Learning ReLU Networks on Linearly Separable Data: Algorithm, Optimality, and Generalization

Gang Wang¹⁰, Member, IEEE, Georgios B. Giannakis¹⁰, Fellow, IEEE, and Jie Chen¹⁰, Fellow, IEEE

Abstract-Neural networks with rectified linear unit (ReLU) activation functions (a.k.a. ReLU networks) have achieved great empirical success in various domains. Nonetheless, existing results for learning ReLU networks either pose assumptions on the underlying data distribution being, e.g., Gaussian, or require the network size and/or training size to be sufficiently large. In this context, the problem of learning a two-layer ReLU network is approached in a binary classification setting, where the data are linearly separable and a hinge loss criterion is adopted. Leveraging the power of random noise perturbation, this paper presents a novel stochastic gradient descent (SGD) algorithm, which can provably train any single-hidden-layer ReLU network to attain global optimality, despite the presence of infinitely many bad local minima, maxima, and saddle points in general. This result is the first of its kind, requiring no assumptions on the data distribution, training/network size, or initialization. Convergence of the resultant iterative algorithm to a global minimum is analyzed by establishing both an upper bound and a lower bound on the number of non-zero updates to be performed. Moreover, generalization guarantees are developed for ReLU networks trained with the novel SGD leveraging classic compression bounds. These guarantees highlight a key difference (at least in the worst case) between reliably learning a ReLU network as well as a leaky ReLU network in terms of sample complexity. Numerical tests using both synthetic data and real images validate the effectiveness of the algorithm and the practical merits of the theory.

Index Terms—Deep learning, stochastic gradient descent, global optimality, escaping local minima, generalization.

I. INTRODUCTION

D EEP neutral networks have recently boosted the notion of "deep learning from data," with field-changing performance improvements reported in numerous machine learning and artificial intelligence tasks [9], [16]. Despite their widespread use as well as numerous recent contributions, our

Manuscript received August 10, 2018; revised December 17, 2018; accepted February 20, 2019. Date of publication March 20, 2019; date of current version April 1, 2019. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Sotirios Chatzis. The work of G. Wang and G. B. Giannakis was supported partially by the National Science Foundation under Grants 1500713, 1514056, 1505970, and 1711471. The work of J. Chen was in part supported by the National Natural Science Foundation of China under Grants U1509215 and 61621063, and in part by the Program for Changjiang Scholars and Innovative Research Team in University (IRT1208). (*Corresponding author: Georgios B. Giannakis.*)

G. Wang and G. B. Giannakis are with the Digital Technology Center and the Department of Electrical and Computer Engineering, University of Minnesota, Minneapolis, MN 55455 USA (e-mail: gangwang@umn.edu; georgios@umn.edu).

J. Chen is with the State Key Laboratory of Intelligent Control and Decision of Complex Systems, Beijing Institute of Technology, Beijing 100081, China, and also with the Tongji University, Shanghai 200092, China (e-mail: chenjie@bit.edu.cn).

Digital Object Identifier 10.1109/TSP.2019.2904921

understanding of how and why neural networks (NNs) achieve this success remains limited. While their expressivity (expressive power) has been well argued [29], [30], the research focus has shifted toward addressing the computational challenges of training such models and understanding their generalization behavior.

From the vantage point of optimization, training deep NNs requires dealing with extremely high-dimensional and non-convex problems, which are NP-hard in the worst case. It has been shown that even training a two-layer NN of three nodes is NP-complete [3], and the loss function associated with even a single neuron exhibits exponentially many local minima [2]. It is therefore not clear whether and how we can provably yet efficiently train a NN to global optimality.

Nevertheless, as often evidenced by empirical tests, these NN architectures can be 'successfully' trained by means of simple local search heuristics, such as 'plain-vanilla' (stochastic) (S) gradient descent (GD) on real or randomly generated data. Considering the over-parameterized setting in particular, where the NNs have far more parameters than training samples, SGD can often successfully train these networks while exhibiting favorable generalization performance without overfitting [23]. As an example, the celebrated VGG19 net with 20 million parameters trained on the CIFAR-10 dataset of 50 thousand data samples achieves state-of-the-art classification accuracy, and also generalizes well to other datasets [33]. In addition, training NNs by e.g., adding noise to the training samples [35], or to the (stochastic) gradients during back-propagation [22], has welldocumented merits in training with enhancing generalization performance, as well as in avoiding bad local minima [35]. In this contribution, we take a further step toward understanding the analytical performance of NNs, by providing fundamental insights into the optimization landscape and generalization capability of NNs trained by means of SGD with properly injected noise.

For concreteness, we address these challenges in a binary classification setting, where the goal is to train a two-layer ReLU network on *linearly separable data*. Although a nonlinear NN is clearly not necessary for classifying linearly separable data, as a linear classifier such as the Perceptron, would do [31], the fundamental question we target here is whether and how one can efficiently train a ReLU network to global optimality, *despite* the presence of infinitely many local minima, maxima, and saddle points [18]. The motivation behind employing separable data is twofold. They can afford a zero training loss, and distinguish whether a NN is successfully trained or not (as most loss functions for training NNs are non-convex, it is in general

1053-587X © 2019 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

difficult to check its global optimum). In addition, separable data enable improvement of the plain-vanilla SGD by leveraging the power of random noise in a principled manner, so that the modified SGD algorithm can provably escape local minima and saddle points efficiently, and converge to a global minimum in a finite number of non-zero updates. We further investigate the generalization capability of successfully trained ReLU networks leveraging compression bounds [21]. Thus, the binary classification setting offers a favorable testbed for studying the effect of training noise on avoiding overfitting when learning ReLU networks, our novel algorithm and theoretical results can shed light on developing reliable training algorithms for as, well as on, understanding generalization of deep networks.

In a nutshell, the main contributions of the present work are:

- c1) A simple SGD algorithm that can provably escape *local minima* and *saddle points* to efficiently train *any* two-layer ReLU network to attain global optimality;
- c2) Theoretical and empirical evidence supporting the injection of noise during training NNs to escape bad local minima and saddle points; and
- c3) Tight generalization error bounds and guarantees for (possibly over-parameterized) ReLU networks optimally trained with the novel SGD algorithm.

The remainder of this paper is structured as follows. Section II reviews related contributions. Section III introduces the binary classification setting, and the problem formulation. Section IV presents the novel SGD algorithm, and establishes its theoretical performance. Section V deals with the generalization behavior of ReLU networks trained with the novel SGD algorithm. Numerical tests on synthetic data and real images are provided in Section VI. The present paper is concluded with research outlook in Section VII, while technical proofs of the main results are delegated to the Appendix.

Notation: Lower- (upper-)case boldface letters denote vectors (matrices), e.g., a(A). Calligraphic letters are reserved for sets, e.g. S, with the exception of D representing some probability distribution. The operation $\lfloor c \rfloor$ returns the largest integer no greater than the given number c > 0, the cardinality |S| counts the number of elements in set S, and $||\mathbf{x}||_2$ denotes the Euclidean norm of \mathbf{x} .

II. RELATED WORK

As mentioned earlier, NN models have lately enjoyed great empirical success in numerous domains [9], [16], [38]. Many contributions have been devoted to explaining such a success; see e.g., [4], [12]–[14], [18], [19], [27], [37], [40]. Recent research efforts have focused on the expressive ability of deep NNs [30], and on the computational tractability of training such models [4], [34]. In fact, training NNs is NP-hard in general, even for small and shallow networks [2], [10]. Under various assumptions (e.g., Gaussian data, and a sufficiently large number of hidden units) as well as different models however, it has been shown that local search heuristics such as (S)GD can efficiently learn two-layer NNs with quadratic or ReLU activations [34].

Another line of research has studied the landscape properties of various loss functions for learning NNs; see e.g. [4], [13], [18], [24], [25], [28], [39], [40]. Generalizing the results for the ℓ_2 loss [13], [40], it has been proved that deep linear networks with arbitrary convex and differentiable losses have no sub-optimal (a.k.a. bad) local minima, that is all local minima are global, when the hidden layers are at least as wide as the input or the output layer [17]. For nonlinear NNs, most results have focused on learning shallow networks. For example, it has been shown that there are no bad local minima in learning two-layer networks with quadratic activations and the ℓ_2 loss, provided that the number of hidden neurons exceeds twice that of inputs [34], [41]. Focusing on a binary classification setting, [4] demonstrated that despite the non-convexity present in learning one-hidden-layer leaky ReLU networks with a hinge loss criterion, all critical points are global minima if the data are linearly separable. Thus, SGD can efficiently find a global optimum of a leaky ReLU network. On the other hand, it has also been shown that there exist infinitely many bad local optima in learning even two-layer ReLU networks under mild conditions; see e.g., [4, Thm. 8], [18]. Interestingly, [18] provided a complete description of all sub-optimal critical points in learning two-layer ReLU networks with a hinge loss on separable data. Yet, it remains unclear whether and how one can efficiently train even a single-hidden-layer ReLU network to global optimality.

Recent efforts have also been centered on understanding generalization behavior of deep NNs by introducing and/or studying different complexity measures. These include Rademacher complexity, uniform stability, and spectral complexity; see [15] for a recent survey. However, the obtained generalization bounds do not account for the underlying training schemes, namely optimization methods. As such, they do not provide tight guarantees for generalization performance of (over-parameterized) networks trained with iterative algorithms [4]. Even though recent work suggested an improved generalization bound by optimizing the PAC-Bayes bound of an over-parameterized network in a binary classification setting [7], this result is meaningful only when the optimization succeeds. Leveraging standard compression bounds, generalization guarantees have been derived for two-layer leaky ReLU networks trained with plain-vanilla SGD [4]. But this bound does not generalize to ReLU networks, due to the challenge and impossibility of using plain-vanilla SGD to train ReLU networks to global optimum.

III. PROBLEM FORMULATION

Consider a binary classification setting, in which the training set $S := \{(x_i, y_i)\}_{i=1}^n$ comprises n data sampled i.i.d. from some unknown distribution \mathcal{D} over $\mathcal{X} \times \mathcal{Y}$, where without loss of generality we assume $\mathcal{X} := \{x \in \mathbb{R}^d : ||x||_2 \le 1\}$ and $\mathcal{Y} := \{-1, 1\}$. We are interested in the linearly separable case, in which there exists an optimal linear classifier vector $\omega^* \in \mathbb{R}^d$ such that $\mathbb{P}_{(x,y)\sim\mathcal{D}}(y \, \omega^{*\top} x \ge 1) = 1$. To allow for affine classifiers, a "bias term" can be appended to the classifier vector by augmenting all data vectors $x \in \mathcal{X}$ with an extra component of 1 accordingly.

We deal with single-hidden-layer NNs having d scalar inputs, k > 0 hidden neurons, and a single output (for binary classification). The overall input-output relationship of such a two-layer NN is

$$\boldsymbol{x} \mapsto f(\boldsymbol{x}) := \sum_{j=1}^{k} v_j \sigma\left(\boldsymbol{w}_j^{\top} \boldsymbol{x}\right)$$
 (1)

which maps each input vector $\boldsymbol{x} \in \mathbb{R}^d$ to a scalar output by combining k nonlinear maps of linearly projected renditions of \boldsymbol{x} , effected via the ReLU activation $\sigma(z) := \max\{0, z\}$. Clearly, due to the non negativity of ReLU outputs, one requires at least $k \ge 2$ hidden units so that the output $f(\cdot)$ can take both positive and negative values to signify the 'positive' and 'negative' classes. Here, $\boldsymbol{w}_j \in \mathbb{R}^d$ stacks up the weights of the links connecting the input \boldsymbol{x} to the *j*-th hidden neuron, and v_j is the weight of the link from the *j*-th hidden neuron to the output. Upon defining $\boldsymbol{W} := [\boldsymbol{w}_1 \cdots \boldsymbol{w}_k]^\top$ and $\boldsymbol{v} := [v_1 \cdots v_k]^\top$, which are henceforth collectively denoted as $\mathcal{W} := \{\boldsymbol{v}, \boldsymbol{W}\}$ for brevity, one can express $f(\boldsymbol{x})$ in a compact matrix-vector representation as

$$f(\boldsymbol{x}; \mathcal{W}) = \boldsymbol{v}^{\top} \sigma(\boldsymbol{W} \boldsymbol{x})$$
(2)

where the ReLU activation $\sigma(z)$ should be understood entrywise when applied to a vector z.

