
Elimination of All Bad Local Minima in Deep Learning

Kenji Kawaguchi
MIT

Leslie Pack Kaelbling
MIT

Abstract

In this paper, we theoretically prove that adding one special neuron per output unit eliminates all suboptimal local minima of any deep neural network, for multi-class classification, binary classification, and regression with an arbitrary loss function, under practical assumptions. At every local minimum of any deep neural network with these added neurons, the set of parameters of the original neural network (without added neurons) is guaranteed to be a global minimum of the original neural network. The effects of the added neurons are proven to automatically vanish at every local minimum. Moreover, we provide a novel theoretical characterization of a failure mode of eliminating suboptimal local minima via an additional theorem and several examples. This paper also introduces a novel proof technique based on the perturbable gradient basis (PGB) necessary condition of local minima, which provides new insight into the elimination of local minima and is applicable to analyze various models and transformations of objective functions beyond the elimination of local minima.

1 Introduction

Deep neural networks have achieved significant practical success in the fields of computer vision, machine learning, and artificial intelligence. However, theoretical understanding of deep neural networks is scarce relative to its empirical success. One of the major difficulties in theoretically understanding deep neural networks lies in the non-convexity and high-dimensionality of the objective functions used to train

the networks. Because of the non-convexity and high-dimensionality, it is often unclear whether a deep neural network will be guaranteed to have a desired property after training, instead of becoming stuck around an arbitrarily poor local minimum. Indeed, it is NP-hard to find a global minimum of a general non-convex function (Murty and Kabadi, 1987), and of non-convex objective functions used to train certain types of neural networks (Blum and Rivest, 1992). In the past, such theoretical concerns were considered one of reasons to prefer classical machine learning models (with or without a kernel approach) that require only convex optimization. Given their recent empirical success, a question remains whether practical deep neural networks can be theoretically guaranteed to avoid poor local minima.

There have been numerous recent studies that have advanced theoretical understanding regarding the optimization of deep neural networks with significant over-parameterization (Nguyen and Hein, 2017, 2018; Allen-Zhu et al., 2018; Du et al., 2018; Zou et al., 2018) and model simplification (Choromanska et al., 2015; Kawaguchi, 2016; Hardt and Ma, 2017; Bartlett et al., 2019; Du and Hu, 2019). For shallow networks with a single hidden layer, there have been many positive results, yet often with strong assumptions, for example, requiring the use of significant over-parameterization, simplification, or Gaussian inputs (Andoni et al., 2014; Sedghi and Anandkumar, 2014; Soltanolkotabi, 2017; Brutzkus and Globerson, 2017; Ge et al., 2017; Soudry and Hoffer, 2017; Goel and Klivans, 2017; Zhong et al., 2017; Li and Yuan, 2017; Du and Lee, 2018).

Instead of using these strong assumptions, adding one neuron to a neural network with a single output unit was recently shown to be capable of eliminating all suboptimal local minima (i.e., all local minima that are not global minima) for binary classification with a special type of smoothed hinge loss function (Liang et al., 2018). However, the restriction to a neural network with a single output unit and a special loss function for binary classification makes it inapplicable to many common and important deep learning problems. Both removing the restriction to networks with a sin-

gle output unit and removing restrictions on loss functions are seen as important open problems in different but related theoretical work on local minima of neural networks (Shamir, 2018; Laurent and Brecht, 2018).

In this paper, we state and prove a novel and significantly stronger theorem: adding one neuron per output unit can eliminate all suboptimal local minima of any deep neural network with an arbitrary loss function for multi-class classification, binary classification, and regression. This paper also introduces a novel proof technique based on the perturbable gradient basis (PGB) condition, which provides new insight into the elimination of local minima and can be used to analyze new models and transformations of objective functions beyond the elimination of local minima. This paper analyzes the problem in the regime without significant over-parameterization or model simplification (except adding the few extra neurons).

While analyzing the properties of local minima in this regime with no strong assumption is a potentially important step in theory, it does not immediately guarantee the efficient solution of general neural network optimization. We explain this phenomenon in another key contribution of this paper, which is a novel characterization of a failure mode of the removal of bad local optima, in terms of its effect on gradient-based optimization methods.

2 Elimination of local minima

The optimization problem for the elimination of local minima is defined in Section 2.1. Our theoretical results on the elimination of local minima are presented in Section 2.2 for arbitrary datasets, and in Section 2.3 for realizable datasets. We discuss these results in terms of non-vacuousness, consistency, and the effect of over-parameterization in Section 2.4.

2.1 Problem description

Let $x \in \mathcal{X} \subseteq \mathbb{R}^{d_x}$ and $y \in \mathcal{Y} \subseteq \mathbb{R}^{d_y}$ be an input vector and a target vector, respectively. Define $((x_i, y_i))_{i=1}^m$ as a training dataset of size m . Given an input x and parameter θ , let $f(x; \theta) \in \mathbb{R}^{d_y}$ be the pre-activation output of the last layer of any arbitrary deep neural network with any structure (e.g., any convolutional neural network with any depth and any width, with or without skip connections). That is, there is no assumption with regard to f except that $f(x; \theta) \in \mathbb{R}^{d_y}$.

We consider the following standard objective function L to train an arbitrary neural network f :

$$L(\theta) = \frac{1}{m} \sum_{i=1}^m \ell(f(x_i; \theta), y_i),$$

where $\ell : \mathbb{R}^{d_y} \times \mathcal{Y} \rightarrow \mathbb{R}$ is an arbitrary loss criterion such as cross entropy loss, smoothed hinge loss, or squared loss.

We then define an auxiliary objective function \tilde{L} :

$$\tilde{L}(\tilde{\theta}) = \frac{1}{m} \sum_{i=1}^m \ell(f(x_i; \theta) + g(x_i; a, b, W), y_i) + \lambda \|a\|_2^2,$$

where $\lambda > 0$, $\tilde{\theta} = (\theta, a, b, W)$, $a, b \in \mathbb{R}^{d_y}$, $W = [w_1 \ w_2 \ \dots \ w_{d_y}] \in \mathbb{R}^{d_x \times d_y}$ with $w_k \in \mathbb{R}^{d_x}$, and $g(x; a, b, W)_k = a_k \exp(w_k^\top x + b_k)$ for all $k \in \{1, \dots, d_y\}$.

We also define a modified neural network \tilde{f} as

$$\tilde{f}(x; \tilde{\theta}) = f(x; \theta) + g(x; a, b, W),$$

which is equivalent to adding one neuron $g(x; a, b, W)_k$ per each output unit $f(x; \theta)_k$ of the original neural network. Because $\tilde{L}(\tilde{\theta}) = \frac{1}{m} \sum_{i=1}^m \ell(\tilde{f}(x_i; \tilde{\theta}), y_i) + \lambda \|a\|_2^2$, the auxiliary objective function \tilde{L} is the standard objective function L with the modified neural network \tilde{f} and a regularizer on a .

2.2 Result for arbitrary datasets

Under only a mild assumption (Assumption 1), Theorem 1 states that at every local minimum (θ, a, b, W) of the modified objective function \tilde{L} , the parameter vector θ achieves a global minimum of the original objective function L , and the modified neural network \tilde{f} automatically becomes the original neural network f . The proof of Theorem 1 is presented in Section 4 and Appendix B.

