

Learning to Reweight with Deep Interactions*

Yang Fan¹, Yingce Xia², Lijun Wu², Shufang Xie², Weiqing Liu², Jiang Bian², Tao Qin²,
Xiang-Yang Li¹

¹University of Science and Technology of China

²Microsoft Research Asia

fyabc@mail.ustc.edu.cn, xiangyangli@ustc.edu.cn

{yingce.xia, lijuwu, shufxi, Weiqing.Liu, Jiang.Bian, taoqin}@microsoft.com

Abstract

Recently, the concept of teaching has been introduced into machine learning, in which a teacher model is used to guide the training of a student model (which will be used in real tasks) through data selection, loss function design, etc. Learning to reweight, which is a specific kind of teaching that reweights training data using a teacher model, receives much attention due to its simplicity and effectiveness. In existing learning to reweight works, the teacher model only utilizes shallow/surface information such as training iteration number and loss/accuracy of the student model from training/validation sets, but ignores the internal states of the student model, which limits the potential of learning to reweight. In this work, we propose an improved data reweighting algorithm, in which the student model provides its internal states to the teacher model, and the teacher model returns adaptive weights of training samples to enhance the training of the student model. The teacher model is jointly trained with the student model using meta gradients propagated from a validation set. Experiments on image classification with clean/noisy labels and neural machine translation empirically demonstrate that our algorithm makes significant improvement over previous methods.

Introduction

Inspired by human education systems, the concept of teaching has been introduced into machine learning, in which a teacher model is employed to teach and assist the training of a student model. Previous work can be categorized into two branches: (1) One is to transfer knowledge (e.g., image classification, machine translation) from the teacher to student (Zhu 2015; Liu and Zhu 2016) and the teacher model aims to teach the student model with minimal cost, like data selection (Liu et al. 2017, 2018) in supervised learning and action modification in reinforcement learning (Zhang et al. 2020). (2) The student model is used for real tasks (e.g., image classification, machine translation) and the teacher is a meta-model that can guide the training of student

model. The teacher model takes the information from the student model and the validation set as inputs and outputs some signals to guide the training of the student model, e.g., adjusting the weights of training data (Fan et al. 2018; Shu et al. 2019; Jiang et al. 2018; Ren et al. 2018), generating better loss functions (Wu et al. 2018), etc. These approaches have shown promising results in image classification (Jiang et al. 2018; Shu et al. 2019), machine translation (Wu et al. 2018), and text classification (Fan et al. 2018). Among the teaching methods, learning to reweight the data is widely adopted due to its simplicity and effectiveness and we focus on this direction in this work.

Previously, the teacher model used for data reweighting only utilizes surface information derived from the student model. In (Fan et al. 2018; Wu et al. 2018; Shu et al. 2019; Jiang et al. 2018), the inputs of the teacher model include training iteration number, training loss (as well as the margin (Schapire et al. 1998)), validation loss, the output of the student model, etc. In those algorithms, the teacher model does not leverage the internal states of the student model, e.g., the values of the hidden neurons of a neural network based student model. We notice that the internal states of a model have been widely investigated and shown its effectiveness in deep learning algorithms and tasks. In ELMo (Peters et al. 2018), a pre-trained LSTM provides its internal states, which are the values of its hidden layers, for downstream tasks as feature representations. In image captioning (Xu et al. 2015; Anderson et al. 2018), a faster RCNN (Ren et al. 2015) pre-trained on ImageNet provides its internal states (i.e., mean-pooled convolutional features) of the selected regions, serving as representations of images (Anderson et al. 2018). In knowledge distillation (Romero et al. 2015; Aguilar et al. 2020), a student model mimics the output of the internal layers of the teacher model so as to achieve comparable performances with the teacher model. However, to the best of our knowledge, this kind of deep information is not extensively investigated in learning to reweight algorithms. The success of leveraging internal states in the above algorithms and applications motivates us to investigate them in learning to reweight, which leads to deep interactions between the teacher and student model.

*This work was done when Yang Fan was an intern at Microsoft Research Asia. Corresponding authors: Yingce Xia and Xiang-Yang Li.

