

Adversarial Transformations for Semi-Supervised Learning

Tepei Suzuki, Ikuro Sato
 DENSO IT LABORATORY, INC.
 2-15-1 Shibuya, Shibuya-ku Tokyo, Japan
 {tsuzuki, isato}@d-itlab.co.jp

Abstract

We propose a **Regularization** framework based on **Adversarial Transformations (RAT)** for semi-supervised learning. RAT is designed to enhance robustness of the output distribution of class prediction for a given data against input perturbation. RAT is an extension of Virtual Adversarial Training (VAT) in such a way that RAT adversarially transforms data along the underlying data distribution by a rich set of data transformation functions that leave class label invariant, whereas VAT simply produces adversarial additive noises. In addition, we verified that a technique of gradually increasing of perturbation region further improves the robustness. In experiments, we show that RAT significantly improves classification performance on CIFAR-10 and SVHN compared to existing regularization methods under standard semi-supervised image classification settings.

Introduction

Semi-supervised learning (SSL) (Chapelle, Scholkopf, and Zien 2006) is an effective learning framework on datasets that have large amounts of the data and few labels. In a practical situations, obtained datasets are often partially labeled, because labeling is more costly than collecting data in many cases. Thus, a powerful SSL framework that enhances the model performance is needed.

Among the many SSL algorithms, virtual adversarial training (VAT) (Miyato et al. 2016; 2018) is a successful one. It regularizes local distributional smoothness, by which we mean the robustness of the output distribution around each input datapoint against local perturbation, by using virtual adversarial perturbation. The perturbation is calculated as the noise, which adversarially changes the output of the model, and VAT imposes the consistency between the outputs of the model for the data and the perturbed data. It indicates that VAT enforces the output distribution to be robust with respect to perturbation within the ϵ -ball centered on the data, where ϵ denotes the norm of the perturbation vector. The specific advantage of VAT is to leverage the adversarial perturbation.

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

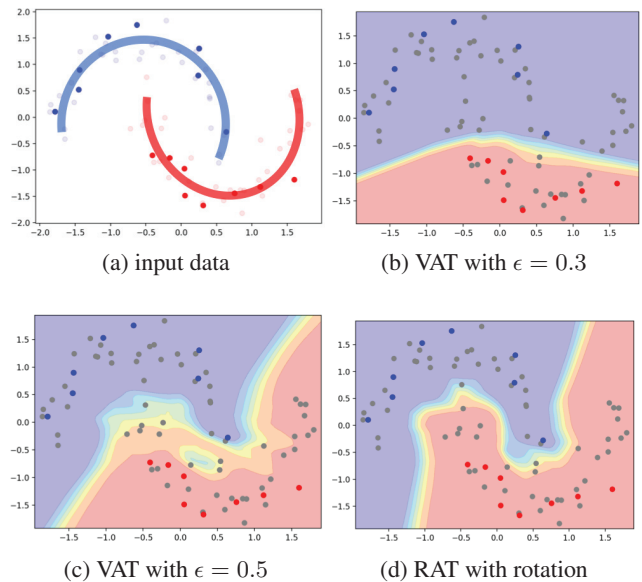


Figure 1: Comparison between VAT and RAT with moons dataset. Colored regions of (b), (c), and (d) correspond to prediction confidence. Each moon has 10 labeled data (dark color points in (a)) and 30 unlabeled data (light color points in (a)), and each data is randomly generated from semicircles in (a) with Gaussian noise. RAT uses noise injection and rotation along the blue and red semicircles in (a) as transformation. The setting of (b) and (d) is the same except for using rotation in (d).

Although VAT demonstrates remarkable accuracy in standard benchmarks, VAT allows the smoothing only within the ϵ -ball. This means that there is no guarantee that two datapoints are classified into the same class even if they belong to the same class, if the distance between them is greater than 2ϵ (see Figure 1 (b)). However, if ϵ is too large, the perturbed data may penetrate the true class boundary; inconsistency within ϵ -ball may occur (see Figure 1 (c)). To summarize, the ϵ -ball, that is an isotropic hypersphere in the input space, is too universal to take the underlying data distribution into

account. Our basic idea is to use some adversarial transformations under which a datapoint transforms to another point within the underlying distribution of the same class.

We propose a regularization framework, RAT (Regularization based on Adversarial Transformations), which regularizes the smoothness of the output distribution by utilizing adversarial transformations. The aim of RAT is to make the outputs with respect to the same class data close.

We justify RAT as the regularization for smoothing the distribution, and provide use of composite functions and a technique, ϵ -rampup, that enlarge the area in which the smoothing is effective. To demonstrate the effectiveness of RAT, we compare it with baseline methods in the experiments following the valid setting proposed by Oliver et al. (2018), and RAT outperforms baseline methods.

We summarize our contributions as follows:

- We propose the regularization framework based on adversarial transformations, which includes VAT as a special case. Unlike VAT, RAT can impose distributional smoothness along the underlying data distribution.
- We claim that use of composite transformations, each of which leaves class label invariant, can further improve the smoothing effect, because it enhances the degrees of freedom.
- Moreover, we provide a technique to enhance the smoothing effect by ramping up ϵ . The technique is common to VAT and RAT, and enlarge the area in which the smoothing is effective.
- RAT outperforms baseline methods in semi-supervised image classification tasks. In particular, RAT is robust against reduction of labeled samples, compared to other methods.

Virtual Adversarial Training: A Review

In this section, we review virtual adversarial training (VAT) (Miyato et al. 2016; 2018). VAT is a similar method of the adversarial training (Goodfellow, Shlens, and Szegedy 2014), but the aim of VAT is to regularize distributional smoothness. Miyato et al. (2016; 2018) claimed importance of the adversarial noise and proved this by comparing it with random perturbation training. VAT has indeed shown state-of-the-art results in valid benchmarks (Oliver et al. 2018).

Let $x \in \mathbb{R}^I$ be a data sample where I is the dimension of the input, and $p_\theta(y|x)$ be a model distribution parameterized by θ . The objective function of VAT in SSL scenario is:

$$\mathcal{L}_{\text{VAT}}(x_l, y, x_u, \theta) \equiv \mathcal{L}(x_l, y, \theta) - \lambda \text{LDS}(x_u, \theta), \quad (1)$$

where \mathcal{L} denotes a supervised loss term. x_l, y and x_u are labeled data, its labels, and unlabeled data. λ is a scaling parameter for regularization. LDS, local distributional smoothness, is defined as follows:

$$\text{LDS}(x, \theta) \equiv -\Delta_{\text{KL}}(r_{\text{v-adv}}, x, \theta), \quad (2)$$

where Δ_{KL} and $r_{\text{v-adv}}$ are KL divergence between output distributions with respect to an original data and a noise-

added data, and virtual adversarial perturbation, respectively. They are defined as follows:

$$\Delta_{\text{KL}}(r, x, \theta) \equiv \text{KL}[p_\theta(y|x) \| p_\theta(y|x+r)], \quad (3)$$

$$r_{\text{v-adv}} \equiv \arg \max_r \{\Delta_{\text{KL}}(r, x, \theta); \|r\|_2 \leq \epsilon\}, \quad (4)$$

where $\text{KL}[\cdot \| \cdot]$ denotes KL divergence and $\|\cdot\|_2$ denotes the L2 norm. ϵ is a hyperparameter to control the range of the smoothing.

$r_{\text{v-adv}}$ emerges as the eigenvector of Hessian matrix $H(x, \theta) \equiv \nabla \nabla_r \Delta_{\text{KL}}(r, x, \theta)|_{r=0}$ from the following derivation. Since $\Delta_{\text{KL}}(r, x, \theta)$ takes a minimum value 0 at $r = 0$, $\nabla_r \Delta_{\text{KL}}(r, x, \theta)|_{r=0}$ is 0. Therefore, the second-order Taylor approximation of $\Delta_{\text{KL}}(r, x, \theta)$ around $r = 0$ is:

$$\Delta_{\text{KL}}(r, x, \theta) \approx \frac{1}{2} r^\top H(x, \theta) r. \quad (5)$$

We can describe (4) using (5) as follows:

$$r_{\text{v-adv}} \approx \arg \max_r \{r^\top H(x, \theta) r; \|r\|_2 \leq \epsilon\}. \quad (6)$$

By this approximation, $r_{\text{v-adv}}$ is parallel to the eigenvector corresponding to the largest eigenvalue of $H(x, \theta)$. Note that VAT assumes that $p_\theta(y|x)$ is differentiable with respect to θ and x almost everywhere, and we also assume it in this paper.

Regularization Based on Adversarial Transformations

We propose the regularization framework, RAT, which imposes the consistency between output distributions with respect to datapoints belonging to the same class. Leveraging the power of adversariality, we introduce adversarial transformations that replace additive adversarial noises $r_{\text{v-adv}}$ in VAT.

To consider imposing the consistency, we assume that each datapoint belonging to k -th class is in a class-specific subspace \mathcal{S}_k . We consider the generic transformation parameterized by ϕ : for any k , $f_\phi : \mathcal{S}_k \rightarrow \mathcal{S}_k$ such that $\|\phi\| \leq \epsilon$, where $\|\cdot\|$ denotes valid norm with respect to f_ϕ such as L1, L2, or operator norm. Our strategy is to regularize the output distribution by utilizing f_ϕ instead of $r_{\text{v-adv}}$. Note that we mainly consider the image classification tasks in this work. Therefore, we deal with the transformation such as the spatial transformation or the color distortion.

We define local distributional smoothness utilizing a transformation as follows:

$$\text{LDS}_T(x, \theta) \equiv -\tilde{\Delta}_{\text{KL}}(f_{\phi_{\text{T-adv}}}, x, \theta), \quad (7)$$

where $\tilde{\Delta}_{\text{KL}}$ and $\phi_{\text{T-adv}}$ are defined as follows:

$$\tilde{\Delta}_{\text{KL}}(f_\phi, x, \theta) \equiv \text{KL}[p_\theta(y|x) \| p_\theta(y|f_\phi(x))], \quad (8)$$

$$\phi_{\text{T-adv}} \equiv \arg \max_\phi \{\tilde{\Delta}_{\text{KL}}(f_\phi, x, \theta); \|\phi\| \leq \epsilon\}. \quad (9)$$

We refer to $f_{\phi_{\text{T-adv}}}$ as an adversarial transformation in this paper. There are some adversarial attacks utilizing functions instead of additive noises (Alaifari, Alberti, and Gauksson

2018; Tsuzuku and Sato 2019), and the relation between adversarial transformation and these attacks is similar to the relation between virtual adversarial perturbation and adversarial perturbation (Goodfellow, Shlens, and Szegedy 2014).

We utilize LDS_T for imposing the consistency. Thus, the objective function of RAT for SSL scenario is represented as follows:

$$\mathcal{L}_{\text{RAT}}(x_l, y, x_u, \theta) \equiv \mathcal{L}(x_l, y, \theta) - \lambda \text{LDS}_T(x_u, \theta). \quad (10)$$

We can identify VAT as the special case of RAT; when $f_\phi(x) = x + \phi$ and $\|\cdot\|$ is the L2 norm, RAT is equal to VAT.

To compute LDS_T , we have to solve the maximization problem (9). However, it is difficult to exactly solve it. Therefore, we consider approximating it by using Taylor approximation.

To efficiently approximate and solve (9), f_ϕ needs to satisfy the following two conditions:

- C1. f_ϕ is differentiable with respect to ϕ almost everywhere.
- C2. There is a parameter ϕ_{id} that makes f_ϕ identity transformation, $f_{\phi_{\text{id}}}(x) = x$.

If f_ϕ satisfies these conditions, $\tilde{\Delta}_{\text{KL}}(f_\phi, x, \theta)$ takes a minimum value at $\phi = \phi_{\text{id}}$ and (8) is written in the same form as (5) by the second-order Taylor approximation around ϕ_{id} as $\tilde{\Delta}_{\text{KL}}(f_\phi, x, \theta) \approx \frac{1}{2} \phi^\top H(f_\phi, x, \theta) \phi$, where $H(f_\phi, x, \theta) \equiv \nabla \nabla_\phi \tilde{\Delta}_{\text{KL}}(f_\phi, x, \theta)|_{\phi=\phi_{\text{id}}}$. Thus, (9) is approximated as follows:

$$\phi_{T\text{-adv}} \approx \arg \max_{\phi} \{\phi^\top H(f_\phi, x, \theta) \phi; \|\phi\| \leq \epsilon\}. \quad (11)$$

$\phi_{T\text{-adv}}$ is also parallel to the eigenvector corresponding to the largest eigenvalue of $H(f_\phi, x, \theta)$.

f_ϕ allows any transformations satisfying $f_\phi : \mathcal{S}_k \rightarrow \mathcal{S}_k$, and the conditions C1 and C2. In this paper, we use hand crafted transformations depending on input data domain such as image coordinate shift, image resizing, or global color change, to name a few for the case of an image classification task. In a next section, we propose use of composite transformations as f_ϕ to further enhance the effect of the smoothing.

Use of Composite Transformations

Let \mathcal{F} be a set of functions, $\{f_\phi : \mathcal{S}_k \rightarrow \mathcal{S}_k; \|\phi\| \leq \epsilon\}$, satisfying the conditions C1 and C2. \mathcal{F} may contain composite functions of the form $\tilde{f}_{\tilde{\phi}} = f_{\phi^{(1)}}^{(1)} \circ f_{\phi^{(2)}}^{(2)} \cdots \circ f_{\phi^{(n)}}^{(n)}$, where $f_{\phi^{(i)}}^{(i)} \in \mathcal{F}; \|\phi^{(i)}\| \leq \epsilon^{(i)}$ and $\tilde{\phi} \equiv \{\phi^{(i)}\}_{i=1}^n$. It is obvious that these composite functions $\tilde{f}_{\tilde{\phi}} : \mathcal{S}_k \rightarrow \mathcal{S}_k (\forall k)$ satisfy the conditions C1 and C2.

By having such composite functions, one can obtain a much richer set of transformations that still yield class-invariant transformations, as given by the relation, $\tilde{f}_{\tilde{\phi}}^{(j)} = f_{\phi^{(j)}}^{(j)}$ when $f_{\phi^{(i)}}^{(i)} = f_{\text{id}}^{(i)}$ for all i except for j . It is reasonable for RAT to utilize the composite functions, because the composite functions leads to the richer transformation and imposing the consistency over a wider range.

Fast Approximation of $f_{\phi_{T\text{-adv}}}$

Although $\phi_{T\text{-adv}}$ emerges as the eigenvector corresponding to the largest eigenvalue of Hessian matrix $\tilde{H} = \nabla \nabla_{\tilde{\phi}} \tilde{\Delta}_{\text{KL}}(\tilde{f}_{\tilde{\phi}}, x, \theta)|_{\tilde{\phi}=\tilde{\phi}_{\text{id}}}$ as already described, the computational costs of the eigenvector are $O(n^3)$. There is a way to approximate the eigenvector with small computational costs (Golub and Van der Vorst 2001; Miyato et al. 2016; 2018).

We approximate the parameters $\tilde{\phi}$ with the power iteration and the finite difference method, just as VAT does. For each transformation $f^{(i)} \in \mathcal{F}$, we sample random unit vectors $\tilde{v} \equiv \{v^{(i)}\}_{i=1}^n$ as initial parameters and calculate $\tilde{v} \leftarrow \tilde{H} \tilde{v}$ iteratively. It makes \tilde{v} converge to the eigenvector corresponding to the largest eigenvalue of \tilde{H} .

The Hessian-vector product is calculated with the finite difference method as follows:

$$\begin{aligned} \tilde{H} \tilde{v} &\approx \frac{\nabla_{\tilde{\phi}} \tilde{\Delta}_{\text{KL}}(\tilde{f}_{\tilde{\phi}}, x, \theta)|_{\tilde{\phi}=\tilde{\phi}_{\text{id}}+\xi \tilde{v}} - \nabla_{\tilde{\phi}} \tilde{\Delta}_{\text{KL}}(\tilde{f}_{\tilde{\phi}}, x, \theta)|_{\tilde{\phi}=\tilde{\phi}_{\text{id}}}}{\xi}, \\ &= \frac{\nabla_{\tilde{\phi}} \tilde{\Delta}_{\text{KL}}(\tilde{f}_{\tilde{\phi}}, x, \theta)|_{\tilde{\phi}=\tilde{\phi}_{\text{id}}+\xi \tilde{v}}}{\xi}, \end{aligned} \quad (12)$$

with $\xi \neq 0$. We used the fact that $\nabla_{\tilde{\phi}} \tilde{\Delta}_{\text{KL}}(\tilde{f}_{\tilde{\phi}}, x, \theta)|_{\tilde{\phi}=\tilde{\phi}_{\text{id}}} = 0$. We approximate $\phi_{T\text{-adv}}$ by normalizing the norm of the approximated eigenvector as described in the next section. Note that we calculate just one iteration for power iteration in this paper, because it is reported in (Miyato et al. 2016; 2018) that one iteration is sufficient for computing accurate $\tilde{H} \tilde{v}$ and increasing the iteration does not have an effects.

ϵ -Rampup

Although $\phi_{T\text{-adv}}$ should satisfy $\|\phi\| \leq \epsilon$, $\phi_{T\text{-adv}}$ is actually given as the parameter satisfying $\|\phi\| = \epsilon$ when solving (11), because the Hessian of KL divergence is semi-positive definite.

In the case of VAT, the smoothing with $\|r\|_2 = \epsilon$ means that the model should satisfy consistency between the outputs with respect to the original data and the data on the surface of the ϵ -ball centered on the original data.

We propose a technique, ϵ -rampup, that enhances the effect of the smoothing not only on the boundary but also inside the boundary with small computational costs. The technique is to ramp up ϵ from 0 to a pre-defined value during training, and the parameters of the adversarial transformation are determined by solving $\arg \max_{\phi} \{\tilde{\Delta}_{\text{KL}}(f_\phi, x, \theta); \|\phi\| = \epsilon\}$. One can approximately solve this by normalizing the approximated eigenvector to satisfy $\|\phi\| = \epsilon$ after the power iteration.

Although there are many techniques to ramp up or anneal parameters (Tarvainen and Valpola 2017; Smith 2017), we adopt the procedure used in Mean Teacher (Tarvainen and Valpola 2017). Mean Teacher utilizes a sigmoid shape function, $\exp(-5(1-x)^2)$; $x = [0, 1]$, for ramping up the regularization coefficient, and we adopt it for ϵ .

We show the pseudocode of the generation process of $\phi_{T\text{-adv}}$ in Algorithm 1.

Algorithm 1 Generation of $\phi_{T\text{-adv}}$

```
1: Input: Data  $x \in \mathcal{X}$ ; transformation functions  $\{f^{(i)}\}_{i=1}^n \in \mathcal{F}$ ; scalar parameters  $\{\epsilon^{(i)}\}_{i=1}^n$  and  $\xi$ 
2: Output: function parameters  $\{\phi_{T\text{-adv}}^{(i)}\}_{i=1}^n$ 
3: Make copy of data  $\hat{x} \leftarrow x$ 
4: for  $i = 1, \dots, n$  do
5:   initialize  $v^{(i)}$  as a random unit vector
6:    $\phi^{(i)} = \phi_{\text{id}}^{(i)} + \xi v^{(i)}$ 
7:    $\hat{x} \leftarrow f_{\phi^{(i)}}^{(i)}(\hat{x})$ 
8: end for
9: for  $i = 1, \dots, n$  do
10:   $v^{(i)} \leftarrow \nabla_{v^{(i)}} \text{KL}[p_{\theta}(y|x) \| p_{\theta}(y|\hat{x})]$ 
11:   $\epsilon^{(i)} \leftarrow \text{Rampup}(\epsilon^{(i)})$ 
12:  Normalize  $v^{(i)}$  to satisfy  $\|\phi^{(i)}\| = \epsilon^{(i)}$ 
13:   $\phi_{T\text{-adv}}^{(i)} = \phi_{\text{id}}^{(i)} + v^{(i)}$ 
14: end for
```

Evaluation on Synthetic Dataset

We show the smoothing effect of RAT with a toy problem, a moons dataset, in Figure 1. We make 10 labeled samples and 30 unlabeled samples for each moon. $p_{\theta}(y|x)$ consists of a three-layer neural network with ReLU non-linearity. All hidden layers have 128 units. θ is optimized with Adam optimizer for 500 iterations with default parameters suggested in (Kingma and Ba 2014). In each iteration, we use all samples for updating θ . We treat this toy problem as if we already know appropriate class-invariant transformations; we adopt class-wise rotation along each moon as $\tilde{f}_{\phi}(x) = R(x, \phi^{(1)}) + \phi^{(2)}$; $\|\phi^{(2)}\|_2 \leq 0.3, |\phi^{(1)}| \leq 10^\circ$ for illustration purpose. Note that we do not ramp up ϵ in this experiment.

VAT with a small ϵ draws the decision boundary crossing \mathcal{S}_k as shown in Figure 1 (b). When we adopt a larger ϵ , VAT cannot smooth the output distribution within the ϵ -ball as shown in Figure 1 (c), because the larger ϵ allows the unexpected transformation $f : \mathcal{S}_k \rightarrow \mathcal{S}_{k'}$, and causes inconsistency. On the other hand, RAT draws the decision boundary along \mathcal{S}_k .

In this toy problem, we utilize \mathcal{S}_k and it is equal to using the label information implicitly. Therefore, in a next section, we evaluate RAT using realistic situations, where we do not know \mathcal{S}_k .

Experiments

We evaluate the effectiveness of RAT and ϵ -rampup through three experiments on a semi-supervised image classification task: (i) evaluation of composite transformations, (ii) evaluation of ϵ -rampup for VAT and RAT, and (iii) comparison of RAT to baseline methods. As the baseline methods, we use Π -Model (Sajjadi, Javanmardi, and Tasdizen 2016; Laine and Aila 2017), Pseudo-label (Lee 2013), Mean Teacher (Tarvainen and Valpola 2017), and VAT (Miyato et al. 2018). Note that VAT utilizes entropy regularization in all experiments. To evaluate on realistic situations, we follow the setting proposed by Oliver et al. (2018). We use Py-

Torch (Paszke et al. 2017) to implement and evaluate SSL algorithms, and we carefully reproduced the results of Oliver et al. (2018). All hyperparameters for SSL algorithms are adopted the same as in Oliver et al. (2018) except that we do not use L1 and L2 regularization.

For all experiments, we used the same Wide ResNet architecture, depth 28 and width 2 (Zagoruyko and Komodakis 2016), and we use the CIFAR-10 (Krizhevsky and Hinton 2009) and SVHN (Netzer et al. 2011) datasets for evaluation. CIFAR-10 has 50,000 training data and 10,000 test data, and we split training data into a train/validation set, 45,000 data for training and 5,000 data for validation. SVHN has 73,257 data for training and 26,032 data for testing. We also split training data into 65,931 data for training and 7,326 data for validation. For the semi-supervised setting, we further split training data into labeled data and unlabeled data.

We utilize standard preprocessing and data augmentations for training, following Oliver et al. (2018). For SVHN, we normalize the pixel value into the range $[-1, 1]$, and use random translation by up to 2 pixels as data augmentation. For CIFAR-10, we apply ZCA normalization (Krizhevsky and Hinton 2009) and global contrast normalization as normalization, and random horizontal flipping, random translation by up to 2 pixels, and Gaussian noise injection with standard deviation 0.15 as data augmentation.

We report the mean and standard deviation of error rates over five trials with test sets. The test error rates are evaluated with the model that has the minimum error rate on validation sets. The evaluation on validation set is executed every 25,000 training iterations.

Implementation Details of RAT

We tested three types of data transformations, all of which are commonly used in data augmentation in image classification tasks. All transformations f_{ϕ} discussed below satisfy the conditions C1 and C2. We evaluated different types of composite transformations as is discussed below.

Noise Injection The noise injection is represented as $f_{\phi}(x) = x + \phi$. We define the norm for the parameters of the noise injection as the L2 norm, $\|\phi\|_2$. This is equal to the formulation of VAT.

Spatial Transformation We consider three spatial transformations with different degrees of freedom: affine transformation, thin plate spline (Bookstein 1989), and flow field. All these transformations shift the pixel position (u, v) by offset vector $(\delta u, \delta v)$ to give $(u + \delta u, v + \delta v)$. The pixel values of the transformed image are calculated by bilinear interpolation. The details of these transformations are provided in (Jaderberg et al. 2015).

The difference between transformations is the degrees of freedom to calculate the offset vectors as follows: affine transformation has six parameters, thin plate spline has parameters proportional to the number of control points m , and the flow field directly has the local offset vectors as parameters, meaning that the number of parameters of the flow field is proportional to the spatial resolution of the image. We set m to 16 in all experiments, which means we employ a 4×4 grid.

Table 1: The parameters ϵ for transformations.

Transformations	ϵ
Channel-wise weighting	0.001
Affine transformation	0.6
Thin plate spline	1
Flow field	0.01

Table 2: VAT and RAT shared parameters. ϵ is required 6.0 for CIFAR-10 and 1.0 for SVHN.

Parameters	values
Initial learning rate	0.003
Max consistency coefficient λ	0.3
noise injection ϵ	6.0 or 1.0
ξ	10^{-6}
Entropy penalty multiplier	0.06

We define the norm for the parameters of the flow field transformations as the L2 norm, $\sqrt{\sum_i^N (\delta u_i^2 + \delta v_i^2)}$, where N is the number of pixels. The norm for thin plate spline is also defined as the L2 norm of offset vectors for the control points.

Affine transformation is a linear operator in homogeneous coordinates, and the norm is given as the operator norm. Thus, $\|\phi\|$ is calculated as the maximum singular value of an affine transformation matrix.

Color Distortion Color distortion is an effective augmentation method for image classification tasks. Among many methods for color distortion, we use a simple way, channel-wise weighting, $f_\phi(x_c) = \phi_c x_{c,i}$, where $x_{c,i}$ is the pixel value of the c -th channel of the i -th pixel, and ϕ_c is the scalar value for each channel. This transformation is described as the linear operator, and we define the norm as the operator norm. Note that channel-wise weighting is represented as the multiplication of a diagonal matrix and a pixel value, and the operator norm is calculated as $\max_c |\phi_c|$.

Evaluation of Composite Transformations

Since the performance of RAT depends on the combination of the transformations, we report the effect of a combination of functions by adding transformations to VAT.

We first seek good ϵ for each transformation with a grid search on CIFAR-10 with 4,000 labeled data, from 0.001 to 0.01 with 0.001 step size for channel-wise weighting, from 0.1 to 1 with a 0.1 step size for affine transformation and thin plate spline, and from 0.01 to 0.1 with a 0.01 step size for flow field. We show the grid search results in Table 1. Other parameters such as ξ , λ , and parameters for optimization are the same as VAT suggested in (Oliver et al. 2018). We summarize the parameters in Table 2. Note that ϵ -rampup is not utilized in this experiment, because this experiment explores the effect of composite transformations.

The results of adopting various transformations with CIFAR-10 with 4,000 labeled data and SVHN with 1,000 labeled data are shown in Table 3. All transformations and combinations improve the performance from supervised

learning. However, channel-wise weighting and flow field increase the test error rates from VAT in CIFAR-10. For CIFAR-10, since we apply the ZCA normalization and the global contrast normalization, channel-wise weighting for the space of normalized data is an unnatural transformation for natural images. On the other hand, flow field is the transformation that can break object structure. Unlike a simple structure like the data in SVHN, the detail structure is the important feature for general objects. Thus, flow field induces the unfavorable effect.

Affine transformation achieves the best performance of all spatial transformations, and the results are induced by the low degree of freedom of the affine transformation. The affine transformation has the lowest degree of freedom among the three. In particular, the affine transformation preserves points, straight lines and planes. In other words, the affine transformation preserves the basic structure of the objects. Therefore, except for extreme cases, the affine transformation ensures that the class of transformed data is the same as the class of original data. This fact matches the strategy of RAT, which is that the output distribution of the data belonging to the same class should be close.

The effect of combining channel-wise weighting and each spatial transformation is less effective. This fact means that combining channel-wise weighting does not expand the smoothing effect to a meaningful range. Indeed, the difference of combining channel-wise weighting is within the standard deviation.

Evaluation of ϵ -Rampup

We evaluate the effectiveness of ϵ -rampup with CIFAR-10 with 4,000 labeled data. Since ϵ -rampup is the technique for VAT and RAT, we compare the results with and without ramping up for VAT and RAT. We ramp up ϵ for 400,000 iterations. We utilize the composite transformations consisting of affine transformation and noise injection for RAT, and the hyperparameters are as in Table 2.

The results are shown in Table 4. For both VAT and RAT, ϵ -rampup brings a positive effect. As interesting effects, ϵ -rampup allows for a large ϵ . Since the smoothing effect reaches within the range of $\|\phi\| \leq \epsilon$ by ramping up, VAT and RAT with ramping up work well with a relatively large ϵ .

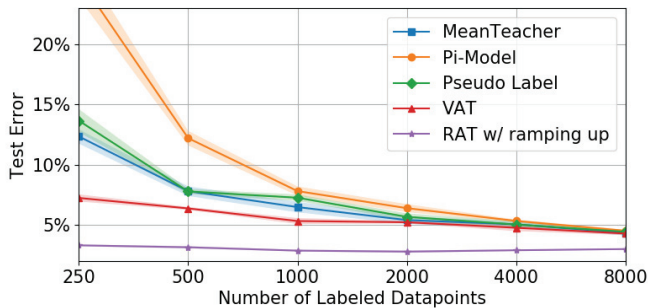
Comparison of RAT to Baselines

We show the effectiveness of RAT by comparing it with baseline methods and non-adversarial version of RAT (Random Transformation). We use the composite transformations consisting of affine transformation and noise injection for RAT. The hyperparameters of RAT and ϵ for transformations are the same as shown in Tables 1 and 2, respectively, except that we set ϵ for affine transformation to 1 for RAT with ramping up with CIFAR-10, and set ϵ for noise injection to 5 for RAT and VAT with ϵ -rampup with SVHN. Note that all the parameters of the random transformation are the same as RAT.

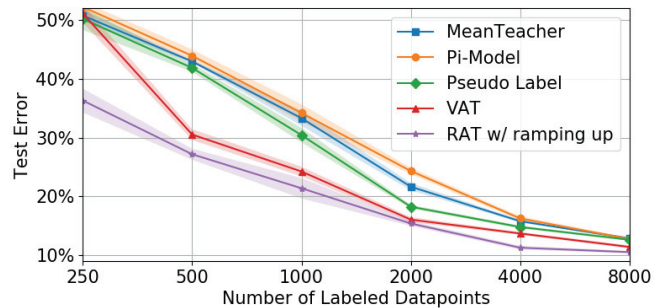
In Table 5, we show the comparison results with standard SSL settings, CIFAR-10 with 4,000 labeled data and SVHN with 1,000 labeled data. On both datasets, RAT improves

Table 3: Test error rates of RAT with various transformations on CIFAR-10 with 4,000 labeled data and SVHN with 1,000 labeled data. All settings of RAT include noise injection as transformation, and all of the hyperparameters and experiment settings of VAT and RAT are the same, except for the inherent parameters of RAT.

Methods	Channel-wise	Affine	Thin Plate Spline	Flow Field	CIFAR-10	SVHN
Supervised					20.35±0.14%	12.33±0.25%
VAT					13.68±0.25%	5.32±0.25%
RAT	✓				14.33±0.44%	5.19±0.29%
		✓			11.70±0.32%	3.10±0.12%
			✓		13.06±0.44%	4.12±0.11%
				✓	14.27±0.41%	4.93±0.35%
	✓	✓			11.32±0.44%	3.14±0.12%
	✓		✓	13.46±0.37%	4.06±0.16%	
	✓			✓	14.17±0.27%	4.85±0.12%



(a) SVHN



(b) CIFAR-10

Figure 2: Test error rates obtained by varying the number of labeled data. Shaded regions indicate standard deviation over five trials.

Table 4: Comparison between with and without ϵ -rampup on CIFAR-10 with 4,000 labeled data. RAT has two ϵ , one for the noise injection and one for the affine transformation.

Methods	maximum ϵ	CIFAR-10
VAT	6	13.68±0.25%
VAT	10	14.31±0.33%
VAT w/ ϵ -rampup	10	13.26±0.20%
RAT	(6, 0.6)	11.70±0.32%
RAT	(6, 1)	12.68±0.29%
RAT w/ ϵ -rampup	(6, 1)	11.26±0.34%

test error rates more than 2% from the best baseline method, VAT. Furthermore, RAT also improves the error rates from the random transformation. The results prove the importance of the adversariality.

We evaluate the test error rates of each method with a varying number of labeled data from 250 to 8,000. The results are shown in Figure 2. RAT consistently outperforms other methods for all the range on both datasets. Remarkably, RAT significantly improves the test error rate in CIFAR-10 with 250 labeled data compared with the best result of baseline methods, from 50.20±1.88% to 36.31±2.03%. We believe that the improvement results from

Table 5: Test error rates of RAT and baseline methods on CIFAR-10 with 4,000 labeled data and SVHN with 1,000 labeled data.

Methods	CIFAR-10	SVHN
	4,000 Labels	1,000 Labels
Supervised	20.35±0.14%	12.33±0.25%
II-Model	16.24±0.38%	7.81±0.39%
Pseudo-Label	14.78±0.26%	7.26±0.27%
Mean Teacher	15.77±0.22%	6.48±0.44%
VAT	13.68±0.25%	5.32±0.25%
VAT w/ ϵ -rampup	13.26±0.20%	5.17±0.26%
Random Transformation	12.71±0.82%	6.06±0.50%
RAT	11.70±0.32%	3.10±0.12%
RAT w/ ϵ -rampup	11.26±0.34%	2.86±0.07%

appropriately smoothed model prediction along the underlying data distribution. The experimental fact of VAT’s underperformance, namely 15% degraded CIFAR-10 test error rate compared to RAT, is a clear indication that adversarial, class-invariant transformation provides far better consistency regularization than isotropic, adversarial noise superposition.

Lastly, we compare our results with very recently reported results of Mixup-based SSL method, MixMatch (Berthelot

et al. 2019). Interestingly, RAT is comparative or superior to MixMatch on SVHN, while MixMatch is superior to RAT on CIFAR-10. But, we should point out that this comparison does not seem fair because the experimental settings of MixMatch are different in several respects from ours and others shown in Figure 2. In MixMatch paper, the authors took exponential moving average of models, and the final results were given by the median of the last 20 test error rates measured every 2^{16} training iterations. These settings, that are missing in our experiments, seem to partly boost their performance. Readers are referred to Appendix for the detailed comparison with MixMatch.

Related Work

There are many SSL algorithms such as graph-based methods (Kipf and Welling 2017; Bengio, Delalleau, and Le Roux 2006), generative model-based methods (Kingma et al. 2014; Kumar, Sattigeri, and Fletcher 2017), and regularization-based methods (Miyato et al. 2016; Sajjadi, Javanmardi, and Tasdizen 2016; Laine and Aila 2017; Tarvainen and Valpola 2017; Berthelot et al. 2019).

Label propagation (Bengio, Delalleau, and Le Roux 2006) is a representative graph-based SSL method. Modern graph-based approaches utilize neural networks for graphs (Kipf and Welling 2017). Graph-based methods demonstrate the effectiveness, when elaborated graph structure is given.

Generative models such as VAE and GAN are now popular frameworks for the SSL setting (Kingma et al. 2014; Kumar, Sattigeri, and Fletcher 2017). Although generative model-based SSL methods typically have to train additional models, they tend to show remarkable gain in test performance.

Regularization-based methods are comparatively much more tractable, and can be utilized in arbitrary models. Next, we review three regularization-based methods closely related to RAT.

Consistency Regularization

Consistency regularization is a method imposing the consistency between the outputs of one model with respect to a typically unlabeled data and its perturbed counterpart, or the outputs of two models with respect to the same input. One of the simplest ways of constructing consistency regularization is to add stochastic perturbation to the data, $x \rightarrow \hat{x}$, as follows:

$$\min_{\theta} d(p_{\theta}(y|x), p_{\theta}(y|\hat{x})) \quad (13)$$

where d is some distance functions; e.g., Euclidean distance or KL divergence. VAT (Miyato et al. 2016; 2018) and RAT are classified in this category.

The random transformation-based consistency regularization techniques, the work (Sajjadi, Javanmardi, and Tasdizen 2016) and ‘‘II-Model’’ (Laine and Aila 2017) are vary similar and famous ones. We refer to these models as II-Model in this paper. One can view II-Model as non-adversarial version of RAT.

The other way of constructing consistency regularization is to utilize dropout (Srivastava et al. 2014). Let θ_1 and θ_2 be

the randomly selected parameters through dropout. Dropout as consistency regularization is represented as follows:

$$\min_{\theta_1, \theta_2 \sim \theta} d(p_{\theta_1}(y|x), p_{\theta_2}(y|x)). \quad (14)$$

Mean Teacher (Tarvainen and Valpola 2017) is a successful method that employs consistency regularization between two models. It makes a teacher model by exponential moving average of the parameters of a student model, and imposes the consistency between the teacher and the student. Although Mean Teacher can be combined with other SSL methods, the combination sometimes impairs the model performance as reported in (Berthelot et al. 2019).

VAT (Miyato et al. 2016; 2018) is a very effective consistency regularization method. The advantage of VAT lies in the generation of adversarial noises, and the adversariality leads to isotropic smoothness around sampled datapoints. VAT also show the effectiveness in natural language processing tasks (Miyato, Dai, and Goodfellow 2016).

Entropy Regularization

Entropy regularization is a way to bring low entropy on $p_{\theta}(y|x)$ to make model prediction more discriminative, and is known to give low-density separation (Grandvalet and Bengio 2005; Chapelle, Scholkopf, and Zien 2006). The entropy regularization term is:

$$\min_{\theta} -p_{\theta}(y|x) \log p_{\theta}(y|x). \quad (15)$$

This regularization is often combined with other SSL algorithms (Sajjadi, Javanmardi, and Tasdizen 2016; Miyato et al. 2018) and a combined method, VAT+entropy regularization, shows the state-of-the-art results in (Oliver et al. 2018).

Mixup

Mixup (Zhang et al. 2018) is a powerful regularization method that is very recently used for SSL (Verma et al. 2019a; Berthelot et al. 2019; Verma et al. 2019b). Mixup blends two different data, x_1 and x_2 , and their labels, y_1 and y_2 , as follows:

$$\hat{x} = \beta x_1 + (1 - \beta)x_2, \hat{y} = \beta y_1 + (1 - \beta)y_2, \quad (16)$$

where β is the scalar value sampled from Beta distribution. In a semi-supervised setting, \hat{y} is calculated as a blend between a label and a prediction or predictions, $\beta y_1 + (1 - \beta)p_{\theta}(y|x_2)$ or $\beta p_{\theta}(y|x_1) + (1 - \beta)p_{\theta}(y|x_2)$, and a regularization term is described as consistency regularization, $\min_{\theta} d(\hat{y}, p_{\theta}(y|\hat{x}))$.

Conclusion

We proposed an SSL framework, called RAT, Regularization framework based on Adversarial Transformation. RAT aims to smooth model output along the underlying data distribution within a given class based on recent advancement of generation of adversarial inputs that stem from unlabeled data. Instead of just superposing adversarial noise, RAT uses a wider range of data transformations, each of which leaves class label invariant. We further propose use of composite

transformations and a technique, called ϵ -rampup, to enlarge the area in which the smoothing is effective without sacrificing computational cost. We experimentally show that RAT significantly outperform the baseline methods including VAT in semi-supervised image classification tasks. RAT is especially robust against reduction of labeled samples, compared to other methods. As a future work, we would like to replace the designing of composite functions by black box function optimization.

References

- Alaifari, R.; Alberti, G. S.; and Gauksson, T. 2018. Adef: An iterative algorithm to construct adversarial deformations. *International Conference on Learning Representations*.
- Bengio, Y.; Delalleau, O.; and Le Roux, N. 2006. *Label Propagation and Quadratic Criterion, chapter 11*. MIT Press.
- Berthelot, D.; Carlini, N.; Goodfellow, I.; Papernot, N.; Oliver, A.; and Raffel, C. 2019. Mixmatch: A holistic approach to semi-supervised learning. *arXiv preprint arXiv:1905.02249*.
- Bookstein, F. L. 1989. Principal warps: Thin-plate splines and the decomposition of deformations. *IEEE Transactions on pattern analysis and machine intelligence* 11(6):567–585.
- Chapelle, O.; Scholkopf, B.; and Zien, A. 2006. *Semi-supervised learning*. MIT Press.
- Golub, G. H., and Van der Vorst, H. A. 2001. Eigenvalue computation in the 20th century. In *Numerical analysis: historical developments in the 20th century*. Elsevier. 209–239.
- Goodfellow, I. J.; Shlens, J.; and Szegedy, C. 2014. Explaining and harnessing adversarial examples. *International Conference on Learning Representations*.
- Grandvalet, Y., and Bengio, Y. 2005. Semi-supervised learning by entropy minimization. In *Advances in neural information processing systems*, 529–536.
- Jaderberg, M.; Simonyan, K.; Zisserman, A.; and Kavukcuoglu, K. 2015. Spatial transformer networks. In *Advances in neural information processing systems*, 2017–2025.
- Kingma, D. P., and Ba, J. 2014. Adam: A method for stochastic optimization. *International Conference on Learning Representations*.
- Kingma, D. P.; Mohamed, S.; Rezende, D. J.; and Welling, M. 2014. Semi-supervised learning with deep generative models. In *Advances in neural information processing systems*, 3581–3589.
- Kipf, T. N., and Welling, M. 2017. Semi-supervised classification with graph convolutional networks. *International Conference on Learning Representations*.
- Krizhevsky, A., and Hinton, G. 2009. Learning multiple layers of features from tiny images. Technical report, Citeseer.
- Kumar, A.; Sattigeri, P.; and Fletcher, T. 2017. Semi-supervised learning with gans: Manifold invariance with improved inference. In *Advances in Neural Information Processing Systems*, 5534–5544.
- Laine, S., and Aila, T. 2017. Temporal ensembling for semi-supervised learning. *International Conference on Learning Representations*.
- Lee, D.-H. 2013. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *ICML Workshop on Challenges in Representation Learning*, volume 3, 2.
- Miyato, T.; Maeda, S.-i.; Koyama, M.; Nakae, K.; and Ishii, S. 2016. Distributional smoothing with virtual adversarial training. *International Conference on Learning Representations*.
- Miyato, T.; Maeda, S.-i.; Koyama, M.; and Ishii, S. 2018. Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 41(8):1979–1993.
- Miyato, T.; Dai, A. M.; and Goodfellow, I. 2016. Adversarial training methods for semi-supervised text classification. *arXiv preprint arXiv:1605.07725*.
- Netzer, Y.; Wang, T.; Coates, A.; Bissacco, A.; Wu, B.; and Ng, A. Y. 2011. Reading digits in natural images with unsupervised feature learning. In *NIPS Workshop on Deep Learning and Unsupervised Feature Learning*.
- Oliver, A.; Odena, A.; Raffel, C. A.; Cubuk, E. D.; and Goodfellow, I. 2018. Realistic evaluation of deep semi-supervised learning algorithms. In *Advances in Neural Information Processing Systems*, 3235–3246.
- Paszke, A.; Gross, S.; Chintala, S.; Chanan, G.; Yang, E.; DeVito, Z.; Lin, Z.; Desmaison, A.; Antiga, L.; and Lerer, A. 2017. Automatic differentiation in PyTorch. In *NIPS Autodiff Workshop*.
- Sajjadi, M.; Javanmardi, M.; and Tasdizen, T. 2016. Regularization with stochastic transformations and perturbations for deep semi-supervised learning. In *Advances in Neural Information Processing Systems*, 1163–1171.
- Smith, L. N. 2017. Cyclical learning rates for training neural networks. In *2017 IEEE Winter Conference on Applications of Computer Vision (WACV)*, 464–472. IEEE.
- Srivastava, N.; Hinton, G.; Krizhevsky, A.; Sutskever, I.; and Salakhutdinov, R. 2014. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research* 15(1):1929–1958.
- Tarvainen, A., and Valpola, H. 2017. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *Advances in Neural Information Processing Systems*, 1195–1204.
- Tsuzuku, Y., and Sato, I. 2019. On the structural sensitivity of deep convolutional networks to the directions of fourier basis functions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 51–60.
- Verma, V.; Lamb, A.; Beckham, C.; Najafi, A.; Mitliagkas, I.; Lopez-Paz, D.; and Bengio, Y. 2019a. Manifold mixup: Better representations by interpolating hidden states. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, 6438–6447. PMLR.
- Verma, V.; Lamb, A.; Kannala, J.; Bengio, Y.; and Lopez-Paz, D. 2019b. Interpolation consistency training for semi-supervised learning. *arXiv preprint arXiv:1903.03825*.
- Zagoruyko, S., and Komodakis, N. 2016. Wide residual networks. In *The British Machine Vision Conference*.
- Zhang, H.; Cisse, M.; Dauphin, Y. N.; and Lopez-Paz, D. 2018. mixup: Beyond empirical risk minimization. *International Conference on Learning Representations*.