Assumption 1. (Use of common loss criteria) *For any $i \in \{1, \dots, m\}$, the function $\ell_{y_i} : q \mapsto \ell(q, y_i) \in \mathbb{R}_{\geq 0}$ is differentiable and convex (e.g., the squared loss, cross entropy loss, and polynomial hinge loss satisfy this assumption).*

Theorem 1. *Let Assumption 1 hold. Then, at every local minimum (θ, a, b, W) of \tilde{L} , the following statements hold:*

- (i) θ is a global minimum of L , and
- (ii) $\tilde{f}(x; \theta, a, b, W) = f(x; \theta)$ for all $x \in \mathbb{R}^{d_x}$, and $\tilde{L}(\theta, a, b, W) = L(\theta)$.

Assumption 1 is satisfied by simply using a common loss criterion, including the squared loss $\ell(q, y) = \|q - y\|_2^2$ or $\ell(q, y) = (1 - yq)^2$ (the latter with $d_y = 1$), cross entropy loss $\ell(q, y) = -\sum_{k=1}^{d_y} y_k \log \frac{\exp(q_k)}{\sum_{k'} \exp(q_{k'})}$, or smoothed hinge loss $\ell(q, y) = (\max\{0, 1 - yq\})^p$ with $p \geq 2$ (the hinge loss with $d_y = 1$). Although the objective function $L : \theta \mapsto L(\theta)$ used to train a neural network is non-convex in θ , the loss criterion $\ell_{y_i} : q \mapsto \ell(q, y_i)$ is usually convex in q . Therefore, Theorem 1 is directly applicable to most common deep learning tasks in practice.

2.3 Result for realizable datasets

Theorem 2 makes a statement similar to Theorem 1 under a weaker assumption on the loss criterion (Assumption 2) but with an additional assumption on the training dataset (Assumption 3). The proof of Theorem 2 is presented in Appendix B.

Assumption 2. (On the loss) *For any $i \in \{1, \dots, m\}$, the function $\ell_{y_i} : q \mapsto \ell(q, y_i)$ is differentiable, and $q \in \mathbb{R}^{d_y}$ is a global minimum of ℓ_{y_i} if $\nabla \ell_{y_i}(q) = 0$.*

Assumption 3. (On the label consistency) *There exists a function f^* such that $f^*(x_i) = y_i$ for all $i \in \{1, \dots, m\}$.*

Theorem 2. *Let Assumptions 2 and 3 hold. Then, at every local minimum (θ, a, b, W) of \tilde{L} , the following statements hold:*

- (i) θ is a global minimum of L ,
- (ii) $\tilde{f}(x; \theta, a, b, W) = f(x; \theta)$ for all $x \in \mathbb{R}^{d_x}$, and $\tilde{L}(\theta, a, b, W) = L(\theta)$, and
- (iii) $f(x_i; \theta)$ is a global minimum of $\ell_{y_i} : q \mapsto \ell(q, y_i)$ for all $i \in \{1, \dots, m\}$.

Assumption 2 is weaker than Assumption 1 in the sense that the former is implied by the latter but not vice versa. However, as discussed above, Assumption 1 already accommodates most common loss criteria. Assumption 3 is automatically satisfied if a target y given an input x is not random, but the non-randomness is not necessary to satisfy Assumption 3. Even if the targets are generated at random, as long as all x_1, x_2, \dots, x_m are distinct (i.e., $x_i \neq x_j$ for all $i \neq j$), Assumption 3 is satisfied.

Therefore, although Theorem 2 might be less applicable in practice when compared to Theorem 1, Theorem 2 can still be applied to many common deep learning tasks with the additional guarantee, as stated in Theorem 2 (iii). By using an appropriate loss criterion for classification, Theorem 2 (iii) implies that the trained neural network $f(\cdot; \theta)$ at every local minimum correctly classifies all training data points.

2.4 Non-vacuousness, consistency, and effect of over-parameterization

Theorems 1 and 2 are both non-vacuous and consistent with pathological cases. For the consistency, Theorems 1 and 2 (vacuously) hold true if there is no local minimum of \tilde{L} , for example, with a pathological case of $\ell(q, y_i) = q - y_i$.

For non-vacuousness, there exists a local minimum of \tilde{L} if there exists a global minimum θ of L such that $f(x_i; \theta)$ achieves a global minimum for each $f(x_i; \theta) \mapsto \ell(f(x_i; \theta), y_i)$ for $i \in \{1, \dots, m\}$ (this is because, given

such a θ , any point with $a = 0$ is a local minimum of \tilde{L}). Therefore, the existence of local minimum for \tilde{L} can be guaranteed by the *weak* degree of over-parameterization that ensures the existence of a global minimum θ for each $f(x_i; \theta) \mapsto \ell(f(x_i; \theta), y_i)$ *only for a given training dataset* (rather than for all datasets).

This is in contrast to the previous papers that require *significant* over-parameterization to ensure interpolation of all datasets and to make the corresponding neural tangent kernel to be approximately unchanged during training (Nguyen and Hein, 2017, 2018; Allen-Zhu et al., 2018; Du et al., 2018; Zou et al., 2018). Our paper does not require those and allow the neural tangent kernel to significantly change during training. Because of this difference, our paper only needs $\tilde{\Omega}(1)$ parameters, whereas the state-of-the-art previous paper requires $\tilde{\Omega}(H^{12}n^8)$ parameters (Zou and Gu, 2019) or $\tilde{\Omega}(n)$ parameters (Kawaguchi and Huang, 2019), .

Because a local minimum does not need to be a *strict* local minimum (i.e., a local minimum with a strictly less value than others in a neighborhood), there are many other cases where there exists a local minimum of \tilde{L} : e.g., Example 5 in Section 6.2 also illustrates a situation where there exists a local minimum of \tilde{L} without the above condition of the existence of sample-wise global minimum of L or weak over-parameterization.

3 PGB necessary condition beyond elimination of local minima

In this section, we introduce a more general result beyond elimination of local minima. Namely, we prove the *perturbable gradient basis (PGB) necessary condition of local minima*, which directly implies the elimination result as a special case. Beyond the specific transformation of the objective function \tilde{L} , the PGB necessary condition of local minima can be applied to other objective functions with various transformations and models.

Theorem 3. (PGB necessary condition of local minima). *Define the objective function Q by*

$$Q(z) = \frac{1}{m} \sum_{i=1}^m Q_i(\phi_i(z)) + R_i(\varphi_i(z)) \quad (1)$$

where for all $i \in \{1, \dots, m\}$, the functions $Q_i : q \in \mathbb{R}^{d_\phi} \mapsto Q_i(q) \in \mathbb{R}_{\geq 0}$ and $R_i : q \in \mathbb{R}^{d_\phi} \mapsto R_i(q) \in \mathbb{R}_{\geq 0}$ are differentiable and convex, and ϕ_i and φ_i are differentiable. Assume that there exists a function $h : \mathbb{R}^{d_z} \rightarrow \mathbb{R}^{d_z}$ and a real number $\rho \neq 0$ such that for all $i \in \{1, \dots, m\}$ and all $z \in \mathbb{R}^{d_z}$, $\phi_i(z) = \sum_{k=1}^{d_z} h(z)_k \partial_k \phi_i(z)$ and $\varphi_i(z) = \rho \sum_{k=1}^{d_z} h(z)_k \partial_k \varphi_i(z)$. Then, for any local minimum

$z \in \mathbb{R}^{d_z}$ of Q , the following holds: there exists $\epsilon_0 > 0$ such that for any $\epsilon \in [0, \epsilon_0)$,

$$Q(z) \leq \inf_{\substack{S \subseteq_{\text{fin}} \mathcal{V}[z, \epsilon], \\ \alpha \in \mathbb{R}^{d_z \times |S|}}} \tilde{Q}_{\epsilon, z}(\alpha, S) + \frac{\rho - 1}{\rho m} \sum_{i=1}^m \partial R_i(\varphi_i(z)) \varphi_i(z),$$

where

$$\tilde{Q}_{\epsilon, z}(\alpha, S) = \frac{1}{m} \sum_{i=1}^m Q_i(\phi_i^z(\alpha, \epsilon, S)) + R_i(\varphi_i^z(\alpha, \epsilon, S)),$$

$$\phi_i^z(\alpha, \epsilon, S) = \sum_{k=1}^{d_z} \sum_{j=1}^{|S|} \alpha_{k,j} \partial_k \phi_i(z + \epsilon S_j),$$

and

$$\varphi_i^z(\alpha, \epsilon, S) = \sum_{k=1}^{d_z} \sum_{j=1}^{|S|} \alpha_{k,j} \partial_k \varphi_i(z + \epsilon S_j).$$

Here, $S \subseteq_{\text{fin}} S'$ denotes a finite subset S of a set S' and $\mathcal{V}[z, \epsilon]$ is the set of all vectors $v \in \mathbb{R}^{d_z}$ such that $\|v\|_2 \leq 1$, $\phi_i(z + \epsilon v) = \phi_i(z)$, and $\varphi_i(z + \epsilon v) = \varphi_i(z)$ for all $i \in \{1, \dots, m\}$. Furthermore, if $\rho = 1$, this statement holds with equality as $Q(z) = \inf_{S \subseteq_{\text{fin}} \mathcal{V}[z, \epsilon], \alpha \in \mathbb{R}^{d_z \times |S|}} \tilde{Q}_{\epsilon, z}(\alpha, S)$.

The PGB necessary condition of local minima states that whether a given objective Q is non-convex or convex, if z is a *local* minimum of the original (potentially non-convex) objective Q , then z must achieve the *global* minimum value of the transformed objective $\tilde{Q}_{\epsilon, z}(\alpha, S)$ for any sufficiently small ϵ . Here, the transformed objective $\tilde{Q}_{\epsilon, z}$ is the original objective Q except that the original functions $\phi_i(z)$ and $\varphi_i(z)$ are replaced by perturbable gradient basis (PGB) functions $\phi_i^z(\alpha, \epsilon, S)$ and $\varphi_i^z(\alpha, \epsilon, S)$. In other words, all local minima with the original functions $\phi_i(z)$ and $\varphi_i(z)$ achieve the global minimum values of PGB functions $\phi_i^z(\alpha, \epsilon, S)$ and $\varphi_i^z(\alpha, \epsilon, S)$.

Here, the original objective Q is a non-convex function in general because both ϕ_i and φ_i can be non-convex functions. If ϕ_i and φ_i are both linear functions, then we have $\phi_i(\alpha) = \phi_i^z(\alpha, 0, \emptyset)$ and $\varphi_i(\alpha) = \varphi_i^z(\alpha, 0, \emptyset)$ (for all α and z) and hence the PGB necessary condition of local minima recovers the following known statement: every local minimum of original Q is a global minimum of original Q .

The PGB condition of local minima can be also understood based on the following geometric interpretation. For the geometric interpretation, we consider two spaces – the parameter space \mathbb{R}^{d_z} and the output space $\mathbb{R}^{m(d_\phi + d_\varphi)}$ – and the map from the parameter space to the output space, which is defined by $\Phi : z \in \mathbb{R}^{d_z} \mapsto (\phi_1(z)^\top, \varphi_1(z)^\top, \dots, \phi_m(z)^\top, \varphi_m(z)^\top)^\top \in$

$\mathbb{R}^{m(d_\phi + d_\varphi)}$. Then, in the output space, we can intuitively consider the “tangent” space $T_{\Phi(z)} = \text{span}(\{\partial_1 \Phi(z), \dots, \partial_{d_z} \Phi(z)\} + \{\Phi(z)\})$, where the sum of the two sets represents the Minkowski sum of the sets.¹ Then, given a $\epsilon (\leq \epsilon_0)$, the span of the set of all vectors of the “tangent” spaces $T_{\Phi(z + \epsilon v)}$ at all perturbed points $z + \epsilon v$, defined by $\tilde{T}_{\Phi(z + \epsilon v)} = \text{span}(\{\mathbf{f} \in \mathbb{R}^{m(d_\phi + d_\varphi)} : (\exists v \in \mathcal{V}[\theta, \epsilon])[\mathbf{f} \in T_{\Phi(z + \epsilon v)}]\})$, is exactly equal to the space of the outputs of the PGB functions.

Therefore, from the geometric viewpoint, the PGB necessary condition of local minima states that the output $\Phi(z)$ at any local minimum z is globally optimal in the span of the “tangent” spaces $\tilde{T}_{\Phi(z)}$. The PGB condition of local minima translates the *local* optimality in the *parameter* space \mathbb{R}^{d_z} into the *global* optimality in the span of the “tangent” spaces in the *output space* $\mathbb{R}^{m(d_\phi + d_\varphi)}$.

The PGB necessary condition of local minima is an extension of theorem 2 in a previous study (Kawaguchi et al., 2019) to the problem with a regularization term and a general transformation of the objective function. Accordingly, beyond the elimination of local minima, the PGB necessary condition of local minima entails the previously proven statements of no bad local minima for deep linear neural networks (Laurent and Brecht, 2018) and deep nonlinear residual neural networks (Kawaguchi and Bengio, 2019) (since the PGB necessary condition of local minima is strictly more general than theorem 2 by Kawaguchi et al. 2019, which was shown to entail those statements).

With our extension, the PGB necessary condition of local minima can be now used to study the effects of various transformations of objective functions and regularization terms, including the elimination of local minima, as shown in the next section.

4 Application of PGB necessary condition for elimination of local minima

In this section, we introduce a novel and concise proof for Theorem 1 based on the PGB necessary condition of local minima, which shows that all suboptimal local minima can be eliminated because the global minimum value of the PGB model is indeed the global minimum value of the original model. From the geometric viewpoint, this is because the PGB model is shown to be expressive in that the span of the “tangent” spaces $T_{\Phi(z + \epsilon v)}$ contains the output space $\mathbb{R}^{m d_\phi} \times \{\mathbf{0}\}$ where

¹In special cases (e.g., when $\text{rank}(\partial \Phi(z))$ is constant in a neighborhood of z), $T_{\Phi(z)}$ is indeed a tangent space of a local manifold embedded in the output space. In general, $T_{\Phi(z)}$ and $\tilde{T}_{\Phi(z)}$ are the affine subspaces of the output space $\mathbb{R}^{m(d_\phi + d_\varphi)}$.