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

We propose a new data reweighting algorithm, in which the teacher model and the student model have deep interactions: the student model provides its internal states (e.g., the values of its internal layers) and optionally surface information (e.g., predictions, classification loss) to the teacher model, and the teacher model outputs adaptive weights of training samples which are used to enhance the training of the student model. A workflow of our method is in Figure 1. We decompose the student model into a feature extractor, which can process the input x to an internal state c (the yellow parts), and a classifier, which is a relatively shallow model (e.g., a linear classifier; blue parts) to map c to the final prediction \hat{y} . In (Fan et al. 2018; Wu et al. 2018), the teacher model only takes the surface information of the student model as inputs like training and validation loss (i.e., the blue parts), which are related to \hat{y} and ground truth label y but not explicitly related to the internal states c . In contrast, the teacher model in our algorithm leverages both the surface information and the internal states c of the student model as inputs. In this way, more information from the student model becomes accessible to the teacher model.

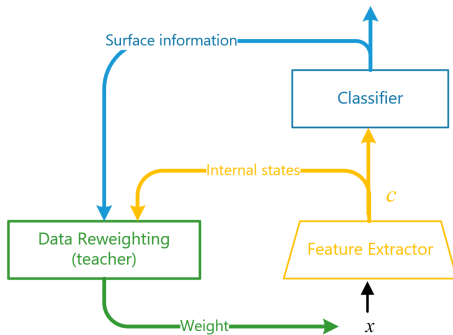


Figure 1: Workflow of our approach.

In our algorithm, the teacher and the student models are jointly optimized in an alternating way, where the teacher model is updated according to the validation loss via reverse-mode differentiation (Maclaurin, Duvenaud, and Adams 2015), and the student model tries to minimize the loss on reweighted data. Experimental results on CIFAR-10 and CIFAR-100 (Krizhevsky, Hinton et al. 2009) with both clean labels and noisy labels demonstrate the effectiveness of our algorithm. We also conduct a group of experiments on IWSLT German→English translation to demonstrate the effectiveness of our method on sequence generation tasks. We achieve promising results over previous methods of learning to teach.

Related Work

Assigning weights to different data points have been widely investigated in literature, where the weights can be either continuous (Friedman et al. 2000; Jiang et al. 2018) or binary (Fan et al. 2018; Bengio et al. 2009). The weights can be explicitly bundled with data, like Boosting and AdaBoost methods (Freung and Shapire 1997; Hastie et al. 2009; Friedman et al. 2000) where the weights of incorrectly classified data are gradually increased, or implicitly achieved by

controlling the sampling probability, like hard negative mining (Malisiewicz, Gupta, and Efros 2011) where the harder examples in a previous round will be sampled again in the next round. As a comparison, in self-paced learning (SPL) (Kumar, Packer, and Koller 2010), weights of hard examples will be assigned to zero in the early stage of training, and the threshold is gradually increased during the training process to control the student model to learn from easy to hard. An important motivation of data weighting is to increase the robustness of training, including addressing the problem of imbalanced data (Sun et al. 2007; Dong, Gong, and Zhu 2017; Khan et al. 2018), biased data (Zadrozny 2004; Ren et al. 2018), noisy data (Angluin and Laird 1988; Reed et al. 2014; Sukhbaatar and Fergus 2014; Koh and Liang 2017). The idea of adjusting weights for the data is also essential in another line of research about combing optimizers with different sampling techniques (Katharopoulos and Fleuret 2018; Liu, Wu, and Mozafari 2020; Namkoong et al. 2017).

Except for manually designing weights for the data, there is another branch of work that leverages a meta model to assign weights. Learning to teach (Fan et al. 2018) is a learning paradigm where there is a student model for the real task, and a teacher model to guide the training of the student model. Based on the collected information, the teacher model provides signals to the student model, which can be the weights of training data (Fan et al. 2018), adaptive loss functions (Wu et al. 2018), etc. The general scheme of machine teaching is discussed and summarized in (Zhu 2015). The concept of teaching can be found in label propagation (Gong et al. 2016a,b), pedagogical teaching (Ho et al. 2016; Shafto, Goodman, and Griffiths 2014), etc. (Liu et al. 2017) leverages a teaching way to speed up the training, where the teacher model selects the training data balancing the trade off between the difficulty and usefulness of the data. (Shu et al. 2019; Ren et al. 2018; Jiang et al. 2018) mainly focuses on the setting that the data is biased or imbalanced.

In machine teaching literature (which focuses on transferring knowledge from teacher models to student models), there are some works that that teacher get more information beyond surface information like loss function. (Liu et al. 2017, 2018) focuses on data selection to speed up the learning of student model. However, their algorithms and analysis are built upon linear models. (Lessard, Zhang, and Zhu 2019) tackles a similar problem, which is about to find the shortest training sequence to drive the student model to the target one. According to our knowledge, there is no extensive study that teacher and student can deeply interact for data reweighting based on deep neural neural networks. We will empirically verify the benefits of our proposals.

