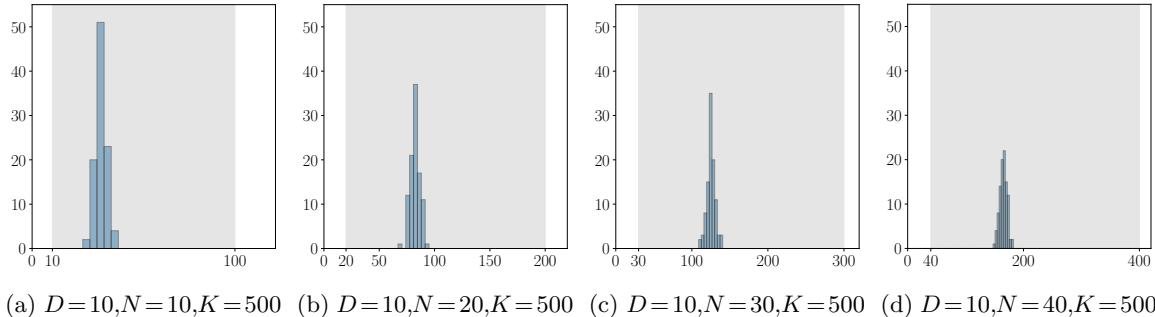
(a) $D=10, N=10, K=500$ (b) $D=10, N=20, K=500$ (c) $D=10, N=30, K=500$ (d) $D=10, N=40, K=500$

Figure 3: Distribution of the number of active columns for the solutions to the multi-task lasso problem on randomly generated matrices of varying sizes. The horizontal axis is the number of nonzero columns in the optimal \mathbf{V} and the vertical axis is the frequency. We ran this experiment for 100 randomly generated matrices. In all cases $r_{\Phi} = N$ and $r_{\Psi} = D$ so by Theorem 11 we expect $N \leq \hat{K} \leq ND$. The shaded region indicates our theoretical bounds. The wide gap suggests that our upper bound can be sharpened.

7.2 Multi-Task Lasso Experiments

We solve the multi-task lasso problem in Theorem 11 on randomly generated matrices Φ and Ψ using CVXPY (Diamond and Boyd, 2016). Our experiments illustrate that, while our bounds hold, the exact number of nonzero columns depends on the data itself. In Figure 3 we show histograms for the distribution of nonzero columns over 100 randomly generated pairs of Φ and Ψ .

In Figure 4 we perform a similar set of experiments but alter the underlying rank of Φ . This validates our bound showing that the sparsity of the solution can be much lower depending on the rank. We again see that the distribution is dependent on the data. We note that, however, we never achieve our upper bound which may indicate that our bounds can be further sharpened.

To demonstrate that the sparsest solutions to the multi-task lasso problem Theorem 11 depends on the data matrices Φ and Ψ we also ran a small-scale experiment similar to Figure 4 and Figure 3. However, in the next set of experiments, we exhaustively searched over all 2^K sparsity patterns that the solution may have. We arrive at the same conclusion as we did before: The sparsest solution can lie anywhere within our bound. This is depicted in Figure 5.

7.3 Compression of Pre-Trained DNNs

We compress DNNs pre-trained on the CIFAR-10 dataset based on the principled approach outlined in Section 6. The first model we consider is a pre-trained VGG-19 architecture trained with weight decay.⁴ This model consists of various convolutional and batch-norm

4. We use the pre-trained VGG-19 model from <https://github.com/hwang595/Pufferfish>.

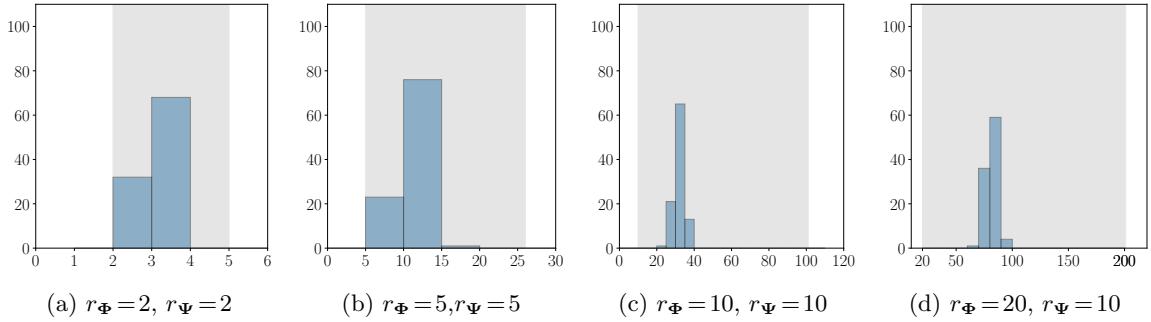


Figure 4: Distribution of the number of active columns for the solutions to the multi-task lasso problem on randomly generated matrices with Φ of various rank. The horizontal axis is the number of nonzero columns in the optimal \mathbf{V} and the vertical axis is the frequency. We ran this experiment for 100 randomly generated matrices, in all cases $D = 10, N = 20$ and $K = 200$. By Theorem 11 we expect $r_{\Phi} \leq \hat{K} \leq r_{\Phi} \cdot r_{\Psi}$. The shaded region indicates our theoretical bounds.

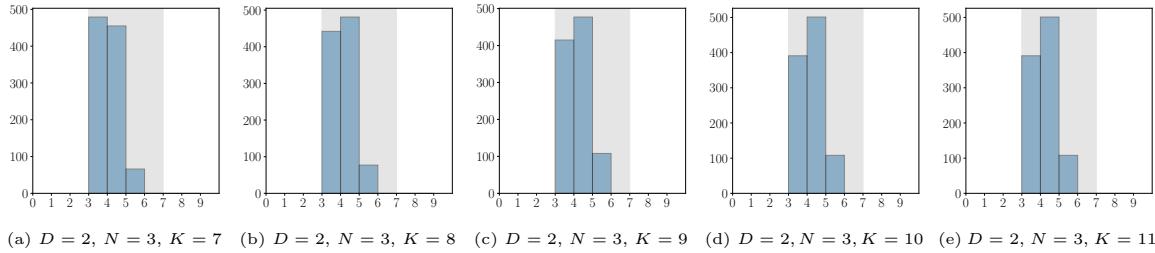


Figure 5: Distribution of number of nonzero columns for the solutions to the multi-task lasso problem. We ran this experiment over 1000 randomly generated matrices Φ and Ψ . The horizontal axis is the number of nonzero columns in the optimal \mathbf{V} where each bin is left-inclusive corresponding to a single integer. The vertical axis is the frequency. We see that the sparsest solution is dependent on the data and can vary between our bounds. The area between the two shaded regions indicate our theoretical bounds. The gap indicates that our upper bound can be sharpened.

layers followed by a fully-connected ReLU layer that contains 512 neurons. We run our compression procedure on this ReLU layer.

The output of this ReLU layer is the output of the entire network. There are 10 outputs that correspond to the 10 classes of CIFAR-10. Therefore, $\Psi \in \mathbb{R}^{10 \times N}$ with $r_{\Psi} \leq 10$. The post-activation feature matrix $\Phi \in \mathbb{R}^{512 \times N}$ has an approximate rank of 10. We approximate the rank following the same procedure of Huh et al. (2023, Appendix D): We threshold the singular values at a value of 10^{-3} .

Thus, Theorem 10 suggests that there exists an alternative optimal representation with no more than 100 neurons in this layer that can be found by solving a convex multi-task

lasso problem. We validate this by minimizing⁵ a regularized version of the constrained multi-task lasso problem

$$\min_{\{\mathbf{v}_k\}} \frac{1}{ND} \|\mathbf{V}\Phi - \Psi\|_2^2 + \lambda \sum_{k=1}^K \|\mathbf{v}_k\|_2, \quad (52)$$

with $\lambda = 3 \times 10^{-3}$.

The second model we consider is a pre-trained AlexNet architecture trained with weight decay.⁶ We use this mode to illustrate how we can compress multiple layers in a DNN in a layer-by-layer fashion. This model has multiple convolutional layers followed by three fully connected ReLU layers with 9216, 4096, and 4096 neurons, respectively. We run our compression procedure on these ReLU layers.

For the first fully-connected layer we found that $r_\Psi \approx 69$ and $r_\Phi \approx 225$. For the second fully-connected layer we found that $r_\Psi \approx 56$ and $r_\Phi \approx 69$. For the last fully-connected layer we found that $r_\Psi \leq 10$ and $r_\Phi \approx 56$. The approximate ranks were again computed with a threshold of 10^{-3} on the singular values. Theorem 10 suggests that we can significantly compress the last two layers. For the first and second fully-connected layers we solve (52) with $\lambda = 10^{-6}$. For the last fully-connected layer we used $\lambda = 10^{-3}$.

The results for the compression of both VGG-19 and AlexNet can be found in Table 1. Observe that there is almost no change in the training loss or the sum of squared weights of the model. The minor discrepancies between the original models and the compressed models are due to the fact that we solve the regularized problem (52).