$\mathbf{0} \in \mathbb{R}^{m d_\varphi}$. In Appendix B, we also provide an alternative proof of Theorem 1 without the PGB necessary condition, which is intended to be more detailed with elementary facts but less elegant than the following proof via the PGB necessary condition.

Proof of Theorem 1. Let θ be fixed. Let (a, b, W) be a local minimum of $\tilde{L}|_\theta(a, b, W) := \tilde{L}(\theta, a, b, W)$. Let $\bar{\theta} \in \mathbb{R}^{d_\theta}$ be the vector containing (a, b, W) , defined by $\bar{\theta} = (a^\top, b^\top, \text{vec}(W)^\top)^\top$. We apply the PGB necessary condition of local minima by setting $Q(\bar{\theta}) = \tilde{L}|_\theta(a, b, W)$ with $\phi_i(\bar{\theta}) = g(x_i; a, b, W)$, $\varphi_i(\bar{\theta}) = (a_1^2, \dots, a_{d_y}^2)^\top$, $Q_i(q) = \ell(f(x_i; \theta) + q, y_i)$, and $R_i(q) = \lambda \sum_{j=1}^{d_y} q_j$. Then, for all $i \in \{1, \dots, m\}$, the functions Q_i and R_i are differentiable and convex, and the functions ϕ_i and φ_i are differentiable. Furthermore, we can rewrite $\phi_i(\bar{\theta}) = \sum_{k=1}^{d_y} a_k \partial_{a_k} \phi_i(\bar{\theta})$ and $\varphi_i(\bar{\theta}) = \rho \sum_{k=1}^{d_y} a_k \partial_{a_k} \varphi_i(\bar{\theta})$ with $\rho = 1/2$ for all $i \in \{1, \dots, m\}$ and all $\bar{\theta} \in \mathbb{R}^{d_\theta}$. These satisfy the assumptions of the PGB necessary condition of local minima of $Q(\bar{\theta}) = \tilde{L}|_\theta(a, b, W)$.

From the PGB necessary condition of local minima of $Q(\bar{\theta}) = \tilde{L}|_\theta(a, b, W)$, there exists $\epsilon_0 > 0$ such that for any $\epsilon \in [0, \epsilon_0)$,

$$\begin{aligned} Q(\bar{\theta}) - \frac{\rho-1}{\rho m} \sum_{i=1}^m \partial R_i(\varphi_i(\bar{\theta})) \varphi_i(\bar{\theta}) \\ \leq \inf_{\substack{S \subseteq_{\text{fin}} \mathcal{V}[\bar{\theta}, \epsilon], \\ \alpha \in \mathbb{R}^{d_z \times |S|}}} \tilde{Q}_{\epsilon, \bar{\theta}}(\alpha, S) \leq \inf_{\substack{S \subseteq_{\text{fin}} \mathcal{V}[\bar{\theta}, \epsilon], \\ \alpha \in \mathbb{R}^{d_z \times |S|}}} \tilde{Q}_{\epsilon, \bar{\theta}}(\alpha, S), \end{aligned} \quad (2)$$

where $\mathcal{V}[\bar{\theta}, \epsilon]$ is the set of all vectors $v \in \mathbb{R}^{d_\theta}$ such that $\|v\|_2 \leq 1$, $\phi_i(\bar{\theta} + \epsilon v) = \phi_i(\bar{\theta})$, and $\varphi_i(\bar{\theta} + \epsilon v) = \varphi_i(\bar{\theta})$ for all $i \in \{1, \dots, m\}$. Then, the subset $\tilde{\mathcal{V}}[\bar{\theta}, \epsilon] \subset \mathcal{V}[\bar{\theta}, \epsilon]$ is defined by $\tilde{\mathcal{V}}[\bar{\theta}, \epsilon] = \{(a^\top, b^\top, \text{vec}(W)^\top)^\top : (a^\top, b^\top, \text{vec}(W)^\top)^\top \in \mathcal{V}[\bar{\theta}, \epsilon], a = 0\}$.

Since (a, b, W) is a local minimum and hence the partial derivatives with respect to (a, b) are zeros, we have that for all $k \in \{1, 2, \dots, d_y\}$, $a_k \frac{\partial \tilde{L}(\theta, a, b, W)}{\partial a_k} = \frac{1}{m} \sum_{i=1}^m (\nabla \ell_{y_i}(f(x_i; \theta) + g(x_i; a, b, W)))_k a_k \exp(w_k^\top x + b_k) + 2\lambda a_k = \frac{\partial \tilde{L}(\theta, a, b, W)}{\partial b_k} + 2\lambda a_k^2 = 2\lambda a_k^2 = 0$, which implies that $a_k = 0$ for all $k \in \{1, 2, \dots, d_y\}$, since $2\lambda \neq 0$. Since $a = 0$, it proves statement (ii), and we have $L(\theta) = \tilde{L}|_\theta(a, b, W)$, $(\partial_{a_k} \varphi_i(\bar{\theta} + \epsilon v))_k = 2a_k = 0$ for all $v \in \tilde{\mathcal{V}}[\bar{\theta}, \epsilon]$, and $\frac{\rho-1}{\rho m} \sum_{i=1}^m \partial R_i(\varphi_i(\bar{\theta})) \varphi_i(\bar{\theta}) = 0$. Since $L(\theta) = \tilde{L}|_\theta(a, b, W) = Q(\bar{\theta}) = Q(\bar{\theta}) - \frac{\rho-1}{\rho m} \sum_{i=1}^m \partial R_i(\varphi_i(\bar{\theta})) \varphi_i(\bar{\theta})$ with (2) and $\partial_{a_k} \varphi_i(\bar{\theta} + \epsilon v) = 0$ (for all $v \in \tilde{\mathcal{V}}[\bar{\theta}, \epsilon]$), we have that for any θ' ,

$$\begin{aligned} L(\theta) - L(\theta') \\ \leq \inf_{\substack{S \subseteq_{\text{fin}} \mathcal{V}[\bar{\theta}, \epsilon], \\ \alpha \in \mathbb{R}^{d_z \times |S|}}} \tilde{Q}_{\epsilon, \bar{\theta}}(\alpha, S) - L(\theta') \end{aligned} \quad (3)$$

$$\begin{aligned} &= \inf_{\substack{S \subseteq_{\text{fin}} \tilde{\mathcal{V}}[\bar{\theta}, \epsilon], \\ \alpha \in \mathbb{R}^{d_\theta \times |S|}}} \frac{1}{m} \sum_{i=1}^m Q_i(\phi_i^{\bar{\theta}}(\alpha, \epsilon, S)) - \ell(f(x_i; \theta'), y_i) \\ &\leq 0, \end{aligned}$$

where the last inequality is to be shown to hold in the following.

Since $\phi_i^{\bar{\theta}}(\alpha, \epsilon, S)$ can differ for different indexes i only through different inputs x_i , we can rewrite $\phi_{x_i}^{\bar{\theta}}(\alpha, \epsilon, S) = \phi_i^{\bar{\theta}}(\alpha, \epsilon, S)$. We then rearrange the sum in the last line of (3):

$$\begin{aligned} &\frac{1}{m} \sum_{i=1}^m Q_i(\phi_{x_i}^{\bar{\theta}}(\alpha, \epsilon, S)) - \ell(f(x_i; \theta'), y_i) \\ &= \frac{1}{m} \sum_{j=1}^{m'} \sum_{i \in \mathcal{I}_j} Q_i(\phi_{\bar{x}_j}^{\bar{\theta}}(\alpha, \epsilon, S)) - \ell(f(\bar{x}_j; \theta'), y_i). \end{aligned}$$

where $\{\mathcal{I}_1, \dots, \mathcal{I}_{m'}\}$ is a partition² of the set $\{1, \dots, m\}$ such that for any $x \in \mathcal{I}_j$ and $x' \in \mathcal{I}_{j'}$, $x = x'$ if $j = j'$, and $x \neq x'$ if $j \neq j'$. Here, we write $\bar{x}_j := x$ with a representative $x \in \mathcal{I}_j$.

Let $S_t \in \mathbb{R}^{d_\theta}$ be the vector containing $(\hat{a}^{(t)}, \hat{b}^{(t)}, \hat{W}^{(t)})$. Since $a = 0$, we have $\phi_i(\bar{\theta} + \epsilon S_t) = \phi_i(\bar{\theta})$ for all vectors S_t containing any $(\hat{a}^{(t)}, \hat{b}^{(t)}, \hat{W}^{(t)})$ with $\hat{a}^{(t)} = 0$. In other words, for any finite $|S|$ and any $((\hat{b}^{(t)}, \hat{W}^{(t)}))_{t=1}^{|S|}$ with $\|S_t\|_2 \leq 1$, there exists $S \subseteq_{\text{fin}} \tilde{\mathcal{V}}[\bar{\theta}, \epsilon]$ such that $\phi_{\bar{x}_j}^{\bar{\theta}}(\alpha, \epsilon, S)_k = \exp(w_k^\top \bar{x}_j + b_k) \sum_{t=1}^{|S|} \alpha_k^{(t)} \exp(\epsilon((\hat{w}_k^{(t)})^\top \bar{x}_j + \hat{b}_k^{(t)}))$ (by letting $\hat{a}^{(t)} = 0$ for all t). Therefore, given m' distinct input points $\bar{x}_1, \dots, \bar{x}_{m'}$, fix $((\hat{b}^{(t)}, \hat{W}^{(t)}))_{t=1}^{|S|}$ such that the rank of the matrix $M \in \mathbb{R}^{m' \times |S|}$ with entries $M_{j,t} = \exp((\hat{w}_k^{(t)})^\top \bar{x}_j + \hat{b}_k^{(t)})$ is m' with a sufficiently large finite $|S|$. Then, by letting $M_{j,t}(\epsilon) \in \mathbb{R}^{m' \times |S|}$ be the matrix with entries $M_{j,t}(\epsilon) = \exp(\epsilon((\hat{w}_k^{(t)})^\top \bar{x}_j + \hat{b}_k^{(t)}))$, the function $\psi(\epsilon) = \det(M_{j,t}(\epsilon) M_{j,t}(\epsilon)^\top)$ is an analytic function of ϵ . Since ψ is analytic and $\psi(1) \neq 0$, either the zeros of ψ are isolated or $\psi(\epsilon) \neq 0$ for all ϵ . In both cases, there exists $\epsilon \in [0, \epsilon_0)$ such that $\psi(\epsilon) \neq 0$. Fix a $\epsilon \in [0, \epsilon_0)$ with $\psi(\epsilon) \neq 0$. Then, since $M_{j,t}(\epsilon)$ has rank m' , for any θ' , there exists α such that for all $j \in \{1, \dots, m'\}$,

$$\phi_{\bar{x}_j}^{\bar{\theta}}(\alpha, \epsilon, S) = f(\bar{x}_j; \theta') - f(\bar{x}_j; \theta).$$

Therefore, the last inequality in (3) holds, and hence for any θ' , $L(\theta) \leq L(\theta')$, which proves statement (i). \square

²That is, $\mathcal{I}_1 \cup \dots \cup \mathcal{I}_{m'} = \{1, \dots, m\}$, $\mathcal{I}_j \cap \mathcal{I}_{j'} = \emptyset$ for all $j \neq j'$, and $\mathcal{I}_j \neq \emptyset$ for all $j \in \{1, \dots, m'\}$.

5 Failure mode of eliminating suboptimal local minima

Our theoretical results in the previous sections have shown that, for a wide range of deep learning tasks, all suboptimal local minima can be removed by adding one neuron per output unit. This might be surprising given the fact that dealing with suboptimal local minima in general is known to be challenging in theory.

However, for the worst case scenario, the following theorem illuminates a novel failure mode for the elimination of suboptimal local minima. Theorem 4 holds true for the previous results (Liang et al., 2018) (as we discuss further in Section 7) and hence provides a novel failure mode for elimination of local minima in general, apart from our Theorems 1 and 2. Our result in this paper is the first result that points out this type of failure mode for elimination of local minima. The proof of Theorem 4 is presented in Appendix C.

Theorem 4. *Let Assumption 1 hold, or let Assumptions 2 and 3 hold. Then, for any θ , if θ is not a global minimum of L , there is no local minimum $(a, b, W) \in \mathbb{R}^{d_y} \times \mathbb{R}^{d_y} \times \mathbb{R}^{d_x \times d_y}$ of $\tilde{L}|_{\theta}(a, b, W) := \tilde{L}(\theta, a, b, W)$. Furthermore, there exists a tuple $(\ell, f, \{(x_i, y_i)\}_{i=1}^m)$ and a suboptimal stationary point θ of L such that $\frac{\partial \tilde{L}(\theta, a, b, W)}{\partial \theta} = 0$ for all $(a, b, W) \in \mathbb{R}^{d_y} \times \mathbb{R}^{d_y} \times \mathbb{R}^{d_x \times d_y}$.*

Therefore, on the one hand, Theorems 1 and 2 state that if an algorithm can find a local minimum of \tilde{L} , then it can find a global minimum of L (via a local minimum of \tilde{L}). On the other hand, Theorem 4 suggests that if an algorithm moves toward a local minimum of \tilde{L} by simply following (negative) gradient directions, then either it moves toward a global minimum θ of L or the norm of (a, b, W) approaches infinity. By monitoring the norm of (a, b, W) , we can detect this failure mode. This suggests a hybrid approach of local and global optimization algorithms beyond a pure local gradient-based method, with a mechanism to monitor the increase in the norm of (a, b, W) .

6 Examples

The previous sections show that one can eliminate suboptimal local minima, and there remains a detectable failure mode for gradient-based optimization methods after elimination. In this section, we provide numerical and analytical examples to illustrate these phenomena. Through these examples, we show that using \tilde{L} instead of L can help training deep neural networks in ‘good’ cases while it may not help in ‘bad’ cases.

6.1 Numerical examples

Figure 1 illustrates the novel failure mode proven by Theorem 4. Note that there exist local minima of \tilde{L}

near $\theta = 0.8$ in bounded (and open) subspaces that are also global minima. Although there exist local minima of \tilde{L} in bounded subspaces, there is no local minimum of $L|_{\theta}(a, b, W)$ with a fixed suboptimal θ . The setting used for plotting Figure 1 is summarized in Example 1, where the dataset consists of only one sample (x_1, y_1) .