Given our NN described by f(x; W) and adopting a hinge loss criterion $\ell(z) := \max\{0, 1 - z\}$, we define the empirical loss as the average loss of f(x; W) over the training set S, that is

$$L_{\mathcal{S}}(\mathcal{W}) := \frac{1}{n} \sum_{i=1}^{n} \ell(y_i f(\boldsymbol{x}_i; \mathcal{W}))$$
$$= \frac{1}{n} \sum_{i=1}^{n} \max\left\{0, 1 - y_i \boldsymbol{v}^{\top} \sigma(\boldsymbol{W} \boldsymbol{x}_i)\right\}$$

With the output $f(x; W) \in \mathbb{R}$, we construct a binary classifier $g_f : \mathbb{R}^d \to \mathcal{Y}$ as $g_f = \operatorname{sgn}(f)$, where the sign function $\operatorname{sgn}(z) = 1$ if $z \ge 0$, and $\operatorname{sgn}(z) = 0$ otherwise. For this classifier, the training error (a.k.a. misclassification rate) $\hat{R}_{\mathcal{S}}(W)$ over \mathcal{S} is

$$\hat{R}_{\mathcal{S}}(\mathcal{W}) = \frac{1}{n} \sum_{i=1}^{n} \mathbb{1}_{\{y_i \neq \text{sgn}(f(\boldsymbol{x}_i; \mathcal{W}))\}}$$
(3)

where $\mathbb{1}_{\{\cdot\}}$ denotes the indicator function taking value 1 if the argument is true, and 0 otherwise.

In this paper, we first fix the second layer of f(x; W) to be some constant vector v given *a priori*, with at least one positive and one negative entry. We will subsequently generalize the results to the case where v and W are updated simultaneously. Thus, training the ReLU network f(x; W) boils down to learning the weight matrix W only. As such, the network is henceforth denoted by f(x; W) := f(x; W), and the goal is to solve the following optimization problem

$$\boldsymbol{W}^* := \arg\min_{\boldsymbol{W} \in \mathbb{R}^{k \times d}} L_{\mathcal{S}}(\boldsymbol{W}) \tag{4}$$

where $L_{\mathcal{S}}(\mathbf{W}) = (1/n) \sum_{i=1}^{n} \ell(y_i f(\mathbf{x}_i; \mathbf{W}))$. Evidently for separable data and the ReLU network considered in this paper, it must hold that $L_{\mathcal{S}}(\mathbf{W}^*) = 0$. Due to piecewise linear (non-smooth) ReLU activations, $L_{\mathcal{S}}(\mathbf{W})$ becomes non-smooth. It can be further shown that $L_{\mathcal{S}}(\mathbf{W})$ is non-convex (e.g., [4, Proposition 5.1]), which indeed admits infinitely many (sub-optimal) local minima [4, Thm. 8].

Interestingly though, it is possible to provide an analytical characterization of all sub-optimal solutions. Specifically, at *any* critical point¹ \mathbf{W}^{\dagger} of $L_{\mathcal{S}}(\mathbf{W})$ that incurs a non-zero loss $\ell(y_i f(\mathbf{x}_i; \mathbf{W}^{\dagger})) > 0$ for a datum $(\mathbf{x}_i, y_i) \in \mathcal{S}$, it holds that $\sigma(\mathbf{W}^{\dagger} \mathbf{x}_i) = \mathbf{0}$ [18, Thm. 6], or entry-wise

$$\sigma\left(\boldsymbol{w}_{j}^{\dagger^{\top}}\boldsymbol{x}_{i}\right)=0, \forall j=1, 2, \dots, k.$$
(7)

Expressed differently, if data pair (\boldsymbol{x}_i, y_i) yields a non-zero loss at a critical point \boldsymbol{W}^{\dagger} , the ReLU output $\sigma(\boldsymbol{w}_j^{\dagger} \boldsymbol{x}_i)$ must vanish at *all* hidden neurons. Building on this observation, we say that a critical point \boldsymbol{W}^{\dagger} of $L_{\mathcal{S}}(\boldsymbol{W})$ is *sub-optimal* if it obeys simultaneously the following two conditions: i) $\ell(y_i f(\boldsymbol{x}_i; \boldsymbol{W}^{\dagger})) > 0$ for some data sample $(\boldsymbol{x}_i, y_i) \in S$, and ii) for which it holds that $\sigma(\boldsymbol{W}^{\dagger}\boldsymbol{x}_i) = \mathbf{0}$. According to these two defining conditions, the set of all sub-optimal critical points includes different local minima, as well as all maxima and saddle points; see Fig. 1 for an illustration. It is also clear from Fig. 1 that the two conditions in certain cases *cannot* be changed by small perturbations on \boldsymbol{W}^{\dagger} , suggesting that there are in general *infinitely* many suboptimal critical points. Therefore, optimally training even such a single-hidden-layer ReLU network is indeed challenging.

Consider minimizing $L_{\mathcal{S}}(W)$ by means of plain-vanilla SGD with constant learning rate $\eta > 0$, as

$$\boldsymbol{W}^{t+1} = \boldsymbol{W}^t - \eta \left. \frac{\partial \ell(y_{i_t} f(\boldsymbol{x}_{i_t}; \boldsymbol{W}))}{\partial \boldsymbol{W}} \right|_{\boldsymbol{W} = \boldsymbol{W}^t}$$
(8)

with the (sub-)gradient of the hinge loss at a randomly sampled datum $(x_{i_t}, y_{i_t}) \in S$ given by

$$\frac{\partial \ell(y_{i_t} f(\boldsymbol{x}_{i_t}; \boldsymbol{W}))}{\partial \boldsymbol{W}} = -\mathbb{1}_{\left\{1 - y_{i_t} \boldsymbol{v}^{\top} \sigma(\boldsymbol{W} \boldsymbol{x}_{i_t}) > 0\right\}} y_{i_t} \operatorname{diag} \left(\mathbb{1}_{\left\{\boldsymbol{W} \boldsymbol{x}_{i_t} \ge \boldsymbol{0}\right\}}\right) \boldsymbol{v} \boldsymbol{x}_{i_t}^{\top} \quad (9)$$

where diag(z) is a diagonal matrix holding entries of vector z on its diagonal, and the indicator function $\mathbb{1}_{\{z \ge 0\}}$ applied to z is understood entry-wise. For any sub-optimal critical point W^{\dagger} incurring a nonzero loss for some (x_{i_t}, y_{i_t}) , it can be readily deduced that diag $(\mathbb{1}_{\{W^{\dagger}x_{i_t} \ge 0\}}) = 0$ [18].

Following the convention [22], we say that a ReLU is *active* if its output is non-zero, and *inactive* otherwise. Furthermore, we denote the state of per *j*-th ReLU by its *activity indicator* function $\mathbb{1}_{\{w_j^\top x_{i_t} \ge 0\}}$. In words, there exists always some data sample(s) for which all hidden neurons become inactive at a sub-optimal critical point. This is corroborated by the fact that under some conditions, plain-vanilla SGD converges to a sub-optimal local minimum with high probability [4]. It will also be

¹The critical point for a general non-convex and non-smooth function is defined invoking the Clarke sub-differential (see e.g., [5], [36]). Precisely, consider a function $h(\boldsymbol{z}) : \boldsymbol{Z} \mapsto \mathbb{R}$, which is locally Lipschitz around $\boldsymbol{z} \in \boldsymbol{Z}$, and differentiable on $\boldsymbol{Z} \backslash \boldsymbol{G}$, with \boldsymbol{G} being a set of Lebesgue measure zero. Then the convex hull of the set of limits of the form $\lim \nabla h(\boldsymbol{z}_k)$, where $\boldsymbol{z}_k \to \boldsymbol{z}$ as $k \to +\infty$, i.e.,

$$\partial_0 h(\boldsymbol{z}) := ext{c.h.} \left\{ \lim_k \nabla h(\boldsymbol{z}_k) : \boldsymbol{z}_k o \boldsymbol{z}, \ \boldsymbol{z}_k \notin \mathcal{G} \right\}$$
 (5)

is the so-termed *Clarke sub-differential* of h at z. Furthermore, if **0** belongs to $\partial_0 h(z)$, namely

$$\mathbf{0} \in \partial_0 h(\boldsymbol{z}) \tag{6}$$

then we say that z is a *critical point* of h in the Clarke sense.



Fig. 1. Four different critical points of the loss function $\hat{L}(\mathbf{W})$ (a non-zero classification error $\hat{R}(\mathbf{W})$ yields a non-zero loss) for a ReLU network with two hidden neurons (corresponding to the hyperplanes $\mathbf{w}_1^\top \mathbf{x} = 0$) (left blue line in each plot, with $v_1 = 1$) and $\mathbf{w}_2^\top \mathbf{x} = 0$ (right blue line, with $v_2 = -1$). The arrows point to the positive cone of each hyperplane, namely $\mathbf{w}_j^\top \mathbf{x} \ge 0$ for j = 1, 2. The training dataset contains 3 samples, two of which belong to class '+1' (colored red) and one of which belong to class '-1' (colored black). Data points with a non-zero classification error (hence non-zero loss) must lie in the negative cone of all hyperplanes.

verified by our numerical tests in Section VI, that SGD can indeed get stuck in sub-optimal local minima when training ReLU networks.

IV. MAIN RESULTS

In this section, we present our main results that include a modified SGD algorithm and theory for efficiently training singlehidden-layer ReLU networks to global optimality. As in the convergence analysis of the Perceptron algorithm (see e.g., [26], [32, Chapter 9]), we define an update at iteration t as *non-zero* or *effective* if the corresponding (modified) stochastic gradient is non-zero, or equivalently, whenever one has $W^{t+1} \neq W^t$.

A. Algorithm

As explained in Section III, plain-vanilla SGD iterations for minimizing $L_{\mathcal{S}}(W)$ can get stuck in sub-optimal critical points. Recall from (9) that whenever this happens, it must hold that $\operatorname{diag}(\mathbb{1}_{\{\boldsymbol{W}\boldsymbol{x}_{i_t}\geq\boldsymbol{0}\}})=\boldsymbol{0}$ for some data sample $(\boldsymbol{x}_{i_t},y_{i_t})\in\mathcal{S}$, or equivalently $\boldsymbol{w}_{i}^{\top}\boldsymbol{x}_{i_{t}} < 0$ for all $j = 1, 2, \ldots, k$. To avoid being trapped in these points, we will endow the algorithm with a non-zero '(sub-)gradient' even at a sub-optimal critical point, so that the algorithm will be able to continue updating, and will have a chance to escape from sub-optimal critical points. If successful, then when the algorithm *converges*, it must hold that $\mathbb{1}_{\{1-y_i \boldsymbol{v}^{\top} \sigma(\boldsymbol{W}^{\dagger} \boldsymbol{x}_i) > 0\}} = 0$ for all data samples $(\boldsymbol{x}_i, y_i) \in \mathcal{S}$ (cf. (9)), or $1 - y_i \boldsymbol{v}^{\top} \sigma(\boldsymbol{W}^{\dagger} \boldsymbol{x}_i) \leq 0$ for all i = 1, 2, ..., n, thanks to linear separability of the data. This in agreement with the definition of the hinge loss function satisfies that $L_{\mathcal{S}}(\mathbf{W}^{\dagger}) = 0$ in (4), which guarantees that the algorithm converges to a global optimum. Two critical questions arise at this point: Q1) How can we endow a non-zero '(sub-)gradient' based search direction even at a sub-optimal critical point, while having the global minima as limiting points of the algorithm? and Q2) How is it possible to guarantee convergence?

Question Q1) can be answered by ensuring that at least one ReLU is active at a non-optimal point. Toward this objective,

motivated by recent efforts in escaping saddle points [1], [8], [11], we are prompted to add a zero-mean random noise vector $\epsilon^t \in \mathbb{R}^k$ to $W^t x_{i_t} \in \mathbb{R}^k$, namely the input vector to the activity indicator function of *all* ReLUs. This would replace $\mathbb{1}_{\{W^t x_{i_t} \geq 0\}}$ in the subgradient (cf. (9)) with $\mathbb{1}_{\{W^t x_{i_t} + \epsilon^t \geq 0\}}$ at every iteration. In practice, Gaussian additive noise $\epsilon^t \sim \mathcal{N}(\mathbf{0}, \gamma^2 \mathbf{I})$ with sufficiently large variance $\gamma^2 > 0$ works well.

Albeit empirically effective in training ReLU networks, SGD with such architecture-agnostic injected noise into all ReLU activity indicator functions cannot guarantee convergence in general, or convergence is difficult or even impossible to establish. We shall take a different route to bypass this hurdle here, which will lead to a simple algorithm provably convergent to a wanted global optimum in a finite number of non-zero updates. This result holds regardless of the data distribution, initialization, network size, or the number of hidden neurons. Toward, to ensure convergence of our modified SGD algorithm, we carefully design the noise injection process by maintaining at least one non-zero ReLU activity indicator variable at every non-optimal critical point.

For the picked data sample $(\boldsymbol{x}_{i_t}, y_{i_t}) \in S$ per iteration $t \geq 0$, we inject Gaussian noise $\epsilon_j^t \sim \mathcal{N}(0, \gamma^2)$ into the *j*-th ReLU activity indicator function $\mathbb{1}_{\{\boldsymbol{w}_j^\top \boldsymbol{x}_{i_t} \geq 0\}}$ in the SGD update of (9), if and only if the corresponding quantity $y_{i_t}v_j \geq 0$ holds, and we repeat this for all neurons $j = 1, 2, \ldots, k$.

Interestingly, the noise variance γ^2 , admits simple choices, so long as it is selected sufficiently large matching the size of the corresponding summands $\{|w_j^t {}^T x|^2\}_{j,t}$. We will build up more intuition and highlight the basic principle behind such a noise injection design shortly in Section IV-B, along with our formal convergence analysis. For implementation purposes, we summarize the novel SGD algorithm with randomly perturbed ReLU activity indicator functions in Algorithm 1. As far as stopping criterion is concerned, it is safe to conclude that the algorithm has converged, if there has been no non-zero update for *a succession of* say, *np* iterations, where p > 0 is some fixed large enough integer. This holds with high probability, which depends on *p*, **Algorithm 1:** Learning Two-Layer ReLU Networks via SGD with Randomly Perturbed ReLU Activity Indicators.

- Input: Training data S = {(x_i, y_i)}ⁿ_{i=1}, second layer weight vector v ∈ ℝ^k with at least one positive and at least one negative entry, initialization parameter ρ ≥ 0, learning rate η > 0, and noise variance γ² > 0.
- 2: **Initialize** W^0 with **0**, or randomly having its rows obey $||w_j||_2 \le \rho$.
- 3: For $t = 0, 1, 2, \dots$ do
- 4: **Pick** i_t uniformly at random from, or
- deterministically cycle through $\{1, 2, \ldots, n\}$.
- 5: Update

$$\boldsymbol{W}^{t+1} = \boldsymbol{W}^{t} + \eta \, \mathbb{1}_{\left\{1 - y_{i_t} \boldsymbol{v}^{\top} \sigma(\boldsymbol{W}^t \boldsymbol{x}_{i_t}) > 0\right\}} \\ \times y_{i_t} \operatorname{diag} \left(\mathbb{1}_{\left\{\boldsymbol{W}^t \boldsymbol{x}_{i_t} + \boldsymbol{\epsilon}^t \ge \mathbf{0}\right\}}\right) \boldsymbol{v} \boldsymbol{x}_{i_t}^{\top} \quad (10)$$

where per *j*-th entry of noise $\epsilon^t \in \mathbb{R}^k$ follows $\epsilon_j^t \sim \mathcal{N}(0, \gamma^2)$, if $y_{i_t} v_j \geq 0$; and $\epsilon_j^t = 0$, otherwise.

6: Output: W^{t+1}

and $|\mathcal{N}_v^+|$ ($|\mathcal{N}_v^-|$), where the latter denotes the number of neurons with $v_j > 0$ ($v_j < 0$). We have the following result, whose proof is provided in Appendix D.

Proposition 1: Let $\|\boldsymbol{w}_{j}^{t}\|_{2} \leq w_{\max}$ for all neurons $j = 1, 2, \ldots, k$, and all iterations $t \geq 0$, and consider i_{t} cycling deterministically through $\{1, 2, \ldots, n\}$. If there is no non-zero update after a succession of np iterations, then Algorithm 1 converges to a global optimum of $L_{\mathcal{S}}(\boldsymbol{W})$ with probability at least $1 - [\Phi(w_{\max}/\gamma)]^{p\min\{|\mathcal{N}_{v}^{+}|, |\mathcal{N}_{v}^{-}|\}}$, where $\Phi(z) := (1/\sqrt{2\pi}) \int_{-\infty}^{z} e^{-s^{2}} ds$ is the cumulative density function of the standardized Gaussian distribution $\mathcal{N}(0, 1)$.

Observe that the probability in Proposition 1 can be made arbitrarily close to 1 by taking sufficiently large p and/or γ . Regarding our proposed approach in Algorithm 1, three remarks are worth making.

Remark 1: With the carefully designed noise injection rule, our algorithm constitutes a non-trivial generalization of the Perceptron or plain-vanilla SGD algorithms to learn ReLU networks. Implementing Algorithm 1 is as easy as plain-vanilla SGD, requiring almost negligible extra computation overhead. Both numerically and analytically, we will demonstrate the power of our principled noise injection into partial ReLU activity indicator functions, as well as establish the *optimality*, *efficiency*, and *generalization* performance of Algorithm 1 in learning two-layer (over-parameterized) ReLU networks on linearly separable data.

Remark 2: It is worth remaking that the random (Gaussian) noise in our proposal is *solely added to* the ReLU activity indicator functions, *rather than* to any of the hidden neurons. This is evident from the first indicator function $\mathbb{1}_{\{1-y_{i_t}v^{\top}\sigma(W^tx_{i_t})>0\}}$ being the (sub)derivative of a hinge loss, in Step 5 of Algorithm 1, which is kept as it is in the plain-vanilla SGD, namely it is *not affected* by the noise. Moreover, our use of random noise in this way distinguishes itself from those in the vast literature for evading saddle points (see e.g., [1], [8], [11]),

which simply add noise to either the iterates or to the (stochastic) (sub)gradients. This distinction endows our approach with the unique capability of also escaping local minima (in addition to saddle points). To the best of our knowledge, our approach is the *first* of its kind in provably yet efficiently escaping local minima under suitable conditions.

Remark 3: Compared with previous efforts in learning ReLU networks (e.g., [34]), our proposed Algorithm 1 provably converges to a global optimum in a finite number of non-zero updates, without any assumptions on the data distribution, training/network size, or initialization. This holds even in the presence of exponentially many local minima and saddle points. To the best of our knowledge, Algorithm 1 provides the *first* solution to efficiently train such a single-hidden-layer ReLU network to global optimality with a hinge loss, so long as the training samples are linearly separable. Generalizations to other objective functions based on e.g., the τ -hinge loss and the smoothed hinge loss (a.k.a. polynomial hinge loss) [20], as well as to multilayer ReLU networks are possible, and they are left for future research.

B. Convergence Analysis

In this section, we analyze the convergence of Algorithm 1 for learning single-hidden-layer ReLU networks with a hinge loss criterion on linearly separable data, namely for minimizing $L_S(W)$ in (4). Recall since we only train the first layer having the second layer weight vector $v \in \mathbb{R}^k$ fixed *a priori*, we can assume without further loss of generality that entries of v are all non-zero. Otherwise, one can exclude the corresponding hidden neurons from the network, yielding an equivalent reduced-size NN whose second layer weight vector has all its entries non-zero.

Before presenting our main convergence results for Algorithm 1, we introduce some notation. To start, let $\mathcal{N}_y^+ \subseteq \{1, 2, \ldots, n\}$ (\mathcal{N}_y^-) be the index set of data samples $\{(\boldsymbol{x}_i, y_i)\}_{1 \leq i \leq n}$ belonging to the 'positive' ('negative') class, namely whose $y_i = +1$ ($y_i = -1$). It is thus self-evident that $\mathcal{N}_y^+ \cup \mathcal{N}_y^- = \{1, 2, \ldots, n\}$ and $\mathcal{N}_v^+ \cup \mathcal{N}_v^- = \{1, 2, \ldots, k\}$ hold under our assumptions. Putting our work in context, it is useful to first formally summarize the landscape properties of the objective function $L_{\mathcal{S}}(\boldsymbol{W}) = (1/n) \sum_{i=1}^n \max\{0, 1 - y_i f(\boldsymbol{x}_i; \boldsymbol{W})\}$, which can help identify the challenges in learning ReLU networks.

Proposition 2: Function $L_{\mathcal{S}}(W)$ has the following properties: i) it is non-convex, and ii) for each sub-optimal local minimum (that incurs a non-zero loss), there exists (at least) a datum $(x_i, y_i) \in S$ for which all ReLUs become inactive.

The proof of Property i) in Proposition 2 can be easily adapted from that of [4, Proposition 5.1], while Property ii) is just a special case of [18, Thm. 5] for a fixed v; hence they are both omitted in this paper.

We will provide an upper bound on the number of non-zero updates that Algorithm 1 performs until no non-zero update occurs after within a succession of say, e.g. np iterations (cf. (10)), where p is a large enough integer. This, together with the fact that all sub-optimal critical points of $L_{\mathcal{S}}(W)$ are not limiting points of Algorithm 1 due to the Gaussian noise injection with a large enough variance $\gamma^2 > 0$ at every iteration, will guarantee convergence of Algorithm 1 to a global optimum of $L_{\mathcal{S}}(W)$. Specifically, the main result is summarized in the following theorem.

Theorem 1 (Optimality): If all rows of the initialization W^0 satisfy $||w_j^0||_2 \le \rho$ for any constant $\rho \ge 0$, and the second layer weight vector $v \in \mathbb{R}^k$ is kept fixed with both positive and negative (but non-zero) entries, then Algorithm 1 with some constant step size $\eta > 0$ converges to a global minimum of $L_S(W)$ after performing at most T_k non-zero updates, where for $v_{\min} = \min_{1 \le j \le k} |v_j|$ it holds that

$$T_{k} := \frac{k}{\eta v_{\min}^{2}} \left[\left(\eta \| \boldsymbol{v} \|_{2}^{2} + 2 \right) \| \boldsymbol{\omega}^{*} \|_{2}^{2} + 2\rho v_{\min} \| \boldsymbol{v} \|_{2}^{2} + \sqrt{2\rho v_{\min} \| \boldsymbol{\omega}^{*} \|_{2} \left(\eta \| \boldsymbol{v} \|_{2}^{2} + 2 \right)} \| \boldsymbol{\omega}^{*} \|_{2} \right].$$
(11)

In particular, if $W^0 = 0$, then Algorithm 1 converges to a global optimum after at most $T_k^0 := \frac{k}{\eta v_{\min}^2} (\eta \| \boldsymbol{v} \|_2^2 + 2) \| \boldsymbol{\omega}^* \|_2^2$ non-zero updates.

Regarding Theorem 1, a couple of observations are of interest. The developed Algorithm 1 converges to a globally optimal solution of the non-convex optimization (4) within a finite number of non-zero updates, which implicitly corroborates the ability of Algorithm 1 to escape sub-optimal local minima, as well as saddle points. This holds regardless of the underlying data distribution \mathcal{D} , the number n of training samples, the number k of hidden neurons, or even the initialization \mathbf{W}^0 . It is also worth highlighting that the number T_k of non-zero updates does not depend on the dimension d of input vectors, but it scales with k (in the worst case), and it is inversely proportional to the step size $\eta > 0$. Recall that the worst-case bound for SGD learning of leaky-ReLU networks with initialization $\mathbf{W}^0 = \mathbf{0}$ is [4, Thm. 2]

$$T_{\text{leaky}}^{\alpha} \leq \frac{\|\boldsymbol{\omega}^*\|_2^2}{\alpha^2} \left(1 + \frac{1}{\eta \|\boldsymbol{v}\|_2^2}\right) \tag{12}$$

where again, ω^* denotes an optimal linear classifier obeying $\mathbb{P}_{(\boldsymbol{x},y)\sim\mathcal{D}}(y \boldsymbol{\omega}^{*\top} \boldsymbol{x} \geq 1) = 1$. Clearly, the upper bound above does not depend on k. This is due to the fact that the loss function corresponding to learning leaky-ReLU networks has no bad local minima, since all critical points are global minima. This is in sharp contrast with the loss function associated with learning ReLU networks investigated here, which generally involves infinitely many bad local minima! On the other hand, the bound in (12) scales inversely proportional with the quadratic 'leaky factor' α^2 of leaky ReLUs. This motivates having $\alpha \to 0$, which corresponds to letting the leaky ReLU approach the ReLU. In such a case, (12) would yield a worst-case bound of infinity for learning ReLU networks, corroborating the challenge and impossibility of learning ReLU networks by 'plain-vanilla' SGD. Indeed, the gap between T_k^0 in Theorem 1 and the bound in (12) is the price for being able to escape local minima and saddle points paid by our noise-injected SGD Algorithm 1. Last but not least, Theorem 1 also suggests that for a given network and a fixed step size η , Algorithm 1 with $W^0 = 0$ works well too.

We briefly present the main ideas behind the proof of Theorem 1 next, but delegate the technical details to Appendix A. Our proof mainly builds upon the convergence proof of the classical Perceptron algorithm (see e.g., [32, Thm. 9.1]), and it is also inspired by that of [4, Thm. 1]. Nonetheless, the novel approach of performing SGD with principled noise injection into the ReLU activity indicator functions distinguishes itself from previous efforts. Since we are mainly interested in the (maximum) number of non-zero updates to be performed until convergence, we will assume for notational convenience that all iterations $t \ge 0$ in (10) of Algorithm 1 perform a non-zero update. This assumption is made without loss of generality. To see this, since after the algorithm converges, one can always re-count the number of effective iterations that correspond to a non-zero update and re-number them by t = 0, 1, ...

Our main idea is to demonstrate that every single non-zero update of the form (10) in Algorithm 1 makes a non-negligible progress in bringing the current iterate $W^t \in \mathbb{R}^{k \times d}$ toward some global optimum $\Omega^* \in \mathbb{R}^{k \times d}$ of $L_{\mathcal{S}}(W)$, constructed based on the linear classifier weight vector ω^* . Specifically, as in the convergence proof of the Perceptron algorithm, we will establish separately a lower bound on the term $\langle W^t, \Omega^* \rangle_F$, which is the so-termed Frobenius inner product, performing a componentwise inner product of two same-size matrices as though they are vectors; and, an upper bound on the norms $||W_t||_F$ and $||\Omega^*||_F$. Both bounds will be carefully expressed as functions of the number t of *performed* non-zero updates. Recalling the Cauchy-Schwartz inequality $|\langle \mathbf{W}^t, \mathbf{\Omega}^* \rangle_F| \leq ||\mathbf{W}_t||_F ||\mathbf{\Omega}^*||_F$, the lower bound of $|\langle W^t, \Omega^* \rangle_F|$ cannot grow larger than the upper bound on $\|\boldsymbol{W}_t\|_F \|\boldsymbol{\Omega}^*\|_F$. Since every non-zero update brings the lower and upper bounds closer by a non-negligible amount, the worst case (in terms of the number of non-zero updates) is to have the two bounds equal at convergence, i.e., $|\langle W^t, \Omega^* \rangle_F| =$ $\|\boldsymbol{W}_t\|_F \|\boldsymbol{\Omega}^*\|_F$. To arrive at this equality, we are able to deduce an upper bound (due to a series of inequalities used in the proof to produce a relatively clean bound) on the number of non-zero updates by solving a univariate quadratic equality.

It will become clear in the proof that injecting random noise into just a subset of (rather than all) ReLU activity indicator functions enables us to leverage two key inequalities, namely, $\sum_{j=1}^{k} y_{it} v_j \sigma(\boldsymbol{w}_j^{T^{\top}} \boldsymbol{x}_{it}) < 1$ and $y_i \boldsymbol{\omega}^{*^{\top}} \boldsymbol{x}_i \geq 1$ for all data samples $(\boldsymbol{x}_i, y_i) \in S$. These inequalities uniquely correspond to whether an update is non-zero or not. In turn, this characterization is indeed the key to establishing the desired lower and upper bounds for the two quantities on the two sides of the Cauchy-Schwartz inequality, a critical ingredient of our convergence analysis.

C. Lower Bound

Besides the worst-case upper bound given in Theorem 1, we also provide a lower bound on the number of non-zero updates required by Algorithm 1 for convergence, which is summarized in the following theorem. The proof is provided in Appendix C.

Theorem 2 (Lower bound): Under the conditions of Theorem 1, consider Algorithm 1 with initialization $W^0 = 0$. Then for any d > 0, there exists a set of linearly separable data samples on which Algorithm 1 performs at least $\|\omega^*\|_2^2/(\eta \|v\|_2^2)$ non-zero updates to optimally train a single-hidden-layer ReLU network.

The lower bound on the number of non-zero updates to be performed in Theorem 2 matches that for learning single-hiddenlayer leaky-ReLU networks initialized from zero [4, Thm. 4]. On the other hand, it is also clear that the worst-case bound established in Theorem 1 is (significantly) loose than the lower bound here. The gap between the two bounds (in learning ReLU versus leaky ReLU networks) is indeed the price we pay for escaping bad local minima and saddle points through our noiseinjected SGD approach.

Remark 4: Considering learning v and W simultaneously, one can include the following update after (10) in Algorithm in the loop

$$oldsymbol{v}^{t+1} = oldsymbol{v}^t + \eta y_{i_t} \sigma(oldsymbol{W}^t oldsymbol{x}_{i_t}) \odot \mathbb{1}_{\{y_{i_t} oldsymbol{v}^t \geq oldsymbol{0}\}}$$

with any random initialization having at least one positive and one negative entries. Our main results in Theorems 2 and 3 still hold with different constants as entries of v^t maintaining signs during the updates.

V. GENERALIZATION

In this section, we investigate the generalization performance of training (possibly over-parameterized) ReLU networks using Algorithm 1 with randomly perturbed ReLU activity indicator functions. Toward this objective, we will rely on compression generalization bounds, specifically for the 0/1 classification error as in (3) [21].

Recall that our ReLU network has k hidden units, and a fixed second-layer weight $v \in \mathbb{R}^k$. Stressing the number of ReLUs in the subscript, let $G_k^{\text{Alg1}}(S; W^0)$ denote the classifier obtained by training the network over training set S using Algorithm 1 with initialization W^0 having rows obeying $\{||w_j^0||_2 \le \rho\}_{j=1}^k$. Let also \mathcal{G}_k denote the set of all classifiers $\{G_k\}$ obtained using any S and any W^0 , not necessarily those employed by Algorithm 1.

Suppose now that Algorithm 1 has converged after $\tau_k \leq T_k$ non-zero updates, as per Theorem 1. And let $(x_{i_1}, x_{i_2}, \ldots, x_{i_{\tau_k}})$ be the τ_k -tuple of training data from S randomly picked by SGD iterations of Algorithm 1. To exemplify the τ_k -tuple used per realization of Algorithm 1, we write $G_k^{\text{Alg1}}(S; W^0)$ $= \Gamma_{W^0}(x_{i_1}, x_{i_2}, \ldots, x_{i_{\tau_k}})$. Since τ_k can be smaller than n, function Γ_{W^0} and thus $G_k^{\text{Alg1}}(S; W^0)$ rely on compressed (down to size τ_k) versions of the n-tuples comprising the set \mathcal{H}_k [32, Definition 30.4]. Let $S_{\tau_k}^c := \{i | i \in \{1, 2, \ldots, n\} \setminus \{i_1, i_2, \ldots, i_{\tau_k}\}\}$ be the subset of training data not picked by SGD to yield $G_k^{\text{Alg1}}(S; W^0)$; and correspondingly, let R_D $(G_k^{\text{Alg1}}(S; W^0))$ denote the ensemble risk associated with G_k^{Alg1} , and $\hat{R}_{S_{\tau_k}^c}(G_k^{\text{Alg1}}(S; W^0))$ the empirical risk associated with the complement training set, namely $S_{\tau_k}^c$. With these notational conventions, our next result follows from [32, Thm. 30.2].

Theorem 3 (Compression bound): If $n \ge 2\tau_k$, then the following inequality holds with probability of at least $1 - \delta$ over the choice of S and W^0

$$R_{\mathcal{D}}(G_k^{\text{Alg1}}(\mathcal{S}; \mathbf{W}^0)) \le \hat{R}_{\mathcal{S}_{\tau_k}^c}(G_k^{\text{Alg1}}(\mathcal{S}; \mathbf{W}^0)) + \sqrt{\hat{R}_{\mathcal{S}_{\tau_k}^c}(G_k^{\text{Alg1}}(\mathcal{S}; \mathbf{W}^0)) \frac{4\tau_k \log(n/\delta)}{n}} + \frac{8\tau_k \log(n/\delta)}{n}.$$
(13)

Regarding Theorem 3, two observations are in order. The bound in (13) is non-asymptotic but as $n \to \infty$, the last two terms on the right-hand-side vanish, implying that the

ensemble risk $R_{\mathcal{D}}(G_k^{\text{Alg1}}(\mathcal{S}; \mathbf{W}^0))$ is upper bounded by the empirical risk $\hat{R}_{\mathcal{S}_{\tau_k}^c}(G_k^{\text{Alg1}}(\mathcal{S}; \mathbf{W}^0))$. Moreover, once the SGD iterations in Algorithm 1 converge, we can find the complement training set $\mathcal{S}_{\tau_k}^c$, and thus $\hat{R}_{\mathcal{S}_{\tau_k}^c}(G_k(\mathcal{S}; \mathbf{W}^0))$ can be determined. After recalling that $\hat{R}_{\mathcal{S}_{\tau_k}^c}(G_k^{\text{Alg1}}(\mathcal{S}; \mathbf{W}^0)) = 0$ holds at a global optimum of $L_{\mathcal{S}}$ by Theorem 1, we obtain from Theorems 1 and 3 the following corollary.

Corollary 1: If $n \ge 2\tau_k$, and all rows of the initialization satisfy $\{\|\boldsymbol{w}_j^0\|_2 \le \rho\}_{j=1}^k$, then the following holds with probability at least $1 - \delta$ over the choice of S

$$R_{\mathcal{D}}(G_k^{\text{Alg1}}(\mathcal{S}; \boldsymbol{W}^0)) \le \frac{8T_k \log(n/\delta)}{n}$$
(14)

where T_k is given in Theorem 1.

Expressed differently, the bound in (14) suggests that in order to guarantee a low generalization error, one requires *in the worst case* about $n = O(k^2 || \boldsymbol{\omega}^* ||_2^2)$ training data to *reliably* learn a two-layer ReLU network of k hidden neurons. This holds true *despite* the fact that Algorithm 1 can achieve a zero training loss *regardless of* the training size n. One implication of Corollary 1 is a fundamental difference in the sample complexity for generalization between training a ReLU network (at least in the worst case), versus training a α -leaky ReLU network ($0 < \alpha < 1$), which at most needs $n = O(|| \boldsymbol{\omega}^* ||_2^2 / \alpha^2)$ data to be trained via SGD-type algorithms.

VI. NUMERICAL TESTS

To validate our theoretical results, this section evaluates the empirical performance of Algorithm 1 using both synthetic data and real data. To benchmark Algorithm 1, we also simulated the plain-vanilla SGD. To compare between the two algorithms as fair as possible, the same initialization W^0 , constant step size $\eta > 0$, and data random sampling scheme were employed. For reproducibility, the Matlab code of Algorithm 1 is publicly available at https://gangwg.github.io/RELUS/.

A. Synthetic Data

We consider first two synthetic tests using data generated from Gaussian as well as uniform distributions. In the first test, feature vectors $\{x_i \in \mathbb{R}^d\}_{i=1}^n$ were sampled i.i.d. from a standardized Gaussian distribution $\mathcal{N}(\mathbf{0}, \mathbf{I}_d)$, and classifier $\omega^* \in \mathbb{R}^d$ was drawn from $\mathcal{N}(\mathbf{0}, \mathbf{I}_d)$. Labels $y_i \in \{+1, -1\}$ were generated according to $y_i = \operatorname{sgn}(\omega^{*\top}x_i)$. To further yield $y_i\omega^{*\top}x_i \ge 1$ for all $i = 1, 2, \ldots, n$, we normalized ω^* by the smallest number among $\{y_i\omega^{*\top}x_i\}_{i=1}^n$. We performed 100 independent experiments with d = 128, and over a varying set of $n \in \{20, 40, \ldots, 200\}$ training samples using ReLU networks comprising $k \in \{2, 4, \ldots, 20\}$ hidden neurons. The second layer weight vector $v \in \mathbb{R}^k$ was kept fixed with the first $\lfloor k/2 \rfloor$ entries being +1 and the remaining being -1. For fixed n and k, each experiment used a random initialization generated from $\mathcal{N}(\mathbf{0}, 0.01\mathbf{I})$, step size $\eta = 0.01$, and noise variance $\gamma = 100$, along with a maximum of 5,000 effective data passes.

Figure 2 depicts our results, where we display success rates of the plain-vanilla SGD (top panel) and our noise-injected SGD in Algorithm 1 (bottom panel); each plot presents results obtained from the 100 experiments. Within each plot, a



Fig. 2. Empirical success rates of plain-vanilla SGD (top panel) and the Algorithm 1 (bottom panel) for learning two-layer ReLU networks of k hidden units on n randomly generated data samples of dimension d = 128.



Fig. 3. Empirical success rates of the 'plain-vanilla' SGD for learning two-layer ReLU networks of k hidden units on n data samples of dimension: (left) d = 128 generated from uniform distribution $\mathcal{U}[-\mathbf{1}_{128}, \mathbf{1}_{128}]$, and (right) d = 4 from the UCI machine learning repository [6].

white square signifies that 100% of the trials were successful, meaning that the learned ReLU network yields a training loss $L_{\{(\boldsymbol{x}_i, y_i)\}_{i=1}^n}(\boldsymbol{W}^T) \leq 10^{-10}$, while black squares indicate 0% success rates. It is evident that the developed Algorithm 1 trained all considered ReLU networks to global optimality, while plainvanilla SGD can get stuck with bad local minima, for small k in particular. The bottom panel confirms that Algorithm 1 achieves optimal learning of single-hidden-layer ReLU networks on separable data, regardless of the network size, the number of training samples, and the initialization. The top panel however, suggests that learning ReLU networks becomes easier with plain-vanilla SGD as k grows larger, namely as the network becomes 'more over-parameterized.'

We repeated the first test using synthetic data $\{x \in \mathbb{R}^{128}\}_{i=1}^{n}$ as well as classifier $\omega \in \mathbb{R}^{128}$ generated i.i.d. from the uniform distribution $\mathcal{U}[-\mathbf{1}_{128}, \mathbf{1}_{128}]$. All other settings were kept the same. Success rates of plain-vanilla SGD are plotted in Fig. 3 (left panel), while those of the proposed Algorithm 1 are omitted, as they are 100% successful in all simulated tests.

B. Real Data

Performance of Algorithm 1 for training (over-)parameterized ReLU networks is further corroborated using two real datasets: *iris* in UCI's machine learning repository [6], and *MNIST* images.² The iris dataset contains 150 four-dimensional feature vectors belonging to three classes. To obtain a two-class linearly separable dataset, the first-class data vectors were relabeled +1, while the remaining were relabeled -1. We performed 100 independent experiments over a varying set of $n \in$

²Downloaded from http://yann.lecun.com/exdb/mnist/



Fig. 4. Empirical success rates of plain-vanilla SGD (top panel) and Algorithm 1 (bottom panel) for learning two-layer ReLU networks of k hidden units on n MNIST images of digits 3 and 5.

{30, 60, 90, 120, 150} training samples using ReLU networks with $k \in \{2, 4, 6, 8, 10\}$ hidden neurons. Gaussian initialization from $\mathcal{N}(\mathbf{0}, \mathbf{I})$, step size $\eta = 0.1$, noise variance $\gamma = 10$, and a maximum of 100 effective data passes were simulated. Success rates of plain-vanilla SGD are given in Fig. 3 (right). Again, Algorithm 1 achieves a 100% success rate in all simulated settings.

The linearly separable MNIST dataset collects 2,000 images of digits 3 (labeled +1) and 5 (labeled -1), each having dimension 784. We performed 100 independent experiments over a varying set of $n \in \{200, 400, \ldots, 2, 000\}$ training samples using ReLU networks with $k \in \{2, 4, \ldots, 40\}$ hidden neurons. The constant step size of both plain-vanilla SGD and Algorithm 1 was set to $\eta = 0.001$ ($\eta = 0.01$) when the ReLU networks have $k \leq 4$ (k > 4) hidden units, while the noise variance in Algorithm 1 was set to $\gamma = 10$. Similar to the first experiment on randomly generated data, we plot success rates of the plain-vanilla SGD (top panel) and our noise-injected SGD (bottom panel) algorithms over training sets of MNIST images in Fig. 4. It is self-evident that Algorithm 1 achieved a 100% success rate under all testing conditions, which confirms our theoretical results in Theorem 1, and it markedly improves upon its plain-vanilla SGD alternative.

VII. CONCLUSIONS

This paper approached the task of training ReLU networks from a non-convex optimization point of view. Focusing on the task of binary classification with a hinge loss criterion, this contribution put forth the first algorithm that can provably yet efficiently train any single-hidden-layer ReLU network to global optimality, provided that the data are linearly separable. The algorithm is as simple as plain-vanilla SGD, but it is able to exploit the power of random additive noise to break 'optimality' of the SGD learning process at any sub-optimal critical point. We established an upper and a lower bound on the number of non-zero updates that the novel algorithm requires for convergence to a global optimum. Our result holds regardless of the underlying data distribution, network/training size, or initialization. We further developed generalization error bounds for two-layer NN classifiers with ReLU activations, which provide the first theoretical guarantee for the generalization behavior of ReLU networks trained with SGD. A comparison of such bounds with those of a leaky ReLU network versus that of a leaky ReLU network in the sample complexity required for generalization.

Since analysis, comparisons, and corroborating tests focus on single-hidden-layer networks with a hinge loss criterion here, our future work will naturally aim at generalizing the novel noise-injection design to SGD for multilayer ReLU networks, and considering alternative loss functions, and generalizations to (multi-)kernel based approaches.

APPENDIX

A. Proof of Theorem 1

Consider Algorithm 1 has performed t > 0 non-zero updates with a sufficiently large noise variance γ^2 . Observe that if all data $(x_{i_t}, y_{i_t}) \in S$ lead to zero update after a succession of say e.g., pn > 0 iterations, then Algorithm 1 has reached a global minimum with high probability as per Proposition 1. Let $\operatorname{vec}(\boldsymbol{W}^{\top}) := [\boldsymbol{w}_1^{\top} \, \boldsymbol{w}_2^{\top} \, \cdots \, \boldsymbol{w}_k^{\top}]^{\top} \in \mathbb{R}^{nk \times 1}, \text{ and define}$

$$\mathbf{\Omega}^{\star} := \frac{1}{v_{\min}} \operatorname{sgn}(\boldsymbol{v}) \otimes \boldsymbol{\omega}^{*\top} \in \mathbb{R}^{k \times n}$$
(15)

which is constructed from the optimum ω^* of the linear classifier, where $v_{\min} := \min_{1 \le j \le k} |v_j| > 0$. Using the definition of Ω^* and (4), it holds that

$$L_{\mathcal{S}}(\mathbf{\Omega}^*) = \frac{1}{n} \sum_{i=1}^{n} \max\left\{ 0, 1 - y_i \sum_{j=1}^{k} v_j \sigma\left(\operatorname{sgn}(v_j) \,\boldsymbol{\omega}^{*\top} \boldsymbol{x}_i \right) \right\}.$$
(16)

For $i \in \mathcal{N}_y^+$, we clearly have y = 1, and $\boldsymbol{\omega}^* \boldsymbol{x}_i > 0$. In addition, for $j \in \mathcal{N}_v^+$, we find $v_j > 0$ and hence $\operatorname{sgn}(v_j) \boldsymbol{\omega}^{*\top} \boldsymbol{x}_i > 0$, which implies $\sigma(\operatorname{sgn}(v_j) \boldsymbol{\omega}^{*\top} \boldsymbol{x}_i) = \boldsymbol{\omega}^{*\top} \boldsymbol{x}_i$; similarly, for $j \in \mathcal{N}_v^-$, we find $v_j < 0$, and thus $\operatorname{sgn}(v_j) \boldsymbol{\omega}^{*\top} \boldsymbol{x}_i < 0$, which yields $\sigma(\operatorname{sgn}(v_j) \boldsymbol{\omega}^{*\top} \boldsymbol{x}_i) = 0$. These considerations show that for $i \in \mathcal{N}_y^+$ in (16), only summands with $j \in \mathcal{N}_v^+$ survive; and arguing along the same lines, we deduce that for $i \in \mathcal{N}_y^-$, only summands with $j \in \mathcal{N}_v^-$ should be present. All in all, (16) reduces to

$$L_{\mathcal{S}}(\mathbf{\Omega}^{*}) = \frac{1}{n} \sum_{i \in \mathcal{N}_{y}^{+}} \max\left\{0, 1 - \sum_{j \in \mathcal{N}_{v}^{+}} \frac{v_{j}}{v_{\min}} y_{i} \boldsymbol{\omega}^{*\top} \boldsymbol{x}_{i}\right\}$$
$$+ \frac{1}{n} \sum_{i \in \mathcal{N}_{y}^{-}} \max\left\{0, 1 + \sum_{j \in \mathcal{N}_{v}^{-}} \frac{v_{j}}{v_{\min}} y_{i} \boldsymbol{\omega}^{*\top} \boldsymbol{x}_{i}\right\}. (17)$$

But since $y_i \boldsymbol{\omega}^{*\top} \boldsymbol{x}_i \geq 1$, and $\sum_{v_j>0} v_j/v_{\min} \geq 1$ as well as $\sum_{v_j<0} v_j/v_{\min} \leq -1$, we infer that $L_{\mathcal{S}}(\boldsymbol{\Omega}^*) = 0$, and hence $\boldsymbol{\Omega}^*$ is indeed a global minimum of $L_{\mathcal{S}}(\boldsymbol{W})$.

The subsequent analysis builds critically on the following two functions (cf. (15))

$$\phi(\boldsymbol{W}^t) :== \left\langle \operatorname{vec}\left(\boldsymbol{W}^{t^{\top}}\right), \operatorname{vec}\left(\boldsymbol{\Omega}^{*^{\top}}\right) \right\rangle$$
(18a)

$$\psi(\boldsymbol{W}^{t}) := \left\| \operatorname{vec} \left(\boldsymbol{W}^{t^{\top}} \right) \right\|_{2} = \left(\sum_{j=1}^{k} \left\| \boldsymbol{w}_{j}^{t} \right\|_{2}^{2} \right)^{1/2}$$
(18b)

Using the Cauchy-Schwartz inequality, we can write

$$\frac{|\phi(\boldsymbol{W}^{t})|}{\psi(\boldsymbol{W}^{t})\psi(\boldsymbol{\Omega}^{*})} = \frac{\left|\left\langle \operatorname{vec}\left(\boldsymbol{W}^{t^{\top}}\right), \operatorname{vec}\left(\boldsymbol{\Omega}^{*^{\top}}\right)\right\rangle\right|}{\left\|\operatorname{vec}\left(\boldsymbol{W}^{t^{\top}}\right)\right\|_{2}\left\|\operatorname{vec}\left(\boldsymbol{\Omega}^{*^{\top}}\right)\right\|_{2}} \le 1.$$
(19)

We will next derive a lower and an upper bound for the numerator and denominator of (19). Consider an iteration t for which Algorithm 1 admits a non-zero update, meaning that $\mathbb{1}_{\{1-y_{i_t}\boldsymbol{v}^{\top}\sigma(\boldsymbol{W}^t\boldsymbol{x}_{i_t})>0\}} = 1$, or equivalently, $y_{i_t}\boldsymbol{v}^{\top}\sigma(\boldsymbol{W}^t\boldsymbol{x}_{i_t}) = \sum_{j=1}^{k} y_{i_t}v_j\sigma(\boldsymbol{w}_j^{\top}\boldsymbol{x}_{i_t}) < 1$. It will also come handy to rewrite (10) row-wise as

$$\boldsymbol{w}_{j}^{t+1} = \boldsymbol{w}_{j}^{t} + \eta y_{i_{t}} v_{j} \mathbb{1}_{\left\{\boldsymbol{w}_{j}^{t^{\top}} \boldsymbol{x}_{i_{t}} + \boldsymbol{\epsilon}_{j}^{t} \ge 0\right\}} \boldsymbol{x}_{i_{t}}, \, j = 1, \, 2, \, \dots, \, k.$$
(20)

Combining (20) with (18b), we can upper bound $\psi^2(\mathbf{W}^t)$ in the denominator of (19) as

$$\psi^{2}(\mathbf{W}^{t+1}) = \sum_{j=1}^{k} \|\mathbf{w}_{j}^{t+1}\|_{2}^{2}$$

$$\stackrel{(a)}{=} \sum_{j=1}^{k} \left(\|\mathbf{w}_{j}^{t}\|_{2}^{2} + \eta^{2}v_{j}^{2}\|\mathbf{x}_{i_{t}}\|_{2}^{2}\mathbb{1}_{\left\{\mathbf{w}_{j}^{t^{\top}}\mathbf{x}_{i_{t}} + \epsilon_{j}^{t} \ge 0\right\}}\right)$$

$$+ 2\eta y_{i_{t}}v_{j} \mathbf{w}_{j}^{t^{\top}}\mathbf{x}_{i_{t}}\mathbb{1}_{\left\{\mathbf{w}_{j}^{t^{\top}}\mathbf{x}_{i_{t}} + \epsilon_{j}^{t} \ge 0\right\}}\right)$$

$$\stackrel{(b)}{\leq} \sum_{j=1}^{k} \|\mathbf{w}_{j}^{t}\|_{2}^{2} + \eta^{2} \sum_{j=1}^{k} v_{j}^{2}$$

$$+ 2\eta \sum_{j=1}^{k} y_{i_{t}}v_{j}\sigma\left(\mathbf{w}_{j}^{t^{\top}}\mathbf{x}_{i_{t}}\right)$$

$$\stackrel{(c)}{\leq} \psi^{2}(\mathbf{W}^{t}) + \eta^{2} \|\mathbf{v}\|_{2}^{2} + 2\eta \qquad (21)$$

where (a) follows directly from (20) after expanding the squares; (b) uses the working condition $\|\boldsymbol{x}_{i_t}\|_2 \leq 1$ adopted without loss of generality, as well as the fact that $\mathbb{1}_{\{\cdot\}} \leq 1$ holds true for any event, and also the inequality $\sum_{j=1}^{k} y_{i_t} v_j \boldsymbol{w}_j^{t^{\top}} \boldsymbol{x}_{i_t}$ $\mathbb{1}_{\{\boldsymbol{w}_j^{t^{\top}} \boldsymbol{x}_{i_t} + \boldsymbol{\epsilon}_j^t \geq 0\}}) \leq \sum_{j=1}^{k} y_{i_t} v_j \sigma(\boldsymbol{w}_j^{t^{\top}} \boldsymbol{x}_{i_t})$ established in Lemma 1 below, whose proof is postponed to Appendix B for readability. Finally, (c) is due to the non-zero update at iteration t, which implies that $\sum_{j=1}^{k} y_{i_t} v_j \sigma(\boldsymbol{w}_j^{t^{\top}} \boldsymbol{x}_{i_t}) < 1$.

Lemma 1: For any $(x, y) \in S$, and any $\{w_j \in \mathbb{R}^d\}_{j=1}^k$ and $v \in \mathbb{R}^k$, it holds that

$$\sum_{j=1}^{k} y v_j \boldsymbol{w}_j^{\top} \boldsymbol{x} \mathbb{1}_{\left\{\boldsymbol{w}_j^{\top} \boldsymbol{x} + \boldsymbol{\epsilon}_j \ge 0\right\}} \le \sum_{j=1}^{k} y v_j \sigma \left(\boldsymbol{w}_j^{\top} \boldsymbol{x}\right)$$
(22)

if the entries of the additive noise $\epsilon \in \mathbb{R}^k$ satisfy $\epsilon_j \sim \mathcal{N}(0, \gamma^2)$, when $yv_j \ge 0$, and $\epsilon_j = 0$, otherwise.

Writing down (21) for the already executed t non-zero updates and by means of telescoping, we obtain

$$\psi^{2}(\mathbf{W}^{t}) \leq \psi^{2}(\mathbf{W}^{0}) + t\left(\eta^{2} \|\boldsymbol{v}\|_{2}^{2} + 2\eta\right).$$
 (23)

We now turn to deriving a lower bound for $\phi(\mathbf{W}^t)$ in (19), starting with (cf. (18a) and (18b))

$$\begin{split} \phi(\boldsymbol{W}^{t+1}) &= \frac{1}{v_{\min}} \sum_{j \in \mathcal{N}_{v}^{+}} \boldsymbol{w}_{j}^{t+1^{\top}} \boldsymbol{\omega}^{*} - \frac{1}{v_{\min}} \sum_{j \in \mathcal{N}_{v}^{-}} \boldsymbol{w}_{j}^{t+1^{\top}} \boldsymbol{\omega}^{*} \\ &\stackrel{(a)}{=} \frac{1}{v_{\min}} \sum_{j \in \mathcal{N}_{v}^{+}} \boldsymbol{w}_{j}^{t^{\top}} \boldsymbol{\omega}^{*} - \frac{1}{v_{\min}} \sum_{j \in \mathcal{N}_{v}^{-}} \boldsymbol{w}_{j}^{t^{\top}} \boldsymbol{\omega}^{*} \\ &+ \frac{1}{v_{\min}} \sum_{j \in \mathcal{N}_{v}^{+}} \eta v_{j} y_{i_{t}} \boldsymbol{x}_{i_{t}}^{\top} \boldsymbol{\omega}^{*} \mathbb{1}_{\left\{\boldsymbol{w}_{j}^{t^{\top}} \boldsymbol{x}_{i_{t}} + \boldsymbol{\epsilon}_{j}^{t} \ge 0\right\}} \\ &- \frac{1}{v_{\min}} \sum_{j \in \mathcal{N}_{v}^{-}} \eta v_{j} y_{i_{t}} \boldsymbol{x}_{i_{t}}^{\top} \boldsymbol{\omega}^{*} \mathbb{1}_{\left\{\boldsymbol{w}_{j}^{t^{\top}} \boldsymbol{x}_{i_{t}} + \boldsymbol{\epsilon}_{j}^{t} \ge 0\right\}} \\ &\stackrel{(b)}{=} \phi(\boldsymbol{W}^{t}) + \eta y_{i_{t}} \boldsymbol{x}_{i_{t}}^{\top} \boldsymbol{\omega}^{*} \sum_{j=1}^{k} \frac{|v_{j}|}{v_{\min}} \mathbb{1}_{\left\{\boldsymbol{w}_{j}^{t^{\top}} \boldsymbol{x}_{i_{t}} + \boldsymbol{\epsilon}_{j}^{t} \ge 0\right\}} \\ &\stackrel{(c)}{\geq} \phi(\boldsymbol{W}^{t}) + \eta \end{split}$$

t

where (a) is derived by plugging in (20); (b) uses $v_j > 0$ (<0) if $j \in \mathcal{N}_v^+$ ($\in \mathcal{N}_v^-$); and (c) follows from the two critical inequalities: i) $y_i \boldsymbol{x}_i^\top \boldsymbol{\omega}^* \ge 1$ for all $(\boldsymbol{x}_i, y_i) \in \mathcal{S}$, and ii) $(|v_j|/v_{\min}) \mathbb{1}_{\{\boldsymbol{w}_j^{\dagger}^\top \boldsymbol{x}_{i_t} + \epsilon_j^t \ge 0\}} \ge 1$, because a non-zero update at iteration t asserts that at least one out of the k ReLU activity indicator functions $\{\mathbb{1}_{\{\boldsymbol{w}_j^{\dagger}^\top \boldsymbol{x}_{i_t} + \epsilon_s^t \ge 0\}}\}_{j=1}^k$ equals one.

Again, telescoping the t recursions (24) for the non-zero updates 0 to (t-1), yields

$$\phi(\boldsymbol{W}^t) \ge \phi(\boldsymbol{W}^0) + t\eta. \tag{25}$$

Substituting the bounds in (23) and (25) into (19), we have that

$$\begin{split} \phi(\boldsymbol{W}^{0}) + t\eta &\leq \left|\phi(\boldsymbol{W}^{t})\right| \\ &\leq \psi(\boldsymbol{W}^{t})\psi(\boldsymbol{\Omega}^{*}) \\ &= \sqrt{\psi^{2}(\boldsymbol{W}^{0}) + t\left(\eta^{2} \|\boldsymbol{v}\|_{2}^{2} + 2\eta\right)} \ \psi(\boldsymbol{\Omega}^{*}). \end{split}$$
(26)

Using further that $\sqrt{p^2 + q^2} \le |p| + |q|$, we arrive at $\phi(\mathbf{W}^0) + t\eta \le |\phi(\mathbf{W}^t)|$

$$\leq \psi(\boldsymbol{W}^{0})\psi(\boldsymbol{\Omega}^{*}) + \sqrt{t} \cdot \sqrt{\eta^{2} \|\boldsymbol{v}\|_{2}^{2} + 2\eta} \ \psi(\boldsymbol{\Omega}^{*}).$$
(27)

Using that $[sgn(v_j)]^2 = 1$, it is easy to verify that

$$\psi(\mathbf{\Omega}^*) := \left\| \operatorname{vec}\left(\mathbf{\Omega}^{*\top}\right) \right\|_2 = \left\|\mathbf{\Omega}^*\right\|_F$$
$$= \frac{1}{v_{\min}} \left\| \operatorname{sgn}(v) \otimes \boldsymbol{\omega}^{*\top} \right\|_F = \frac{\sqrt{k}}{v_{\min}} \|\boldsymbol{\omega}^*\|_2.$$
(28)

Under our assumption that all rows of W^0 satisfy $||w_j^0||_2 \le \rho$, we have for $\psi(W^0) := ||W^0||_F$ that

$$\psi(\boldsymbol{W}^0) \le \sqrt{k}\rho \tag{29}$$

Using (18a) along with (28) and (29), we find

$$\phi(\boldsymbol{W}^{0}) = \left\langle \operatorname{vec}\left(\boldsymbol{W}^{0^{\top}}\right), \operatorname{vec}\left(\boldsymbol{\Omega}^{*^{\top}}\right) \right\rangle$$

$$\geq - \left\| \operatorname{vec}\left(\boldsymbol{W}^{0^{\top}}\right) \right\|_{2} \left\| \operatorname{vec}\left(\boldsymbol{\Omega}^{*^{\top}}\right) \right\|_{2}$$

$$= -\psi(\boldsymbol{W}^{0})\psi(\boldsymbol{\Omega}^{*})$$

$$\geq -\frac{k\rho}{v_{\min}} \left\| \boldsymbol{\omega}^{*} \right\|_{2}.$$
(30)

Substituting the bounds in (28), (29), and (30) into (27) and re-arranging terms, we further arrive at

$$\eta v_{\min} t \le \|\boldsymbol{\omega}^*\|_2 \sqrt{k \left(\eta^2 \|\boldsymbol{v}\|_2^2 + 2\eta\right)} \sqrt{t} + 2k\rho \|\boldsymbol{\omega}^*\|_2 \quad (31)$$

which upon letting $z := \sqrt{t} \ge 0$, boils down to the quadratic inequality

$$az^2 + bz + c \le 0 \qquad \text{s. to} \qquad z \ge 0 \tag{32}$$

where the coefficients are given by $a = \eta v_{\min} > 0, b = -\|\boldsymbol{\omega}^*\|_2$ $\sqrt{k(\eta^2 \|\boldsymbol{v}\|_2^2 + 2\eta)}$, and $c = -2k\rho \|\boldsymbol{\omega}^*\|_2 < 0$. Because c < 0 and $b^2 - 4ac > 0$, we have real roots of opposite sign, which implies that (32) is satisfied for

$$z \in \left[0, \ \frac{-b + \sqrt{b^2 - 4ac}}{2a}\right]. \tag{33}$$

Plugging in those coefficients and appealing again to the inequality $\sqrt{p^2+q^2} \le |p|+|q|$, we deduce that

$$= z^{2} \leq \frac{b^{2} + b^{2} - 4ac - 2b\sqrt{b^{2} - 4ac}}{4a^{2}}$$

$$\leq \frac{b^{2}}{2a^{2}} - \frac{c}{a} + \frac{b^{2}}{2a^{2}} - \frac{b\sqrt{-ac}}{a^{2}}$$

$$\leq \frac{k\left(\eta^{2} \|v\|_{2}^{2} + 2\eta\right) \|\omega^{*}\|_{2}^{2}}{2\eta^{2}v_{\min}^{2}} + \frac{2k\rho\eta v_{\min} \|\omega^{*}\|_{2}^{2}}{\eta^{2}v_{\min}^{2}}$$

$$+ \frac{k\left(\eta^{2} \|v\|_{2}^{2} + 2\eta\right) \|\omega^{*}\|_{2}^{2}}{2\eta^{2}v_{\min}^{2}}$$

$$+ \frac{\|\omega^{*}\|_{2}\sqrt{k\left(\eta^{2} \|v\|_{2}^{2} + 2\eta\right)}\sqrt{2k\rho\eta v_{\min} \|\omega^{*}\|_{2}}}{\eta^{2}v_{\min}^{2}}$$

$$= \frac{k}{\eta v_{\min}^{2}} \left[\left(\eta \|v\|_{2}^{2} + 2\right) \|\omega^{*}\|_{2}^{2} + 2\rho v_{\min} \|v\|_{2}^{2} + \sqrt{2\rho v_{\min} \|\omega^{*}\|_{2} \left(\eta \|v\|_{2}^{2} + 2\right)} \|\omega^{*}\|_{2} \right]$$

$$\stackrel{\triangle}{=} T_{k}. \qquad (34)$$

By taking $\rho = 0$ in (34), one finally confirms that the maximum number of non-zero updates for Algorithm 1 initialized with $W^0 = 0$ until convergence, is

$$T_{k}^{0} := \frac{k}{\eta v_{\min}^{2}} \left(\eta \left\| \boldsymbol{v} \right\|_{2}^{2} + 2 \right) \left\| \boldsymbol{\omega}^{*} \right\|_{2}^{2}$$
(35)

which completes the proof.

B. Proof of Lemma 1

We first prove that the following inequality holds per hidden neuron j = 1, 2, ..., k

$$yv_j \boldsymbol{w}_j^{\top} \boldsymbol{x} \mathbb{1}_{\left\{\boldsymbol{w}_j^{\top} \boldsymbol{x} + \epsilon_j \ge 0\right\}} \le yv_j \sigma \left(\boldsymbol{w}_j^{\top} \boldsymbol{x}\right).$$
 (36)

Depending on whether the *j*-th ReLU is active or not $(\boldsymbol{w}_j^{\top} \boldsymbol{x} \ge 0)$ and (Gaussian) noise is injected or not $(v_j y \ge 0)$, we consider separately the following four cases:

- c1) $\boldsymbol{w}_{j}^{\top}\boldsymbol{x} \geq 0$ and $v_{j}y \geq 0$ (ReLU active and noise injected);
- c2) $\boldsymbol{w}_{j}^{\top}\boldsymbol{x} \geq 0$ and $v_{j}\boldsymbol{y} < 0$ (ReLU active and no noise);
- c3) $\boldsymbol{w}_j^\top \boldsymbol{x} < 0$ and $v_j y \ge 0$ (ReLU inactive and noise injected); and,
- c4) $\boldsymbol{w}_{j}^{\top}\boldsymbol{x} < 0$ and $v_{j}y < 0$ (ReLU inactive and no noise).

For c1), the right-hand-side of (36) satisfies

$$yv_j\sigma(\boldsymbol{w}_j^{\top}\boldsymbol{x}) = yv_j\boldsymbol{w}_j^{\top}\boldsymbol{x}.$$
 (37)

Regarding the left-hand-side, it takes values depending on whether ϵ_j changes the state of the ReLU activity indicator function, it leads to a two-branch inequality

$$yv_j \boldsymbol{w}_j^{\top} \boldsymbol{x} \mathbb{1}_{\left\{\boldsymbol{w}_j^{\top} \boldsymbol{x} + \epsilon_j \ge 0\right\}} = \begin{cases} yv_j \boldsymbol{w}_j^{\top} \boldsymbol{x}, \ \epsilon_j \ge -\boldsymbol{w}_j^{\top} \boldsymbol{x} \\ 0, \qquad \epsilon_j < -\boldsymbol{w}_j^{\top} \boldsymbol{x} \end{cases}.$$
(38)

Combining (37) with (38), we deduce that $yv_j \boldsymbol{w}_j^{\top} \boldsymbol{x} \mathbb{1}_{\{\boldsymbol{w}_j^{\top} \boldsymbol{x} + \epsilon_j \ge 0\}}$ $\leq yv_j \sigma(\boldsymbol{w}_j^{\top} \boldsymbol{x})$ holds under c1).

For c2), the *j*-th ReLU is active too, but there is no noise injection, namely $\epsilon_j = 0$. The right-hand-side of (37) still holds however. It is also not difficult to check the left-hand-side term $yv_j \boldsymbol{w}_j^\top \boldsymbol{x} \mathbb{1}_{\{\boldsymbol{w}_j^\top \boldsymbol{x} + \epsilon_j \ge 0\}} = yv_j \boldsymbol{w}_j^\top \boldsymbol{x}$. Evidently, the desired inequality holds with equality in this case.

For c3), we have $\boldsymbol{w}_{j}^{\top}\boldsymbol{x} < 0$, meaning that the *j*-th ReLU is inactive, and therefore, the right-hand-side of (36) becomes $yv_{j}\sigma(\boldsymbol{w}_{j}^{\top}\boldsymbol{x}) = 0$. However, given $v_{j}y \ge 0$, there is a noise injection. Hence, the left-hand-side can be similarly treated as in c2), to infer that (38) remains valid. Recalling again that $\boldsymbol{w}_{j}^{\top}\boldsymbol{x} < 0$ and $yv_{j} \ge 0$, one deduces that $yv_{j}\boldsymbol{w}_{j}^{\top}\boldsymbol{x}\mathbb{1}_{\{\boldsymbol{w}_{j}^{\top}\boldsymbol{x}+\epsilon_{j}\ge 0\}} \le 0$, regardless of ϵ_{j} . Thus, the inequality under consideration is also true under c2).

Finally, for c4), the j-th ReLU is inactive, and there is no noise injection. It is straightforward to verify that both the left-hand-side and right-hand-side equal zero, and (36) holds with equality as well.

Putting together c1)-c4), we have established that $yv_jw_j^{\dagger}x$ $\mathbb{1}_{\{w_j^{\top}x+\epsilon_j\geq 0\}} \leq yv_j\sigma(w_j^{\top}x)$ for j = 1, 2, ..., k. Summing up such inequalities for all k hidden neurons completes the proof.

C. Proof of Theorem 2

Let $\{e_i \in \mathbb{R}^d\}_{i=1}^d$ be the canonical basis of \mathbb{R}^d . Consider the following set $S_1 \subseteq \mathcal{X} \times \mathcal{Y}$ of d training data from the 'positive' class

$$S_1 := \{ (e_1, 1), (e_2, 1), \dots, (e_d, 1) \}.$$
 (39)

Letting $\boldsymbol{\omega}^* := [1 \ 1 \ \cdots \ 1]^\top \in \mathbb{R}^d$, it is clear that the linear classifier $y = \boldsymbol{w}_j^\top \boldsymbol{x}$ correctly classifies all the *d* data points in S. Note also that $\|\boldsymbol{\omega}^*\|_2^2 = d$ in this case.

Consider the update in (10) initialized with $W^0 = 0$, and telescope each row to obtain

$$\boldsymbol{w}_{j}^{t+1} = \eta v_{j} \sum_{i=1}^{t} \mathbb{1}_{\left\{\boldsymbol{w}_{j}^{t^{\top}} \boldsymbol{e}_{i_{t}} + \boldsymbol{\epsilon}_{j}^{t} \ge 0\right\}} \boldsymbol{e}_{i_{t}}, \ j = 1, 2, \dots, k \quad (40)$$

where i_t deterministically cycles through $\{1, 2, \ldots, d\}$.

At the global optimum, say W^{τ} for some iteration number τ , it holds for $(e_s, 1) \in S_1$ that

$$y_s f(\boldsymbol{e}_s; \boldsymbol{W}) = \sum_{j=1}^k v_j \sigma\left(\boldsymbol{w}_j^{\tau \top} \boldsymbol{e}_s\right)$$

$$=\sum_{j\in\mathcal{N}_{v}^{+}}v_{j}\sigma\left(\boldsymbol{w}_{j}^{\tau\top}\boldsymbol{e}_{s}\right)-\sum_{j\in\mathcal{N}_{v}^{-}}|v_{j}|\sigma\left(\boldsymbol{w}_{j}^{\tau\top}\boldsymbol{e}_{s}\right)$$

$$\geq1.$$
(41)

Since $|v_j|\sigma(w_j^{\tau \top} e) \ge 0$ in (41), a necessary condition for W^{τ} to be a global minimum is (cf. (40))

$$1 \leq \sum_{j \in \mathcal{N}_{v}^{+}} v_{j} \sigma \left(\boldsymbol{w}_{j}^{\tau-1^{\top}} \boldsymbol{e}_{s} \right) = \sum_{j \in \mathcal{N}_{v}^{+}} v_{j} \left\langle \boldsymbol{w}_{j}^{\tau}, \boldsymbol{e}_{s} \right\rangle$$
$$= \sum_{j \in \mathcal{N}_{v}^{+}} v_{j} \left\langle \eta v_{j} \sum_{t=1}^{\tau-1} \mathbb{1}_{\left\{ \boldsymbol{w}_{j}^{t^{\top}} \boldsymbol{e}_{i_{t}} + \boldsymbol{\epsilon}_{j}^{t} \geq 0 \right\}} \boldsymbol{e}_{i_{t}}, \boldsymbol{e}_{s} \right\rangle$$
$$\geq 1. \tag{42}$$

Assume for simplicity that $\tau - 1$ is a multiple of d, namely $\tau - 1 = dp$ for some integer $p \ge 1$. On the other hand, we have from (42) that

$$\sum_{j \in \mathcal{N}_{v}^{+}} v_{j} \left\langle \eta v_{j} \sum_{i=1}^{\tau-1} \mathbb{1}_{\left\{ \boldsymbol{w}_{j}^{t^{\top}} \boldsymbol{e}_{i_{t}} + \boldsymbol{\epsilon}_{j}^{t} \ge 0 \right\}} \boldsymbol{e}_{i_{t}}, \boldsymbol{e}_{s} \right\rangle$$

$$\leq \sum_{j \in \mathcal{N}_{v}^{+}} \eta v_{j}^{2} \left\langle \sum_{i=1}^{\tau-1} \boldsymbol{e}_{i_{t}}, \boldsymbol{e}_{s} \right\rangle$$

$$= \eta p \sum_{j \in \mathcal{N}_{v}^{+}} v_{j}^{2}$$

$$\leq p \eta \| \boldsymbol{v} \|_{2}^{2}$$
(43)

where we have used the following inequalities: i) $\mathbb{1}_{\{\boldsymbol{w}_{j}^{t^{\top}}\boldsymbol{e}_{i_{t}}+\boldsymbol{\epsilon}_{j}^{t}\geq 0\}} \leq 1$, ii) $\sum_{i=1}^{\tau-1} \boldsymbol{e}_{i_{t}} = \sum_{i=1}^{p} \mathbf{1}$ with $\mathbf{1}$ being an all-one vector of suitable dimension that is clear from the context, and iii) $\sum_{i \in \mathcal{N}_{r}^{+}} v_{i}^{2} \leq \|\boldsymbol{v}\|_{2}^{2}$.

Combing the bounds in (42) and (43), we obtain that $p\eta \|\boldsymbol{v}\|_2^2 \ge 1$, or equivalently $p \ge 1/(\eta \|\boldsymbol{v}\|_2^2)$. Hence, to find a global optimum, Algorithm 1 initialized from $\boldsymbol{W}^0 = \boldsymbol{0}$ makes at least

$$M_k^0 \ge \frac{d}{\eta \|\boldsymbol{v}\|_2^2} = \frac{\|\boldsymbol{\omega}^*\|_2^2}{\eta \|\boldsymbol{v}\|_2^2}$$
(44)

mistakes. This concludes the proof.

D. Proof of Proposition 1

We have established in (23) that

$$\psi^{2}(\boldsymbol{W}^{t}) = \sum_{j=1}^{\kappa} \left\| \boldsymbol{w}_{j}^{t} \right\|_{2}^{2} \leq \psi^{2}(\boldsymbol{W}^{0}) + t \left(\eta^{2} \| \boldsymbol{v} \|_{2}^{2} + 2\eta \right).$$
(45)

We have also proved in Theorem 1 that Algorithm 1 performs at most T_k non-zero updates regardless of γ^2 (cf. (34)). Hence, as long as the initialization W^0 is bounded, all iterates W^t will be bounded; that is, there exists some constant $w_{\max} > 0$ such that $\|\boldsymbol{w}_j^t\|_2 \le w_{\max}$ holds for all $j = 1, 2, \ldots, k$ and $t = 0, 1, \ldots, T_k$.

For notational brevity, we drop the iteration index t, and let the current iterate be denoted by W. If there is no non-zero update

with the current sampled data point $(\boldsymbol{x}_i, y_i) \in S$, then one of the following two cases must be true: c1) $\mathbb{1}_{\{1-y_i\boldsymbol{v}^{\top}\sigma(\boldsymbol{W}\boldsymbol{x}_i)>0\}} = 0$, or equivalently, $\max(0, 1 - y_i\boldsymbol{v}^{\top}\sigma(\boldsymbol{W}\boldsymbol{x}_i)) = 0$ implying that (\boldsymbol{x}_i, y_i) is correctly classified; and, c2) $\mathbb{1}_{\{\boldsymbol{W}\boldsymbol{x}_i + \boldsymbol{\epsilon} \ge 0\}} = \mathbf{0}$, or equivalently, $\epsilon_j < -\boldsymbol{w}_j^{\top}\boldsymbol{x}_i$ for all neurons j = 1, 2, ..., k.

When *i* cycles through $\{1, 2, ..., n\}$ in a deterministic manner (with each integer drawn exactly once every *n* iterations), then within every succession of *np* iterations, the noise-injected stochastic gradient term $\mathbb{1}_{\{1-y_iv^{\top}\sigma(Wx_i)>0\}}$. $y_iv \operatorname{diag}(\mathbb{1}_{\{Wx_i+\epsilon\geq 0\}})x_i^{\top}$ in (10) will be evaluated exactly *p* times at every datum $(x_i, y_i) \in S$, but with *p* different random noise realizations. Hence, if there is no non-zero update within a succession of *np* iterations, the probability of event c2) occurring *np* times is at most

$$\left[\prod_{j\in\min\left\{\mathcal{N}_{v}^{+},\mathcal{N}_{v}^{-}\right\}}\Phi\left(\frac{\max_{1\leq i\leq n}-\boldsymbol{w}_{j}^{\top}\boldsymbol{x}_{i}}{\gamma}\right)\right]^{p} \leq \left[\Phi\left(\frac{w_{\max}}{\gamma}\right)\right]^{p\min\left\{\left|\mathcal{N}_{v}^{+}\right|,\left|\mathcal{N}_{v}^{-}\right|\right\}}$$
(46)

when there is only one out of the *n* data points that is left incorrectly classified. Here, by a slight abuse of notation, we use $j \in \min\{\mathcal{N}_v^+, \mathcal{N}_v^-\}$ to mean $j \in \mathcal{N}_v^+$ if $|\mathcal{N}_v^+| \leq |\mathcal{N}_v^-|$, and $j \in \mathcal{N}_v^-$ otherwise. Furthermore, to obtain the inequality in (46), we have used $\max_{1 \leq i \leq n} (-\boldsymbol{w}_j^\top \boldsymbol{x}_i) \leq \|\boldsymbol{w}_j\|_2 \|\boldsymbol{x}_i\|_2 \leq w_{\max}$ under our assumptions that $\|\boldsymbol{w}_j\|_2 \leq w_{\max}$ and $\|\boldsymbol{x}_i\|_2 \leq 1$ for all $j = 1, 2, \ldots, k$ and $i = 1, 2, \ldots, n$. Therefore, by the total probability theorem, the probability of having c1) hold for all data points is at least

$$1 - \left[\Phi\left(\frac{w_{\max}}{\gamma}\right)\right]^{p\min\left\{\left|\mathcal{N}_{v}^{+}\right|,\left|\mathcal{N}_{v}^{-}\right|\right\}}$$
(47)

which can be made arbitrarily close to 1 by taking either a large enough p and/or $\gamma > 0$. The case of picking i_t uniformly at random from $\{1, 2, ..., n\}$ can be discussed in a similar fashion, but it is omitted here. This completes the proof.

REFERENCES

- G. An, "The effects of adding noise during backpropagation training on a generalization performance," *Neural Comput.*, vol. 8, no. 3, pp. 643–674, Apr. 1996.
- [2] P. Auer, M. Herbster, and M. K. Warmuth, "Exponentially many local minima for single neurons," in *Adv. Neural Inf. Process. Syst.*, Denver, CO, USA, Nov. 27–Dec. 2, 1995, pp. 316–322.
- [3] A. Blum and R. L. Rivest, "Training a 3-node neural network is NPcomplete," in Adv. Neural Inf. Process. Syst., Cambridge, MA, USA, Aug. 3–5, 1988, pp. 494–501.
- [4] A. Brutzkus, A. Globerson, E. Malach, and S. Shalev-Shwartz, "SGD learns over-parameterized networks that provably generalize on linearly separable data," in *Proc. Int. Conf. Learn. Rep.*, Vancouver, BC, Canada, Apr. 30–May 3, 2018, pp. 1–9.
- [5] F. H. Clarke, Optimization and Nonsmooth Analysis. New York, NY, USA: Wiley, 1990, vol. 5.
- [6] D. Dheeru and E. Karra Taniskidou, "UCI machine learning repository," 2017. [Online]. Available: http://archive.ics.uci.edu/ml
- [7] G. K. Dziugaite and D. M. Roy, "Computing nonvacuous generalization bounds for deep (stochastic) neural networks with many more parameters than training data," 2007, arXiv:1703.11008.

- [8] R. Ge, F. Huang, C. Jin, and Y. Yuan, "Escaping from saddle points— Online stochastic gradient for tensor decomposition," in *Proc. Conf. Learn. Theory*, vol. 40, Paris, France, Jul. 3–6, 2015, pp. 797–842.
- [9] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. Cambridge: MIT Press, 2016, vol. 1.
- [10] M. Gori and A. Tesi, "On the problem of local minima in backpropagation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 14, no. 1, pp. 76–86, Jan. 1992.
- [11] L. Holmstrom and P. Koistinen, "Using additive noise in back-propagation training," *IEEE Trans. Neural Netw.*, vol. 3, no. 1, pp. 24–38, Jan. 1992.
- [12] S. M. M. Kalan, M. Soltanolkotabi, and A. S. Avestimehr, "Fitting ReLUs via SGD and Quantized SGD," 2019, arXiv:1901.06587.
- [13] K. Kawaguchi, "Deep learning without poor local minima," in Adv. Neural Inf. Process. Syst., Barcelona, Spain, Dec. 5–10, 2016, pp. 586–594.
- [14] K. Kawaguchi and L. P. Kaelbling, "Elimination of all bad local minima in deep learning," 2019, arXiv:1901.00279.
- [15] K. Kawaguchi, L. P. Kaelbling, and Y. Bengio, "Generalization in deep learning," 2017, arXiv:1710.05468.
- [16] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," in *Adv. Neural Inf. Process. Syst.*, Lake Tahoe, Nevada, Dec. 3–6, 2012, pp. 1097–1105.
- [17] T. Laurent and J. von Brecht, "Deep linear neural networks with arbitrary loss: All local minima are global," in *Int. Conf. Mach. Learn.*, Stockholm, Sweden, Jul. 10–15, 2018, pp. 2908–2913.
- [18] T. Laurent and J. von Brecht, "The multilinear structure of ReLU networks," in *Proc. Int. Conf. Mach. Learn.*, vol. 80, Stockholm, Sweden, Jul. 10–15, 2018, pp. 2914–2922.
- [19] D. Li, T. Ding, and R. Sun, "Over-parameterized deep neural networks have no strict local minima for any continuous activations," 2018, arXiv:1812.11039.
- [20] S. Liang, R. Sun, J. D. Lee, and R. Srikant, "Adding one neuron can eliminate all bad local minima," 2018, arXiv:1805.08671.
- [21] N. Littlestone and M. Warmuth, "Relating data compression and learnability," Dept. Comput. Inf. Sci., Univ. California, Santa Cruz, CA, USA, 1986.
- [22] V. Nair and G. E. Hinton, "Rectified linear units improve restricted Boltzmann machines," in *Proc. Int. Conf. Mach. Learn.*, Haifa, Israel, Jun. 21– 24, 2010, pp. 807–814.
- [23] B. Neyshabur, R. Tomioka, and N. Srebro, "In search of the real inductive bias: On the role of implicit regularization in deep learning," 2014, arXiv:1412.6614.
- [24] Q. Nguyen, "On connected sublevel sets in deep learning," 2019, arXiv:1901.07417.
- [25] Q. Nguyen and M. Hein, "Optimization landscape and expressivity of deep CNNs," in *Int. Conf. Mach. Learn.*, 2018, pp. 3727–3736.
- [26] A. B. Novikoff, "On convergence proofs for perceptrons," in *Proc. Symp. Math. Theory Automata*, vol. 12, 1963, pp. 615–622.
- [27] S. Oymak, "Stochastic gradient descent learns state equations with nonlinear activations," 2018, arXiv:1809.03019.
- [28] S. Oymak and M. Soltanolkotabi, "Towards moderate overparameterization: Global convergence guarantees for training shallow neural networks," 2019, arXiv:1902.04674.
- [29] B. Poole, S. Lahiri, M. Raghu, J. Sohl-Dickstein, and S. Ganguli, "Exponential expressivity in deep neural networks through transient chaos," in *Adv. Neural Inf. Process. Syst.*, Barcelona, Spain, Dec. 5–10, 2016, pp. 3360–3368.
- [30] M. Raghu, B. Poole, J. Kleinberg, S. Ganguli, and J. Sohl-Dickstein, "On the expressive power of deep neural networks," in *Proc. Int. Conf. Mach. Learn.*, vol. 70, Sydney, Australia, Aug. 6–11, 2017, pp. 2847–2854.
- [31] F. Rosenblatt, "The perceptron: A probabilistic model for information storage and organization in the brain," *Psychol. Rev.*, vol. 65, no. 6, pp. 386– 408, Nov. 1958.
- [32] S. Shalev-Shwartz and S. Ben-David, Understanding Machine Learning: From Theory to Algorithms. New York, NY, USA: Cambridge Univ. Press, 2014.
- [33] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," 2014, arXiv:1409.1556.
- [34] M. Soltanolkotabi, "Learning ReLU via gradient descent," in Adv. Neural Inf. Process. Syst., Long Beach, CA, USA, Dec. 4–9, 2017, pp. 2007–2017.
- [35] C. Wang and J. C. Principe, "Training neural networks with additive noise in the desired signal," *IEEE Trans. Neural Netw.*, vol. 10, no. 6, pp. 1511– 1517, Nov. 1999.
- [36] G. Wang, G. B. Giannakis, and Y. C. Eldar, "Solving systems of random quadratic equations via truncated amplitude flow," *IEEE Trans. Inf. The*ory, vol. 64, no. 2, pp. 773–794, Feb. 2018.

- [37] T. Xu, Y. Zhou, K. Ji, and Y. Liang, "Convergence of SGD in learning ReLU models with separable data," 2018, arXiv:1806.04339.
- [38] L. Zhang, G. Wang, and G. B. Giannakis, "Real-time power system state estimation and forecasting via deep neural networks," 2018, arXiv:1811.06146.
- [39] K. Zhong, Z. Song, P. Jain, P. L. Bartlett, and I. S. Dhillon, "Recovery guarantees for one-hidden-layer neural networks," in *Proc. Int. Conf. Mach. Learn.*, vol. 70, Sydney, Australia, Aug. 6–11, 2017, pp. 4140–4149.
- [40] Y. Zhou and Y. Liang, "Critical points of linear neural networks: Analytical forms and landscape properties," in *Proc. Int. Conf. Learn. Rep.*, Vancouver, BC, Canada, Apr. 30–May 3, 2018, pp. 1–9.
- [41] S. S. Du, J. D. Lee, Y. Tian, B. Poczos, and A. Singh, "Gradient descent learns one-hidden-layer CNN: Don't be afraid of spurious local minima," in *Proc. Int. Conf. Mach. Learn.*, vol. 80, Stockholm, Sweden, pp. 1338– 1347, Jul. 10–15, 2018.



Gang Wang (M'18) received the B.Eng. degree in electrical engineering and automation from the Beijing Institute of Technology, Beijing, China, in 2011, and the Ph.D. degree in electrical and computer engineering from the University of Minnesota, Minneapolis, MN, USA, in 2018. He is currently a Postdoctoral Associate with the Department of Electrical and Computer Engineering, University of Minnesota. His research interests focus on the areas of statis-

tical learning, optimization, and deep learning with applications to data science and smart grids. He re-

ceived a National Scholarship (2014), a Guo Rui Scholarship (2017), and an Innovation Scholarship (first place in 2017), all from China, as well as a Best Student Paper Award at the 2017 European Signal Processing Conference.



Georgios B. Giannakis (F'97) received the diploma degree in electrical engineering from the National Technical University of Athens, Athens, Greece, 1981. From 1982 to 1986 he was with the University of Southern California (USC), Los Angeles, CA, USA, where he received the M.Sc. degree in electrical engineering in 1983, the M.Sc. degree in electrical engineering in 1986, and the Ph.D. degree in electrical engineering in 1986. He was a Faculty Member with the University of Virginia, Charlottesville, VI, USA, from 1987 to 1998, and since 1999 he has been a Professor with

the University of Minnesota, Minneapolis, MN, USA, where he holds an ADC Endowed Chair, a University of Minnesota McKnight Presidential Chair in ECE, and serves as the Director of the Digital Technology Center.

His general interests span the areas of statistical learning, communications, and networking—subjects on which he has published more than 450 journal papers, 750 conference papers, 25 book chapters, 2 edited books, and 2 research monographs (h-index 140). Current research focuses on data science, and network science with applications to the Internet of Things, social, brain, and power networks with renewables. He is the (co-) inventor of 32 patents issued and the (co-) recipient of 9 best journal paper awards from the IEEE Signal Processing (SP) and Communications. He also received Technical Achievement Awards from the SP Society (2000), from EURASIP (2005), a Young Faculty Teaching Award, the G. W. Taylor Award for Distinguished Research from the University of Minnesota, and the IEEE Fourier Technical Field Award (inaugur al recipient in 2015). He is a Fellow of EURASIP and has served the IEEE SiPs.



Jie Chen (F'19) received the B.Sc., M.Sc., and Ph.D. degrees in control theory and control engineering from the Beijing Institute of Technology, Beijing, China, in 1986, 1996, and 2001, respectively. From 1989 to 1990, he was a visiting scholar at the California State University, Long Beach, CA, USA. From 1996 to 1997, he was a Research Fellow with the School of Engineering, University of Birmingham, Birmingham, U.K. Since 1986, he has been with the Beijing Institute of Technology, where he is now a Professor in control science and engineering, and the

Head of the State Key Laboratory of Intelligent Control and Decision of Complex Systems, China. He currently serves as the President of Tongji University, Shanghai, China.

His main research interests include intelligent control and decision in complex systems, multi-agent systems, nonlinear control, and optimization. He has (co-)authored 4 books and more than 200 research papers. He served as a Managing Editor for the *Journal of Systems Science and Complexity*, and an Associate/Subject Editor for several other international journals, including the IEEE TRANSACTIONS ON CYBERNETICS and the *International Journal of Robust and Nonlinear Control*. He was the recipient of several prestigious awards, including one second-grade prize of the National Natural Science and Technology Progress Award of China (2009 and 2011). He is a member of the Chinese Academy of Engineering.