Table 1: Compression results for various layers of VGG-19 and AlexNet models pre-trained on the CIFAR-10 dataset. FC stands for fully connected.

	Width		Train Loss (Test Acc.)		$\ \theta\ _2^2$	
	Original	New	Original	New	Original	New
VGG-19 (last layer)	512	99	0.0002 (93.91%)	0.0002 (93.90%)	3823.41	3816.70
AlexNet (last layer)	4096	555	0.2203 (72.17%)	0.2201 (72.14%)	1141.50	961.46
AlexNet (penultimate layer)	4096	2369	0.2203 (72.17%)	0.2158 (72.16%)	1141.50	1137.65
AlexNet (first FC layer)	9216	8208	0.2203 (72.17%)	0.2190 (72.18%)	1141.50	1137.33
AlexNet (last 3 layers)	(4096,4096,9216)	(555,2369,8208)	0.2203 (72.17%)	0.21521(72.13%)	1141.50	956.5

8. Related Work

In this section, we highlight some related works and discuss how our work fits into the current literature.

5. We minimize the objective with proximal gradient methods.
6. We trained this model ourselves based on the implementation in <https://github.com/rasbt/deeplearning-models>.

8.1 Variation Spaces and Representer Theorems for Neural Networks

Understanding neural networks by studying their associated variation space has been explored in the past (Barron, 1993; Kůrková and Sanguineti, 2001; Mhaskar, 2004; Bach, 2017). Representer theorems that show that finite-width shallow networks are solutions to data-fitting problems posed over these spaces have also been developed (Parhi and Nowak, 2021). These results have been extended to a wide variety of activation functions (Parhi and Nowak, 2020), more general settings (Korolev, 2022; Bartolucci et al., 2023; Spek et al., 2022; Wang et al., 2024b), as well as deep networks (Parhi and Nowak, 2022; Bartolucci et al., 2024; Wang et al., 2024a). The variation norm has also shown practical benefits in learning sparse neural networks (Yang et al., 2022) and implicit neural representations (Shenouda et al., 2024).

It is important to note that these prior works primarily investigate scalar-output neural networks or simple Cartesian product constructions to account for multiple outputs. As we saw in Section 3.2 (see also Appendix A), the vector-valued setting brings unique difficulties in the choice of the norm associated with the variation space, and many choices do not always coincide with regularizers used in practice such as weight decay. While we note that there has been work devoted to the investigation of vector-valued RKBSSs for multi-task learning (see Zhang and Zhang (2013) and references therein), we are the first to study and propose the $\mathcal{V}_\sigma(\mathbb{R}^d; \mathbb{R}^D)$ space and norm. Furthermore, our new representer theorem (Theorem 5) improves many results in the representer theorem literature since the number of neurons does not depend on the output dimension (see Bredies et al. (2024) for more details on the output-dimension dependence).

8.2 Weight Decay and Sparsity

Numerous works have observed that weight decay is biased towards solutions with fewer neurons (Savarese et al., 2019; Parhi and Nowak, 2021; Yang et al., 2022). Empirically, Yang et al. (2022) have shown that training DNNs with weight decay for many iterations induces neuron-wise sparsity on the trained network. Furthermore, Parhi and Nowak (2022); Jacot et al. (2022) provide bounds on the widths for DNNs that solve the weight decay objective. Our bounds improve upon the ones presented in both of these works. In particular, we develop the first *data-dependent* bounds. Moreover, our setting is also more general: Our theory holds for *any* DNN architecture that has homogeneous activation functions.

Recently, the works of Ergen and Pilancı (2021); Mishkin et al. (2022) have shown that training shallow ReLU neural networks with weight decay can be recast as a (very large) convex program. Their reformulation reveals that weight decay induces a sparsity-promoting regularizer. A similar observation was developed for the vector-valued case by Sahiner et al. (2021).

8.3 Low-Rank Features

It has been observed that DNNs are biased toward learning low-rank features. Theoretical evidence for this phenomenon has been developed by Du et al. (2018); Ji and Telgarsky (2019); Radhakrishnan et al. (2021); Le and Jegelka (2022); Mousavi-Hosseini et al. (2023). This has also been observed empirically in Nar et al. (2019); Waleffe and Rekatsinas (2022);

Feng et al. (2022); Huh et al. (2023); Kwon et al. (2024); Yaras et al. (2024). We highlight that some works such as Waleffe and Rekatsinas (2022); Kwon et al. (2024) use this insight to develop specialized algorithms that compress models during training. In contrast, our theory and experiments indicate that narrower networks can be found regardless of the training procedure by solving a simple convex multi-task lasso. Moreover, Theorem 10 guarantees that the compressed network has the same sum of squared weights and minimizes the same objective as the original network.

9. Conclusion

In this work we proposed the $\mathcal{V}_\sigma(\mathbb{R}^d; \mathbb{R}^D)$ function space which gives insights in the inductive bias of vector-valued neural networks. This is a critical step towards understanding DNNs, the theory of which significantly lags behind that of shallow neural networks. We proved a new representer theorem for this space showing that finite-width vector-valued neural networks are solutions to the data-fitting problem over this space. For DNNs with homogeneous activation functions we developed a novel connection between training DNNs with weight decay and multi-task lasso to prove the sharpest known bounds on network widths. This result motivated the design of a principled and computationally efficient procedure to compress pre-trained DNNs. Finally, we presented some experimental results showing the validity and practical utility of our theory.

Acknowledgments

The authors thank Ryan Tibshirani for helpful discussions regarding the lasso and related methods, especially the use of Carathéodory’s theorem for characterizing the sparsity of their solutions. JS was supported by the ONR MURI grant N00014-20-1-2787. RP was supported by the National Science Foundation (NSF) Graduate Research Fellowship Program under grant DGE-1747503 while he was with the University of Wisconsin–Madison and the European Research Council under grant 101020573 while he was with the École polytechnique fédérale de Lausanne. KL was supported by the NSF grant DMS-2023239. RN was supported in part by the NSF grants DMS-2134140 and DMS-2023239, the ONR MURI grant N00014-20-1-2787, and the AFOSR/AFRL grant FA9550-18-1-0166.

Appendix A. Different Norms on Vector-Valued Measures

In this appendix, we will discuss different choices of norms one can equip the space of vector-valued measures $\mathcal{M}(\Omega; \mathbb{R}^D)$, where Ω is any locally compact Hausdorff space. Eventually, we will show in Lemma 14 that these choices are equivalent Banach norms, though only one choice (the one proposed in Section 3.2) corresponds to weight-decay regularization in neural networks with homogeneous activation functions.

By viewing $\mathcal{M}(\Omega; \mathbb{R}^D)$ as the D -fold Cartesian product $\times_{j=1}^D \mathcal{M}(\Omega)$, a naïve choice of norm would be the mixed norm

$$\|\boldsymbol{\nu}\|_{\mathcal{M},p} := \left(\sum_{j=1}^D \|\nu_j\|_{\mathcal{M}(\Omega)}^p \right)^{1/p} = \left(\sum_{j=1}^D \left(\sup_{\substack{\Omega = \bigcup_{i=1}^n A_i \\ n \in \mathbb{N}}} \sum_{i=1}^n |\nu_j(A_i)| \right)^p \right)^{1/p}, \quad (53)$$

with $p \geq 1$ and $\boldsymbol{\nu} = (\nu_1, \dots, \nu_D)$ where each $\nu_j \in \mathcal{M}(\Omega)$, $j = 1, \dots, D$. In this scenario, it is clear that $\mathcal{M}(\Omega; \mathbb{R}^D)$ is a Banach space when equipped with the norm $\|\cdot\|_{\mathcal{M},p}$, $p \geq 1$. Alternatively, one could consider the norm

$$\|\boldsymbol{\nu}\|_{p,\mathcal{M}} := \sup_{\substack{\Omega = \bigcup_{i=1}^n A_i \\ n \in \mathbb{N}}} \sum_{i=1}^n \|\boldsymbol{\nu}(A_i)\|_p = \sup_{\substack{\Omega = \bigcup_{i=1}^n A_i \\ n \in \mathbb{N}}} \sum_{i=1}^n \left(\sum_{j=1}^D |\nu_j(A_i)|^p \right)^{1/p}, \quad (54)$$

with $p \geq 1$. We have that $(\mathcal{M}(\Omega; \mathbb{R}^D), \|\cdot\|_{p,\mathcal{M}})$ is also Banach space (cf. Diestel and Uhl, 1977, pg. 29). When $\Omega = \mathbb{S}^d$ and $p = 2$, this coincides with (14). The next lemma shows that all of these norms are actually equivalent Banach norms.

Lemma 14 *Let Ω be any locally compact Hausdorff space. The norms $\|\cdot\|_{p,\mathcal{M}}$, $p \geq 1$, and $\|\cdot\|_{\mathcal{M},q}$, $q \geq 1$ are all equivalent Banach norms for $\mathcal{M}(\Omega; \mathbb{R}^D)$.*

Proof We first observe that for any $p, q \geq 1$, the norms $\|\cdot\|_{\mathcal{M},p}$ and $\|\cdot\|_{\mathcal{M},q}$ are equivalent by the equivalence of p -norms in finite dimensions. Similarly, for any $p, q \geq 1$, the norms $\|\cdot\|_{p,\mathcal{M}}$ and $\|\cdot\|_{q,\mathcal{M}}$ are equivalent. Thus, it suffices to show that $\|\cdot\|_{\mathcal{M},p}$ and $\|\cdot\|_{q,\mathcal{M}}$ are equivalent for some $p, q \geq 1$. We have for any $\boldsymbol{\nu} \in \mathcal{M}(\Omega; \mathbb{R}^D)$

$$\begin{aligned} \|\boldsymbol{\nu}\|_{\mathcal{M},1} &= \sum_{j=1}^D \|\nu_j\|_{\mathcal{M}(\Omega)} \\ &= \sum_{j=1}^D \sup_{\substack{\Omega = \bigcup_{i=1}^n A_i \\ n \in \mathbb{N}}} \sum_{i=1}^n |\nu_j(A_i)| \\ &\geq \sup_{\substack{\Omega = \bigcup_{i=1}^n A_i \\ n \in \mathbb{N}}} \sum_{j=1}^D \sum_{i=1}^n |\nu_j(A_i)| \\ &= \sup_{\substack{\Omega = \bigcup_{i=1}^n A_i \\ n \in \mathbb{N}}} \sum_{i=1}^n \sum_{j=1}^D |\nu_j(A_i)| \\ &= \sup_{\substack{\Omega = \bigcup_{i=1}^n A_i \\ n \in \mathbb{N}}} \sum_{i=1}^n \|\boldsymbol{\nu}(A_i)\|_1 = \|\boldsymbol{\nu}\|_{1,\mathcal{M}}. \end{aligned} \quad (55)$$

For the reverse inequality, observe that

$$\begin{aligned}
 \|\boldsymbol{\nu}\|_{1,\mathcal{M}} &= \sup_{\substack{\Omega = \bigcup_{i=1}^n A_i \\ n \in \mathbb{N}}} \sum_{i=1}^n \sum_{j=1}^D |\nu_j(A_i)| \\
 &\geq \sup_{\substack{\Omega = \bigcup_{i=1}^n A_i \\ n \in \mathbb{N}}} \sum_{i=1}^n |\nu_{j_0}(A_i)| \\
 &= \|\nu_{j_0}\|_{\mathcal{M}(\Omega)} \\
 &\geq \frac{1}{D} \sum_{j=1}^D \|\nu_j\|_{\mathcal{M}(\Omega)} \\
 &= \frac{1}{D} \|\boldsymbol{\nu}\|_{\mathcal{M},1},
 \end{aligned} \tag{56}$$

where $j_0 := \arg \max_{j \in [D]} \|\nu_j\|_{\mathcal{M}(\Omega)}$. Therefore, we have that

$$\|\boldsymbol{\nu}\|_{1,\mathcal{M}} \leq \|\boldsymbol{\nu}\|_{\mathcal{M},1} \leq D \|\boldsymbol{\nu}\|_{1,\mathcal{M}}, \tag{57}$$

which proves the lemma. ■

Remark 15 While the norms may be equivalent, only the $\|\cdot\|_{2,\mathcal{M}}$ -norm corresponds to weight-decay regularization as shown in Section 3.2. The use of some of these other norms have been explored in the literature. For example, Parhi and Nowak (2022) define a similar vector-valued variation space for ReLU networks with the $\|\cdot\|_{\mathcal{M},1}$ -norm, which does not correspond to weight-decay regularization.

A.1 Connection to the Total Variation of Function

The total variation of a measure is different than the total variation of a function, but the ideas are tightly linked. In the univariate case, consider a Radon measure $\mu \in \mathcal{M}(\mathbb{R})$ and suppose there exists a function $g_\mu : \mathbb{R} \rightarrow \mathbb{R}$ such that

$$\int_{\mathbb{R}} f(x) d\mu(x) = \int_{\mathbb{R}} f(x) dg_\mu(x), \tag{58}$$

where f is any bounded continuous function and the integral on the right-hand side is a Riemann–Stieltjes integral. If g_μ is differentiable and its derivative, denoted by g'_μ , is in $L^1(\mathbb{R})$, then we have the equality

$$\|\mu\|_{\mathcal{M}(\mathbb{R})} = \|g'_\mu\|_{L^1(\mathbb{R})}. \tag{59}$$

The quantity on the right-hand side is referred to as the *total variation* of the function g_μ . Furthermore, if g_μ is not differentiable in the classical sense, its distributional derivative can be identified with the Radon measure μ , i.e., $g'_\mu = \mu$, where equality is understood in $\mathcal{M}(\mathbb{R})$. In this case, $\|g'_\mu\|_{\mathcal{M}(\mathbb{R})} = \|\mu\|_{\mathcal{M}(\mathbb{R})}$ is the total variation of the function g_μ .

This correspondence extends to higher dimensions. If the vector measure $\mu = \nabla g$ for some function $g : \mathbb{R}^d \rightarrow \mathbb{R}$, then the norm we are using $\|\mu\|_{2,\mathcal{M}} = \|\nabla g\|_{2,\mathcal{M}}$, is the *isotropic total variation* of g . On the other hand, the norm $\|\mu\|_{1,\mathcal{M}} = \|\nabla g\|_{1,\mathcal{M}}$ corresponds to the *anisotropic total variation* of g . These two notions of total variation in multiple dimensions are often used in image processing.

The isotropic total variation is equivalently specified by in the dual form

$$\text{TV}_{\text{iso}}(g) = \sup_{\substack{\varphi \in \mathcal{D}(\mathbb{R}^d; \mathbb{R}^d) \\ \|\varphi\|_{L^\infty(\mathbb{R}^d; \mathbb{R}^d)} = 1}} \int_{\mathbb{R}^d} g(\mathbf{x}) \operatorname{div} \varphi(\mathbf{x}) \, d\mathbf{x}, \quad (60)$$

where

$$\|\varphi\|_{L^\infty(\mathbb{R}^d; \mathbb{R}^d)} := \operatorname{ess\,sup}_{\mathbf{x} \in \mathbb{R}^d} \|\varphi(\mathbf{x})\|_2, \quad (61)$$

and $\mathcal{D}(\mathbb{R}^d; \mathbb{R}^d)$ denotes the space of infinitely differentiable compactly supported functions mapping $\mathbb{R}^d \rightarrow \mathbb{R}^d$. We refer the reader to the book of Evans and Gariepy (2015) for more details about the total variation of a function.

Appendix B. Proof of Theorem 3

Proof First notice that for any $f = (f_1, \dots, f_D) \in L^2(\mathbb{R}^d; \mathbb{R}^D)$ we have that

$$\begin{aligned} \|f\|_{L^2(Q^d; \mathbb{R}^D)}^2 &= \int_{Q^d} \|f(\mathbf{x})\|_2^2 \, d\mathbf{x} \\ &= \int_{Q^d} \sum_{j=1}^D |f_j(\mathbf{x})|^2 \, d\mathbf{x} \\ &= \sum_{j=1}^D \int_{Q^d} |f_j(\mathbf{x})|^2 \, d\mathbf{x} \\ &= \sum_{j=1}^D \|f_j\|_{L^2(Q^d)}^2. \end{aligned} \quad (62)$$

For any $f \in \mathcal{V}_\sigma(Q^d; \mathbb{R}^D)$ we have that

$$\sum_{j=1}^D \|f_j\|_{\mathcal{V}_\sigma(Q^d)}^2 \leq \left(\sum_{j=1}^D \|f_j\|_{\mathcal{V}_\sigma(Q^d)} \right)^2. \quad (63)$$

Next,

$$\begin{aligned}
 \sum_{j=1}^D \|f_j\|_{\mathcal{V}_\sigma(Q^d)} &= \sum_{j=1}^D \inf_{\substack{\nu_j \in \mathcal{M}(\mathbb{S}^d) \\ f_j = f_{\nu_j}}} \|\nu_j\|_{\mathcal{M}(\mathbb{S}^d)} \\
 &\leq \inf_{\substack{\nu \in \mathcal{M}(\mathbb{S}^d; \mathbb{R}^D) \\ f = f_\nu}} \sum_{j=1}^D \|\nu_j\|_{\mathcal{M}(\mathbb{S}^d)} \\
 &= \inf_{\substack{\nu \in \mathcal{M}(\mathbb{S}^d; \mathbb{R}^D) \\ f = f_\nu}} \|\nu\|_{\mathcal{M},1} \\
 &\leq D^{\frac{3}{2}} \inf_{\substack{\nu \in \mathcal{M}(\mathbb{S}^d; \mathbb{R}^D) \\ f = f_\nu}} \|\nu\|_{2, \mathcal{M}} = D^{\frac{3}{2}} \|f\|_{\mathcal{V}_\sigma(Q^d; \mathbb{R}^D)}, \tag{64}
 \end{aligned}$$

where the equalities of the form $f_j = f_{\nu_j}$ or $f = f_\nu$ are understood as a function of $\mathbf{x} \in Q^d$. The fourth line follows from (57) combined with the inequality between $\|\cdot\|_1$ and $\|\cdot\|_2$.