Example 1. Let $m = 1$, $d_y = 1$, $x_1 = 0$, and $y_1 = -1$. In addition, let $L(\theta) = \ell(f(x_1; \theta), y_1) = (\max(0, 1 - y_1 f(x_1; \theta)))^3$. Let $f(x_1; \theta) = 5(-0.3e^{-16*(\theta-0.2)^2} - 0.7e^{-32*(\theta-0.8)^2} + 0.5)$ for a simple illustration. Because $x_1 = 0$, we can think of this function as a model with an extra parameter θ' , the effect of which disappears as $\theta'x_1 = 0$ (e.g., $f(x_1; \theta) = \bar{f}(x_1; \theta, \theta') = 5(-0.3e^{-16*(\theta'x+\theta-0.2)^2} - 0.7e^{-32*(\theta'x+\theta-0.8)^2} + 0.5)$).

A classical proof using the Weierstrass theorem to guarantee the existence of the optimal solutions in a (nonempty) subspace $S \subseteq \mathbb{R}^d$ requires a lower semi-continuity of the objective function \tilde{L} and the existence of a $q \in S$ for which the set $\{q' \in S : \tilde{L}(q') \leq \tilde{L}(q)\}$ is compact (e.g., see Bertsekas 1999 for more discussion on the existence of optimal solutions). In the above example, for the function $\tilde{L}|_{\theta}(a, b, W)$ with a fixed suboptimal θ , the former condition of lower semi-continuity is satisfied, whereas the latter condition of compactness is not.

While Figure 1 provides one of the ‘bad-case’ examples, Appendix D and the following section provide some of the ‘good-case’ examples where using \tilde{L} instead of L helps optimization of L .

6.2 Analytical examples

To further understand the properties of eliminating suboptimal local minima in an analytical manner, this section presents several analytical examples. Example 2 shows a general case where using \tilde{L} instead of L helps a gradient-based optimization method. For the simple illustration of the failure mode, Example 3 uses a single data point and squared loss. Example 4 is the version of Example 3 with two data points and shows that the value of \tilde{L} can also approach a suboptimal value. Finally, Example 5 illustrates the existence of a local minimum of \tilde{L} via only the existence of a standard global minimum θ of L . In Appendix E, Examples 6 and 7 show the same phenomena as those in Examples 3 and 4 with a smoothed hinge loss.

Example 2. Let $A[\theta] = \frac{1}{m}[(\frac{\partial f(x_1; \theta)}{\partial \theta})^\top, \dots, (\frac{\partial f(x_m; \theta)}{\partial \theta})^\top] \in \mathbb{R}^{d_\theta \times (md_y)}$ be a matrix, and $r[\varphi] = [\nabla_{y_1}(\varphi(x_1))^\top, \dots, \nabla_{y_m}(\varphi(x_m))^\top]^\top \in \mathbb{R}^{md_y}$ be a column vector given a function $\varphi : \mathbb{R}^{d_x} \rightarrow \mathbb{R}^{d_y}$. The modified objective \tilde{L} helps a gradient-based method by creating extra decreasing directions as $r[f(\cdot; \theta) + g(\cdot; a, b, W)] \notin \text{Null}(A[\theta])$ even when

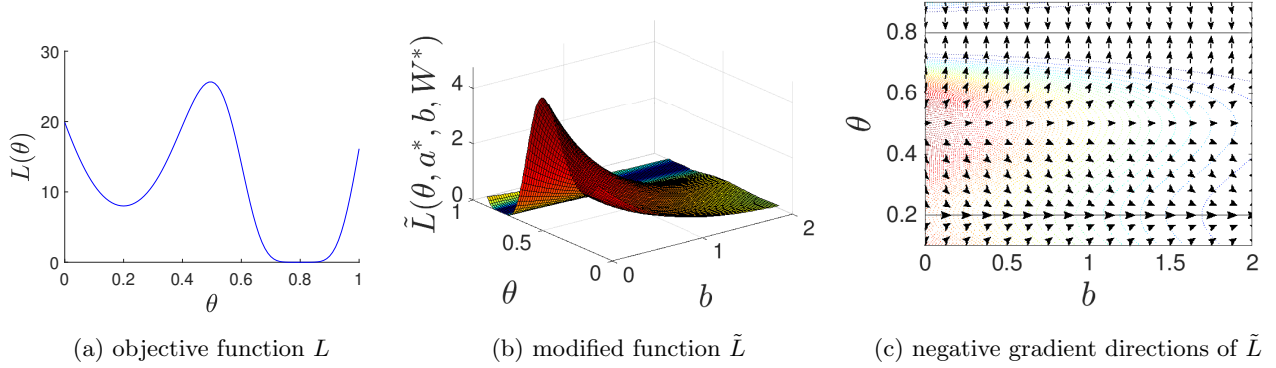


Figure 1: Illustration of the failure mode suggested by Theorem 4. In sub-figure (a), the original objective function L has a suboptimal local minimum near $\theta = 0.2$ and global minimum near $\theta = 0.8$. In sub-figures (b) and (c), it can be observed that even with the modified objective function \tilde{L} , if θ is initially near the suboptimal local minimum (0.2), a pure gradient-based local optimization algorithm can still converge to the suboptimal local minimum as $\theta \rightarrow 0.2$ and can diverge in b as $b \rightarrow \infty$. In sub-figure (c), the arrows represent the negative normalized gradient vectors at each point. In sub-figures (b) and (c), the function \tilde{L} is plotted along the coordinates (θ, b) by setting other parameters to be solutions (a^*, W^*) of each objective, $\text{minimize}_{a,W} \tilde{L}|_{\theta,b}(a, W) = \tilde{L}(\theta, a, b, W)$, at each given point (θ, b) .

$r[f(\cdot; \theta)] \in \text{Null}(A[\theta])$. This also helps optimization when $r[f(\cdot; \theta)]$ is approximately in $\text{Null}(A[\theta])$ while $r[f(\cdot; \theta) + g(\cdot; a, b, W)]$ is not.

Example 3. Let $m = 1$ and $d_y = 1$. In addition, let $L(\theta) = \ell(f(x_1; \theta), y_1) = (f(x_1; \theta) - y_1)^2$. Accordingly, $\tilde{L}(\theta, a, b, W) = (f(x_1; \theta) + a \exp(w^\top x_1 + b) - y_1)^2 + \lambda a^2$. Let θ be a non-global minimum of L as $f(x_1; \theta) \neq y_1$. In particular, let us first consider the case of $f(x_1; \theta) = 2$ and $y_1 = 1$. Then, $L(\theta) = 1$ and $\tilde{L}(\theta, a, b, W) = 1 + 2a \exp(w^\top x_1 + b) + a^2 \exp(2w^\top x_1 + 2b) + \lambda a^2$. If (a, b, W) is a local minimum, from the stationary point conditions of $\frac{\partial \tilde{L}(\theta, a, b, W)}{\partial a} = 0$ and $\frac{\partial \tilde{L}(\theta, a, b, W)}{\partial b} = 0$, we must have $a = 0$, yielding that $\tilde{L}(\theta, a, b, W) = 1$. However, a point with $a = 0$ is not a local minimum (with finite (b, w)), because with $a < 0$ and $|a| > 0$ being sufficiently small, $\tilde{L}(\theta, a, b, W) = 1 - 2|a| \cdot \exp(w^\top x_1 + b) + |a|^2 (\exp(2w^\top x_1 + 2b) + \lambda) < 1$. Hence, there is no local minimum $(a, b, W) \in \mathbb{R} \times \mathbb{R} \times \mathbb{R}^{d_x}$ of $\tilde{L}|_{\theta}$. Indeed, if we set $a = -\exp(-1/\epsilon)$ and $b = 1/\epsilon - w^\top x_1$, $\tilde{L}(\theta, a, b, W) = \lambda \exp(-2/\epsilon) \rightarrow 0$ as $\epsilon \rightarrow 0$, and hence as $a \rightarrow 0^-$ and $b \rightarrow \infty$. This illustrates the case in which (a, b) does not attain a solution in $\mathbb{R} \times \mathbb{R}$. The identical conclusion holds with the general case of $f(x_1; \theta) \neq y_1$ by following the same steps of reasoning.

