

Unveiling Hidden Convexity in Deep Learning: A Sparse Signal Processing Perspective

Emi Zeger and Mert Pilanci

Department of Electrical Engineering, Stanford University

INTRODUCTION

Deep neural networks (DNNs), particularly those using Rectified Linear Unit (ReLU) activation functions, have achieved remarkable success across diverse machine learning tasks, including image recognition, audio processing, and language modeling [1]. Despite this success, the non-convex nature of DNN loss functions complicates optimization and limits theoretical understanding. These challenges are especially pertinent for signal processing applications where stability, robustness, and interpretability are crucial.

Various theoretical approaches have been developed for analyzing neural networks (NNs). In this paper, we highlight how recently developed convex equivalences of ReLU NNs and their connections to sparse signal processing models can address the challenges of training and understanding NNs. Recent research has uncovered several hidden convexities in the loss landscapes of certain NN architectures, notably two-layer ReLU networks and other deeper or varied architectures. By reframing the training process as a convex optimization task, it becomes possible to efficiently find globally optimal solutions. This approach offers new perspectives on the network's generalization and robustness characteristics while facilitating interpretability. Leveraging Lasso-type models and structure-inducing regularization frameworks, which are fundamental tools in sparse representation modeling and compressed sensing (CS), NN training can be approached as a convex optimization problem, enabling the interpretation of both globally and locally optimal solutions.

This paper seeks to provide an accessible and educational overview that bridges recent advances in the mathematics of deep learning with traditional signal processing, encouraging broader signal processing applications. The paper is organized as follows. We first give a brief background on NNs and approaches to analyzing them using convex optimization. We then present an equivalence theorem between a two-layer ReLU NN and a convex group Lasso problem. We describe how deeper networks and alternative architectures can also be formulated as convex problems, discuss geometric insights, and present experimental results demonstrating performance benefits of training NNs as convex models. Finally, we discuss remaining challenges and research directions for convex analysis of NNs.

Related tutorial literature: A number of recent works touch on aspects of regularized neural network training from a signal processing perspective. One such line of work, by Parhi and Nowak [2], studies neural weight balancing, analyzing how parameter magnitudes are distributed under a fixed total ℓ_2 norm constraint. While this viewpoint is related to the convex duality formulations referenced here, the present

article adopts a broader perspective that brings together multiple convex formulations and highlights their implications across a range of neural network models and learning settings. A distinct and complementary direction, explored in Sun et al. [3], investigates how regularization and the addition of special neurons can make the loss landscape benign, where local minima are globally optimal. While regularization is also key to improving the loss landscape in the present paper, we focus on approaches to guarantee that the NN loss is not just benign, but convex.

Notation: Denote $\mathbf{1}$ as the vector of ones and $[n] = \{1, \dots, n\}$. The Boolean function $1\{x\}$ returns 1 if x is true, and 0 otherwise. Functions and operations such as $1\{x\}$, \geq , and activations $\sigma(x)$ apply element-wise to vector inputs. The $d \times d$ identity matrix is \mathbf{I}_d . There are m hidden neurons and n training samples \mathbf{x}_i , each of dimension d , which are stacked column-wise into the data matrix $\mathbf{X} \in \mathbb{R}^{d \times n}$ of rank r . NNs $f : \mathbb{R}^d \rightarrow \mathbb{R}$ are defined for row vectors and extend to data matrix inputs as $f(\mathbf{X}) = (f(\mathbf{x}_1), \dots, f(\mathbf{x}_n))^T$.

BACKGROUND: NEURAL NETWORKS

NNs are parameterized functions used for supervised learning, composed of functions called neurons:

$$f_{\text{neuron}}(\mathbf{x}) = \sigma(\mathbf{x}^T \mathbf{w}) \quad (1)$$

mapping a vector input to a scalar, where $\mathbf{w} \in \mathbb{R}^d$ is a *weight* parameter and $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ is a nonlinear *activation* function. A common activation function is the Rectified Linear Unit (ReLU): $\sigma(x) = \max\{x, 0\}$. An activation function is *active* at x if $\sigma(x) \neq 0$. The activation σ is (positively) *homogeneous* (of degree one) if $\sigma(ax) = a\sigma(x)$ for $a > 0$ [4]. For example, non-negative inputs activate a ReLU. The nonlinearity of an activation function distinguishes a neuron from a traditional linear model. The neuron (1) is inspired by a biological neuron in the brain, which receives synaptic inputs (represented by \mathbf{x}) whose intensities are modulated by the number of receptors (\mathbf{w}), and then fires action potentials as outputs. In the brain, neurons can operate in series, feedback, and parallel pathways [5]. Motivated by these biological features, a *hidden layer* consists of m parallel neurons: $f_{\text{layer}} : \mathbb{R}^d \rightarrow \mathbb{R}^{1 \times m}$,

$$f_{\text{layer}}(\mathbf{x}) = \sigma(\mathbf{x}^T \mathbf{W}), \quad (2)$$

where $\mathbf{W} \in \mathbb{R}^{d \times m}$ is called a weight matrix. An L -layer NN consists of composing $L - 1$ nonlinear hidden layers (2), followed by a linear output layer. The depth is L , and the width of a layer is the number of neurons in that layer. A standard two-layer NN $f : \mathbb{R}^d \rightarrow \mathbb{R}$ takes the form

$$f(\mathbf{x}) = \sigma(\mathbf{x}^T \mathbf{W}) \boldsymbol{\alpha} \quad (3)$$

where (2) is the first hidden layer of (3), and $\boldsymbol{\alpha} \in \mathbb{R}^m$ is the weight vector of the outer linear layer. Bias parameters can be added implicitly by appending a 1 to \mathbf{x} and an extra row to \mathbf{W} . An L -layer NN is said to be shallow if $L = 2$ and deep if $L > 2$. The NN (3) is a two-layer, *fully connected, feed-forward* network. There are many types of NNs consisting of variations on (3), as well as deeper networks. NNs fit data to labels/targets based on known pairs $(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_n, y_n)$ that are given, where $y_i \in \mathbb{R}$ is the *target*

of training sample \mathbf{x}_i . A NN is *trained* by finding parameters (weights and biases) so that $f(\mathbf{x}_i) \approx y_i$, formulated as a *training problem*

$$\min_{\theta} \ell(f(\mathbf{X}), \mathbf{y}) + R(\theta) \quad (4)$$

where $\mathbf{X} \in \mathbb{R}^{d \times n}$ contains \mathbf{x}_i as columns, $\mathbf{y} \in \mathbb{R}^n$ contains y_i , ℓ is a convex loss function quantifying the error of the NN fit, and θ is the parameter set, e.g., $\theta = \{\mathbf{W}, \boldsymbol{\alpha}\}$ for two-layer NNs (2). R is a regularization function that penalizes large parameter magnitudes, encouraging simpler solutions, and will be specified for each NN. The training problem (4) is non-convex due to the product of inner and outer weights. This poses a major challenge for analyzing NNs. Traditional approaches to train NNs involve gradient descent algorithms, but these methods can converge to suboptimal local minima. In contrast, if the training loss is convex, all stationary points would be globally optimal. The next section discusses approaches to improve the loss landscape.

Improving the Loss Landscape

Beyond reducing training complexity, regularization can also improve the geometry of the training objective, or *loss landscape*. In this context, a common goal is to enforce *benign* behavior, meaning that optimization methods are less likely to encounter undesirable stationary points. A representative example, surveyed in [3] and references therein, shows that augmenting a network with additional “special” neurons together with appropriate regularization can eliminate suboptimal local minima (under suitable assumptions). These ideas extend beyond binary classification to multiclass and regression settings, and illustrate how architectural modifications and regularization can jointly shape the loss landscape. We refer the reader to [3] for precise statements and conditions.

Computational Complexity

A well-known result is that the global optimization of two-layer ReLU NNs requires exponential complexity in the input dimension d (see e.g. [6] and the references therein). A relatively less known result is that, even for fixed d , training is NP-hard, implying no known polynomial-time algorithm exists unless $P = NP$ [6]. However, this hardness result applies specifically to unregularized training. In contrast, as discussed below, convex formulations for networks incorporating various regularizations can circumvent NP-hardness, achieving polynomial-time complexity for fixed d [7]. While the convex formulations remain computationally demanding in high-dimensional settings, randomized approximation schemes [8] related to 1-bit compressed sensing [9] can guarantee near-global optimality. In addition to convex approximations, exact convex formulations for NNs enjoy close connections to sparse signal processing. The next section reviews some of these signal processing methods.

BACKGROUND ON SPARSE SIGNAL PROCESSING, COMPRESSED SENSING, AND GEOMETRIC ALGEBRA

Sparse signal processing analyzes signals that admit a sparse representation in some basis or dictionary. Such sparsity enables efficient storage, transmission, and processing of signals. A key application of

sparse signal processing is CS [10, 11], which studies how sparse signals can be recovered from far fewer measurements than traditionally required by Nyquist sampling theory. Various algorithms such as matching pursuit achieve sparsity; we focus on Lasso-like models for their connection to NNs.

Lasso

The Least Absolute Shrinkage and Selection Operator (Lasso) [12] is a convex model widely used in sparse signal processing. Lasso performs variable selection and regularization simultaneously, making it an effective tool for sparse representation in CS applications. The Lasso optimization problem is defined as:

$$\min_{\mathbf{z}} \frac{1}{2} \|\mathbf{A}\mathbf{z} - \mathbf{y}\|_2^2 + \beta \|\mathbf{z}\|_1 \quad (5)$$

where \mathbf{A} is the measurement (or dictionary) matrix whose columns $\mathbf{A}_i \in \mathbb{R}^n$ are called feature vectors or dictionary atoms, $\mathbf{y} \in \mathbb{R}^n$ is the observation vector, $\mathbf{z} \in \mathbb{R}^d$ and $\beta > 0$ is a regularization parameter controlling the sparsity, i.e., the number of non-zero elements in a solution for \mathbf{z} . The ℓ_1 -norm $\|\mathbf{z}\|_1$ induces a sparse support of \mathbf{z} by penalizing the absolute sum of the coefficients. In CS, Lasso serves as a convex relaxation of the NP-hard ℓ_0 -norm minimization problem, providing an efficient computational framework for inducing sparse representations and variable selection with theoretical guarantees under certain conditions on the measurement matrix \mathbf{A} .

Group Lasso

Group Lasso generalizes Lasso to promote sparse representations at the group level [13]. In many signal processing applications, signals exhibit group sparsity, where non-zero coefficients occur in clusters. Group Lasso accounts for this structure by grouping variables and using an ℓ_1/ℓ_2 -norm regularization:

$$\min_{\mathbf{z}} \frac{1}{2} \|\mathbf{A}\mathbf{z} - \mathbf{y}\|_2^2 + \beta \sum_{g=1}^G \|\mathbf{z}^{(g)}\|_2 \quad (6)$$

where $\mathbf{z}^{(g)}$ denotes the coefficients in group g , G is the number of groups, and $\|\mathbf{z}^{(g)}\|_2$ is the ℓ_2 -norm of the group coefficients. This encourages entire groups of coefficients to be zero, promoting structured sparsity that aligns with the underlying signal characteristics. When each group contains a single coefficient z_i , the term $\sum_{g=1}^G \|\mathbf{z}^{(g)}\|_2 = \sum_j |z_j|$ reduces to the ℓ_1 -norm, recovering the standard Lasso. Lasso and group Lasso models are used in CS, described next.

Compressed Sensing (CS)

Let $\mathbf{z} \in \mathbb{R}^d$ be an unknown signal that we can indirectly observe through linear measurements

$$y_i = \mathbf{z}^T \mathbf{a}_i \quad (7)$$

where $\mathbf{a}_i \in \mathbb{R}^d$ are known *measurement vectors*. This yields a system of equations $\mathbf{A}\mathbf{z} = \mathbf{y}$ where the measurement matrix $\mathbf{A} \in \mathbb{R}^{M \times d}$ has rows \mathbf{a}_i^T and the observation vector $\mathbf{y} \in \mathbb{R}^M$ contains corresponding

observations y_i . We choose the measurement vectors \mathbf{a}_i^T and the number of measurements M . If $M = d$ and the measurement vectors \mathbf{a}_i are linearly independent, then $\mathbf{z} = \mathbf{A}^{-1}\mathbf{y}$ can be uniquely recovered. However, in many applications (such as MRI image reconstruction), each measurement incurs costs such as latency or financial expense. CS estimates \mathbf{z} from fewer than d measurements, effectively compressing the signal by projecting it onto measurement vectors. The measurement vectors can be random, e.g., $\mathbf{a}_i \sim \mathcal{N}(0, \mathbf{I}_d)$ [11, 14]. With fewer than d measurements, multiple candidate solutions satisfy $\mathbf{Az} = \mathbf{y}$; CS selects a parsimonious solution by solving a Lasso (5) or group Lasso (6) depending on the sparsity structure of \mathbf{z} . A variant, 1-bit compressive sensing [9], uses quantized measurements $y_i = Q(\mathbf{z}^T \mathbf{a}_i)$, where $Q(x) = \text{sign}(x) \in \{-1, 1\}$, compressing each measurement to only one bit of information. Solving

$$\begin{aligned} \min_{\mathbf{z}} \|\mathbf{z}\|_1 \\ \text{s.t. } \text{Diag}(\mathbf{y})(\mathbf{Az}) \geq 0, \|\mathbf{z}\|_2 = 1 \end{aligned}$$

recovers \mathbf{z} , where the ℓ_2 constraint resolves ambiguity in the solution. Both CS and 1-bit CS models connect deeply to convex NNs as we will show in future sections.

Nuclear Norm Regularization

The *nuclear norm* of a matrix \mathbf{W} is the ℓ_1 -norm of its singular values: $\|\mathbf{W}\|_* = \sum_{k=1}^r \sigma_k(\mathbf{W}) = \|\sigma(\mathbf{W})\|_1$, where r is the rank of \mathbf{W} [15]. As a special case, the nuclear norm of a positive semidefinite matrix is its trace. The nuclear norm is often used in optimization problems to search for low-rank matrices [15]. We discuss two examples, robust PCA and matrix completion.

1) *Robust PCA*: The robust Principal Component Analysis (PCA) problem [16] for a matrix \mathbf{X} is

$$\min_{\mathbf{W}, \mathbf{S}: \mathbf{X} = \mathbf{S} + \mathbf{W}} \|\mathbf{W}\|_* + \beta \|\mathbf{S}\|_1. \quad (8)$$

The robust PCA problem is a convex heuristic to decompose \mathbf{X} into the sum of a low-rank matrix \mathbf{W} and a sparse matrix \mathbf{S} by penalizing the singular values of \mathbf{W} and all elements of \mathbf{S} . The low-rank matrix \mathbf{W} can represent the underlying low-dimensional subspace and the sparse matrix \mathbf{S} represents outliers.

2) *Matrix completion*: The matrix completion problem [17] is as follows. Given a $n \times n$ matrix where only some of the values are known, fill in the rest of the matrix so that it has the lowest rank possible, consistent with the given elements. This problem can be approximately solved by filling in the unknown values that give the lowest nuclear norm [17].

Geometric Algebra

Geometric algebra is a mathematical framework that generalizes complex numbers to higher-dimensional vector spaces, including quaternions and hypercomplex numbers. It provides a unified way to represent and perform operations on geometric objects [18]. Recent work has applied geometric algebra to develop new signal processing techniques for image, audio, and video processing [19], as well as deep learning by leveraging geometric representations to enable more compact and expressive models [19]. The following sections relate signal processing methods to convex optimization of NNs.

Convexity w.r.t:	Approach	Motivation/Technique	Relation to SP
outer-layer weights [20]	column generation	boosting	matching pursuit
parameters [7]	convex duality, Lasso	hyperplane arrangements zonotopes	sparse recovery, 1-bit CS [9]
probability measures [21, 22]	mean field integration	neural tangent kernel (NTK) [23]	kernel methods [24]
network output [4]	shift complexity to regularization	matrix factorization, nuclear norm [15]	CS: robust PCA [16] matrix completion [17]
quantized weights [25]	linear program	intersection graph analysis	graphical models [26]

TABLE I: Comparison of convex perspectives on NNs. Each approach is motivated by, or involves, techniques related to signal processing (SP). The approaches in the first two rows also involve ℓ_1 regularizations, linking them to sparse representations in signal processing.

A CONVEX VIEW OF SHALLOW NEURAL NETWORKS

Since convex functions have been well-studied theoretically, one major approach to study NNs is to reformulate them via convex optimization. Such convex formulations provide valuable intuition into the underlying structure of optimal network weights and the associated model classes. They enable rigorous analysis of optimality guarantees, principled stopping criteria, numerical stability, and reliability—all critical features for mission-critical applications common in signal processing, such as real-time control systems in autonomous vehicles, power grid management, and healthcare monitoring. Moreover, convex programs make training NNs agnostic to hyperparameters such as initialization and mini-batching, which often exert a significant influence on the performance of local optimization methods. We discuss various approaches to represent NN training problems as convex problems. This section assumes the networks are two-layer ReLU networks unless otherwise stated. Deeper networks follow in later sections.

Convexity with Respect to the Outer-Layer Weights

Bengio et al. [20] and Bach [27] show that training a two-layer, infinitely wide NN with a convex loss and sparsity-promoting regularization on the outer layer can be cast as a convex optimization problem. (In their setup, an ℓ_1 norm is immediate when the hidden-unit set is countable.) Let

$$\Phi = \{ \sigma(\mathbf{X}^\top \mathbf{w}) : \mathbf{w} \in \mathbb{R}^d \}$$

denote the (uncountable) set of hidden-layer outputs evaluated on the training data matrix $\mathbf{X} \in \mathbb{R}^{d \times n}$. Training reduces to selecting outer-layer weights α by solving [20, 27]

$$\min_{\alpha} \ell \left(\int_{\Phi} \phi d\alpha(\phi), \mathbf{y} \right) + \|\alpha\|_{\text{TV}}, \quad (9)$$

where ℓ is any convex loss, the integral integrates over all possible $\phi \in \Phi \subset \mathbb{R}^n$, the final-layer weight $\alpha : \Phi \rightarrow \mathbb{R}$ is a finite signed measure on Φ so that $\int_{\Phi} \phi d\alpha(\phi) \in \mathbb{R}^n$ is well-defined, and the regularization

is the total variation norm [27]. While [27] develops a function-space analysis of (9) by viewing the outer-layer weights as a signed measure over a continuous feature family and regularizing its total variation, the earlier work [20] highlights that the sparsity penalty admits an optimal solution supported on finitely many hidden units (at most $n + 1$, by their Theorem 3.1). In this case, the NN reduces to a discrete and finite linear combination of features ϕ , the continuous distribution $\alpha(\phi)$ in (9) is identified with a vector of coefficients $\alpha \in \mathbb{R}^m$, and the regularization $\|\alpha\|_{\text{TV}}$ simplifies to the usual ℓ_1 -norm penalty. Bengio et al. [20] adopts a *column generation* strategy to optimize this NN: begin with an empty hidden layer and iteratively (i) identify the neuron that most improves the objective and add it to the hidden layer (ii) re-optimize (9) over the current coefficients. Step (i) consists of solving a weighted linear classification problem, which is non-convex in general. With hinge loss, the procedure is shown to reach a global optimum in at most $n + 1$ iterations, corresponding to $n + 1$ neurons [20]. This incremental scheme links NNs to *boosting*—both build predictors by sequentially adding simple components—and, more broadly, to *greedy approximation* methods in signal processing (e.g., matching pursuit and related procedures) [28]. The extension to continuous distributions in [27] also connects naturally to *Frank–Wolfe* style incremental methods. The main computational bottleneck is the non-convex subproblem used to insert each neuron, which remains challenging in high dimensions. From a signal processing perspective, the features $\phi \in \Phi$ can be viewed as basis functions for the network [27]. The approach in [27] is used to analyze generalization properties of infinite-width, two-layer NNs with homogeneous activations (such as ReLU).

Convexity with Respect to the Network Output

In this section we review the convex formulation introduced by Haeffele and Vidal [4]. Their central idea is to treat the *entire network output* as the decision variable, which converts the learning task into a convex optimization problem. Accordingly, the two-layer network in (3) can be rewritten as

$$f(\mathbf{x}) = \sum_{i=1}^m \sigma(\mathbf{x}^T \mathbf{w}^{(i)}) \alpha_i \quad (10)$$

where $\mathbf{w}^{(i)} \in \mathbb{R}^d$ is the i^{th} column of weight matrix $\mathbf{W} \in \mathbb{R}^{d \times m}$ (2). For a wide class of regularization functions $R(\theta)$ including norm penalties, the training problem (4) is equivalent to the convex problem

$$\min_{\mathbf{F} \in \mathbb{R}^n} \ell(\mathbf{F}, \mathbf{y}) + \tilde{R}(\mathbf{F}) \quad (11)$$

where $\tilde{R}(\mathbf{F})$ is the minimum parameter regularization such that the corresponding network outputs \mathbf{F} [4]:

$$\tilde{R}(\mathbf{F}) = \min_{\theta} R(\theta) \text{ s.t. } f(\mathbf{X}) = \mathbf{F} \quad (12)$$

The modified regularization \tilde{R} (12) is convex [4], making (11) convex. In contrast to [4], the convex problem (11) is finite-dimensional, as the optimization variable \mathbf{F} is finite-dimensional. This convex reformulation shifts the complexity from optimizing over all weights to optimizing the regularization term. The approach

in [4] is motivated by matrix factorization, as follows. Suppose $\mathbf{A} \in \mathbb{R}^{d_1 \times r}$, $\mathbf{B} \in \mathbb{R}^{r \times d_2}$, $\mathbf{Y} \in \mathbb{R}^{d_1 \times d_2}$, where \mathbf{Y} , d_1 , d_2 are known and \mathbf{A} , \mathbf{B} , r are unknown. The matrix factorization problem

$$\min_{r, \mathbf{A}, \mathbf{B}} \frac{1}{2} \|\mathbf{AB} - \mathbf{Y}\|_F^2 + \beta(\|\mathbf{A}\|_F^2 + \|\mathbf{B}\|_F^2)$$

where the loss function between the matrices is $\ell(\mathbf{Z}, \mathbf{Y}) = \frac{1}{2} \|\mathbf{Z} - \mathbf{Y}\|_F^2$ and the regularization is $R(\mathbf{A}, \mathbf{B}) = \beta(\|\mathbf{A}\|_F^2 + \|\mathbf{B}\|_F^2)$, can be analyzed similarly to (11) as

$$\min_{\mathbf{F} \in \mathbb{R}^{d_1 \times d_2}} \frac{1}{2} \|\mathbf{F} - \mathbf{Y}\|_F^2 + \tilde{R}(\mathbf{F})$$

where similarly to (12), the modified regularization is

$$\tilde{R}(\mathbf{F}) = \min_{r, \mathbf{A}, \mathbf{B}} \beta(\|\mathbf{A}\|_F^2 + \|\mathbf{B}\|_F^2) \text{ s.t. } \mathbf{AB} = \mathbf{F}$$

For matrix factorization, this modified regularization turns out to be the nuclear norm: $\tilde{R}(\mathbf{F}) = \|\mathbf{F}\|_*$, a regularization employed in sparse signal processing (see *Nuclear Norm Regularization* above) [4]. However, for general NNs, the definition of $\tilde{R}(F)$ (12) is not explicit and leaves (11) intractable. Despite this limitation, convexity enables useful theoretical insights. Consider an augmented training problem where the network is expanded to include one more neuron. A local optimum of the original training problem is globally optimal if adding a neuron with zero-valued weights is locally optimal in the augmented training problem [4]. In particular, the loss landscape is benign [3] if for some $i \in [m]$, the i^{th} parameters are zero: $\mathbf{w}^{(i)} = \mathbf{0}$, $\alpha_i = 0$ in (10) [4]. This perspective on benign landscapes via the addition of neurons is related to the approaches discussed in [3]: see *Improving the Loss Landscape* above.

Building off of the benign landscape results, [4] describes an algorithm to find a globally optimal NN by repeating a local descent step that finds a local optimum and an expansion step that adds another neuron to the NN. A local descent path is a sequence of parameter values with a non-increasing objective in the non-convex training problem. If the network starts with at least n neurons (i.e., the number of training samples), a local descent path to a global optimum exists [4]. For a sufficiently smooth and strongly convex objective, gradient descent with sufficiently small step size would be an example of local descent [4]. However, for non-convex problems, finding a local descent path can be NP-hard [4]. In addition, the upper bound on the number of neurons needed to guarantee a local descent path can be large for practical algorithms [4]. A future area of work is to analyze cases where this can be reduced [4].

Convexity with Respect to Probability Measures

A complementary route is to lift the *inner* parameters and optimize over a *probability distribution* on weights (see e.g. Bach and Chizat [21], Mei et al. [22]). The starting point is to express a finite two-layer network of width m as an expectation. If p_m denotes the empirical measure that assigns mass $1/m$ to each parameter pair $(\mathbf{w}^{(i)}, \alpha_i)$, then a 2-layer NN can be represented as

$$f_m(\mathbf{x}) = \frac{1}{m} \sum_{i=1}^m \alpha_i \sigma(\mathbf{x}^\top \mathbf{w}^{(i)}) = \mathbb{E}_{(\mathbf{W}, \alpha) \sim p_m} [\sigma(\mathbf{x}^\top \mathbf{W}) \alpha] \quad (13)$$

Letting the width $m \rightarrow \infty$ while keeping the parameters bounded, the discrete measures p_m converge (along subsequences) to a probability distribution p on \mathbb{R}^{d+1} . The network output converges accordingly to the mean-field expression

$$f_p(\mathbf{x}) = \int \sigma(\mathbf{x}^\top \mathbf{W}) \boldsymbol{\alpha} dp(\mathbf{W}, \boldsymbol{\alpha}). \quad (14)$$

Because $f_p(\mathbf{x})$ depends *linearly* on p , any convex loss $\ell(f(\mathbf{X}), \mathbf{y})$ yields a convex training problem over probability measures,

$$\min_{p \in \mathcal{P}} \ell(f_p(\mathbf{X}), \mathbf{y}),$$

where \mathcal{P} denotes the set of all non-negative Borel measures on \mathbb{R}^{d+1} whose total mass is one. Therefore, the training problem is revealed as a convex optimization in function space.

Kernel (lazy-training) Regime and the Neural Tangent Kernel

Let $\phi_m(\mathbf{x}) = (\sigma(\mathbf{x}^\top \mathbf{w}^{(1)}), \dots, \sigma(\mathbf{x}^\top \mathbf{w}^{(m)}))^\top$. If we sample the inner weights i.i.d. from a random distribution, then *freeze* them and optimize only the outer coefficients, optimizing the discrete mean-field model (13) reduces to linear regression: $\min_{\boldsymbol{\alpha}} \frac{1}{2} \|\mathbf{y} - \frac{1}{m} \phi_m(\mathbf{X})^\top \boldsymbol{\alpha}\|_2^2$. Gradient descent for this problem is known as *kernel gradient descent* with *empirical kernel* $\hat{k}_m(\mathbf{x}, \mathbf{x}') = \frac{1}{m} \phi_m(\mathbf{x})^\top \phi_m(\mathbf{x}')$. As the width m tends to infinity the empirical kernel $\hat{k}_m(\mathbf{x}, \mathbf{x}') = \frac{1}{m} \sum_{i=1}^m \sigma(\mathbf{x}^\top \mathbf{w}^{(i)}) \sigma(\mathbf{x}'^\top \mathbf{w}^{(i)})$ converges to a deterministic limit $k_{\text{RF}}(\mathbf{x}, \mathbf{x}')$, where the subscript ‘RF’ denotes that this limit stems from random features. Jacot et al. [23] show that if one trains the *full* network but stays in the *lazy-training* regime, the Jacobian of the loss stays nearly constant and training is equivalent to kernel gradient descent with a (generally different but related) limiting kernel $k_{\text{NTK}}(\mathbf{x}, \mathbf{x}')$, called the *Neural Tangent Kernel* (NTK). The NTK connects NNs to kernel methods, a tool used in signal processing to lift signals into higher dimensions for linear separation, among other applications [24].

Convexity via Quantization and Polyhedral Representations

In another approach, Bienstock et al. [25] shows that a linear program can approximate NN training through weight quantization and polyhedral representations. This formulation assumes the training data lies within the box $[-1, 1]^d$, and then quantizes the data and weights to finite numbers of bits. This transforms the training problem into an optimization program with a linear objective, binary variables, and linear constraints, known as a binary program. An *intersection graph* represents the optimization variables of the binary program as nodes, and the constraints as edges between them. Relaxing the binary variables to continuous variables converts the binary program into a linear program. The treewidth of the intersection graph measures how closely the graph resembles a tree, defined as the minimal width among all possible tree decompositions of the graph. This treewidth characterizes the computational complexity involved in solving the linear program. Specifically, the binary program analyzed in [25] possesses low treewidth, leading to a linear program whose constraints grow linearly with the number of training samples. Therefore, global NN training can be approximated via low treewidth optimization. Treewidth indicates the

sparsity of the intersection graph relative to a tree structure, with exact trees having the lowest treewidth. In general, structured sparsity simplifies optimization; in this case, tree-structured sparsity restricts supports to rooted subtrees (connected subgraphs) of a tree. Such signal representations arise in graphical models for image and graph signal processing [26], and decision or regression tree models [28]. Although these problems are not typically expressed with explicit constraints, they are often formulated as optimization problems that induce implicit intersection graphs capturing parent-child branch dependencies among the signal components. NNs may therefore be interpreted as looser relatives of tree-sparse models with richer parameter interactions that still preserve tractability properties associated with low treewidth. In [25], the polyhedral constraint set captures the NN structure. Training on different datasets corresponds to solving linear programs constrained by distinct faces of this polyhedron. Importantly, this linear program representation does not require the loss function to be convex. This approach applies to both unregularized and regularized training, including weight norm penalties as regularization.

Convexity with Respect to the Network Parameters

The previous sections describe approaches that cast training problems as implicit convex programs, either in terms of the network output or hidden layer activations, or by modifying the weights through quantization, continuous distributions, or infinite-width limits. In contrast, a line of recent work [7, 18] introduces explicit convex formulations that establish direct correspondences between the variables of the convex programs and the network weights. These formulations involve a finite number of optimization variables, making them more representative of NNs in practice. These convex models closely parallel sparse modeling methods in signal processing such as Lasso, group Lasso, one-bit CS, and nuclear-norm-based matrix completion. This convex approach uncovers certain variable selection properties of weight-decay regularization, a characteristic of NNs also studied in [27]. Building on these ideas, Bai et al. [8] develops an adversarial training strategy, while Prakhya et al. [29] constructs semidefinite relaxations for vector-output, two-layer ReLU networks with ℓ_2 regularization. The mathematical proof techniques used in the proofs of these results are convex analytic in nature, and include convex geometry, polar duality and analysis of extreme points. We explore the key elements of this convex equivalence, starting with two-layer networks.

EXPLICIT CONVEX PROGRAMS FOR TWO-LAYER NETWORKS

This section examines two-layer NNs with scalar outputs and ReLU activations:

$$f(\mathbf{x}) = (\mathbf{x}^T \mathbf{W})_+ \alpha, \quad (15)$$

and presents an equivalent Lasso-type explicit formulation of the training problem that establishes a connection between sparsity-inducing signal processing and NNs. The Lasso equivalence gives both practical approaches for training and theoretical insights into NN representation power through ideas from

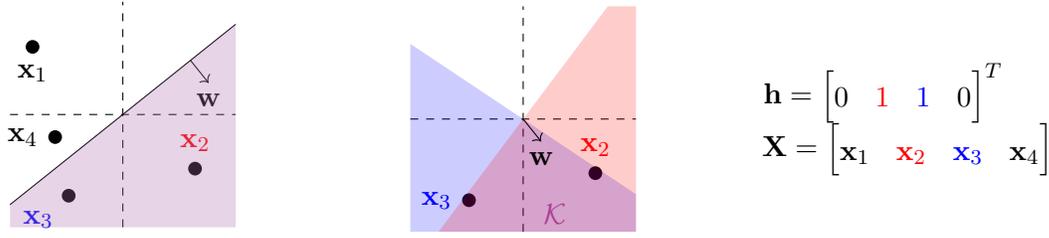


Fig. 1: Illustration of a separating hyperplane (left plot) and an activation chamber (right plot).

signal processing. The training problem (4) for the network (15) using ℓ_2 loss and ℓ_2 regularization is

$$\min_{\mathbf{W}, \alpha} \frac{1}{2} \|f(\mathbf{X}) - \mathbf{y}\|_2^2 + \beta (\|\mathbf{W}\|_F^2 + \|\alpha\|_2^2). \quad (16)$$

The term $\|\mathbf{W}\|_F^2 + \|\alpha\|_2^2$ represents *weight decay* regularization. We will recast the non-convex training problem as a convex group Lasso problem via hyperplane arrangements, which we define next.

Hyperplane Arrangement Patterns

A *hyperplane arrangement* is a finite collection of hyperplanes in a given vector space. For simplicity, we first consider hyperplane arrangements in \mathbb{R}^2 . In \mathbb{R}^2 , a *separation pattern* or *hyperplane arrangement pattern* of a dataset is a signed partition of the data by a line through the origin, called a separating line. A separation pattern is represented by a vector $\mathbf{h} \in \{0, 1\}^n$, where n is the number of training samples. Specifically, \mathbf{h} is a separation pattern if there exists a separating line such that each data point \mathbf{x}_i is assigned $h_i = 0$ if it lies on one side of the line and $h_i = 1$ if it lies on the opposite side. Each separating line has an associated orthogonal vector whose orientation determines which side corresponds to $h_i = 0$ versus $h_i = 1$. The orthogonal vectors associated with separating lines that give rise to the same separation pattern form a cone, called an *activation chamber*; an example is given below.

Example: Consider 2-D training data $\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3, \mathbf{x}_4$ shown in the left plot of Figure 1. An example of a separation (or hyperplane arrangement) pattern is $\mathbf{h} = (0, 1, 1, 0)^T$. This separation pattern corresponds to \mathbf{x}_2 and \mathbf{x}_3 being on the same side of a separating line of the form $\mathbf{w}^T \mathbf{x} = 0$. As shown in the figure, $\mathbf{x}_2, \mathbf{x}_3$ lie on the side of this line in the direction of \mathbf{w} (halfspace shaded in purple), and $\mathbf{x}_1, \mathbf{x}_4$ are on the other side. In the right plot, the magenta-shaded activation chamber \mathcal{K} illustrates the cone of all possible vectors \mathbf{w} that are orthogonal to separating lines producing pattern \mathbf{h} . The red and blue regions are the halfspaces $\mathbf{x}_2^T \mathbf{w} \geq 0$ and $\mathbf{x}_3^T \mathbf{w} \geq 0$, respectively, and the chamber \mathcal{K} is their intersection. The hyperplane arrangement definitions extend to higher dimensions by considering separating hyperplanes instead of separating lines. The set of *hyperplane arrangement patterns* of a data matrix $\mathbf{X} \in \mathbb{R}^{d \times n}$ is

$$\mathcal{H} = \left\{ \mathbf{1}\{\mathbf{X}^T \mathbf{w} \geq 0\} : \mathbf{w} \in \mathbb{R}^d \right\}. \quad (17)$$

Training data on the positive side of the hyperplane $\mathbf{X}^T \mathbf{w} = 0$ satisfy $\mathbf{x}^T \mathbf{w} \geq 0$ and thus activate the neuron $(\mathbf{x}^T \mathbf{w})_+$, while those on the negative side do not. The number of patterns is $|\mathcal{H}| \leq 2r \left(\frac{e(n-1)}{r}\right)^r$, where $r = \text{rank}(\mathbf{X}) \leq \min\{n, d\}$ and $e = 2.718\dots$ denotes Euler's constant [7]. Let $G = |\mathcal{H}|$ and enumerate the patterns as $\mathcal{H} = \{\mathbf{h}_1, \dots, \mathbf{h}_G\}$. The *activation chamber* for a pattern $\mathbf{h}_g \in \mathcal{H}$ is $\mathcal{K}^{(g)} = \{\mathbf{w} \in \mathbb{R}^d : 1\{\mathbf{X}^T \mathbf{w} \geq \mathbf{0}\} = \mathbf{h}_g\}$, which is the cone of all weights that induce an activation pattern \mathbf{h}_g . The *hyperplane encoding matrix* is $\mathbf{D}_g = \text{Diag}(\mathbf{h}_g)$ for $g = 1, \dots, G$. For fixed \mathbf{X} , a neuron's output $(\mathbf{X}^T \mathbf{w})_+$ as a function of \mathbf{w} is linear over $\mathbf{w} \in \mathcal{K}^{(g)}$, as $(\mathbf{X}^T \mathbf{w})_+ = \mathbf{D}_g \mathbf{X}^T \mathbf{w}$. The activation chambers partition $\mathbb{R}^d = \bigcup_{g=1}^G \mathcal{K}^{(g)}$, and so the matrices $\mathbf{D}_1, \dots, \mathbf{D}_G$ completely characterize the piecewise linearity of the neuron output $(\mathbf{X}^T \mathbf{w})_+$. For additional mathematical properties and a thorough introduction to the rich theory of hyperplane arrangements, we refer the reader to [30]. The training problem is *equivalent* to another optimization problem P if they share the same optimal value, and if an optimal network to the training problem can be reconstructed from an optimal solution to P .

Theorem 1 ([7]). *The non-convex training problem (16) with weight decay regularization for a two-layer ReLU network is equivalent to the convex group Lasso problem*

$$\min_{\mathbf{u}^{(g)}, \mathbf{v}^{(g)} \in \mathcal{K}^{(g)}} \frac{1}{2} \left\| \sum_{g=1}^G \mathbf{D}_g \mathbf{X}^T \left(\mathbf{u}^{(g)} - \mathbf{v}^{(g)} \right) - \mathbf{y} \right\|_2^2 + \beta \sum_{g=1}^G \left(\|\mathbf{u}^{(g)}\|_2 + \|\mathbf{v}^{(g)}\|_2 \right) \quad (18)$$

where $\mathcal{K}^{(g)} = \{\mathbf{z}^{(g)} : (2\mathbf{D}_g - \mathbf{I})\mathbf{X}^T \mathbf{z}^{(g)} \geq \mathbf{0}\}$, provided the number of neurons satisfies $m \geq m^*$ where m^* is the number of non-zero $\mathbf{u}^{(g)}, \mathbf{v}^{(g)}$.

Theorem 1 enables global optimization of ReLU NNs via convex optimization and an interpretation of the network as a signal processing model that promotes sparsity, with the dictionary $\mathbf{A} = [\mathbf{D}_1 \mathbf{X}^T, \dots, \mathbf{D}_G \mathbf{X}^T]$. The number m^* is a critical threshold on the number of neurons (number of columns of \mathbf{W}) necessary to allow the network to be sufficiently expressive to model the data. The group Lasso variable $\mathbf{z} \in \mathbb{R}^{Gd}$ is partitioned into G consecutive subvectors $\mathbf{z}^{(g)} = \mathbf{u}^{(g)} - \mathbf{v}^{(g)} \in \mathbb{R}^d$, each corresponding to a single neuron. Given an optimal solution to the group Lasso problem (18), the weights for an optimal NN are

$$\begin{aligned} \mathbf{w}^{(g)} &= \frac{\mathbf{u}^{(g)}}{\sqrt{\|\mathbf{u}^{(g)}\|_2}}, \alpha_g = \sqrt{\|\mathbf{u}^{(g)}\|_2} \\ \text{or } \mathbf{w}^{(g)} &= \frac{\mathbf{v}^{(g)}}{\sqrt{\|\mathbf{v}^{(g)}\|_2}}, \alpha_g = -\sqrt{\|\mathbf{v}^{(g)}\|_2} \end{aligned} \quad (19)$$

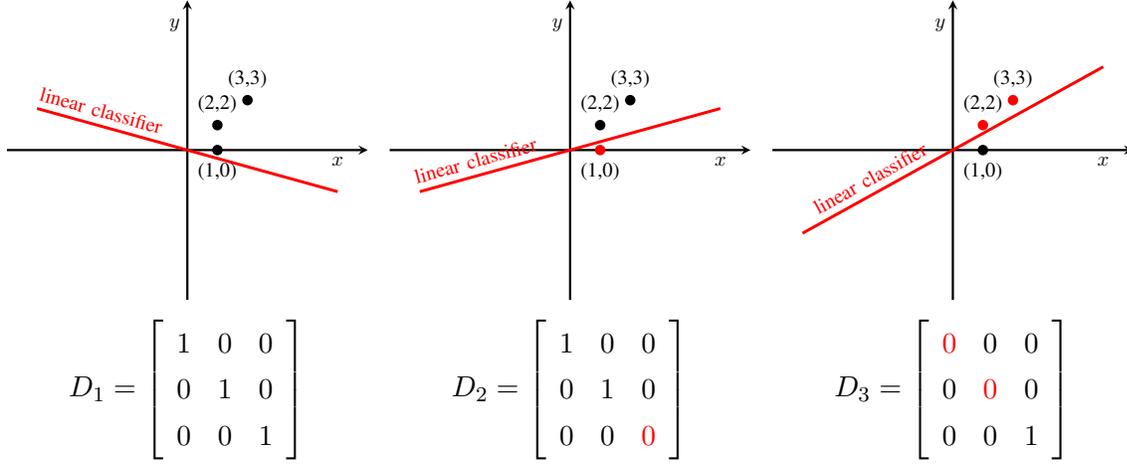
for all non-zero $\mathbf{u}^{(g)}, \mathbf{v}^{(g)}$. The weight vectors $\mathbf{w}^{(g)} \in \mathbb{R}^d$ are columns of the weight matrix $\mathbf{W} \in \mathbb{R}^{d \times m}$ and $\alpha_g \in \mathbb{R}$ are scalar components of $\boldsymbol{\alpha} \in \mathbb{R}^m$. The Lasso variables $\mathbf{u}^{(g)}$ and $\mathbf{v}^{(g)}$ represent weights corresponding to positive versus negative final-layer weights α_g .

Example [31]: Consider the training data matrix

$$\mathbf{X} = \begin{bmatrix} \mathbf{x}_1 & \mathbf{x}_2 & \mathbf{x}_3 \end{bmatrix} = \begin{bmatrix} 2 & 3 & 1 \\ 2 & 3 & 0 \end{bmatrix}. \quad (20)$$

Although there are $2^3 = 8$ distinct binary sequences of length 3, \mathbf{X} has $G = 3$ hyperplane arrangement

patterns in this case excluding the 0 matrix, as illustrated below from [31].



Consider training a two-layer ReLU NN on the data matrix in (20). For an arbitrary label vector $\mathbf{y} \in \mathbb{R}^3$ and the squared loss, the network in the equivalent convex program (18) is written piecewise linearly as

$$f(\mathbf{X}) = \mathbf{D}_1 \mathbf{X}^T (\mathbf{u}^{(1)} - \mathbf{v}^{(1)}) + \mathbf{D}_2 \mathbf{X}^T (\mathbf{u}^{(2)} - \mathbf{v}^{(2)}) + \mathbf{D}_3 \mathbf{X}^T (\mathbf{u}^{(3)} - \mathbf{v}^{(3)}). \quad (21)$$

This NN is interpretable from a signal processing perspective [13]: it looks for a group sparse model to explain the response \mathbf{y} via a mixture of linear models $\mathbf{u}^{(i)} - \mathbf{v}^{(i)}$ tuned to each activation chamber. The linear term $\mathbf{u}^{(2)} - \mathbf{v}^{(2)}$ predicts on $\{\mathbf{x}_1, \mathbf{x}_2\}$, and $\mathbf{u}^{(3)} - \mathbf{v}^{(3)}$ on $\{\mathbf{x}^{(3)}\}$, etc. Due to the regularization term $\sum_{g=1}^3 \|\mathbf{u}^{(g)}\|_2 + \|\mathbf{v}^{(g)}\|_2$ in (18), only a few of these linear terms will be non-zero at the optimum, showing a bias towards simple solutions among all piecewise linear models. The convex formulation therefore offers insights into the role of regularization in preventing overfitting. While the 2-D example (20) enumerates hyperplanes by hand, this becomes impractical in high dimensions and is addressed next.

Computational Complexity of Global Optimization and Randomized Sampling for Polynomial-time Approximation: Various approaches have analyzed the complexity resulting from training NNs with a convex program. The complexity of solving the convex problem (18) is proportional to $\binom{n}{d}^d$, where n is the number of samples and d is the dimension of the input. Although the exponential dependence with respect to d is unavoidable, the complexity of (18) is polynomial in the number of neurons and samples, which is significantly lower than brute-force search over the linear regions of ReLUs ($O(2^{md})$) [7]. By leveraging weight decay regularization, Theorem 1 overcomes the NP-hardness of unregularized NN training for fixed d [6]. In [4], the maximum width needed for a two-layer NN (trained with a broad class of regularizations) to reach global optimality corresponds to the number of local descent searches their algorithm must perform, which is linear in n . However, each local descent step can be NP-hard [4]. Theorem 1 requires a sufficient width of $m \geq m^*$ for the convex equivalence to hold. This is consistent with [6], which shows that training has polynomial time complexity when the number of neurons exceeds a critical width, and can become NP-hard otherwise. The exact width thresholds in [6] and Theorem 1

differ because Theorem 1 considers parameter regularization but [6] does not.

In [25], polytopes geometrically encode discretized training problems, both with and without regularization. While [7] characterizes a NN for fixed training data as an exact piecewise linear model defined over a collection of polytopes that is exponential in d and polynomial in n , [25] shows that a single polytope, with a number of faces that is exponential in d and linear in n , encodes approximations of training problems for all data sets in $[-1, 1]^d$. Specifically, training a network on a given data set amounts to solving a linear program over a face of the polytope, up to an additive quantization error [25].

Hyperplane Arrangements and Zonotopes

Hyperplane arrangements (17) can be described geometrically. Specifically, the hyperplane arrangements of \mathbf{X} correspond to vertices of the *zonotope* (Figure 2) of \mathbf{X} , defined as

$$\mathcal{Z} = \text{Conv} \left\{ \sum_{i=1}^n h_i \mathbf{x}_i : h_i \in \{0, 1\} \right\} = \{ \mathbf{X} \mathbf{h} : \mathbf{h} \in [0, 1]^n \}. \quad (22)$$

The correspondence is as follows [31]. Since \mathcal{Z} is a polytope, for every $\mathbf{w} \in \mathbb{R}^d$, there is a vertex \mathbf{z}^* of \mathcal{Z} such that

$$(\mathbf{z}^*)^T \mathbf{w} = S(\mathbf{w}) = \max_{\mathbf{z} \in \mathcal{Z}} \mathbf{z}^T \mathbf{w}. \quad (23)$$

The line $\mathbf{z}^T \mathbf{w} = (\mathbf{z}^*)^T \mathbf{w}$ is a supporting hyperplane of \mathcal{Z} , shown in Figure 3 of *Sampling Arrangement Patterns* below, and $S(\mathbf{w})$ is the support function of \mathcal{Z} . We have shown that vertices maximize the support function of \mathcal{Z} . Conversely, for every vertex \mathbf{z}^* , there is a \mathbf{w} such that that $\mathbf{z}^* = \arg \max_{\mathbf{z} \in \mathcal{Z}} \mathbf{z}^T \mathbf{w}$. Now, (23) is equivalent to

$$\max_{\mathbf{h}_i \in [0, 1]} \sum_{i=1}^n h_i \mathbf{x}_i^T \mathbf{w}$$

whose solution is a hyperplane arrangement pattern $\mathbf{h} = \mathbf{1}\{\mathbf{X}^T \mathbf{w} \geq 0\}$. Therefore each vertex \mathbf{z}^* corresponds to an activation chamber $\{\mathbf{w} : \text{sign}(\mathbf{X}^T \mathbf{w}) = \mathbf{h}\}$.

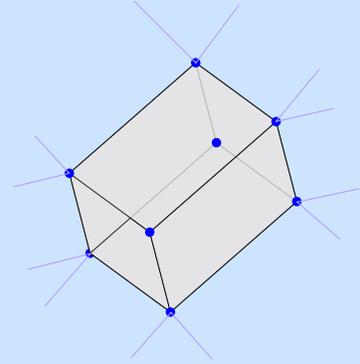


Fig. 2: Zonotope example. Lines indicate normal cones.

The inset “Sampling Arrangement Patterns” discusses hyperplane sampling to reduce computational complexity. Bai et al. [8] presents and analyzes a procedure to sample a collection of hyperplanes that scales linearly with n . As shown in [8], [31], it is possible to control the optimization error resulting from randomly sampling arrangements for the Lasso problem. In particular, sampling $O(P \log(P))$ random arrangements yields a globally optimal solution with high probability, where P is the total number of hyperplane arrangements [31]. In [32], it was proven that the patterns can be randomly subsampled to lower the complexity to $\frac{n}{d} \log n$ with a guaranteed $\sqrt{\log n}$ relative approximation of the objective. This enables fully polynomial-time approximation schemes for the convex NN program.

Sampling Arrangement Patterns

Using a subset of hyperplane arrangements (17) in the Lasso problem (18) can reduce the exponential complexity of enumerating all patterns. This yields a subsampled Lasso problem, and is proven to correspond to stationary points of the non-convex training problem [31]. A practical way to sample hyperplane arrangements is to generate random $\mathbf{w} \sim \mathcal{N}(0, \mathbf{I}_d)$ and evaluate $\mathbf{h}_g = 1\{\mathbf{X}^T \mathbf{w} \geq 0\}$ [7],[8]. Then $\mathbf{D}_g = \text{Diag}(\mathbf{h}_g)$ are used to solve an estimated group Lasso problem (18). An estimated optimal network is reconstructed from the estimated Lasso solution via (19).

The hyperplane sampling implicitly samples vertices of the data's zonotope, whose normal cone solid angles (Figure 3) are proportional to the probability of sampling each vertex [31]. Zonotopes are described in (22) in the inset "Hyperplane Arrangements and Zonotopes". Zonotopes have various signal processing applications [33].

The hyperplane sampling approach is directly connected to 1-bit compressive sensing [9], which randomly samples a signal $\mathbf{x} \in \mathbb{R}^d$ as $\text{sign}(\mathbf{x}^T \mathbf{w}_i)$ where $\mathbf{w}_i = \mathbf{a}_i, \mathbf{x}_i = \mathbf{z}_i$ in (7). Multiple signals stacked into a data matrix \mathbf{X} can be randomly sampled at once as $\text{sign}(\mathbf{X}^T \mathbf{w})$ where $\mathbf{w} \sim \mathcal{N}(0, \mathbf{I}_d)$, which is precisely sampling activation chambers, or equivalently, the vertices of a zonotope.

Sampling patterns is also related to Locality-Sensitive Hashing (LSH), a method for efficiently finding nearest neighbors (in Euclidean distance) of an entry $\mathbf{x} \in \mathbb{R}^d$ in a database [34]. LSH places \mathbf{x} in a database bin according to the sign of $\mathbf{x}^T \mathbf{w}$ where $\mathbf{w} \sim \mathcal{N}(0, \mathbf{I}_d)$, similar to 1-bit CS.

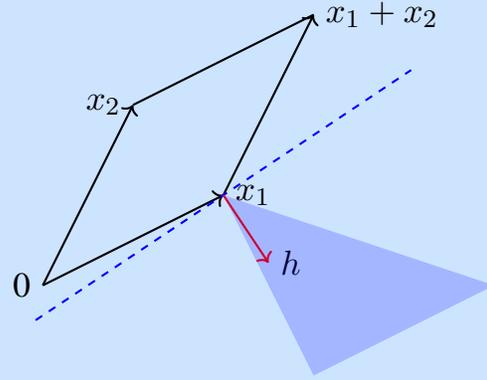


Fig. 3: Zonotope normal cones

The group Lasso formulation in Theorem 1 enables model pruning algorithms to generate minimal models [35]. A minimal model is a minimum-width, globally optimal NN. Starting an optimal convex solution (18), [35] successively removes 'redundant' neurons corresponding to group Lasso variables $\mathbf{u}^{(g)}$ whose contributions $\mathbf{D}_g \mathbf{X}^T \mathbf{u}^{(g)}$ to the prediction are linear combinations of that from other group variables. The algorithm also extends to constructing approximate minimal models that use fewer neurons than required for global optimality [35].

The proof of Theorem 1 involves showing that the bidual of the non-convex training problem is equivalent to the Lasso problem and reconstructing a network from the Lasso problem that achieves the same objective in the training problem as the Lasso problem, thus closing the duality gap and proving equivalence. We note that there are alternative ways to obtain the convex formulation. One approach involves showing that neurons in the same arrangement can be merged and then optimizing over neurons constrained to the

activation chambers. Here, we focus on the convex duality approach, which parallels the duality theory of traditional Lasso methods used in signal processing. The next section describes a neural weight property related to the proof.

Gaining Insight into Network Weights: Optimal Neural Scaling. Yang et al. [36] and Parhi and Nowak [2] describe a “Neural balance Theorem” for training with weight decay regularization. The theorem states that the magnitude of the input and output weights of any neuron with homogeneous activation must be equal, satisfying a “balancing constraint.” Specifically, the minimum-weight representation of a given function by a two-layer NN satisfies $|\alpha_j| = \|\mathbf{w}^{(j)}\|_2$ for all j . The network must have a homogeneous activation such as ReLU. This weight scaling strategy is used as an important first step in the convexification approach of [7]. The ReLU NN (10) $f(\mathbf{x}) = \sum_{j=1}^m (\mathbf{x}^T \mathbf{w}^{(j)})_+ \alpha_j$ is invariant to multiplying $\mathbf{w}^{(j)}$ and dividing α_j by any positive scalar γ_j . On the other hand, for fixed $\mathbf{w}^{(j)}, \alpha_j$, their training regularization $\|\gamma_j \mathbf{w}^{(j)}\|_2^2 + |\frac{1}{\gamma_j} \alpha_j|^2$ is minimized over γ_j when $\|\gamma_j \mathbf{w}^{(j)}\|_2 = |\frac{1}{\gamma_j} \alpha_j|$, i.e., $\gamma_j^* = \frac{|\alpha_j|}{\|\mathbf{w}^{(j)}\|_2}$. Therefore an optimal NN that minimizes the regularized training problem will have equal magnitude inner and outer weights $|\alpha_j| = \|\mathbf{w}^{(j)}\|_2$ for all j . Intuitively, this means that an optimal network balances weights evenly between layers. Similar scaling properties hold for deeper networks and other activations. Next we discuss the derivation of the explicit convex program (18) for ReLU NNs using duality theory.

Convex duality of two-layer ReLU networks. The previous discussion shows that the value of the weight decay regularization with optimal parameters in the training problem (16) is equivalent to $\beta \sum_{j=1}^m \|\mathbf{w}^{(j)}\|_2 |\alpha_j|$ and the network can be written as $f(\mathbf{x}) = \sum_{j=1}^m (\mathbf{x}^T \frac{\mathbf{w}^{(j)}}{\|\mathbf{w}^{(j)}\|_2})_+ \alpha_j$. Substitute $\|\mathbf{w}^{(j)}\|_2 \alpha_j \rightarrow \alpha_j$ and rewrite the training problem (16) as

$$\begin{aligned} \min_{\|\mathbf{w}^{(j)}\|_2=1, \alpha} \quad & \frac{1}{2} \|\mathbf{u} - \mathbf{y}\|_2^2 + \beta \|\alpha\|_1 \\ \text{s.t. } \quad & \mathbf{u} = \sum_{j=1}^m (\mathbf{X}^T \mathbf{w}^{(j)})_+ \alpha_j. \end{aligned} \tag{24}$$

The Lagrangian of (24) has linear and ℓ_1 norm terms of α , so minimizing it over α gives the dual of (24):

$$\begin{aligned} \max_{\mathbf{v}} \quad & -\frac{1}{2} \|\mathbf{v} - \mathbf{y}\|_2^2 \\ \text{s.t. } \quad & \max_{\|\mathbf{w}\|_2 \leq 1} |\mathbf{v}^T (\mathbf{X}^T \mathbf{w})_+| \leq \beta. \end{aligned} \tag{25}$$

While (25) has a semi-infinite constraint over all $\|\mathbf{w}\|_2 \leq 1$, there are in fact a finite number of unique possible vectors $(\mathbf{X}^T \mathbf{w})_+$, corresponding to activation chambers $\mathcal{K}^{(g)}$, as shown in “Hyperplane Arrangement Patterns.” Maximizing over each $\mathcal{K}^{(g)}$ makes the constraint in (24) finite and linear:

$$\begin{aligned} \max_{\mathbf{v}} \quad & -\frac{1}{2} \|\mathbf{v} - \mathbf{y}\|_2^2 \\ \text{s.t. } \quad & \max_{\|\mathbf{w}\|_2 \leq 1, \mathbf{w} \in \mathcal{K}^{(g)}} |\mathbf{v}^T \mathbf{D}_g \mathbf{X}^T \mathbf{w}| \leq \beta \text{ for all } g = 1, \dots, G \end{aligned} \tag{26}$$

Taking the dual of the (26), assigning $\mathbf{u}^{(g)}, \mathbf{v}^{(g)}$ to correspond to positive and negative signs inside the absolute value, and simplifying gives the Lasso problem. The Lasso problem therefore presents a lower bound on the training problem. However, the reconstruction (19) gives a network that achieves the same objective, and thus the Lasso problem and training problem are equivalent. **Hyperplane enumeration gives the key equivalence of (25) and (26), and the key Lasso duality is the equivalence of (26) and (18).** Most importantly, the active constraints in (26) identify exactly which groups are selected at an optimum, and hence which neurons appear in the recovered network. This mirrors the Lasso dual, where constraints that are tight pinpoint the active primal coefficients, equivalently the support of the sparse solution. Similar approaches can extend to deeper architectures, discussed next.

DEEPER RELU NETWORKS AND CONVEXITY

While shallow networks suffice for simpler applications, deeper networks can capture more complex relationships between data. Deep networks can also admit convex programs depending on the architecture.

Importance of Parallel Architecture

We consider two architectures for a deep NN. First, a L -layer *standard* network extends the two-layer network (3) by composing more layers as $f : \mathbb{R}^d \rightarrow \mathbb{R}$,

$$f(\mathbf{x}) = \left(f^{(L-1)} \circ f^{(L-2)} \circ \dots \circ f^{(1)}(\mathbf{x}) \right) \alpha \quad (27)$$

where $f^{(l)}(\mathbf{x}) = \sigma(\mathbf{x}^T \mathbf{W}^{(l)})$ is the l^{th} layer (2) and the trainable parameters are weight matrices $\mathbf{W}^{(l)}$ for $l \in [L-1]$ and final layer parameter α . An alternative architecture is a L -layer *parallel network* which is a sum of K standard networks as $f : \mathbb{R}^d \rightarrow \mathbb{R}$,

$$f(\mathbf{x}) = \sum_{k=1}^K \left(f^{(k,L-1)} \circ f^{(k,L-2)} \circ \dots \circ f^{(k,1)}(\mathbf{x}) \right) \alpha^{(k)} \quad (28)$$

where $f^{(k,l)}(\mathbf{x}) = \sigma(\mathbf{x}^T \mathbf{W}^{(k,l)})$ is the l^{th} layer (2) of the k^{th} unit, which consists of m_l neurons. The trainable parameters are the outermost weights $\alpha^{(k)} \in \mathbb{R}^{m_{L-1}}$ and the inner weights $\mathbf{W}^{(k,l)} \in \mathbb{R}^{m_{l-1} \times m_l}$ for each layer $l \in [L-1]$ and unit $k \in [K]$, where $m_0 = d$ and $m_L = 1$ [37]. The standard (27) and parallel (28) architectures both appear in literature [4, 27, 37, 38], and coincide for two-layer NNs. ‘‘Convex Duality of 2-layer ReLU Networks’’ shows that [7] constructs the convex NN formulation as the bidual of, and hence a lower bound on, the non-convex training problem. This lower bound is tight for two-layer networks: strong duality holds with no duality gap, so the convex formulation is equivalent to the training problem [38]. The duality gap is the difference between the optimal value of the original problem and its dual (which is shown to be is equivalent to its bidual) [7]. Applying the same duality approach to deeper parallel NNs also produces an equivalent convex problem with no duality gap [38]. However, this approach for standard networks, even with a linear activation $\sigma(x) = x$, produces a non-zero duality gap [38], necessitating parallel architectures to extend this convex analysis for more layers.

The linear programs representing discretized NN training in [25] hold for multiple layers, although solving the linear programs remains intractable. In [4], the global optimality results for a two-layer NN analogously extend to deep parallel networks: a locally optimal parallel NN is globally optimal if at least one of the parallel units is zero (all weights in the unit are zero matrices). The training problem is equivalent to a convex program analogous to (11), but as in the two-layer case, the convex program is intractable with an implicitly formulated regularization $\tilde{R}(\mathbf{F})$ and cannot be typically solved in polynomial time [4].

Neural Feature Repopulation

The mean-field convexification approach [21, 22] extends to multiple layers in [39] through a *neural feature repopulation* strategy which considers each layer as a feature. The NNs use a class of regularizers including $\ell_{1,2}$ regularization, where inner and outer weights incur ℓ_2 and ℓ_1 norm penalties respectively. For a network with one neuron in each layer, let $x^{(l)}$ be the vector-valued output of the neuron in layer l . [39]. Let $w^{(2)}$ be a function that takes in $\mathbf{w}^{(1)}$ and the second layer output $x^{(2)}$ and outputs the second layer weight $\mathbf{w}^{(2)}$. Let $p^{(1)}$ be the distribution of $\mathbf{w}^{(1)}$. Every possible output of the second layer $x^{(2)}$ is formed by integrating $x^{(1)} = \mathbf{X}^T \mathbf{w}^{(1)}$ passed through an activation and scaled by the second layer weight [39]:

$$x^{(2)} = \int_{\mathbf{w}^{(1)}} w^{(2)}(\mathbf{w}^{(1)}, x^{(2)}) \sigma(x^{(1)}) dp^{(1)}(\mathbf{w}^{(1)}). \quad (29)$$

For deeper layers, the vector weight $\mathbf{w}^{(l)}$ between layers l and $l+1$ is viewed as the output of a weight function $w^{(l)}$ that takes as input $(x^{(l)}, x^{(l+1)})$ [39]. Also let $p^{(l)}$ be the probability distribution of the output of the l^{th} layer. Every intermediate layer between the first and last hidden layer satisfies

$$x^{(l+1)} = \int_{x^{(l)}} w^{(l)}(x^{(l)}, x^{(l+1)}) \sigma(x^{(l)}) dp^{(l)}(x^{(l)}). \quad (30)$$

Finally, the last layer satisfies

$$x^{(L+1)} = \int_{x^{(L)}} \alpha(x^{(L)}) \sigma(x^{(L)}) dp^{(L)}(x^{(L)}) \quad (31)$$

where α is the outer weight as a function of the output of the last hidden layer. The composition of layers is converted to constraints in a convex optimization problem with an objective that is linear in the feature corresponding to the output layer [39]. The complexity is confined to the regularization. While this approach is effective for a finite, arbitrary number of layers, as with the two-layer case, the convex problem is infinite-dimensional (in the number of optimization variables). On the other hand, the convex equivalences formulated in [7] extends to three-layer ReLU parallel networks, giving explicit convex training programs as detailed next.

Explicit Three-layer Convex Networks with Path Regularization

The convex formulation (18) extends to deeper networks [37]. The training problem for a three-layer parallel ReLU network $f(\mathbf{x}) = \sum_{k=1}^K \left((\mathbf{x}^T \mathbf{W}^{(k,1)})_+ \mathbf{W}^{(k,2)} \right)_+ \boldsymbol{\alpha}^{(k)}$ with *path regularization* is

$$\min_{\theta} \frac{1}{2} \|f(\mathbf{X}) - \mathbf{y}\|_2^2 + \beta \sum_{k=1}^K \sqrt{\sum_{i,j} \|\mathbf{w}_i^{(k,1)}\|_2^2 (w_{i,j}^{(k,2)})^2 (\alpha_j^{(k)})^2} \quad (32)$$

where the parameter set is $\theta = \{\mathbf{W}^{(k,1)}, \mathbf{W}^{(k,2)}, \boldsymbol{\alpha}^{(k)} : k \in [K]\}$, $\mathbf{w}_i^{(k,1)}$ is the i^{th} column of $\mathbf{W}^{(k,1)} \in \mathbb{R}^{d \times m_1}$, $w_{i,j}^{(k,2)}$ is the $(i, j)^{\text{th}}$ element of $\mathbf{W}^{(k,2)} \in \mathbb{R}^{m_1 \times m_2}$ and $\alpha_j^{(k)}$ is the j^{th} element of $\boldsymbol{\alpha}^{(k)} \in \mathbb{R}^{m_2}$. The regularization penalizes all of the paths from the input through the parallel networks to the output. The three-layer ReLU training problem (32) is equivalent to the convex group Lasso problem

$$\min_{\mathbf{u}^{(i,j,k)}, \mathbf{v}^{(i,j,k)} \in \mathcal{K}^{(i,j)}} \frac{1}{2} \left\| \sum_{i=1}^{G_1} \sum_{j,k=1}^{G_2} \mathbf{D}_i \mathbf{D}'_{j,k} \mathbf{X}^T (\mathbf{u}^{(i,j,k)} - \mathbf{v}^{(i,j,k)}) - \mathbf{y} \right\|_2^2 + \beta \sum_{i,j,k} \left(\|\mathbf{u}^{(i,j,k)}\|_2 + \|\mathbf{v}^{(i,j,k)}\|_2 \right) \quad (33)$$

where $\mathbf{D}_i, \mathbf{D}'_{j,k}$ are diagonal matrices encoding hyperplane arrangement patterns in the first and second layer, and $\mathcal{K}^{(i,j)}$ are multilayer activation chambers [37]. Multilayer hyperplanes divide the two-layer activation chambers $\mathcal{K}^{(i)}$ into subchambers $\mathcal{K}^{(i,j)}$, partitioning the data into finer regions where the NN acts as a local linear model. The convex perspective shows that NNs gain representation power with depth by learning activation patterns that capture more complex structures in the data and tailoring linear models more locally within each activation chamber. The convex formulation provides a geometric interpretation of the network operations, revealing the underlying sparsity structure. A network reconstruction formula is given in [37]. For the deep learning community, the convex recasting of training eliminates the need for reducing local minima, determining initialization, and other hyperparameter tuning and heuristics used in training. A limitation of this method is the restriction on network architecture. In deeper models, the number of features increases significantly, making it more difficult to use the convex formulations.

INTERPRETABILITY OF NEURAL NETWORKS VIA CONVEXITY

As described earlier, [25] identifies a geometric connection between training a NN and analyzing the faces of a polyhedron using linear programming. The previous sections also described a geometric relationship between NNs and hyperplane arrangements and zonotopes through the convex equivalence in [7]. Here, we describe another geometric interpretation in terms of volumes of polytopes that arises from a variation of the Lasso problem in [7].

1-D data: features encode distances between data points

Thus far, the input dimension to a NN has been arbitrary. However, the special case of 1-D data offers concrete insights on NN structure. Consider a two-layer network with bias parameters $f(\mathbf{x}) = \sigma(\mathbf{x}^T \mathbf{W} + \mathbf{b}) \boldsymbol{\alpha}$ where \mathbf{b} is a m -dimensional row vector representing trainable but unregularized bias

parameters. The training problem for this two-layer network with 1-D inputs, ReLU activation, and weight decay regularization $\beta (\|\mathbf{W}\|_F^2 + \|\boldsymbol{\alpha}\|_2^2)$ is shown in [40] to be equivalent to the Lasso problem

$$\min_{\mathbf{z}} \frac{1}{2} \|\mathbf{A}^+ \mathbf{z}^+ + \mathbf{A}^- \mathbf{z}^- - \mathbf{y}\|_2^2 + \beta \|\mathbf{z}\|_1. \quad (34)$$

The submatrices of the dictionary in (34) are $\mathbf{A}^+, \mathbf{A}^- \in \mathbb{R}^{n \times n}$ with $A_{i,j}^+ = \sigma(x_i - x_j)$, $A_{i,j}^- = \sigma(x_j - x_i)$. For symmetric activations such as $\sigma(x) = |x|$, we can replace \mathbf{A} by \mathbf{A}^+ and \mathbf{z} by \mathbf{z}^+ . In contrast to the constrained group Lasso (18) which requires enumerating hyperplane arrangement patterns, the convex problem (34) is a standard, unconstrained Lasso problem (5) with a straightforward dictionary to construct. This Lasso formulation extends to higher-dimensional data by identifying dictionary matrices \mathbf{A}^+ and \mathbf{A}^- in (34) with volumes of oriented simplices formed by the data points [18], discussed next.

High dimensional data: geometric algebra features encode volumes spanned by data

A key tool in generalizing Lasso results to higher dimensions is Clifford Algebra, and in particular, its wedge product operation. Clifford Algebra is a mathematical framework that has been recently explored in signal processing [19]. The *wedge product* of vectors \mathbf{a} and \mathbf{b} is denoted as $\mathbf{a} \wedge \mathbf{b}$, which measures the signed area of the parallelogram spanned by \mathbf{a} and \mathbf{b} . The wedge product of two 2-D vectors $\mathbf{a} = (a_1, a_2)^T$ and $\mathbf{b} = (b_1, b_2)^T$ is the determinant of the matrix they form: $\mathbf{a} \wedge \mathbf{b} = \det \begin{pmatrix} a_1 & b_1 \\ a_2 & b_2 \end{pmatrix} = a_1 b_2 - b_1 a_2 = 2 \mathbf{Vol}(\triangle(0, \mathbf{a}, \mathbf{b}))$, where $\triangle(\mathbf{c}, \mathbf{a}, \mathbf{b})$ is the triangle with vertices $\mathbf{c}, \mathbf{a}, \mathbf{b}$ and \mathbf{Vol} denotes the 2-volume (area). The two-layer, univariate-input ReLU dictionary elements (34) $A_{i,j}^+ = (x_i - x_j)_+$ represent the positive length of the interval $[x_i, x_j]$. On the other hand [18], we can also write $(x_i - x_j)_+ = \left(\begin{bmatrix} x_i \\ 1 \end{bmatrix} \wedge \begin{bmatrix} x_j \\ 1 \end{bmatrix} \right)_+$, which equals twice the positive area of the triangle with vertices $(x_i, 1)^T, (x_j, 1)^T$ and the origin. This triangle has unit height and base $x_i - x_j$. Lifting the training points $x_i, x_j \in \mathbb{R}$ to $(x_i, 1)^T, (x_j, 1)^T \in \mathbb{R}^2$ corresponds to appending a vector of ones to the data matrix to represent a bias parameter in the network.

Now consider a two-layer ReLU network $f(\mathbf{x}) = \sigma(\mathbf{x}^T \mathbf{W}) \boldsymbol{\alpha}$ (without a bias parameter \mathbf{b}) trained on $\mathbf{x}_i \in \mathbb{R}^2, y_i \in \mathbb{R}$ and the non-convex training problem (4) with ℓ_2 -norm loss and ℓ_p -norm weight decay regularization [18] $\|\mathbf{W}\|_p^2 + \|\boldsymbol{\alpha}\|_p^2$. Consider a dictionary for the Lasso problem (5) with elements $A_{ij} = \frac{2}{\|\mathbf{x}_j\|_p} \mathbf{Vol}(\triangle(0, \mathbf{x}_i, \mathbf{x}_j))_+$ [18]. If we add a bias term to the neurons, then let $A_{ij} = \frac{2}{\|\mathbf{x}_{j_1} - \mathbf{x}_{j_2}\|_p} \mathbf{Vol}(\triangle(\mathbf{x}_i, \mathbf{x}_{j_1}, \mathbf{x}_{j_2}))_+$ where $j = (j_1, j_2)$ represents a multi-index. Here, $\mathbf{Vol}(\cdot)_+$ distinguishes triangles based on the orientation, i.e., whether their vertices form clockwise or counter-clockwise loops. The Lasso problem (5) is equivalent to the non-convex training problem when the regularization is ℓ_1 -norm ($p = 1$), and produces an ϵ -optimal network when the regularization is ℓ_2 -norm ($p = 2$) [18]. These results generalize to higher dimensions, with wedge products corresponding to higher-dimensional volumes of parallelotopes spanned by vectors. Figure 4 from [18] illustrates a wedge product in \mathbf{R}^3 as the volume of a paralleliped spanned by 3 vectors. The higher-dimensional dictionary elements are $A_{ij} = \frac{(\mathbf{x}_i \wedge \mathbf{x}_{j_1} \wedge \dots \wedge \mathbf{x}_{j_{d-1}})_+}{\|\mathbf{x}_{j_1} \wedge \dots \wedge \mathbf{x}_{j_{d-1}}\|_2}$,

leading to a convex program that captures the geometric relationships in the data [18]. Notably, A_{ij} is a ratio of volumes of parallelotopes, where the positive part of the signed volume encodes the order of the indices.

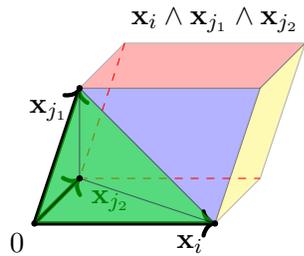


Fig. 4: Wedge product in geometric algebra [18].

The geometric algebra approach provides a unified framework for interpreting deep NNs as a dictionary of wedge product features, revealing the underlying geometric structures inherited from the training data. The volumes of parallelotopes spanned by the training data determine the features. In 1-D, the two-layer NN training problem is equivalent to a classical linear spline Lasso model where the features are hinge functions. The wedge product dictionary shows that in higher dimensions, these hinge functions generalize to signed distances to affine hulls spanned by training data points. The wedge product

formulation also shows that optimal weights are orthogonal to training data [18]. This result establishes a theoretical foundation for the success of deep NNs by linking them to equivalent Lasso formulations whose dictionaries capture geometric and symmetry properties, providing novel insights into the impact of depth on NN expressivity and generalization.

EXTENSIONS TO OTHER NEURAL NETWORK ARCHITECTURES

The convex approaches described in this paper extend to various architectures beyond two-layer ReLU networks. The convex training with respect to outer-layer weights in [20] generalizes to multiple layers, and the parameter distribution approach in [21] extends to homogeneous activations. While [21] qualitatively analyzes training convergence, guarantees on convergence rate and the network width are areas of future work. The implicit convex model with respect to NN output in [4] also extends to multiple layers and a wide family of homogeneous activations. A limitation in [4] is that a potentially large number of neurons in two-layer NNs (or parallel units in deeper networks) are required to guarantee global optimality of training, and a future direction is reducing this size. The linear programs in [25] apply to a variety of architectures including deep networks, convolutional networks, residual networks, vector outputs, and activations. The convex representation approach in [7] extends to modern architectures revealing new forms of convex regularizers. The following contains a brief overview (see [31] and the references therein). The convex Lasso approach extends to polynomial and threshold activation networks, networks with quantized weights, generative adversarial networks and diffusion models used to produce synthetic images and other data, and Convolutional NNs (CNNs), which are popular for image and audio processing problems. Transformers and attention mechanisms are used in Large Language Models for understanding human language, and are convexified by leveraging nuclear norm regularization, a tool used in sparse signal processing. Convex reformulation techniques also reveal that batch normalization, an algorithmic trick essential for the success of local search heuristics, corresponds to applying a whitening matrix to the data, a preprocessing technique used in signal estimation [41]. Next, we discuss numerical experiments.

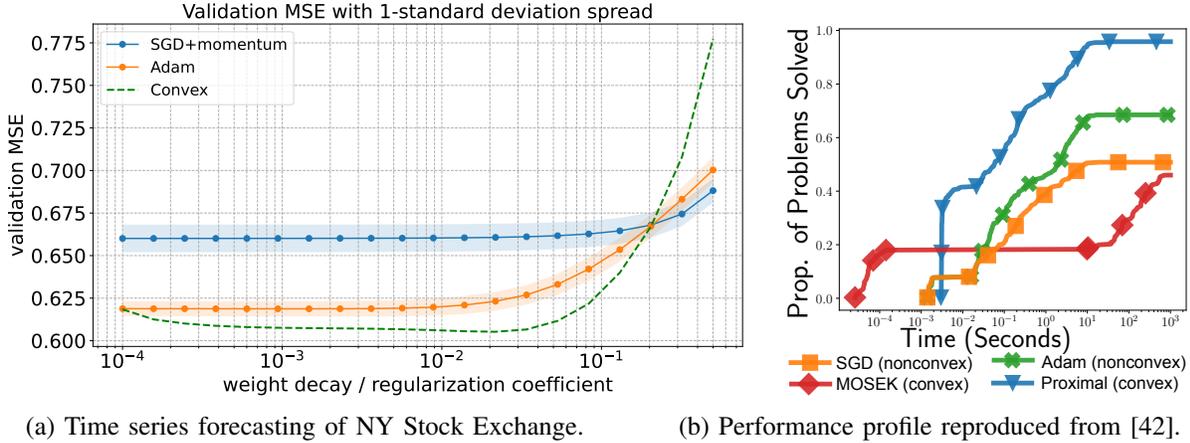


Fig. 5: Experimental results comparing 2-layer NN training with non-convex and convex formulations.

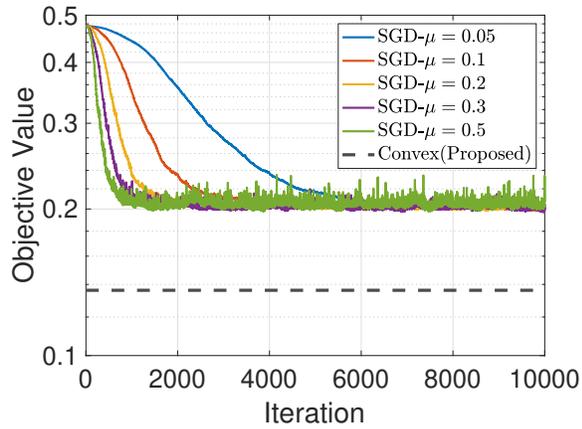
SIMULATIONS

In addition to theoretical guarantees, simulations demonstrate performance advantages of training NNs by using their convex models. In [43], networks trained on convex programs from [7] are shown to perform better or comparably to networks trained on non-convex training problems for standard autoregressive benchmarking tests including MNIST. We focus on two additional experiments using the convex approach in [7].

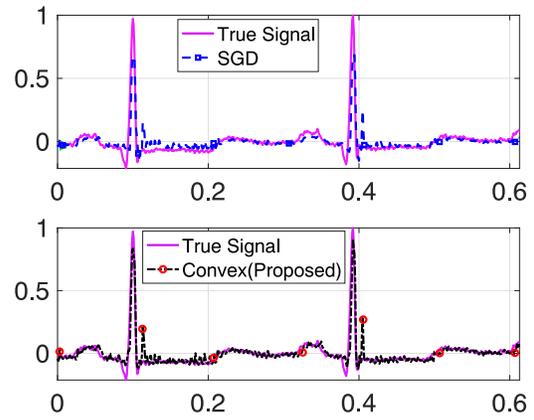
Autoregressive Signal Prediction

Autoregression models a signal $\mathbf{x} \in \mathbb{R}^d$ at time t as a function of its past T samples $\mathbf{x}_t = f_\theta(\mathbf{x}_{t-1}, \dots, \mathbf{x}_{t-T})$ where θ is a parameter set. A linear autoregressive model is of the form $f_\theta(\mathbf{x}_{t-1}, \dots, \mathbf{x}_{t-T}) = a_1\mathbf{x}_{t-1} + \dots + a_T\mathbf{x}_{t-T}$ where the set of coefficients $\theta = \{a_1, \dots, a_T\}$ is estimated from the data. Alternatively, a NN can replace the linear model f_θ to increase expressivity of the predictor. The following two experiments on financial and ECG data investigate the use of convex models to train two-layer NNs for autoregression.

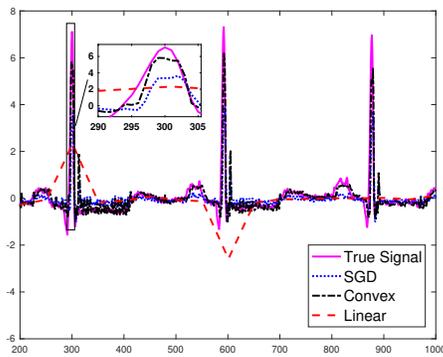
The first experiment performs time series forecasting on the New York stock exchange dataset [44]. Non-linear autoregressive models predict the log volume of exchange from the log volume, Dow Jones return, and log volatility in the past 5 time steps. The models use 4276 training samples and 1770 validation samples. Figure 5a illustrates the results, plotting the mean squared error (MSE) over validation data for different values of β . The NNs trained with convex programs (green dashed line) perform the best overall. Here, a custom proximal algorithm was developed to solve the convex program [42]. The experiment compares different optimizers for non-convex training: stochastic gradient descent (SGD) with momentum and Adam. These non-convex optimizers result in variability between runs, indicated by the shaded area representing one standard deviation across 5 random initializations. In contrast, the convex solver produces consistent and lower error predictions regardless of the initialization. The second experiment, from [31], measures the performance of NN training algorithms in a time series prediction



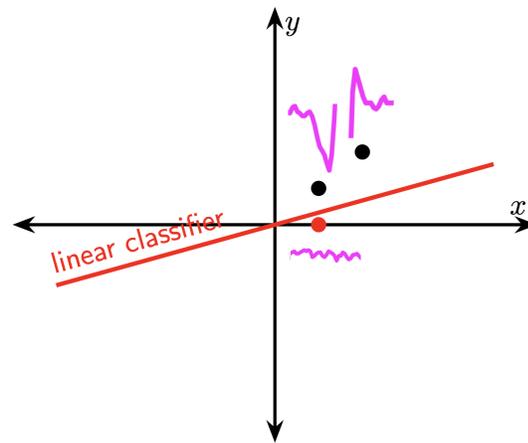
(a) Training objective value



(b) Test predictions



(c) Linear versus network predictions



(d) Lasso shows that NNs are implicitly local linear classifiers.

Fig. 6: Comparison of a two-layer ReLU network trained with SGD and the convex program on ECG data [31]. The convex method achieves lower training loss (6a) and better test performance (6b) compared to SGD with various learning rates μ . The convex program shows that NN are local linear models and can therefore adapt to different regions of data better than standard linear models (6c), (6d).

task using ECG voltage data. Figure 6 illustrates results. Each training sample consists of three consecutive voltage values ($d = 3$), and the model aims to predict the next voltage value. The dataset contains $n = 2393$ observations. A two-layer ReLU network is trained using both SGD and the convex Lasso method. SGD used a batch size of 100 and a grid of learning rates. As shown in Fig. 6, the convex optimization method outperforms SGD in both training loss and test prediction accuracy. The convex method, which yields a globally optimal NN, obtains a lower training objective than SGD, which can get stuck in local minima. The figure shows that SGD demonstrates poorer generalization, demonstrating practical advantages of the convex approach in signal processing applications. The Lasso formulation also reveals why NNs perform better than linear methods. In contrast to a linear classifier which treats all

signals with the same model, whether they are spiking or in a rest phase, the Lasso model acts as a local linear classifier, treating different types of signals with different models, as visualized in Figure 6 (d). The convex models enable direct application of specialized solvers used in signal processing such as proximal gradient methods, which are used in total variation denoising, audio recovery, and other applications [45]. These solvers achieve faster convergence and robustness compared to traditional methods [42]. Figure 5b from [42] compares the proportion of problems solved to 10^{-3} relative training accuracy over 400 UC Irvine datasets. The blue and red curves in Figure 5b plot timing profiles for the convex problem, comparing a custom proximal solver [42] and MOSEK. SGD and ADAM are applied to the non-convex objective, and the experiment selects the best NN over a grid of 3 random initializations and 7 step sizes. The convex version achieves a global optimum and better training performance compared to training with the non-convex problem.

CONCLUSION

The training problem of a two-layer ReLU NN can be equivalently formulated as a convex optimization problem in a variety of ways, including within the frameworks of Lasso and group Lasso from sparse signal processing and CS. Networks can be expressed as convex programs with respect to the outermost weights [20], network outputs [4], or parameter sets [7]. Convexification approaches range from treating the weights as probability distributions and the network as an expectation [21, 22], to quantizing weights and formulating the training problem as a linear program [25], to leveraging convex duality with hyperplane arrangements [7]. Various approaches [20], [7] center on a convex program with a l_1 penalty on the outermost weights, which can be solved or estimated with techniques related to boosting [20] or CS [7] [31]. The convex networks can be interpreted geometrically, as polyhedra representing classes of linear programs [25], or via hyperplane arrangements and zonotope vertices [7], [31], or as high-dimensional volumes using geometric algebra [18]. Convex approaches offer theoretical insights and practical advantages, including stability, interpretability, efficient optimization and improved generalization. The convex programs for regularized training also overcome NP-hard complexities of unregularized training. Future work includes extending these methods to deeper networks with modern activations and exploring their implications in various signal processing and machine learning applications.

REFERENCES

- [1] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, p. 436, 2015.
- [2] R. Parhi and R. D. Nowak, "Deep learning meets sparse regularization: A signal processing perspective," *IEEE Signal Processing Magazine*, vol. 40, no. 6, pp. 63–74, 2023.
- [3] R. Sun, D. Li, S. Liang, T. Ding, and R. Srikant, "The global landscape of neural networks: An overview," *IEEE Signal Processing Magazine*, vol. 37, no. 5, pp. 95–108, 2020.
- [4] B. D. Haeffele and R. Vidal, "Global optimality in neural network training," in *2017 IEEE CVPR*, 2017, pp. 4390–4398.
- [5] L. Luo, "Architectures of neuronal circuits," *Science*, vol. 373, no. 6559, p. eabg7285, 2021.
- [6] V. Froese and C. Hertrich, "Training neural networks is np-hard in fixed dimension," in *NeurIPS*, 2023.
- [7] M. Pilanci and T. Ergen, "Neural networks are convex regularizers: Exact polynomial-time convex optimization formulations for 2-layer networks," in *International Conference on Machine Learning*, 2020.

- [8] Y. Bai, T. Gautam, Y. Gai, and S. Sojoudi, “Practical convex formulations of one-hidden-layer neural network adversarial training,” in *2022 American Control Conference (ACC)*, 2022, pp. 1535–1542.
- [9] P. T. Boufounos and R. G. Baraniuk, “1-bit compressive sensing,” in *2008 42nd Annual Conference on Information Sciences and Systems*, 2008, pp. 16–21.
- [10] E. J. Candès, J. Romberg, and T. Tao, “Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information,” *IEEE Transactions on Information Theory*, vol. 52, no. 2, pp. 489–509, 2006.
- [11] D. L. Donoho, “Compressed sensing,” *IEEE Transactions on information theory*, vol. 52, no. 4, pp. 1289–1306, 2006.
- [12] R. Tibshirani, “Regression shrinkage and selection via the lasso,” *Journal of the Royal Statistical Society Series B: Statistical Methodology*, vol. 58, no. 1, pp. 267–288, 1996.
- [13] M. Yuan and Y. Lin, “Model selection and estimation in regression with grouped variables,” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 68, no. 1, pp. 49–67, 2006.
- [14] E. J. Candès *et al.*, “Compressive sampling,” in *Proceedings of the international congress of mathematicians*, vol. 3. Madrid, Spain, 2006, pp. 1433–1452.
- [15] B. Recht, M. Fazel, and P. A. Parrilo, “Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization,” *SIAM Review*, vol. 52, no. 3, pp. 471–501, 2010. [Online]. Available: <https://doi.org/10.1137/070697835>
- [16] E. J. Candes and T. Tao, “The power of convex relaxation: Near-optimal matrix completion,” *IEEE Transactions on Information Theory*, vol. 56, no. 5, pp. 2053–2080, 2010.
- [17] E. J. Candès, X. Li, Y. Ma, and J. Wright, “Robust principal component analysis?” *J. ACM*, vol. 58, no. 3, Jun. 2011.
- [18] M. Pilanci, “From complexity to clarity: Analytical expressions of deep neural network weights via clifford algebra and convexity,” *Transactions on Machine Learning Research*, 2024.
- [19] N. A. Valous, E. Hitzer, S. Vitabile, S. Bernstein, C. Lavor, D. Abbott, M. E. Luna-Elizarrarás, and W. Lopes, “Hypercomplex signal and image processing: Part 1,” *IEEE Signal Processing Magazine*, vol. 41, no. 2, pp. 11–13, 2024.
- [20] Y. Bengio, N. Roux, P. Vincent, O. Delalleau, and P. Marcotte, “Convex neural networks,” in *Advances in Neural Information Processing Systems*, vol. 18. MIT Press, 2005.
- [21] F. Bach and L. Chizat, “Gradient Descent on Infinitely Wide Neural Networks: Global Convergence and Generalization,” in *International Congress of Mathematicians*, Saint-Petersbourg, Russia, Jul. 2022.
- [22] S. Mei, A. Montanari, and P.-M. Nguyen, “A mean field view of the landscape of two-layer neural networks,” *Proceedings of the National Academy of Sciences*, vol. 115, no. 33, pp. E7665–E7671, 2018.
- [23] A. Jacot, F. Gabriel, and C. Hongler, “Neural tangent kernel: Convergence and generalization in neural networks,” in *Advances in Neural Information Processing Systems*, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, Eds., vol. 31. Curran Associates, Inc., 2018.
- [24] F. Perez-Cruz and O. Bousquet, “Kernel methods and their potential use in signal processing,” *IEEE Signal Processing Magazine*, vol. 21, no. 3, pp. 57–65, 2004.
- [25] D. Bienstock, G. Muñoz, and S. Pokutta, “Principled deep neural network training through linear programming,” *Discrete Optimization*, vol. 49, p. 100795, 2023. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1572528623000373>
- [26] M. J. Wainwright, M. I. Jordan *et al.*, “Graphical models, exponential families, and variational inference,” *Foundations and Trends® in Machine Learning*, vol. 1, no. 1–2, pp. 1–305, 2008.
- [27] F. Bach, “Breaking the curse of dimensionality with convex neural networks,” *JMLR*, vol. 18, no. 1, pp. 629–681, 2017.
- [28] P. Bühlmann, “Consistency for l_2 boosting and matching pursuit with trees and tree-type basis functions,” Seminar für Statistik, Eidgenössische Technische Hochschule (ETH) Zürich, Technical Report 109, 2002. [Online]. Available: <https://doi.org/10.3929/ETHZ-A-004430028>
- [29] K. Prakhya, T. Birdal, and A. Yurtsever, “Convex formulations for training two-layer relu neural networks,” in *ICLR*, 2025.
- [30] R. P. Stanley, “An introduction to hyperplane arrangements,” in *Geometric Combinatorics*, ser. IAS/Park City Mathematics Series, E. Miller, V. Reiner, and B. Sturmfels, Eds. Providence, RI: American Mathematical Society, 2007, vol. 13, pp. 389–496.
- [31] T. Ergen and M. Pilanci, “The convex landscape of neural networks: Characterizing global optima and stationary points via lasso models,” *ArXiv*, vol. abs/2312.12657, 2023. [Online]. Available: <https://api.semanticscholar.org/CorpusID:266374487>

- [32] S. Kim and M. Pilanci, “Convex relaxations of relu neural networks approximate global optima in polynomial time,” *ICML*, 2024.
- [33] K. Stinson, D. F. Gleich, and P. G. Constantine, “A randomized algorithm for enumerating zonotope vertices,” 2016. [Online]. Available: <https://arxiv.org/pdf/1602.06620v1.pdf>
- [34] M. Slaney and M. Casey, “Locality-sensitive hashing for finding nearest neighbors [lecture notes],” *IEEE Signal Processing Magazine*, vol. 25, no. 2, pp. 128–131, 2008.
- [35] A. Mishkin and M. Pilanci, “Optimal sets and solution paths of relu networks,” *ICML*, 2023.
- [36] L. Yang, J. Zhang, J. Shenouda, D. Papailiopoulos, K. Lee, and R. D. Nowak, “A better way to decay: Proximal gradient training algorithms for neural nets,” *OPT2022: 14th Annual Workshop on Optimization for Machine Learning*, vol. abs/2210.03069, 2022. [Online]. Available: <https://api.semanticscholar.org/CorpusID:252735239>
- [37] T. Ergen and M. Pilanci, “Path regularization: A convexity and sparsity inducing regularization for parallel relu networks,” in *NeurIPS*, 2023.
- [38] Y. Wang, T. Ergen, and M. Pilanci, “Parallel deep neural networks have zero duality gap,” *ICLR*, 2023.
- [39] C. Fang, Y. Gu, W. Zhang, and T. Zhang, “Convex formulation of overparameterized deep neural networks,” *IEEE Transactions on Information Theory*, vol. 68, no. 8, pp. 5340–5352, 2022.
- [40] E. Zeger, Y. Wang, A. Mishkin, T. Ergen, E. Candès, and M. Pilanci, “A library of mirrors: Deep neural nets in low dimensions are convex lasso models with reflection features,” *arXiv preprint arXiv:2403.01046*, 2024.
- [41] T. Felhauer, “Digital signal processing for optimum wideband channel estimation in the presence of noise,” *IEE Proceedings F (Radar and Signal Processing)*, vol. 140, pp. 179–186, 1993.
- [42] A. Mishkin, A. Sahiner, and M. Pilanci, “Fast convex optimization for two-layer relu networks: Equivalent model classes and cone decompositions,” *International Conference on Machine Learning*, 2022.
- [43] V. Gupta, B. Bartan, T. Ergen, and M. Pilanci, “Convex neural autoregressive models: Towards tractable, expressive, and theoretically-backed models for sequential forecasting and generation,” in *ICASSP. IEEE*, 2021, pp. 3890–3894.
- [44] G. James, D. Witten, T. Hastie, R. Tibshirani, and J. Taylor, *An Introduction to Statistical Learning with Applications in Python*, ser. Springer Texts in Statistics. Cham: Springer, 2023.
- [45] N. Antonello, L. Stella, P. Patrinos, and T. Waterschoot, “Proximal gradient algorithms: Applications in signal processing,” 03 2018.

ACKNOWLEDGEMENT

Mert Pilanci was supported in part by the National Science Foundation (NSF) CAREER Award under Grant CCF-2236829; in part by the U.S. Army Research Office Early Career Award under Grant W911NF-21-1-0242; and in part by the Office of Naval Research under Grant N00014-24-1-2164.

BIOGRAPHIES

Emi Zeger (*Student Member, IEEE*) received the B.S. degree in mathematics, the B.S. degree in electrical engineering (EE), and the M.S. degree in EE from UCLA, CA, USA, in 2021. She is pursuing the Ph.D. degree in EE at Stanford University, CA. Her research interests are in optimization, neural networks and sustainable computing. She is the recipient of a best paper award at the IEEE International Conference on Communications (ICC) in 2024.

Mert Pilanci (*Member, IEEE*) is an Assistant Professor in the Department of Electrical Engineering at Stanford University. He received his Ph.D. in Electrical Engineering and Computer Sciences from the University of California, Berkeley in 2016. In 2017, he was a postdoctoral fellow working with Emmanuel Candès at Stanford University. His research interests include optimization, machine learning, and signal processing, with a focus on developing theory and efficient algorithms for neural networks. He is a recipient of the NSF CAREER Award and the Army Research Office Early Career Award. He is a Member of the IEEE and a Member of SIAM.