
Optimal Transport Model Distributional Robustness

Van-Anh Nguyen¹ Trung Le¹ Anh Tuan Bui¹ Thanh-Toan Do¹
Dinh Phung^{1,2}

¹Department of Data Science and AI, Monash University, Australia

²VinAI, Vietnam

{van-anh.nguyen, trunglm, tuan.bui, toan.do, dinh.phung}@monash.edu

Abstract

Distributional robustness is a promising framework for training deep learning models that are less vulnerable to adversarial examples and data distribution shifts. Previous works have mainly focused on exploiting distributional robustness in the data space. In this work, we explore an optimal transport-based distributional robustness framework in model spaces. Specifically, we examine a model distribution within a Wasserstein ball centered on a given model distribution that maximizes the loss. We have developed theories that enable us to learn the optimal robust center model distribution. Interestingly, our developed theories allow us to flexibly incorporate the concept of sharpness awareness into training, whether it's a single model, ensemble models, or Bayesian Neural Networks, by considering specific forms of the center model distribution. These forms include a Dirac delta distribution over a single model, a uniform distribution over several models, and a general Bayesian Neural Network. Furthermore, we demonstrate that Sharpness-Aware Minimization (SAM) is a specific case of our framework when using a Dirac delta distribution over a single model, while our framework can be seen as a probabilistic extension of SAM. To validate the effectiveness of our framework in the aforementioned settings, we conducted extensive experiments, and the results reveal remarkable improvements compared to the baselines.

1 Introduction

Distributional robustness (DR) is a promising framework for learning and decision-making under uncertainty, which has gained increasing attention in recent years [4, 15, 16, 6]. The primary objective of DR is to identify the worst-case data distribution within the vicinity of the ground-truth data distribution, thereby challenging the model's robustness to distributional shifts. DR has been widely applied to various fields, including semi-supervised learning [5, 11, 66], transfer learning and domain adaptation [39, 16, 69, 50, 51, 38, 52, 57], domain generalization [57, 70], and improving model robustness [57, 8, 61]. Although the principle of DR can be applied to either data space or model space, the majority of previous works on DR have primarily concentrated on exploring its applications in data space.

Sharpness-aware minimization (SAM) [18] has emerged as an effective technique for enhancing the generalization ability of deep learning models. SAM aims to find a perturbed model within the vicinity of a current model that maximizes the loss over a training set. The success of SAM and its variants [36, 33, 62] has inspired further investigation into its formulation and behavior, as evidenced by recent works such as [31, 45, 2]. While [30] empirically studied the difference in sharpness obtained by SAM [18] and SWA [26], and [46] demonstrated that SAM is an optimal Bayes relaxation of standard Bayesian inference with a normal posterior, none of the existing works have explored the connection between SAM and distributional robustness.

In this work, we study the theoretical connection between distributional robustness in model space and sharpness-aware minimization (SAM), as they share a conceptual similarity. We examine Optimal Transport-based distributional robustness in model space by considering a Wasserstein ball centered around a model distribution and searching for the worst-case distribution that maximizes the empirical loss on a training set. By controlling the worst-case performance, it is expected to have a smaller generalization error, as demonstrated by a smaller empirical loss and sharpness. We then develop rigorous theories that suggest us the strategy to learn the center model distribution. We demonstrate the effectiveness of our framework by devising the practical methods for three cases of model distribution: (i) a Dirac delta distribution over a single model, (ii) a uniform distribution over several models, and (iii) a general model distribution (i.e., a Bayesian Neural Network [48, 56]). Furthermore, we show that SAM is a specific case of our framework when using a Dirac delta distribution over a single model, and our framework can be regarded as a probabilistic extension of SAM.

In summary, our contributions in this work are as follows:

- We propose a framework for enhancing model generalization by introducing an Optimal Transport (OT)-based model distribution robustness approach (named OT-MDR). To the best of our knowledge, this is the first work that considers distributional robustness within the model space.
- We have devised three practical methods tailored to different types of model distributions. Through extensive experiments, we demonstrate that our practical methods effectively improve the generalization of models, resulting in higher natural accuracy and better uncertainty estimation.
- Our theoretical findings reveal that our framework can be considered as a probabilistic extension of the widely-used sharpness-aware minimization (SAM) technique. In fact, SAM can be viewed as a specific case within our comprehensive framework. This observation not only explains the outperformance of our practical method over SAM but also sheds light on future research directions for improving model generalization through the perspective of distributional robustness.

2 Related Work

2.1 Distributional Robustness

Distributional Robustness (DR) is a promising framework for enhancing machine learning models in terms of robustness and generalization. The main idea behind DR is to identify the most challenging distribution that is in close proximity to a given distribution and then evaluate the model's performance on this distribution. The proximity between two distributions can be measured using either a f -divergence [4, 15, 16, 43, 47] or Wasserstein distance [6, 21, 35, 44, 60]. In addition, some studies [7, 61] have developed a dual form for DR, which allows it to be integrated into the training process of deep learning models. DR has been applied to a wide range of domains, including semi-supervised learning [5, 11, 66], domain adaptation [57, 50], domain generalization [57, 70], and improving model robustness [57, 8, 61].

2.2 Flat Minima

The generalization ability of neural networks can be improved by finding flat minimizers, which allow models to find wider local minima and increase their robustness to shifts between training and test sets [29, 55, 17]. The relationship between the width of minima and generalization ability has been studied theoretically and empirically in many works, including [24, 49, 13, 19, 53]. Various methods have been proposed to seek flat minima, such as those presented in [54, 10, 32, 27, 18]. Studies such as [32, 28, 64] have investigated the effects of different training factors, including batch-size, learning rate, gradient covariance, and dropout, on the flatness of found minima. In addition, some approaches introduce regularization terms into the loss function to pursue wide local minima, such as low-entropy penalties for softmax outputs [54] and distillation losses [68, 67, 10].

SAM [18] is a method that seeks flat regions by explicitly minimizing the worst-case loss around the current model. It has received significant attention recently for its effectiveness and scalability compared to previous methods. SAM has been successfully applied to various tasks and domains, in-

cluding meta-learning bi-level optimization [1], federated learning [59], vision models [12], language models [3], domain generalization [9], Bayesian Neural Networks [53], and multi-task learning [58]. For instance, SAM has demonstrated its capability to improve meta-learning bi-level optimization in [1], while in federated learning, SAM achieved tighter convergence rates than existing works and proposed a generalization bound for the global model [59]. Additionally, SAM has shown its ability to generalize well across different applications such as vision models, language models, domain generalization, and multi-task learning. Some recent works have attempted to enhance SAM’s performance by exploiting its geometry [36, 33], minimizing surrogate gap [71], and speeding up its training time [14, 41]. Moreover, [30] empirically studied the difference in sharpness obtained by SAM and SWA [26], while [46] demonstrated that SAM is an optimal Bayes relaxation of the standard Bayesian inference with a normal posterior.

3 Distributional Robustness

In this section, we present the background on the OT-based distributional robustness that serves our theory development in the sequel. Distributional robustness (DR) is an emerging framework for learning and decision-making under uncertainty, which seeks the worst-case expected loss among a ball of distributions, containing all distributions that are close to the empirical distribution [20].

Here we consider a generic Polish space S endowed with a distribution \mathbb{Q} . Let $f : S \rightarrow \mathbb{R}$ be a real-valued (risk) function and $c : S \times S \rightarrow \mathbb{R}_+$ be a cost function. Distributional robustness setting aims to find the distribution $\tilde{\mathbb{Q}}$ in the vicinity of \mathbb{Q} and maximizes the risk in the \mathbb{E} form [61, 7]:

$$\max_{\tilde{\mathbb{Q}}: \mathcal{W}_c(\tilde{\mathbb{Q}}, \mathbb{Q}) < \epsilon} \mathbb{E}_{\tilde{\mathbb{Q}}} [f(z)], \quad (1)$$

where $\epsilon > 0$ and \mathcal{W}_c denotes the optimal transport (OT) or a Wasserstein distance [63] for a metric c , which is defined as:

$$\mathcal{W}_c(\tilde{\mathbb{Q}}, \mathbb{Q}) := \inf_{\gamma \in \Gamma(\tilde{\mathbb{Q}}, \mathbb{Q})} \int cd\gamma, \quad (2)$$

where $\Gamma(\tilde{\mathbb{Q}}, \mathbb{Q})$ is the set of couplings whose marginals are $\tilde{\mathbb{Q}}$ and \mathbb{Q} .

With the assumption that $f \in L^1(\mathbb{Q})$ is upper semi-continuous and the cost c is a non-negative lower semi-continuous satisfying $c(z, z') = 0$ iff $z = z'$, [7] shows that the *dual* form for Eq. (1) is:

$$\min_{\lambda \geq 0} \left\{ \lambda \epsilon + \mathbb{E}_{z \sim \mathbb{Q}} [\max_{z'} \{f(z') - \lambda c(z', z)\}] \right\}. \quad (3)$$

[61] further employs a Lagrangian for Wasserstein-based uncertainty sets to arrive at a relaxed version with $\lambda \geq 0$:

$$\max_{\tilde{\mathbb{Q}}} \left\{ \mathbb{E}_{\tilde{\mathbb{Q}}} [f(z)] - \lambda \mathcal{W}_c(\tilde{\mathbb{Q}}, \mathbb{Q}) \right\} = \mathbb{E}_{z \sim \mathbb{Q}} [\max_{z'} \{f(z') - \lambda c(z', z)\}]. \quad (4)$$

4 Proposed Framework

4.1 OT based Sharpness-aware Distribution Robustness

Given a family of deep nets f_θ where $\theta \in \Theta$, let \mathbb{Q}_ϕ with the density function $q_\phi(\theta)$ where $\phi \in \Phi$ be a family of distributions over the parameter space Θ . To improve the generalization ability of the optimal \mathbb{Q}_{ϕ^*} , we propose the following distributional robustness formulation on the model space:

$$\min_{\phi \in \Phi} \max_{\tilde{\mathbb{Q}}: \mathcal{W}_d(\tilde{\mathbb{Q}}, \mathbb{Q}_\phi) \leq \rho} \mathcal{L}_S(\tilde{\mathbb{Q}}), \quad (5)$$

where $S = \{(x_1, y_1), \dots, (x_N, y_N)\}$ is a training set, $d(\theta, \tilde{\theta}) = \|\theta - \tilde{\theta}\|_2^p$ ($p \geq 1$) is a distance on the model space, and we have defined

$$\mathcal{L}_S(\tilde{\mathbb{Q}}) = \mathbb{E}_{\theta \sim \tilde{\mathbb{Q}}} [\mathcal{L}_S(\theta)] = \mathbb{E}_{\theta \sim \tilde{\mathbb{Q}}} \left[\frac{1}{N} \sum_{n=1}^N \ell(f_\theta(x_n), y_n) \right]$$

with the loss function ℓ .

The OP in (5) seeks the most challenging model distribution $\tilde{\mathbb{Q}}$ in the WS ball around \mathbb{Q}_ϕ and then finds $\mathbb{Q}_\phi, \phi \in \Phi$ which minimizes the worst loss. To derive a solution for the OP in (5), we define

$$\Gamma_{\rho, \phi} = \left\{ \gamma : \gamma \in \cup_{\tilde{\mathbb{Q}}} \Gamma(\tilde{\mathbb{Q}}, \mathbb{Q}_\phi), \mathbb{E}_{(\theta, \tilde{\theta}) \sim \gamma} \left[d(\theta, \tilde{\theta}) \right]^{1/p} \leq \rho \right\}.$$

Moreover, the following theorem characterizes the solution of the OP in (5).

Theorem 4.1. *The OP in (5) is equivalent to the following OP:*

$$\min_{\phi \in \Phi} \max_{\gamma \in \Gamma_{\rho, \phi}} \mathcal{L}_S(\gamma), \quad (6)$$

where $\mathcal{L}_S(\gamma) = \mathbb{E}_{(\theta, \tilde{\theta}) \sim \gamma} \left[\frac{1}{N} \sum_{n=1}^N \ell(f_{\tilde{\theta}}(x_n), y_n) \right]$.

We now need to solve the OP in (6). To make it solvable, we add the entropic regularization term as

$$\min_{\phi \in \Phi} \max_{\gamma \in \Gamma_{\rho, \phi}} \left\{ \mathcal{L}_S(\gamma) + \frac{1}{\lambda} \mathbb{H}(\gamma) \right\}, \quad (7)$$

where $\mathbb{H}(\gamma)$ returns the entropy of the distribution γ with the trade-off parameter $1/\lambda$. We note that when λ approaches $+\infty$, the OP in (7) becomes equivalent to the OP in (6). The following theorem indicates the solution of the OP in (7).

Theorem 4.2. *When $p = +\infty$, the inner max in the OP in (7) has the solution which is a distribution with the density function*

$$\gamma^*(\theta, \tilde{\theta}) = q_\phi(\theta) \gamma^*(\tilde{\theta} | \theta),$$

where $\gamma^*(\tilde{\theta} | \theta) = \frac{\exp\{\lambda \mathcal{L}_S(\tilde{\theta})\}}{\int_{B_\rho(\theta)} \exp\{\lambda \mathcal{L}_S(\theta')\} d\theta'}$, $q_\phi(\theta)$ is the density function of the distribution \mathbb{Q}_ϕ , and $B_\rho(\theta) = \{\theta' : \|\theta' - \theta\|_2 \leq \rho\}$ is the ρ -ball around θ .

Referring to Theorem 4.2, the OP in (7) hence becomes:

$$\min_{\phi \in \Phi} \mathbb{E}_{\theta \sim \mathbb{Q}_\phi, \tilde{\theta} \sim \gamma^*(\tilde{\theta} | \theta)} \left[\mathcal{L}_S(\tilde{\theta}) \right]. \quad (8)$$

The OP in (8) implies that given a model distribution \mathbb{Q}_ϕ , we sample models $\theta \sim \mathbb{Q}_\phi$. For each individual model θ , we further sample the particle models $\tilde{\theta} \sim \gamma^*(\tilde{\theta} | \theta)$ where $\gamma^*(\tilde{\theta} | \theta) \propto \exp\{\lambda \mathcal{L}_S(\tilde{\theta})\}$. Subsequently, we update \mathbb{Q}_ϕ to minimize the average of $\mathcal{L}_S(\tilde{\theta})$. It is worth noting that the particle models $\tilde{\theta} \sim \gamma^*(\tilde{\theta} | \theta) \propto \exp\{\lambda \mathcal{L}_S(\tilde{\theta})\}$ seek the modes of the distribution $\gamma^*(\tilde{\theta} | \theta)$, aiming to obtain high and highest likelihoods (i.e., $\exp\{\lambda \mathcal{L}_S(\tilde{\theta})\}$). Additionally, in the implementation, we employ stochastic gradient Langevin dynamics (SGLD) [65] to sample the particle models $\tilde{\theta}$.

In what follows, we consider three cases where \mathbb{Q}_ϕ is (i) a Dirac delta distribution over a single model, (ii) a uniform distribution over several models, and (iii) a general distribution over the model space (i.e., a Bayesian Neural Network (BNN)) and further devise the practical methods for them.

4.2 Practical Methods

4.2.1 Single-Model OT-based Distributional Robustness

We first examine the case where \mathbb{Q}_ϕ is a Dirac delta distribution over a single model θ , i.e., $\mathbb{Q}_\phi = \delta_\theta$ where δ is the Dirac delta distribution. Given the current model θ , we sample K particles $\tilde{\theta}_{1:K}$ using SGLD with only two-step sampling. To diversify the particles, in addition to adding a small Gaussian

noise to each particle, given a mini-batch B , for each particle $\tilde{\theta}_k$, we randomly split $B = [B_k^1, B_k^2]$ into two equal halves and update the particle models as follows:

$$\begin{aligned}\tilde{\theta}_k^1 &= \theta + \rho \frac{\nabla_{\theta} \mathcal{L}_{B_k^1}(\theta)}{\|\nabla_{\theta} \mathcal{L}_{B_k^1}(\theta)\|_2} + \epsilon_k^1, \\ \tilde{\theta}_k^2 &= \tilde{\theta}_k^1 + \rho \frac{\nabla_{\theta} \mathcal{L}_{B_k^2}(\tilde{\theta}_k^1)}{\|\nabla_{\theta} \mathcal{L}_{B_k^2}(\tilde{\theta}_k^1)\|_2} + \epsilon_k^2,\end{aligned}\tag{9}$$

where \mathbb{I} is the identity matrix and $\epsilon_k^1, \epsilon_k^2 \sim \mathcal{N}(0, \rho \mathbb{I})$.

Furthermore, we base on the particle models to update the next model as follows:

$$\theta = \theta - \frac{\eta}{K} \sum_{k=1}^K \nabla_{\theta} \mathcal{L}_B(\tilde{\theta}_k^2),\tag{10}$$

where $\eta > 0$ is a learning rate.

It is worth noting that in the update formula (9), we use different random splits $B = [B_k^1, B_k^2]$, $k = 1, \dots, K$ to encourage the diversity of the particles $\tilde{\theta}_{1:K}$. Moreover, we can benefit from the parallel computing to estimate these particles in parallel, which costs $|B|$ (i.e., the batch size) gradient operations. Additionally, similar to SAM [18], when computing the gradient $\nabla_{\theta} \mathcal{L}_{B_k^2}(\tilde{\theta}_k^1)$, we set the corresponding Hessian matrix to the identity one.

Particularly, the gradient $\nabla_{\theta} \mathcal{L}_{B_k^2}(\tilde{\theta}_k^1)$ is evaluated on the second half of the current batch so that the entire batch is used for the update. Again, we can take advantage of the parallel computing to evaluate $\nabla_{\theta} \mathcal{L}_{B_k^2}(\tilde{\theta}_k^1)$, $k = 1, \dots, K$ all at once, which costs $|B|$ (i.e., the batch size) gradient operations. Eventually, with the aid of the parallel computing, the total gradient operations in our approach is $2|B|$, which is similar to SAM.

4.2.2 Ensemble OT-based Distributional Robustness

We now examine the case where \mathbb{Q}_{ϕ} is a uniform distribution over several models, i.e., $\mathbb{Q}_{\phi} = \frac{1}{M} \sum_{m=1}^M \delta_{\theta_m}$. In the context of ensemble learning, for each base learner θ_m , $m = 1, \dots, M$, we seek K particle models $\tilde{\theta}_{mk}$, $k = 1, \dots, K$ as in the case of single model.

$$\begin{aligned}\tilde{\theta}_{mk}^1 &= \theta + \rho \frac{\nabla_{\theta} \mathcal{L}_{B_{mk}^1}(\theta_m)}{\|\nabla_{\theta} \mathcal{L}_{B_{mk}^1}(\theta_m)\|_2} + \epsilon_{mk}^1, \\ \tilde{\theta}_{mk}^2 &= \tilde{\theta}_{mk}^1 + \rho \frac{\nabla_{\theta} \mathcal{L}_{B_{mk}^2}(\tilde{\theta}_{mk}^1)}{\|\nabla_{\theta} \mathcal{L}_{B_{mk}^2}(\tilde{\theta}_{mk}^1)\|_2} + \epsilon_{mk}^2,\end{aligned}\tag{11}$$

where \mathbb{I} is the identity matrix and $\epsilon_{mk}^1, \epsilon_{mk}^2 \sim \mathcal{N}(0, \rho \mathbb{I})$.

Furthermore, we base on the particles to update the next base learners as follows:

$$\theta_m = \theta_m - \frac{\eta}{K} \sum_{k=1}^K \nabla_{\theta} \mathcal{L}_B(\tilde{\theta}_{mk}^2).\tag{12}$$

It is worth noting that the random splits $B = [B_{mk}^1, B_{mk}^2]$, $m = 1, \dots, M$, $k = 1, \dots, K$ of the current batch and the added Gaussian noise constitute the diversity of the base learners $\theta_{1:M}$. In our developed ensemble model, we do not invoke any term to explicitly encourage the model diversity.

4.2.3 BNN OT-based Distributional Robustness

We finally examine the case where \mathbb{Q}_{ϕ} is an approximate posterior or a BNN. To simplify the context, we assume that \mathbb{Q}_{ϕ} consists of Gaussian distributions $\mathcal{N}(\mu_l, \text{diag}(\sigma_l^2))$, $l = 1, \dots, L$ over

the weight matrices $\theta = W_{1:L}$ ¹. Given $\theta = W_{1:L} \sim \mathbb{Q}_\phi$, the reparameterization trick reads $W_l = \mu_l + \text{diag}(\sigma_l)\kappa_l$ with $\kappa_l \sim \mathcal{N}(0, \mathbb{I})$. We next sample K particle models $\tilde{\theta}_k = [\tilde{W}_{lk}]_{lk}$, $l = 1, \dots, L$ and $k = 1, \dots, K$ from $\gamma^*(\tilde{\theta} | \theta)$ as in Theorem 4.2. To sample $\tilde{\theta}_k$ in the ball around θ , for each layer l , we indeed sample $\tilde{\mu}_{lk}$ in the ball around μ_l and then form $\tilde{W}_{lk} = \tilde{\mu}_{lk} + \text{diag}(\sigma_l)\kappa_l$. We randomly split $B = [B_k^1, B_k^2]$, $k = 1, \dots, K$ and update the approximate Gaussian distribution as

$$\begin{aligned}\tilde{\mu}_{lk}^1 &= \mu_l + \rho \frac{\nabla_{\mu_l} \mathcal{L}_{B_k^1}([\mu_l + \text{diag}(\sigma_l)\kappa_l]_l)}{\|\nabla_{\mu_l} \mathcal{L}_{B_k^1}([\mu_l + \text{diag}(\sigma_l)\kappa_l]_l)\|_2} + \epsilon_k^1, \\ \tilde{\mu}_{lk}^2 &= \tilde{\mu}_{lk}^1 + \rho \frac{\nabla_{\mu_l} \mathcal{L}_{B_k^2}([\tilde{\mu}_{lk}^1 + \text{diag}(\sigma_l)\kappa_l]_l)}{\|\nabla_{\mu_l} \mathcal{L}_{B_k^2}([\tilde{\mu}_{lk}^1 + \text{diag}(\sigma_l)\kappa_l]_l)\|_2} + \epsilon_k^2, \\ \tilde{\theta}_k &= [\tilde{\mu}_{lk}^2 + \text{diag}(\sigma_l)\kappa_l]_l, k = 1, \dots, K, \\ \mu &= \mu - \frac{\eta}{K} \sum_{k=1}^K \nabla_{\mu} \mathcal{L}_B(\tilde{\theta}_k) \text{ and } \sigma = \sigma - \eta \nabla_{\sigma} \mathcal{L}_B(\theta),\end{aligned}$$

where $\epsilon_k^1, \epsilon_k^2 \sim \mathcal{N}(0, \rho \mathbb{I})$.

4.3 Connection of SAM and OT-based Model Distribution Robustness

In what follows, we show a connection between OT-based model distributional robustness and SAM. Specifically, we prove that SAM is a specific case of OT-based distributional robustness on the model space with a particular distance metric. To depart, we first recap the formulation of the OT-based distributional robustness on the model space in (5):

$$\min_{\phi \in \Phi} \max_{\tilde{\mathbb{Q}}: \mathcal{W}_d(\tilde{\mathbb{Q}}, \mathbb{Q}_\phi) \leq \rho} \mathcal{L}_S(\tilde{\mathbb{Q}}) = \min_{\phi \in \Phi} \max_{\tilde{\mathbb{Q}}: \mathcal{W}_d(\tilde{\mathbb{Q}}, \mathbb{Q}_\phi) \leq \rho} \mathbb{E}_{\theta \sim \tilde{\mathbb{Q}}} [\mathcal{L}_S(\theta)]. \quad (13)$$

By linking to the dual form in (3), we reach the following equivalent OP:

$$\min_{\phi \in \Phi} \min_{\lambda > 0} \left\{ \lambda \rho + \mathbb{E}_{\theta \sim \mathbb{Q}_\phi} \left[\max_{\tilde{\theta}} \left\{ \mathcal{L}_S(\tilde{\theta}) - \lambda d(\tilde{\theta}, \theta) \right\} \right] \right\}. \quad (14)$$

Considering the simple case wherein $\mathbb{Q}_\phi = \delta_\theta$ is a Dirac delta distribution. The OPs in (14) equivalently entails

$$\min_{\theta} \min_{\lambda > 0} \left\{ \lambda \rho + \max_{\tilde{\theta}} \left\{ \mathcal{L}_S(\tilde{\theta}) - \lambda d(\tilde{\theta}, \theta) \right\} \right\}. \quad (15)$$

The optimization problem in (15) can be viewed as a probabilistic extension of SAM. Specifically, for each $\theta \sim \mathbb{Q}_\phi$, the inner max: $\max_{\tilde{\theta}} \left\{ \mathcal{L}_S(\tilde{\theta}) - \lambda d(\tilde{\theta}, \theta) \right\}$ seeks a model $\tilde{\theta}$ maximizing the loss $\mathcal{L}_S(\tilde{\theta})$ on a soft ball around θ controlled by $\lambda d(\tilde{\theta}, \theta)$. Particularly, a higher value of λ seeks the optimal $\tilde{\theta}$ in a smaller ball. Moreover, in the outer min, the term $\lambda \rho$ trades off between the value of λ and the radius of the soft ball, aiming to find out an optimal λ^* and the optimal model $\tilde{\theta}^*$ maximizing the loss function over an appropriate soft ball. Here we note that SAM also seeks maximizing the loss function but over the ball with the radius ρ around the model θ .

Interestingly, by appointing a particular distance metric between two models, we can exactly recover the SAM formulation as shown in the following theorem.

Theorem 4.3. *With the distance metric d defined as*

$$d(\theta, \tilde{\theta}) = \begin{cases} \|\tilde{\theta} - \theta\|_2 & \|\tilde{\theta} - \theta\|_2 \leq \rho \\ +\infty & \text{otherwise} \end{cases}, \quad (16)$$

the OPs in (13), (14) with $\mathbb{Q}_\phi = \delta_\theta$, and (15) equivalently reduce to the OP of SAM as

$$\min_{\theta} \max_{\tilde{\theta}: \|\tilde{\theta} - \theta\|_2 \leq \rho} \mathcal{L}_S(\tilde{\theta}).$$

¹For simplicity, we absorb the biases to the weight matrices

Table 1: Classification accuracy on the CIFAR datasets of the single model setting with *one particle*. All experiments are trained three times with different random seeds.

Dataset	Method	WideResnet28x10	Pyramid101	Densenet121
CIFAR-10	SAM	96.72 ± 0.007	96.20 ± 0.134	91.16 ± 0.240
	OT-MDR (Ours)	96.97 ± 0.009	96.61 ± 0.063	91.44 ± 0.113
CIFAR-100	SAM	82.69 ± 0.035	81.26 ± 0.636	68.09 ± 0.403
	OT-MDR (Ours)	84.14 ± 0.172	82.28 ± 0.183	69.84 ± 0.176

Table 2: Classification score on Resnet18. The results of baselines are taken from [46]

Method	CIFAR-10		CIFAR-100	
	ACC ↑	AUROC ↑	ACC ↑	AUROC ↑
SGD	94.76 ± 0.11	0.926 ± 0.006	76.54 ± 0.26	0.869 ± 0.003
SAM	95.72 ± 0.14	0.949 ± 0.003	78.74 ± 0.19	0.887 ± 0.003
bSAM	96.15 ± 0.08	0.954 ± 0.001	80.22 ± 0.28	0.892 ± 0.003
OT-MDR (Ours)	96.59 ± 0.07	0.992 ± 0.004	81.23 ± 0.13	0.991 ± 0.001

Theorem 4.3 suggests that the OT-based model distributional robustness on the model space is a probabilistic relaxation of SAM in general, while SAM is also a specific crisp case of the OT-based model distributional robustness on the model space. This connection between the OT-based model distributional robustness and SAM is intriguing and may open doors to propose other sharpness-aware training approaches and leverage adversarial training [23, 42] to improve model robustness. Moreover, [46] has shown a connection between SAM and the standard Bayesian inference with a normal posterior. Using the Gaussian approximate posterior $N(\omega, \nu I)$, it is demonstrated that the maximum of the likelihood loss $L(q_\mu)$ (i.e., μ is the expectation parameter of the Gaussian $N(\omega, \nu I)$) can be lower-bounded by a relaxed-Bayes objective which is relevant to the SAM loss. Our findings are supplementary but independent of the above finding. Furthermore, by linking SAM to the OT-based model distributional robustness on the model space, we can expect to leverage the rich body theory of distributional robustness for new discoveries about improving the model generalization ability.

5 Experiments

In this section, we present the results of various experiments² to evaluate the effectiveness of our proposed method in achieving distribution robustness. These experiments are conducted in three main settings: a single model, ensemble models, and Bayesian Neural Networks. To ensure the reliability and generalizability of our findings, we employ multiple architectures and evaluate their performance using the CIFAR-10 and CIFAR-100 datasets. For each experiment, we report specific metrics that capture the performance of each model in its respective setting.

5.1 Experiments on a Single Model

To evaluate the performance of our proposed method for one particle training, we conducted experiments using three different architectures: WideResNet28x10, Pyramid101, and Densenet121. We compared our approach’s results against models trained with SAM optimizer as our baseline. For consistency with the original SAM paper, we adopted their setting, using $\rho = 0.05$ for CIFAR-10 experiments and $\rho = 0.1$ for CIFAR-100 and report the result in Table 1. In our OT-MDR method, we chose different values of ρ for each half of the mini-batch B , and denoted ρ_1 for B_1 and ρ_2 for B_2 , where $\rho_2 = 2\rho_1$ to be simplified (this simplified setting is used in all experiments). To ensure a fair comparison with SAM, we also set $\rho_1 = 0.05$ for CIFAR-10 experiments and $\rho_1 = 0.1$ for CIFAR-100. Our approach outperformed the baseline with significant gaps, as indicated in Table 1. On average, our method achieved a 0.73% improvement on CIFAR-10 and a 1.68% improvement on CIFAR-100 compared to SAM. These results demonstrate the effectiveness of our proposed method for achieving higher accuracy in one-particle model training.

²The implementation is provided in https://github.com/anh-ntv/OT_MDR.git

Table 3: Evaluation of the ensemble **Accuracy** (%) on the CIFAR-10/100 datasets. We reproduce all baselines with the same hyperparameter for a fair comparison.

Method	CIFAR-10					CIFAR-100				
	ACC \uparrow	Brier \downarrow	NLL \downarrow	ECE \downarrow	AAC \downarrow	ACC \uparrow	Brier \downarrow	NLL \downarrow	ECE \downarrow	AAC \downarrow
Ensemble of five Resnet10 models										
Deep Ensemble	92.7	0.091	0.272	0.072	0.108	73.7	0.329	0.87	0.145	0.162
Fast Geometric	92.5	0.251	0.531	0.121	0.144	63.2	0.606	1.723	0.149	0.162
Snapshot	93.6	0.083	0.249	0.065	0.107	72.8	0.338	0.929	0.153	0.338
EDST	92.0	0.122	0.301	0.078	0.112	68.4	0.427	1.151	0.155	0.427
DST	93.2	0.102	0.261	0.067	0.108	70.8	0.396	1.076	0.150	0.396
SGD	95.1	0.078	0.264	-	0.108	75.9	0.346	1.001	-	0.346
SAM	95.4	0.073	0.268	0.050	0.107	77.7	0.321	0.892	0.136	0.321
OT-MDR (Ours)	95.4	0.069	0.145	0.021	0.004	79.1	0.059	0.745	0.043	0.054
Ensemble of three Resnet18 models										
Deep Ensemble	93.7	0.079	0.273	0.064	0.107	75.4	0.308	0.822	0.14	0.155
Fast Geometric	93.3	0.087	0.261	0.068	0.108	72.3	0.344	0.95	0.15	0.169
Snapshot	94.8	0.071	0.27	0.054	0.108	75.7	0.311	0.903	0.147	0.153
EDST	92.8	0.113	0.281	0.074	0.11	69.6	0.412	1.123	0.151	0.197
DST	94.7	0.083	0.253	0.057	0.107	70.4	0.405	1.153	0.155	0.194
SGD	95.2	0.076	0.282	-	0.108	78.9	0.304	0.919	-	0.156
SAM	95.8	0.067	0.261	0.044	0.107	80.1	0.285	0.808	0.127	0.151
OT-MDR (Ours)	96.2	0.059	0.134	0.018	0.005	81.0	0.268	0.693	0.045	0.045
Ensemble of ResNet18, MobileNet and EfficientNet										
Deep Ensemble	89.0	0.153	0.395	0.111	0.126	62.7	0.433	1.267	0.176	0.209
DST	93.4	0.102	0.282	0.070	0.109	71.7	0.393	1.066	0.148	0.187
SGD	92.6	0.113	0.317	-	0.112	72.6	0.403	1.192	-	0.201
SAM	93.8	0.094	0.280	0.060	0.110	76.4	0.347	1.005	0.142	0.177
OT-MDR (Ours)	94.8	0.078	0.176	0.021	0.007	78.3	0.310	0.828	0.047	0.063

We conduct experiments to compare our OT-MDR with bSAM [46], SGD, and SAM [18] on Resnet18. The results, shown in Table 2, demonstrate that the OT-MDR approach consistently outperforms all baselines by a substantial margin. Here we note that we cannot evaluate bSAM on the architectures used in Table 1 because the authors did not release the code. Instead, we run our OT-MDR with the setting mentioned in the original bSAM paper.

5.2 Experiments on Ensemble Models

To investigate the effectiveness of our approach in the context of a uniform distribution over a model space, we examine the ensemble inference of multiple base models trained independently. The ensemble prediction is obtained by averaging the prediction probabilities of all base models, following the standard process of ensemble methods. We compare our approach against several state-of-the-art ensemble methods, including Deep Ensembles [37], Snapshot Ensembles [25], Fast Geometric Ensemble (FGE) [22], and sparse ensembles EDST and DST [40]. In addition, we compare our approach with another ensemble method that utilizes SAM as an optimizer to improve the generalization ability, as discussed in Section 4.3. The value of ρ for SAM and ρ_1, ρ_2 for OT-MDR is the same as in the single model setting.

To evaluate the performance of each method, we measure five metrics over the average prediction, which represent both predictive performance (Accuracy - ACC) and uncertainty estimation (Brier score, Negative Log-Likelihood - NLL, Expected Calibration Error - ECE, and Average across all calibrated confidence - AAC) on the CIFAR dataset, as shown in Table 3.

Notably, our OT-MDR approach consistently outperforms all baselines across all metrics, demonstrating the benefits of incorporating diversity across base models to achieve distributional robustness. Remarkably, OT-MDR even surpasses SAM, the runner-up baseline, by a significant margin, indicating a better generalization capability.

5.3 Experiment on Bayesian Neural Networks

We now assess the effectiveness of OT-MDR in the context of variational inference, where model parameters are sampled from a Gaussian distribution. Specifically, we apply our proposed method to

the widely-used variational technique as SGVB [34], and compare its performance with the original approach. We conduct experiments on two different architectures, Resnet10 and Resnet18 using the CIFAR dataset, and report the results in Table 4. It is clear that our approach outperforms the original SGVB method in all metrics, showcasing significant improvements. These findings underscore OT-MDR ability to increase accuracy, better calibration, and improve uncertainty estimation.

Table 4: Classification scores of approximate the Gaussian posterior on the CIFAR datasets. All experiments are trained three times with different random seeds.

Dataset	Method	Resnet10			Resnet18		
		ACC \uparrow	NLL \downarrow	ECE \downarrow	ACC \uparrow	NLL \downarrow	ECE \downarrow
CIFAR-10	SGVB	80.52 \pm 2.10	0.78 \pm 0.23	0.23 \pm 0.06	86.74 \pm 1.25	0.54 \pm 0.01	0.18 \pm 0.02
	OT-MDR (Ours)	81.26 \pm 0.06	0.81 \pm 0.12	0.26 \pm 0.08	87.55 \pm 0.14	0.52 \pm 0.01	0.17 \pm 0.01
CIFAR-100	SGVB	54.40 \pm 0.98	1.96 \pm 0.05	0.21 \pm 0.00	60.91 \pm 2.31	1.74 \pm 0.15	0.24 \pm 0.03
	OT-MDR (Ours)	55.33 \pm 0.11	1.85 \pm 0.06	0.18 \pm 0.03	63.17 \pm 0.04	1.55 \pm 0.05	0.20 \pm 0.03

5.4 Ablation Study

Effect of Number of Particle Models. For multiple-particles setting on a single model as mentioned in Section 4.2.1, we investigate the effectiveness of diversity in achieving distributional robustness. We conduct experiments using the WideResnet28x10 model on the CIFAR datasets, training with $K \in \{1, 2, 3, 4\}$ particles. The results are presented in Figure 1. Note that in this setup, we utilize the same hyper-parameters as the one-particle setting, but only train for 100 epochs to save time. Interestingly, we observe that using two-particles achieved higher accuracy compared to one particle. However, as we increase the number of particles, the difference between them also increases, resulting in worse performance. These results suggest that while diversity can be beneficial in achieving distributional robustness, increasing the number of particles beyond a certain threshold may have diminishing returns and potentially lead to performance deterioration.

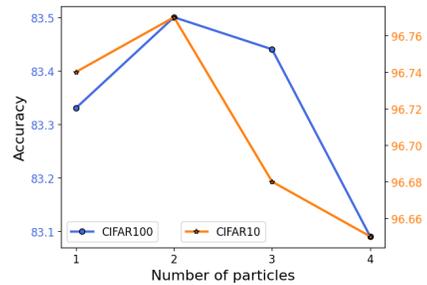


Figure 1: Multiple particle classification accuracies on the CIFAR datasets with WideResnet28x10. Note that we train 100 epochs for each experiment and report the average accuracy

Loss landscape. We depict the loss landscape for ensemble inference using different architectures on the CIFAR 100 dataset and make a comparison with the SAM method, which serves as our runner-up. As demonstrated in Figure 2, our approach guides the model towards a lower and flatter loss region compared to SAM, which improves the model’s performance. This is important because a lower loss signifies better optimization, while a flatter region indicates improved generalization and robustness. By attaining both these characteristics, our approach enhances the model’s ability to achieve high accuracy and stability (Table 3).

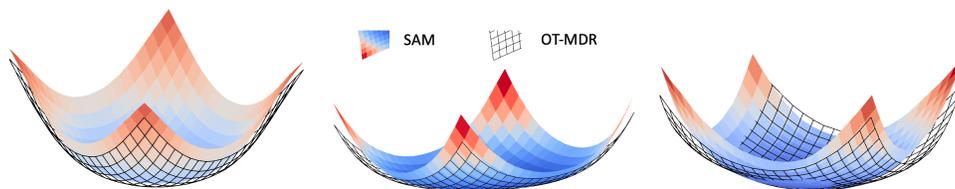


Figure 2: Comparing loss landscape of (left) Ensemble of 5 Resnet10, (middle) Ensemble of 3 Resnet18, and (right) Ensemble of ResNet18, MobileNet and EfficientNet on CIFAR-100 dataset training with SAM and OT-MDR. Evidently, OT-MDR leads the model to a flatter and lower loss area

6 Conclusion

In this paper, we explore the relationship between OT-based distributional robustness and sharpness-aware minimization (SAM), and show that SAM is a special case of our framework when a Dirac delta distribution is used over a single model. Our proposed framework can be seen as a probabilistic extension of SAM. Additionally, we extend the OT-based distributional robustness framework to propose a practical method that can be applied to (i) a Dirac delta distribution over a single model, (ii) a uniform distribution over several models, and (iii) a general distribution over the model space (i.e., a Bayesian Neural Network). To demonstrate the effectiveness of our approach, we conduct experiments that show significant improvements over the baselines. We believe that the theoretical connection between the OT-based distributional robustness and SAM could be valuable for future research, such as exploring the dual form in Eq. (15) to adapt the perturbed radius ρ .

Acknowledgements. This work was partly supported by ARC DP23 grant DP230101176 and by the Air Force Office of Scientific Research under award number FA2386-23-1-4044.

References

- [1] Momin Abbas, Quan Xiao, Lisha Chen, Pin-Yu Chen, and Tianyi Chen. Sharp-maml: Sharpness-aware model-agnostic meta learning. *arXiv preprint arXiv:2206.03996*, 2022. 3
- [2] Maksym Andriushchenko and Nicolas Flammarion. Towards understanding sharpness-aware minimization. In *International Conference on Machine Learning*, pages 639–668. PMLR, 2022. 1
- [3] Dara Bahri, Hossein Mobahi, and Yi Tay. Sharpness-aware minimization improves language model generalization. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7360–7371, Dublin, Ireland, May 2022. Association for Computational Linguistics. 3
- [4] Aharon Ben-Tal, Dick Den Hertog, Anja De Waegenaere, Bertrand Melenberg, and Gijs Rennen. Robust solutions of optimization problems affected by uncertain probabilities. *Management Science*, 59(2):341–357, 2013. 1, 2
- [5] Jose Blanchet and Yang Kang. Semi-supervised learning based on distributionally robust optimization. *Data Analysis and Applications 3: Computational, Classification, Financial, Statistical and Stochastic Methods*, 5:1–33, 2020. 1, 2
- [6] Jose Blanchet, Yang Kang, and Karthyek Murthy. Robust wasserstein profile inference and applications to machine learning. *Journal of Applied Probability*, 56(3):830–857, 2019. 1, 2
- [7] Jose Blanchet and Karthyek Murthy. Quantifying distributional model risk via optimal transport. *Mathematics of Operations Research*, 44(2):565–600, 2019. 2, 3
- [8] Tuan Anh Bui, Trung Le, Quan Tran, He Zhao, and Dinh Phung. A unified wasserstein distributional robustness framework for adversarial training. *arXiv preprint arXiv:2202.13437*, 2022. 1, 2
- [9] Junbum Cha, Sanghyuk Chun, Kyungjae Lee, Han-Cheol Cho, Seunghyun Park, Yunsung Lee, and Sungrae Park. Swad: Domain generalization by seeking flat minima. *Advances in Neural Information Processing Systems*, 34:22405–22418, 2021. 3
- [10] Pratik Chaudhari, Anna Choromańska, Stefano Soatto, Yann LeCun, Carlo Baldassi, Christian Borgs, Jennifer T. Chayes, Levent Sagun, and Riccardo Zecchina. Entropy-sgd: biasing gradient descent into wide valleys. *Journal of Statistical Mechanics: Theory and Experiment*, 2017. 2

- [11] Ruidi Chen and Ioannis C Paschalidis. A robust learning approach for regression models based on distributionally robust optimization. *Journal of Machine Learning Research*, 19(13), 2018. 1, 2
- [12] Xiangning Chen, Cho-Jui Hsieh, and Boqing Gong. When vision transformers outperform resnets without pre-training or strong data augmentations. *arXiv preprint arXiv:2106.01548*, 2021. 3
- [13] Laurent Dinh, Razvan Pascanu, Samy Bengio, and Yoshua Bengio. Sharp minima can generalize for deep nets. In *International Conference on Machine Learning*, pages 1019–1028. PMLR, 2017. 2
- [14] Jiawei Du, Daquan Zhou, Jiashi Feng, Vincent YF Tan, and Joey Tianyi Zhou. Sharpness-aware training for free. *arXiv preprint arXiv:2205.14083*, 2022. 3
- [15] John C Duchi, Peter W Glynn, and Hongseok Namkoong. Statistics of robust optimization: A generalized empirical likelihood approach. *Mathematics of Operations Research*, 2021. 1, 2
- [16] John C Duchi, Tatsunori Hashimoto, and Hongseok Namkoong. Distributionally robust losses against mixture covariate shifts. *Under review*, 2019. 1, 2
- [17] Gintare Karolina Dziugaite and Daniel M. Roy. Computing nonvacuous generalization bounds for deep (stochastic) neural networks with many more parameters than training data. In *UAI*. AUAI Press, 2017. 2
- [18] Pierre Foret, Ariel Kleiner, Hossein Mobahi, and Behnam Neyshabur. Sharpness-aware minimization for efficiently improving generalization. In *International Conference on Learning Representations*, 2021. 1, 2, 5, 8
- [19] Stanislav Fort and Surya Ganguli. Emergent properties of the local geometry of neural loss landscapes. *arXiv preprint arXiv:1910.05929*, 2019. 2
- [20] Rui Gao, Xi Chen, and Anton J Kleywegt. Wasserstein distributional robustness and regularization in statistical learning. *arXiv e-prints*, pages arXiv–1712, 2017. 3
- [21] Rui Gao and Anton J Kleywegt. Distributionally robust stochastic optimization with wasserstein distance. *arXiv preprint arXiv:1604.02199*, 2016. 2
- [22] Timur Garipov, Pavel Izmailov, Dmitrii Podoprikin, Dmitry P Vetrov, and Andrew G Wilson. Loss surfaces, mode connectivity, and fast ensembling of dnns. *Advances in neural information processing systems*, 31, 2018. 8
- [23] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014. 7
- [24] Sepp Hochreiter and Jürgen Schmidhuber. Simplifying neural nets by discovering flat minima. In *NIPS*, pages 529–536. MIT Press, 1994. 2
- [25] Gao Huang, Yixuan Li, Geoff Pleiss, Zhuang Liu, John E. Hopcroft, and Kilian Q. Weinberger. Snapshot ensembles: Train 1, get m for free. In *International Conference on Learning Representations*, 2017. 8
- [26] Pavel Izmailov, Dmitrii Podoprikin, Timur Garipov, Dmitry Vetrov, and Andrew Gordon Wilson. Averaging weights leads to wider optima and better generalization. *arXiv preprint arXiv:1803.05407*, 2018. 1, 3
- [27] Pavel Izmailov, Dmitrii Podoprikin, Timur Garipov, Dmitry P. Vetrov, and Andrew Gordon Wilson. Averaging weights leads to wider optima and better generalization. In *UAI*, pages 876–885. AUAI Press, 2018. 2
- [28] Stanislaw Jastrzebski, Zachary Kenton, Devansh Arpit, Nicolas Ballas, Asja Fischer, Yoshua Bengio, and Amos J. Storkey. Three factors influencing minima in sgd. *ArXiv*, abs/1711.04623, 2017. 2

- [29] Yiding Jiang, Behnam Neyshabur, Hossein Mobahi, Dilip Krishnan, and Samy Bengio. Fantastic generalization measures and where to find them. In *ICLR*. OpenReview.net, 2020. [2](#)
- [30] Jean Kaddour, Linqing Liu, Ricardo Silva, and Matt Kusner. When do flat minima optimizers work? In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, *Advances in Neural Information Processing Systems*, 2022. [1](#), [3](#)
- [31] Jean Kaddour, Linqing Liu, Ricardo Silva, and Matt J Kusner. A fair comparison of two popular flat minima optimizers: Stochastic weight averaging vs. sharpness-aware minimization. *arXiv preprint arXiv:2202.00661*, 1, 2022. [1](#)
- [32] 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 *ICLR*. OpenReview.net, 2017. [2](#)
- [33] Minyoung Kim, Da Li, Shell X Hu, and Timothy Hospedales. Fisher SAM: Information geometry and sharpness aware minimisation. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato, editors, *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 11148–11161. PMLR, 17–23 Jul 2022. [1](#), [3](#)
- [34] Durk P Kingma, Tim Salimans, and Max Welling. Variational dropout and the local reparameterization trick. *Advances in neural information processing systems*, 28, 2015. [9](#)
- [35] Daniel Kuhn, Peyman Mohajerin Esfahani, Viet Anh Nguyen, and Soroosh Shafieezadeh-Abadeh. Wasserstein distributionally robust optimization: Theory and applications in machine learning. In *Operations Research & Management Science in the Age of Analytics*, pages 130–166. INFORMS, 2019. [2](#)
- [36] Jungmin Kwon, Jeongseop Kim, Hyunseo Park, and In Kwon Choi. Asam: Adaptive sharpness-aware minimization for scale-invariant learning of deep neural networks. In *International Conference on Machine Learning*, pages 5905–5914. PMLR, 2021. [1](#), [3](#)
- [37] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in neural information processing systems*, 30, 2017. [8](#)
- [38] Trung Le, Tuan Nguyen, Nhat Ho, Hung Bui, and Dinh Phung. Lamda: Label matching deep domain adaptation. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 6043–6054. PMLR, 18–24 Jul 2021. [1](#)
- [39] Jaeho Lee and Maxim Raginsky. Minimax statistical learning with wasserstein distances. In *NeurIPS*, pages 2692–2701, 2018. [1](#)
- [40] Shiwei Liu, Tianlong Chen, Zahra Atashgahi, Xiaohan Chen, Ghada Sokar, Elena Mocanu, Mykola Pechenizkiy, Zhangyang Wang, and Decebal Constantin Mocanu. Deep ensembling with no overhead for either training or testing: The all-round blessings of dynamic sparsity. In *International Conference on Learning Representations*, 2022. [8](#)
- [41] Yong Liu, Siqi Mai, Xiangning Chen, Cho-Jui Hsieh, and Yang You. Towards efficient and scalable sharpness-aware minimization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12360–12370, 2022. [3](#)
- [42] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In *International Conference on Learning Representations*, 2018. [7](#)
- [43] Takeru Miyato, Shin-ichi Maeda, Masanori Koyama, Ken Nakae, and Shin Ishii. Distributional smoothing with virtual adversarial training. *arXiv preprint arXiv:1507.00677*, 2015. [2](#)
- [44] Peyman Mohajerin Esfahani and Daniel Kuhn. Data-driven distributionally robust optimization using the wasserstein metric: Performance guarantees and tractable reformulations. *arXiv e-prints*, pages arXiv–1505, 2015. [2](#)

- [45] Thomas Möllenhoff and Mohammad Emtiyaz Khan. Sam as an optimal relaxation of bayes. *arXiv preprint arXiv:2210.01620*, 2022. [1](#)
- [46] Thomas Möllenhoff and Mohammad Emtiyaz Khan. SAM as an optimal relaxation of bayes. In *The Eleventh International Conference on Learning Representations*, 2023. [1](#), [3](#), [7](#), [8](#)
- [47] Hongseok Namkoong and John C Duchi. Stochastic gradient methods for distributionally robust optimization with f-divergences. In *NIPS*, volume 29, pages 2208–2216, 2016. [2](#)
- [48] Radford M Neal. *Bayesian learning for neural networks*, volume 118. Springer Science & Business Media, 2012. [2](#)
- [49] Behnam Neyshabur, Srinadh Bhojanapalli, David McAllester, and Nati Srebro. Exploring generalization in deep learning. *Advances in neural information processing systems*, 30, 2017. [2](#)
- [50] Tuan Nguyen, Trung Le, Nhan Dam, Quan Hung Tran, Truyen Nguyen, and Dinh Q Phung. Tidot: A teacher imitation learning approach for domain adaptation with optimal transport. In *IJCAI*, pages 2862–2868, 2021. [1](#), [2](#)
- [51] Tuan Nguyen, Trung Le, He Zhao, Quan Hung Tran, Truyen Nguyen, and Dinh Phung. Most: multi-source domain adaptation via optimal transport for student-teacher learning. In Cassio de Campos and Marloes H. Maathuis, editors, *Proceedings of the Thirty-Seventh Conference on Uncertainty in Artificial Intelligence*, volume 161 of *Proceedings of Machine Learning Research*, pages 225–235. PMLR, 27–30 Jul 2021. [1](#)
- [52] Tuan Nguyen, Van Nguyen, Trung Le, He Zhao, Quan Hung Tran, and Dinh Phung. Cycle class consistency with distributional optimal transport and knowledge distillation for unsupervised domain adaptation. In James Cussens and Kun Zhang, editors, *Proceedings of the Thirty-Eighth Conference on Uncertainty in Artificial Intelligence*, volume 180 of *Proceedings of Machine Learning Research*, pages 1519–1529. PMLR, 01–05 Aug 2022. [1](#)
- [53] Van-Anh Nguyen, Tung-Long Vuong, Hoang Phan, Thanh-Toan Do, Dinh Phung, and Trung Le. Flat seeking bayesian neural network. In *Advances in Neural Information Processing Systems*, 2023. [2](#), [3](#)
- [54] Gabriel Pereyra, George Tucker, Jan Chorowski, Lukasz Kaiser, and Geoffrey E. Hinton. Regularizing neural networks by penalizing confident output distributions. In *ICLR (Workshop)*. OpenReview.net, 2017. [2](#)
- [55] Henning Petzka, Michael Kamp, Linara Adilova, Cristian Sminchisescu, and Mario Boley. Relative flatness and generalization. In *NeurIPS*, pages 18420–18432, 2021. [2](#)
- [56] Cuong Pham, C. Cuong Nguyen, Trung Le, Phung Dinh, Gustavo Carneiro, and Thanh-Toan Do. Model and feature diversity for bayesian neural networks in mutual learning. In *Advances in Neural Information Processing Systems*, 2023. [2](#)
- [57] Hoang Phan, Trung Le, Trung Phung, Anh Tuan Bui, Nhat Ho, and Dinh Phung. Global-local regularization via distributional robustness. In Francisco Ruiz, Jennifer Dy, and Jan-Willem van de Meent, editors, *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics*, volume 206 of *Proceedings of Machine Learning Research*, pages 7644–7664. PMLR, 25–27 Apr 2023. [1](#), [2](#)
- [58] Hoang Phan, Ngoc Tran, Trung Le, Toan Tran, Nhat Ho, and Dinh Phung. Stochastic multiple target sampling gradient descent. *Advances in neural information processing systems*, 2022. [3](#)
- [59] Zhe Qu, Xingyu Li, Rui Duan, Yao Liu, Bo Tang, and Zhuo Lu. Generalized federated learning via sharpness aware minimization. *arXiv preprint arXiv:2206.02618*, 2022. [3](#)
- [60] Soroosh Shafieezadeh-Abadeh, Peyman Mohajerin Esfahani, and Daniel Kuhn. Distributionally robust logistic regression. *arXiv preprint arXiv:1509.09259*, 2015. [2](#)
- [61] Aman Sinha, Hongseok Namkoong, and John Duchi. Certifying some distributional robustness with principled adversarial training. In *International Conference on Learning Representations*, 2018. [1](#), [2](#), [3](#)

- [62] Tuan Truong, Hoang-Phi Nguyen, Tung Pham, Minh-Tuan Tran, Mehrtash Harandi, Dinh Phung, and Trung Le. Rsam: Learning on manifolds with riemannian sharpness-aware minimization, 2023. [1](#)
- [63] Cédric Villani. Optimal transport: Old and new. 2008. [3](#)
- [64] Colin Wei, Sham Kakade, and Tengyu Ma. The implicit and explicit regularization effects of dropout. In *International conference on machine learning*, pages 10181–10192. PMLR, 2020. [2](#)
- [65] Max Welling and Yee W Teh. Bayesian learning via stochastic gradient langevin dynamics. In *Proceedings of the 28th international conference on machine learning (ICML-11)*, pages 681–688, 2011. [4](#)
- [66] Insoon Yang. Wasserstein distributionally robust stochastic control: A data-driven approach. *IEEE Transactions on Automatic Control*, 2020. [1](#), [2](#)
- [67] Linfeng Zhang, Jiebo Song, Anni Gao, Jingwei Chen, Chenglong Bao, and Kaisheng Ma. Be your own teacher: Improve the performance of convolutional neural networks via self distillation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3713–3722, 2019. [2](#)
- [68] Ying Zhang, Tao Xiang, Timothy M. Hospedales, and Huchuan Lu. Deep mutual learning. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4320–4328, 2018. [2](#)
- [69] Han Zhao, Remi Tachet Des Combes, Kun Zhang, and Geoffrey Gordon. On learning invariant representations for domain adaptation. In *International Conference on Machine Learning*, pages 7523–7532. PMLR, 2019. [1](#)
- [70] Long Zhao, Ting Liu, Xi Peng, and Dimitris Metaxas. Maximum-entropy adversarial data augmentation for improved generalization and robustness. *arXiv preprint arXiv:2010.08001*, 2020. [1](#), [2](#)
- [71] Juntang Zhuang, Boqing Gong, Liangzhe Yuan, Yin Cui, Hartwig Adam, Nicha Dvornek, Sekhar Tatikonda, James Duncan, and Ting Liu. Surrogate gap minimization improves sharpness-aware training. *arXiv preprint arXiv:2203.08065*, 2022. [3](#)