To prove the claim, given any $f \in \mathcal{V}_\sigma(Q^d; \mathbb{R}^D)$, we construct a K -term approximant f_j^K for each component f_j , $j = 1, \dots, D$, as in (23). We then construct the vector-valued function $f_{DK} = (f_1^K, \dots, f_D^K)$ which has, at most, DK terms. This approximant satisfies

$$\begin{aligned}
 \|f - f_{DK}\|_{L^2(Q^d; \mathbb{R}^D)}^2 &= \sum_{j=1}^D \|f_j - f_j^K\|_{L^2(Q^d)}^2 \\
 &\leq \sum_{j=1}^D C_0^2 C_\sigma^2 \|f_j\|_{\mathcal{V}_\sigma(Q^d)}^2 K^{-1} \\
 &\leq C_0^2 C_\sigma^2 K^{-1} \sum_{j=1}^D \|f_j\|_{\mathcal{V}_\sigma(Q^d)}^2 \\
 &\leq C_0^2 C_\sigma^2 K^{-1} \left(\sum_{j=1}^D \|f_j\|_{\mathcal{V}_\sigma(Q^d)} \right)^2 \\
 &\leq C_0^2 C_\sigma^2 D^3 \|f\|_{\mathcal{V}_\sigma(Q^d; \mathbb{R}^D)}^2 K^{-1}. \tag{65}
 \end{aligned}$$

Therefore,

$$\|f - f_{DK}\|_{L^2(Q^d; \mathbb{R}^D)} \leq C_0 C_\sigma D^{\frac{3}{2}} \|f\|_{\mathcal{V}_\sigma(Q^d; \mathbb{R}^D)} K^{-1/2}, \tag{66}$$

which proves the theorem. ■

Appendix C. Proof of Theorem 5

Proof From the definition of the $\mathcal{V}_\sigma(\mathbb{R}^d; \mathbb{R}^D)$ -norm (16), it follows from a standard argument (e.g., Bartolucci et al., 2023, Proposition 3.7) that the problem in (28) is equivalent

to the problem

$$\inf_{\nu \in \mathcal{M}(\mathbb{S}^d; \mathbb{R}^D)} \sum_{i=1}^N \mathcal{L}\left(\mathbf{y}_i, \int_{\mathbb{S}^d} \sigma(\mathbf{w}^T \bar{\mathbf{x}}_i) d\nu(\mathbf{w})\right) + \lambda \|\nu\|_{2,\mathcal{M}} \quad (67)$$

in the sense that their infimal values are the same and if ν^* is a solution to (67), then

$$f_{\nu^*}(\mathbf{x}) = \int_{\mathbb{S}^d} \sigma(\mathbf{w}^T \bar{\mathbf{x}}_i) d\nu^*(\mathbf{w}) \quad (68)$$

is a solution to (28). We now proceed in four steps to prove the theorem.

Step (i): Existence of solutions to (67). Define

$$\mathcal{J}(\nu) := \sum_{i=1}^N \mathcal{L}\left(\mathbf{y}_i, \int_{\mathbb{S}^d} \sigma(\mathbf{w}^T \bar{\mathbf{x}}_i) d\nu(\mathbf{w})\right) + \lambda \|\nu\|_{2,\mathcal{M}}. \quad (69)$$

Given an arbitrary $\nu_0 \in \mathcal{M}(\mathbb{S}^d; \mathbb{R}^D)$, let $C_0 := \mathcal{J}(\nu_0)$. Then, we can transform (67) into the constrained problem

$$\inf_{\nu \in \mathcal{M}(\mathbb{S}^d; \mathbb{R}^D)} \mathcal{J}(\nu) \quad \text{s.t.} \quad \|\nu\|_{2,\mathcal{M}} \leq C_0/\lambda. \quad (70)$$

This transformation is valid since any measure that does not satisfy the constraint will have an objective value strictly larger than ν_0 and therefore will not be in the solution set. Next, we note that we can write

$$\mathcal{J}(\nu) = \sum_{i=1}^N \mathcal{L}(\mathbf{y}_i, \mathbf{H}_i\{\nu\}) + \lambda \|\nu\|_{2,\mathcal{M}} \quad (71)$$

where

$$\mathbf{H}_i\{\nu\} := \begin{bmatrix} \langle \nu, h_{i,j} \rangle \\ \vdots \\ \langle \nu, h_{i,D} \rangle \end{bmatrix} \in \mathbb{R}^D \quad (72)$$

with $h_{i,j}(\mathbf{w}) := \sigma(\mathbf{w}^T \bar{\mathbf{x}}_i) \mathbf{e}_j$, where $\mathbf{e}_j \in \mathbb{R}^D$ is the j th canonical unit vector.

Note that $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ is continuous by assumption. Therefore, we readily observe that $h_{i,j} \in C(\mathbb{S}^d; \mathbb{R}^D)$, the space of continuous functions on \mathbb{S}^d taking values in \mathbb{R}^D . In (72), $\langle \cdot, \cdot \rangle$ denotes the duality pairing⁷ between $C(\mathbb{S}^d; \mathbb{R}^D)$ and $\mathcal{M}(\mathbb{S}^d; \mathbb{R}^D)$. Thus, for $i = 1, \dots, N$, \mathbf{H}_i is component-wise weak* continuous on $\mathcal{M}(\mathbb{S}^d; \mathbb{R}^D)$. Since $\mathcal{L}(\cdot, \cdot)$ is lower semicontinuous in its second argument combined with the fact that every norm is weak* continuous on its corresponding Banach space, we have that \mathcal{J} is weak* continuous on $\mathcal{M}(\mathbb{S}^d; \mathbb{R}^D)$. By the Banach–Alaoglu theorem (Rudin, 1991, Chapter 3), the constraint set of (70) is weak* compact. Thus, (70) is the minimization of a weak* continuous functional over a weak* compact set. By the Weierstrass extreme value theorem on general topological spaces (Kurdila and Zabarankin, 2006, Chapter 5), there exists a solution to (70) (and subsequently of (67)).

7. The continuous dual of $C(\mathbb{S}^d; \mathbb{R}^D)$ can be identified with $\mathcal{M}(\mathbb{S}^d; \mathbb{R}^D)$ by Singer's representation theorem (Singer, 1957, 1959) (see also Hensgen, 1996; Bredies and Holler, 2020).

Step (ii): Recasting (67) as an interpolation problem. Let $\tilde{\nu}$ be a (not necessarily unique) solution to (67), which is guaranteed to exist by the previous argumentation. For $i = 1, \dots, D$, define

$$\mathbf{z}_i = \mathbf{H}_i\{\tilde{\nu}\} \in \mathbb{R}^D. \quad (73)$$

Then, $\tilde{\nu}$ must satisfy

$$\tilde{\nu} \in \arg \min_{\nu \in \mathcal{M}(\mathbb{S}^d; \mathbb{R}^D)} \|\nu\|_{2, \mathcal{M}} \quad \text{s.t.} \quad \mathbf{H}_i\{\nu\} = \mathbf{z}_i, \quad i = 1, \dots, N. \quad (74)$$

To see this, we note that if this were not the case, it would contradict the optimality of $\tilde{\nu}$. This reduction implies that any solution to the interpolation problem (74) will also be a solution to (67).

Step (iii): The form of the solution. We can rewrite the interpolation problem as

$$\min_{\nu \in \mathcal{M}(\mathbb{S}^d; \mathbb{R}^D)} \|\nu\|_{2, \mathcal{M}} \quad \text{s.t.} \quad \langle \nu, h_{i,j} \rangle = z_{i,j}, \quad i = 1, \dots, N \text{ and } j = 1, \dots, D. \quad (75)$$

This is the vector-valued analogue of the classical (Radon) measure recovery problem with ND weak* continuous measurements (cf., Zuhovickiĭ, 1948). By the abstract representer theorems of Boyer et al. (2019); Bredies and Carioni (2020); Unser (2021), there always exists a solution to (75) that takes the form

$$\nu^* = \sum_{k=1}^K c_k \mathbf{e}_k \quad (76)$$

with $K \leq ND$, $c_k \in \mathbb{R} \setminus \{0\}$ where for $k = 1, \dots, K$, \mathbf{e}_k is an extreme point of the unit regularization ball

$$B := \{\nu \in \mathcal{M}(\mathbb{S}^d; \mathbb{R}^D) : \|\nu\|_{2, \mathcal{M}} \leq 1\}. \quad (77)$$

From Werner (1984, Theorem 2), the extreme points of B take the form $\mathbf{a}\delta_{\mathbf{w}}$ with $\mathbf{a} \in \mathbb{R}^D$, $\|\mathbf{a}\|_2 = 1$, and $\mathbf{w} \in \mathbb{S}^d$. Thus, we can write

$$\nu^* = \sum_{k=1}^K c_k \mathbf{a}_k \delta_{\mathbf{w}_k} \quad (78)$$

with $\mathbf{a}_k \in \mathbb{R}^D$ and $\mathbf{w}_k \in \mathbb{S}^d$. By the equivalence between (67) and (28), we find that there exists a solution to (28) that takes the form

$$\hat{f}(\mathbf{x}) = \sum_{k=1}^{\hat{K}} \hat{\mathbf{v}}_k \sigma(\hat{\mathbf{w}}_k^T \bar{\mathbf{x}}), \quad \mathbf{x} \in \mathbb{R}^d, \quad (79)$$

where $\hat{K} \leq ND$, $\hat{\mathbf{v}}_k := c_k \mathbf{a}_k \in \mathbb{R}^D$ and $\hat{\mathbf{w}}_k \in \mathbb{S}^d$.

Step (iv): Sharpening the sparsity bound. Step (iii) establishes that there exists an optimal solution with $\hat{K} \leq ND$ neurons. Thus, we can apply the same argumentation as in the proof of Theorem 10 to find another solution where the number of neurons is bounded by the product of the rank of the labels $\mathbf{Y} \in \mathbb{R}^{D \times N}$ with the rank of the post-activation feature matrix $\Phi \in \mathbb{R}^{\hat{K} \times N}$. From the dimensions of these matrices, we see that the product of the ranks is $\leq N^2$. Therefore, there exists a solution to (28) that takes the form

$$f^*(\mathbf{x}) = \sum_{k=1}^{K_0} \mathbf{v}_k \sigma(\mathbf{w}_k^T \bar{\mathbf{x}}), \quad \mathbf{x} \in \mathbb{R}^d, \quad (80)$$

where $K_0 \leq N^2$, $\mathbf{v}_k \in \mathbb{R}^D$ and $\mathbf{w}_k \in \mathbb{S}^d$. Thus, we can always find a solution to (28) with $\min\{N^2, ND\}$ neurons improving upon the bound of $ND + 1$ predicted by Carathéodory's theorem. \blacksquare

Appendix D. Proof of Theorem 8

Proof Given $f = f^{(L)} \circ \dots \circ f^{(1)}$ such that $f^{(\ell)} \in \mathcal{V}_\sigma(\mathbb{R}^{d_{\ell-1}}, \mathbb{R}^{d_\ell})$, $\ell = 1, \dots, L$, let the functional \mathcal{J} denote the objective value of the optimization problem, i.e.,

$$\mathcal{J}(f) := \mathcal{J}(f^{(L)}, \dots, f^{(1)}) := \sum_{i=1}^N \mathcal{L}(\mathbf{y}_i, f(\mathbf{x}_i)) + \lambda \sum_{\ell=1}^L \|f^{(\ell)}\|_{\mathcal{V}_\sigma(\mathbb{R}^{d_{\ell-1}}; \mathbb{R}^{d_\ell})}. \quad (81)$$

Next, following the same approach in the proof of Theorem 5, for an arbitrary $g = g^{(L)} \circ \dots \circ g^{(1)}$, where $g^{(\ell)} \in \mathcal{V}_\sigma(\mathbb{R}^{d_{\ell-1}}, \mathbb{R}^{d_\ell})$, $\ell = 1, \dots, L$, we define its objective value as $C := \mathcal{J}(g)$. The unconstrained problem (32) can then be transformed into the equivalent constrained problem

$$\inf_{\substack{f^{(\ell)} \in \mathcal{V}_\sigma(\mathbb{R}^{d_{\ell-1}}; \mathbb{R}^{d_\ell}) \\ \ell=1, \dots, L \\ f=f^{(L)} \circ \dots \circ f^{(1)}}} \mathcal{J}(f) \quad \text{s.t.} \quad \|f^{(\ell)}\|_{\mathcal{V}_\sigma(\mathbb{R}^{d_{\ell-1}}; \mathbb{R}^{d_\ell})} \leq C/\lambda, \quad \ell = 1, \dots, L. \quad (82)$$

This transformation is valid since any collection of functions $f^{(\ell)}$, $\ell = 1, \dots, L$, that do not satisfy the constraints would result in an objective value that is strictly larger than that of g , and would therefore not be in the solution set.

For any $f_0 = f_0^{(L)} \circ \dots \circ f_0^{(1)}$, where $f^{(\ell)} \in \mathcal{V}_\sigma(\mathbb{R}^{d_{\ell-1}}; \mathbb{R}^{d_\ell})$, $\ell = 1, \dots, L$, we will show that, for any fixed $\tilde{\ell} \in \{1, \dots, L\}$, the map $f_0^{(\tilde{\ell})} \mapsto \mathcal{J}(f_0)$ is weak* lower semicontinuous on $\mathcal{V}_\sigma(\mathbb{R}^{d_{\tilde{\ell}-1}}; \mathbb{R}^{d_{\tilde{\ell}}})$. First, observe that the map $f_0^{(\tilde{\ell})} \mapsto f_0(\mathbf{x}_0)$, for any $\mathbf{x}_0 \in \mathbb{R}^d$, is component-wise weak* continuous on $\mathcal{V}_\sigma(\mathbb{R}^{d_{\tilde{\ell}-1}}; \mathbb{R}^{d_{\tilde{\ell}}})$. Indeed, this follows since for any $\mathbf{x}_0 \in \mathbb{R}^d$ the point evaluation operator

$$\tilde{\mathbf{x}}_0 : f \mapsto f(\mathbf{x}_0) = \begin{bmatrix} f_1(\mathbf{x}_0) \\ \vdots \\ f_D(\mathbf{x}_0) \end{bmatrix} \quad (83)$$

is component-wise weak* continuous by Lemma 2.9 of Parhi and Nowak (2022) combined with the equivalence of norms in Lemma 14. Thus, since $f_0^{(\tilde{\ell})} \mapsto f_0^{(L)} \circ \dots \circ f_0^{(1)}(\mathbf{x}_0)$ is made up of compositions of component-wise continuous and component-wise weak* continuous functions, it is therefore itself component-wise weak* continuous on $\mathcal{V}_\sigma(\mathbb{R}^{d_{\tilde{\ell}-1}}; \mathbb{R}^{d_{\tilde{\ell}}})$. Therefore, the map $(f_0^{(1)}, \dots, f_0^{(L)}) \mapsto \mathcal{J}(f_0)$ is weak* lower semicontinuous on $\mathcal{V}_\sigma(\mathbb{R}^{d_{\tilde{\ell}-1}}; \mathbb{R}^{d_{\tilde{\ell}}})$. Finally, by the Banach–Alaoglu theorem (Rudin, 1991, Chapter 3), the feasible set in (82) is weak* compact. Therefore there exists a solution to (82), and thus (32), by the Weierstrass extreme value theorem on general topological spaces (Kurdila and Zabarankin, 2006, Chapter 5).

To complete the proof, let $\tilde{f} = \tilde{f}^{(L)} \circ \dots \circ \tilde{f}^{(1)}$ be any solution to (32). By applying \tilde{f} to each data point \mathbf{x}_i , $i = 1, \dots, N$, we can recursively compute the intermediate vectors $\mathbf{z}_{i,\ell} \in \mathbb{R}^{d_\ell}$ as follows:

- Initialize $\mathbf{z}_{i,0} := \mathbf{x}_i$.
- For each $\ell = 1, \dots, L$, recursively update $\mathbf{z}_{i,\ell} := \tilde{f}^\ell(\mathbf{z}_{i,\ell-1})$.

The solution \tilde{f} must satisfy

$$\tilde{f}^{(\ell)} \in \arg \min_{f \in \mathcal{V}_\sigma(\mathbb{R}^{d_{\ell-1}}; \mathbb{R}^{d_\ell})} \|f\|_{\mathcal{V}_\sigma(\mathbb{R}^{d_{\ell-1}}; \mathbb{R}^{d_\ell})} \quad \text{s.t.} \quad f(\mathbf{z}_{i,\ell-1}) = \mathbf{z}_{i,\ell}, i = 1, \dots, N, \quad (84)$$

for $\ell = 1, \dots, L$ (since otherwise, it would contradict the optimality of \tilde{f}). By Theorem 5, there always exists a solution to (84) that takes the form of a shallow vector-valued neural network. Hence, there always exists a solution to (32) of the form in (33). \blacksquare

Appendix E. Proof of Theorem 11

Our proof relies extensively on Carathéodory’s theorem (Clarke, 2013, Proposition 2.6) which we first state here.

Theorem 16 (Carathéodory’s theorem) *Let S be a subset of a normed vector space with finite dimension R . Then every point $\mathbf{x} \in \text{Conv}(S)$, the convex hull of S , can be represented by a convex combination of at most $R + 1$ points from S .*

Proof [Proof of Theorem 11] The constraint $\Psi = \mathbf{V}\Phi$ can be satisfied if and only if the row space of Ψ is contained in the row space of Φ . This assumption, and the fact that the objective is continuous and coercive, implies a solution exists. Suppose that \mathbf{V} is a solution to our problem. Then we will show that there exists a (possibly different) solution $\hat{\mathbf{V}}$ with no more than $r_\Phi r_\Psi$ nonzero columns.

Let $\text{col}(\mathbf{V})$ and $\text{col}(\Psi)$ denote the column space of \mathbf{V} and Ψ respectively. We first show that $\text{col}(\mathbf{V}) = \text{col}(\Psi)$. Let $\mathbf{V} = \mathbf{A} + \mathbf{B}$, where $\mathbf{A} \in \text{col}(\Psi)$ and $\mathbf{B} \in \text{col}(\Psi)^\perp$, the subspace orthogonal to $\text{col}(\Psi)$. Then,

$$\sum_{k=1}^K \|\mathbf{v}_k\|_2 = \sum_{k=1}^K \sqrt{\|\mathbf{a}\|_2^2 + \|\mathbf{b}\|_2^2}. \quad (85)$$

Let $\mathbf{P}_{\text{col}(\Psi)}$ be the orthogonal projection onto $\text{col}(\Psi)$. We can then express the constraint as $\Psi = \mathbf{P}_{\text{col}(\Psi)}\Psi = \mathbf{P}_{\text{col}(\Psi)}(\mathbf{A} + \mathbf{B})\Phi = \mathbf{A}\Phi$. Thus the solution must have $\mathbf{B} = 0$, since anything nonzero would increase the objective without contributing to the constraint. Therefore, $\text{col}(\mathbf{V}) = \text{col}(\Psi)$ and $\text{rank}(\mathbf{V}) = \text{rank}(\Psi)$. Now observe that we can also express the constraint as a sum of outer products,

$$\Psi = \sum_{k=1}^K \mathbf{v}_k \phi_k^T \quad (86)$$

where \mathbf{v}_k are the columns of \mathbf{V} and ϕ_k^T are the rows of Φ . Let $\mathbf{M}_k = \mathbf{v}_k \phi_k^T$. Since the \mathbf{v}_k belong to an r_Ψ dimensional subspace and the ϕ_k belong to an r_Φ dimensional subspace, the \mathbf{M}_k all belong to a subspace of dimension at most $r_\Phi r_\Psi$. For ease of notation we let $R = r_\Phi r_\Psi$. Now, define the optimal objective value $\gamma = \sum_{k=1}^K \|\mathbf{v}_k\|_2$ and $\alpha_k = \|\mathbf{v}_k\|_2 / \gamma$ so that $\sum_{k=1}^K \alpha_k = 1$. We can then write Ψ as

$$\Psi = \sum_{k=1}^K \alpha_k \left(\frac{1}{\alpha_k} \mathbf{M}_k \right) = \sum_{k=1}^K \alpha_k \widetilde{\mathbf{M}}_k. \quad (87)$$

This shows that Ψ is in the convex hull of matrices $\widetilde{\mathbf{M}}_k$ with dimension at most R . Carathéodory's theorem implies that we can represent Ψ by a convex combination of a subset $\{\widetilde{\mathbf{M}}_j\}_{j \in J}$ where $J \subset \{1, \dots, K\}$ and $|J| \leq R + 1$. Thus we can represent Ψ with no more than $R + 1$ nonzero columns vectors in the solution

$$\Psi = \sum_{j \in J} \beta_j \widetilde{\mathbf{M}}_j = \sum_{j \in J} \beta_j \left(\frac{1}{\alpha_j} \mathbf{v}_j \phi_j^T \right). \quad (88)$$

Now to show that a solution exists with no more than R nonzero columns we study the KKT conditions and apply Carathéodory's theorem a second time. Assume that $|J| = R + 1$ and define $\tilde{\mathbf{v}}_j = \frac{1}{\alpha_j} \mathbf{v}_j$. Thus Ψ is in the convex hull of

$$\{\tilde{\mathbf{v}}_j \phi_j^T\}_{j \in J}. \quad (89)$$

Each matrix $\tilde{\mathbf{v}}_j \phi_j^T$ belongs to a subspace of dimension at most R , therefore, any matrix in this set can be expressed as a linear combination of the others

$$\tilde{\mathbf{v}}_i \phi_i^T = \sum_{j \in J; j \neq i} c_j \tilde{\mathbf{v}}_j \phi_j^T. \quad (90)$$

By the subgradient optimality conditions, we will prove that we must have $\sum_{j \in J; j \neq i} c_j = 1$. This in turn implies that the set of matrices $\{\tilde{\mathbf{v}}_j \phi_j^T\}_{j \in J}$ are not only linearly dependent but they also span an $R - 1$ dimensional affine space (i.e., an $R - 1$ dimensional hyperplane not including the origin (Rockafellar, 1970)). We can then apply Carathéodory's again to show that a convex combination of R matrices from this set suffices to satisfy the constraint.

The Lagrangian of our optimization problem (51) is

$$\mathcal{L}(\mathbf{V}, \boldsymbol{\nu}_1, \boldsymbol{\nu}_2) = \sum_{k=1}^K \|\mathbf{v}_k\|_2 + \boldsymbol{\nu}_1^T (\Psi - \mathbf{V}\Phi) \boldsymbol{\nu}_2 \quad (91)$$

with Lagrange multipliers $\boldsymbol{\nu}_1 \in \mathbb{R}^D$ and $\boldsymbol{\nu}_2 \in \mathbb{R}^N$. By the KKT conditions, any solution must satisfy

$$0 \in \partial_{\mathbf{V}} \mathcal{L}, \quad (92)$$

$$0 = \nabla_{\boldsymbol{\nu}_1, \boldsymbol{\nu}_2} \mathcal{L}. \quad (93)$$

Now for our original solution \mathbf{V} we have,

$$\partial_{\mathbf{V}} \mathcal{L} = \partial \left(\sum_{k=1}^K \|\mathbf{v}_k\|_2 \right) - \boldsymbol{\nu}_1 (\boldsymbol{\Phi} \boldsymbol{\nu}_2)^T = 0 \quad (94)$$

$$\nabla_{\boldsymbol{\nu}_1, \boldsymbol{\nu}_2} \mathcal{L} = \boldsymbol{\Psi} - \mathbf{V} \boldsymbol{\Phi} = 0. \quad (95)$$

Note we have no restriction of $\boldsymbol{\nu}_1$ and $\boldsymbol{\nu}_2$. Then for any $j \in J$ (the nonzero columns of \mathbf{V}) the j th column of the subgradient of our objective is

$$\partial_{\mathbf{V}} \left(\sum_{k=1}^K \|\mathbf{v}_k\|_2 \right)_j = \frac{\mathbf{v}_j}{\|\mathbf{v}_j\|_2}. \quad (96)$$

Therefore, by (94) for the j th column of the subgradient we have

$$\partial \left(\sum_{k=1}^K \|\mathbf{v}_k\|_2 \right)_j = \boldsymbol{\nu}_1 \boldsymbol{\nu}_2^T \boldsymbol{\phi}_j \implies \frac{\mathbf{v}_j}{\|\mathbf{v}_j\|_2} = \boldsymbol{\nu}_1 \boldsymbol{\nu}_2^T \boldsymbol{\phi}_j \quad (97)$$

where $\boldsymbol{\phi}_j \in \mathbb{R}^{N \times 1}$ is the j th row of $\boldsymbol{\Phi}$. Now consider the i th column of (94) and right-multiply both sides by $\tilde{\mathbf{v}}_i^T$ to obtain

$$\frac{\mathbf{v}_i \tilde{\mathbf{v}}_i^T}{\|\mathbf{v}_i\|_2} = \boldsymbol{\nu}_1 \boldsymbol{\nu}_2^T \boldsymbol{\phi}_i \tilde{\mathbf{v}}_i^T. \quad (98)$$

By the linear dependence of the matrices $\{\tilde{\mathbf{v}}_j \boldsymbol{\phi}_j^T\}_{j \in J}$ for any $i \in J$ we have $\boldsymbol{\phi}_i \tilde{\mathbf{v}}_i^T = \sum_{j \in J; j \neq i} c_j \boldsymbol{\phi}_j \tilde{\mathbf{v}}_j^T$ for some constants c_j , substituting this in we get

$$\begin{aligned} \frac{\mathbf{v}_i \tilde{\mathbf{v}}_i^T}{\|\mathbf{v}_i\|_2} &= \boldsymbol{\nu}_1 \boldsymbol{\nu}_2^T \sum_{j \in J; j \neq i} c_j \boldsymbol{\phi}_j \tilde{\mathbf{v}}_j^T \\ &= \sum_{j \in J; j \neq i} c_j \boldsymbol{\nu}_1 \boldsymbol{\nu}_2^T \boldsymbol{\phi}_j \tilde{\mathbf{v}}_j^T \\ &= \sum_{j \in J; j \neq i} c_j \frac{\mathbf{v}_j \tilde{\mathbf{v}}_j^T}{\|\mathbf{v}_j\|_2}, \end{aligned} \quad (99)$$

where the final equality follows from (97). Since $\tilde{\mathbf{v}}_j = \frac{1}{\alpha_j} \mathbf{v}_j = \frac{\gamma}{\|\mathbf{v}_j\|_2} \mathbf{v}_j$ plugging this in on both sides and cancelling out γ we have

$$\frac{\mathbf{v}_i \mathbf{v}_i^T}{\|\mathbf{v}_i\|_2^2} = \sum_{j \in J; j \neq i} c_j \frac{\mathbf{v}_j \mathbf{v}_j^T}{\|\mathbf{v}_j\|_2^2}. \quad (100)$$

Taking the trace of the left and right hand sides reveals

$$\sum_{j \neq i} c_j = 1. \quad (101)$$

Therefore any $\phi_i \tilde{\mathbf{v}}_i^T$ is in the affine hull of the other matrices and thus the matrices $\{\tilde{\mathbf{v}}_j \phi_j^T\}_{j \in J}$ form an affine set of dimension $R - 1$ (Rockafellar, 1970). This implies that the matrices do not just span an R subspace but rather an $R - 1$ dimensional hyperplane that does not pass through the origin. Thus, the vector space corresponding to this set of matrices has dimension $R - 1$. By applying Carathéodory's theorem again we can represent the constraint as a convex combination of just R nonzero vectors

$$\Psi = \sum_{j \in J'} \hat{\beta}_j \tilde{\mathbf{v}}_j \phi_j^T = \sum_{j \in J'} \mathbf{v}_j^* \phi_j^T, \quad (102)$$

where $\mathbf{v}_j^* := \hat{\beta}_j \tilde{\mathbf{v}}_j$, $J' \subset \{1, \dots, K\}$ and $|J'| \leq R = r_\Phi r_\Psi$. The coefficients $\hat{\beta}_j$ are nonnegative and sum to 1. Finally, it remains to show that this new representation of Ψ has the same objective value as the original solution, indeed we have

$$\sum_{j \in J'} \|\mathbf{v}_j^*\|_2 = \sum_{j \in J'} \frac{\hat{\beta}_j}{\alpha_j} \|\mathbf{v}_j\|_2 = \sum_{j \in J'} \gamma \hat{\beta}_j = \gamma \quad (103)$$

where we recall that γ is the objective value obtained by our original assumed solution.

The lower bound follows solely from the constraint. First we note that for any Ψ we always need at least r_Ψ nonzero columns in the solution \mathbf{V} . If this were not the case we would have $\text{rank}(\mathbf{V}) < r_\Psi$ and therefore $\text{rank}(\Psi) \leq \min\{\text{rank}(\mathbf{V}), r_\Phi\} < r_\Psi$ contradicting the fact that $\text{rank}(\Psi) = r_\Psi$.

Now for any Φ the number of subspaces spanned by $p < r_\Phi$ of the rows of Φ is at most

$$\binom{K}{1} + \binom{K}{2} + \dots + \binom{K}{r_\Phi - 1}. \quad (104)$$

The union of these subspaces has Lebesgue measure zero in the $\text{row}(\Phi)$. Therefore, there exist uncountably many points in $\text{row}(\Phi)$ that cannot be represented with fewer than r_Φ columns of Φ . ■

References

Andreas Argyriou, Theodoros Evgeniou, and Massimiliano Pontil. Convex multi-task feature learning. *Machine Learning*, 73:243–272, 2008.

Francis Bach. Breaking the curse of dimensionality with convex neural networks. *Journal of Machine Learning Research*, 18(1):629–681, 2017.

Waheed U. Bajwa, Marco F. Duarte, and Robert Calderbank. Conditioning of random block subdictionaries with applications to block-sparse recovery and regression. *IEEE Transactions on Information Theory*, 61(7):4060–4079, 2015.

Andrew R. Barron. Universal approximation bounds for superpositions of a sigmoidal function. *IEEE Transactions on Information Theory*, 39(3):930–945, 1993.

Andrew R. Barron, Albert Cohen, Wolfgang Dahmen, and Ronald A. DeVore. Approximation and learning by greedy algorithms. *Annals of Statistics*, 36(1):64–94, 2008.

Francesca Bartolucci, Ernesto De Vito, Lorenzo Rosasco, and Stefano Vigogna. Understanding neural networks with reproducing kernel Banach spaces. *Applied and Computational Harmonic Analysis*, 62:194–236, 2023.

Francesca Bartolucci, Ernesto De Vito, Lorenzo Rosasco, and Stefano Vigogna. Neural reproducing kernel Banach spaces and representer theorems for deep networks. *arXiv preprint arXiv:2403.08750*, 2024.

Yoshua Bengio, Nicolas Roux, Pascal Vincent, Olivier Delalleau, and Patrice Marcotte. Convex neural networks. *Advances in Neural Information Processing Systems*, 18, 2005.

Claire Boyer, Antonin Chambolle, Yohann De Castro, Vincent Duval, Frédéric de Gournay, and Pierre Weiss. On representer theorems and convex regularization. *SIAM Journal on Optimization*, 29(2):1260–1281, 2019.

Kristian Bredies and Marcello Carioni. Sparsity of solutions for variational inverse problems with finite-dimensional data. *Calculus of Variations and Partial Differential Equations*, 59(1):Paper No. 14, 26, 2020.

Kristian Bredies and Martin Holler. Higher-order total variation approaches and generalisations. *Inverse Problems*, 36(12):123001, 2020.

Kristian Bredies, Jonathan Chirinos Rodriguez, and Emanuele Naldi. On extreme points and representer theorems for the Lipschitz unit ball on finite metric spaces. *Archiv der Mathematik*, pages 1–8, 2024.

Jie Chen and Xiaoming Huo. Theoretical results on sparse representations of multiple-measurement vectors. *IEEE Transactions on Signal Processing*, 54(12):4634–4643, 2006.

Francis Clarke. *Functional analysis, calculus of variations and optimal control*, volume 264. Springer, 2013.

Shane F. Cotter, Bhaskar D. Rao, Kjersti Engan, and Kenneth Kreutz-Delgado. Sparse solutions to linear inverse problems with multiple measurement vectors. *IEEE Transactions on Signal Processing*, 53(7):2477–2488, 2005.

Steven Diamond and Stephen Boyd. CVXPY: A python-embedded modeling language for convex optimization. *Journal of Machine Learning Research*, 17(1):2909–2913, 2016.

J. Diestel and J.J. Uhl. *Vector Measures*. Mathematical surveys and monographs. American Mathematical Society, 1977. ISBN 9780821815151.

Simon S. Du, Wei Hu, and Jason D. Lee. Algorithmic regularization in learning deep homogeneous models: Layers are automatically balanced. *Advances in Neural Information Processing Systems*, 31, 2018.

Weinan E, Chao Ma, and Lei Wu. The Barron space and the flow-induced function spaces for neural network models. *Constructive Approximation*, 55(1):369–406, 2022.

Tolga Ergen and Mert Pilanci. Convex geometry and duality of over-parameterized neural networks. *Journal of Machine Learning Research*, 2021.

Lawrence Craig Evans and Ronald F. Gariepy. *Measure theory and fine properties of functions*. CRC press, 2015.

Ruili Feng, Kecheng Zheng, Yukun Huang, Deli Zhao, Michael Jordan, and Zheng-Jun Zha. Rank diminishing in deep neural networks. *Advances in Neural Information Processing Systems*, 35, 2022.

Yves Grandvalet. Least absolute shrinkage is equivalent to quadratic penalization. In *International Conference on Artificial Neural Networks*, pages 201–206. Springer, 1998.

Yves Grandvalet and Stéphane Canu. Outcomes of the equivalence of adaptive ridge with least absolute shrinkage. *Advances in Neural Information Processing Systems*, 11, 1998.

Stephen Hanson and Lorien Pratt. Comparing biases for minimal network construction with back-propagation. *Advances in Neural Information Processing Systems*, 1, 1988.

Wolfgang Hensgen. A simple proof of Singer’s representation theorem. *Proceedings of the American Mathematical Society*, 124(10):3211–3212, 1996.

Minyoung Huh, Hossein Mobahi, Richard Zhang, Brian Cheung, Pulkit Agrawal, and Phillip Isola. The low-rank simplicity bias in deep networks. *Transactions on Machine Learning Research*, 2023. ISSN 2835-8856.

Arthur Jacot, Eugene Golikov, Clément Hongler, and Franck Gabriel. Feature learning in L_2 -regularized DNNs: Attraction/repulsion and sparsity. In *Advances in Neural Information Processing Systems*, volume 35, 2022.

Ziwei Ji and Matus Telgarsky. Gradient descent aligns the layers of deep linear networks. In *International Conference on Learning Representations*, 2019.

Lee K. Jones. A simple lemma on greedy approximation in Hilbert space and convergence rates for projection pursuit regression and neural network training. *The Annals of Statistics*, pages 608–613, 1992.

Yury Korolev. Two-layer neural networks with values in a Banach space. *SIAM Journal on Mathematical Analysis*, 54(6):6358–6389, 2022.

Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Communications of the ACM*, 60(6):84–90, 2017.

Daniel Kunin, Javier Sagastuy-Brena, Surya Ganguli, Daniel L.K. Yamins, and Hidenori Tanaka. Neural mechanics: Symmetry and broken conservation laws in deep learning dynamics. In *International Conference on Learning Representations*, 2021.