Example 4. Let $m = 2$ and $d_y = 1$. In addition, $L(\theta) = (f(x_1; \theta) - y_1)^2 + (f(x_2; \theta) - y_2)^2$. Let us consider the case of $f(x_1; \theta) = f(x_2; \theta) = 0$, $y_1 = 1$, and $y_2 = -1$. Then, $L(\theta) = 2$. If (a, b, W) is a local minimum, we must have $a = 0$ similarly to Example 3, yielding that $\tilde{L}(\theta, a, b, W) = 2$. On the other hand,

$$\begin{aligned} \tilde{L}(\theta, a, b, W) & \\ &= 2 - 2a(\exp(w^\top x_1 + b) - \exp(w^\top x_2 + b)) + \varphi(a^2), \end{aligned} \quad (4)$$

where $\varphi(a^2) = a^2 \exp(2w^\top x_1 + 2b) + a^2 \exp(2w^\top x_2 + 2b) + \lambda a^2$. Note that, with a sufficiently small $|a| > 0$, the term $\varphi(a^2)$ becomes negligible. Let $x_1 \neq x_2$. In this case, our θ with $f(x_1; \theta) = f(x_2; \theta) = 0$ is not a global minimum. Then, a point with $a = 0$ can be shown to be not a local minimum as follows. If $\exp(w^\top x_1 + b) > \exp(w^\top x_2 + b)$, with $a > 0$ being sufficiently small, $\tilde{L}(\theta, a, b, W) < 2$. If $\exp(w^\top x_1 + b) < \exp(w^\top x_2 + b)$, with $a < 0$ and $|a|$ being sufficiently small, $\tilde{L}(\theta, a, b, W) < 2$. If $\exp(w^\top x_1 + b) = \exp(w^\top x_2 + b)$, since $x_1 \neq x_2$, we can perturb w with an arbitrarily small magnitude to make $\exp(w^\top x_1 + b) \neq \exp(w^\top x_2 + b)$, and hence we can yield the above cases. Thus, a point with $a = 0$ is not a local minimum. Therefore, there is no local minimum (a, b, W) of $\tilde{L}|_{\theta}$. Indeed, since $x_1 \neq x_2$, if we set $a = \exp(-1/\epsilon)$, $b = 1/\epsilon - w^\top x_1$, and $w = -\frac{1}{\epsilon}(x_2 - x_1)$, $\tilde{L}(\theta, a, b, W) = (\exp(-\|x_2 - x_1\|_2^2/\epsilon) + 1)^2 + \lambda \exp(-2/\epsilon) \rightarrow 1$, as $\epsilon \rightarrow 0$, and hence as $a \rightarrow 0^+$, $b \rightarrow \infty$ and $\|w\| \rightarrow \infty$, illustrating the case in which (a, b, W) does not attain a solution in $\mathbb{R} \times \mathbb{R} \times \mathbb{R}^{d_x}$.

Example 5. Consider the exact same example as in Example 4, with the exception that $x_1 = x_2$. In this case, a θ with $f(x_1; \theta) = f(x_2; \theta) = 0$ is a global minimum unlike in Example 4. A point with $a = 0$ is indeed a local minimum of \tilde{L} , which can be seen in Equation (4) where for all a , $2 - 2a(\exp(w^\top x_1 + b) - \exp(w^\top x_1 + b)) + \varphi(a^2) = 2 + \varphi(a^2) \geq 2$.

7 Related work

There have been many analyses regarding the optimization of deep neural networks with significant over-parameterization (Nguyen and Hein, 2017; Allen-Zhu et al., 2018; Du et al., 2018; Zou et al., 2018) and model simplification (Choromanska et al., 2015; Kawaguchi, 2016; Hardt and Ma, 2017; Bartlett et al., 2019; Du and Hu, 2019). In contrast, this paper studies the problem in the regime without significant over-parameterization or model simplification.

Our results are also different when compared to the previous study by Liang et al. (2018). First, we have solved the open problem left in the previous study; i.e., our theoretical results are applicable to practical settings and neural networks with multiple output units and common loss functions. The results in the previous study are applicable to a neural network with a single output unit for binary classification with particular smoothed hinge loss functions that are not used in common practice. In particular, the previous results are not applicable to multi-class classification or regression with any loss criteria, or binary classification with standard loss criteria (e.g., cross entropy loss and smoothed hinge loss without twice differentiability).

Second, we proved and demonstrated the failure mode of eliminating local minima as a key contribution via Theorem 4 as well as numerical and analytical examples. The failure mode proven in Theorem 4 also holds true for the results in the previous study.³ Indeed, Examples 1, 6, and 7 as well as Figure 1 directly illustrate the failure mode of the results in the previous paper. The previous study does not discuss any possible failure mode of eliminating local minima and, in fact, states that good neural networks are “just one neuron away” from bad neural networks (with suboptimal local optima). Our Theorem 4 together with analytical examples prove that such “good” neural networks with an added neuron are still subject to their own failure mode, opening up the need of future research.

Third, this paper has introduced a novel and concise proof based on the PGB necessary condition as well as a longer but more elementary proof. Our proof introduced new insight into why we can eliminate sub-optimal local minima; i.e., the global minimum value of the perturbable gradient basis of an added network is indeed the global minimum value of $L(\theta)$ (see Section 4 for more details). Beyond the elimination of local minima and the particular modification \tilde{L} , the PGB condition can be used to analyze other models and modifications, and it might be helpful to design

³This is because the assumptions of Theorem 4 are implied by the assumptions used in the previous study, and the construction of the tuple $(\ell, f, \{(x_i, y_i)\}_{i=1}^m, \theta)$ in the proof also accommodates the setting in the previous study.

new modifications of the objective functions.

In addition to the use of the PGB necessary condition, our proofs also differ from the previous proofs because the scope and the assumptions of the results are different. Indeed, the analyses of one dimensional output with $y \in \{-1, +1\}$ (the previous study) and high-dimensional output with $y \in \mathbb{R}^{d_y}$ (this paper) are naturally different. For example, when the matrix $f(X; \theta) = [f(x_1; \theta), \dots, f(x_m; \theta)]$ is rank-deficient, all outputs must be simply zero as $f(X; \theta) = 0$ in the previous study, whereas $f(X; \theta)$ can be any one of infinitely many non-zero (rank-deficient) matrices in this paper. Unlike the previous study, our proofs also cannot invoke properties of discrete points and second-order Taylor expansions because we do not assume $y \in \{-1, +1\}$ (together with the particular smoothed hinge loss) and twice differentiability.

