

Regularizing Neural Networks via Adversarial Model Perturbation

Yaowei Zheng
BDBC and SKLSDE
Beihang University, China
hiyouga@buaa.edu.cn

Richong Zhang*
BDBC and SKLSDE
Beihang University, China
zhangrc@act.buaa.edu.cn

Yongyi Mao
School of EECS
University of Ottawa, Canada
ymao@uottawa.ca

Abstract

Effective regularization techniques are highly desired in deep learning for alleviating overfitting and improving generalization. This work proposes a new regularization scheme, based on the understanding that the flat local minima of the empirical risk cause the model to generalize better. This scheme is referred to as adversarial model perturbation (AMP), where instead of directly minimizing the empirical risk, an alternative “AMP loss” is minimized via SGD. Specifically, the AMP loss is obtained from the empirical risk by applying the “worst” norm-bounded perturbation on each point in the parameter space. Comparing with most existing regularization schemes, AMP has strong theoretical justifications, in that minimizing the AMP loss can be shown theoretically to favour flat local minima of the empirical risk. Extensive experiments on various modern deep architectures establish AMP as a new state of the art among regularization schemes. Our code is available at <https://github.com/hiyouga/AMP-Regularizer>.

1. Introduction

To date, the generalization behaviour of deep neural networks is still a mystery, despite some recent progress (see, e.g., [1, 3, 19, 24, 26, 46, 48]). A commonly accepted and empirically verified understanding in this regard is that the model parameter that corresponds to a flat minimum of the empirical risk tends to generalize better. For example, the authors of [19, 26] argue that flat minima correspond to simple models, which are less likely to overfit. This understanding has inspired great effort studying factors in the optimization process (such as learning rate and batch size) that impacting the flatness of the found minima [14, 20, 25] so as to better understand the generalization behaviour of deep networks.

Meanwhile developing effective regularization techniques remains as the most important approach in practice to alleviate overfitting and force model towards better generalization

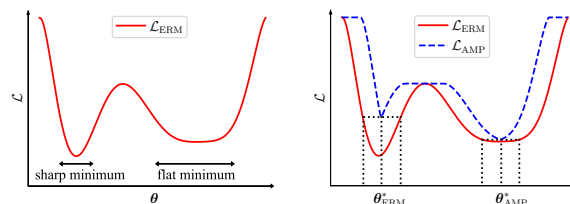


Figure 1: An example showing an empirical risk curve (left) and its corresponding AMP loss curve (right, blue).

(e.g., [16, 28, 29, 39, 41, 49]). Some recent research in fact suggests that the effectiveness of certain regularization techniques is due to their ability to find flatter minima [22, 44].

Additionally, there have been significant research advances in recent years in developing more effective regularization schemes, which include, for example, MixUp, Flooding [22, 49]. Despite their great success, these techniques usually fall short of strong principles or theoretical justifications. Thus one expects more principled and more powerful regularization schemes are yet to be discovered.

This work sets out to develop a powerful regularization scheme under the principle of finding flat local minima of the empirical risk. To that end, we propose a novel regularization scheme which can be strongly justified in terms of its ability to finding flat minima. This scheme is referred to as *Adversarial Model Perturbation* or AMP, where instead of minimizing the empirical risk $\mathcal{L}_{\text{ERM}}(\theta)$ over model parameter θ , it minimizes an alternative “AMP loss”. Briefly, the AMP loss $\mathcal{L}_{\text{AMP}}(\theta)$ at a parameter setting θ is the worst (or highest) empirical risk of all perturbations of θ with the perturbation norm no greater than a small value ϵ , namely,

$$\mathcal{L}_{\text{AMP}}(\theta) := \max_{\Delta: \|\Delta\| \leq \epsilon} \mathcal{L}_{\text{ERM}}(\theta + \Delta) \quad (1)$$

To see why minimizing the AMP loss provides opportunities to find flat local minima of the empirical risk, consider the example in Figure 1. Figure 1 (left) sketches an empirical risk curve \mathcal{L}_{ERM} , which contains two local minima, a sharp one on the left and a flat one on the right. The process of obtaining the AMP loss from the empirical risk can be seen as

*Corresponding author

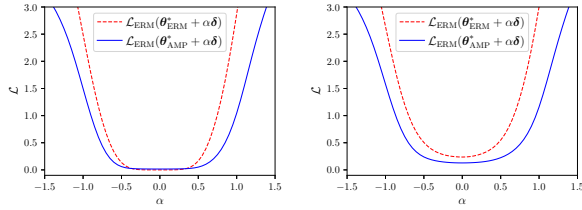


Figure 2: Landscapes of the empirical risks obtained from the PreActResNet18 [18] models trained with ERM (red) and AMP (blue) on CIFAR-10. Left: on training set; right: on test set. These curves are computed using the technique presented in [31], where δ indicates a random direction and α is a displacement in that direction.

a “max-pooling” operation, which slides a window of width 2ϵ (in high dimension, more precisely, a sphere with radius ϵ) across the parameter space and, at each location, returns the maximum value inside the window (resp. sphere). The resulting AMP loss is shown as the blue curve in Figure 1 (right). Since the right minimum in the AMP loss is lower than the left one, minimizing the AMP loss gives the right minimum as its solution.

In this paper, we formally analyze the AMP loss minimization problem and its preference of flat local minima. Specifically, we show that this minimization problem implicitly uses the “narrowest width” of a local minimum as a notion of flatness, and tries to penalize the minima that are not flat in this sense.

We derive a mini-batch SGD algorithm for solving this minimization problem, which gives rise to the proposed AMP regularization scheme. Interestingly, we show that this algorithm can also be seen as the regular empirical risk minimization with an additional penalty term on the gradient norm. This provides an alternative justification of the AMP scheme. Figure 2 contains an experimental result suggesting that AMP indeed selects flatter minima than ERM does.

We conduct experiments on several benchmark image classification datasets (SVHN, CIFAR-10, CIFAR-100) to validate the effectiveness of the proposed AMP scheme. Compared with other popular regularization schemes, AMP demonstrates remarkable regularization performance, establishing itself as a new state of the art.

Our contributions can be summarized as follows.

- 1) Motivated by the understanding that flat minima help generalization, we propose adversarial model perturbation (AMP) as a novel and efficient regularization scheme.
- 2) We theoretically justify that AMP is capable of finding flatter local minima, thereby improving generalization.
- 3) Extensive experiments on the benchmark datasets demonstrate that AMP achieves the best performance among the compared regularization schemes on various modern neural network architectures.

2. Related Work

2.1. Flat Minima and Generalization

There has been a rich body of works that investigate the relationship between the flatness of the local minima and the generalization performance of a deep neural network [4, 19, 26, 31]. In [19], it is suggested that flat minima correspond to low-complexity networks, which tend to generalize well under the principle of minimum description length [37]. [4] presents another explanation supporting this argument through the lens of Gibbs free energy. [31] demonstrates the better generalization of flat minima by visualizing the loss landscape. The empirical results in [26] show that gradient descent with large batch size finds sharp minima while small batch sizes lead to flatter minima and provide better generalization. On the other hand, [8] argues that sharp minima do not necessarily lead to poor generalization. The argument is that due to the parameterization redundancy in deep networks and under a certain notion of flatness, one can transform a flat minimum to an equivalent sharp one. Nonetheless, it is in general accepted and empirically verified that flat minima tend to give better generalization performance, and this understanding underlies the design of several recent regularization techniques (*e.g.*, [10, 22, 23, 44]). A concurrent work of [9] further provides a PAC-Bayesian justification as to why the flatness of the minima helps generalization.

2.2. Regularization

Regularization may broadly refer to any training techniques that help to improve generalization. Despite the well-known regularization techniques such as weight decay [29], Dropout [39], normalization tricks [2, 21] and data augmentation [6, 28, 32], various new techniques have been developed. Label smoothing [41] mixes the one-hot label of the training example with a uniform distribution over the label space. Shake-Shake regularization [11] combines parallel branches with a stochastic affine function in multi-branch networks. ShakeDrop regularization [45] extends Shake-Shake to single-branch deep architectures. Cutout [7] randomly masks out some regions of input images in training. MixUp [49] regularizes deep networks by perturbing training samples along the direction of other samples. Flooding [22] forces training loss to stay above zero to avoid overfitting. Adversarial training techniques [13, 33], originally designed for improving the model’s adversarial robustness, are also shown to have great regularization effect when their parameters are carefully chosen [33]. But it is also observed that training the model excessively towards adversarial robustness may hurt generalization [42]. Some recent progresses (*e.g.*, [10, 23]) exploit multiple states of the model parameter while training and ensembling them to improve generalization. The concurrent work [9] also independently discovered a similar regularization scheme as we present in this paper.

3. AMP: Adversarial Model Perturbation

3.1. From Empirical Risk to AMP Loss

Consider a classification setting, where we aim at finding a classifier $f : \mathcal{X} \rightarrow \mathcal{Y}$ that maps the input space \mathcal{X} to the label space \mathcal{Y} . Note that we may take \mathcal{Y} as the set of all distributions over the set of possible labels, so that not only $f(\mathbf{x})$ belongs to \mathcal{Y} , each ground-truth label also belongs to \mathcal{Y} since it can be written as a degenerated distribution or a “one-hot” vector.

We will take f as a neural network parameterized by θ , taking values from its weight space Θ . For each training example $(\mathbf{x}, \mathbf{y}) \in \mathcal{X} \times \mathcal{Y}$, we will denote by $\ell(\mathbf{x}, \mathbf{y}; \theta)$ the loss of the prediction $f(\mathbf{x})$ by the model with respect to the true label \mathbf{y} . That is, the function ℓ absorbs the function f within.

Under the empirical risk minimization (ERM) principle [43], the training of the neural network using a training set \mathcal{D} is performed by minimizing the following *empirical risk* (which we also refer to as “ERM loss”) \mathcal{L}_{ERM} over parameter θ :

$$\mathcal{L}_{\text{ERM}}(\theta) := \frac{1}{|\mathcal{D}|} \sum_{(\mathbf{x}, \mathbf{y}) \in \mathcal{D}} \ell(\mathbf{x}, \mathbf{y}; \theta) \quad (2)$$

It is however well known that ERM training is prone to overfitting [43] and that the learned model often fails to generalize well on the unseen examples.

This work is motivated by the need of effective regularization schemes to improve the model’s generalization capability. Specifically, our objective is to develop a technique that forces the training process to find flatter minima or low-norm solutions when minimizing the empirical risk. The technique we develop is termed “adversarial model perturbation” or AMP, which we now elaborate.

For any positive ϵ and any $\mu \in \Theta$, let $\mathbf{B}(\mu; \epsilon)$ denote the norm ball in Θ with radius ϵ centred at μ . That is

$$\mathbf{B}(\mu; \epsilon) := \{\theta \in \Theta : \|\theta - \mu\| \leq \epsilon\} \quad (3)$$

We note that the norm used in defining the norm ball is chosen as L_2 norm. However, we expect that the norm can be extended beyond this choice.

We now define an “AMP loss” \mathcal{L}_{AMP} as follows.

$$\mathcal{L}_{\text{AMP}}(\theta) := \max_{\Delta \in \mathbf{B}(\mathbf{0}; \epsilon)} \frac{1}{|\mathcal{D}|} \sum_{(\mathbf{x}, \mathbf{y}) \in \mathcal{D}} \ell(\mathbf{x}, \mathbf{y}; \theta + \Delta) \quad (4)$$

where the ϵ is a small positive value serving as a hyperparameter.

Figure 1 sketches an example demonstrating that minimizing the AMP loss \mathcal{L}_{AMP} over θ presents opportunities of finding flatter minima of the ERM loss \mathcal{L}_{ERM} . Theoretical justifications for AMP’s capability of finding flatter minima are given in Section 4.

3.2. Training Algorithm

Practical considerations on computation complexity and training speed motivate a mini-batched training approach to minimizing the AMP loss \mathcal{L}_{AMP} . Specifically, we may approximate the AMP loss in (4) using the corresponding loss obtained from a random mini-batch \mathcal{B} , namely,

$$\begin{aligned} \mathcal{L}_{\text{AMP}}(\theta) &\approx \max_{\Delta \in \mathbf{B}(\mathbf{0}; \epsilon)} \frac{1}{|\mathcal{B}|} \sum_{(\mathbf{x}, \mathbf{y}) \in \mathcal{B}} \ell(\mathbf{x}, \mathbf{y}; \theta + \Delta) \\ &:= \mathcal{J}_{\text{AMP}, \mathcal{B}}(\theta) \end{aligned} \quad (5)$$

Then the AMP loss minimization problem can be approximated as minimizing the expected batch-level AMP loss $\mathcal{J}_{\text{AMP}, \mathcal{B}}$, formally as finding

$$\theta_{\text{AMP}}^* := \arg \min_{\theta} \mathbb{E}_{\mathcal{B}} \mathcal{J}_{\text{AMP}, \mathcal{B}}(\theta) \quad (6)$$

Under this formulation, each batch \mathcal{B} is associated with a perturbation vector $\Delta_{\mathcal{B}}$ on the model parameter θ . The training involves an inner maximization nested inside and outer minimization: in the inner maximization, a fixed number, say N , of steps are used to update $\Delta_{\mathcal{B}}$ in the direction of increasing the ERM loss so as to obtain $\mathcal{J}_{\text{AMP}, \mathcal{B}}$; the outer minimization loops over random batches and minimizes $\mathbb{E}_{\mathcal{B}} \mathcal{J}_{\text{AMP}, \mathcal{B}}$ using mini-batched SGD. The precise algorithm is described in Algorithm 1. Note that at the end of training, the learned θ_{AMP}^* is used as model parameter for predictions; that is, no perturbation is applied in testing.

Algorithm 1 Adversarial Model Perturbation Training

Require: Training set $\mathcal{D} = \{(\mathbf{x}, \mathbf{y})\}$, Batch size m , Loss function ℓ , Initial model parameter θ_0 , Outer learning rate η , Inner learning rate ζ , Inner iteration number N , L_2 norm ball radius ϵ

- 1: **while** θ_k not converged **do**
 - 2: Update iteration: $k \leftarrow k + 1$
 - 3: Sample $\mathcal{B} = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^m$ from training set \mathcal{D}
 - 4: Initialize perturbation: $\Delta_{\mathcal{B}} \leftarrow \mathbf{0}$
 - 5: **for** $n \leftarrow 1$ to N **do**
 - 6: Compute gradient:

$$\nabla \mathcal{J}_{\text{AMP}, \mathcal{B}} \leftarrow \sum_{i=1}^m \nabla_{\theta} \ell(\mathbf{x}_i, \mathbf{y}_i; \theta_k + \Delta_{\mathcal{B}}) / m$$
 - 7: Update perturbation: $\Delta_{\mathcal{B}} \leftarrow \Delta_{\mathcal{B}} + \zeta \nabla \mathcal{J}_{\text{AMP}, \mathcal{B}}$
 - 8: **if** $\|\Delta_{\mathcal{B}}\|_2 > \epsilon$ **then**
 - 9: Normalize perturbation: $\Delta_{\mathcal{B}} \leftarrow \epsilon \Delta_{\mathcal{B}} / \|\Delta_{\mathcal{B}}\|_2$
 - 10: **end if**
 - 11: **end for**
 - 12: Compute gradient:

$$\nabla \mathcal{J}_{\text{AMP}, \mathcal{B}} \leftarrow \sum_{i=1}^m \nabla_{\theta} \ell(\mathbf{x}_i, \mathbf{y}_i; \theta_k + \Delta_{\mathcal{B}}) / m$$
 - 13: Update parameter: $\theta_{k+1} \leftarrow \theta_k - \eta \nabla \mathcal{J}_{\text{AMP}, \mathcal{B}}$
 - 14: **end while**
-

4. Theoretical Justifications of AMP

4.1. AMP Finds Flatter Local Minima

For the ease of obtaining analytic results, we will assume that the loss surface of each local minimum in \mathcal{L}_{ERM} can be approximated as an inverted Gaussian surface.

More precisely, suppose $\Theta = \mathbb{R}^K$. For any scalar C , any positive scalar A , any $\boldsymbol{\mu} \in \Theta$ and any $K \times K$ positive definite matrix $\boldsymbol{\kappa}$, let

$$\gamma(\boldsymbol{\theta}; \boldsymbol{\mu}, \boldsymbol{\kappa}, A, C) := C - A \exp\left(-(\boldsymbol{\theta} - \boldsymbol{\mu})^T \boldsymbol{\kappa}^{-1} (\boldsymbol{\theta} - \boldsymbol{\mu}) / 2\right) \quad (7)$$

Note that this function is merely an inverted Gaussian surface over the space Θ .

Locally Gaussian Assumption of Empirical Risk Using the above notation, we assume that each minimum of the empirical risk can be *locally* approximated by such an inverted Gaussian surface, namely, that if $\boldsymbol{\mu} \in \Theta$ gives a local minimum of \mathcal{L}_{ERM} , then there exist some $\epsilon > 0$, a positive definite matrix $\boldsymbol{\kappa}$ and scalars A and C with $A > 0$, such that

$$\mathcal{L}_{\text{ERM}}(\boldsymbol{\theta}) = \gamma(\boldsymbol{\theta}; \boldsymbol{\mu}, \boldsymbol{\kappa}, A, C) \quad (8)$$

at any $\boldsymbol{\theta} \in \mathbf{B}(\boldsymbol{\mu}; 2\epsilon)$. We note that in general, such an assumption is only an approximation. But for small ϵ , the approximation is arguably accurate.

We will use $\gamma^*(\boldsymbol{\mu}, \boldsymbol{\kappa}, A, C)$ to denote the minimum value of function $\gamma(\boldsymbol{\theta}; \boldsymbol{\mu}, \boldsymbol{\kappa}, A, C)$ (obtained by minimizing over $\boldsymbol{\theta}$). It is apparent that

$$\gamma^*(\boldsymbol{\mu}, \boldsymbol{\kappa}, A, C) = C - A \quad (9)$$

Let γ_{AMP} denote the AMP loss derived from γ , namely,

$$\gamma_{\text{AMP}}(\boldsymbol{\theta}; \boldsymbol{\mu}, \boldsymbol{\kappa}, A, C) := \max_{\Delta \in \mathbf{B}(\mathbf{0}; \epsilon)} \gamma(\boldsymbol{\theta} + \Delta; \boldsymbol{\mu}, \boldsymbol{\kappa}, A, C) \quad (10)$$

We will use $\gamma_{\text{AMP}}^*(\boldsymbol{\mu}, \boldsymbol{\kappa}, A, C)$ to denote the minimum value of $\gamma_{\text{AMP}}(\boldsymbol{\theta}; \boldsymbol{\mu}, \boldsymbol{\kappa}, A, C)$ (minimized over $\boldsymbol{\theta}$).

Theorem 1. *For any given $(\boldsymbol{\mu}, \boldsymbol{\kappa}, A, C)$, the function $\gamma_{\text{AMP}}(\boldsymbol{\theta}; \boldsymbol{\mu}, \boldsymbol{\kappa}, A, C)$ is minimized when $\boldsymbol{\theta} = \boldsymbol{\mu}$ and the minimum value is*

$$\gamma_{\text{AMP}}^*(\boldsymbol{\mu}, \boldsymbol{\kappa}, A, C) = C - A \exp\left(-\frac{\epsilon^2}{2\sigma^2}\right) \quad (11)$$

where σ^2 is the smallest eigenvalue of $\boldsymbol{\kappa}$.

Proof. The properties of a Gaussian surface suggest that the minimum of $\gamma_{\text{AMP}} = (\boldsymbol{\theta}; \boldsymbol{\mu}, \boldsymbol{\kappa}, A, C)$ is given by $\boldsymbol{\theta} = \boldsymbol{\mu}$ and the inner maximum is reached at $\Delta = \epsilon \hat{\mathbf{q}}$, where $\hat{\mathbf{q}}$ is the normalized eigenvector corresponding to the smallest eigenvalue of $\boldsymbol{\kappa}$. That is because that the direction of $\hat{\mathbf{q}}$ is one of the fastest increase of γ_{AMP} . Such a direction can also be regarded as the direction of the ‘‘narrowest width’’ of the Gaussian surface, which is illustrated in Figure 3.

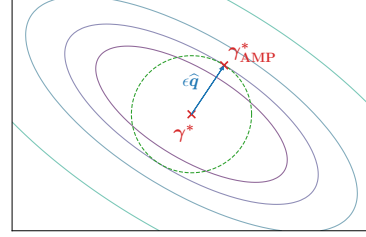


Figure 3: Locally Gaussian assumption and the minimum values of γ and γ_{AMP} , respectively.

The symmetric positive definite matrix $\boldsymbol{\kappa}$ can be factorized as $\mathbf{Q}\boldsymbol{\Lambda}\mathbf{Q}^T$, where \mathbf{Q} is the $K \times K$ orthogonal matrix whose i -th column is the normalized eigenvector \mathbf{q}_i of the matrix $\boldsymbol{\kappa}$, and $\boldsymbol{\Lambda}$ is the diagonal matrix whose diagonal elements are the corresponding eigenvalues. Let σ^2 denote the smallest eigenvalue of $\boldsymbol{\kappa}$, and suppose that the eigenvalues are in ascending order along the diagonal of $\boldsymbol{\Lambda}$. Let \mathbf{e}_1 denote the vector $(1, 0, \dots, 0)^T$ whose first element is 1 and others are 0. We have:

$$\begin{aligned} \gamma_{\text{AMP}}^*(\boldsymbol{\mu}, \boldsymbol{\kappa}, A, C) &= C - A \exp\left(-\frac{\epsilon^2 \hat{\mathbf{q}}^T \mathbf{Q}^T \boldsymbol{\Lambda}^{-1} \mathbf{Q} \hat{\mathbf{q}}}{2}\right) \\ &= C - A \exp\left(-\frac{\epsilon^2 \mathbf{e}_1^T \boldsymbol{\Lambda}^{-1} \mathbf{e}_1}{2}\right) \\ &= C - A \exp\left(-\frac{\epsilon^2}{2\sigma^2}\right) \end{aligned} \quad (12)$$

This proves the theorem. \square

From this theorem, it is clear that the minimum value achieved with γ_{AMP}^* although related to the minimum ERM loss value γ^* through A and C , it also takes into account the curvature of the surface around the local minimum. Specifically, the smallest eigenvalue σ^2 measures the ‘‘width’’ of the surface along its narrowest principal direction (noting that the cross-section of the surface is an ellipsoid). Thus the value σ^2 can be treated as a ‘‘worst-case’’ measure of the flatness or width of the surface. The larger is σ^2 , the flatter or wider is the local minimum.

Following this theorem, we next show that minimizing the AMP loss \mathcal{L}_{AMP} favours the solutions corresponding to ‘‘flatter’’ local minima of the empirical risk \mathcal{L}_{ERM} .

Corollary 1. *Let $\boldsymbol{\mu}_1, \boldsymbol{\mu}_2 \in \Theta$ be two local minima of \mathcal{L}_{ERM} . Assume the locally Gaussian assumption hold such that the surface of the two local minima follow respectively $\gamma_1(\boldsymbol{\theta}; \boldsymbol{\mu}_1, \boldsymbol{\kappa}_1, A_1, C_1)$ and $\gamma_2(\boldsymbol{\theta}; \boldsymbol{\mu}_2, \boldsymbol{\kappa}_2, A_2, C_2)$. Then $\gamma_1^* < \gamma_2^*$ but $\gamma_{1,\text{AMP}}^* > \gamma_{2,\text{AMP}}^*$ if and only if*

$$\begin{aligned} A_1 - A_2 &> C_1 - C_2 \\ &> A_1 \exp\left(-\frac{\epsilon^2}{2\sigma_1^2}\right) - A_2 \exp\left(-\frac{\epsilon^2}{2\sigma_2^2}\right) \end{aligned} \quad (13)$$

where σ_1^2 and σ_2^2 are the smallest eigenvalues of κ_1 and κ_2 respectively.

Under the condition of this corollary, although μ_1 gives a lower empirical risk than μ_2 , but when we minimize the AMP loss, μ_2 is a more preferred solution since the local curvatures of the two minima in \mathcal{L}_{ERM} are also considered. We next show, using the special case of $C_1 = C_2$, that indeed minimizing AMP loss favours the local minima with a flatter surface.

Corollary 2. *Suppose that $C_1 = C_2$. Let $A_2 = \beta A_1$ for some $\beta < 1$. Note that in this setting, $\gamma_1^* < \gamma_2^*$. Suppose that $\sigma_2^2 = r\sigma_1^2$ for some positive r . Then*

$$\gamma_{2,\text{AMP}}^* < \gamma_{1,\text{AMP}}^* \quad (14)$$

if and only if

$$\beta > \exp\left(-\frac{\epsilon^2}{2\sigma_1^2}\right) \text{ and } r > \frac{1}{1 + \frac{2\sigma_1^2}{\epsilon^2} \log \beta} \quad (15)$$

Please refer to Appendix A for the details of the proof.

Note that in this setting, the local minimum in the empirical risk corresponding to μ_1 is lower than that corresponding to μ_2 . The value β governs how close the two minimum values are; the closer to 1 is β , the closer two minimum values are. The value r governs the flatness of the second local minimum relative to the first: $r > 1$ indicates the second is flatter than the first, and the larger is r , the flatter is the second minimum. This corollary presents a sufficient and necessary condition for the second minimum to be preferred to the first when the AMP loss is minimized. Specifically, we may refer to the set of all (β, r) pairs that satisfy the condition as the ‘‘operational region of AMP’’, since the region specifies all points on which minimizing the AMP loss will give a solution that deviates from that given by minimizing the ERM loss. The general shape of such a region is shown in Figure 4 (left). In Figure 4 (right), the regions are plotted for different values of $\epsilon^2/2\sigma_1^2$. Assuming σ_1^2 to be a fixed value, it can then be seen that as ϵ decreases, the operational region of AMP shrinks, namely that the minimization of the AMP loss has decreased opportunity to deviate from minimizing the empirical risk; in the limit when ϵ approaches 0, minimizing the AMP loss reduces to minimizing the ERM loss. On the other hand, for a large value of ϵ , the operational region of AMP is large, then minimizing the AMP loss may frequently find different solutions from those obtained from minimizing the empirical risk. In this case however, the operational region includes points (β, r) with small β and relatively small r . Such points, for example the one marked with ‘‘ \times ’’ in Figure 4 (right), corresponds to local minima not flatter than μ_1 by much but having much higher empirical risk values. When such solutions are obtained by

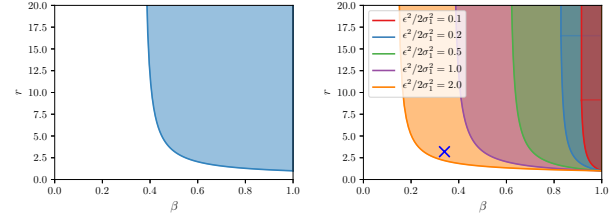


Figure 4: The operational region of AMP (left) and how it varies with varying values of $\epsilon^2/2\sigma_1^2$ (right).

minimizing the AMP loss, the learned model risks significant underfitting. Thus, in general there is a sweet spot of ϵ setting that gives the optimal tradeoff between the flatness and depth of the selected local minima.

4.2. AMP Regularizes Gradient Norm

The fact that AMP favours flatter minima can also be seen from the AMP training algorithm, independent of the above analysis. Specifically, we now show that the AMP algorithm may be viewed as minimizing the empirical risk with a certain penalty on its gradient norm.

To see this, consider that the number N of inner updates per batch is 1 (this is in fact used in our experiments). Denote by $\mathcal{J}_{\text{ERM},\mathcal{B}}$ the approximation of the empirical loss \mathcal{L}_{ERM} using batch \mathcal{B} , that is,

$$\mathcal{J}_{\text{ERM},\mathcal{B}}(\theta) := \frac{1}{|\mathcal{B}|} \sum_{(x,y) \in \mathcal{B}} \ell(x,y;\theta) \quad (16)$$

Theorem 2. *Let $N = 1$. Then for a sufficiently small inner learning rate ζ , a minimization update step in AMP training for batch \mathcal{B} is equivalent to a gradient-descent step on the following loss function with learning rate η :*

$$\tilde{\mathcal{J}}_{\text{ERM},\mathcal{B}}(\theta) := \mathcal{J}_{\text{ERM},\mathcal{B}}(\theta) + \Omega(\theta) \quad (17)$$

where

$$\Omega(\theta) := \begin{cases} \zeta \|\nabla_{\theta} \mathcal{J}_{\text{ERM},\mathcal{B}}(\theta)\|_2^2, & \|\zeta \nabla_{\theta} \mathcal{J}_{\text{ERM},\mathcal{B}}(\theta)\|_2 \leq \epsilon \\ \epsilon \|\nabla_{\theta} \mathcal{J}_{\text{ERM},\mathcal{B}}(\theta)\|_2, & \|\zeta \nabla_{\theta} \mathcal{J}_{\text{ERM},\mathcal{B}}(\theta)\|_2 > \epsilon \end{cases} \quad (18)$$

Please refer to Appendix B for the details of the proof.

We note that the regularization term $\Omega(\theta)$, in either one of the two cases, penalizes the gradient norm of $\mathcal{J}_{\text{ERM},\mathcal{B}}$. Thus the AMP training algorithm effectively tries to find local minima of $\mathcal{J}_{\text{ERM},\mathcal{B}}$ (and hence of \mathcal{L}_{ERM}) that not only have low values, but also have small gradient norm near the minima. Note that a minimum with smaller gradient norms around it is a flatter minimum. This theorem therefore provides another justification of the AMP training algorithm, in addition to our development from constructing the AMP loss \mathcal{L}_{AMP} .

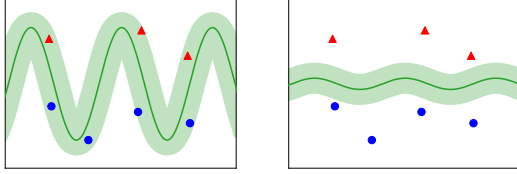


Figure 5: A bad classification boundary (left) usually has a poor generalization performance and is fragile to adversarial attacks. A better classification boundary (right) keeps itself far from the training examples, and tends to generalize well.

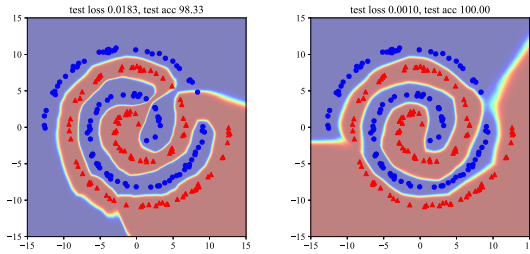


Figure 6: AMP yields a better classification boundary (right) than ERM does (left). The experiment is conducted on a spiral dataset [40] using a feed-forward network.

4.3. Perspectives from the Input Space

The construction and analysis so far have focused on the parameter space. Further insights may be obtained by inspecting AMP in the input space.

It is evident that the effect of parameter θ on the input space is no more than defining the class boundaries and specifying how the loss $\ell(\mathbf{x}, \mathbf{y})$ changes with input \mathbf{x} . Then minimizing the AMP loss can be seen as finding a class boundary (and loss function $\ell(\mathbf{x}, \mathbf{y})$) which has the lowest average loss over the training examples, even when the worst ϵ -bounded perturbation is applied. A consequence of such a minimization is arguably creating smoother class boundaries and keeping the training examples not too close to the boundaries. This is illustrated in Figure 5 and experimentally validated in Figure 6.

When viewed from input space, the adversarial perturbation of the model parameter in AMP shares some similarity with adversarial training [13]. There is also a significant difference between the two: adversarial training defends the model against adversarial attacks whereas AMP defends it against overfitting. We further elaborate why AMP is different from adversarial training in Appendix C.

5. Experiment

We empirically investigate the performance of AMP in various perspectives. Firstly, we compare the generalization ability of AMP on benchmark image classification datasets

with several popular regularization schemes in Section 5.1, including Dropout [39], label smoothing [41], Flooding [22], MixUp [49] and adversarial training [13]. We also compare our scheme with ERM [43], which does not utilize any regularization and optimizes the neural network with \mathcal{L}_{ERM} . We include a baseline named random model perturbation (RMP) for comparison. Specifically, RMP applies a random perturbation (instead of the “worst perturbation”) to the parameter within a small range to help the model to find a better minimum. Then, we study the performance of AMP on more complex deep architectures with powerful data augmentation techniques in Section 5.2. In addition, we investigate the calibration effect of AMP in Section 5.3 and demonstrate the influence of perturbation in Section 5.4. Finally, we compare the computational cost of AMP with ERM in Section 5.5. The implementation is on PyTorch framework [35], and the experiments are carried out on NVIDIA Tesla V100 GPUs.

5.1. Comparison of the Generalization Ability

Three publicly available benchmark image datasets are used for performance evaluation. The SVHN dataset [34] has 10 classes containing 73,257 digits for training and 26,032 digits for testing. Limited by the computing resource, we did not use the additional 531,131 images in SVHN training. The CIFAR datasets contain 32×32 -pixel colour images, where CIFAR-10 has 10 classes containing 5,000 images for training and 1,000 images for testing per class, CIFAR-100 has 100 classes containing 500 images for training and 100 images for testing per class [27].

Two representative deep architectures for image classification, PreActResNet18 [18] and VGG16 [38], are taken as the underlying classifier. In the training procedure, random crops and horizontal flips are adopted as data augmentation schemes. We compute the mean and standard deviation on the training set to normalize the input images. SGD with momentum is exploited as the optimizer with a step-wise learning rate decay. Specifically, the outer learning rate is initialized as 0.1 and divided by 10 after 100 and 150 epochs. We train each model for 200 epochs on the training set with 50 examples per mini-batch. Weight decay is set to 10^{-4} for all compared models. We tune hyper-parameters on each dataset using 10% of the training set as the validation set. For Dropout, we randomly choose 10% of the neurons in each layer after ReLU activation and deactivate them at each training iteration. The label smoothing coefficient is set to 0.2 and the flooding level is set to 0.02. For MixUp, we follow the original study [49] and linearly combine random pairs of training examples by using coefficient variables drawn from Beta(1, 1). For adversarial training, we set the perturbation size to 1 for each pixel and take one single step to generate adversarial examples. For RMP, we set the L_2 norm ball radius to 0.1. For AMP, we fix the number of inner iteration as $N = 1$, and adopt $\epsilon = 0.5, \gamma = 1$ for PreActResNet18

PreActResNet18	Test Error (%)	Test NLL	PreActResNet18	Test Error (%)	Test NLL	PreActResNet18	Test Error (%)	Test NLL
ERM	2.95±0.063	0.166±0.004	ERM	5.02±0.212	0.239±0.009	ERM	24.31±0.303	1.056±0.013
Dropout	2.80±0.065	0.156±0.012	Dropout	4.86±0.148	0.223±0.009	Dropout	24.48±0.351	1.110±0.021
Label Smoothing	2.78±0.087	0.998±0.002	Label Smoothing	4.85±0.115	1.038±0.003	Label Smoothing	22.07±0.256	2.099±0.005
Flooding	2.84±0.047	0.130±0.003	Flooding	4.97±0.082	0.166±0.003	Flooding	24.50±0.234	0.950±0.011
MixUp	2.74±0.044	0.146±0.004	MixUp	4.09±0.117	0.198±0.004	MixUp	21.78±0.210	0.910±0.007
Adv. Training	2.77±0.080	0.151±0.018	Adv. Training	4.99±0.085	0.247±0.006	Adv. Training	25.23±0.229	1.110±0.012
RMP	2.93±0.066	0.161±0.010	RMP	4.97±0.167	0.239±0.008	RMP	24.28±0.138	1.059±0.011
AMP	2.30±0.025	0.096±0.002	AMP	3.97±0.091	0.129±0.003	AMP	21.51±0.308	0.774±0.016

VGG16	Test Error (%)	Test NLL	VGG16	Test Error (%)	Test NLL	VGG16	Test Error (%)	Test NLL
ERM	3.14±0.060	0.140±0.027	ERM	6.32±0.193	0.361±0.012	ERM	27.84±0.297	1.827±0.209
Dropout	2.96±0.049	0.134±0.027	Dropout	6.22±0.147	0.314±0.009	Dropout	27.72±0.337	1.605±0.062
Label Smoothing	3.07±0.070	1.004±0.002	Label Smoothing	6.29±0.158	1.076±0.003	Label Smoothing	27.49±0.179	2.310±0.005
Flooding	3.15±0.085	0.128±0.003	Flooding	6.26±0.145	0.234±0.005	Flooding	27.93±0.271	1.221±0.037
MixUp	3.09±0.057	0.160±0.003	MixUp	5.48±0.112	0.251±0.003	MixUp	26.81±0.254	1.136±0.013
Adv. Training	2.94±0.091	0.122±0.003	Adv. Training	6.49±0.130	0.380±0.010	Adv. Training	29.12±0.145	1.535±0.389
RMP	3.19±0.052	0.134±0.004	RMP	6.30±0.109	0.363±0.010	RMP	27.81±0.327	1.873±0.035
AMP	2.73±0.015	0.116±0.006	AMP	5.65±0.147	0.207±0.005	AMP	25.60±0.168	1.049±0.049

PreActResNet18	Test Error (%)	Test NLL	PreActResNet18	Test Error (%)	Test NLL	PreActResNet18	Test Error (%)	Test NLL
ERM	2.95±0.063	0.166±0.004	ERM	5.02±0.212	0.239±0.009	ERM	24.31±0.303	1.056±0.013
Dropout	2.80±0.065	0.156±0.012	Dropout	4.86±0.148	0.223±0.009	Dropout	24.48±0.351	1.110±0.021
Label Smoothing	2.78±0.087	0.998±0.002	Label Smoothing	4.85±0.115	1.038±0.003	Label Smoothing	22.07±0.256	2.099±0.005
Flooding	2.84±0.047	0.130±0.003	Flooding	4.97±0.082	0.166±0.003	Flooding	24.50±0.234	0.950±0.011
MixUp	2.74±0.044	0.146±0.004	MixUp	4.09±0.117	0.198±0.004	MixUp	21.78±0.210	0.910±0.007
Adv. Training	2.77±0.080	0.151±0.018	Adv. Training	4.99±0.085	0.247±0.006	Adv. Training	25.23±0.229	1.110±0.012
RMP	2.93±0.066	0.161±0.010	RMP	4.97±0.167	0.239±0.008	RMP	24.28±0.138	1.059±0.011
AMP	2.30±0.025	0.096±0.002	AMP	3.97±0.091	0.129±0.003	AMP	21.51±0.308	0.774±0.016

VGG16	Test Error (%)	Test NLL	VGG16	Test Error (%)	Test NLL	VGG16	Test Error (%)	Test NLL
ERM	3.14±0.060	0.140±0.027	ERM	6.32±0.193	0.361±0.012	ERM	27.84±0.297	1.827±0.209
Dropout	2.96±0.049	0.134±0.027	Dropout	6.22±0.147	0.314±0.009	Dropout	27.72±0.337	1.605±0.062
Label Smoothing	3.07±0.070	1.004±0.002	Label Smoothing	6.29±0.158	1.076±0.003	Label Smoothing	27.49±0.179	2.310±0.005
Flooding	3.15±0.085	0.128±0.003	Flooding	6.26±0.145	0.234±0.005	Flooding	27.93±0.271	1.221±0.037
MixUp	3.09±0.057	0.160±0.003	MixUp	5.48±0.112	0.251±0.003	MixUp	26.81±0.254	1.136±0.013
Adv. Training	2.94±0.091	0.122±0.003	Adv. Training	6.49±0.130	0.380±0.010	Adv. Training	29.12±0.145	1.535±0.389
RMP	3.19±0.052	0.134±0.004	RMP	6.30±0.109	0.363±0.010	RMP	27.81±0.327	1.873±0.035
AMP	2.73±0.015	0.116±0.006	AMP	5.65±0.147	0.207±0.005	AMP	25.60±0.168	1.049±0.049

(a) SVHN

(b) CIFAR-10

(c) CIFAR-100

Table 1: Top-1 classification errors and test neg-log-likelihoods on (a) SVHN, (b) CIFAR-10 and (c) CIFAR-100. We run experiments 10 times to report the mean and the standard deviation of errors and neg-log-likelihoods.

and $\epsilon = 0.1, \gamma = 0.2$ for VGG16. Top-1 classification error and test neg-log-likelihood are reported in Table 1.

From Table 1, it is evident that AMP outperforms the baseline methods in various settings, in terms of both classification error and test neg-log-likelihood, except on the CIFAR-10 dataset where VGG16 is employed. Despite the remarkable performance of AMP, MixUp also demonstrates competitive improvement in classification accuracy, and Flooding achieves small testing neg-log-likelihood in many settings. We note that, compared with AMP, the results of RMP suggest that randomly perturbing parameters cannot obtain comparable performance to the AMP. This result confirms that the adversarial perturbation provides the most useful information to the regularizer. Above results describe the efficiency of AMP in regularizing deep networks.

5.2. Improvement over Data Augmentation

Data augmentation techniques can be viewed as regularization schemes since through introducing additional training examples, they impose additional constraints on the model parameter thereby improving generalization. To validate the effectiveness of AMP over other data augmentation techniques, we choose vanilla augmentation [28], Cutout [7] and AutoAugment [6] as underlying augmentation methods and compare the classification accuracy of AMP with ERM. The vanilla augmentation exploits manually designed policies including random crops and horizontal flips. We use the same Cutout configuration and AutoAugment policy as their corresponding original studies. For the hyper-parameters of AMP, we fix $N = 1$ and adopt $\epsilon = 0.5, \gamma = 1$ for vanilla augmentation, $\epsilon = 0.3, \gamma = 0.5$ for Cutout, and $\epsilon = 0.1, \gamma = 0.1$ for AutoAugment. We employ two recent powerful deep architectures, WideResNet [47] and PyramidNet [17], with the

		WideResNet-28-10		PyramidNet-164-270	
		ERM	AMP	ERM	AMP
SVHN	Vanilla	2.57±0.067	2.19±0.036	2.47±0.034	2.11±0.041
	Cutout	2.27±0.085	1.83±0.018	2.19±0.021	1.82±0.023
	AutoAug	1.91±0.059	1.61±0.024	1.80±0.044	1.35±0.056
CIFAR-10	Vanilla	3.87±0.167	3.00±0.059	3.60±0.197	2.75±0.040
	Cutout	3.38±0.081	2.67±0.043	2.83±0.102	2.27±0.034
	AutoAug	2.78±0.134	2.32±0.097	2.49±0.128	1.98±0.062
CIFAR-100	Vanilla	19.17±0.270	17.33±0.110	17.13±0.210	15.09±0.092
	Cutout	18.12±0.114	16.04±0.071	16.45±0.136	14.34±0.153
	AutoAug	17.79±0.185	14.95±0.088	15.43±0.269	13.36±0.245

Table 2: Mean and standard deviation of top-1 errors (%) on SVHN, CIFAR-10 and CIFAR-100 over 10 trials.

compared data augmentation techniques on SVHN, CIFAR-10 and CIFAR-100. The top-1 classification errors are shown in Table 2. The results suggest AMP’s regularization effect in the presence of advanced data augmentation techniques.

5.3. Calibration Effect

A well calibrated neural network is one in which the predicted softmax scores give better indicators of the actual likelihood of a correct prediction. Deep neural networks without any regularizer are prone to overconfidence on incorrect predictions, and a well calibrated network is required especially in some application areas like object detection [12, 36] and autonomous vehicle control [5, 30]. AMP chooses a flatter minimum of the empirical risk, which gives less confidence on possibly misclassified examples, ensuring the neural network to be better calibrated. On the contrary, the sharp minima given by ERM may contain incorrectly classified examples, which makes overconfident predictions on misclassified examples. In this section, we demonstrate that AMP can improve the calibration effect of neural networks.

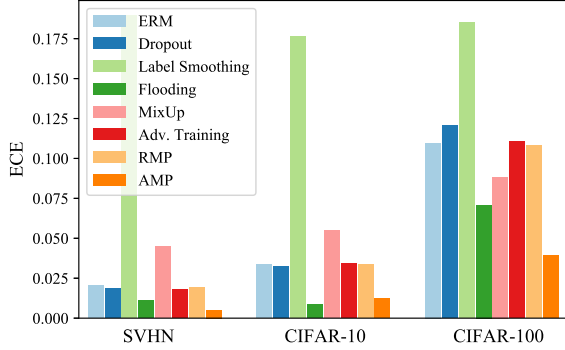


Figure 7: Expected calibration errors (ECEs) of AMP and other baseline methods. Results are averaged over 10 trials.

We adopt the measurement of calibration described in [15], namely, the Expected Calibration Error (ECE) (see Appendix D for definition). To evaluate the calibration effect of different regularization schemes, we compute ECEs of the pretrained PreActResNet18 models on the SVHN, CIFAR-10 and CIFAR-100 datasets. The results are shown in Figure 7. From the results, AMP achieves excellent calibration performance on various datasets. We can also find that Flooding outperforms other methods in calibration error on the CIFAR-10 dataset. Additionally, label smoothing degrades the calibration effect of neural networks, since it excessively biases the labels to the uniform distribution.

5.4. Influence of Perturbation

To investigate the relationship between the perturbation magnitudes and the geometry of the selected local minima, we compare the empirical risks of models trained with ERM and AMP. To clearly illustrate this, we adopt $\eta = 2$ and $N = 2$. Figure 8 shows the empirical risks on CIFAR-10 training and test set by varying the perturbation radius ϵ . It can be seen that $\mathcal{L}_{\text{ERM}}(\theta_{\text{AMP}}^*)$ on the training set tends to be high when radius is large. This indicates that the selected minimum has a smaller depth. Moreover, when evaluating on the test set, $\mathcal{L}_{\text{ERM}}(\theta_{\text{AMP}}^*)$ arrives at minimum when the perturbation radius is around 0.06. This suggests that an appropriate magnitude of perturbation regularizes networks efficiently, corresponds to the good properties of the selected minimum both in flatness and depth. Results on more datasets are given in Appendix E.

5.5. Computational Cost

AMP computes the adversarial perturbation of the model parameter in the training phase, such an operation arguably increases the computational cost. In addition to the gradient computation for updating the parameter, each stochastic gradient descent iteration requires multiple gradient computations to produce the adversarial perturbations. The compu-

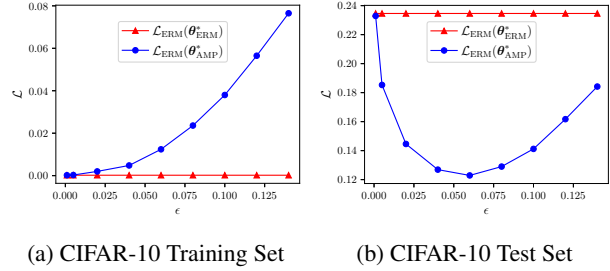


Figure 8: The empirical risks of the model trained with ERM (red) and AMP (blue) on the CIFAR-10 training set and test set with varying perturbation radius.

tational cost will significantly increase as the inner iteration number N grows. However, we find that $N = 1$ is sufficient to regularize the neural networks. Under this setting, we evaluate the practical computational cost of AMP compared to the ERM method. From our observation, AMP usually takes around $1.8\times$ that of ERM training. Therefore, the extra effort for adversarial perturbation is affordable.

6. Conclusion

Regularization is the main tool for deep learning practitioners to combat overfitting. In this work, we propose a novel regularization scheme, Adversarial Model Perturbation (AMP), built upon the understanding that flat local minima lead to better generalization. Unlike many other data-dependent regularization schemes, which are large of a heuristic nature, AMP has strong theoretical justifications under a certain approximating assumption. These justifications also allow us to predict its behaviour with respect to varying hyper-parameters. Our theoretical analysis and the regularization performance of AMP are confirmed through extensive experiments on image classification datasets. It is also observed that AMP helps the model to better calibrate itself. The outstanding performance of AMP arguably makes it into the current state of the art among all regularization schemes. The empirical validation of AMP presented in this paper appears to further confirm the connection between flat minima and generalization.

Acknowledgements. This work is supported partly by the National Key Research and Development Program of China (2016YFB1000103), by the National Natural Science Foundation of China (No. 61772059), by the Beijing Advanced Innovation Center for Big Data and Brain Computing, by the Fundamental Research Funds for the Central Universities, by the Beijing S&T Committee (No. Z191100008619007) and by the State Key Laboratory of Software Development Environment (No. SKLSDE-2020ZX-14). The authors specially thank Linfang Hou for helpful discussions.

References

- [1] Sanjeev Arora, Simon Du, Wei Hu, Zhiyuan Li, and Ruosong Wang. Fine-grained analysis of optimization and generalization for overparameterized two-layer neural networks. In *Proceedings of the 36th International Conference on Machine Learning, ICML*, pages 322–332, 2019.
- [2] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.
- [3] Mikhail Belkin, Daniel Hsu, Siyuan Ma, and Soumik Mandal. Reconciling modern machine-learning practice and the classical bias–variance trade-off. *Proceedings of the National Academy of Sciences*, 116(32):15849–15854, 2019.
- [4] Pratik Chaudhari, Anna Choromanska, Stefano Soatto, Yann LeCun, Carlo Baldassi, Christian Borgs, Jennifer Chayes, Levent Sagun, and Riccardo Zecchina. Entropy-sgd: Biasing gradient descent into wide valleys. In *5th International Conference on Learning Representations, ICLR*, 2017.
- [5] Chenyi Chen, Ari Seff, Alain Kornhauser, and Jianxiong Xiao. Deepdriving: Learning affordance for direct perception in autonomous driving. In *Proceedings of the IEEE International Conference on Computer Vision, ICCV*, pages 2722–2730, 2015.
- [6] Ekin D Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V Le. Autoaugment: Learning augmentation strategies from data. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, pages 113–123, 2019.
- [7] Terrance DeVries and Graham W Taylor. Improved regularization of convolutional neural networks with cutout. *arXiv preprint arXiv:1708.04552*, 2017.
- [8] Laurent Dinh, Razvan Pascanu, Samy Bengio, and Yoshua Bengio. Sharp minima can generalize for deep nets. In *Proceedings of the 34th International Conference on Machine Learning, ICML*, pages 1019–1028, 2017.
- [9] Pierre Foret, Ariel Kleiner, Hossein Mobahi, and Behnam Neyshabur. Sharpness-aware minimization for efficiently improving generalization. In *9th International Conference on Learning Representations, ICLR*, 2021.
- [10] Timur Garipov, Pavel Izmailov, Dmitrii Podoprikin, Dmitry P Vetrov, and Andrew G Wilson. Loss surfaces, mode connectivity, and fast ensembling of dnns. In *Advances in Neural Information Processing Systems*, pages 8789–8798, 2018.
- [11] Xavier Gastaldi. Shake-shake regularization of 3-branch residual networks. In *Workshop Track Proceedings in International Conference on Learning Representations, ICLR*, 2017.
- [12] Ross Girshick. Fast R-CNN. In *Proceedings of the IEEE International Conference on Computer Vision, ICCV*, pages 1440–1448, 2015.
- [13] Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In *3rd International Conference on Learning Representations, ICLR*, 2015.
- [14] Priya Goyal, Piotr Dollár, Ross Girshick, Pieter Noordhuis, Lukasz Wesolowski, Aapo Kyrola, Andrew Tulloch, Yangqing Jia, and Kaiming He. Accurate, large mini-batch SGD: Training imagenet in 1 hour. *arXiv preprint arXiv:1706.02677*, 2017.
- [15] Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. On calibration of modern neural networks. In *Proceedings of the 34th International Conference on Machine Learning, ICML*, pages 1321–1330, 2017.
- [16] Hongyu Guo, Yongyi Mao, and Richong Zhang. Mixup as locally linear out-of-manifold regularization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 3714–3722, 2019.
- [17] Dongyoon Han, Jihwan Kim, and Junmo Kim. Deep pyramidal residual networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, pages 5927–5935, 2017.
- [18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *European Conference on Computer Vision*, pages 630–645, 2016.
- [19] Sepp Hochreiter and Jürgen Schmidhuber. Flat minima. *Neural Computation*, 9(1):1–42, 1997.
- [20] Elad Hoffer, Itay Hubara, and Daniel Soudry. Train longer, generalize better: closing the generalization gap in large batch training of neural networks. In *Advances in Neural Information Processing Systems*, pages 1731–1741, 2017.
- [21] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *Proceedings of the 32nd International Conference on Machine Learning, ICML*, pages 448–456, 2015.
- [22] Takashi Ishida, Ikko Yamane, Tomoya Sakai, Gang Niu, and Masashi Sugiyama. Do we need zero training loss after achieving zero training error? In *Proceedings of the 37th International Conference on Machine Learning, ICML*, 2020.
- [23] Pavel Izmailov, Dmitrii Podoprikin, Timur Garipov, Dmitry Vetrov, and Andrew Gordon Wilson. Averaging weights leads to wider optima and better generalization. In *Proceedings of the Thirty-Fourth Conference on Uncertainty in Artificial Intelligence, UAI*, pages 876–885, 2018.
- [24] Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. In *Advances in Neural Information Processing Systems*, pages 8571–8580, 2018.
- [25] Stanislaw Jastrzebski, Zachary Kenton, Devansh Arpit, Nicolas Ballas, Asja Fischer, Yoshua Bengio, and Amos Storkey. Three factors influencing minima in SGD. *arXiv preprint arXiv:1711.04623*, 2017.
- [26] Nitish Shirish Keskar, Dheevatsa Mudigere, Jorge Nocedal, Mikhail Smelyanskiy, and Ping Tak Peter Tang. On large-batch training for deep learning: Generalization gap and sharp minima. In *5th International Conference on Learning Representations, ICLR*, 2017.
- [27] Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009.
- [28] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems*, pages 1097–1105, 2012.
- [29] Anders Krogh and John A Hertz. A simple weight decay can improve generalization. In *Advances in Neural Information Processing Systems*, pages 950–957, 1992.

- [30] Jesse Levinson, Jake Askeland, Jan Becker, Jennifer Dolson, David Held, Soeren Kammel, J Zico Kolter, Dirk Langer, Oliver Pink, Vaughan Pratt, et al. Towards fully autonomous driving: Systems and algorithms. In *2011 IEEE Intelligent Vehicles Symposium (IV)*, pages 163–168. IEEE, 2011.
- [31] Hao Li, Zheng Xu, Gavin Taylor, Christoph Studer, and Tom Goldstein. Visualizing the loss landscape of neural nets. In *Advances in Neural Information Processing Systems*, pages 6389–6399, 2018.
- [32] Sungbin Lim, Ildoo Kim, Taesup Kim, Chiheon Kim, and Sungwoong Kim. Fast autoaugment. In *Advances in Neural Information Processing Systems*, pages 6665–6675, 2019.
- [33] Takeru Miyato, Andrew M. Dai, and Ian J. Goodfellow. Adversarial training methods for semi-supervised text classification. In *5th International Conference on Learning Representations, ICLR, 2017*.
- [34] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Bo Wu, and Andrew Y Ng. Reading digits in natural images with unsupervised feature learning. In *NeurIPS Workshop on Deep Learning and Unsupervised Feature Learning*, 2011.
- [35] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*, pages 8026–8037, 2019.
- [36] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. In *Advances in Neural Information Processing Systems*, pages 91–99, 2015.
- [37] Jorma Rissanen. Modeling by shortest data description. *Automatica*, 14(5):465–471, 1978.
- [38] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *3rd International Conference on Learning Representations, ICLR, 2015*.
- [39] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958, 2014.
- [40] Masashi Sugiyama. *Introduction to statistical machine learning*. Morgan Kaufmann, 2015.
- [41] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, pages 2818–2826, 2016.
- [42] Dimitris Tsipras, Shibani Santurkar, Logan Engstrom, Alexander Turner, and Aleksander Madry. Robustness may be at odds with accuracy. In *7th International Conference on Learning Representations, ICLR, 2019*.
- [43] Vladimir Vapnik. *Statistical learning theory*. Wiley, 1998.
- [44] Colin Wei, Sham M. Kakade, and Tengyu Ma. The implicit and explicit regularization effects of dropout. In *Proceedings of the 37th International Conference on Machine Learning, ICML, 2020*.
- [45] Yoshihiro Yamada, Masakazu Iwamura, Takuya Akiba, and Koichi Kise. Shakedrop regularization for deep residual learning. *IEEE Access*, 7:186126–186136, 2019.
- [46] Zitong Yang, Yaodong Yu, Chong You, Jacob Steinhardt, and Yi Ma. Rethinking bias-variance trade-off for generalization of neural networks. In *Proceedings of the 37th International Conference on Machine Learning, ICML, 2020*.
- [47] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. In *Proceedings of the British Machine Vision Conference, BMVC, 2016*.
- [48] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. In *5th International Conference on Learning Representations, ICLR, 2017*.
- [49] Hongyi Zhang, Moustapha Cisse, Yann N. Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. In *6th International Conference on Learning Representations, ICLR, 2018*.