Andrew J. Kurdila and Michael Zabarankin. *Convex Functional Analysis*. Systems & Control: Foundations & Applications. Birkhäuser Basel, 2006.

Soo Min Kwon, Zekai Zhang, Dogyoon Song, Laura Balzano, and Qing Qu. Efficient low-dimensional compression of overparameterized models. In *International Conference on Artificial Intelligence and Statistics*, pages 1009–1017. PMLR, 2024.

Věra Kůrková and Marcello Sanguineti. Bounds on rates of variable-basis and neural-network approximation. *IEEE Transactions on Information Theory*, 47(6):2659–2665, 2001.

Věra Kůrková and Marcello Sanguineti. Comparison of worst case errors in linear and neural network approximation. *IEEE Transactions on Information Theory*, 48(1):264–275, 2002.

Thien Le and Stefanie Jegelka. Training invariances and the low-rank phenomenon: beyond linear networks. In *International Conference on Learning Representations*, 2022.

Rong Rong Lin, Hai Zhang Zhang, and Jun Zhang. On reproducing kernel Banach spaces: Generic definitions and unified framework of constructions. *Acta Mathematica Sinica, English Series*, 38(8):1459–1483, 2022.

Hrushikesh N. Mhaskar. On the tractability of multivariate integration and approximation by neural networks. *Journal of Complexity*, 20(4):561–590, 2004.

Aaron Mishkin, Arda Sahiner, and Mert Pilanci. Fast convex optimization for two-layer ReLU networks: Equivalent model classes and cone decompositions. In *International Conference on Machine Learning*, 2022.

Alireza Mousavi-Hosseini, Sejun Park, Manuela Girotti, Ioannis Mitliagkas, and Murat A. Erdogdu. Neural networks efficiently learn low-dimensional representations with SGD. In *The Eleventh International Conference on Learning Representations*, 2023.

Kamil Nar, Orhan Ocal, S. Shankar Sastry, and Kannan Ramchandran. Cross-entropy loss leads to poor margins, 2019.

Behnam Neyshabur, Ryota Tomioka, and Nathan Srebro. In search of the real inductive bias: On the role of implicit regularization in deep learning. In *International Conference on Learning Representations (Workshop)*, 2015.

Guillaume Obozinski, Ben Taskar, and Michael Jordan. Multi-task feature selection. *Statistics Department, UC Berkeley, Tech. Rep*, 2(2.2):2, 2006.

Guillaume Obozinski, Ben Taskar, and Michael I. Jordan. Joint covariate selection and joint subspace selection for multiple classification problems. *Statistics and Computing*, 20:231–252, 2010.

Guillaume Obozinski, Martin J. Wainwright, and Michael I. Jordan. Support union recovery in high-dimensional multivariate regression. *The Annals of Statistics*, 39(1):1–47, 2011.

Greg Ongie, Rebecca Willett, Daniel Soudry, and Nathan Srebro. A function space view of bounded norm infinite width ReLU nets: The multivariate case. In *International Conference on Learning Representations*, 2020.

Vardan Papyan, XY Han, and David L. Donoho. Prevalence of neural collapse during the terminal phase of deep learning training. *Proceedings of the National Academy of Sciences*, 117(40):24652–24663, 2020.

Rahul Parhi and Robert D. Nowak. The role of neural network activation functions. *IEEE Signal Processing Letters*, 27:1779–1783, 2020.

Rahul Parhi and Robert D. Nowak. Banach space representer theorems for neural networks and ridge splines. *Journal of Machine Learning Research*, 22(43):1–40, 2021.

Rahul Parhi and Robert D. Nowak. What kinds of functions do deep neural networks learn? Insights from variational spline theory. *SIAM Journal on Mathematics of Data Science*, 4(2):464–489, 2022.

Rahul Parhi and Robert D. Nowak. Near-minimax optimal estimation with shallow ReLU neural networks. *IEEE Transactions on Information Theory*, 69(2):1125–1140, 2023a.

Rahul Parhi and Robert D. Nowak. Deep learning meets sparse regularization: A signal processing perspective. *IEEE Signal Processing Magazine*, 40(6):63–74, 2023b.

Gilles Pisier. Remarques sur un résultat non publié de B. Maurey. *Séminaire d’Analyse Fonctionnelle (dit “Maurey-Schwartz”)*, pages 1–12, April 1981.

Adityanarayanan Radhakrishnan, Eshaan Nichani, Daniel Bernstein, and Caroline Uhler. On alignment in deep linear neural networks. *ICML Workshop on Over-parameterization: Pitfalls and Opportunities*, 2021.

R. Tyrrell Rockafellar. *Convex analysis*. Princeton University Press, 1970.

Saharon Rosset, Ji Zhu, and Trevor Hastie. Boosting as a regularized path to a maximum margin classifier. *Journal of Machine Learning Research*, 5:941–973, 2004.

Walter Rudin. *Functional Analysis*. International series in pure and applied mathematics. McGraw-Hill, 1991.

Arda Sahiner, Tolga Ergen, John M. Pauly, and Mert Pilanci. Vector-output ReLU neural network problems are copositive programs: Convex analysis of two layer networks and polynomial-time algorithms. In *International Conference on Learning Representations*, 2021.

Pedro Savarese, Itay Evron, Daniel Soudry, and Nathan Srebro. How do infinite width bounded norm networks look in function space? In *Conference on Learning Theory*, pages 2667–2690. PMLR, 2019.

Joseph Shenouda and Waheed U. Bajwa. A guide to computational reproducibility in signal processing and machine learning [tips & tricks]. *IEEE Signal Processing Magazine*, 40(2):141–151, 2023.

Joseph Shenouda, Yamin Zhou, and Robert D. Nowak. ReLUs are sufficient for learning implicit neural representations. In *International Conference on Machine Learning*, 2024.

Jonathan W. Siegel and Jinchao Xu. Characterization of the variation spaces corresponding to shallow neural networks. *Constructive Approximation*, pages 1–24, 2023.

Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *CoRR*, abs/1409.1556, 2014.

Ivan Singer. Linear functionals on the space of continuous mappings of a compact Hausdorff space into a Banach spaces. *Rev. Math. Pures Appl.*, 2:301–315, 1957.

Ivan Singer. Sur les applications linéaires intégrales des espaces de fonctions continues. i. *Rev. Math. Pures Appl.*, 4(1959):391–401, 1959.

Len Spek, Tjeerd Jan Heeringa, and Christoph Brune. Duality for neural networks through reproducing kernel Banach spaces. *arXiv preprint arXiv:2211.05020*, 2022.

Liang Sun, Jun Liu, Jianhui Chen, and Jieping Ye. Efficient recovery of jointly sparse vectors. *Advances in Neural Information Processing Systems*, 22, 2009.

Ryan J. Tibshirani. The lasso problem and uniqueness. *Electronic Journal of Statistics*, 7: 1456–1490, 2013.

Michael Unser. A unifying representer theorem for inverse problems and machine learning. *Foundations of Computational Mathematics*, 21(4):941–960, 2021.

Roger Waleffe and Theodoros Rekatsinas. Principal component networks: Parameter reduction early in training. *Hardware Aware Efficient Training (HAET) Workshop at the 39th ICML*, 2022.

Hongyi Wang, Saurabh Agarwal, and Dimitris Papailiopoulos. Pufferfish: Communication-efficient models at no extra cost. *Proceedings of Machine Learning and Systems*, 3:365–386, 2021.

Rui Wang, Yuesheng Xu, and Mingsong Yan. Hypothesis spaces for deep learning. *arXiv preprint arXiv:2403.03353*, 2024a.

Rui Wang, Yuesheng Xu, and Mingsong Yan. Sparse representer theorems for learning in reproducing kernel Banach spaces. *Journal of Machine Learning Research*, 25(93):1–45, 2024b.

Dirk Werner. Extreme points in spaces of operators and vector-valued measures. *Proceedings of the 12th Winter School on Abstract Analysis*, pages 135–143, 1984.

Liu Yang, Jifan Zhang, Joseph Shenouda, Dimitris Papailiopoulos, Kangwook Lee, and Robert D. Nowak. A better way to decay: Proximal gradient training algorithms for neural nets. In *OPT 2022: Optimization for Machine Learning (NeurIPS Workshop)*, 2022.

Can Yaras, Peng Wang, Laura Balzano, and Qing Qu. Compressible dynamics in deep over-parameterized low-rank learning & adaptation. In *International Conference on Machine Learning*, 2024.

Haizhang Zhang and Jun Zhang. Vector-valued reproducing kernel Banach spaces with applications to multi-task learning. *Journal of Complexity*, 29(2):195–215, 2013.

Haizhang Zhang, Yuesheng Xu, and Jun Zhang. Reproducing kernel Banach spaces for machine learning. *Journal of Machine Learning Research*, 10(12), 2009.

S. Zuhovicki. Remarks on problems in approximation theory. *Mat. Zbirnik KDU*, pages 169–183, 1948.