8 Conclusion

In this paper, we proved that if an algorithm finds a local minimum of a modified objective function \tilde{L} , then it immediately recovers a global minimum of the original objective function L of an arbitrary deep neural network. However, Theorem 4 together with analytical examples showed that if an algorithm simply follows negative gradient directions toward a local minimum of \tilde{L} , either it moves toward a global minimum θ of L or the norm of (a, b, W) approaches infinity. This suggested a hybrid approach of local and global optimization algorithms, with a mechanism to monitor the norm of (a, b, W) .

From a theoretical viewpoint, we have shown a reduction of the problem of getting stuck around an arbitrarily poor local minimum into the detectable problem of the divergence of the norm. This proven reduction might be useful as a future proof technique in a theoretical literature and as a foundation of a future algorithm in practice.

In summary, this paper has advanced theoretical understanding of the properties of the optimization landscape in the regime that has not been studied well by previous research with significant over-parameterization or model simplification. Beyond the elimination of local minima, this paper has introduced the proof technique based the PGB necessary condition of local minima that can be used to study general machine learning models and transformations of objective functions.

Acknowledgements

We gratefully acknowledge support from NSF grants 1523767 and 1723381, AFOSR grant FA9550-17-1-0165, ONR grant N00014-18-1-2847, Honda Research, and the MIT-Sensetime Alliance on AI.

References

- Allen-Zhu, Z., Li, Y., and Song, Z. (2018). A convergence theory for deep learning via overparameterization. *arXiv preprint arXiv:1811.03962*.
- Andoni, A., Panigrahy, R., Valiant, G., and Zhang, L. (2014). Learning polynomials with neural networks. In *International Conference on Machine Learning*, pages 1908–1916.
- Bartlett, P. L., Helmbold, D. P., and Long, P. M. (2019). Gradient descent with identity initialization efficiently learns positive-definite linear transformations by deep residual networks. *Neural computation*, 31(3):477–502.
- Bertsekas, D. P. (1999). *Nonlinear programming*. Athena scientific Belmont.
- Blum, A. L. and Rivest, R. L. (1992). Training a 3-node neural network is NP-complete. *Neural Networks*, 5(1):117–127.
- Brescia, M. (1994). Semeion handwritten digit data set. *Semeion Research Center of Sciences of Communication*.
- Brutzkus, A. and Globerson, A. (2017). Globally optimal gradient descent for a convnet with gaussian inputs. In *International Conference on Machine Learning*, pages 605–614.
- Choromanska, A., Henaff, M., Mathieu, M., Ben Arous, G., and LeCun, Y. (2015). The loss surfaces of multilayer networks. In *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Statistics*, pages 192–204.
- Clanuwat, T., Bober-Irizar, M., Kitamoto, A., Lamb, A., Yamamoto, K., and Ha, D. (2018). Deep learning for classical japanese literature. *arXiv preprint arXiv:1812.01718*.
- Du, S. S. and Hu, W. (2019). Width provably matters in optimization for deep linear neural networks. *arXiv preprint arXiv:1901.08572*.
- Du, S. S. and Lee, J. D. (2018). On the power of overparametrization in neural networks with quadratic activation. *arXiv preprint arXiv:1803.01206*.
- Du, S. S., Lee, J. D., Li, H., Wang, L., and Zhai, X. (2018). Gradient descent finds global minima of deep neural networks. *arXiv preprint arXiv:1811.03804*.
- Gasca, M. and Sauer, T. (2000). Polynomial interpolation in several variables. *Advances in Computational Mathematics*, 12(4):377.
- Ge, R., Lee, J. D., and Ma, T. (2017). Learning one-hidden-layer neural networks with landscape design. *arXiv preprint arXiv:1711.00501*.
- Goel, S. and Klivans, A. (2017). Learning depth-three neural networks in polynomial time. *arXiv preprint arXiv:1709.06010*.
- Hardt, M. and Ma, T. (2017). Identity matters in deep learning. In *International Conference on Learning Representations*.
- Kawaguchi, K. (2016). Deep learning without poor local minima. In *Advances in Neural Information Processing Systems*, pages 586–594.
- Kawaguchi, K. and Bengio, Y. (2019). Depth with nonlinearity creates no bad local minima in ResNets. *Neural Networks*, 118:167–174.
- Kawaguchi, K. and Huang, J. (2019). Gradient descent finds global minima for generalizable deep neural networks of practical sizes. In *57th Allerton Conference on Communication, Control, and Computing (Allerton)*. IEEE.
- Kawaguchi, K., Huang, J., and Kaelbling, L. P. (2019). Every local minimum value is the global minimum value of induced model in nonconvex machine learning. *Neural Computation*, 31(12):2293–2323.
- Laurent, T. and Brecht, J. (2018). Deep linear networks with arbitrary loss: All local minima are global. In *International Conference on Machine Learning*, pages 2908–2913.
- LeCun, Y., Bottou, L., Bengio, Y., and Haffner, P. (1998). Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324.
- Li, Y. and Yuan, Y. (2017). Convergence analysis of two-layer neural networks with relu activation. In *Advances in Neural Information Processing Systems*, pages 597–607.
- Liang, S., Sun, R., Lee, J. D., and Srikant, R. (2018). Adding one neuron can eliminate all bad local minima. In *Advances in Neural Information Processing Systems*.
- Murty, K. G. and Kabadi, S. N. (1987). Some NP-complete problems in quadratic and nonlinear programming. *Mathematical programming*, 39(2):117–129.
- Nguyen, Q. and Hein, M. (2017). The loss surface of deep and wide neural networks. In *International Conference on Machine Learning*, pages 2603–2612.
- Nguyen, Q. and Hein, M. (2018). Optimization landscape and expressivity of deep cnns. In *International Conference on Machine Learning*, pages 3727–3736.
- Sedghi, H. and Anandkumar, A. (2014). Provable methods for training neural networks with sparse connectivity. *arXiv preprint arXiv:1412.2693*.

- Shamir, O. (2018). Are ResNets provably better than linear predictors? In *Advances in Neural Information Processing Systems*.
- Soltanolkotabi, M. (2017). Learning relus via gradient descent. In *Advances in Neural Information Processing Systems*, pages 2007–2017.
- Soudry, D. and Hoffer, E. (2017). Exponentially vanishing sub-optimal local minima in multilayer neural networks. *arXiv preprint arXiv:1702.05777*.
- Zhang, X., Ling, C., and Qi, L. (2012). The best rank-1 approximation of a symmetric tensor and related spherical optimization problems. *SIAM Journal on Matrix Analysis and Applications*, 33(3):806–821.
- Zhong, K., Song, Z., Jain, P., Bartlett, P. L., and Dhillon, I. S. (2017). Recovery guarantees for one-hidden-layer neural networks. In *International Conference on Machine Learning*, pages 4140–4149.
- Zou, D., Cao, Y., Zhou, D., and Gu, Q. (2018). Stochastic gradient descent optimizes over-parameterized deep relu networks. *arXiv preprint arXiv:1811.08888*.
- Zou, D. and Gu, Q. (2019). An improved analysis of training over-parameterized deep neural networks. *arXiv preprint arXiv:1906.04688*.