

Dual Student: Breaking the Limits of the Teacher in Semi-supervised Learning

Zhanghan Ke^{1,2} * Daoye Wang² Qiong Yan² Jimmy Ren² Rynson W.H. Lau¹

¹ City University of Hong Kong

² SenseTime Research

Abstract

Recently, consistency-based methods have achieved state-of-the-art results in semi-supervised learning (SSL). These methods always involve two roles, an explicit or implicit teacher model and a student model, and penalize predictions under different perturbations by a consistency constraint. However, the weights of these two roles are tightly coupled since the teacher is essentially an exponential moving average (EMA) of the student. In this work, we show that the coupled EMA teacher causes a performance bottleneck. To address this problem, we introduce Dual Student, which replaces the teacher with another student. We also define a novel concept, stable sample, following which a stabilization constraint is designed for our structure to be trainable. Further, we discuss two variants of our method, which produce even higher performance. Extensive experiments show that our method improves the classification performance significantly on several main SSL benchmarks. Specifically, it reduces the error rate of the 13-layer CNN from 16.84% to 12.39% on CIFAR-10 with 1k labels and from 34.10% to 31.56% on CIFAR-100 with 10k labels. In addition, our method also achieves a clear improvement in domain adaptation.

1. Introduction

Deep supervised learning has gained significant success in computer vision tasks, which leads the community to challenge larger and more complicated datasets like ImageNet [28] and WebVision [17]. However, obtaining full labels for a huge dataset is usually a very costly task. Hence, more attention is now drawn on deep semi-supervised learning (SSL). In order to utilize unlabeled data, many methods in traditional machine learning have been proposed [35], and some of them are successfully adapted to deep learning. In addition, some latest techniques like self-training [5] and Generative Adversarial Networks (GANs) [30, 22, 19] have been utilized for deep SSL with promising results. A primary track of recent deep semi-supervised methods [27, 16, 7, 32] can be summarized as consistency-based




Figure 1: Teacher-Student versus Dual Student. The teacher (T) in Teacher-Student is an EMA of the student (S), imposing a consistency constraint on the student. Their weights are tightly coupled. In contrast, a bidirectional stabilization constraint is applied between the two students (S and S') in Dual Student. Their weights are loosely coupled.

methods. In this type of methods, two roles are commonly created, either explicitly or implicitly: a teacher model and a student model (i.e., a Teacher-Student structure). The teacher guides the student to approximate its performance under perturbations. The perturbations could come from the noise of the input or the dropout layer [31], etc. A consistency constraint is then imposed on the predictions between two roles, and forces the unlabeled data to meet the *smoothness assumption* of semi-supervised learning.

The teacher in the Teacher-Student structure can be summarized as being generated by an exponential moving average (EMA) of the student model. In the VAT Model [20] and the Π Model [16], the teacher shares the same weights as the student, which is equivalent to setting the averaging coefficient to zero. The Temporal Model [16] is similar to Π Model except that it also applies an EMA to accumulate the historical predictions. The Mean Teacher [32] applies an EMA to the student to obtain an ensemble teacher. In this work, we show that the two roles in the Teacher-Student structure are tightly coupled and the degree of the coupling increases as the training goes on. This phenomenon leads to a performance bottleneck since a coupled EMA teacher is not sufficient for the student.

To overcome this problem, the knowledge coming from another independent model should help. Motivated by this observation, we replace the EMA teacher by another student model. The two students start from different initial

*kezhanghan@outlook.com

states and are optimized through individual paths during training. Hence, their weights will not be tightly coupled and each learns its own knowledge. What remains unclear is how to extract and exchange knowledge between the students. Naively, adding a consistency constraint may lead to the two models collapsing into each other. Thus, we define the *stable sample* and propose a stabilization constraint for effective knowledge exchange. Our method improves the performance significantly on several main SSL benchmarks. Fig. 1 demonstrates the Teacher-Student structure and our Dual Student structure.

In summary, the main contributions of this work include:

- We demonstrate that the coupled EMA teacher causes a performance bottleneck of the existing Teacher-Student methods.
- We define the *stable samples* of a model and propose a novel stabilization constraint between models.
- We propose a new SSL structure, Dual Student, and discuss two variants of Dual Student with higher performances.
- Extensive experiments are conducted to evaluate the performance of our method on several benchmarks and in different tasks.

2. Related Work

2.1. Overview

Consistency-based SSL methods are derived from the network noise regularization [29]. Goodfellow *et al.* [7] first showed the advantage of adversarial noise over random noise. Miyato *et al.* [20] further explored this idea for unlabeled data and generated virtual adversarial samples for the implicit teacher, while Park *et al.* [24] proposed a virtual adversarial dropout based on [31]. In addition to noise, the quality of targets for the consistency constraint is also vital in this process. Bachman *et al.* [2] and Rasmus *et al.* [27] showed the effectiveness of regularizing the targets. Laine *et al.* then proposed the internally consistent II Model and Temporal Model in [16]. Tarvainen took advantage of averaging model weights [25] to obtain an explicit ensemble teacher [32] for generating targets. Some works derived from the traditional methods also improve the consistency-based SSL. Smooth Neighbor by Luo *et al.* [18] utilized the connection between data points and built a neighbor graph to cluster data more tightly. Athiwaratkun *et al.* [1] modified the stochastic weight averaging (SWA) [13] to obtain a stronger ensemble teacher faster. Qiao *et al.* [26] proposed Deep Co-Training [26], by adding a consistency constraint between independent models.

2.2. Teacher-Student Structure

The most common structure of recent SSL methods is the Teacher-Student structure. It applies a consistency con-

straint between a teacher model and a student model to learn knowledge from unlabeled data. Formally, we assume that a dataset \mathcal{D} consists of an unlabeled subset and a labeled subset. Let θ' denote the weights of the teacher, and θ denote the weights of the student. The consistency constraint is defined as:

$$\mathcal{L}_{con} = \mathbb{E}_{x \in \mathcal{D}} \mathcal{R}(f(\theta, x + \zeta), \mathcal{T}_x), \quad (1)$$

where $f(\theta, x + \zeta)$ is the prediction from model $f(\theta)$ for input x with noise ζ . \mathcal{T}_x is the consistency target from the teacher. $\mathcal{R}(\cdot, \cdot)$ measures the distance between two vectors, and is usually set to mean squared error (MSE) or KL-divergence. Previous works have proposed several ways to generate \mathcal{T}_x .

II Model: In II Model, the implicit teacher shares parameters with the student. It forwards a sample x twice with different random noise ζ and ζ' in each iteration, and treats the prediction of $x + \zeta'$ as \mathcal{T}_x .

Temporal Model: While II Model needs to forward a sample twice in each iteration, Temporal Model reduces this computational overhead by using EMA to accumulate the predictions over epochs as \mathcal{T}_x . This approach could reduce the prediction variance and stabilize the training process.

Mean Teacher: Temporal Model needs to store a record for each sample, and the target \mathcal{T}_x gets updated only once per epoch while the student is updated multiple times. Hence, Mean Teacher defines an explicit teacher by an EMA of the student and update its weights in each iteration before generating \mathcal{T}_x .

VAT Model: Although random noise is effective in previous methods, VAT Model adopts the adversarial noise to generate better \mathcal{T}_x for the consistency constraint.

2.3. Deep Co-Training

It is known that fusing knowledge from multiple models could improve performance in SSL [34]. However, directly adding the consistency constraint between models results in the models collapsing into each other. Deep Co-Training addressed this issue by utilizing the *Co-Training assumption* from the traditional Co-Training algorithm [3]. It treats the features from the convolutional layers as a view of the input and uses the adversarial samples from other collaborators to ensure that view differences exist among the models. Consistent predictions can then be used for training. However, this strategy requires generating adversarial samples of each model in the whole process, which is complicated and time-consuming.

Our method also has interactions between models to break the limits of the EMA teacher, but there are two major differences between our method and Deep Co-Training. First, instead of enforcing the consistency constraint and the different-views constraint, we only extract reliable knowledge of the models and exchange them by a more effective

stabilization constraint. Second, our method is more efficient since we do not need the adversarial samples.

3. Limits of the EMA Teacher

One fundamental assumption in SSL is the *smoothness assumption* - “If two data points in a high-density region are close, then so should be the corresponding outputs” [4]. All existing Teacher-Student methods utilize unlabeled data according to this assumption. In practice, if x and \bar{x} are generated from a sample with different small perturbations, they should have consistent predictions by the corresponding teacher and student. Previous methods achieve this by the consistency constraint and have mainly focused on generating more meaningful targets through ensemble or well-designed noise.

However, previous works neglect that the teacher is essentially an EMA of the student. Hence, their weights are tightly coupled. Formally, the teacher weights θ' are an ensemble of the student weights θ in a successive training step t with a smoothing coefficient $\alpha \in [0, 1]$:

$$\theta'_t = \alpha \theta'_{t-1} + (1 - \alpha) \theta_t. \quad (2)$$

In II Model and VAT Model, as α is set to zero, θ' is equal to θ . Temporal Model improves II Model by an EMA on historical predictions, but its teacher still shares weights with the student. As for Mean Teacher, the updates of the student weights decreases as the model converges, i.e., $|\theta_t - \theta_{t-1}|$ becomes smaller and smaller as the number of training steps t increases. Theoretically, it can be proved that the EMA of a converging sequence converges to the same limit as the sequence, which is shown in Appendix A (Supplementary). Thus, the teacher will be very close to the student when the training process converges. In all the above cases, the coupling fact between the teacher and the student is obvious.

To further visualize it, we train two structures on the CIFAR-10 SSL benchmark. One contains a student and an EMA teacher (named S_{ema}) while the other contains two independent models (named S_{split}). We then calculate the Euclidean distance of the weights and predictions between the two models in each structure. Fig. 2 shows the results. As expected, the EMA teacher in S_{ema} is very close to the student, and their distance approaches zero with increasing epochs. In contrast, the two models in S_{split} always keep a larger distance from each other. These results confirm our conjecture that the EMA teacher is tightly coupled with the student. In addition, they also demonstrate that the two independent models are loosely coupled.

Due to the coupling effect between the two roles in the existing Teacher-Student methods, the teacher does not have more meaningful knowledge compared to the student. In addition, if the student has biased predictions for specific samples, the EMA teacher is most likely to accumulate the




Figure 2: Left: S_{ema} contains two models with similar weights, while the weights of the two models in S_{split} keep a certain distance. Right: The predictions of the two models in S_{split} keep a larger distance than those of S_{ema} .




Figure 3: Our method can alleviate the confirmation bias. f_1 and f_2 are the independent students from our Dual Student, while f_s is the student guided by the Mean Teacher. For a misclassified sample (belonging to *class1*), f_1 can correct it quickly with the knowledge from f_2 . However, f_s is unable to correct its prediction due to the wrong guidance from the EMA teacher.

mistakes and to enforce the student to follow, making the misclassification irreversible. This is a case of the confirmation bias [32]. Most methods apply a ramp-up operation for the consistency constraint to alleviate the bias, but it is inadequate to solve the problem. From this perspective, training independent models are also beneficial. Fig. 3 visualizes this inability of the EMA teacher. Three models, f_1 , f_2 , and f_s , are trained on a two-category task simultaneously. f_s is the student from Mean Teacher. f_1 and f_2 are two relatively independent but interactive models, representing the two students from our Dual Student structure (Section 4). They have the same initialization, while f_2 is different from them. The plot shows how the predictions of a sample from *class1* changes with epochs for these three models, which demonstrates that our method can alleviate the confirmation bias.

4. Dual Student

As analyzed above, the targets from an EMA teacher are not adequate to guide the student when the number of training steps t is large. Therefore, our method gains the loosely coupled targets by training two independent models simultaneously. However, the outputs of these two models

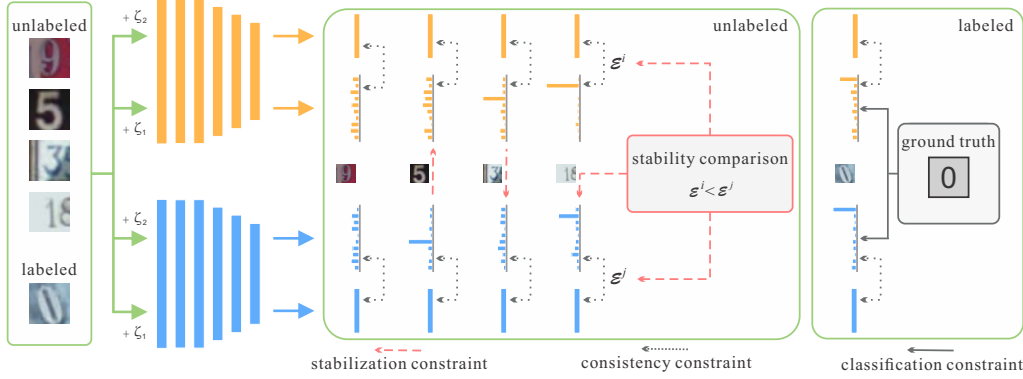


Figure 4: Dual Student structure overview. We train two student models separately. Each batch includes labeled and unlabeled data and is forwarded twice. The stabilization constraint based on the *stable samples* is enforced between the students. Each student also learns labeled data by the classification constraint and meets the *smooth assumption* by the consistency constraint.

may vary widely, and applying the consistency constraint directly will cause them to collapse into each other by exchanging the wrong knowledge. The EMA teacher does not suffer from this problem due to the coupling effect.

We propose an efficient way to overcome this problem, which is to exchange only reliable knowledge of the models. To put this idea into practice, we need to solve two problems. One is how to define and acquire reliable knowledge of a model. Another is how to exchange the knowledge mutually. To address them, we define the *stable sample* in Section 4.1 and then elaborate the derived stabilization constraint for training in Section 4.2.

4.1. Stable Sample

A model can be regarded as a decision function that can make reliable predictions for some samples but not for the others. We define the *stable sample* and treat it as the reliable knowledge of a model. A *stable sample* satisfies two conditions. First, according to the *smoothness assumption*, a small perturbation should not affect the prediction of this sample, i.e., the model should be smooth in the neighborhood of this sample. Second, the prediction of this sample is far from the decision boundary. This means that this sample has a high probability for the predicted label.

Definition 4.1 (Stable sample). Given a constant $\xi \in [0, 1)$, a dataset $\mathcal{D} \subseteq \mathbb{R}^m$ that satisfies the *smoothness assumption* and a model $f : \mathcal{D} \rightarrow [0, 1]^n$ that satisfies $\|f(x)\|_1 = 1$ for all $x \in \mathcal{D}$, x is a *stable sample* with respect to f if:

1. $\forall \bar{x} \in \mathcal{D}$ near x , their predicted labels are the same.
2. x satisfies the inequality: $\|f(x)\|_\infty > \xi$.¹

Def. 4.1 defines the *stable sample*, and Fig. 5 illustrates its conditions in details. Notice that the concept of the *stable*

¹ $\|\mathbf{a}\|_1 := \sum_{i=1}^n |a_i|$, $\|\mathbf{a}\|_\infty := \max_{i=1..n} |a_i|$, $\mathbf{a} = (a_1, a_2, \dots, a_n)$

sample is specific to the models. A data point x can be stable with respect to any one model but may not be to the others. This fact is a key to our stabilization constraint, and will be elaborated in Section 4.2. In addition to the criterion of whether a sample point x is stable or not, we would also like to know the degree of stability of a *stable sample* x . This can be reflected by the prediction consistency in its neighborhood. The more consistent the predictions are, the more stable x is.

4.2. Training by the Stabilization Constraint

We briefly introduce Dual Student structure before explaining the details on training. It contains two independent student models, which share the same network architecture with different initial states and are updated separately (Fig. 4). For our structure to be trainable, we derive a novel stabilization constraint from the *stable sample*.

In practice, we only utilize two close samples to approximate the conditions of the *stable sample* to reduce the computational overhead. Formally, we use θ^i and θ^j to represent weights of the two students. We first define a boolean function $\{condition\}_1$, which outputs 1 when the condition is true and 0 otherwise. Suppose \bar{x} is a noisy augmentation of a sample x . We then check whether x is a *stable sample* for student i :

$$\mathcal{R}_x^i = \{\mathcal{P}_x^i = \mathcal{P}_{\bar{x}}^i\}_1 \& (\{\mathcal{M}_x^i > \xi\}_1 \| \{\mathcal{M}_{\bar{x}}^i > \xi\}_1), \quad (3)$$

where $\mathcal{M}_x^i = \|f(\theta^i, x)\|_\infty$.

\mathcal{P}_x^i and $\mathcal{P}_{\bar{x}}^i$ are the predicted labels of x and \bar{x} , respectively, by student i . Hyperparameter ξ is a confidence threshold in $[0, 1)$. If the maximum prediction probability of sample x exceeds ξ , x is considered to be far enough from the classification boundary. We then use the Euclidean distance to measure the prediction consistency, to indicate the stability of x , as:

$$\mathcal{E}_x^i = \|f(\theta^i, x) - f(\theta^i, \bar{x})\|^2. \quad (4)$$




Figure 5: Illustration of the conditions for a *stable sample*. Consider three pairs of adjacent data points: (1) x_1 and \bar{x}_1 do not satisfy the 1st condition, (2) x_2 and \bar{x}_2 do not satisfy the 2nd condition, and (3) x_3 and \bar{x}_3 satisfy both conditions.

A smaller \mathcal{E}_x^i means that x is more stable to student i . The distance between the predictions of students i and j can be measured using the mean squared error (MSE) as:

$$\mathcal{L}_{mse}(x) = \|f(\theta^i, x) - f(\theta^j, x)\|^2. \quad (5)$$

Finally, the stabilization constraint for the student i on sample x is written as:

$$\mathcal{L}_{sta}^i(x) = \begin{cases} \{\mathcal{E}_x^i > \mathcal{E}_x^j\} \mathcal{L}_{mse}(x), & \mathcal{R}_x^i = \mathcal{R}_x^j = 1, \\ \mathcal{R}_x^j \mathcal{L}_{mse}(x), & \text{otherwise.} \end{cases} \quad (6)$$

We calculate the stabilization constraint for the student j in the same way. As we can see, the stabilization constraint changes dynamically depending on the outputs of the two students. There are three cases: (1) No constraint is applied if x is unstable for both students. (2) If x is stable only for student i , it can guide student j . (3) If x is stable for both students, the stability is calculated, and the constraint is applied from the more stable one to the other.

Following previous works, our Dual Student structure also imposes the consistency constraint in each student to meet the *smoothness assumption*. We also apply the decoupled top layers trick from the Mean Teacher, which splits the constraints for the classification and the smoothness.

To train Dual Student, the final constraint for student i is a combination of three parts: the classification constraint, consistency constraint in each model, and stabilization constraint between models, as:

$$\mathcal{L}^i = \mathcal{L}_{cls}^i + \lambda_1 \mathcal{L}_{con}^i + \lambda_2 \mathcal{L}_{sta}^i, \quad (7)$$

where λ_1 and λ_2 are hyperparameters to balance the constraints. Algorithm 1 summarizes the optimization process.

4.3. Variants of Dual Student

Here, we briefly discuss two variants of Dual Student, named *Multiple Student* and *Imbalanced Student*. Both of

Algorithm 1 Training of Dual Student for SSL.