Our Method

We focus on data teaching in this work, where the teacher model assigns an adaptive weight to each sample. We first introduce the notations used in this work, then describe our algorithm, and finally provide some discussions.

Notations

Let \mathcal{X} and \mathcal{Y} denote the source domain and the target domain respectively. We want to learn a mapping f , i.e., the student

model, from \mathcal{X} and \mathcal{Y} . W.l.o.g, we can decompose f into a feature extractor and a decision maker, denoted as φ_f and φ_d respectively, where $\varphi_f : \mathcal{X} \mapsto \mathbb{R}^d$, $\varphi_d : \mathbb{R}^d \mapsto \mathcal{Y}$, and d is the dimension of the extracted feature. That is, for any $x \in \mathcal{X}$, $f(x) = \varphi_d(\varphi_f(x))$. We denote the parameters of f as θ . In this section, we take image classification as an example and present our algorithm. Our approach can be easily adopted into other tasks like sequence generation, which is shown at Section 6.

Given a classification network f , we manually define $\varphi_f(\cdot)$ is the output of the second-to-last layer, and φ_d is a linear classifier taking $\varphi_f(x)$ as input. Let $\phi(I, M; \omega)$ denote the teacher model parameterized by ω , where I is the internal states of a student model and M is the surface information like training iteration, training loss, labels of the samples, etc. ϕ can map an input sample $(x, y) \in \mathcal{X} \times \mathcal{Y}$ to a non-negative scalar, representing the weight of the sample. Let $\ell(f(x), y; \theta)$ denote the training loss on sample pair (x, y) , and $R(\theta)$ is a regularization term on θ , independent of the training samples.

Let D_{train} and D_{valid} denote the training and validation sets respectively, both of which are subsets of $\mathcal{X} \times \mathcal{Y}$ with N_T and N_V samples. Denote the validation metric as $m(y, \hat{y})$, where y and \hat{y} are the ground truth label and predicted label respectively. We require that $m(\cdot, \cdot)$ should be a differentiable function w.r.t. the second input. $m(y, \hat{y})$ can be specialized as the expected accuracy (Wu et al. 2018) or the log-likelihood on the validation set. Define $\mathcal{M}(D_{\text{valid}}; \theta)$ as $\frac{1}{N_V} \sum_{(x, y) \in D_{\text{valid}}} m(y, f(x; \theta))$.

Algorithm

The teacher model outputs a weight for any input data. When facing a real-world machine learning problem, we need to fit a student model on the training data, select the best model according to validation performance, and apply it to the test set. Since the test set is not accessible during training and model selection, we need to maximize the validation performance of the student model. This can be formulated as a bi-level optimization problem:

$$\begin{aligned} & \max_{\omega, \theta^*(\omega)} \mathcal{M}(D_{\text{valid}}; \theta^*(\omega)) \quad \text{s.t.} \\ & \theta^*(\omega) = \operatorname{argmin}_{\theta} \frac{1}{N_T} \sum_{i=1}^{N_T} w(x_i, y_i) \ell(f(x_i), y_i; \theta) + \lambda R(\theta), \\ & w(x_i, y_i) = \phi(\varphi_f(x_i), M; \omega), \end{aligned} \quad (1)$$

where λ is a hyperparameter, and $w(x_i)$ represents the weight of data x_i . The task of the student model is to minimize the loss on weighted data, as shown in the second line of Eqn.(1). Without a teacher, all $w(x_i)$'s are fixed as one. In a learning-to-teach framework, the parameters of the teacher model (i.e., ω) and the student model (i.e., θ) are jointly optimized. Eqn.(1) is optimized in an iterative way, where we calculate $\theta^*(\omega)$ based on a given ω , then we update ω based on the obtained $\theta^*(\omega)$. We need to figure out how to obtain θ^* , and how to calculate $\frac{\partial}{\partial \omega} \mathcal{M}(D_{\text{valid}}; \theta^*(\omega))$.

Obtaining $\theta^*(\omega)$: Considering a deep neural network is highly non-convex, generally, we are not able to get the

closed-form solution of the θ^* in Eqn.(1). We choose stochastic gradient descend method (briefly, SGD) with momentum for optimization (Polyak 1964), which is an iterative algorithm. We use a subscript t to denote the t -th step in optimization. D_t is the data of the t -th minibatch, with the k -th sample $(x_{t,k}, y_{t,k})$ in it. For ease of reference, denote w_t as a column vector, where the k -th element is the weight for sample $(x_{t,k}, y_{t,k})$, and $\ell(D_t; \theta_t)$ is another column vector with the k -element $\ell(f(x_{t,k}), y_{t,k}; \theta_t)$, both of which are defined in Eqn.(1). Following the implementation of PyTorch (Paszke et al. 2019), the update rule of momentum SGD is:

$$\begin{aligned} v_{t+1} &= \mu v_t + \frac{\partial}{\partial \theta_t} \left[\frac{1}{|D_t|} w_t^\top \ell(D_t; \theta_t) + \lambda R(\theta_t) \right]; \\ \theta_{t+1} &= \theta_t - \eta_t v_{t+1}, \end{aligned} \quad (2)$$

where $\theta_0 = v_0 = 0$. η_t is the learning rate at the t -th step, and μ is the momentum coefficient. Assume we update the model for K steps. We can eventually obtain θ_K , which serves as the proxy for θ^* . To stabilize the training, we will set $\frac{\partial w_t}{\partial \theta_t} = 0$.

Calculating $\frac{\partial}{\partial \omega} \mathcal{M}(D_{\text{valid}}; \theta_K)$: Motivated by reverse-mode differentiation (Maclaurin, Duvenaud, and Adams 2015), we use a recursive way to calculate gradients. For ease of reference, let $d\theta_t$ and dv_t denote $\frac{\partial}{\partial \theta_t} \mathcal{M}(D_{\text{valid}}; \theta_K)$ and $\frac{\partial}{\partial v_t} \mathcal{M}(D_{\text{valid}}; \theta_K)$ respectively. According to the chain rule to compute derivative, for any $t \in \{0, 1, 2, \dots, K-1\}$, we have

$$d\theta_t = \left(\frac{\partial \theta_{t+1}}{\partial \theta_t} \right)^\top d\theta_{t+1} + \left(\frac{\partial v_{t+1}}{\partial \theta_t} \right)^\top dv_{t+1} \quad (3)$$

$$= d\theta_{t+1} + \frac{\partial^2}{\partial \theta_t^2} \left[\frac{1}{|D_t|} w_t^\top \ell(D_t; \theta_t) + \lambda R(\theta_t) \right] dv_{t+1};$$

$$dv_t = \left(\frac{\partial \theta_t}{\partial v_t} \right)^\top d\theta_t + \left(\frac{\partial v_{t+1}}{\partial v_t} \right)^\top dv_{t+1} \quad (4)$$

$$= -\eta_t d\theta_t + \mu dv_{t+1}; \quad (5)$$

$$\frac{\partial v_{t+1}}{\partial \omega} = \frac{\partial^2}{\partial \theta_t \partial \omega} w_t^\top \ell(D_t; \theta_t); \quad (6)$$

$$\frac{\partial}{\partial \omega} \mathcal{M}(D_{\text{valid}}; \theta_K) = \sum_{t=1}^K \left(\frac{\partial v_t}{\partial \omega} \right)^\top dv_t.$$

According to Eqn.(5), we can design Algorithm 1 to calculate the gradients of the teacher model.

In Algorithm 1, we can see that we need a backpropagation interval B as an input, indicating how many internal θ 's are used to calculate the gradients of the teacher. When $B = K$, all student models on the optimization trajectory will be leveraged. B balances the tradeoff between efficiency and accuracy. To use Algorithm 1, we require $\mu > 0$.

As shown in step 2, we first calculate $\frac{\partial}{\partial \theta_K} \mathcal{M}(D_{\text{valid}}; \theta_K)$, with which we can initialize $d\theta$, dv and $d\omega$. We then recover the θ , v and the gradients at the previous step (see step 4). Based on Eqn.(5), we recover the corresponding $d\theta$, dv and $d\omega$. We repeat step 4 and step 5 until getting the eventual $d\omega$, which is the gradient of the validation metric w.r.t. the parameters of the teacher model. Finally, we can leverage any gradient-based algorithm to update the teacher model. With

Algorithm 1: The gradients of the validation metric w.r.t. the parameters of the teacher.

- 1 *Input:* Teacher model backpropagation interval B ; parameters and momentum of the student model θ_K and v_K ; learning rates $\{\eta_t\}_{t=K-B}^{K-1}$; momentum coefficient $\mu (> 0)$; minibatches of data $\{D_t\}_{t=K-B}^{K-1}$;
 - 2 *Initialization:* $d\theta = \frac{\partial}{\partial \theta_K} \mathcal{M}(D_{\text{valid}}; \theta_K)$; $dv = -\eta_K d\theta_K$; $d\omega \leftarrow 0$; $\theta \leftarrow \theta_K$; $v \leftarrow v_K$;
 - 3 **for** $t \leftarrow K - 1 : -1 : K - B$ **do**
 - 4 $\theta \leftarrow \theta + \eta_t v$; $g \leftarrow \frac{\partial}{\partial \theta} [w_t^\top \ell(D_t; \theta) + \lambda R(\theta)]$;
 $v \leftarrow \frac{v-g}{\mu}$;
 - 5 $d\omega \leftarrow d\omega + \frac{\partial}{\partial \omega} (g^\top dv)$; $d\theta \leftarrow d\theta + \frac{\partial}{\partial \theta} (g^\top dv)$;
 $dv \leftarrow -\eta_t d\theta + \mu dv$;
 - 6 *Return* $d\omega$.
-

the new teacher model, we can iteratively update θ^* and ω until reaching the stopping criteria.

In order to avoid calculating Hessian matrix, which needs to store $O(|\theta|^2)$ parameters, we leverage the property that $\frac{\partial^2 \ell}{\partial \theta^2} v = \frac{\partial}{\partial \theta} g^\top v$, where ℓ is the loss function related to θ , v is a vector with size $|\theta| \times 1$, and $g = \frac{\partial \ell}{\partial \theta}$. With this trick, we only require $O(|\theta|)$ GPU memory.

Discussions: Compared to previous work (Jiang et al. 2018; Shu et al. 2019; Fan et al. 2018; Wu et al. 2018), except for the key differences that we use internal states as features, there are some differences in optimization. In (Fan et al. 2018), the teacher is learned in a reinforcement learning manner, which is relatively hard to optimize. In (Wu et al. 2018), the student model is optimized with vanilla SGD, by which all the intermediate θ_t should be stored. In our algorithm, we use momentum SGD, where we only need to store the final θ_K and v_K , by which we can recover all intermediate parameters. We will study how to effectively apply our derivations to more optimizers and more applications in the future.

Teacher Model

We introduce the default network architecture of the teacher model used in experiments. We use a linear model with sigmoid activation. Given a pair (x, y) , we first use φ_f to extract the output of the second-to-last layer, i.e., $I = \varphi_f(x)$. The surface feature M we choose is the one-hot representation of the label, i.e., $M = y$. Then weight of the data (x, y) is $\varphi(I, M) = \sigma(W_I I + E M + b)$, where $\sigma(\cdot)$ denotes the sigmoid function, W_I , E , and b are the parameters to be learned. E can be regarded as an embedding matrix, which enriches the representations of the labels. One can easily extend the teacher model to a multi-layer feed-forward network by replacing σ with a deeper network.

We need to normalize the weights within a minibatch. When a minibatch D_t comes, after calculating the weight $w_{t,k}$ for the data $(x_{t,k}, y_{t,k}) \in D_t$, it is normalized as $\tilde{w}_{t,k} = w_{t,k} / \sum_{j=1}^{|D_t|} w_{t,j}$. This is to ensure that the sum of weights within a batch D_t is always 1.

Experiments on Image Classification

In this section, we conduct experiments on CIFAR-10 and CIFAR-100 image classification. We first show the overall results and then provide several analysis. Finally, we apply our algorithm to the image classification with noised labels.

Settings

There are 50000 and 10000 images in the training and test sets. CIFAR-10 and CIFAR-100 are a 10-class and a 100-class classification tasks respectively. We split 5000 samples from the training dataset as D_{valid} and the remaining 45000 samples are used as D_{train} . Following (He et al. 2016), we use momentum SGD with learning rate 0.1 and divide the learning rate by 10 at the 80-th and 120-th epoch. The momentum coefficient μ is 0.9. The K and B in Algorithm 1 are set as 20 and 2 respectively. We train the models for 300 epochs to ensure convergence. The minibatch size is 128. We conducted experiments on ResNet-32, ResNet-110 and Wide ResNet-28-10 (WRN-28-10) (Zagoruyko and Komodakis 2016). All the models are trained on a single P40 GPU.

We compare the results with the following baselines: (1) The baseline of data teaching (Fan et al. 2018) and loss function teaching (Wu et al. 2018). They are denoted as L2T-data and L2T-loss respectively. (2) Focal loss (Lin et al. 2017), where each data is weighted by $(1 - p)^\gamma$, p is the probability that the data is correctly classified, and γ is a hyperparameter. We search γ on $\{0.5, 1, 2\}$ suggested by (Lin et al. 2017). (3) Self-paced learning (SPL) (Kumar, Packer, and Koller 2010), where we start from easy samples first and then move to harder examples.

Results

The test error rates of different settings are reported in Table 1. For CIFAR-10, we can see that the baseline results of ResNet-32, ResNet-110 and WRN-28-10 are 7.22, 6.38 and 4.27 respectively. With our method, we can obtain 6.20, 5.65 and 3.72 test error rates, which are the best among all listed algorithms. For CIFAR-100, our approach can improve the baseline by 0.92, 1.67 and 1.11 points. These consistent improvements demonstrate the effectiveness of our method. We have the following observations: (1) L2T-data is proposed to speed up the training. Therefore, we can see that the error rates are almost the same as the baselines. (2) For L2T-loss, on CIFAR-10 and CIFAR-100, it can achieve 0.27 and 0.32 points improvements, which are far behind of our proposed method. This shows the great advantage of our method than the previous learning to teach algorithms. (3) Focal loss sets weights to the data according to the hardness only, which does not leverage internal states neither. There exists non-negligible gap between focal loss and our method. (4) For SPL, the results are similar (even worse) to the baseline. This shows the importance of a learning based scheme for data selection.

Analysis

To further verify how our method works, we conduct several ablation studies. All experiments are conducted on CIFAR-10 with ResNet-32.

CIFAR-10	Baseline	L2T-data	L2T-loss	Focal loss	SPL	Ours
ResNet-32	7.22	7.16	6.95	6.60	11.48	6.20
ResNet-110	6.38	6.10	6.02	6.19	11.06	5.65
WRN-28-10	4.27	4.09	3.97	4.57	4.25	3.72
CIFAR-100	Baseline	L2T-data	L2T-loss	Focal loss	SPL	Ours
ResNet-32	29.57	29.54	29.25	28.85	29.98	28.65
ResNet-110	27.69	27.02	26.61	26.55	27.91	26.02
WRN-28-10	20.49	19.92	19.93	19.86	20.56	19.38

Table 1: Results on CIFAR-10/CIFAR-100. The labels are clean.

Comparison with surface information: The features of the teacher model used in Table 1 are the output of the second-to-last layer of the network (denoted as I_0), and the label embedding (denoted as M_0). Based on (Shu et al. 2019; Ren et al. 2018; Wu et al. 2018; Fan et al. 2018), we define another group of features about surface information. Five components are included: the training iteration (normalized by the total number of iteration), average training loss until the current iteration, best validation accuracy until the current iteration, the predicted label of the current input, and the margin values. These surface features are denoted as M_1 .

For the teacher model, We try different combinations of the internal states and surface features. The settings and results are shown in Table 2.

Setting	Error rate
$I_0 + M_0$	6.20
I_0	6.34
M_0	6.50
M_1	6.54
$M_0 + M_1$	6.50
$I_0 + M_0 + M_1$	6.30

Table 2: Ablation study on the usage of features.

As shown in Table 2, we can see that the results of using surface features only (i.e., the settings without I_0) cannot catch up with those with internal states of the network (i.e., the settings with I_0). This shows the effectiveness of the internal states for learning to teach. We do not observe significant differences among the settings M_0 , M_1 and $M_0 + M_1$. Using I_0 only can result in less improvement than using $I_0 + M_0$. Combining I_0 , M_0 and M_1 also slightly hurts the result. Therefore, we choose $I_0 + M_0$ as the default setting. *Internal states from different levels:* By default, we use the output of second-to-last layer as the features of internal states. We also try several other variants, naming I_1 , I_2 and I_3 , which are the outputs of the last convolutional layer with size 8×8 , 16×16 and 32×32 . A larger subscript represents that the corresponding features are more similar to the raw input. We explore the setting $I_i + M_0$, $i \in \{0, 1, 2, 3\}$. Results are reported in Table 3. We can see that leveraging internal states (i.e., I_i) can achieve lower test error rates than those without such features. Currently, there is not significant difference on

where the internal states are from. Therefore, by default, we recommend to use the states from the second-to-last layer.

Setting ($I_i + M_0$)	0	1	2	3
Error rate	6.20	6.22	6.31	6.21

Table 3: Features from different levels.

Architectures of the teacher models: We explore the teacher networks with different number of hidden layers. Each hidden layer is followed by a ReLU activation (denoted as MLP-#layer). The dimension of the hidden states are the same as the input. Results are in Table 4.

Using a more complex teacher model will not bring improvement to the simplest one as we used in the default setting. Our conjecture is that more complex models are harder to optimize, which can not provide accurate signals for the student models.

Analysis on the weights: We take comparison between the weights output by the teacher model leveraging surface features M_1 only (denoted as \mathcal{T}_0) and those output by our teacher leveraging internal features (denoted as \mathcal{T}_1). The results are shown in Figure 2, where the top row represents the results of \mathcal{T}_0 and the bottom row for \mathcal{T}_1 . In Figure 2(a), (b), (d), (e), the data points in the same category are painted with the same color. The first column shows the correlation between the output data weight (y -axis) and the training loss (x -axis); the second column is used to visualize the internal states through t-SNE (Maaten and Hinton 2008); the third column plots heatmaps regarding output weights of all data points (red means large weight and blue means smaller), in accordance with those in the second column. We have the following observations:

(1) As shown in the first column, \mathcal{T}_0 tries to assign lower weights to the data with higher loss, regardless of the category the image belongs to. In contrast, the weights set by \mathcal{T}_1 heavily rely on the category information. For example, the

Setting	MLP-0	MLP-1	MLP-2
Error rate	6.20	6.48	6.59

Table 4: Teacher with various hidden layers.

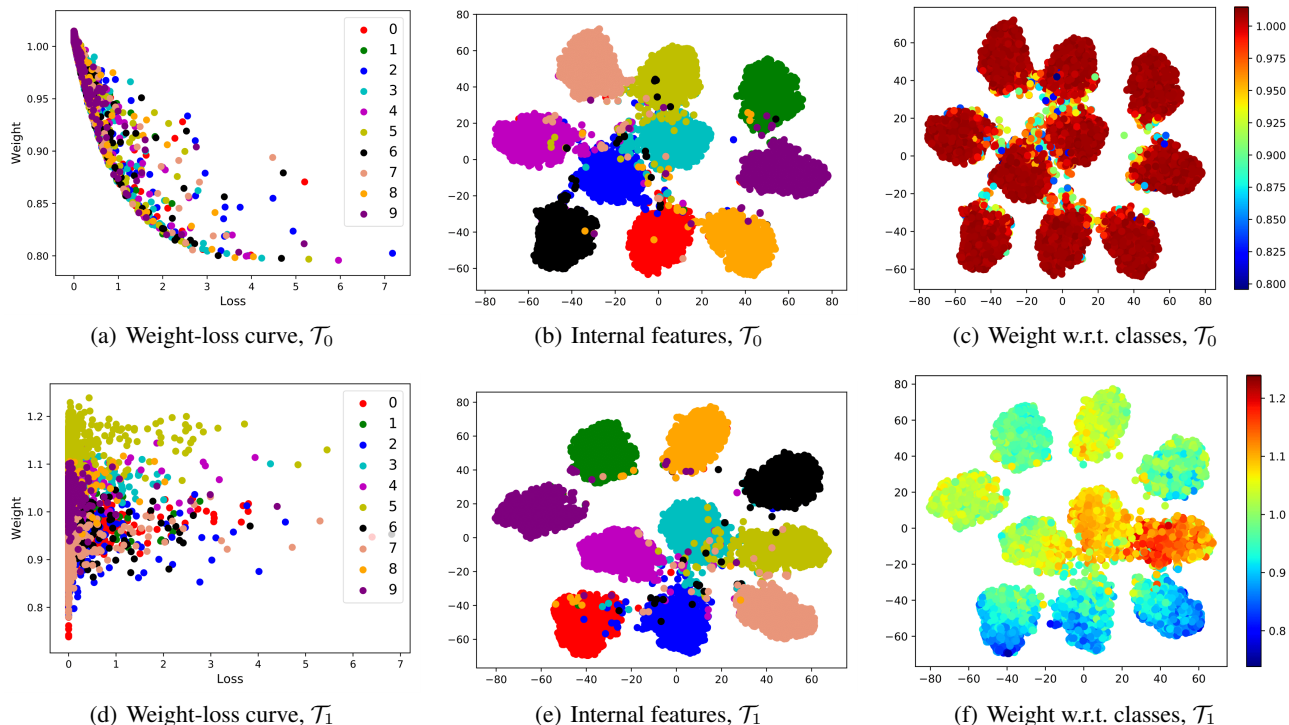


Figure 2: Visualization of weights and loss values of \mathcal{T}_0 and \mathcal{T}_1 .

data points with label 5 have the highest weights regardless of the training loss, followed by those with label 3, where label 3 and 5 correspond to the “cat” and “dog” in CIFAR-10, respectively.

(2) To further investigate the reason why the data of cat class and dog class are assigned with larger weights by \mathcal{T}_1 , we turn to Figure 2(e), from which we can find that the internal states of dog and cat are much overlapped. We therefore hypothesize that, since the dog and cat are somewhat similar to each other, \mathcal{T}_1 is learned to separate these two classes by assigning large weights to them. Yet, this phenomenon cannot be observed in \mathcal{T}_0 .

Preliminary exploration on deeper interactions: To stabilize training, we do not backpropagate the gradients to the student model via the weights, i.e., $\frac{\partial w_t}{\partial \theta_t}$ is set as zero. If we enable $\frac{\partial w_t}{\partial \theta_t}$, the teacher model will have another path to pass the supervision signal to the student model, which has great potential to improve the performances. We quickly verify this variant on CIFAR-10 using ResNet-32. We choose $I_0 + M_0$ as the features of the teacher model. We find that with this technique, we can further lower the test error rate to 6.08%, another 0.12 improvement compared to the current methods. We will further explore this direction in the future.

Image Classification with Noisy Labels

To verify the ability of our proposed method to deal with the noisy data, we conduct several experiments on CIFAR-10/100 datasets with noisy labels.

We derive most of the settings from (Shu et al. 2019). The

images remain the same as those in standard CIFAR-10/100, but we introduce noise to their labels, including the uniform noise and flip noise. For the validation and test sets, both the images and the labels are clean.

1. *Uniform noise:* We follow a common setting from (Zhang et al. 2017). The label of each image is uniformly mapped to a random class with probability p . In our experiments, we set the probability p as 40% and 60%. Following (Shu et al. 2019), the network architecture of the student network is WRN-28-10. We use momentum SGD with learning rate 0.1 and divide the learning rate by 10 at 36-th epoch and 38-th epoch (40 epoch in total).
2. *Flip noise:* We follow (Shu et al. 2019) to set flip noise. The label of each image is independently flipped to two similar classes with probability p . The two similar classes are randomly chosen, and we flip labels to them with equal probability. In our experiments, we set probability p as 20% and 40% and adopt ResNet-32 as the student model. We use momentum SGD with learning rate 0.1 and divide the learning rate by 10 at 40-th epoch and 50-th epoch (60 epoch in total).

For the teacher model, we follow settings as those used for clean data. We compare the results with MentorNet (Jiang et al. 2018) and Meta-Weight-Net (Shu et al. 2019).

The results are shown in Table 5 and Table 6. Our results are better than previous baselines like MentorNet and Meta-Weight-Net, regardless of the type and magnitude. When the noise type is uniform, we can improve Meta-Weight-Net by about 0.5 point. On flip noise with ResNet-32 network, the

improvement is more significant, where in most cases, we can improve the baseline by more than one point.

The experiment results demonstrate that leveraging internal states is also useful for the datasets with noisy labels. This shows the generality of our proposed method.

Method	CIFAR-10	
	$p = 40\%$	$p = 60\%$
Baseline	31.93	46.88
MentorNet (Jiang et al. 2018)	12.67	17.20
Meta-Weight-Net (Shu et al. 2019)	10.73	15.93
Ours	10.29	15.37

Method	CIFAR-100	
	$p = 40\%$	$p = 60\%$
Baseline	48.89	69.08
MentorNet	38.61	63.13
Meta-Weight-Net	32.27	41.25
Ours	31.36	40.62

Table 5: Results of WRN-28-10 with uniform noise labels.

Method	CIFAR-10	
	$p = 20\%$	$p = 40\%$
Baseline	23.17	29.23
MentorNet (Jiang et al. 2018)	13.64	18.24
Meta-Weight-Net (Shu et al. 2019)	9.67	12.46
Ours	8.95	11.29

Method	CIFAR-100	
	$p = 20\%$	$p = 40\%$
Baseline	49.14	56.99
MentorNet	38.03	47.34
Meta-Weight-Net	35.78	41.36
Ours	33.92	39.49

Table 6: Results of ResNet-32 with flip noise labels.

Experiments on Machine Translation

In this section, we verify our algorithm on neural machine translation (NMT). We conduct experiments on IWSLT’14 German-to-English (briefly, De→En) translation, with both cleaned and noisy data.

Settings

There are 153k/7k/7k sentence pairs in the training/valid/test sets of IWSLT’14 De→En translation dataset. We first tokenize the words and then leverage byte-pair-encoding (BPE) (Sennrich, