
How do Minimum-Norm Shallow Denoisers Look in Function Space?

Chen Zeno

Electrical and Computer Engineering
Technion

Greg Ongie

Department Mathematical and Statistical Sciences
Marquette University

Yaniv Blumenfeld, Nir Weinberger, Daniel Soudry

Electrical and Computer Engineering
Technion

{chenzeno, yanivbl}@campus.technion.ac.il, gregory.ongie@marquette.edu
nirwein@technion.ac.il, daniel.soudry@gmail.com

Abstract

Neural network (NN) denoisers are an essential building block in many common tasks, ranging from image reconstruction to image generation. However, the success of these models is not well understood from a theoretical perspective. In this paper, we aim to characterize the functions realized by shallow ReLU NN denoisers — in the common theoretical setting of interpolation (i.e., zero training loss) with a minimal representation cost (i.e., minimal ℓ^2 norm weights). First, for univariate data, we derive a closed form for the NN denoiser function, find it is contractive toward the clean data points, and prove it generalizes better than the empirical MMSE estimator at a low noise level. Next, for multivariate data, we find the NN denoiser functions in a closed form under various geometric assumptions on the training data: data contained in a low-dimensional subspace, data contained in a union of one-sided rays, or several types of simplexes. These functions decompose into a sum of simple rank-one piecewise linear interpolations aligned with edges and/or faces connecting training samples. We empirically verify this alignment phenomenon on synthetic data and real images.

1 Introduction

The ability to reconstruct an image from a noisy observation has been studied extensively in the last decades, as it is useful for many practical applications (e.g., Hasinoff et al. [2010]). In recent years, Neural Network (NN) denoisers commonly replace classical expert-based approaches as they achieve substantially better results than the classical approaches (e.g., Zhang et al. [2017]). Beyond this natural usage, NN denoisers also serve as essential building blocks in a variety of common computer vision tasks, such as solving inverse problems [Zhang et al., 2021] and image generation [Song and Ermon, 2019, Ho et al., 2020]. To better understand the role of NN denoisers in such complex applications, we first wish to theoretically understand the type of solutions they converge to.

In practice, when training denoisers, we sample multiple noisy samples for each clean image and minimize the Mean Squared Error (MSE) loss for recovering the clean image. Since we sample numerous noisy samples per clean sample, the number of training samples is typically larger than the number of parameters in the network. Interestingly, even in such an under-parameterized regime, the loss has multiple global minima corresponding to distinct denoiser functions which achieve zero

loss on the observed data. To characterize these functions, we study, similarly to previous works [Savarese et al., 2019, Ongie et al., 2020], the shallow NN solutions that interpolate the training data with minimal representation cost, i.e., where the ℓ^2 -norm of the weights (without biases and skip connections) is as small as possible. This is because we converge to such *min-cost* solutions when we minimize the loss with a vanishingly small ℓ^2 regularization on these weights.

We first examine the univariate input case: building on existing results [Hanin, 2021], we characterize the min-cost interpolating solution and its generalization to unseen data. Next, we aim to extend this analysis to the multivariate case. However, this is challenging, since, to the best of our knowledge, there are no results that explicitly characterize these min-cost solutions for general multivariate shallow NNs — except in two basic cases. In the first case, the input data is co-linear [Ergen and Pilanci, 2021]. In the second case, the input samples are identical to their target outputs, so the trivial min-cost solution is the identity function. The NN denoisers’ training regime is ‘near’ the second case: there, the input samples are noisy versions of the clean target outputs. Interestingly, we find that this regime leads to non-trivial min-cost solutions far from identity — even with an infinitesimally small input noise. We analytically investigate these solutions here.

Our Contributions. We study the NN solutions in the setting of interpolation of noisy samples with min-cost, in a practically relevant “low noise regime” where the noisy samples are well clustered. In the univariate case,

- We find a closed-form solution for the minimum representation cost NN denoiser. Then, we prove this solution generalizes better than the empirical minimum MSE (eMMSE) denoiser.
- We prove this min-cost NN solution is contractive toward the clean data points, that is, applying the denoiser necessarily reduces the distance of a noisy sample to one of the clean samples.

In the multivariate case,

- We derive a closed-form solution for the min-cost NN denoiser in multivariate case under various assumptions on the geometric configuration of the clean training samples. To the best of our knowledge, this is the first set of results to explicitly characterize a min-cost interpolating NN in a non-basic multivariate setting.
- We illustrate a general alignment phenomenon of min-cost NN denoisers in the multivariate setting: the optimal NN denoiser decomposes into a sum of simple rank-one piecewise linear interpolations aligned with edges and/or faces connecting clean training samples.

2 Preliminaries and problem setting

The denoising problem. Let $\mathbf{y} \in \mathbb{R}^d$ be a noisy observation of $\mathbf{x} \in \mathbb{R}^d$, such that $\mathbf{y} = \mathbf{x} + \epsilon$ where \mathbf{x} and ϵ are statistically independent, and $\mathbb{E}[\epsilon] = \mathbf{0}$. Commonly, this noise is Gaussian with covariance matrix $\sigma^2\mathbf{I}$. The ultimate goal of a denoiser $\hat{\mathbf{x}}(\mathbf{y})$ is to minimize the MSE loss over the joint probability distribution of the data and the noisy observation (“population distribution”), i.e., to minimize

$$\mathcal{L}(\hat{\mathbf{x}}) = \mathbb{E}_{\mathbf{x}, \mathbf{y}} \|\hat{\mathbf{x}}(\mathbf{y}) - \mathbf{x}\|^2. \quad (1)$$

The well-known optimal solution for (1) is the minimum mean square error (MMSE) denoiser, i.e.,

$$\hat{\mathbf{x}}^*(\mathbf{y}) = \mathbb{E}_{\mathbf{x}|\mathbf{y}}[\mathbf{x} | \mathbf{y}] \in \arg \min_{\hat{\mathbf{x}}(\mathbf{y})} \mathbb{E}_{\mathbf{x}, \mathbf{y}} \|\mathbf{x} - \hat{\mathbf{x}}(\mathbf{y})\|^2. \quad (2)$$

Since we do not have access to the distribution of the data, and hence not to the posterior distribution, we rely on a finite amount of clean data $\{\mathbf{x}_n\}_{n=1}^N$ in order to learn a good approximation for the MMSE estimator. One approach is to assume an empirical data distribution and derive the optimal solution of (1), i.e., the empirical minimum mean square error (eMMSE) denoiser,

$$\hat{\mathbf{x}}^{\text{eMMSE}}(\mathbf{y}) \in \arg \min_{\hat{\mathbf{x}}(\mathbf{y})} \frac{1}{N} \sum_{n=1}^N \mathbb{E}_{\mathbf{y}|\mathbf{x}_n} \|\hat{\mathbf{x}}(\mathbf{y}) - \mathbf{x}_n\|^2. \quad (3)$$

If the noise is Gaussian with a covariance of $\sigma^2\mathbf{I}$, an explicit solution to the eMMSE is given by

$$\hat{\mathbf{x}}^{\text{eMMSE}}(\mathbf{y}) = \frac{\sum_{n=1}^N \mathbf{x}_n \exp\left(-\frac{\|\mathbf{y} - \mathbf{x}_n\|^2}{2\sigma^2}\right)}{\sum_{n=1}^N \exp\left(-\frac{\|\mathbf{y} - \mathbf{x}_n\|^2}{2\sigma^2}\right)}. \quad (4)$$

An alternative approach to computing the eMMSE directly is to draw M noisy samples for each clean data point, as $\mathbf{y}_{n,m} = \mathbf{x}_n + \boldsymbol{\epsilon}_{n,m}$, where $\boldsymbol{\epsilon}_{n,m} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$ are independent and identically distributed, and to minimize the following loss function

$$\mathcal{L}_{\text{offline},M}(\hat{\mathbf{x}}) = \frac{1}{MN} \sum_{m=1}^M \sum_{n=1}^N \|\hat{\mathbf{x}}(\mathbf{y}_{n,m}) - \mathbf{x}_n\|^2. \quad (5)$$

Denoiser model and algorithms. In practice, we approximate the optimal denoiser $\hat{\mathbf{x}}(\mathbf{y})$ using a parametric model $\mathbf{h}_{\boldsymbol{\theta}}(\mathbf{y})$, typically a NN. We focus on a shallow ReLU network model with a skip connection of the form

$$\mathbf{h}_{\boldsymbol{\theta}}(\mathbf{y}) = \sum_{k=1}^K \mathbf{a}_k [\mathbf{w}_k^\top \mathbf{y} + b_k]_+ + \mathbf{V} \mathbf{y} + \mathbf{c} \quad (6)$$

where $\boldsymbol{\theta} = ((\theta_k)_{k=1}^K; \mathbf{c}, \mathbf{V})$ with $\theta_k = (b_k, \mathbf{a}_k, \mathbf{w}_k) \in \mathbb{R} \times \mathbb{R}^d \times \mathbb{R}^d$ and $\mathbf{c} \in \mathbb{R}^d, \mathbf{V} \in \mathbb{R}^{d \times d}$. We train the model on a finite set of clean data points $\{\mathbf{x}_n\}_{n=1}^N$. The common practical training method is based on an online approach. First, we sample a random batch (with replacement) from the data $\mathcal{B} \subseteq \{\mathbf{x}_n\}_{n=1}^N$. Then, for each clean data point $\mathbf{x}_n \in \mathcal{B}$, we draw a noisy sample $\mathbf{y}_n = \mathbf{x}_n + \boldsymbol{\epsilon}_n$, where $\boldsymbol{\epsilon}_n \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$ are independent of the clean data points and other noise samples. At each iteration t out of T iterations, we update the model parameters according to a stochastic gradient descent rule, with a vanishingly small regularization term $\lambda C(\boldsymbol{\theta})$, that is,

$$\boldsymbol{\theta}_{t+1} = \boldsymbol{\theta}_t - \eta \nabla_{\boldsymbol{\theta}_t} \frac{1}{|\mathcal{B}|} \sum_{n \in \mathcal{B}} \|\mathbf{h}_{\boldsymbol{\theta}_t}(\mathbf{y}_n) - \mathbf{x}_n\|^2 - \eta \lambda \nabla_{\boldsymbol{\theta}_t} C(\boldsymbol{\theta}_t). \quad (7)$$

Another training method [Chen et al., 2014] is based on an offline approach. We sample M noisy sample for each clean data point and minimize (5) plus a regularization term

$$\mathcal{L}_{\text{offline},M}(\boldsymbol{\theta}) = \frac{1}{MN} \sum_{m=1}^M \sum_{n=1}^N \|\mathbf{h}_{\boldsymbol{\theta}}(\mathbf{y}_{n,m}) - \mathbf{x}_n\|^2 + \lambda C(\boldsymbol{\theta}). \quad (8)$$

Similarly to previous works [Savarese et al., 2019, Ongie et al., 2020], we assume an ℓ^2 penalty on the weights, but not on the biases and skip connections, i.e.,

$$C(\boldsymbol{\theta}) = \frac{1}{2} \sum_{k=1}^K (\|\mathbf{a}_k\|^2 + \|\mathbf{w}_k\|^2). \quad (9)$$

Low noise regime. In this paper, we study the solution of the NN denoiser when the clusters of noisy samples around each clean point are well-separated, a setting which we refer to as the “low noise regime”. This is a rather relevant regime since denoisers are practically used when the noise level is mild. Indeed, common image-denoising benchmarks test on low (but not negligible) noise levels. For instance, in the commonly used denoising benchmark BSD68 [Roth and Black, 2009], the noise level $\sigma = 0.1$ is in the low noise regime.¹ Moreover, this setting is important, for example, in diffusion-based image generation, since at the end of the reverse denoising process, new images are sampled by denoising smaller and smaller noise levels.²

3 Basic properties of neural network denoisers

Offline v.s. online NN solutions. NN denoisers are traditionally trained in an online fashion (7), using a finite amount of T iterations. Consequently, only a finite number of noisy samples are used for each clean data point. We empirically observe that the solutions in the offline and online settings are similar. Specifically, in the univariate case, we show in Figure 1 that denoisers based on offline and

¹The minimum distance between two images in BSD68 is about 97 while the image resolution is $d = 481 \times 321$. Also, the norm of the noise concentrates around the value of $\sqrt{d}\sigma \approx \sqrt{481 \cdot 321} \cdot 0.1 \approx 40 < 97$. Therefore, the clusters of noisy samples around each clean point are generally well-separated.

²Interestingly, it was suggested that the “useful” part of the diffusion dynamics happens only below some critical noise level [Raya and Ambrogioni, 2023].




Figure 1: NN denoiser vs eMMSE denoiser. We trained a one-hidden-layer ReLU network with a skip connection on a denoising task. The clean dataset has four points equally spaced in the interval $[-5, 5]$, and the noisy samples are generated by adding zero-mean Gaussian noise with $\sigma = 1.5$. We use $\lambda = 10^{-5}$ in both setting. The Figure shows the denoiser output as a function of its input for: (1) NN denoiser trained online using (7) for 100K iterations, (2) NN denoiser trained offline using (8) with $M = 9000$ and 20K epochs, and (3) the eMMSE denoiser (4).

online loss functions converge to indistinguishable solutions. For the multivariate case, we observe (Figure 11 in Appendix D) that the offline and online solutions achieve approximately the same test MSE when trained on a subset of the MNIST dataset. The comparison is made using the same number of iterations for both training methods, while using much less noisy samples in the offline setting. Evidently, this lower number of samples does not significantly affect the generalization error. Hence, in the rest of the paper, we focus on offline training (i.e., minimizing the offline loss $\mathcal{L}_{\text{offline},M}$), as it defines an explicit loss function with solutions that can be theoretically analyzed, as in [Savarese et al., 2019, Ongie et al., 2020].

The empirical MMSE denoiser. The law of large numbers implies that the denoiser minimizing the offline loss $\mathcal{L}_{\text{offline},M}$ approaches the eMMSE estimator in the limit of infinitely many noisy samples,

$$\hat{\mathbf{x}}^{\text{eMMSE}}(\mathbf{y}) \in \arg \min_{\hat{\mathbf{x}}} \lim_{M \rightarrow \infty} \mathcal{L}_{\text{offline},M}(\hat{\mathbf{x}}). \quad (10)$$

Therefore, it may seem that for a reasonable number of noise samples M , a large enough model, and small enough regularization, the denoiser we get by minimizing the offline loss (6) will also be similar to the eMMSE estimator. However, Figure 1 shows that the eMMSE solution and the NN solutions (both online and offline) are quite different. The eMMSE denoiser has a much sharper transition and maps almost all inputs to a value of a clean data point. This is because in the case of low noise the eMMSE denoiser (4) approximates the one nearest-neighbor (1-NN) classifier, i.e.,

$$\lim_{\sigma \rightarrow 0^+} \hat{\mathbf{x}}^{\text{eMMSE}}(\mathbf{y}) = \arg \min_{\mathbf{x} \in \{\mathbf{x}_i\}_{i=1}^N} \|\mathbf{y} - \mathbf{x}\|. \quad (11)$$

In contrast, the NN denoiser maps each noisy sample to its corresponding clean sample only in a limited “noise ball” around the clean point, and interpolates near-linearly between the “noise balls”. Hence, we may expect that the smoother NN denoiser typically generalizes better than the eMMSE denoiser. We prove that this is indeed true for the univariate case in Section 4.

Why the NN denoiser does not converge to the eMMSE denoiser? Note that the limit in (10) is not practically relevant for the low-level noise regime, since we need an exponentially large M in order to converge in this limit. For example, in the case of univariate Gaussian noise, we have that $P(|\epsilon| > t) \leq 2 \exp(-\frac{t^2}{2\sigma^2})$, $\forall t > 0$. Therefore, during training, we effectively observe only noisy samples that are in a bounded interval of size $2\sigma\sqrt{\log M}$ around each clean sample (see Figure 1). In other words, in the low-noise regime and for non-exponential M , there is no way to distinguish if the noise is sampled from some distribution with limited support or from a Gaussian distribution. The denoiser minimizing the loss with respect to a bounded-support distribution can be radically different from the eMMSE denoiser in the regions outside the “noise balls” surrounding the clean samples, where the denoiser function is not constrained by the loss. This leads to a large difference between the NN denoiser and the MMSE estimator.

Alternatively, one may suggest that the NN denoiser does not converge to the eMMSE denoiser due to an approximation error (i.e., the shallow NN’s model capacity is too small to approximate the MMSE denoiser). Nevertheless, we provide empirical evidence indicating it is not the case. Specifically, recall that in the low noise regime, the eMMSE denoiser tends to the nearest-neighbor classifier, and

such a solution does not generalize well to test data. Thus, if the NN denoiser would have converged to the eMMSE solution, then its test error would have increased with the network size, in contrast to what we observe in Figure 12 (Appendix D).

Therefore, in order to approximate the eMMSE with a NN, it seems we must have an exponentially large M . Alternatively, we may converge to the eMMSE if we use a loss function marginalized over the Gaussian noise. This idea was previously suggested by Chen et al. [2014], with the goal of effectively increasing the number of noisy samples and thus improving the training performance of denoising autoencoders. Therein, this improvement was obtained by approximating the marginalized loss function by a Taylor series expansion. However, for shallow denoisers, we may actually obtain an explicit expression for this marginalized loss, without any approximation. Specifically, if we assume for simplicity, that the network does not have a linear unit ($\mathbf{V} = \mathbf{0}$) and its bias terms are zero ($\mathbf{c} = \mathbf{0}, b_k = 0$), then the marginalized loss for Gaussian noise, derived in Appendix A, is given by

$$\begin{aligned} \mathcal{L}(\boldsymbol{\theta}, \sigma) &= \mathbb{E}_{\mathbf{x}, \mathbf{y}} \|\mathbf{h}_{\boldsymbol{\theta}}(\mathbf{y}) - \mathbf{x}\|^2 = \mathbb{E}_{\mathbf{x}, \mathbf{y}} \left\| \sum_{k=1}^K \mathbf{a}_k [\mathbf{w}_k^\top \mathbf{y}]_+ - \mathbf{x} \right\|^2 = \mathbb{E}_{\mathbf{x}} \mathbb{E}_{\mathbf{y}|\mathbf{x}} \left\| \sum_{k=1}^K \mathbf{a}_k [\mathbf{w}_k^\top \mathbf{y}]_+ - \mathbf{x} \right\|^2 \\ &= \mathbb{E}_{\mathbf{x}} \left\| \sum_{k=1}^K \mathbf{a}_k \tilde{\phi}(\hat{\mathbf{w}}_k^\top \mathbf{x}, \|\mathbf{w}_k\|, \sigma) - \mathbf{x} \right\|^2 + \sum_{i=1}^K \sum_{j=1}^K \mathbf{a}_i^\top \mathbf{a}_j \mathbf{H}_{ij}(\mathbf{w}_i, \mathbf{w}_j, \sigma^2), \end{aligned} \quad (12)$$

where $\mathbf{w}_k = \hat{\mathbf{w}}_k \|\mathbf{w}_k\|$ and $\mathbf{H} \succeq 0$, $\tilde{\phi}(\cdot)$ are defined in Appendix A. NN denoisers trained over this loss function will thus tend to the eMMSE solution as the network size is increased. However, as we explained above, this is not necessarily desirable, so we only mention (12) to show exact marginalization is feasible.

Regularization biases toward specific neural network denoisers. To further explore the converged solution for offline training, we note that the offline loss function $\mathcal{L}_{\text{offline}, M}(\boldsymbol{\theta})$ allows the network to converge to a zero-loss solution. This is in contrast to online training for which each batch leads to new realizations of noisy samples, and thus the training error is never exactly zero. Specifically, consider the low noise regime (well-separated noisy clusters). Then, the network can perfectly fit all the noisy samples using a finite number of neurons (see Section 4 for a more accurate description in the univariate case). Importantly, there are multiple ways to cluster the noisy data points with such neurons, and so there are multiple global training loss minima that the network can achieve with zero loss, each with a different generalization capability. In contrast to the standard case considered in the literature, this holds even in the under-parameterized case (where NM , the total number of noisy samples, is larger than the number of parameters).

Since there are many minima that perfectly fit the training data, we converge to specific minima which also minimize the ℓ^2 regularization we use (even though we assumed it is vanishing). Specifically, in the limit of vanishing regularization $C(\boldsymbol{\theta})$, the minimizers of $\mathcal{L}_{\text{offline}, M}(\boldsymbol{\theta})$ also minimize the representation cost.

Definition 1. Let $\mathbf{h}_{\boldsymbol{\theta}} : \mathbb{R}^d \rightarrow \mathbb{R}^d$ denote a shallow ReLU network of the form (6). For any function $\mathbf{f} : \mathbb{R}^d \rightarrow \mathbb{R}^d$ realizable as a shallow ReLU network, we define its **representation cost** as

$$R(\mathbf{f}) = \inf_{\boldsymbol{\theta}: \mathbf{f} = \mathbf{h}_{\boldsymbol{\theta}}} C(\boldsymbol{\theta}) = \inf_{\boldsymbol{\theta}} \sum_{k=1}^K \|\mathbf{a}_k\| \quad \text{s.t.} \quad \|\mathbf{w}_k\| = 1 \quad \forall k, \mathbf{f} = \mathbf{h}_{\boldsymbol{\theta}}, \quad (13)$$

and a **minimizer** of this cost, i.e. the ‘min-cost’ solution, as

$$\mathbf{f}^* \in \underset{\mathbf{f}}{\text{argmin}} R(\mathbf{f}) \quad \text{s.t.} \quad \mathbf{f}(\mathbf{y}_{n,m}) = \mathbf{x}_n \quad \forall n, m, \quad (14)$$

where the second equality in (13) holds due to the 1-homogeneity of the ReLU activation function, and since the bias terms are not regularized (see [Savarese et al., 2019, Appendix A]). In the next sections, we examine which function we obtain by minimizing the representation cost $R(\mathbf{f})$ in various settings.

4 Closed form solution for the NN denoiser function — univariate data

In this section, we prove that NN denoisers that minimize $R(\mathbf{f})$ for univariate data have the specific piecewise linear form observed in Figure 1, and we discuss the properties of this form. We observe N

clean univariate data points $\{x_n\}_{n=1}^N$, s.t. $-\infty < x_1 < x_2 < \dots < x_N < \infty$, and M noisy samples (drawn from some known distribution) for each clean data point, such that $y_{n,m} = x_n + \epsilon_{n,m}$. We denote by ϵ_n^{\max} the maximal noise seen for data point x_n , and by ϵ_n^{\min} the minimal noise seen for data point x_n , i.e.,

$$\epsilon_n^{\max} \equiv \max_m \epsilon_{n,m}, \quad \epsilon_n^{\min} \equiv \min_m \epsilon_{n,m}, \quad (15)$$

and assume the following,

Assumption 1. Assume the data $\{x_n\}_{n=1}^N$ is well-separated after the addition of noise, i.e.,

$$\forall n \in [N - 1] : x_n + \epsilon_n^{\max} < x_{n+1} + \epsilon_{n+1}^{\min}, \quad (16)$$

and $\epsilon_n^{\max} > 0, \epsilon_n^{\min} < 0$.

So we can state the following,

Proposition 1. For all datasets such that Assumption 1 holds, the unique minimizer of $R(f)$ is

$$f_{1D}^*(y) = \begin{cases} x_1, & y < x_1 + \epsilon_1^{\min} \\ x_n, & x_n + \epsilon_n^{\min} \leq y \leq x_n + \epsilon_n^{\max} \\ \frac{x_{n+1} - x_n}{x_{n+1} + \epsilon_{n+1}^{\min} - (x_n + \epsilon_n^{\max})} (y - (x_n + \epsilon_n^{\max})) + x_n, & x_n + \epsilon_n^{\max} < y < x_{n+1} + \epsilon_{n+1}^{\min} \\ x_N, & y > x_N + \epsilon_N^{\max} \end{cases} \quad (17)$$

The proof (which is based on Theorem 1.2. in [Hanin, 2021]) can be found in Appendix B.1. As can be seen from Figure 1, the empirical simulation matches Proposition 1.³ Proposition 1 states that (17) is a closed-form solution for (8) with minimal representation cost. Notice that the *minimal* number of neurons needed to represent f_{1D}^* using $h_\theta(y)$ is $2N - 2$, which is less than the number of the total training samples NM for $M \geq 2$.

In the case of univariate data, we can prove that the representation cost minimizer f_{1D}^* (linear interpolation) generalizes better than the optimal estimator over the empirical distribution (eMMSE) for low noise levels.

Theorem 1. Let $y = x + \epsilon$ where $x \sim p_x(x)$ and $\epsilon \sim \mathcal{N}(0, \sigma^2)$ where x and ϵ are statistically independent. Then for all datasets such that Assumption 1 holds, and for all density probability distributions $p_x(x)$ with bounded second moment such that $p_x(x) > 0$ for all $x \in [\min_n x_n, \max_n x_n]$, the following holds,

$$\lim_{\sigma \rightarrow 0^+} \text{MSE}(\hat{x}^{\text{eMMSE}}(y)) > \lim_{\sigma \rightarrow 0^+} \text{MSE}(f_{1D}^*(y)).$$

See Appendix B.2 for the proof. We may deduce from Theorem 1 that for each density probability distribution $p(x)$ there exists a critical noise level for which the the representation cost minimizer f_{1D}^* has strictly lower MSE than the eMMSE for all smaller noise levels (this is because the MSE is a continuous function of σ). The critical noise level can change significantly depending on $p(x)$. For example, if $p(x)$ has a high “mass” in between the training points then the critical noise level is large. However, if the density function has a low “mass” between the training points then the critical noise level is small. In Appendix D we show the MSE vs. the noise level on MNIST denoiser for NN denoiser and eMMSE denoiser (Figure 13). As can be seen there, the critical noise level in this case is not small (~ 5).

Intuitively, the difference between the NN denoiser and the eMMSE denoiser is how they operate on inputs that are not close to any of the clean samples (compared to the noise standard deviation). For such a point, the eMMSE denoiser does not take into account that the empirical distribution of the clean samples does not approximate well their true distribution. Thus, for small noise, it insists on “assigning” it to the closest clean sample point. By contrast, the NN denoiser generalizes better since it takes into account that, far from the clean samples, the data distribution is not well approximated by the empirical sample distribution. Thus, its operation there is near the identity function, with a small contraction toward the clean points, as we discuss next.

³Notice that the training points in Figure 1 are used in the online setting (7) and in the offline setting (8) we observe less noisy samples.

Minimal norm leads to contractive solutions on univariate data. Radhakrishnan et al. [2018] have empirically shown that Auto-Encoders (AE, i.e. NN denoisers without input noise), are *locally* contractive toward the training samples. Specifically, they showed that the clean dataset can be recovered when iterating the AE output multiple times until convergence. Additionally, they showed that, as we increase the width or the depth of the NN, the network becomes more contractive toward the training examples. In addition, Radhakrishnan et al. [2018] proved that 2-layer AE models are locally contractive under strong assumptions (the weights of the input layer are fixed and the number of neurons goes to infinity). Next, we prove that a univariate shallow NN denoiser is *globally* contractive toward the clean data points without using the assumptions used by Radhakrishnan et al. [2018] (i.e., the minimizer optimizes over both layers and has a finite number of neurons).

Definition 2. We say that $\mathbf{f} : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is contractive toward a set of points $\{\mathbf{x}_n\}_{n=1}^N$ on $\mathcal{Y} \subseteq \mathbb{R}^d$ if there exists a real number $0 \leq \alpha < 1$ such that for any $\mathbf{y} \in \mathcal{Y}$ there exists $i \in [N]$ so that

$$\|\mathbf{f}(\mathbf{y}) - \mathbf{f}(\mathbf{x}_i)\| \leq \alpha \|\mathbf{y} - \mathbf{x}_i\|. \quad (18)$$

Lemma 1. $f_{1D}^*(\mathbf{y})$ is contractive toward the clean training points $\{\mathbf{x}_n\}_{n=1}^N$ on $\mathcal{Y} = \mathbb{R} \setminus \cup_{n \in [N-1]} \left\{ \frac{x_{n+1}\epsilon_n^{\max} - x_n\epsilon_{n+1}^{\min}}{\epsilon_n^{\max} - \epsilon_{n+1}^{\min}} \right\}$.

The proof can be found in Appendix B.3.

5 Minimal norm leads to alignment phenomenon on multivariate data

In the multivariate case, min-cost solutions of (14) are difficult to explicitly characterize. Even in the setting of fitting scalar-valued shallow ReLU networks, explicitly characterizing min-cost solutions under interpolation constraints remains an open problem, except in some basic cases (e.g., co-linear data [Ergen and Pilanci, 2021]).

As an approximation to (14) that is more mathematically tractable, we assume the functions being fit are constant and equal to \mathbf{x}_n on a closed ball of radius ρ centered at each \mathbf{x}_n , i.e., $\mathbf{f}(\mathbf{y}) = \mathbf{x}_n$ for all $\|\mathbf{y} - \mathbf{x}_n\| \leq \rho$, such that the balls do not overlap. Letting $B(\mathbf{x}_n, \rho)$ denote the ball of radius ρ centered at \mathbf{x}_n , we can write this constraint more compactly as $\mathbf{f}(B(\mathbf{x}_n, \rho)) = \{\mathbf{x}_n\}$. Consider minimizing the representation cost under this constraint:

$$\min_{\mathbf{f}} R(\mathbf{f}) \quad \text{s.t.} \quad \mathbf{f}(B(\mathbf{x}_n, \rho)) = \{\mathbf{x}_n\} \quad \forall n \in [N]. \quad (19)$$

However, even with this approximation, explicitly describing minimizers of (19) for an arbitrary collection of training samples remains challenging. Instead, to gain intuition, we describe minimizers of (19) assuming the training samples belong to simple geometric structures that yield explicit solutions. Our results reveal a general alignment phenomenon, such that the weights of the representation cost minimizer align themselves with edges and/or faces connecting data points. We also show that approximate solutions of (14) obtained numerically by training a NN denoiser with weight decay match well with the solutions of (19) having exact closed-form expressions.

5.1 Training data on a subspace

In the event that the clean training samples belong to a subspace, we show the representation cost minimizer depends only on the projection of the inputs onto the subspace containing the training data, and its output is also constrained to this subspace.

Theorem 2. Assume the training samples $\{\mathbf{x}_n\}_{n=1}^N$ belong to a linear subspace $\mathcal{S} \subset \mathbb{R}^d$, and let $\mathbf{P}_{\mathcal{S}} \in \mathbb{R}^{d \times d}$ denote the orthogonal projector onto \mathcal{S} . Then any minimizer \mathbf{f}^* of (19) satisfies $\mathbf{f}^*(\mathbf{y}) = \mathbf{P}_{\mathcal{S}} \mathbf{f}^*(\mathbf{P}_{\mathcal{S}} \mathbf{y})$ for all $\mathbf{y} \in \mathbb{R}^d$.

The proof of this result and all others in this section is given in Appendix C.

Note the assumption that the dataset lies on a subspace is practically relevant, since, in general, large datasets are (approximately) low rank, i.e., lie on a linear subspace [Udell and Townsend, 2019]. In Appendix D we also validated that common image datasets are (approximately) low rank (Table 1).

Specializing to the case co-linear training data (i.e., training samples belonging to a one-dimensional subspace) the min-cost solution is unique and is described by the following corollary:




Figure 2: Predicted (top row) and empirical (bottom row) min-cost NN denoisers for $N = 3$ clean training samples in $d = 2$ dimensions. The empirical NN denoisers were trained with weight decay parameter $\lambda = 10^{-5}$ and $M = 100$ noisy samples. As predicted by our theory, the ReLU boundaries align either perpendicular to the triangle edges in the obtuse case (left panel), or parallel to the triangle edges (right panel).

Corollary 1. Assume the training samples $\{\mathbf{x}_n\}_{n=1}^N$ are co-linear, i.e., $\mathbf{x}_n = c_n \mathbf{u}$ for some scalars $c_1 < c_2 < \dots < c_n$ where $\mathbf{u} \in \mathbb{R}^d$ is a unit-vector. Then the minimizer \mathbf{f}^* of (19) is unique and given by $\mathbf{f}^*(\mathbf{y}) = \mathbf{u}\phi(\mathbf{u}^\top \mathbf{y})$ where $\phi: \mathbb{R} \rightarrow \mathbb{R}$ has the same form as the 1-D minimizer (17) f_{1D}^* with $x_n = c_n$ and $\epsilon_n^{\max} = -\epsilon_n^{\min} = \rho$.

In other words, the min-cost solution has a particularly simple form in this case: $\mathbf{f}^*(\mathbf{y}) = \mathbf{u}\phi(\mathbf{u}^\top \mathbf{y})$, where ϕ is a monotonic piecewise linear function. We call any function of this form a *rank-one piecewise linear interpolator*. Below we show that many other min-cost solutions can be expressed as superpositions of rank-one piecewise linear interpolators.

5.2 Training data on rays

As an extension of the previous setting, we now consider training data belonging to a union of one-sided rays sharing a common origin. Assuming the rays are well-separated (in a sense made precise below) we prove that the representation cost minimizer decomposes into a sum of rank-one piecewise linear interpolators aligned with each ray.

Theorem 3. Suppose the training samples X belong to a union of L rays plus a sample at the origin: $X = \{\mathbf{0}\} \cup \{\mathbf{x}_n^{(1)}\}_{n=1}^{N_1} \cup \dots \cup \{\mathbf{x}_n^{(L)}\}_{n=1}^{N_L}$ where $\mathbf{x}_n^{(\ell)} = c_n^{(\ell)} \mathbf{u}_\ell$ for some unit vector \mathbf{u}_ℓ and constants $0 < c_1^{(\ell)} < c_2^{(\ell)} < \dots < c_{N_\ell}^{(\ell)}$. Assume that the rays make obtuse angles with each other (i.e., $\mathbf{u}_\ell^\top \mathbf{u}_k < 0$ for all $\ell \neq k$). Then the minimizer \mathbf{f}^* of (19) is unique and is given by

$$\mathbf{f}^*(\mathbf{y}) = \mathbf{u}_1 \phi_1(\mathbf{u}_1^\top \mathbf{y}) + \dots + \mathbf{u}_L \phi_L(\mathbf{u}_L^\top \mathbf{y}), \quad (20)$$

where $\phi_\ell: \mathbb{R} \rightarrow \mathbb{R}$ has the form of the 1-D minimizer (17) f_{1D}^* with $x_n = c_n^{(\ell)}$, $\epsilon_n^{\max} = -\epsilon_n^{\min} = \rho$.

Additionally, in the Appendix C.2.1 we show that this min-cost solution is stable with respect to small perturbations of the data. In particular, if the training data is perturbed from the rays, the functional form of the min-cost solution only changes slightly, such that the inner and outer-layer weight vectors align with the line segments connecting consecutive data points.

5.3 Special case: training data forming a simplex

Here, we study the representation cost minimizers for $N \leq d + 1$ training points that form the vertices of a $(N - 1)$ -simplex, i.e., a $(N - 1)$ -dimensional simplex in \mathbb{R}^d (e.g., a 2-simplex is a triangle, a 3-simplex is a tetrahedron, etc.). As we will show, the angles between vertices of the simplex (e.g., an acute versus obtuse triangle in $N = 3$) influences the functional form of the min-cost solution.

Our first result considers one extreme where the simplex has one vertex that makes an obtuse angle with *all* other vertices (e.g., an obtuse triangle for $N = 3$).

Proposition 2. Suppose the convex hull of the training points $\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N\} \subset \mathbb{R}^d$ is a $(N - 1)$ -simplex such that \mathbf{x}_1 forms an obtuse angle with all other vertices, i.e., $(\mathbf{x}_j - \mathbf{x}_1)^\top (\mathbf{x}_i - \mathbf{x}_1) < 0$

for all $i \neq j$ with $i, j > 1$. Then the minimizer \mathbf{f}^* of (19) is unique, and is given by

$$\mathbf{f}^*(\mathbf{y}) = \mathbf{x}_1 + \sum_{n=2}^N \mathbf{u}_n \phi_n(\mathbf{u}_n^\top(\mathbf{y} - \mathbf{x}_1)) \quad (21)$$

where $\mathbf{u}_n = \frac{\mathbf{x}_n - \mathbf{x}_1}{\|\mathbf{x}_n - \mathbf{x}_1\|}$, $\phi_n(t) = s_n([t - a_n]_+ - [t - b_n]_+)$, with $a_n = \rho$, $b_n = \|\mathbf{x}_n - \mathbf{x}_1\| - \rho$, and $s_n = \|\mathbf{x}_n - \mathbf{x}_1\| / (b_n - a_n)$ for all $n = 2, \dots, N$.

This result is essentially a corollary of Theorem 3, since after translating \mathbf{x}_1 to be the origin, the vertices of the simplex belong to rays making obtuse angles with each other, where there is exactly one sample per ray. Details of the proof are given in Appendix C.2.

At the opposite extreme, we consider the case where every vertex of the simplex is acute, meaning for all $n = 1, \dots, N$ we have $(\mathbf{x}_i - \mathbf{x}_n)^\top(\mathbf{x}_j - \mathbf{x}_n) > 0$ for all $i, j \neq n$. In this case, we make following conjecture: the min-cost solution is instead a sum of N rank-one piecewise linear interpolators, each aligned orthogonal to a different $(N - 2)$ -dimensional face of the simplex.

Conjecture 1. *Suppose the convex hull of the training points $\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N\} \subset \mathbb{R}^d$ is a $(N - 1)$ -simplex where every vertex of the simplex is acute. Then the minimizer \mathbf{f}^* of (19) is unique, and is given by*

$$\mathbf{f}^*(\mathbf{y}) = \bar{\mathbf{x}} + \sum_{n=1}^N \mathbf{v}_n \phi_n(\mathbf{u}_n^\top(\mathbf{y} - \mathbf{z}_n)) \quad (22)$$

where: \mathbf{z}_n is the projection of \mathbf{x}_n onto the unique $(N - 2)$ -dimensional face of the simplex not containing \mathbf{x}_n ; $\bar{\mathbf{x}}$ is the weighted geometric median of the vertices specified by

$$\bar{\mathbf{x}} = \arg \min_{\mathbf{x} \in \mathbb{R}^d} \sum_{n=1}^N \frac{\|\mathbf{x}_n - \mathbf{x}\|}{\|\mathbf{x}_n - \mathbf{z}_n\|};$$

$\mathbf{u}_n = \frac{\mathbf{x}_n - \mathbf{z}_n}{\|\mathbf{x}_n - \mathbf{z}_n\|}$, $\mathbf{v}_n = \frac{\mathbf{x}_n - \bar{\mathbf{x}}}{\|\mathbf{x}_n - \bar{\mathbf{x}}\|}$; and $\phi_n(t) = s_n([t - a_n]_+ - [t - b_n]_+)$ with $a_n = \rho$, $b_n = \|\mathbf{x}_n - \mathbf{z}_n\| - \rho$, and $s_n = \|\mathbf{x}_n - \bar{\mathbf{x}}\| / (b_n - a_n)$.

Justification for this conjecture is given in Appendix C.3.2. In particular, we prove that the interpolator \mathbf{f}^* given in (86) is a min-cost solution in the special case of three training points whose convex hull is an equilateral triangle. If true in general, this would imply a phase transition behavior in the min-cost solution when the simplex changes from having one obtuse vertex to all acute vertices, such that ReLU boundaries go from being aligned orthogonal to the edges connecting vertices, to being aligned parallel with the simplex faces. Figure 2 illustrates this for $N = 3$ training points forming a triangle in $d = 2$ dimensions. Moreover, Figure 2 shows that the empirical minimizer obtained using noisy samples and weight decay regularization agrees well with the form of the exact min-cost solution predicted by Proposition 2 and Conjecture 1.

In general, any given vertex of a simplex may make acute angles with some vertices and obtuse angles with others. This case is not covered by the above results. Currently, we do not have a conjectured form of the min-cost solution in this case, and we leave this as an open problem for future work.

6 Related works

Numerous methods have been proposed for image denoising. In last decade NN-based methods achieve state-of-the-art results [Zhang et al., 2017, 2021]. See [Elad et al., 2023] for a comprehensive review of image denoising. Sonthalia and Nadakuditi [2023] empirically showed a double decent behavior in NN denoisers, and theoretically proved it in a linear model. Similar to a denoiser, an Auto-Encoder (AE) is a NN model whose output dimension equals its input dimension, and is trained to match the output to the input. For AE, the typical goal is to learn an efficient lower-dimensional representation of the samples. Radhakrishnan et al. [2018] proved that a single hidden-layer AE that interpolates the training data (i.e., achieves zero loss), projects the input onto a nonlinear span of the training data. In addition, Radhakrishnan et al. [2018] empirically demonstrated that a multi-layer ReLU AE is locally contractive toward the training samples by iterating the AE and showing that the points converge to one of the training samples. Denoising autoencoders inject noise into the input

data in order to learn a good representation [Alain and Bengio, 2014]. The marginalized denoising autoencoder, proposed by Chen et al. [2014], approximates the marginalized loss over the noise (which is equivalent to observing infinitely many noisy samples) by using a Taylor approximation. Chen et al. [2014] demonstrated that by using the approximate marginalized loss we can achieve a substantial speedup in training and improved representation compared to standard denoising AE.

Many recent works aim to characterize function space properties of interpolating NN with minimal representation cost (i.e., min-cost solutions). Building off of the connection between weight decay and path-norm regularization identified in Neyshabur et al. [2015, 2017], Savarese et al. [2019] showed that the representation cost of a function realizable as a univariate two-layer ReLU network coincides with the L^1 -norm of the second derivative of the function. Extensions to the multivariate setting were studied in Ongie et al. [2020], which identified the representation cost of a multivariate function with its R -norm, a Banach space semi-norm defined in terms of the Radon transform. Related work has extended the R -norm to other activation functions Parhi and Nowak [2021], vector-valued networks Shenouda et al. [2023], and deeper architectures Parhi and Nowak [2022]. A separate line of research studies min-cost solutions from a convex duality perspective Ergen and Pilanci [2021], including two-layer CNN denoising AEs Sahiner et al. [2021]. Recent work also studies properties of min-cost solutions in the case of arbitrarily deep NNs with ReLU activation Jacot [2022], Jacot et al. [2022].

Despite these advances in understanding min-cost solutions, there are few results explicitly characterizing their functional form. One exception is Hanin [2021], which gives a complete characterization of min-cost solutions in the case of shallow univariate ReLU networks with unregularized bias. This characterization is possible because the univariate representation cost is defined in terms of the 2nd derivative, which acts locally. Therefore, global minimizers can be found by minimizing the representation cost locally over intervals between data points. An extension of these results to the case of regularized bias is studied in Boursier and Flammarion [2023]. In the multivariate setting, the representation cost involves the Radon transform of the function – a highly non-local operation – that complicates the analysis. Parhi and Nowak [2021] prove a representer theorem showing that there always exists a min-cost solution realizable as a shallow ReLU network with finitely many neurons, and Ergen and Pilanci [2021] give an implicit characterization of min-cost NNs the solution to a convex optimization problem, and give explicit solutions in the case of co-linear training features. However, to the best of our knowledge, there are no results explicitly characterizing min-cost solutions in the case of non-colinear multivariate inputs, even for networks having scalar outputs.

7 Discussions

Conclusions. We have explored the elementary properties of NN solutions for the denoising problem, while focusing on offline training of a one hidden-layer ReLU network. When the noisy clusters of the data samples are well-separated, there are multiple networks with zero loss, even in the case of under-parameterization, while having a different representation cost. In contrast, previous theoretical works focused on the over-parametrized regime. In the univariate case, we have derived a closed-form solution to such global minima with minimum representation cost. We also showed that the univariate NN solution generalizes better than the eMMSE denoiser. In the multivariate case, we showed that the interpolating solution with minimal representation cost is aligned with the edges and/or faces connecting the clean data points in several cases.

Limitations. One limitation of our analysis in the multivariate case is that we assume the denoiser interpolates data on a full d -dimensional ball centered at each clean training sample, where d is the input dimension. In practical settings, often the number of noisy samples $M \ll d$. A more accurate model would be to assume that denoiser interpolates over an $(M - 1)$ -dimensional disc centered at each training sample. This model may still be a tractable alternative to assuming interpolation of finitely many noisy samples. Also, our results relate to NN denoisers trained with explicit weight decay regularization, which is not always used in practice. However, recent work shows that stable minimizers of SGD must have low representation cost Mulayoff et al. [2021], Nacson et al. [2023], and so some of our analysis may provide insight for unregularized training, as well. Finally, for mathematical tractability, we focussed on the case of fully-connected ReLU networks with one hidden-layer. Extending our analysis to deeper architectures and convolutional neural networks is an important direction for future work.

Acknowledgments

We thank Itay Hubara for his technical advice and valuable comments on the manuscript. The research of DS was funded by the European Union (ERC, A-B-C-Deep, 101039436). Views and opinions expressed are however those of the author only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency (ERCEA). Neither the European Union nor the granting authority can be held responsible for them. DS also acknowledges the support of Schmidt Career Advancement Chair in AI. The research of NW was supported by the Israel Science Foundation (ISF), grant no. 1782/22. GO was supported by the National Science Foundation (NSF) CRII award CCF-2153371.

References

- Guillaume Alain and Yoshua Bengio. What regularized auto-encoders learn from the data-generating distribution. *The Journal of Machine Learning Research*, 15(1):3563–3593, 2014.
- Etienne Boursier and Nicolas Flammarion. Penalising the biases in norm regularisation enforces sparsity. *arXiv preprint arXiv:2303.01353*, 2023.
- Minmin Chen, Kilian Weinberger, Fei Sha, and Yoshua Bengio. Marginalized denoising auto-encoders for nonlinear representations. In Eric P. Xing and Tony Jebara, editors, *Proceedings of the 31st International Conference on Machine Learning*, volume 32 of *Proceedings of Machine Learning Research*, pages 1476–1484, Beijing, China, 22–24 Jun 2014. PMLR. URL <https://proceedings.mlr.press/v32/cheng14.html>.
- Michael Elad, Bahjat Kawar, and Gregory Vaksman. Image denoising: The deep learning revolution and beyond—a survey paper. *SIAM Journal on Imaging Sciences*, 16(3):1594–1654, 2023. doi: 10.1137/23M1545859. URL <https://doi.org/10.1137/23M1545859>.
- Tolga Ergen and Mert Pilanci. Convex geometry and duality of over-parameterized neural networks. *The Journal of Machine Learning Research*, 22(1):9646–9708, 2021.
- Boris Hanin. Ridgeless interpolation with shallow relu networks in $1d$ is nearest neighbor curvature extrapolation and provably generalizes on lipschitz functions. *arXiv preprint arXiv:2109.12960*, 2021.
- Samuel W Hasinoff, Frédo Durand, and William T Freeman. Noise-optimal capture for high dynamic range photography. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 553–560. IEEE, 2010.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020.
- William C Horrace. Moments of the truncated normal distribution. *Journal of Productivity Analysis*, 43:133–138, 2015.
- Arthur Jacot. Implicit bias of large depth networks: a notion of rank for nonlinear functions. In *International Conference on Learning Representations (ICLR)*, 2022.
- Arthur Jacot, Eugene Golikov, Clement Hongler, and Franck Gabriel. Feature learning in l_2 -regularized dnns: Attraction/repulsion and sparsity. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 35, pages 6763–6774, 2022.
- B. G. Manjunath and Stefan Wilhelm. Moments calculation for the doubly truncated multivariate normal density. *Journal of Behavioral Data Science*, 1(1):17–33, May 2021. doi: 10.35566/jbds/v1n1/p2. URL <https://jbds.isdsa.org/index.php/jbds/article/view/9>.
- Rotem Mulayoff, Tomer Michaeli, and Daniel Soudry. The implicit bias of minima stability: A view from function space. *Advances in Neural Information Processing Systems (NeurIPS)*, 34: 17749–17761, 2021.

- Mor Shpigel Nacson, Rotem Mulyoff, Greg Ongie, Tomer Michaeli, and Daniel Soudry. The implicit bias of minima stability in multivariate shallow relu networks. In *International Conference on Learning Representations (ICLR)*, 2023.
- Behnam Neyshabur, Ryota Tomioka, and Nathan Srebro. Norm-based capacity control in neural networks. In *Conference on Learning Theory (COLT)*, pages 1376–1401. PMLR, 2015.
- Behnam Neyshabur, Srinadh Bhojanapalli, David McAllester, and Nati Srebro. Exploring generalization in deep learning. *Advances in Neural Information Processing Systems (NeurIPS)*, 30, 2017.
- 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. URL <https://openreview.net/forum?id=H11NPxHKDH>.
- Rahul Parhi and Robert D Nowak. Banach space representer theorems for neural networks and ridge splines. *Journal of Machine Learning Research*, 22(1):1960–1999, 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.
- Adityanarayanan Radhakrishnan, Karren Yang, Mikhail Belkin, and Caroline Uhler. Memorization in overparameterized autoencoders. *arXiv preprint arXiv:1810.10333*, 2018.
- Gabriel Raya and Luca Ambrogioni. Spontaneous symmetry breaking in generative diffusion models. *arXiv preprint arXiv:2305.19693*, 2023.
- Stefan Roth and Michael J. Black. Fields of experts. *International Journal of Computer Vision*, 82(2):205–229, 2009. doi: 10.1007/s11263-008-0197-6. URL <https://doi.org/10.1007/s11263-008-0197-6>.
- Arda Sahiner, Morteza Mardani, Batu Ozturkler, Mert Pilanci, and John M. Pauly. Convex regularization behind neural reconstruction. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=VErQxgyrbfn>.
- 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, Rahul Parhi, Kangwook Lee, and Robert D Nowak. Vector-valued variation spaces and width bounds for DNNs: Insights on weight decay regularization. *arXiv preprint arXiv:2305.16534*, 2023.
- Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems*, 32, 2019.
- Rishi Sonthalia and Raj Rao Nadakuditi. Training data size induced double descent for denoising feedforward neural networks and the role of training noise. *Transactions on Machine Learning Research*, 2023. ISSN 2835-8856. URL <https://openreview.net/forum?id=FdMWtpVT1I>.
- Madeleine Udell and Alex Townsend. Why are big data matrices approximately low rank? *SIAM Journal on Mathematics of Data Science*, 1(1):144–160, 2019.
- Kai Zhang, Wangmeng Zuo, Yunjin Chen, Deyu Meng, and Lei Zhang. Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising. *IEEE transactions on image processing*, 26(7): 3142–3155, 2017.
- Kai Zhang, Yawei Li, Wangmeng Zuo, Lei Zhang, Luc Van Gool, and Radu Timofte. Plug-and-play image restoration with deep denoiser prior. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(10):6360–6376, 2021.

A Marginalized loss

In this section, we derive the marginal loss for the case of a 1-hidden layer ReLU neural network. The loss function is,

$$\begin{aligned}
 \mathcal{L}(\boldsymbol{\theta}, \sigma) &= \mathbb{E}_{\mathbf{x}, \mathbf{y}} \left\| \sum_{k=1}^K \mathbf{a}_k [\mathbf{w}_k^\top \mathbf{y}]_+ - \mathbf{x} \right\|^2 \\
 &= \mathbb{E}_{\mathbf{x}, \mathbf{y}} \left[\sum_{i=1}^K \sum_{j=1}^K \mathbf{a}_i^\top \mathbf{a}_j [\mathbf{w}_i^\top \mathbf{y}]_+ [\mathbf{w}_j^\top \mathbf{y}]_+ + \|\mathbf{x}\|^2 - 2 \sum_{i=1}^K \mathbf{x}^\top \mathbf{a}_i [\mathbf{w}_i^\top \mathbf{y}]_+ \right] \\
 &= \sum_{i=1}^K \sum_{j=1}^K \mathbf{a}_i^\top \mathbf{a}_j \mathbb{E}_{\mathbf{x}, \mathbf{y}} \left[[\mathbf{w}_i^\top \mathbf{y}]_+ [\mathbf{w}_j^\top \mathbf{y}]_+ \right] - 2 \sum_{i=1}^K \mathbb{E}_{\mathbf{x}, \mathbf{y}} \left[\mathbf{x}^\top \mathbf{a}_i [\mathbf{w}_i^\top \mathbf{y}]_+ \right] + \mathbb{E} \|\mathbf{x}\|^2 \\
 &= \sum_{i=1}^K \sum_{j=1}^K \mathbf{a}_i^\top \mathbf{a}_j \mathbb{E}_{\mathbf{x}} \left[\mathbb{E}_{\mathbf{y}|\mathbf{x}} \left[[\mathbf{w}_i^\top \mathbf{y}]_+ [\mathbf{w}_j^\top \mathbf{y}]_+ \right] \right] - 2 \sum_{i=1}^K \mathbb{E}_{\mathbf{x}} \left[\mathbb{E}_{\mathbf{y}|\mathbf{x}} \left[\mathbf{x}^\top \mathbf{a}_i [\mathbf{w}_i^\top \mathbf{y}]_+ \right] \right] + \mathbb{E} \|\mathbf{x}\|^2 \\
 &= \sum_{i=1}^K \sum_{j=1}^K \mathbf{a}_i^\top \mathbf{a}_j \mathbb{E}_{\mathbf{x}} \left[\mathbb{E}_{\mathbf{y}|\mathbf{x}} \left[[\mathbf{w}_i^\top \mathbf{y}]_+ \right] \mathbb{E}_{\mathbf{y}|\mathbf{x}} \left[[\mathbf{w}_j^\top \mathbf{y}]_+ \right] \right] - 2 \sum_{i=1}^K \mathbb{E}_{\mathbf{x}} \left[\mathbb{E}_{\mathbf{y}|\mathbf{x}} \left[\mathbf{x}^\top \mathbf{a}_i [\mathbf{w}_i^\top \mathbf{y}]_+ \right] \right] + \mathbb{E} \|\mathbf{x}\|^2 \\
 &\quad + \sum_{i=1}^K \sum_{j=1}^K \mathbf{a}_i^\top \mathbf{a}_j \mathbb{E}_{\mathbf{x}} \left[\mathbb{E}_{\mathbf{y}|\mathbf{x}} \left[[\mathbf{w}_i^\top \mathbf{y}]_+ [\mathbf{w}_j^\top \mathbf{y}]_+ \right] \right] - \sum_{i=1}^K \sum_{j=1}^K \mathbf{a}_i^\top \mathbf{a}_j \mathbb{E}_{\mathbf{x}} \left[\mathbb{E}_{\mathbf{y}|\mathbf{x}} \left[[\mathbf{w}_i^\top \mathbf{y}]_+ \right] \mathbb{E}_{\mathbf{y}|\mathbf{x}} \left[[\mathbf{w}_j^\top \mathbf{y}]_+ \right] \right] \\
 &= \mathbb{E}_{(\mathbf{x})} \left\| \sum_{i=1}^K \mathbf{a}_i \tilde{\phi}(\mathbf{w}_i^\top \mathbf{x}) - \mathbf{x} \right\|^2 + \sum_{i=1}^K \sum_{j=1}^K \mathbf{a}_i^\top \mathbf{a}_j \mathbb{E}_{\mathbf{x}} \left[\mathbb{E}_{\mathbf{y}|\mathbf{x}} \left[[\mathbf{w}_i^\top \mathbf{y}]_+ [\mathbf{w}_j^\top \mathbf{y}]_+ \right] \right] \\
 &\quad - \sum_{i=1}^K \sum_{j=1}^K \mathbf{a}_i^\top \mathbf{a}_j \mathbb{E}_{\mathbf{x}} \left[\mathbb{E}_{\mathbf{y}|\mathbf{x}} \left[[\mathbf{w}_i^\top \mathbf{y}]_+ \right] \mathbb{E}_{\mathbf{y}|\mathbf{x}} \left[[\mathbf{w}_j^\top \mathbf{y}]_+ \right] \right].
 \end{aligned}$$

We denote by

$$\mathbf{H}_{ij}(\mathbf{w}_i, \mathbf{w}_j, \sigma^2) = \mathbb{E}_{\mathbf{x}} \left[\mathbb{E}_{\mathbf{y}|\mathbf{x}} \left[[\mathbf{w}_i^\top \mathbf{y}]_+ [\mathbf{w}_j^\top \mathbf{y}]_+ \right] - \mathbb{E}_{\mathbf{y}|\mathbf{x}} \left[[\mathbf{w}_i^\top \mathbf{y}]_+ \right] \mathbb{E}_{\mathbf{y}|\mathbf{x}} \left[[\mathbf{w}_j^\top \mathbf{y}]_+ \right] \right].$$

Note that $\mathbf{H} \succeq 0$ since \mathbf{H} is a covariance matrix. Thus we get,

$$\mathcal{L}(\boldsymbol{\theta}, \sigma) = \mathbb{E}_{\mathbf{x}} \left\| \sum_{k=1}^K \mathbf{a}_k \tilde{\phi}(\hat{\mathbf{w}}_k^\top \mathbf{x}, \|\mathbf{w}_k\|, \sigma) - \mathbf{x} \right\|^2 + \sum_{i=1}^K \sum_{j=1}^K \mathbf{a}_i^\top \mathbf{a}_j \mathbf{H}_{ij}(\mathbf{w}_i, \mathbf{w}_j, \sigma^2).$$

Lemma 2. *In the case of the ReLU activation function and Gaussian noise the following holds,*

$$\tilde{\phi}(\hat{\mathbf{w}}_i^\top \mathbf{x}, \|\mathbf{w}_i\|, \sigma) = \|\mathbf{w}_i\| \left(\left(1 - \Phi \left(-\frac{\hat{\mathbf{w}}_i^\top \mathbf{x}}{\sigma} \right) \right) \hat{\mathbf{w}}_i^\top \mathbf{x} + \varphi \left(-\frac{\hat{\mathbf{w}}_i^\top \mathbf{x}}{\sigma} \right) \right)$$

where φ, Φ are the density and cumulative distribution of standard normal distribution, and

$$\mathbf{H}_{ij} = \mathbb{E}_{\mathbf{x}} \left[\Psi(\mathbf{x}, \mathbf{w}_i, \mathbf{w}_j, \sigma^2) - \tilde{\phi}(\hat{\mathbf{w}}_i^\top \mathbf{x}, \|\mathbf{w}_i\|, \sigma) \tilde{\phi}(\hat{\mathbf{w}}_j^\top \mathbf{x}, \|\mathbf{w}_j\|, \sigma) \right]$$

where,

$$\begin{aligned}
 \Psi(\mathbf{x}, \mathbf{w}_i, \mathbf{w}_j, \sigma^2) &= P(z_1 > 0, z_2 > 0; \boldsymbol{\mu}^{(i,j)}, \boldsymbol{\Sigma}^{(i,j)}) \left(\sigma_{12} + \det(\boldsymbol{\Sigma}) \frac{f(-\boldsymbol{\mu}^{(i,j)}; \mathbf{0}, \boldsymbol{\Sigma}^{(i,j)})}{P(z_1 > -\mu_1, z_2 > -\mu_2; \mathbf{0}, \boldsymbol{\Sigma}^{(i,j)})} \right. \\
 &\quad \left. + \sigma_{11} F_1(-\mu_1) \mu_2 + \mu_1 \sigma_{22} F_2(-\mu_2) + \mu_1 \mu_2 \right) \\
 F_2(-\mu_2) &= \frac{\varphi\left(-\frac{\mu_2}{\sqrt{\sigma_{22}}}\right) P\left(z > -\mu_1; \frac{\Lambda_{12}\mu_2}{\Lambda_{11}}, \frac{1}{\Lambda_{11}}\right)}{P(z_1 > -\mu_1, z_2 > -\mu_2; \mathbf{0}, \boldsymbol{\Sigma})} \\
 F_1(-\mu_1) &= \frac{\varphi\left(-\frac{\mu_1}{\sqrt{\sigma_{11}}}\right) P\left(z > -\mu_2; \frac{\Lambda_{21}\mu_1}{\Lambda_{22}}, \frac{1}{\Lambda_{22}}\right)}{P(z_1 > -\mu_1, z_2 > -\mu_2; \mathbf{0}, \boldsymbol{\Sigma})}
 \end{aligned}$$




Figure 3: **Numerical evaluation of (23)** Histogram of the sample average of $\text{ReLU}(x)$ for 10000 Monte-Carlo samples. We denote by E the analytical expectation and by \bar{E} the sample average. Figure (a) is for $\mu = 1, \sigma = 5$, the normalized error is $\frac{|E-\bar{E}|}{E} = 0.0032\%$. Figure (b) is for $\mu = -1, \sigma = 5$, the normalized error is $\frac{|E-\bar{E}|}{E} = 0.0059\%$.

and

$$\begin{aligned}\boldsymbol{\mu}^{(i,j)} &= \begin{pmatrix} \mu_1 \\ \mu_2 \end{pmatrix} = \begin{pmatrix} \mathbf{w}_i^\top \mathbf{x} \\ \mathbf{w}_j^\top \mathbf{x} \end{pmatrix} \\ \boldsymbol{\Sigma}^{(i,j)} &= \begin{pmatrix} \sigma_{11} & \sigma_{12} \\ \sigma_{12} & \sigma_{22} \end{pmatrix} = \begin{pmatrix} \sigma^2 \|\mathbf{w}_i\|^2 & \sigma^2 \mathbf{w}_i^\top \mathbf{w}_j \\ \sigma^2 \mathbf{w}_j^\top \mathbf{w}_i & \sigma^2 \|\mathbf{w}_j\|^2 \end{pmatrix} \\ \boldsymbol{\Lambda} &= \boldsymbol{\Sigma}^{-1}.\end{aligned}$$

Notice that we denote by $P(z_1 > 0, z_2 > 0; \boldsymbol{\mu}^{(i,j)}, \boldsymbol{\Sigma}^{(i,j)})$ the probability that $z_1 > 0, z_2 > 0$ where $(z_1, z_2)^\top \sim \mathcal{N}(\boldsymbol{\mu}^{(i,j)}, \boldsymbol{\Sigma}^{(i,j)})$.

Proof. Let $x \sim \mathcal{N}(\mu, \sigma^2)$, then

$$\begin{aligned}\mathbb{E}[[x]_+] &= E[[x]_+ | x > 0] P(x > 0) + E[[x]_+ | x < 0] P(x < 0) \\ &= E[x | x > 0] P(x > 0) \\ &= E[x | x > 0] \left(1 - \Phi\left(-\frac{\mu}{\sigma}\right)\right)\end{aligned}$$

Note that given $x > 0$ the distribution of x is truncated normal [Horrace, 2015]. Therefore,

$$\begin{aligned}\mathbb{E}[x | x > 0] &= \mu + \sigma \frac{\varphi\left(-\frac{\mu}{\sigma}\right)}{1 - \Phi\left(-\frac{\mu}{\sigma}\right)} \\ \mathbb{E}[[x]_+] &= \left(1 - \Phi\left(-\frac{\mu}{\sigma}\right)\right) \mu + \sigma \varphi\left(-\frac{\mu}{\sigma}\right).\end{aligned}\tag{23}$$

Note that given \mathbf{x} , \mathbf{y} is a Gaussian random vector. Therefore, given \mathbf{x} ,

$$\mathbf{w}_k^\top \mathbf{y} \sim \mathcal{N}\left(\mathbf{w}_k^\top \mathbf{x}, \sigma^2 \|\mathbf{w}_k\|^2\right)$$

and we obtain,

$$\begin{aligned}\mathbb{E}_{\mathbf{y}|\mathbf{x}} [[\mathbf{w}_k^\top \mathbf{y}]_+] &= \|\mathbf{w}_k\| \mathbb{E}_{\mathbf{y}|\mathbf{x}} [[\hat{\mathbf{w}}_k^\top \mathbf{y}]_+] \\ &= \|\mathbf{w}_k\| \left(1 - \Phi\left(-\frac{\hat{\mathbf{w}}_k^\top \mathbf{x}}{\sigma}\right)\right) \hat{\mathbf{w}}_k^\top \mathbf{x} + \sigma \varphi\left(-\frac{\hat{\mathbf{w}}_k^\top \mathbf{x}}{\sigma}\right).\end{aligned}$$

Let $(z_1, z_2)^\top$ be a Gaussian random vector with⁴

$$\begin{aligned}\boldsymbol{\mu} &= \begin{pmatrix} \mu_1 \\ \mu_2 \end{pmatrix} \\ \boldsymbol{\Sigma} &= \begin{pmatrix} \sigma_{11} & \sigma_{12} \\ \sigma_{12} & \sigma_{22} \end{pmatrix}\end{aligned}$$

⁴Note that $\sigma_{ii} = \sigma_i^2$.

then,

$$\begin{aligned}\mathbb{E}[[z_1]_+[z_2]_+] &= \mathbb{E}[[z_1]_+[z_2]_+|z_1 > 0, z_2 > 0] P(z_1 > 0, z_2 > 0) \\ &= \mathbb{E}[z_1 z_2 | z_1 > 0, z_2 > 0] P(z_1 > 0, z_2 > 0).\end{aligned}$$

Given $z_1 > 0, z_2 > 0$ the distribution of $(z_1, z_2)^\top$ is truncated multivariate normal distribution [Manjunath and Wilhelm, 2021], therefore,

$$\begin{aligned}\mathbb{E}[z_1 z_2 | z_1 > 0, z_2 > 0] &= \sigma_{12} + \sigma_{21}(-\mu_1) F_1(-\mu_1) + \sigma_{12}(-\mu_2) F_2(-\mu_2) \\ &\quad + (\sigma_{11}\sigma_{22} - \sigma_{12}\sigma_{21}) \frac{f(-\boldsymbol{\mu}; \mathbf{0}, \boldsymbol{\Sigma})}{p(z_1 > -\mu_1, z_2 > -\mu_2; \mathbf{0}, \boldsymbol{\Sigma})} \\ &\quad - (\sigma_{11}F_1(-\mu_1) + \sigma_{12}F_2(-\mu_2)) (\sigma_{21}F_1(-\mu_1) + \sigma_{22}F_2(-\mu_2)) \\ &\quad + (\sigma_{11}F_1(-\mu_1) + \sigma_{12}F_2(-\mu_2) + \mu_1) (\sigma_{21}F_1(-\mu_1) + \sigma_{22}F_2(-\mu_2) + \mu_2) \\ &= \sigma_{12} + \det(\boldsymbol{\Sigma}) \frac{f(-\boldsymbol{\mu}; \mathbf{0}, \boldsymbol{\Sigma})}{p(z_1 > -\mu_1, z_2 > -\mu_2; \mathbf{0}, \boldsymbol{\Sigma})} \\ &\quad + \sigma_{11}F_1(-\mu_1)\mu_2 + \mu_1\sigma_{22}F_2(-\mu_2) + \mu_1\mu_2\end{aligned}\tag{24}$$

where $f(-\boldsymbol{\mu}; \mathbf{0}, \boldsymbol{\Sigma})$ is a density function of Gaussian random vector with mean vector $\mathbf{0}$ and covariance matrix $\boldsymbol{\Sigma}$ at the point $-\boldsymbol{\mu}$, and

$$\begin{aligned}F_2(z_2) &= \frac{\int_{-\mu_1}^{\infty} f(z_1, z_2; \mathbf{0}, \boldsymbol{\Sigma}) dz_1}{P(Z_1 > -\mu_1, Z_2 > -\mu_2; \mathbf{0}, \boldsymbol{\Sigma})} \\ F_1(z_1) &= \frac{\int_{-\mu_2}^{\infty} f(z_1, z_2; \mathbf{0}, \boldsymbol{\Sigma}) dz_2}{P(Z_1 > -\mu_1, Z_2 > -\mu_2; \mathbf{0}, \boldsymbol{\Sigma})}.\end{aligned}$$

We denote by $\boldsymbol{\Lambda} = \boldsymbol{\Sigma}^{-1}$ thus,

$$\begin{aligned}&\int_{-\mu_1}^{\infty} \exp\left(-\frac{1}{2}(\Lambda_{11}z_1^2 + 2\Lambda_{12}z_1z_2 + \Lambda_{22}z_2^2)\right) dz_1 = \\ &\exp\left(-\frac{1}{2}\Lambda_{22}z_2^2\right) \int_{-\mu_1}^{\infty} \exp\left(-\frac{1}{2}\Lambda_{11}z_1^2 - \Lambda_{12}z_1z_2\right) dz_1 = \\ &\exp\left(-\frac{1}{2}\Lambda_{22}z_2^2\right) \int_{-\mu_1}^{\infty} \exp\left(-\frac{1}{2\frac{1}{\Lambda_{11}}}\left(z_1 + \frac{\Lambda_{12}z_2}{\Lambda_{11}}\right)^2 + \frac{\Lambda_{12}^2z_2^2}{2\Lambda_{11}}\right) dz_1 = \\ &\exp\left(-\frac{1}{2}\Lambda_{22}z_2^2 + \frac{\Lambda_{12}^2z_2^2}{2\Lambda_{11}}\right) \int_{-\mu_1}^{\infty} \exp\left(-\frac{1}{2\frac{1}{\Lambda_{11}}}\left(z_1 + \frac{\Lambda_{12}z_2}{\Lambda_{11}}\right)^2\right) dz_1 = \\ &\exp\left(-\frac{1}{2}\Lambda_{22}z_2^2 + \frac{\Lambda_{12}^2z_2^2}{2\Lambda_{11}}\right) \sqrt{\frac{2\pi}{\Lambda_{11}}} P\left(z > -\mu_1; -\frac{\Lambda_{12}z_2}{\Lambda_{11}}, \frac{1}{\Lambda_{11}}\right).\end{aligned}$$

Thus,

$$\begin{aligned}\int_{-\mu_1}^{\infty} f(z_1, z_2; \mathbf{0}, \boldsymbol{\Sigma}) dz_1 &= \frac{1}{2\pi\sqrt{\det(\boldsymbol{\Sigma})}} \exp\left(-\frac{1}{2}\Lambda_{22}z_2^2 + \frac{\Lambda_{12}^2z_2^2}{2\Lambda_{11}}\right) \sqrt{\frac{2\pi}{\Lambda_{11}}} P\left(z > -\mu_1; -\frac{\Lambda_{12}z_2}{\Lambda_{11}}, \frac{1}{\Lambda_{11}}\right) \\ &= \frac{1}{\sqrt{2\pi\sigma_{22}}} \exp\left(-\frac{1}{2\sigma_{22}}z_2^2\right) P\left(z > -\mu_1; -\frac{\Lambda_{12}z_2}{\Lambda_{11}}, \frac{1}{\Lambda_{11}}\right).\end{aligned}$$

Similarly, we obtain,

$$\int_{-\mu_2}^{\infty} f(z_1, z_2; \mathbf{0}, \boldsymbol{\Sigma}) dz_2 = \frac{1}{\sqrt{2\pi\sigma_{11}}} \exp\left(-\frac{1}{2\sigma_{11}}z_1^2\right) P\left(z > -\mu_2; -\frac{\Lambda_{21}z_1}{\Lambda_{22}}, \frac{1}{\Lambda_{22}}\right).$$

Therefore,

$$\begin{aligned}F_2(z_2) &= \frac{\varphi\left(\frac{z_2}{\sqrt{\sigma_{22}}}\right) P\left(z > -\mu_1; -\frac{\Lambda_{12}z_2}{\Lambda_{11}}, \frac{1}{\Lambda_{11}}\right)}{P(Z_1 > -\mu_1, Z_2 > -\mu_2; \mathbf{0}, \boldsymbol{\Sigma})} \\ F_1(z_1) &= \frac{\varphi\left(\frac{z_1}{\sqrt{\sigma_{11}}}\right) P\left(z > -\mu_2; -\frac{\Lambda_{21}z_1}{\Lambda_{22}}, \frac{1}{\Lambda_{22}}\right)}{P(Z_1 > -\mu_1, Z_2 > -\mu_2; \mathbf{0}, \boldsymbol{\Sigma})}.\end{aligned}$$

□




Figure 4: **Numerical evaluation of (24)** Histogram of the sample average of $\text{ReLU}(z_1)\text{ReLU}(z_2)$ for 10000 Monte-Carlo samples. We denote by E the analytical expectation and by \bar{E} the sample average. Figure (a) is for $\mu = \begin{pmatrix} -4 \\ 17 \end{pmatrix}$, $\Sigma = \begin{pmatrix} 13 & -9 \\ -9 & 8 \end{pmatrix}$, the normalized error is $\frac{|E-\bar{E}|}{E} = 0.0093\%$. Figure (b) is for $\mu = \begin{pmatrix} 6 \\ 2 \end{pmatrix}$, $\Sigma = \begin{pmatrix} 10 & 2 \\ 2 & 1 \end{pmatrix}$, the normalized error is $\frac{|E-\bar{E}|}{E} = 0.0044\%$.

Next, we present in Figures 3 and 4 a numerical evaluation of (23) and (24). As can be seen, the Monte-Carlo simulations verify the analytical results.

B Proofs of Results in Section 4

B.1 Proof of Proposition 1

Proof. Let

$$\mathcal{D} = \{(y_n = x_n + \epsilon_{n,m}, x_n)\}, n \in [N], m \in [M], \quad -\infty < x_1 < x_2 < \dots < x_N < \infty$$

such that Assumption 1 holds. We define a reduced dataset, which only contains the noisy samples with the most extreme noises

$$\bar{\mathcal{D}} = \{(x_n + \epsilon_n^{\min}, x_n), (x_n + \epsilon_n^{\max}, x_n)\}, n \in [N].$$

We define the ℓ^2 penalty on the weights,

$$C(\theta, K) = \sum_{k=1}^K (\|a_k\|^2 + \|w_k\|^2).$$

According to Theorem 1.2. in [Hanin, 2021], since we have opposite discrete curvature on the intervals $[x_1 + \epsilon_1^{\max}, x_2 + \epsilon_1^{\min}] \dots [x_{N-1} + \epsilon_1^{\max}, x_N + \epsilon_1^{\min}]$,

$$\{h_\theta(y) \mid h_\theta(y) = x \forall (y, x) \in \bar{\mathcal{D}}, C(\theta, K) = C_*\} = \{f_{1D}^*(y)\},$$

where

$$C_* = \inf_{\theta, K} \{C(\theta, K) \mid \forall (y, x) \in \bar{\mathcal{D}} \quad h_\theta(y) = x\}.$$

Also note that,

$$\{h_\theta(y) \mid h_\theta(y) = x \forall (y, x) \in \mathcal{D}, C(\theta, K) = C_*\} = \{f_{1D}^*(y)\}$$

since $f_{1D}^*(y)$ interpolates all the points in \mathcal{D} , and if we assume by contradiction that $C_* > \inf_{\theta, K} \{C(\theta, K) \mid \forall (y, x) \in \mathcal{D} \quad h_\theta(y) = x\}$ then $C_* \neq \inf_{\theta, K} \{C(\theta, K) \mid \forall (y, x) \in \bar{\mathcal{D}} \quad h_\theta(y) = x\}$, thus contradicting Theorem 1.2. in [Hanin, 2021]. Note that the minimal number of neurons needed to represent $f_{1D}^*(y)$ is $2N - 2$. So, if $K \geq 2N - 2$ then f_{1D}^* is the minimizer of the representation cost (i.e., $f_{1D}^* \in \arg \min_f R(f)$). \square

B.2 Proof of Theorem 1

First we prove the following lemmas.

Lemma 3. Let $y = x + \sigma\epsilon$ where $x, \epsilon \in \mathbb{R}, \sigma > 0$ then,

$$\lim_{\sigma \rightarrow 0^+} \hat{x}^{\text{eMMSE}}(y) = \hat{x}^{1-\text{NN}}(y)$$

Proof. Notice that the following holds,

$$\begin{aligned} \frac{\exp\left(-\frac{|y-x_i|^2}{2\sigma^2}\right)}{\sum_{i=1}^N \exp\left(-\frac{|y-x_i|^2}{2\sigma^2}\right)} &= \frac{\exp\left(-\frac{|y-x_i|^2}{2\sigma^2}\right) \exp\left(\frac{\min_n |y-x_n|^2}{2\sigma^2}\right)}{\sum_{i=1}^N \exp\left(\frac{|y-x_i|^2}{2\sigma^2}\right) \exp\left(\frac{\min_n |y-x_n|^2}{2\sigma^2}\right)} \\ &= \frac{\exp\left(-\frac{|y-x_i|^2 - \min_n |y-x_n|^2}{2\sigma^2}\right)}{\sum_{i=1}^N \exp\left(-\frac{|y-x_i|^2 - \min_n |y-x_n|^2}{2\sigma^2}\right)}. \end{aligned}$$

In addition,

$$\lim_{\sigma \rightarrow 0^+} \exp\left(-\frac{|y-x_i|^2 - \min_n |y-x_n|^2}{2\sigma^2}\right) = \begin{cases} 1 & |y-x_i|^2 = \min_n |y-x_n|^2 \\ 0 & |y-x_i|^2 \neq \min_n |y-x_n|^2 \end{cases}$$

so we obtain,

$$\lim_{\sigma \rightarrow 0^+} \frac{\exp\left(-\frac{|y-x_i|^2}{2\sigma^2}\right)}{\sum_{i=1}^N \exp\left(-\frac{|y-x_i|^2}{2\sigma^2}\right)} = \begin{cases} 1 & |y-x_i|^2 = \min_n |y-x_n|^2 \\ 0 & |y-x_i|^2 \neq \min_n |y-x_n|^2 \end{cases}.$$

Therefore,

$$\begin{aligned} \lim_{\sigma \rightarrow 0^+} \hat{x}^{\text{eMMSE}}(y) &= \lim_{\sigma \rightarrow 0^+} \frac{\sum_{i=1}^N x_i \exp\left(-\frac{|y-x_i|^2}{2\sigma^2}\right)}{\sum_{i=1}^N \exp\left(-\frac{|y-x_i|^2}{2\sigma^2}\right)} = \arg \min_n |y-x_n|^2 \\ &= \arg \min_n |y-x_n| = \hat{x}^{1-\text{NN}}(y). \end{aligned}$$

□

Lemma 4. Let $y = x + \sigma\epsilon$ where $x \in \mathbb{R}, \sigma > 0$ and $\epsilon \sim \mathcal{N}(0, 1)$ then,

$$\lim_{\sigma \rightarrow 0^+} \text{MSE}(\hat{x}^{\text{eMMSE}}(y)) = \lim_{\sigma \rightarrow 0^+} \text{MSE}(\hat{x}^{1-\text{NN}}(y))$$

Proof.

$$\begin{aligned} \text{MSE}(\hat{x}^{\text{eMMSE}}(y)) &= \mathbb{E} \left[(\hat{x}^{\text{eMMSE}}(y) - x)^2 \right] \\ &= \mathbb{E} \left[(\hat{x}^{\text{eMMSE}}(y) - \hat{x}^{1-\text{NN}}(y) + \hat{x}^{1-\text{NN}}(y) - x)^2 \right] \\ &= \text{MSE}(\hat{x}^{1-\text{NN}}(y)) + \mathbb{E} \left[(\hat{x}^{\text{eMMSE}}(y) - \hat{x}^{1-\text{NN}}(y))^2 \right] \\ &\quad + 2\mathbb{E} \left[(\hat{x}^{1-\text{NN}}(y) - x) (\hat{x}^{\text{eMMSE}}(y) - \hat{x}^{1-\text{NN}}(y)) \right] \end{aligned}$$

Note that,

$$\begin{aligned} \hat{x}^{\text{eMMSE}}(y) &= \frac{\sum_{i=1}^N x_i \exp\left(-\frac{|y-x_i|^2}{2\sigma^2}\right)}{\sum_{i=1}^N \exp\left(-\frac{|y-x_i|^2}{2\sigma^2}\right)} \leq \frac{\sum_{i=1}^N \max_j \{x_j\} \exp\left(-\frac{|y-x_i|^2}{2\sigma^2}\right)}{\sum_{i=1}^N \exp\left(-\frac{|y-x_i|^2}{2\sigma^2}\right)} \\ &= \max_j \{x_j\} \frac{\sum_{i=1}^N \exp\left(-\frac{|y-x_i|^2}{2\sigma^2}\right)}{\sum_{i=1}^N \exp\left(-\frac{|y-x_i|^2}{2\sigma^2}\right)} = \max_j \{x_j\}, \end{aligned}$$

$$\begin{aligned}\hat{x}^{\text{eMMSE}}(y) &= \frac{\sum_{i=1}^N x_i \exp\left(-\frac{|y-x_i|^2}{2\sigma^2}\right)}{\sum_{i=1}^N \exp\left(-\frac{|y-x_i|^2}{2\sigma^2}\right)} \geq \frac{\sum_{i=1}^N \min_j\{x_j\} \exp\left(-\frac{|y-x_i|^2}{2\sigma^2}\right)}{\sum_{i=1}^N \exp\left(-\frac{|y-x_i|^2}{2\sigma^2}\right)} \\ &= \min_j\{x_j\} \frac{\sum_{i=1}^N \exp\left(-\frac{|y-x_i|^2}{2\sigma^2}\right)}{\sum_{i=1}^N \exp\left(-\frac{|y-x_i|^2}{2\sigma^2}\right)} = \min_j\{x_j\}.\end{aligned}$$

Similarly,

$$\begin{aligned}\hat{x}^{1-\text{NN}}(y) &= \arg \min_i |y - x_i| \leq \max_i\{x_i\} \\ \hat{x}^{1-\text{NN}}(y) &= \arg \min_i |y - x_i| \geq \min_i\{x_i\}\end{aligned}$$

thus $\exists M_0 > 0 \forall \sigma > 0$

$$|\hat{x}^{\text{eMMSE}}(y) - \hat{x}^{1-\text{NN}}(y)| < M_0.$$

According to Lemma 3

$$\lim_{\sigma \rightarrow 0^+} \hat{x}^{\text{eMMSE}}(y) = \hat{x}^{1-\text{NN}}(y)$$

almost surely. Note that,

$$\mathbb{E}_{x,y} \left[(\hat{x}^{\text{eMMSE}}(y) - \hat{x}^{1-\text{NN}}(y))^2 \right] = \mathbb{E}_{x,\epsilon} \left[(\hat{x}^{\text{eMMSE}}(x + \sigma\epsilon) - \hat{x}^{1-\text{NN}}(x + \sigma\epsilon))^2 \right]$$

Therefore, by the Dominated convergence theorem we obtain

$$\begin{aligned}\lim_{\sigma \rightarrow 0^+} \mathbb{E}_{x,\epsilon} \left[(\hat{x}^{\text{eMMSE}}(x + \sigma\epsilon) - \hat{x}^{1-\text{NN}}(x + \sigma\epsilon))^2 \right] &= \\ \mathbb{E}_{x,\epsilon} \left[\lim_{\sigma \rightarrow 0^+} (\hat{x}^{\text{eMMSE}}(x + \sigma\epsilon) - \hat{x}^{1-\text{NN}}(x + \sigma\epsilon))^2 \right] &= 0\end{aligned}$$

Similarly,

$$\begin{aligned}\lim_{\sigma \rightarrow 0^+} \mathbb{E}_{x,\epsilon} \left[(\hat{x}^{1-\text{NN}}(x + \sigma\epsilon) - x) (\hat{x}^{\text{eMMSE}}(x + \sigma\epsilon) - \hat{x}^{1-\text{NN}}(x + \sigma\epsilon)) \right] &= \\ \mathbb{E}_{x,\epsilon} \left[\lim_{\sigma \rightarrow 0^+} (\hat{x}^{1-\text{NN}}(x + \sigma\epsilon) - x) (\hat{x}^{\text{eMMSE}}(x + \sigma\epsilon) - \hat{x}^{1-\text{NN}}(x + \sigma\epsilon)) \right] &= 0\end{aligned}$$

Since $\hat{x}^{\text{eMMSE}}(y)$ and $\hat{x}^{1-\text{NN}}(y)$ are bounded and $\mathbb{E}[x] < \infty$. Therefore, we get

$$\lim_{\sigma \rightarrow 0^+} \text{MSE}(\hat{x}^{\text{eMMSE}}(y)) = \lim_{\sigma \rightarrow 0^+} \text{MSE}(\hat{x}^{1-\text{NN}}(y))$$

□

Next, we prove Theorem 1.

Proof. Assume, without loss of generality, that $N = 2$. Let $p_x(x)$ be a probability density function with bounded second moment such that $p_x(x) > 0$ for all $x \in [x_1, x_2]$. According to Lemma 4

$$\lim_{\sigma \rightarrow 0^+} \text{MSE}(\hat{x}^{\text{eMMSE}}(y)) = \lim_{\sigma \rightarrow 0^+} \text{MSE}(\hat{x}^{1-\text{NN}}(y))$$

So we need to prove that

$$\lim_{\sigma \rightarrow 0^+} \text{MSE}(\hat{x}^{1-\text{NN}}(y)) > \lim_{\sigma \rightarrow 0^+} \text{MSE}(f_{1D}^*(y))$$

For the case of $N = 2$, the training set includes two data points $\{x_1, x_2\}$. So we get,

$$\begin{aligned}\hat{x}^{1-\text{NN}}(y) &= \begin{cases} x_1 & y < \frac{x_1+x_2}{2} \\ x_2 & \frac{x_1+x_2}{2} \leq y \end{cases} \\ f_{1D}^*(y) &= \begin{cases} x_1 & y < x_1 + \Delta_1 \\ \frac{x_2-x_1}{x_2-x_1+\Delta_2-\Delta_1} (y-x_1-\Delta_1) + x_1 & x_1 + \Delta_1 \leq y \leq x_2 + \Delta_2 \\ x_2 & y > x_2 + \Delta_2 \end{cases}\end{aligned}$$

where $\max_{m \in [M]} \epsilon_{1,m} = \Delta_1 > 0$, $\min_{m \in [M]} \epsilon_{2,m} = \Delta_2 < 0$. Note that $\hat{x}^{1-\text{NN}}(y) = f_{1D}^*(y)$ for all $y \in (-\infty, x_1 + \Delta_1) \cup (x_2 + \Delta, \infty)$ and $\mathbb{E}[x] < \infty$, $\mathbb{E}[x^2] < \infty$. Therefore,

$$\begin{aligned} \text{MSE}(\hat{x}^{1-\text{NN}}(y)) - \text{MSE}(f_{1D}^*(y)) &= \\ &= \int_{-\infty}^{\infty} dx \int_{x_1 + \Delta_1}^{x_2 + \Delta_2} dy (\hat{x}^{1-\text{NN}}(y) - x)^2 p_{y|x}(y|x) p_x(x) \end{aligned} \quad (25)$$

$$- \int_{-\infty}^{\infty} dx \int_{x_1 + \Delta_1}^{x_2 + \Delta_2} dy (f_{1D}^*(y) - x)^2 p_{y|x}(y|x) p_x(x) \quad (26)$$

First, we derive (25):

$$\begin{aligned} & \int_{-\infty}^{\infty} dx \int_{x_1 + \Delta_1}^{x_2 + \Delta_2} dy (\hat{x}^{1-\text{NN}}(y) - x)^2 p_{y|x}(y|x) p_x(x) = \\ & \int_{-\infty}^{\infty} \int_{x_1 + \Delta_1}^{\frac{x_1 + x_2}{2}} (x_1 - x)^2 p_{y|x}(y|x) p_x(x) dy dx + \int_{-\infty}^{\infty} \int_{\frac{x_1 + x_2}{2}}^{x_2 + \Delta_2} (x_2 - x)^2 p_{y|x}(y|x) p_x(x) dy dx = \\ & \mathbb{E}_x \left[P \left(y \in \left[x_1 + \Delta_1, \frac{x_1 + x_2}{2} \right] \middle| x \right) (x_1 - x)^2 \right] + \mathbb{E}_x \left[P \left(y \in \left[\frac{x_1 + x_2}{2}, x_2 + \Delta_2 \right] \middle| x \right) (x_2 - x)^2 \right]. \end{aligned}$$

Note that,

$$\begin{aligned} & \mathbb{E} \left[P \left(y \in \left[x_1 + \Delta_1, \frac{x_1 + x_2}{2} \right] \middle| x \right) (x_1 - x)^2 \right] < \mathbb{E} \left[(x_1 - x)^2 \right] < \infty \\ & \mathbb{E} \left[P \left(y \in \left[\frac{x_1 + x_2}{2}, x_2 + \Delta_2 \right] \middle| x \right) (x_2 - x)^2 \right] < \mathbb{E} \left[(x_2 - x)^2 \right] < \infty \end{aligned}$$

thus by the Dominated convergence theorem we obtain

$$\begin{aligned} & \lim_{\sigma \rightarrow 0^+} \int_{-\infty}^{\infty} dx \int_{x_1 + \Delta_1}^{x_2 + \Delta_2} dy (\hat{x}^{1-\text{NN}}(y) - x)^2 p_{y|x}(y|x) p_x(x) = \\ & \mathbb{E}_x \left[\lim_{\sigma \rightarrow 0^+} P \left(y \in \left[x_1 + \Delta_1, \frac{x_1 + x_2}{2} \right] \middle| x \right) (x_1 - x)^2 \right] + \mathbb{E}_x \left[\lim_{\sigma \rightarrow 0^+} P \left(y \in \left[\frac{x_1 + x_2}{2}, x_2 + \Delta_2 \right] \middle| x \right) (x_2 - x)^2 \right] = \\ & \mathbb{E}_x \left[(x - x_1)^2 1 \left\{ x \in \left[x_1, \frac{x_1 + x_2}{2} \right] \right\} \right] + \mathbb{E}_x \left[(x - x_2)^2 1 \left\{ x \in \left[\frac{x_1 + x_2}{2}, x_2 \right] \right\} \right] > C > 0 \end{aligned}$$

since $p_x(x) > 0$ for all $x \in [x_1, x_2]$. Next, we derive (26)

$$\begin{aligned} & \int_{-\infty}^{\infty} dx \int_{x_1 + \Delta_1}^{x_2 + \Delta_2} dy (f_{1D}^*(y) - x)^2 p_{y|x}(y|x) p_x(x) = \\ & \mathbb{E}_{x,y} \left[1 \{ y \in [x_1 + \Delta_1, x_2 + \Delta_2] \} (f_{1D}^*(y) - x)^2 \right] = \\ & \mathbb{E}_{x,\epsilon} \left[1 \{ x + \sigma\epsilon \in [x_1 + \Delta_1, x_2 + \Delta_2] \} (f_{1D}^*(x + \sigma\epsilon) - x)^2 \right] \end{aligned}$$

Note that,

$$\mathbb{E}_{x,\epsilon} \left[1 \{ x + \sigma\epsilon \in [x_1 + \Delta_1, x_2 + \Delta_2] \} (f_{1D}^*(x + \sigma\epsilon) - x)^2 \right] < \mathbb{E}_{x,\epsilon} \left[(f_{1D}^*(x + \sigma\epsilon) - x)^2 \right] < \infty$$

thus by the Dominated convergence theorem we obtain

$$\begin{aligned} & \lim_{\epsilon \rightarrow 0^+} \int_{-\infty}^{\infty} dx \int_{x_1 + \Delta_1}^{x_2 + \Delta_2} dy (f_{1D}^*(y) - x)^2 p_{y|x}(y|x) p_x(x) = \\ & \lim_{\epsilon \rightarrow 0^+} \mathbb{E}_{x,\epsilon} \left[1 \{ x + \sigma\epsilon \in [x_1 + \Delta_1, x_2 + \Delta_2] \} (f_{1D}^*(x + \sigma\epsilon) - x)^2 \right] = \\ & \mathbb{E}_{x,\epsilon} \left[\lim_{\epsilon \rightarrow 0^+} 1 \{ x + \sigma\epsilon \in [x_1 + \Delta_1, x_2 + \Delta_2] \} (f_{1D}^*(x + \sigma\epsilon) - x)^2 \right] = 0 \end{aligned}$$

since

$$\lim_{\epsilon \rightarrow 0^+} 1 \{x + \sigma\epsilon \in [x_1 + \Delta_1, x_2 + \Delta_2]\} = 1 \{x \in [x_1, x_2]\}$$

$$\lim_{\epsilon \rightarrow 0^+} f_{1D}^*(x + \sigma\epsilon) = \begin{cases} x_1 & x < x_1 \\ x & x_1 \leq x \leq x_2 \\ x_2 & x > x_2 \end{cases}$$

□

B.3 Proof of Lemma 1

Proof. First we prove that for all $y \in \cup_{n \in [N-1]} \left\{ \frac{x_{n+1}\epsilon_n^{\max} - x_n\epsilon_{n+1}^{\min}}{\epsilon_n^{\max} - \epsilon_{n+1}^{\min}} \right\}$ $f_{1D}^*(y) = y$. We find the intersection between the line $f_{1D}^*(y) = y$ and the linear part of (17) (the third branch)

$$\frac{x_{n+1} - x_n}{x_{n+1} + \epsilon_{n+1}^{\min} - (x_n + \epsilon_n^{\max})} (y - (x_n + \epsilon_n^{\max})) + x_n = y \quad (27)$$

$$(x_{n+1} - x_n)(y - (x_n + \epsilon_n^{\max})) + x_n(x_{n+1} + \epsilon_{n+1}^{\min} - (x_n + \epsilon_n^{\max})) = y(x_{n+1} + \epsilon_{n+1}^{\min} - (x_n + \epsilon_n^{\max})) \quad (28)$$

$$x_n(x_{n+1} + \epsilon_{n+1}^{\min} - (x_n + \epsilon_n^{\max})) - (x_{n+1} - x_n)(x_n + \epsilon_n^{\max}) = y(\epsilon_{n+1}^{\min} - \epsilon_n^{\max}) \quad (29)$$

$$x_n(x_{n+1} + \epsilon_{n+1}^{\min}) - x_{n+1}(x_n + \epsilon_n^{\max}) = y(\epsilon_{n+1}^{\min} - \epsilon_n^{\max}) \quad (30)$$

$$x_n\epsilon_{n+1}^{\min} - x_{n+1}\epsilon_n^{\max} = y(\epsilon_{n+1}^{\min} - \epsilon_n^{\max}) \quad (31)$$

$$y = \frac{x_{n+1}\epsilon_n^{\max} - x_n\epsilon_{n+1}^{\min}}{\epsilon_n^{\max} - \epsilon_{n+1}^{\min}}. \quad (32)$$

Note that

$$x_n + \epsilon_n^{\max} < \frac{x_{n+1}\epsilon_n^{\max} - x_n\epsilon_{n+1}^{\min}}{\epsilon_n^{\max} - \epsilon_{n+1}^{\min}} < x_{n+1} + \epsilon_{n+1}^{\min} \quad (33)$$

Since,

$$x_n + \epsilon_n^{\max} < \frac{x_{n+1}\epsilon_n^{\max} - x_n\epsilon_{n+1}^{\min}}{\epsilon_n^{\max} - \epsilon_{n+1}^{\min}} \quad (34)$$

$$(x_n + \epsilon_n^{\max})(\epsilon_n^{\max} - \epsilon_{n+1}^{\min}) < x_{n+1}\epsilon_n^{\max} - x_n\epsilon_{n+1}^{\min} \quad (35)$$

$$(x_n + \epsilon_n^{\max} - \epsilon_{n+1}^{\min})\epsilon_n^{\max} < x_{n+1}\epsilon_n^{\max} \quad (36)$$

$$x_n + \epsilon_n^{\max} - \epsilon_{n+1}^{\min} < x_{n+1} \quad (37)$$

$$x_n + \epsilon_n^{\max} < x_{n+1} + \epsilon_{n+1}^{\min} \quad (38)$$

which holds according to Assumption 1.

$$\frac{x_{n+1}\epsilon_n^{\max} - x_n\epsilon_{n+1}^{\min}}{\epsilon_n^{\max} - \epsilon_{n+1}^{\min}} < x_{n+1} + \epsilon_{n+1}^{\min} \quad (39)$$

$$x_{n+1}\epsilon_n^{\max} - x_n\epsilon_{n+1}^{\min} < (x_{n+1} + \epsilon_{n+1}^{\min})(\epsilon_n^{\max} - \epsilon_{n+1}^{\min}) \quad (40)$$

$$-x_n\epsilon_{n+1}^{\min} < -\epsilon_{n+1}^{\min}(x_{n+1} - \epsilon_n^{\max} + \epsilon_{n+1}^{\min}) \quad (41)$$

$$x_n < x_{n+1} - \epsilon_n^{\max} + \epsilon_{n+1}^{\min} \quad (42)$$

$$x_n + \epsilon_n^{\max} < x_{n+1} + \epsilon_{n+1}^{\min} \quad (43)$$

which holds according to Assumption 1.

Next we prove that $f_{1D}^*(y)$ is contractive toward a set of the clean datapoints $\{x_n\}_{n=1}^N$ on $\mathbb{R} \setminus \cup_{n \in [N-1]} \left\{ \frac{x_{n+1}\epsilon_n^{\max} - x_n\epsilon_{n+1}^{\min}}{\epsilon_n^{\max} - \epsilon_{n+1}^{\min}} \right\}$. In the case where $y \in (-\infty, x_1 + \epsilon_1^{\min}]$ or $y \in \cup_{n \in [N]} [x_n + \epsilon_n^{\min}, x_n + \epsilon_n^{\max}]$ or $y \in [x_N + \epsilon_N^{\max}, \infty)$ $f_{1D}^*(y) \in \{x_n\}_{n=1}^N$ therefore (18) holds for all $0 \leq \alpha < 1$.

In the case where $y \in \cup_{n \in [N-1]} [x_n + \epsilon_n^{\max}, \frac{x_{n+1}\epsilon_n^{\max} - x_n\epsilon_{n+1}^{\min}}{\epsilon_n^{\max} - \epsilon_{n+1}^{\min}})$ we choose $i = n$,

$$|f_{1D}^*(y) - f(x_n)| = \frac{x_{n+1} - x_n}{x_{n+1} + \epsilon_{n+1}^{\min} - (x_n + \epsilon_n^{\max})} (y - (x_n + \epsilon_n^{\max})) \quad (44)$$




Figure 5: Illustration of $f_{1D}^*(y)$.

There exists $0 < \gamma_1 < 1$ such that

$$x_n \leq \frac{x_{n+1} - x_n}{x_{n+1} + \epsilon_{n+1}^{\min} - (x_n + \epsilon_n^{\max})} (y - (x_n + \epsilon_n^{\max})) + x_n \leq \gamma_1 y \quad (45)$$

since for $y \in \cup_{n \in [N-1]} [x_n + \epsilon_n^{\max}, \frac{x_{n+1}\epsilon_n^{\max} - x_n\epsilon_{n+1}^{\min}}{\epsilon_n^{\max} - \epsilon_{n+1}^{\min}}]$ $f_{1D}^*(y)$ is below the line $f(y) = y$ because

$f_{1D}^*(y)$ is an affine function with slope larger than 1 and $f_{1D}^*(\frac{x_{n+1}\epsilon_n^{\max} - x_n\epsilon_{n+1}^{\min}}{\epsilon_n^{\max} - \epsilon_{n+1}^{\min}}) = \frac{x_{n+1}\epsilon_n^{\max} - x_n\epsilon_{n+1}^{\min}}{\epsilon_n^{\max} - \epsilon_{n+1}^{\min}}$.

Therefore there exists $0 < \alpha_1 < 1$

$$|f_{1D}^*(y) - f(x_n)| = \frac{x_{n+1} - x_n}{x_{n+1} + \epsilon_{n+1}^{\min} - (x_n + \epsilon_n^{\max})} (y - (x_n + \epsilon_n^{\max})) \leq \gamma_1 y - x_n \leq \alpha_1 (y - x_n) \quad (46)$$

since,

$$\gamma_1 y - x_n \leq \alpha_1 (y - x_n) \quad (47)$$

$$\alpha_1 \geq \frac{\gamma_1 y - x_n}{y - x_n}. \quad (48)$$

In the case where $y \in \cup_{n \in [N-1]} (\frac{x_{n+1}\epsilon_n^{\max} - x_n\epsilon_{n+1}^{\min}}{\epsilon_n^{\max} - \epsilon_{n+1}^{\min}}, x_{n+1} + \epsilon_{n+1}^{\min}]$ we choose $i = n + 1$,

$$|f_{1D}^*(y) - f(x_{n+1})| = x_{n+1} - \frac{x_{n+1} - x_n}{x_{n+1} + \epsilon_{n+1}^{\min} - (x_n + \epsilon_n^{\max})} (y - (x_n + \epsilon_n^{\max})) - x_n. \quad (49)$$

There exists $0 < \gamma_2 < 1$ such that

$$x_{n+1} \geq \frac{x_{n+1} - x_n}{x_{n+1} + \epsilon_{n+1}^{\min} - (x_n + \epsilon_n^{\max})} (y - (x_n + \epsilon_n^{\max})) + x_n \geq \frac{1}{\gamma_2} y \quad (50)$$

since for $y \in \cup_{n \in [N-1]} (\frac{x_{n+1}\epsilon_n^{\max} - x_n\epsilon_{n+1}^{\min}}{\epsilon_n^{\max} - \epsilon_{n+1}^{\min}}, x_{n+1} + \epsilon_{n+1}^{\min}]$ $f_{1D}^*(y)$ is above the line $f(y) = y$ because

$f_{1D}^*(y)$ is an affine function with slope larger than 1 and $f_{1D}^*(\frac{x_{n+1}\epsilon_n^{\max} - x_n\epsilon_{n+1}^{\min}}{\epsilon_n^{\max} - \epsilon_{n+1}^{\min}}) = \frac{x_{n+1}\epsilon_n^{\max} - x_n\epsilon_{n+1}^{\min}}{\epsilon_n^{\max} - \epsilon_{n+1}^{\min}}$.

Therefore there exists $0 < \alpha_2 < 1$

$$|f_{1D}^*(y) - f(x_{n+1})| = x_{n+1} - \frac{x_{n+1} - x_n}{x_{n+1} + \epsilon_{n+1}^{\min} - (x_n + \epsilon_n^{\max})} (y - (x_n + \epsilon_n^{\max})) - x_n \quad (51)$$

$$\leq x_{n+1} - \frac{1}{\gamma_2} y \leq \alpha_2 (x_{n+1} - y) \quad (52)$$

since,

$$x_{n+1} - \frac{1}{\gamma_2} y \leq \alpha_2 (x_{n+1} - y) \quad (53)$$

$$\alpha_2 \geq \frac{x_{n+1} - \frac{1}{\gamma_2} y}{x_{n+1} - y}. \quad (54)$$

Therefore (18) holds for $\alpha = \max\{\alpha_1, \alpha_2\}$. \square

C Proofs of Results in Section 5

Let $\mathbf{f} : \mathbb{R}^d \rightarrow \mathbb{R}^d$ be any function realizable as a shallow ReLU network. Consider any parametrization of \mathbf{f} given by $\mathbf{f}(\mathbf{y}) = \sum_{k=1}^K \mathbf{a}_k [\mathbf{w}_k^\top \mathbf{y} + b_k]_+ + \mathbf{V} \mathbf{y} + \mathbf{c}$. We say such a parametrization is a *minimal representative of \mathbf{f}* if $\|\mathbf{w}_k\|_2 = 1$ and $\mathbf{a}_k \neq 0$ for all $k = 1, \dots, K$, and $R(\mathbf{f}) = \sum_{k=1}^K \|\mathbf{a}_k\|_2$. In particular, the units making up a minimal representative must be distinct in the sense that the hyperplanes describing ReLU boundaries $H_k = \{\mathbf{x} \in \mathbb{R}^d : \mathbf{w}_k^\top \mathbf{x} + b_k = 0\}$ are distinct, which implies that no units can be cancelled or combined.

We will also make use of the following lemma, which shows that representation costs are invariant to a translation of the training samples, assuming the function is suitably translated.

Lemma 5. *Let $\mathbf{f} : \mathbb{R}^d \rightarrow \mathbb{R}^d$ be any function realizable as a shallow ReLU net satisfying norm-ball interpolation constraints $\mathbf{f}(B(\mathbf{x}_n; \rho)) = \{\mathbf{x}_n\}$ for all $n = 1, \dots, N$. Let $\mathbf{x}_0 \in \mathbb{R}^d$. Then the function $\mathbf{g}(\mathbf{y}) = \mathbf{f}(\mathbf{y} - \mathbf{x}_0) + \mathbf{x}_0$ is such that $R(\mathbf{g}) = R(\mathbf{f})$ and $\mathbf{g}(B(\mathbf{x}_n + \mathbf{x}_0; \rho)) = \{\mathbf{x}_n + \mathbf{x}_0\}$ for all $n = 1, \dots, N$.*

Proof. First we show $\mathbf{g}(B(\mathbf{x}_n + \mathbf{x}_0; \rho)) = \{\mathbf{x}_n + \mathbf{x}_0\}$ for all $n = 1, \dots, N$. Fix any n , and let $\mathbf{y} \in B(\mathbf{x}_n + \mathbf{x}_0; \rho)$. Then $\mathbf{y} = \mathbf{y}' + \mathbf{x}_0$ for some $\mathbf{y}' \in B(\mathbf{x}_n; \rho)$, and so $\mathbf{g}(\mathbf{y}) = \mathbf{f}(\mathbf{y}') + \mathbf{x}_0 = \mathbf{x}_n + \mathbf{x}_0$, as claimed.

To show that $R(\mathbf{g}) = R(\mathbf{f})$, let $\mathbf{f}(\mathbf{y}) = \sum_{k=1}^K \mathbf{a}_k [\mathbf{w}_k^\top \mathbf{y} + b_k]_+ + \mathbf{V} \mathbf{y} + \mathbf{c}$ be any minimal representative of \mathbf{f} . Then

$$\mathbf{g}(\mathbf{y}) = \sum_{k=1}^K \mathbf{a}_k [\mathbf{w}_k^\top (\mathbf{y} - \mathbf{x}_0) + b_k]_+ + \mathbf{V} (\mathbf{y} - \mathbf{x}_0) + \mathbf{c} + \mathbf{x}_0 \quad (55)$$

$$= \sum_{k=1}^K \mathbf{a}_k [\mathbf{w}_k^\top \mathbf{y} + \tilde{b}_k]_+ + \mathbf{V} \mathbf{y} + \tilde{\mathbf{c}} \quad (56)$$

with $\tilde{b}_k = b_k - \mathbf{w}_k^\top \mathbf{x}_0$ and $\tilde{\mathbf{c}} = \mathbf{c} + \mathbf{x}_0 - \mathbf{V} \mathbf{x}_0$. And so $R(\mathbf{g}) \leq \sum_{k=1}^K \|\mathbf{a}_k\|_2 = R(\mathbf{f})$. A parallel argument with the roles of \mathbf{g} and \mathbf{f} reversed shows that $R(\mathbf{f}) \leq R(\mathbf{g})$, and so $R(\mathbf{f}) = R(\mathbf{g})$. \square

In particular, this lemma shows that if $\mathbf{f}^*(\mathbf{y})$ is a norm-ball interpolating representation cost minimizer of the training samples $\{\mathbf{x}_n\}_{n=1}^N$, then $\mathbf{g}^*(\mathbf{y}) = \mathbf{f}^*(\mathbf{y} - \mathbf{x}_0) + \mathbf{x}_0$ is a norm-ball interpolating min-cost solution of the translated training samples $\{\mathbf{x}_n + \mathbf{x}_0\}_{n=1}^N$.

C.1 Training data belonging to a subspace

The proof of Theorem 2 is a direct consequence of the following two lemmas:

Lemma 6. *Suppose the clean training samples $\{\mathbf{x}_n\}_{n=1}^N$ belong to a r -dimensional subspace $\mathcal{S} \subset \mathbb{R}^d$, and let $\mathbf{P}_\mathcal{S}$ denote the orthogonal projector onto \mathcal{S} . Let $\mathbf{f}(\mathbf{y})$ be any function realizable as a shallow ReLU network satisfying $\mathbf{f}(B(\mathbf{x}_n; \rho)) = \{\mathbf{x}_n\}$ for all $n = 1, \dots, N$. Define $\mathbf{g}(\mathbf{y}) = \mathbf{P}_\mathcal{S} \mathbf{f}(\mathbf{x})$. Then $\mathbf{g}(B(\mathbf{x}_n; \rho)) = \{\mathbf{x}_n\}$ for all $n = 1, \dots, N$, and $R(\mathbf{g}) \leq R(\mathbf{f})$, with strict inequality if $\mathbf{f} \neq \mathbf{g}$.*

Proof. First, for any $\mathbf{y} \in B(\mathbf{x}_n; \rho)$ we have $\mathbf{g}(\mathbf{y}) = \mathbf{P}_\mathcal{S} \mathbf{f}(\mathbf{y}) = \mathbf{P}_\mathcal{S} \mathbf{x}_n = \mathbf{x}_n$. Therefore, $\mathbf{g}(B(\mathbf{x}_n; \rho)) = \{\mathbf{x}_n\}$ for all $n = 1, \dots, N$.

Now, let $\mathbf{f}(\mathbf{y}) = \sum_{k=1}^K \mathbf{a}_k [\mathbf{w}_k^\top \mathbf{y} + b_k]_+ + \mathbf{V} \mathbf{y} + \mathbf{c}$ be a minimal representative of \mathbf{f} . Then $\mathbf{g}(\mathbf{y}) = \sum_{k=1}^K \mathbf{P}_\mathcal{S} \mathbf{a}_k [\mathbf{w}_k^\top \mathbf{y} + b_k]_+ + \mathbf{P}(\mathbf{V} \mathbf{y} + \mathbf{c})$. Since $\|\mathbf{P}_\mathcal{S} \mathbf{a}_k\| \leq \|\mathbf{a}_k\|$ for all k , we have $R(\mathbf{g}) \leq \sum_{k=1}^K \|\mathbf{P}_\mathcal{S} \mathbf{a}_k\| \leq \sum_{k=1}^K \|\mathbf{a}_k\| = R(\mathbf{f})$.

Now we show $\mathbf{f} \neq \mathbf{g}$ implies $R(\mathbf{g}) < R(\mathbf{f})$. Observe that if any of the outer-layer weight vectors $\mathbf{a}_k \notin \mathcal{S}$ then $\|\mathbf{P}_\mathcal{S} \mathbf{a}_k\| < \|\mathbf{a}_k\|$, which implies $R(\mathbf{g}) < R(\mathbf{f})$. Hence, it suffices to show that $\mathbf{f} \neq \mathbf{g}$ implies some $\mathbf{a}_k \notin \mathcal{S}$.

First, consider the case where $\mathbf{P}_\mathcal{S} \mathbf{V} = \mathbf{V}$ and $\mathbf{P}_\mathcal{S} \mathbf{c} = \mathbf{c}$. Then in this case $\mathbf{f} \neq \mathbf{g}$ if only if $\sum_{k=1}^K (\mathbf{P}_\mathcal{S} \mathbf{a}_k - \mathbf{a}_k) [\mathbf{w}_k^\top \mathbf{y} + b_k]_+ \neq 0$ for all $\mathbf{y} \in \mathbb{R}^d$, which implies $\mathbf{P}_\mathcal{S} \mathbf{a}_k \neq \mathbf{a}_k$ for some k , or equivalently $\mathbf{a}_k \notin \mathcal{S}$.

Next, assume either $P_S V \neq V$ or $P_S c \neq c$. Fix any training sample index n , and let A_n denote the index set of units active over the ball $B(\mathbf{x}_n, \rho)$. Then, since $\mathbf{f}(\mathbf{y})$ is constant for all $\mathbf{y} \in B(\mathbf{x}_n, \rho)$, the Jacobian $\partial \mathbf{f}(\mathbf{y}) = \sum_{k \in A_n} \mathbf{a}_k \mathbf{w}_k^\top + \mathbf{V} = \mathbf{0}$ for all $\mathbf{y} \in B(\mathbf{x}_n, \rho)$. This gives $\mathbf{V} = -\sum_{k \in A_n} \mathbf{a}_k \mathbf{w}_k^\top$. Therefore, if $P_S V \neq V$ then at least one of the $\mathbf{a}_k \notin \mathcal{S}$. On the other hand, if $P_S c \neq c$ then for all $\mathbf{y} \in B(\mathbf{x}_n, \rho)$ we have

$$\mathbf{f}(\mathbf{y}) = \sum_{k \in A_n} \mathbf{a}_k (\mathbf{w}_k^\top \mathbf{y} + b_k) + \mathbf{V} \mathbf{y} + \mathbf{c} = \sum_{k \in A_n} (\mathbf{a}_k \mathbf{w}_k^\top + \mathbf{V}) \mathbf{y} + \sum_{k \in A_n} b_k \mathbf{a}_k + \mathbf{c} = \sum_{k \in A_n} b_k \mathbf{a}_k + \mathbf{c} = \mathbf{x}_n,$$

which implies $\mathbf{c} = \mathbf{x}_n - \sum_{k \in A_n} b_k \mathbf{a}_k$, and since $\mathbf{x}_n \in \mathcal{S}$ this implies some $\mathbf{a}_k \notin \mathcal{S}$, proving the claim. \square

Lemma 7. Suppose the clean training samples $\{\mathbf{x}_n\}_{n=1}^N$ belong to an r -dimensional subspace $\mathcal{S} \subset \mathbb{R}^d$, and let P_S denote the orthogonal projector onto \mathcal{S} . Let \mathbf{f} be any network satisfying $\mathbf{f}(B(\mathbf{x}_n, \rho)) = \{\mathbf{x}_n\}$. Define $\mathbf{h}(\mathbf{y}) = \mathbf{f}(P_S \mathbf{y})$. Then $\mathbf{h}(B(\mathbf{x}_n, \rho)) = \{\mathbf{x}_n\}$ and $R(\mathbf{h}) \leq R(\mathbf{f})$, with strict inequality if $\mathbf{h} \neq \mathbf{f}$.

Proof. Define $P_S^{-1}(B(\mathbf{x}_n, \rho)) := \{\mathbf{y} \in \mathbb{R}^d : P_S \mathbf{y} \in B(\mathbf{x}_n, \rho)\}$, i.e., the set of points in \mathbb{R}^d whose projection onto \mathcal{S} is contained in $B(\mathbf{x}_n, \rho)$. Then clearly $\mathbf{h}(P_S^{-1}(B(\mathbf{x}_n, \rho))) = \{\mathbf{x}_n\}$. Also, by properties of norm-balls, we have $B(\mathbf{x}_n, \rho) \subset P_S^{-1}(B(\mathbf{x}_n, \rho))$. Therefore, $\mathbf{h}(B(\mathbf{x}_n, \rho)) = \{\mathbf{x}_n\}$ for all $n = 1, \dots, N$.

Next, let $\mathbf{f}(\mathbf{y}) = \sum_{k=1}^K \mathbf{a}_k [\mathbf{w}_k^\top \mathbf{y} + b_k]_+ + \mathbf{V} \mathbf{y} + \mathbf{c}$ be a minimal representative of \mathbf{f} . Then $\mathbf{h}(\mathbf{y}) = \sum_{k=1}^K \mathbf{a}_k [(P_S \mathbf{w}_k)^\top \mathbf{y} + b_k]_+ + \mathbf{V} P_S \mathbf{y} + \mathbf{c}$. Since $\|P_S \mathbf{w}_k\| \leq \|\mathbf{w}_k\|$ for all k , we have $R(\mathbf{h}) \leq \sum_{k=1}^K \|\mathbf{a}_k\| \|P_S \mathbf{w}_k\| \leq \sum_{k=1}^K \|\mathbf{a}_k\| \|\mathbf{w}_k\| = R(\mathbf{f})$.

Finally, we show that if $\mathbf{f} \neq \mathbf{h}$, then $R(\mathbf{h}) < R(\mathbf{f})$. Observe that if any of the inner-layer weight vectors $\mathbf{w}_k \notin \mathcal{S}$ then $\|P_S \mathbf{w}_k\| < \|\mathbf{w}_k\|$, which implies $R(\mathbf{h}) < R(\mathbf{f})$. Hence, it suffices to show that $\mathbf{f} \neq \mathbf{h}$ implies some $\mathbf{w}_k \notin \mathcal{S}$. First, consider the case $\mathbf{V} P_S = \mathbf{V}$. Then $\mathbf{f} \neq \mathbf{h}$ if and only if $P_S \mathbf{w}_k \neq \mathbf{w}_k$ for some k , or equivalently, $\mathbf{w}_k \notin \mathcal{S}$. Next, consider the case $\mathbf{V} P_S \neq \mathbf{V}$. Fix any training sample index n , and let A_n denote the index set of units active over $B(\mathbf{x}_n, \rho)$. Then, since $\mathbf{f}(\mathbf{y})$ is constant for all $\mathbf{y} \in B(\mathbf{x}_n, \rho)$, the Jacobian $\partial \mathbf{f}(\mathbf{x}) = \sum_{k \in A_n} \mathbf{a}_k \mathbf{w}_k^\top + \mathbf{V} = \mathbf{0}$ for all $\mathbf{y} \in B_n(\mathbf{x}_n, \rho)$. This gives $\mathbf{V} = -\sum_{k \in A_n} \mathbf{a}_k \mathbf{w}_k^\top$. Therefore, if $\mathbf{V} P_S \neq \mathbf{V}$ (i.e., at least one row of \mathbf{V} is not contained in \mathcal{S}) then at least one of the $\mathbf{w}_k \notin \mathcal{S}$, proving the claim. \square

Finally, we now give the proof Theorem 2 and Corollary 1.

Proof of Theorem 2. Let \mathbf{f}^* be any min-cost solution. Applying both Lemma 6 and 7 we see it must be the case that $\mathbf{f}^*(\mathbf{y}) = P_S \mathbf{f}^*(P_S \mathbf{y})$ for all $\mathbf{y} \in \mathbb{R}^d$, since otherwise the representation cost of \mathbf{f}^* could be reduced. \square

Proof of Corollary 1. Suppose \mathcal{S} is one-dimensional, i.e., $\mathcal{S} = \text{span}\{\mathbf{u}\}$ for some unit vector $\mathbf{u} \in \mathbb{R}^d$. Theorem 2 shows we can express any min-cost solution \mathbf{f}^* as

$$\mathbf{f}^*(\mathbf{y}) = \sum_{k=1}^K a_k \mathbf{u} [s_k \mathbf{u}^\top \mathbf{y} + b_k]_+ + v \mathbf{u} \mathbf{u}^\top \mathbf{y} + c \mathbf{u} = \mathbf{u} \phi(\mathbf{u}^\top \mathbf{y})$$

where

$$\phi(t) = \sum_k a_k [s_k t + b_k]_+ + v t + c$$

is such that $R(\mathbf{f}^*) = R(\phi)$. Therefore, minimizing the representation cost subject to norm-ball interpolation constraints is reduced to a univariate problem:

$$\min_{\phi} R(\phi) \quad \text{s.t.} \quad \phi((c_n - \rho, c_n + \rho)) = c_n$$

where $c_n = \mathbf{u}^\top \mathbf{x}_n$. The minimizing ϕ^* is unique and coincides with the 1-D denoiser f_{1D}^* in (17) with $x_n = c_n$ and $\epsilon_n^{\max} = -\epsilon_n^{\min} = \rho$. \square

C.2 Data along rays

We begin with a key lemma that is central to the proof of Theorem 3.

Lemma 8. *Suppose $\{\mathbf{u}_i\}_{i=1}^n \subset \mathbb{R}^d$ is a collection of $n > 1$ unit vectors such that $\mathbf{u}_i^\top \mathbf{u}_j < 0$ for all $i \neq j$, and \mathbf{w} is a unit vector such that $\mathbf{u}_i^\top \mathbf{w} > 0$ for all $i = 1, \dots, n$. Let $\mathbf{a} \in \mathbb{R}^d$ be any vector. Then,*

$$\sum_{i=1}^n |\mathbf{u}_i^\top \mathbf{a}| (\mathbf{u}_i^\top \mathbf{w}) < \|\mathbf{a}\|.$$

Before giving the proof of Lemma 8, we first prove an auxiliary result.

Lemma 9. *Let $\mathbf{a} \in \mathbb{R}^d$ be a unit vector, and suppose $\{\mathbf{u}_i\}_{i=1}^n \subset \mathbb{R}^d$ is a collection of unit vectors such that $\mathbf{u}_i^\top \mathbf{a} > 0$ for all i and $\mathbf{u}_i^\top \mathbf{u}_j < 0$ for all $i \neq j$. Let $\mathbf{b} = \sum_{i=1}^n \mathbf{u}_i \mathbf{u}_i^\top \mathbf{a}$. Then $\|\mathbf{b} - \mathbf{a}/2\| \leq 1/2$ with strict inequality if $n > 1$.*

Proof. It suffices to show $\mathbf{b}^\top \mathbf{a} \geq \mathbf{b}^\top \mathbf{b}$, since if this is the case then

$$\|\mathbf{b} - \mathbf{a}/2\| = \sqrt{(\mathbf{b} - \mathbf{a}/2)^\top (\mathbf{b} - \mathbf{a}/2)} = \sqrt{\mathbf{b}^\top \mathbf{b} - \mathbf{b}^\top \mathbf{a} + \frac{1}{4} \mathbf{a}^\top \mathbf{a}} \leq \frac{1}{2}, \quad (57)$$

which also holds with strict inequality when $\mathbf{b}^\top \mathbf{a} > \mathbf{b}^\top \mathbf{b}$.

First, if $n = 1$, then $\mathbf{b} = \mathbf{a}^\top \mathbf{u}_1 \mathbf{u}_1^\top$, and so $\mathbf{b}^\top \mathbf{b} = \mathbf{a}^\top \mathbf{u}_1 \mathbf{u}_1^\top \mathbf{u}_1 \mathbf{u}_1^\top \mathbf{a} = \mathbf{a}^\top \mathbf{u}_1 \mathbf{u}_1^\top \mathbf{a} = \mathbf{b}^\top \mathbf{a}$. Now assume $n > 1$. Then we have

$$\mathbf{b}^\top \mathbf{a} = \sum_{i=1}^n \mathbf{a}^\top \mathbf{u}_i \mathbf{u}_i^\top \mathbf{a} \quad (58)$$

$$= \sum_{i=1}^n \mathbf{a}^\top \mathbf{u}_i \mathbf{u}_i^\top \mathbf{u}_i \mathbf{u}_i^\top \mathbf{a} \quad (59)$$

$$> \sum_{i=1}^n \sum_{j=1}^n \mathbf{a}^\top \mathbf{u}_i \mathbf{u}_i^\top \mathbf{u}_j \mathbf{u}_j^\top \mathbf{a} \quad (60)$$

$$= \mathbf{b}^\top \mathbf{b} \quad (61)$$

The first and last equalities hold by definition. The second equality holds because each \mathbf{u}_i is a unit vector. The last inequality holds because for $i \neq j$

$$(\mathbf{a}^\top \mathbf{u}_i)(\mathbf{u}_i^\top \mathbf{u}_j)(\mathbf{u}_j^\top \mathbf{a}) < 0,$$

since $\mathbf{u}_i^\top \mathbf{u}_j < 0$ and $\mathbf{a}^\top \mathbf{u}_i, \mathbf{a}^\top \mathbf{u}_j > 0$ by assumption. \square

Now we give the proof of Lemma 8.

Proof of Lemma 8. Let $\mathbf{v} = \sum_{i=1}^n |\mathbf{u}_i^\top \mathbf{a}| \mathbf{u}_i$. Then we have

$$S = \sum_{i=1}^n |\mathbf{u}_i^\top \mathbf{a}| (\mathbf{u}_i^\top \mathbf{w}) = \mathbf{v}^\top \mathbf{w}$$

WLOG, we may assume $\|\mathbf{a}\| = 1$, in which case the lemma reduces to showing $S = \mathbf{v}^\top \mathbf{w} < 1$. By the Cauchy-Schwartz inequality, it suffices to show $\|\mathbf{v}\| < 1$. Toward this end, let us write $\mathbf{v} = \mathbf{v}_+ + \mathbf{v}_-$ where $\mathbf{v}_+ = \sum_{i:\mathbf{u}_i^\top \mathbf{a} > 0} |\mathbf{u}_i^\top \mathbf{a}| \mathbf{u}_i = \sum_{i:\mathbf{u}_i^\top \mathbf{a} > 0} \mathbf{u}_i \mathbf{u}_i^\top \mathbf{a}$ and $\mathbf{v}_- = \sum_{i:\mathbf{u}_i^\top \mathbf{a} < 0} |\mathbf{u}_i^\top \mathbf{a}| \mathbf{u}_i = \sum_{i:\mathbf{u}_i^\top \mathbf{a} < 0} \mathbf{u}_i \mathbf{u}_i^\top (-\mathbf{a})$. By Lemma 9 we have $\|\mathbf{v}_+ - \mathbf{a}/2\| \leq 1/2$ and $\|\mathbf{v}_- + \mathbf{a}/2\| \leq 1/2$, and so

$$\|\mathbf{v}\| = \|\mathbf{v}_+ + \mathbf{v}_-\| \quad (62)$$

$$= \|\mathbf{v}_+ - \mathbf{a}/2 + \mathbf{v}_- + \mathbf{a}/2\| \quad (63)$$

$$\leq \|\mathbf{v}_+ - \mathbf{a}/2\| + \|\mathbf{v}_- + \mathbf{a}/2\| \quad (64)$$

$$\leq 1. \quad (65)$$

Also, if either of the index sets $I_+ := \{i : \mathbf{u}_i^\top \mathbf{a} > 0\}$ or $I_- := \{i : \mathbf{u}_i^\top \mathbf{a} < 0\}$ has cardinality greater than one then the lemma above guarantees $\|\mathbf{v}_+ - \mathbf{a}/2\| < 1/2$ or $\|\mathbf{v}_+ - \mathbf{a}/2\| < 1/2$, and so $\|\mathbf{v}\| < 1$, which gives $S < 1$.

It remains to show $S < 1$ when I_+ or I_- has cardinality less than or equal to one, i.e., $|I_+| \leq 1$ or $|I_-| \leq 1$. We consider the various possibilities:

Case 1: $|I_+| = 0$ & $|I_-| = 0$. Then $\mathbf{v} = \mathbf{0}$ and so $S = \mathbf{v}^\top \mathbf{w} = 0 < 1$.

Case 2: $|I_+| = 1$ & $|I_-| = 0$ or $|I_+| = 0$ & $|I_-| = 1$. Then $\mathbf{v} = |\mathbf{u}_i^\top \mathbf{a}| \mathbf{u}_i$ for some i , and $|\mathbf{u}_j^\top \mathbf{a}| = 0$ for all $j \neq i$. By way of contradiction, assume that $S = \mathbf{v}^\top \mathbf{w} = |\mathbf{u}_i^\top \mathbf{a}| \mathbf{u}_i^\top \mathbf{w} = 1$. Since \mathbf{u}_i , \mathbf{a} , and \mathbf{w} are all unit vectors, the only way this is possible is if $|\mathbf{u}_i^\top \mathbf{a}| = 1$ and $\mathbf{u}_i^\top \mathbf{w} = 1$, which implies $\mathbf{a} = \pm \mathbf{w}$ and $\mathbf{u}_i = \mathbf{w}$, and so $\mathbf{a} = \pm \mathbf{u}_i$. However, this shows that $\mathbf{u}_j^\top \mathbf{u}_i = 0$ for all $j \neq i$, contradicting our assumption that $\mathbf{u}_j^\top \mathbf{u}_i < 0$ for all $i \neq j$. Therefore, $S < 1$ in this case as well.

Case 3: $|I_+| = 1$ & $|I_-| = 1$. Then $\mathbf{v} = |\mathbf{u}_i^\top \mathbf{a}| \mathbf{u}_i + |\mathbf{u}_j^\top \mathbf{a}| \mathbf{u}_j$ for some $i \in I_+$ and $j \in I_-$, and so

$$\|\mathbf{v}\| = \||\mathbf{u}_i^\top \mathbf{a}| \mathbf{u}_i + |\mathbf{u}_j^\top \mathbf{a}| \mathbf{u}_j\| \quad (66)$$

$$= \|(\mathbf{u}_i \mathbf{u}_i^\top - \mathbf{u}_j \mathbf{u}_j^\top) \mathbf{a}\| \quad (67)$$

$$\leq \|\mathbf{u}_i \mathbf{u}_i^\top - \mathbf{u}_j \mathbf{u}_j^\top\| \quad (68)$$

where the last inequality follows by the fact that $\|\mathbf{a}\| = 1$. Finally, since the matrix $\mathbf{u}_i \mathbf{u}_i^\top - \mathbf{u}_j \mathbf{u}_j^\top$ is symmetric, its operator norm $\|\mathbf{u}_i \mathbf{u}_i^\top - \mathbf{u}_j \mathbf{u}_j^\top\|$ coincides with the maximum eigenvalue (in absolute value). Eigenvectors in the span of \mathbf{u}_i and \mathbf{u}_j have the form $\mathbf{v} = c_i \mathbf{u}_i + c_j \mathbf{u}_j$ where c_i, c_j satisfy the equation

$$(\mathbf{u}_i \mathbf{u}_i^\top - \mathbf{u}_j \mathbf{u}_j^\top)(c_i \mathbf{u}_i + c_j \mathbf{u}_j) = \lambda(c_i \mathbf{u}_i + c_j \mathbf{u}_j) \quad (69)$$

where $\lambda \in \mathbb{R}$ is the corresponding eigenvalue. Expanding and equating coefficients gives the system

$$\begin{bmatrix} 1 - \lambda & \mathbf{u}_i^\top \mathbf{u}_j \\ \mathbf{u}_i^\top \mathbf{u}_j & 1 + \lambda \end{bmatrix} \begin{bmatrix} c_i \\ c_j \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix} \quad (70)$$

which has non-trivial solutions iff

$$(1 - \lambda)(1 + \lambda) - (\mathbf{u}_i^\top \mathbf{u}_j)^2 = 0 \iff \lambda = \pm \sqrt{1 - (\mathbf{u}_i^\top \mathbf{u}_j)^2} \quad (71)$$

and so $|\lambda| < 1$. Therefore, $\|\mathbf{v}\| \leq \|\mathbf{u}_i \mathbf{u}_i^\top - \mathbf{u}_j \mathbf{u}_j^\top\| < 1$ as claimed. And so $S \leq \|\mathbf{v}\| < 1$. \square




Figure 6: Illustration of “unit splitting” technique used in proof of Theorem 3. If a ReLU unit is active along two rays (left panel), it can be split into two units with lower representation cost by projecting its inner-layer weight vector \mathbf{w} onto the two rays separately (right panel).

Proof of Theorem 3. Let \mathbf{f}^* be any representation cost minimizer. We prove the claim by showing that if \mathbf{f}^* fails to have the form specified in (20) then it is possible to construct a norm-ball interpolant \mathbf{h} whose units are aligned with the rays that has strictly smaller representation cost than \mathbf{f}^* .

First, let $f^*(\mathbf{y}) = \sum_{k=1}^K \mathbf{a}_k[\mathbf{w}_k^\top \mathbf{y} + b_k] + \mathbf{V}\mathbf{x} + \mathbf{c}$ be any minimal representative of f^* . By properties of minimal representatives, none of the ReLU boundary sets $\{\mathbf{y} \in \mathbb{R}^d : \mathbf{w}_k^\top \mathbf{y} + b_k = 0\}$ can intersect any of the norm-balls centered at the training samples, since otherwise f^* would be non-constant on one of the norm-balls. Also, we may alter this parameterization in such a way that the active set of every unit avoids the ball centered at the origin $B_0 := B(\mathbf{0}, \rho)$ without changing the representation cost. In particular, suppose the active set of the k th unit contains B_0 . Then, using the identity $[t]_+ = t - [-t]_+$ for all $t \in \mathbb{R}$, we have

$$\mathbf{a}_k[\mathbf{w}_k^\top \mathbf{y} + b_k]_+ = \mathbf{a}_k(\mathbf{w}_k^\top \mathbf{y} + b_k) - \mathbf{a}_k[-\mathbf{w}_k^\top \mathbf{y} - b_k]_+$$

for all $\mathbf{y} \in \mathbb{R}^d$. The active set of the reversed unit $-\mathbf{a}_k[-\mathbf{w}_k^\top \mathbf{y} - b_k]_+$ does not intersect B_0 . Therefore, after applying this transformation to all units whose active sets contain B_0 , we can write $f^* = \mathbf{g}^* + \ell$ where \mathbf{g}^* is a sum of ReLU units whose active sets do not contain B_0 and ℓ is an affine function that combines the original affine part $\mathbf{V}\mathbf{y} + \mathbf{c}$ with a sum of linear units of the form $\mathbf{a}_k(\mathbf{w}_k^\top \mathbf{y} + b_k \mathbf{a}_k)$ arising from the transformation above. However, because of the interpolation constraint $f(B_0) = \{\mathbf{0}\}$, and by the fact that no units in \mathbf{g}^* are active over B_0 , for all $\mathbf{y} \in B_0$ we have

$$f^*(\mathbf{y}) = \mathbf{g}^*(\mathbf{y}) + \ell(\mathbf{y}) = \ell(\mathbf{y}) = \mathbf{0} \quad (72)$$

which implies $\ell \equiv \mathbf{0}$, i.e., $f^*(\mathbf{y}) = \mathbf{g}^*(\mathbf{y})$ for all \mathbf{y} . Finally, note that the inner- and outer-layer weight vectors on the reversed units change only in sign. Therefore, this re-parameterization does not change the representation cost, and so is also a minimal representative.

Now we construct a new norm-ball interpolant \mathbf{h} as follows. Let f_ℓ be the function defined as the sum of all units making up the re-parametrized f^* that are active on at least one norm-ball centered on the ℓ th ray. Also, let $\mathbf{P}_\ell = \mathbf{u}_\ell \mathbf{u}_\ell^\top \in \mathbb{R}^{d \times d}$ denote the orthogonal projector onto the ℓ th ray. Define $\mathbf{h} = \sum_{\ell=1}^L \mathbf{h}_\ell$ where $\mathbf{h}_\ell(\mathbf{y}) := \mathbf{P}_\ell f_\ell(\mathbf{P}_\ell \mathbf{y})$. Put in words, \mathbf{h} is constructed by “splitting” any units active over multiple rays into a sum of multiple units aligned with each ray; see Figure 6 for an illustration.

First, we prove that \mathbf{h} satisfies norm-ball interpolation constraints. Let $\mathbf{a}[\mathbf{w}^\top \mathbf{y} + b]_+$ denote any unit belonging to f_ℓ . Since this unit is active on a norm-ball centered on ray ℓ , this implies that $\mathbf{w}^\top \mathbf{u}_\ell > 0$, and since the unit is inactive on B_0 we must have $b < 0$. Also, because we assume $\mathbf{u}_\ell^\top \mathbf{u}_m < 0$ for any $m \neq \ell$, we see that the projected unit $\mathbf{P}_\ell \mathbf{a}[\mathbf{w}^\top \mathbf{P}_\ell \mathbf{y} + b]_+ = \mathbf{P}_\ell \mathbf{a}[(\mathbf{w}^\top \mathbf{u}_\ell) \mathbf{u}_\ell^\top \mathbf{y} + b]_+$ is active on norm-balls centered on ray ℓ and inactive on norm-balls centered on any other ray. In particular, if \mathbf{y} belongs to a norm-ball centered on ray ℓ , then $\mathbf{h}_m(\mathbf{y}) = 0$ for all $m \neq \ell$. Therefore, if $\mathbf{y} \in B(\mathbf{x}_n^{(\ell)}, \rho)$ where $\mathbf{x}_n^{(\ell)}$ denotes a training sample along ray ℓ , then

$$\mathbf{h}(\mathbf{y}) = \sum_{m=1}^L \mathbf{h}_m(\mathbf{y}) \quad (73)$$

$$= \mathbf{h}_\ell(\mathbf{y}) \quad (74)$$

$$= \mathbf{P}_\ell f_\ell(\mathbf{P}_\ell \mathbf{y}) \quad (75)$$

$$= \mathbf{P}_\ell \mathbf{f}(\mathbf{P}_\ell \mathbf{y}) \quad (76)$$

$$= \mathbf{P}_\ell \mathbf{x}_n^{(\ell)} \quad (77)$$

$$= \mathbf{x}_n^{(\ell)} \quad (78)$$

which shows that \mathbf{h} satisfies norm-ball interpolation constraints, as claimed.

Next, we show that if $\mathbf{h} \neq f^*$, then \mathbf{h} has strictly smaller representation cost. Let $\mathbf{u}(\mathbf{y}) = \mathbf{a}[\mathbf{w}^\top \mathbf{y} + b]_+$ denote any ReLU unit belonging to f^* . Since we assume $\|\mathbf{w}\| = 1$, the contribution of unit \mathbf{u} to the representation cost of f^* is $\|\mathbf{a}\|$. Let $\mathcal{R} \subset \{1, 2, \dots, L\}$ index the subset of rays ℓ for which unit \mathbf{u} is active over at least one norm-ball centered on that ray. In constructing the interpolant \mathbf{h} , the unit \mathbf{u} get mapped to a sum of $|\mathcal{R}|$ units:

$$\sum_{\ell \in \mathcal{R}} \mathbf{P}_\ell \mathbf{a}[(\mathbf{P}_\ell \mathbf{w})^\top \mathbf{y} + b]_+$$

whose contribution to the representation cost of \mathbf{h} is bounded above by

$$C = \sum_{\ell \in \mathcal{R}} \|\mathbf{P}_\ell \mathbf{a}\| \|\mathbf{P}_\ell \mathbf{w}\| = \sum_{\ell \in \mathcal{R}} |\mathbf{u}_\ell^\top \mathbf{a}| |\mathbf{u}_\ell^\top \mathbf{w}|.$$




Figure 7: Change in unit alignment of minimal representation cost solutions under perturbations of training samples belonging to rays. The red lines represent the ReLU boundaries of the representation cost minimizing solution satisfying norm-ball interpolation constraints. For the unperturbed samples in (a), all ReLU boundaries align perpendicular to the rays. For the perturbed samples in (b), the ReLU boundaries align perpendicular to the line segments connecting successive samples.

If $|\mathcal{R}| > 1$, then by Lemma 8 guarantees $C < \|\mathbf{a}\|$. And if $|\mathcal{R}| = 1$ then $C = |\mathbf{u}_\ell^\top \mathbf{a}| |\mathbf{u}_\ell^\top \mathbf{w}| \leq \|\mathbf{a}\|$, with strict inequality unless \mathbf{w}, \mathbf{a} belong to the span of \mathbf{u}_ℓ . This implies that any min-cost solution must have all its units aligned with the rays, i.e., \mathbf{f}^* must have the form $\mathbf{f}^*(\mathbf{y}) = \sum_{\ell=1}^L \mathbf{u}_\ell \phi_\ell(\mathbf{u}_\ell^\top \mathbf{y})$, where each ϕ_ℓ is a univariate function. The representation cost of any function \mathbf{f}^* in this form is the sum of the (1-D) representation costs of the ϕ_ℓ . Therefore, ϕ_ℓ must also be the 1-D min-cost solution of the norm-ball constraints projected onto ray ℓ , and so must have the form specified by (17). \square

C.2.1 Perturbations of samples along rays

As an extension of the above setting, suppose the training samples along rays $\mathbf{x}_n^{(\ell)}$ have been slightly perturbed, i.e., we consider training samples $\tilde{\mathbf{x}}_n^{(\ell)} = \mathbf{x}_n^{(\ell)} + \boldsymbol{\epsilon}_n^{(\ell)}$ for some vectors $\boldsymbol{\epsilon}_n^{(\ell)}$ with $\|\boldsymbol{\epsilon}_n^{(\ell)}\| < \delta$ for some sufficiently small $\delta > 0$. In particular, we make the following two assumptions:

- (A1) Let $\Delta_\ell := \{\tilde{\mathbf{x}}_n^{(\ell)} - \tilde{\mathbf{x}}_{n-1}^{(\ell)}\}_{n=1}^{N_\ell}$ be the collection of difference vectors between successive perturbed points along the ℓ th ray (with $\tilde{\mathbf{x}}_0^{(\ell)} := \mathbf{0}$). Assume that for all $\mathbf{v}_\ell \in \Delta_\ell$ and for all $\mathbf{v}_k \in \Delta_k$ with $k \neq \ell$, we have $\mathbf{v}_\ell^\top \mathbf{v}_k < 0$.
- (A2) If H is any halfspace not containing the origin that contains the ball $B(\tilde{\mathbf{x}}_n^{(\ell)}; \rho)$, then H also contains all successive balls along that ray, i.e., H contains $B(\tilde{\mathbf{x}}_m^{(\ell)}; \rho)$ for $m \geq n$.

Note that if the original points $\mathbf{x}_n^{(\ell)}$ belong to rays that satisfy the conditions of Theorem 3, then for sufficiently small δ the perturbed samples $\tilde{\mathbf{x}}_n^{(\ell)}$ will satisfy assumptions (A1)&(A2) above.

We show in this case the norm-ball interpolating representation cost minimizer is a perturbed version of the min-cost solution for the unperturbed samples as identified in Theorem 3. In particular, the ReLU boundaries align with the line segments connecting successive points along the rays. See Figure 7 for an illustration in the case of $L = 2$ rays.

Proposition 3. *Suppose the training samples X are a perturbation of data belong to a union of L rays plus a sample at the origin, i.e., $X = \{\mathbf{0}\} \cup \{\tilde{\mathbf{x}}_n^{(1)}\}_{n=1}^{N_1} \cup \dots \cup \{\tilde{\mathbf{x}}_n^{(L)}\}_{n=1}^{N_L}$ where $\tilde{\mathbf{x}}_n^{(\ell)} = c_n^{(\ell)} \mathbf{u}_\ell + \boldsymbol{\epsilon}_n^{(\ell)}$ for some unit vector \mathbf{u}_ℓ , constants $0 < c_1^{(\ell)} < c_2^{(\ell)} < \dots < c_{N_\ell}^{(\ell)}$, and vectors $\boldsymbol{\epsilon}_n^{(\ell)}$. Assume (A1)&(A2) above hold. Then the minimizer \mathbf{f}^* of (19) is unique and is given by*

$$\mathbf{f}^*(\mathbf{y}) = \sum_{\ell=1}^L \sum_{n=1}^{N_\ell} \mathbf{u}_n^{(\ell)} \phi_n^{(\ell)}((\mathbf{u}_n^{(\ell)})^\top (\mathbf{y} - \tilde{\mathbf{x}}_{n-1}^{(\ell)})) \quad (79)$$

where $\mathbf{u}_n^{(\ell)} = \frac{\tilde{\mathbf{x}}_n^{(\ell)} - \tilde{\mathbf{x}}_{n-1}^{(\ell)}}{\|\tilde{\mathbf{x}}_n^{(\ell)} - \tilde{\mathbf{x}}_{n-1}^{(\ell)}\|}$ and $\phi_n^{(\ell)}(t) = s_n^{(\ell)}([t - a_n^{(\ell)}] - [t - b_n^{(\ell)}]_+)$ with $a_n^{(\ell)} = \rho$, $b_n^{(\ell)} = \|\mathbf{x}_n^{(\ell)} - \mathbf{x}_{n-1}^{(\ell)}\| - \rho$, and $s_n^{(\ell)} = \|\tilde{\mathbf{x}}_n^{(\ell)} - \tilde{\mathbf{x}}_{n-1}^{(\ell)}\| / (b_n^{(\ell)} - a_n^{(\ell)})$.

Proof. Let \mathbf{f}^* be any min-cost solution. Following the same steps as in the proof of Theorem 3, there is a minimal representative of \mathbf{f}^* having the form $\mathbf{f}^*(\mathbf{y}) = \sum_k \mathbf{a}_k [\mathbf{w}_k^\top \mathbf{y} + b_k]_+$ where the active sets of all units in this representation have empty intersection with the ball $B(\mathbf{0}, \rho)$.

Let $\mathbf{f}_n^{(\ell)}$ denote the sum of all units in \mathbf{f}^* whose active set boundary $\{\mathbf{y} \in \mathbb{R}^d : \mathbf{w}_k^\top \mathbf{y} + b_k = 0\}$ intersects the line segment connecting the training samples $\tilde{\mathbf{x}}_{n-1}^{(\ell)}$ and $\tilde{\mathbf{x}}_n^{(\ell)}$. By assumption (A2), this is equivalent to the sum of units that are active over balls $B(\tilde{\mathbf{x}}_m^{(\ell)}, \rho)$ for $m \geq n$, and inactive for the balls with $m < n$. In particular, for $\mathbf{y} \in B(\tilde{\mathbf{x}}_1^{(\ell)}, \rho)$, we have $\mathbf{f}^*(\mathbf{y}) = \sum_{m \leq n} \mathbf{f}_m^{(\ell)}(\mathbf{y})$. This implies $\mathbf{f}_1^{(\ell)}(\mathbf{y}) = \tilde{\mathbf{x}}_1$ for all $\mathbf{y} \in B(\tilde{\mathbf{x}}_1^{(\ell)}, \rho)$. Likewise, $\mathbf{f}_2^{(\ell)}(\mathbf{y}) = \tilde{\mathbf{x}}_2^{(\ell)} - \tilde{\mathbf{x}}_1^{(\ell)}$ for all $\mathbf{y} \in B(\tilde{\mathbf{x}}_2^{(\ell)}; \rho)$, and so on, such that for all $n = 1, \dots, N_\ell$ we have $\mathbf{f}_n^{(\ell)}(\mathbf{y}) = \tilde{\mathbf{x}}_n^{(\ell)} - \tilde{\mathbf{x}}_{n-1}^{(\ell)}$ for all $\mathbf{y} \in B(\tilde{\mathbf{x}}_n^{(\ell)}; \rho)$.

Now we show how to construct a new interpolating function \mathbf{h} having representation cost less than or equal to \mathbf{f}^* . Let $\mathbf{u}_n^{(\ell)} = \frac{\tilde{\mathbf{x}}_n^{(\ell)} - \tilde{\mathbf{x}}_{n-1}^{(\ell)}}{\|\tilde{\mathbf{x}}_n^{(\ell)} - \tilde{\mathbf{x}}_{n-1}^{(\ell)}\|}$ and define $\mathbf{P}_n^{(\ell)} = \mathbf{u}_n^{(\ell)} (\mathbf{u}_n^{(\ell)})^\top$, the orthogonal projector onto the span of the difference vector $\tilde{\mathbf{x}}_n^{(\ell)} - \tilde{\mathbf{x}}_{n-1}^{(\ell)}$. Note that the map $\mathbf{y} \mapsto \mathbf{P}_n^{(\ell)} \mathbf{y} + \tilde{\mathbf{x}}_{n-1}^{(\ell)}$ is projection onto the affine line connecting samples $\tilde{\mathbf{x}}_{n-1}^{(\ell)}$ and $\tilde{\mathbf{x}}_n^{(\ell)}$. Consider the function

$$\mathbf{h} = \sum_{\ell=1}^L \sum_{n=1}^{N_\ell} \mathbf{h}_n^{(\ell)}$$

where

$$\mathbf{h}_n^{(\ell)}(\mathbf{y}) = \mathbf{P}_n^{(\ell)} \mathbf{f}_n^{(\ell)}(\mathbf{P}_n^{(\ell)} \mathbf{y} + \tilde{\mathbf{x}}_{n-1}^{(\ell)}).$$

Put in words, \mathbf{h} is constructed by aligning all inner-layer and outer-layer weights of units in \mathbf{f}^* with the line segment connecting successive data points over which that unit first activates. See Figure 8 for an illustration.




Figure 8: Illustration of “unit alignment” technique used in the proof of Proposition 3. The left panel shows an interpolant \mathbf{f} whose ReLU boundaries (shown here as dark blue and dark red lines) are not aligned with the line segments S_1 and S_2 connecting successive training samples. The right panel shows how a new interpolant \mathbf{h} can be constructed by projecting units in \mathbf{f} onto the line segments S_1 and S_2 , which reduces the representation cost.

Now we show that \mathbf{h} satisfies norm-ball interpolation constraints. By assumption (A1), the function $\mathbf{h}_n^{(\ell)}$ is non-zero only on the norm-balls $B(\tilde{\mathbf{x}}_m^{(\ell)}, \rho)$ for $m \geq n$. This implies that for all $\mathbf{y} \in B(\tilde{\mathbf{x}}_n^{(\ell)}; \rho)$ we have $\mathbf{h}(\mathbf{y}) = \sum_{m \leq n} \mathbf{h}_m^{(\ell)}(\mathbf{y})$.

Observe that $\mathbf{h}(B(\mathbf{0}, \rho)) = \{\mathbf{0}\}$ since all functions $\mathbf{h}_n^{(\ell)}$ are zero on $B(\mathbf{0}, \rho)$. Also, for all $\mathbf{y} \in B(\tilde{\mathbf{x}}_1^{(\ell)}; \rho)$ we have $\mathbf{P}_1^{(\ell)} \mathbf{y} \in B(\tilde{\mathbf{x}}_1^{(\ell)}; \rho)$, and so

$$\mathbf{h}(\mathbf{y}) = \mathbf{h}_1^{(\ell)}(\mathbf{y}) = \mathbf{P}_1^{(\ell)} \mathbf{f}_1^{(\ell)}(\mathbf{P}_1^{(\ell)} \mathbf{y}) = \mathbf{P}_1^{(\ell)} \tilde{\mathbf{x}}_1^{(\ell)} = \tilde{\mathbf{x}}_1^{(\ell)}.$$

Similarly, the only terms in \mathbf{h} active for $\mathbf{y} \in B(\tilde{\mathbf{x}}_2^{(\ell)}; \rho)$ are $\mathbf{h}_1^{(\ell)}$ and $\mathbf{h}_2^{(\ell)}$, for which $\mathbf{h}_1^{(\ell)}(\mathbf{y}) = \mathbf{x}_1^{(\ell)}$, and $\mathbf{h}_2^{(\ell)}(\mathbf{y}) = \mathbf{P}_2^{(\ell)} \mathbf{f}_2^{(\ell)}(\mathbf{P}_2^{(\ell)} \mathbf{y} + \tilde{\mathbf{x}}_1^{(\ell)}) = \mathbf{P}_2^{(\ell)}(\tilde{\mathbf{x}}_2^{(\ell)} - \tilde{\mathbf{x}}_1^{(\ell)}) = \tilde{\mathbf{x}}_2^{(\ell)} - \tilde{\mathbf{x}}_1^{(\ell)}$. This gives

$$\mathbf{h}(\mathbf{y}) = \mathbf{h}_1^{(\ell)}(\mathbf{y}) + \mathbf{h}_2^{(\ell)}(\mathbf{y}) = \tilde{\mathbf{x}}_2^{(\ell)}.$$

Iterating this procedure for all $n = 1, \dots, N_\ell$ we see that if $\mathbf{y} \in B(\tilde{\mathbf{x}}_n^{(\ell)}, \rho)$ then $\mathbf{h}_n^{(\ell)}(\mathbf{y}) = \tilde{\mathbf{x}}_n^{(\ell)} - \tilde{\mathbf{x}}_{n-1}^{(\ell)}$, and so $\mathbf{h}(\mathbf{y}) = \sum_{m \leq n} \mathbf{h}_m^{(\ell)}(\mathbf{y}) = \mathbf{x}_n^{(\ell)}$, as claimed.

Following steps similar to the proof of Theorem 3, we can show that $R(\mathbf{h}) \leq R(\mathbf{f}^*)$ and with strict inequality if $\mathbf{h} \neq \mathbf{f}^*$. First, if any unit $\mathbf{u}(\mathbf{y}) = \mathbf{a}[\mathbf{w}^\top \mathbf{y} + b]_+$ in \mathbf{f}^* is active over balls centered on multiple rays $\mathcal{R} \subset \{1, 2, \dots, L\}$, then in the construction of \mathbf{h} the unit \mathbf{u} gets mapped to a sum of multiple units:

$$\sum_{\ell \in \mathcal{R}} \mathbf{u}_{n_\ell}^{(\ell)} (\mathbf{u}_{n_\ell}^{(\ell)})^\top \mathbf{a} [\mathbf{w}^\top \mathbf{u}_{n_\ell}^{(\ell)} (\mathbf{u}_{n_\ell}^{(\ell)})^\top (\mathbf{y} + \tilde{\mathbf{x}}_{n_\ell-1}^{(\ell)}) + b]_+$$

for some n_ℓ with $1 \leq n_\ell \leq N_\ell$. The contribution of the sum of these units to the representation cost of \mathbf{h} is less than or equal to $C = \sum_{\ell \in \mathcal{R}} |(\mathbf{u}_{n_\ell}^{(\ell)})^\top \mathbf{a}| |(\mathbf{u}_{n_\ell}^{(\ell)})^\top \mathbf{w}|$. By assumption (A1), we know $(\mathbf{u}_m^{(\ell)})^\top \mathbf{u}_n^{(p)} < 0$ for all $p \neq \ell$ and for all m, n . Therefore, by Lemma 8 we know $C < \|\mathbf{a}\|$. This shows \mathbf{f}^* cannot have any units active over balls centered on different rays, since otherwise \mathbf{h} has strictly smaller representation cost. Therefore, we have shown \mathbf{f}^* as $\mathbf{f}^* = \sum_{\ell=1}^L \sum_{n=1}^{N_\ell} \mathbf{f}_n^{(\ell)}$. Additionally, we see that $\mathbf{h}_n^{(\ell)} = \mathbf{f}_n^{(\ell)}$, i.e., all inner- and outer-layer weight vectors of units in $\mathbf{f}_n^{(\ell)}$ are aligned with $\mathbf{u}_n^{(\ell)}$, since otherwise the representation cost of \mathbf{f}^* could be reduced. Therefore, each $\mathbf{f}_n^{(\ell)}$ must have the form $\mathbf{f}_n^{(\ell)}(\mathbf{y}) = \mathbf{u}_n^{(\ell)} \psi_n^{(\ell)}((\mathbf{u}_n^{(\ell)})^\top \mathbf{y})$. Minimizing the $\psi_n^{(\ell)}$ separately under interpolation constraints, we arrive at the unique solution given in (79). \square

C.3 Simplex Data

C.3.1 Simplex with one obtuse vertex

Proof of Proposition 2. Consider training samples $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N \in \mathbb{R}^d$ whose convex hull is a $(N - 1)$ -simplex such that \mathbf{x}_1 makes an obtuse angle with all other vertices, i.e., $(\mathbf{x}_n - \mathbf{x}_1)^\top \mathbf{x}_1 < 0$ for all $n = 2, \dots, N$. By Lemma 5, \mathbf{f}^* is a norm-ball interpolating min-cost solution if and only if $\mathbf{g}^*(\mathbf{y}) = \mathbf{f}^*(\mathbf{y} - \mathbf{x}_1) + \mathbf{x}_1$ is a norm-ball interpolating min-cost solution for the translated points $\mathbf{0}, \mathbf{x}_2 - \mathbf{x}_1, \dots, \mathbf{x}_N - \mathbf{x}_1 \in \mathbb{R}^d$. The latter configuration satisfies the hypotheses of Theorem 3 with a single training sample per ray. Therefore, the min-cost solution \mathbf{g}^* of the translated points has units whose inner- and outer-layer weight vectors are aligned with $\mathbf{x}_n - \mathbf{x}_1$, $n = 2, \dots, N$, and likewise for \mathbf{f}^* , since it is translation of \mathbf{g}^* . \square

C.3.2 Simplex with all acute vertices

Justification of Conjecture 1. For concreteness we focus on the case of three points $\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3 \in \mathbb{R}^d$ whose convex hull is an acute triangle, and let B_1, B_2, B_3 be open balls of radius ρ centered at $\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3$ (respectively).

Let \mathbf{f} be any norm-ball interpolating min-cost solution, and let $\mathbf{f}(\mathbf{y}) = \sum_k \mathbf{a}_k [\mathbf{w}_k^\top \mathbf{y} + b_k] + \mathbf{V} \mathbf{x} + \mathbf{c}$ be any minimal representative of \mathbf{f} .

First, by properties of minimal representatives, none of the ReLU boundary sets $H_k = \{\mathbf{y} \in \mathbb{R}^d : \mathbf{w}_k^\top \mathbf{y} + b_k = 0\}$ intersect any of the balls centered at the training samples. Also, the active set of each unit in \mathbf{f} must contain either one or two norm-balls, since otherwise the unit is either inactive over all balls or active over all balls, in which cases the unit can be removed or absorbed into unregularized linear part while strictly reducing the representation cost. By ‘‘reversing units’’, as in the proof of Theorem 3, we may transform the parameterization in such a way that the active set of every unit contains exactly one ball, and do this without changing the representation cost.

After this transformation, we may write

$$\mathbf{f} = \mathbf{f}_1 + \mathbf{f}_2 + \mathbf{f}_3 + \ell$$

where f_i is a sum of units active only on the ball B_i and no others, and where $\ell(\mathbf{y}) = \mathbf{A}\mathbf{y} + \mathbf{b}$ is an affine function. Then we have

$$\forall \mathbf{y}_1 \in B_1, \quad \mathbf{f}(\mathbf{y}_1) = \mathbf{f}_1(\mathbf{y}_1) + \mathbf{A}\mathbf{y}_1 + \mathbf{b} = \mathbf{x}_1 \quad (80)$$

$$\forall \mathbf{y}_2 \in B_2, \quad \mathbf{f}(\mathbf{y}_2) = \mathbf{f}_2(\mathbf{y}_2) + \mathbf{A}\mathbf{y}_2 + \mathbf{b} = \mathbf{x}_2 \quad (81)$$

$$\forall \mathbf{y}_3 \in B_3, \quad \mathbf{f}(\mathbf{y}_3) = \mathbf{f}_3(\mathbf{y}_3) + \mathbf{A}\mathbf{y}_3 + \mathbf{b} = \mathbf{x}_3 \quad (82)$$

and so,

$$\forall \mathbf{y}_1 \in B_1, \quad \mathbf{f}_1(\mathbf{y}_1) = \mathbf{x}_1 - \mathbf{b} - \mathbf{A}\mathbf{y}_1 \quad (83)$$

$$\forall \mathbf{y}_2 \in B_2, \quad \mathbf{f}_2(\mathbf{y}_2) = \mathbf{x}_2 - \mathbf{b} - \mathbf{A}\mathbf{y}_2 \quad (84)$$

$$\forall \mathbf{y}_3 \in B_3, \quad \mathbf{f}_3(\mathbf{y}_3) = \mathbf{x}_3 - \mathbf{b} - \mathbf{A}\mathbf{y}_3 \quad (85)$$

Finding the min-cost solution then amounts to minimizing of the representation cost of f_1, f_2, f_3 subject to the above constraints. This is made challenging by the fact that the constraints are coupled together by the parameters \mathbf{A} and \mathbf{b} of the affine part. However, *under the assumption $\mathbf{A} = \mathbf{0}$* , we prove below that the min-cost solution must have the conjectured form \mathbf{f}^* as given in (86). However, since we cannot *a priori* assume $\mathbf{A} = \mathbf{0}$, this does not constitute a full proof.

Before proceeding, we give a lemma that will be used to lower bound $R(\mathbf{f})$ (assuming $\mathbf{A} = \mathbf{0}$).

Lemma 10. *Suppose \mathbf{g} is a sum of ReLU units such that $\mathbf{g} = \mathbf{0}$ on a closed convex region $C \subset \mathbb{R}^d$ and \mathbf{g} is constant and equal to \mathbf{c} on a closed ball $B \subset \mathbb{R}^d$, and \mathbf{g} has a minimal representative where all its units are active on B . Then*

$$R(\mathbf{g}) \geq \frac{2\|\mathbf{c}\|}{\text{dist}(B, C)}$$

where $\text{dist}(B, C) = \min_{\mathbf{y} \in B, \mathbf{x} \in C} \|\mathbf{y} - \mathbf{x}\|$.

Proof. Let $\mathbf{g}(\mathbf{y}) = \sum_k \mathbf{a}_k [\mathbf{w}_k^\top \mathbf{y} + b_k]_+$ be a minimal representative where each unit is active on B . For \mathbf{g} to be constant on B it must be the case that $\sum_k \mathbf{a}_k \mathbf{w}_k^\top = \mathbf{0}$. Let $\mathbf{y} = \mathbf{y}_0$ be any value for which $\|\partial \mathbf{g}(\mathbf{y}_0)\|$, the operator norm of the Jacobian of \mathbf{g} , is maximized. Since \mathbf{g} is piecewise linear, we see that its Lipschitz constant $\text{Lip}(\mathbf{g}) = \|\partial \mathbf{g}(\mathbf{y}_0)\|$. Let I_0 be the set of indices of active units at \mathbf{y}_0 so that $\partial \mathbf{g}(\mathbf{y}_0) = \sum_{k \in I_0} \mathbf{a}_k \mathbf{w}_k^\top$, and let I_1 be the complementary index set. Then because $\sum_k \mathbf{a}_k \mathbf{w}_k^\top = \mathbf{0}$, we have $\sum_{k \in I_0} \mathbf{a}_k \mathbf{w}_k^\top = -\sum_{k \in I_1} \mathbf{a}_k \mathbf{w}_k^\top$, and so $\|\sum_{k \in I_0} \mathbf{a}_k \mathbf{w}_k^\top\| = \|\sum_{k \in I_1} \mathbf{a}_k \mathbf{w}_k^\top\|$. Therefore,

$$\begin{aligned} 2\text{Lip}(\mathbf{g}) &= 2\|\partial \mathbf{g}(\mathbf{y}_0)\| = 2\left\| \sum_{k \in I_0} \mathbf{a}_k \mathbf{w}_k^\top \right\| = \left\| \sum_{k \in I_0} \mathbf{a}_k \mathbf{w}_k^\top \right\| + \left\| \sum_{k \in I_1} \mathbf{a}_k \mathbf{w}_k^\top \right\| \\ &\leq \sum_k \|\mathbf{a}_k\| \|\mathbf{w}_k\| = R(\mathbf{g}). \end{aligned}$$

Finally,

$$\text{Lip}(\mathbf{g}) \geq \max_{\mathbf{x} \in C, \mathbf{y} \in B} \frac{\|\mathbf{g}(\mathbf{y}) - \mathbf{g}(\mathbf{x})\|}{\|\mathbf{y} - \mathbf{x}\|} = \frac{\|\mathbf{c}\|}{\min_{\mathbf{x} \in C, \mathbf{y} \in B} \|\mathbf{y} - \mathbf{x}\|} = \frac{\|\mathbf{c}\|}{\text{dist}(B, C)}.$$

Combining this with the previous inequality gives the claim. \square

Now, returning to the proof of the main claim, for $n = 1, 2, 3$, let C_n be the closed convex region given by the intersection of all closed half-planes containing the balls B_j and B_k , $j \neq n$, and $k \neq n$, $j \neq k$. By assumption, \mathbf{f}_n vanishes on C_n and \mathbf{f}_n is constant and equal to $\mathbf{x}_n - \mathbf{b}$ on the closed ball B_n . So, by Lemma 10, we see that

$$R(\mathbf{f}_n) \geq \frac{2\|\mathbf{x}_n - \mathbf{b}\|}{\delta_n}$$

where $\delta_n = \text{dist}(B_n, C_n)$. Therefore, for all $\mathbf{b} \in \mathbb{R}^d$ we have

$$R(\mathbf{f}) = R(\mathbf{f}_1) + R(\mathbf{f}_2) + R(\mathbf{f}_3) \geq \sum_{n=1}^3 \frac{2\|\mathbf{x}_n - \mathbf{b}\|}{\delta_n}$$

Let $\bar{\mathbf{x}} \in \mathbb{R}^d$ be the minimizer of

$$\min_{\mathbf{b} \in \mathbb{R}^d} \sum_{n=1}^3 \frac{2\|\mathbf{x}_n - \mathbf{b}\|}{\delta_n}.$$

Then plugging in $\mathbf{b} = \bar{\mathbf{x}}$ above, we have the lower bound

$$R(\mathbf{f}) \geq \sum_{n=1}^3 \frac{2\|\mathbf{x}_n - \bar{\mathbf{x}}\|}{\delta_n},$$

which is independent of \mathbf{b} .

Finally, simple calculations show that the conjectured min-cost solution \mathbf{f}^* specified in (86) has representation cost achieving this lower bound. Therefore, under the assumption $\mathbf{A} = \mathbf{0}$, \mathbf{f}^* is a min-cost solution. □

Now we prove that, in the special case where the convex hull of the training points is an equilateral triangle, the norm-ball interpolator \mathbf{f}^* identified in Conjecture 1 is a min-cost solution (though not necessarily the unique min-cost solution). In this case, we may also give \mathbf{f}^* a more explicit form, as detailed below.

Proposition 4. *Suppose the convex hull of the training points $\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3 \in \mathbb{R}^d$ is an equilateral triangle. Assume the norm-balls $B_n := B(\mathbf{x}_n, \rho)$ centered at each training point have radius $\rho < \|\mathbf{x}_n - \mathbf{x}_0\|/2$, $n = 1, 2, 3$, where $\mathbf{x}_0 = \frac{1}{3}(\mathbf{x}_1 + \mathbf{x}_2 + \mathbf{x}_3)$ is the centroid of the triangle. Then a minimizer \mathbf{f}^* of (19) is given by*

$$\mathbf{f}^*(\mathbf{y}) = \mathbf{u}_1 \phi_1(\mathbf{u}_1^\top(\mathbf{y} - \mathbf{x}_0)) + \mathbf{u}_2 \phi_2(\mathbf{u}_2^\top(\mathbf{y} - \mathbf{x}_0)) + \mathbf{u}_3 \phi_3(\mathbf{u}_3^\top(\mathbf{y} - \mathbf{x}_0)) + \mathbf{x}_0, \quad (86)$$

where $\phi_n(t) = s_n([t - a_n]_+ - [t - b_n]_+)$ with $\mathbf{u}_n = \frac{\mathbf{x}_n - \mathbf{x}_0}{\|\mathbf{x}_n - \mathbf{x}_0\|}$, $a_n = -\frac{1}{2}\|\mathbf{x}_n - \mathbf{x}_0\| + \rho$, $b_n = \|\mathbf{x}_n - \mathbf{x}_0\| - \rho$, and $s_n = \|\mathbf{x}_n - \mathbf{x}_0\|/(b_n - a_n)$.

Proof. By translation and scale invariance of min-cost solutions, without loss of generality we may assume $\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3 \in \mathbb{R}^d$ are unit-norm vectors and mean-zero (i.e., the triangle centroid $\mathbf{x}_0 = \mathbf{0}$). In this case, the assumption on ρ translates to $\rho < 1/2$. Additionally, it suffices to prove the claim in the case $d = 2$. This is because, if $\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3 \in \mathbb{R}^d$ are the vertices of an equilateral triangle whose centroid is at the origin, then these points are contained in a two-dimensional subspace $\mathcal{S} \subset \mathbb{R}^d$. And if we let $\mathbf{P} \in \mathbb{R}^{d \times 2}$ be a matrix whose columns are an orthonormal basis for \mathcal{S} , then by Theorem 2, \mathbf{f} is a min-cost solution if and only if $\mathbf{f}(\mathbf{y}) = \mathbf{P}\mathbf{f}_0(\mathbf{P}^\top \mathbf{y})$, where $\mathbf{f}_0 : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ is a min-cost solution under the constraints $\mathbf{f}_0(B(\mathbf{P}^\top \mathbf{x}_n, \rho)) = \{\mathbf{P}^\top \mathbf{x}_n\}$ for all $n = 1, 2, 3$. Therefore, the problem reduces to finding a min-cost solution of the projected points $\mathbf{P}^\top \mathbf{x}_1, \mathbf{P}^\top \mathbf{x}_2, \mathbf{P}^\top \mathbf{x}_3 \in \mathbb{R}^2$ whose convex hull is an equilateral triangle in \mathbb{R}^2 .

So now let $\mathbf{f} : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be any min-cost solution under the assumption $\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3 \in \mathbb{R}^2$ are unit-norm, have zero mean, and $\rho < 1/2$. By reversing units as necessary, there exists a minimal representative of \mathbf{f} that can be put in the form $\mathbf{f}(\mathbf{y}) = \mathbf{f}_1(\mathbf{y}) + \mathbf{f}_2(\mathbf{y}) + \mathbf{f}_3(\mathbf{y}) + \mathbf{A}\mathbf{y} + \mathbf{c}$, such that $\mathbf{f}_n(\mathbf{y})$ is a sum of ReLU units, all of which are active on B_n , and all of which are inactive on B_j for $j \neq n$.

Let \mathbf{Q} be the reflection matrix $\mathbf{Q} = 2\mathbf{x}_1\mathbf{x}_1^\top - \mathbf{I}$, which reflects points across the line spanned by \mathbf{x}_1 . In particular, $\mathbf{y} \in B_1$ then $\mathbf{Q}\mathbf{y} \in B_1$ while if $\mathbf{y} \in B_2$ then $\mathbf{Q}\mathbf{y} \in B_3$ (and vice versa). Also, $\mathbf{Q}\mathbf{x}_1 = \mathbf{x}_1$, $\mathbf{Q}\mathbf{x}_2 = \mathbf{x}_3$, and $\mathbf{Q}\mathbf{x}_3 = \mathbf{x}_2$. Define $\mathbf{g}(\mathbf{y}) = \frac{1}{2}(\mathbf{f}(\mathbf{y}) + \mathbf{Q}^{-1}\mathbf{f}(\mathbf{Q}\mathbf{y}))$. Then it is easy to check that \mathbf{g} satisfies interpolation constraints, and since \mathbf{Q} is unitary, we have $R(\mathbf{g}) \leq \frac{1}{2}R(\mathbf{f}) + \frac{1}{2}R(\mathbf{Q} \circ \mathbf{f} \circ \mathbf{Q}^{-1}) = \frac{1}{2}R(\mathbf{f}) + \frac{1}{2}R(\mathbf{f}) = R(\mathbf{f})$. Furthermore, since \mathbf{f} is a min-cost solution, we must have $R(\mathbf{g}) = R(\mathbf{f})$. Additionally, since none of the units making up \mathbf{f} are active over more than one ball, neither are the units making up \mathbf{g} , which implies no pair of units belonging to \mathbf{g} combine to form an affine function. Therefore, we may write $\mathbf{g}(\mathbf{y}) = \mathbf{g}_1(\mathbf{y}) + \mathbf{g}_2(\mathbf{y}) + \mathbf{g}_3(\mathbf{y}) + \mathbf{B}\mathbf{y} + \mathbf{v}$, where $\mathbf{B} = \mathbf{A} + \mathbf{Q}^{-1}\mathbf{A}\mathbf{Q}$, such that each \mathbf{g}_n is a sum of ReLU units, all of which are active on B_n , and all of which are inactive on B_j for $j \neq n$.

Let \mathbf{U} be a rotation matrix by 120 degrees such that $\mathbf{x}_2 = \mathbf{U}\mathbf{x}_1$, $\mathbf{x}_3 = \mathbf{U}\mathbf{x}_2$, and $\mathbf{x}_1 = \mathbf{U}\mathbf{x}_3$. Consider the symmetrized version of \mathbf{g} given by $\mathbf{h}(\mathbf{y}) := \frac{1}{3}(\mathbf{g}(\mathbf{y}) + \mathbf{U}^{-1}\mathbf{g}(\mathbf{U}\mathbf{y}) + \mathbf{U}^{-2}\mathbf{g}(\mathbf{U}^2\mathbf{y}))$.

Since for all $\mathbf{y} \in B_n$ we have $U\mathbf{y} \in B_{n+1}$ (with indices understood modulo 3), it is easy to verify that \mathbf{h} satisfies interpolation constraints. Also, since U is unitary, we have $R(\mathbf{h}) \leq R(\mathbf{g}) \leq R(\mathbf{f})$, which implies $R(\mathbf{h}) = R(\mathbf{f})$ since \mathbf{f} is a min-cost solution. Again, since none of the units making up \mathbf{g} are active over more than one ball, neither are the units making up \mathbf{h} , which implies no pair of units belonging to \mathbf{h} combine to form an affine function. And so, we may write $\mathbf{h}(\mathbf{y}) = \mathbf{h}_1(\mathbf{y}) + \mathbf{h}_2(\mathbf{y}) + \mathbf{h}_3(\mathbf{y}) + \mathbf{C} + \mathbf{u}$, such that each \mathbf{h}_n is a sum of ReLU units, all of which are active on B_n , and all of which are inactive on B_j for $j \neq n$.

Observe that $\mathbf{u} = \frac{1}{3}(\mathbf{I} + U^{-1} + U^{-2})\mathbf{v} = \mathbf{0}$. Also, we have $\mathbf{C} = \frac{1}{3}(\mathbf{B} + U^{-1}\mathbf{B}U + U^{-2}\mathbf{B}U^2)$. This implies \mathbf{C} commutes with U , because it is easy to see that $\mathbf{C} = U^{-1}\mathbf{C}U$. Additionally, \mathbf{C} commutes with Q , because it is easy to see $\mathbf{B} = Q^{-1}\mathbf{B}Q$, and by properties of rotations/reflections we have $UQ = QU^{-1}$, which implies

$$\begin{aligned} Q^{-1}\mathbf{C}Q &= \frac{1}{3}(Q^{-1}\mathbf{B}Q + Q^{-1}U^{-1}\mathbf{B}UQ + Q^{-1}U^{-2}\mathbf{B}U^2Q) \\ &= \frac{1}{3}(Q^{-1}\mathbf{B}Q + Q^{-1}U^{-1}\mathbf{B}UQ + Q^{-1}U^{-2}\mathbf{B}U^2Q) \\ &= \frac{1}{3}(Q^{-1}\mathbf{B}Q + (UQ)^{-1}\mathbf{B}(UQ) + (UQ)^{-1}U^{-1}\mathbf{B}U(UQ)) \\ &= \frac{1}{3}(Q^{-1}\mathbf{B}Q + UQ^{-1}\mathbf{B}QU^{-1} + UQ^{-1}U^{-1}\mathbf{B}UQU^{-1}) \\ &= \frac{1}{3}(Q^{-1}\mathbf{B}Q + UQ^{-1}\mathbf{B}QU^{-1} + U^2Q^{-1}\mathbf{B}QU^{-2}) \\ &= \frac{1}{3}(\mathbf{B} + \mathbf{B}U^{-1} + U^2\mathbf{B}U^{-2}) \\ &= \frac{1}{3}(\mathbf{B} + U^{-2}\mathbf{B}U^2 + U^{-1}\mathbf{B}U) \\ &= \mathbf{C}, \end{aligned}$$

which shows \mathbf{C} commutes with Q . Now we show that any matrix \mathbf{C} that commutes with both U and Q is a scaled identity matrix. Let \mathbf{x}_1^\perp denote a unit-vector perpendicular to \mathbf{x}_1 , such that $\{\mathbf{x}_1, \mathbf{x}_1^\perp\}$ form an orthonormal basis for \mathbb{R}^2 . First, we know Q has eigenvectors \mathbf{x}_1 and \mathbf{x}_1^\perp with eigenvalues 1 and -1 , respectively. And so $QC\mathbf{x}_1 = CQ\mathbf{x}_1 = C\mathbf{x}_1$, while $QC\mathbf{x}_1^\perp = CQ\mathbf{x}_1^\perp = -C\mathbf{x}_1^\perp$, which implies $C\mathbf{x}_1 = a\mathbf{x}_1$ and $C\mathbf{x}_1^\perp = b\mathbf{x}_1^\perp$ for some $a, b \in \mathbb{R}$, i.e., \mathbf{x}_1 and \mathbf{x}_1^\perp are eigenvectors of \mathbf{C} with eigenvalues a and b , respectively. Furthermore, we have $CU\mathbf{x}_1 = UC\mathbf{x}_1 = aU\mathbf{x}_1$ and $CU\mathbf{x}_1^\perp = UC\mathbf{x}_1^\perp = bU\mathbf{x}_1^\perp$. This implies $U\mathbf{x}_1$ and $U\mathbf{x}_1^\perp$ are also eigenvectors of \mathbf{C} with eigenvalues a and b , respectively. But since \mathbf{x}_1 and $U\mathbf{x}_1$ are linearly independent, and likewise so are \mathbf{x}_1^\perp and $U\mathbf{x}_1^\perp$, the only way this could be true is if $a = b$, i.e., every vector in \mathbb{R}^2 is an eigenvector of \mathbf{C} for the same eigenvalue $\lambda \in \mathbb{R}$, which implies $\mathbf{C} = \lambda\mathbf{I}$ for some $\lambda \in \mathbb{R}$.




Figure 9: Illustration of ReLU boundaries (in red) of original interpolant \mathbf{f} (left), the function \mathbf{g} (middle) obtained by enforcing reflection symmetry, and the function \mathbf{h} (right) obtained by enforcing both reflection and rotation symmetry.

Therefore, we have shown that every min-cost solution \mathbf{f} maps to a min-cost solution \mathbf{h} of the form

$$\mathbf{h}(\mathbf{y}) = \mathbf{h}_1(\mathbf{y}) + \mathbf{h}_2(\mathbf{y}) + \mathbf{h}_3(\mathbf{y}) + \lambda\mathbf{I}$$

for some $\lambda \in \mathbb{R}$, where each \mathbf{h}_n is a sum of ReLU units, all of which are active on B_n and all of which are inactive on B_j , $j \neq n$. In particular, $\mathbf{h}_n(\mathbf{y}) = \mathbf{x}_n - \lambda \mathbf{y}$ for all $\mathbf{y} \in B_n$ and $\mathbf{h}_n(\mathbf{y}) = \mathbf{0}$ for all $\mathbf{y} \in C_n$, where C_n is the intersection of all half-planes containing the balls B_j , $j \neq n$.

This means we can write $\mathbf{h}(\mathbf{y}) = \sum_{k=1}^K \mathbf{a}_k [\mathbf{w}_k^\top \mathbf{y} + b_k]_+ + \lambda \mathbf{I}$ and $\mathbf{h}_n(\mathbf{y}) = \sum_{k \in \mathcal{A}_n} \mathbf{a}_k [\mathbf{w}_k^\top \mathbf{y} + b_k]_+$ for some index sets $\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3$ partitioning $\{1, \dots, K\}$ so that $R(\mathbf{h}) = \sum_{k=1}^K \|\mathbf{a}_k\| = \sum_{n=1}^3 \sum_{k \in \mathcal{A}_n} \|\mathbf{a}_k\|$. Let $R_0(\cdot)$ denote the representation cost of a function computed without an unregularized linear part, i.e., define R_0 analogously to R except where the model class \mathbf{h}_θ in (6) is constrained to have $\mathbf{V} = \mathbf{0}$. Then, since the realizations of the \mathbf{h}_n functions considered above do not have a linear part, we see that $\sum_{k \in \mathcal{A}_n} \|\mathbf{a}_k\| \geq R_0(\mathbf{h}_n)$ and so $R(\mathbf{h}) = \sum_{n=1}^3 \sum_{k \in \mathcal{A}_n} \|\mathbf{a}_k\| \geq R_0(\mathbf{h}_1) + R_0(\mathbf{h}_2) + R_0(\mathbf{h}_3)$.

Now we show how lower bound $R_0(\mathbf{h}_n)$ for all $n = 1, 2, 3$. Let \mathbf{x}_n^\perp denote a unit-vector perpendicular to \mathbf{x}_n , such that $\{\mathbf{x}_n, \mathbf{x}_n^\perp\}$ form an orthonormal basis for \mathbb{R}^2 . Consider the univariate functions $h_n^\parallel(t) := \mathbf{x}_n^\top \mathbf{h}_n(\mathbf{x}_n t)$ and $h_n^\perp(t) := (\mathbf{x}_n^\perp)^\top \mathbf{h}_n(\mathbf{x}_n^\perp t + \mathbf{x}_n)$. Here h_n^\parallel is the projection of \mathbf{h} onto the line spanned by \mathbf{x}_n , and h_n^\perp is a projection onto the line perpendicular to \mathbf{x}_n passing through the point \mathbf{x}_n . In particular, by the constraints on \mathbf{h}_n , we see that h_n^\parallel and h_n^\perp satisfy the constraints

$$h_n^\parallel(t) = \begin{cases} 0 & \text{if } t < -1/2 + \rho \\ 1 - \lambda t & \text{if } t > 1 - \rho \end{cases} \quad (87)$$

and $h_n^\perp(t) = -\lambda t$ if $|t| \leq \rho$.

Claim 1: For all $n = 1, 2, 3$, $R_0(\mathbf{h}_n) \geq R_0(h_n^\parallel) + R_0(h_n^\perp)$. *Proof:* Let $\mathbf{h}_n(\mathbf{y}) = \sum_k \mathbf{a}_k [\mathbf{w}_k^\top \mathbf{y} + b_k]_+ + \mathbf{c}$ be any realization of \mathbf{h}_n , whose representation cost is $C = \sum_k \frac{1}{2} (\|\mathbf{a}_k\|^2 + \|\mathbf{w}_k\|^2)$. Then realizations of h_n^\parallel and h_n^\perp are given by

$$h_n^\parallel(t) = \mathbf{x}_n^\top \mathbf{h}_i(\mathbf{x}_i t) = \sum_k (\mathbf{x}_n^\top \mathbf{a}_k) [(\mathbf{x}_n^\top \mathbf{w}_k) t + b_k]_+ + \mathbf{x}_n^\top \mathbf{c}. \quad (88)$$

$$h_n^\perp(t) = (\mathbf{x}_n^\perp)^\top \mathbf{h}_n(\mathbf{x}_n^\perp t + \mathbf{x}_n) = \sum_k ((\mathbf{x}_n^\perp)^\top \mathbf{a}_k) [((\mathbf{x}_n^\perp)^\top \mathbf{w}_k) t + b_k + \mathbf{w}_k^\top \mathbf{x}_n]_+ + (\mathbf{x}_n^\perp)^\top \mathbf{c}, \quad (89)$$

whose representation costs C^\parallel and C^\perp , respectively, are given by

$$C^\parallel = \sum_k \frac{1}{2} ((\mathbf{x}_n^\top \mathbf{a}_k)^2 + (\mathbf{x}_n^\top \mathbf{w}_k)^2) \quad (90)$$

$$C^\perp = \sum_k \frac{1}{2} (((\mathbf{x}_n^\perp)^\top \mathbf{a}_k)^2 + ((\mathbf{x}_n^\perp)^\top \mathbf{w}_k)^2) \quad (91)$$

and by the Pythagorean Theorem we see that $C = C^\parallel + C^\perp$. Therefore, $C \geq R_0(h_n^\parallel) + R_0(h_n^\perp)$ and finally minimizing over all realizations of \mathbf{h}_n gives the claim.

Claim 2: For all $n = 1, 2, 3$, $R_0(h_n^\parallel) \geq \left| \frac{1-\lambda\beta}{\beta-\alpha} \right| + \left| \frac{1-\lambda\alpha}{\beta-\alpha} \right|$, where $\beta = 1 - \rho$ and $\alpha = -1/2 + \rho$.

Proof: By results in Savarese et al. [2019], $R_0(h_n^\parallel) \geq R_0(p)$ where $p(t)$ is the function satisfying the same constraints as h_n^\parallel given in (87) while linearly interpolating over the interval $t \in [\alpha, \beta]$. From the formula $R_0(p) = \max\{\int |p''(t)| dt, |p'(\infty)| + |p'(-\infty)|\}$ as established in Savarese et al. [2019], we can show directly that $R_0(p) = \left| \frac{1-\lambda\beta}{\beta-\alpha} \right| + \left| \frac{1-\lambda\alpha}{\beta-\alpha} \right|$, which gives the claimed bound.

Claim 3: For all $n = 1, 2, 3$, $R_0(h_n^\perp) \geq |\lambda|$. *Proof:* The function $q(t)$ with minimal R_0 -cost satisfying the constraint $q(t) = -\lambda t$ for $|t| \leq \rho$ is a single ReLU unit plus a constant: $q(t) = -\lambda[t + \rho]_+ + \lambda\rho$, which has R_0 -cost $|\lambda|$.

Putting the above claims together, we see that

$$R_0(\mathbf{h}_n) \geq \left| \frac{1 - \lambda\beta}{\beta - \alpha} \right| + \left| \frac{1 - \lambda\alpha}{\beta - \alpha} \right| + |\lambda| \quad (92)$$

$$\geq \left| \frac{2 - \lambda(\beta + \alpha)}{\beta - \alpha} \right| + |\lambda| \quad (93)$$

$$\geq \frac{2 - |\lambda|(\beta + \alpha)}{\beta - \alpha} + |\lambda| \quad (94)$$

$$= \frac{2}{\beta - \alpha} + \frac{-2\alpha}{\beta - \alpha} |\lambda| \quad (95)$$

$$\geq \frac{2}{\beta - \alpha} \quad (96)$$

where in the last inequality we used the fact that $\frac{-2\alpha}{\beta - \alpha} > 0$ since $\alpha = -1/2 + \rho < 0$ and $\beta - \alpha = 3/2 - 2\rho > 0$.

Therefore, $R(\mathbf{f}) = R(\mathbf{h}) \geq R_0(\mathbf{h}_1) + R_0(\mathbf{h}_2) + R_0(\mathbf{h}_3) \geq \frac{6}{\beta - \alpha}$. Also, the function \mathbf{f}^* given by

$$\mathbf{f}^* = \mathbf{f}_1^* + \mathbf{f}_2^* + \mathbf{f}_3^*$$

where

$$\mathbf{f}_n^*(\mathbf{y}) = \frac{\mathbf{x}_n}{\beta - \alpha} ([\mathbf{x}_n^\top \mathbf{y} - \alpha]_+ - [\mathbf{x}_n^\top \mathbf{y} - \beta]_+)$$

satisfies norm-ball interpolation constraints and $R(\mathbf{f}^*) = \frac{6}{\beta - \alpha}$. Hence, \mathbf{f}^* is a min-cost solution. \square

D Additional simulations

We train a one-hidden-layer ReLU network *without* a skip connection using the setting we used in Figure 1. As can be seen from Figure 10 we get a similar result for the NN denoiser with and without the skip connection. Therefore in Appendix D.1 we use one-hidden-layer ReLU network *without* a skip connection.




Figure 10: **NN denoiser vs eMMSE denoiser.** We trained a one-hidden-layer ReLU network on a denoising task. The clean dataset has four points equally spaced in the interval $[-5, 5]$, and the noisy samples are generated by adding zero-mean Gaussian noise with $\sigma = 1.5$. We use $\lambda = 10^{-5}$ in both setting. The figure shows the denoiser output as a function of its input for: (1) NN denoiser trained online using (7) for $100K$ iterations, (2) NN denoiser trained offline using (8) with $M = 9000$ and $20K$ epochs, and (3) the eMMSE denoiser (4).

D.1 MNIST

We use the MNIST dataset to verify various properties. First, the offline and online solutions achieve approximately the same test MSE when trained on a subset of the MNIST dataset (Figure 11). Second, to show that the fact that NN denoiser does not converge to the eMMSE denoiser is not due to approximation error (Figure 12). Lastly, to present the critical noise level in which representation cost minimizer f_{1D}^* has strictly lower MSE than the eMMSE, for all smaller noise levels (Figure 13).




Figure 11: **Online setting vs offline setting for MNIST denoiser.** We train a one-hidden layer ReLU network on a subset of $N = 100$ MNIST images for $10K$ iterations. We use a Gaussian noise with zero mean and $\sigma = 0.1$.




Figure 12: **Test loss vs. layer width for MNIST denoiser.** We train a one hidden layer ReLU network on MNIST denoiser task using 7 for 93K iterations with a fixed learning rate. We use a Gaussian noise with zero mean and $\sigma = 0.1$. The figure shows the test loss vs. layer width.




Figure 13: **MSE vs. Noise std.** We train a one-hidden layer ReLU network on a subset of $N = 100$ MNIST images (the range of each pixel is $[0, 1]$) for 10K iterations with fixed LR. We use a Gaussian noise with zero mean. The figure shows the MSE vs noise std (σ) for NN denoiser (orange line) and for eMMSE denoiser (blue line). Note that the eMMSE is dependent on σ (4). For low noise levels, the eMMSE output is one of the training set images. For moderate noise levels, the eMMSE output is a weighted sum of the training set images. For high noise levels, the eMMSE output is the mean of the training set images.

D.2 Three non-colinear training samples

We show in Figures 14 and 15 that for $N = 3$ training points from the MNIST dataset that forming a triangle in $d = 2$ dimensions the empirical minimizer obtained using noisy samples and weight decay regularization agrees well with the form of the exact representation cost minimizer predicted by Proposition 2 and Conjecture 4.

D.3 Empirical validation of the subspace assumption

We validated that the following image datasets are (approximately) low rank:

- CIFAR10
- CINIC10
- Tiny ImageNet (a lower resolution version of ImageNet, enabling us to use SVD)
- BSD (a denoising benchmark composed from 128×1600 patches of size 40×40 cropped from 400 images [Zhang et al., 2017])




Figure 14: **Function space view of a denoiser, trained for 3 MNIST data points (Acute Angle).** Here, we compare empirical results (left) with our theoretical results (right). We compare the function-space view for inputs from the data plane, with respect to the model output in each of the data directions. For the empirical results, we choose 3 random MNIST data points under the same label, and under the condition that they form an acute triangle (64° , for this figure). We trained a single-layer FC ReLU network with linear residual connection for 1M epochs, with weight decay of $1E - 8$ (as described in our model), and ADAM optimizer with learning rate $1E - 5$.

Table 1: We applied a Singular Value Decomposition (SVD) for each of the above datasets, and calculated the relative number of Singular Values (SV) needed to achieve a given percentile of the energy (for the average vector).

Dataset	95%	99%	99.9%
CIFAR10	0.8%	7.5%	30%
CINIC10	1%	23%	41%
Tiny ImageNet	1.6%	20%	36%
BSD	0.1%	1.6%	4.5%

As can be seen from Table 1 all the datasets that we used are (approximately) low rank.




Figure 15: **Function space view of a denoiser, trained for 3 MNIST data points (Obtuse Angle).** Here, we compare empirical results (left) with our theoretical results (right). We compare the function-space view for inputs from the data plane, with respect to the model output in each of the data directions. For the empirical results, we added a single data point to the two previously chosen for Figure 14 under the condition that the three points form an obtuse triangle (95° , for this figure). We trained a single-layer FC as described in Figure 14. As predicted, the function we have converged to for data forming an obtuse triangle is noticeably different from the function we converged to when the data was forming an acute triangle.