Require: Batch \mathcal{B} containing labeled and unlabeled samples

Require: Two independent models $f(\theta)$ and $f(\theta')$

```

1: for each batch  $\mathcal{B}$  do
2:   Get  $\mathcal{B}_1, \mathcal{B}_2$  from  $\mathcal{B}$  by data augmentation
3:   for each model in  $\{f(\theta), f(\theta')\}$  do
4:     Calculate  $\mathcal{L}_{cls}$  on labeled samples
5:     Calculate  $\mathcal{L}_{con}$  by Eq. 1 between  $\mathcal{B}_1$  and  $\mathcal{B}_2$ 
6:   end for
7:   for each unlabeled sample  $x$  do
8:     for each model in  $\{f(\theta), f(\theta')\}$  do
9:       Determine whether  $x$  is stable by Eq. 3
10:    end for
11:    if both  $f(\theta)$  and  $f(\theta')$  are stable for  $x$  then
12:      Calculate the stability of  $x$  by Eq. 5
13:    end if
14:    Calculate  $\mathcal{L}_{sta}$  for  $f(\theta)$  and  $f(\theta')$  by Eq. 6
15:  end for
16:  Update  $f(\theta)$  and  $f(\theta')$  by the loss in Eq. 7
17: end for

```

them have higher performances than the standard Dual Student. They do not increase the inference time, even though more computations are required during training.

Multiple Student: Our Dual Student can be easily extended to Multiple Student. We followed the same strategy as the Deep Co-Training. We assume that our Multiple Student contains $2n$ student models. At each iteration, we randomly divide these students into n pairs. Each pair is then updated like Dual Student. Since our method does not require models to have view differences, the data stream can be shared among the students. This is different from Deep Co-Training, which requires an exclusive data stream for each pair. In practice, four students ($n = 2$) achieve a notable improvement over two students. However, having more than four students do not further improve the performance, as demonstrated in Section 5.2.

Imbalanced Student: Since a well-designed architecture with more parameters usually has better performance, a pre-trained high-performance teacher can be used to improve the light-weight student in knowledge distillation task [9, 10]. Based on the same idea, we extend Dual Student to Imbalanced Student by enhancing the capability of one student. However, we do not consider the sophisticated model as a teacher, since the knowledge will be exchanged mutually. We find that the improvement of the weak student is proportional to the capability of the strong student.

5. Experiments

We first evaluate Dual Student on several common SSL benchmarks, including CIFAR, SVHN, and ImageNet. We then evaluate the performances of the two variants of Dual Student. We further analyze various aspects of our method

Table 1: Test error rate on CIFAR-10 averaged over 5 runs. Parentheses show numbers of training epochs (default 300).

Model	1k labels	2k labels	4k labels	all labels
Π [16]	$31.65 \pm 1.20^\dagger$	$17.57 \pm 0.44^\dagger$	12.36 ± 0.31	5.56 ± 0.10
Π + SN [18]	21.23 ± 1.27	14.65 ± 0.31	11.00 ± 0.13	5.19 ± 0.14
Temp [16]	$23.31 \pm 1.01^\dagger$	$15.64 \pm 0.39^\dagger$	12.16 ± 0.24	5.60 ± 0.10
Temp + SN [18]	18.41 ± 0.52	13.64 ± 0.32	10.93 ± 0.34	5.20 ± 0.14
MT [32]	$18.78 \pm 0.31^\dagger$	$14.43 \pm 0.20^\dagger$	$11.41 \pm 0.27^\dagger$	$5.98 \pm 0.21^\dagger$
MT + FSWA [1]	16.84 ± 0.62	12.24 ± 0.31	9.86 ± 0.27	5.14 ± 0.07
CS	17.38 ± 0.52	13.76 ± 0.27	10.24 ± 0.20	5.18 ± 0.11
DS	15.74 ± 0.45	11.47 ± 0.14	9.65 ± 0.12	5.20 ± 0.03
MT + FSWA (1200) [1]	15.58 ± 0.12	11.02 ± 0.23	9.05 ± 0.21	4.73 ± 0.18
Deep CT (600) [26]	-	-	9.03 ± 0.18	-
DS (600)	14.17 ± 0.38	10.72 ± 0.19	8.89 ± 0.09	4.66 ± 0.07

through ablation experiments. Finally, we demonstrate the application of Dual Student in a domain adaptation task.

Unless specified otherwise, the architecture used in our experiments is a same 13-layer convolutional neural network (CNN), following previous works [16, 20, 32]. Its details are described in Appendix B (Supplementary). As reported in [23], the implementations of recent SSL methods are not exactly same, and the training details (*e.g.*, number of training epochs, optimizer and augmentation) may also be different. For a fair comparison, we implement our method following the previous state-of-the-art [1], which uses the standard Batch Norm [12] instead of the mean-only Batch Norm [11]. The stochastic gradient descent optimizer is adopted with the learning rate adjustment function $\gamma = \gamma_0 * (0.5 + \cos((t - 1) * \pi / N))$, where t is the current training step, N is the total number of steps, and γ_0 is the initial learning rate. These settings provide better baselines for Π Model and Mean Teacher. For other methods, we use the results from the original papers. More training details are provided in Appendix C (Supplementary).

5.1. SSL Benchmarks

We first evaluate Dual Student on the CIFAR benchmark, including CIFAR-10 [14] and CIFAR-100 [15]. CIFAR-10 has 50k training samples and 10k testing samples, from 10 categories. Each sample is a 32×32 RGB image. We extract 1k, 2k, and 4k balanced labels randomly. CIFAR-100 [15] is a more complex dataset including 100 categories. Each category contains only 500 training samples, together with 100 test samples. We extract 10k balanced labels from it randomly. Besides, we also run experiments with full labels on both datasets. We compare our Dual Student (DS) with some recent consistency-based models, including Π Model (Π), Temporal Model (Temp), Mean Teacher (MT), Smooth Neighbor (SN), FastSWA based on Mean Teacher (MT+FSWA), and Deep Co-Training (Deep CT). We also

Table 2: Test error rate on CIFAR-100 averaged over 5 runs.

Model	10k labels	all labels
Temp [16]	38.65 ± 0.51	26.30 ± 0.15
Π [16]	39.19 ± 0.36	26.32 ± 0.04
Π + FSWA [1]	35.14 ± 0.71	22.00 ± 0.21
MT [32]	$35.96 \pm 0.77^\dagger$	$23.37 \pm 0.16^\dagger$
MT + FSWA [1]	34.10 ± 0.31	21.84 ± 0.12
DS	33.08 ± 0.27	21.90 ± 0.14
MT + FSWA (1200) [1]	33.62 ± 0.54	21.52 ± 0.12
Deep CT (600) [26]	34.63 ± 0.14	-
DS (480)	32.77 ± 0.24	21.79 ± 0.11

replace the stabilization constraint in our structure with the consistency constraint (CS) as a baseline.

Table 1 shows the results on CIFAR-10. All models are trained for 300 epochs, except for those specified with parentheses. Results marked with a \dagger are obtained from other works that published better performances than the original ones. We can see that our Dual Student boosts the performance on all semi-supervised settings. The results reveal that as the number of labeled samples decreases, our method can gain more significant improvements. Specifically, Dual Student improves the result with 1k labels to 14.17% with only half of training epochs comparing to FastSWA. Similar results can also be observed in the experiments with 2k and 4k labels. Fig. 6 shows that the accuracy on only the *stable samples* is higher than that on all samples, which proves that the *stable samples* represent the relatively more reliable knowledge of a model. This justifies why our DS with stabilization constraint achieves much better results than the CS. Our result on full labels shows less advantages since the labels play a much more important role in the fully supervised case. Table 2 lists the results on CIFAR-100. Especially, in 10k label experiments, Dual Student records a new state-of-the-art 32.77% with




Figure 6: Test accuracy of each category on the *stable samples* and on all samples of CIFAR-10. The performance gap indicates that the *stable samples* represent relatively more reliable knowledge of a model. The average ratio of the *stable samples* on the test set is about 85% w.r.t. the model.

Table 3: Test error rate on SVHN averaged over 5 runs.

Model	250 labels	500 labels
Supervised [32]	27.77 ± 3.18	16.88 ± 1.30
MT [32]	4.35 ± 0.50	4.18 ± 0.27
DS	4.24 ± 0.10	3.96 ± 0.15

Table 4: Test error rate on ImageNet averaged over 2 runs.

Model	10% labels-top1	10% labels-top5
Supervised	42.15 ± 0.09	19.76 ± 0.11
MT [32]	37.83 ± 0.12	16.65 ± 0.08
DS	36.48 ± 0.05	16.42 ± 0.07

less training epochs than FastSWA and Deep Co-Training.

To evaluate the generalization ability of Dual Student, we also conduct experiments on both SVHN [21] and ImageNet [28]. Street View House Numbers (SVHN) is a dataset containing 73,257 training samples and 26,032 testing samples. Each sample is a 32×32 RGB image with a center close-up of a house number. We only experiment with 250 and 500 labels on SVHN. ImageNet contains more than 10 million RGB images belonging to 1k categories. We extract 10% balanced labels and train a 50-layer ResNeXt model [33]. Tables 3 and 4 show that Dual Student could improve the results on these datasets of various scales.

5.2. Performance of Variants

We evaluate Multiple Student and Imbalanced Student on the CIFAR benchmark. Table 5 compares them with the standard Dual Student, all using the same 13-layer CNN trained for 300 epochs. For Multiple Student (MS), we train both the four students and the eight students. The performance improvement is limited when more than four students are trained simultaneously. For Imbalanced Student

Table 5: Test error rate of two variants of Dual Student (all using the 13-layer CNN) on the CIFAR benchmark averaged over 3 runs. Parentheses of Multiple Student (MS) indicate the numbers of students. Parentheses of Imbalanced Student (IS) indicate the numbers of parameters for the strong student.

Model	CIFAR-10 1k labels	CIFAR-100 10k labels
DS	15.74 ± 0.45	33.08 ± 0.27
MS (4 models)	14.97 ± 0.36	32.89 ± 0.32
MS (8 models)	14.77 ± 0.33	32.83 ± 0.28
IS (3.53M params)	13.43 ± 0.24	32.59 ± 0.27
IS (11.6M params)	12.39 ± 0.26	31.56 ± 0.22




Figure 7: Test accuracy on CIFAR-10 with 1k labels. Left: Combining our method with Mean Teacher can improve its performance. Right: The effectiveness of our stabilization constraint.

(IS), we replace one student by a ResNet [8] with Shake-Shake regularization. We then conduct the experiments on two different model sizes. In particular, a small one with 3.53 million parameters and a large one with 11.65 million parameters. The small ResNet has almost no increase in computational cost, as its number of parameters is similar to that of the 13-layer CNN (3.13 million parameters). Imbalanced Student achieves a significant performance improvement by distilling the knowledge from a more powerful student. Notably, the large ResNet improves the result from 15.74% to 12.39% on CIFAR-10 with 1k labels.

Our structure can also be combined with existing methods easily to further improve the performance. We replace the consistency constraint inside the model by Mean Teacher. Fig. 7 (left) shows the accuracy curves. The obvious performance improvement shows the ability of Dual Student in breaking the limits of the EMA teacher. The accuracy of the combination is similar to that using Dual Student only, which means that our method is insensitive to the type of consistency constraint inside each model.

Table 6: Mean test error rate on the CIFAR benchmark averaged over 5 runs, with different confidence threshold values, ξ . Parentheses show the numbers of the labeled samples.

Dataset (Labels)	$\xi = 0.0$	$\xi = 0.4$	$\xi = 0.6$	$\xi = 0.8$
CIFAR-10 (1k)	16.49	16.12	15.92	15.74
CIFAR-100 (10k)	33.67	33.08	33.23	33.54

5.3. Ablation Experiments

We conduct the ablation experiments on CIFAR-10 with 1k labels to analyze the impact of the confidence threshold and various constraints in our structure.

Confidence threshold: The confidence threshold ξ controls the 2nd condition in Def. 4.1 of the *stable sample* by filtering out samples near to the boundary. Its actual value can be set approximately, since our method is robust to it. Typically, ξ is related to the complexity of the task, e.g., the number of categories to predict or the size of the given dataset. More categories or a smaller size would require a smaller ξ . Table 6 compares different ξ values on the CIFAR benchmark. The results show that ξ is necessary for a better performance, and a meticulous tuning may only help improve the performance slightly.

Effect of the constraints: Dual Student learns the unlabeled data by both \mathcal{L}_{sta} between models and the \mathcal{L}_{con} inside each model. We also study their individual impacts. Besides, we compare the results with the experiment where only the consistency constraint is applied between models (named \mathcal{L}_{cs}). Fig. 7 (right) shows that \mathcal{L}_{cs} reduces the accuracy in the late stage while \mathcal{L}_{sta} helps improve the performance continuously. This demonstrates that our \mathcal{L}_{sta} is better than \mathcal{L}_{cs} . In addition, \mathcal{L}_{con} inside the model also plays a role in boosting the performance further.

5.4. Domain Adaptation

Domain adaptation aims to transfer knowledge learned from a labeled dataset to an unlabeled one. French *et al.* [6] modified Mean Teacher and Temporal Model to enable domain adaptation and showed the effectiveness of the Teacher-Student structure. In this section, we apply Dual Student for adapting the digit recognition model from USPS to MNIST and show that it could be applied to this kind of task with great advantages over the EMA teacher based methods.

Both USPS and MNIST are greyscale hand-written number dataset. USPS consists of 7,000 images of 16×16 , and MNIST contains 60,000 images of 28×28 . To match the image resolution, we resize all images from USPS to 28×28 by cubic spline interpolation. Fig. 8 shows the domain difference between the two datasets. In our experiments, we set USPS as the source domain and MNIST as the target domain. We compare our method with Mean Teacher,



Figure 8: Domain difference between USPS and MNIST. The numbers in USPS are in bold font face and span all over the images without border.




Figure 9: Test curves of domain adaptation from USPS to MNIST versus the number of epochs. Dual Student avoids overfitting and improves the result remarkably.

source domain (USPS) supervised model, and target domain (MNIST) supervised model (trained on 7k balanced labels). All experiments use a small architecture simplified from the above 13-layer CNN. More details are available in Appendix D (Supplementary).

Fig. 9 shows the test accuracy versus the number of epochs. We can see that naively using supervision from USPS would result in overfitting. Mean Teacher avoids it to some extent and improves the top1 accuracy from 69.09% to 80.41%, but it overfits when the number of training epochs is large. Our Dual Student avoids overfitting and boosts the accuracy to 91.50%, which is much closer to the result obtained by supervision from the target domain.

6. Conclusion

In this paper, we have studied the coupling effect of the existing Teacher-Student methods and shown that it sets a performance bottleneck for the structure. We have proposed a new structure, Dual Student, to break limits of the EMA teacher, and a novel stabilization constraint, which provides an effective way to train independent models (either with the same architecture or not). The stabilization constraint is bidirectional overall but is unidirectional for each *stable sample*. The improved performance is notable across datasets and tasks. Besides, we have also discussed two variants of Dual Student, with even better results. However, our method still shares similar limitations as existing methods, e.g., increased memory usage during training and performance degradation on increasing number of labels. In the future, we plan to address these issues and extend our structure to other applications.

References

- [1] Ben Athiwaratkun, Marc Finzi, Pavel Izmailov, and Andrew Gordon Wilson. There are many consistent explanations of unlabeled data: Why you should average. In *Proc ICLR*, 2019. 2, 6
- [2] Philip Bachman, Ouais Alsharif, and Doina Precup. Learning with pseudo-ensembles. In *Proc. NIPS*. 2014. 2
- [3] Avrim Blum and Tom Mitchell. Combining labeled and unlabeled data with co-training. In *Proc. Annual Conference on Computational Learning Theory*, 1998. 2
- [4] Olivier Chapelle, Bernhard Schölkopf, and Alexander Zien. *Semi-Supervised Learning*. The MIT Press, 2006. 3
- [5] Dong-Dong Chen, Wei Wang, Wei Gao, and Zhi-Hua Zhou. Tri-net for semi-supervised deep learning. In *Proc. IJCAI*, 2018. 1
- [6] Geoffrey French, Michal Mackiewicz, and Mark Fisher. Self-ensembling for domain adaptation. In *Proc. ICLR*. 2018. 8
- [7] Ian Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In *Proc. ICLR*, 2015. 1, 2
- [8] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *Proc. CVPR*, 2016. 7
- [9] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. In *NIPS Workshop on Deep Learning and Unsupervised Feature Learning*, 2014. 5
- [10] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Efficient knowledge distillation from an ensemble of teachers. In *Annual Conference of the International Speech Communication Association*, 2017. 5
- [11] Sergey Ioffe. Batch renormalization: Towards reducing minibatch dependence in batch-normalized models. In *Proc. NIPS*. 2017. 6
- [12] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *Proc. ICML*, 2015. 6
- [13] Pavel Izmailov, Dmitrii Podoprikin, Timur Garipov, Dmitry P. Vetrov, and Andrew Gordon Wilson. Averaging weights leads to wider optima and better generalization. In *Proc. UAI*. 2018. 2
- [14] Alex Krizhevsky, Vinod Nair, and Geoffrey Hinton. Cifar-10 (canadian institute for advanced research). 6
- [15] Alex Krizhevsky, Vinod Nair, and Geoffrey Hinton. Cifar-100 (canadian institute for advanced research). 6
- [16] Samuli Laine and Timo Aila. Temporal ensembling for semi-supervised learning. In *Proc. ICLR*. 2017. 1, 2, 6
- [17] Wen Li, Limin Wang, Wei Li, Eirikur Agustsson, and Luc Van Gool. Webvision database: Visual learning and understanding from web data. 2017. 1
- [18] Yucen Luo, Jun Zhu, Mengxi Li, Yong Ren, and Bo Zhang. Smooth neighbors on teacher graphs for semi-supervised learning. In *Proc. CVPR*. 2018. 2, 6
- [19] Lars Maale, Casper Kaae Snderby, Sren Kaae Snderby, and Ole Winther. Auxiliary deep generative models. In *Proc. ICML*. 2016. 1
- [20] Takeru Miyato, Shin-ichi Maeda, Shin Ishii, and Masanori Koyama. Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *IEEE TPAMI*, 2018. 1, 2, 6
- [21] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bischoff, Bo Wu, and Andrew Ng. Reading digits in natural images with unsupervised feature learning. In *NIPS Workshop on Deep Learning and Unsupervised Feature Learning*, 2011. 7
- [22] Augustus Odena. Semi-supervised learning with generative adversarial networks. In *Data Efficient Machine Learning workshop at ICML*. 2016. 1
- [23] Avital Oliver, Augustus Odena, Colin Raffel, Ekin Cubuk, and Ian Goodfellow. Realistic evaluation of semi-supervised learning algorithms. In *Proc. NeurIPS*. 2018. 6
- [24] Sungrae Park, Jun-Keon Park, Su-Jin Shin, and Il-Chul Moon. Adversarial dropout for supervised and semi-supervised learning. In *Proc. AAAI*. 2018. 2
- [25] Boris T. Polyak and Anatoli Juditsky. Acceleration of stochastic approximation by averaging. *SIAM Journal on Control and Optimization*, 1992. 2
- [26] Siyuan Qiao, Wei Shen, Zhishuai Zhang, Bo Wang, and Alan L. Yuille. Deep co-training for semi-supervised image recognition. In *Proc. ECCV*. 2018. 2, 6
- [27] Antti Rasmus, Mathias Berglund, Mikko Honkala, Harri Valpola, and Tapani Raiko. Semi-supervised learning with ladder networks. In *Proc. NIPS*. 2015. 1, 2
- [28] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *IJCV*, 2015. 1, 7
- [29] Jocelyn Sietsma and Robert Dow. Creating artificial neural networks that generalize. *Neural Networks*, 1991. 2
- [30] Jost Tobias Springenberg. Unsupervised and semi-supervised learning with categorical generative adversarial networks. *arXiv:1511.06390*, 2015. 1
- [31] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 2014. 1, 2
- [32] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *Proc. NIPS*. 2017. 1, 2, 3, 6, 7
- [33] Saining Xie, Ross B. Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *Proc. CVPR*. 2017. 7
- [34] Zhi-Hua Zhou. When semi-supervised learning meets ensemble learning. In *Frontiers of Electrical and Electronic Engineering in China*, 2011. 2
- [35] Xiaojin Zhu. Semi-supervised learning literature survey. *TR 1530, University of Wisconsin, Madison*, 2006. 1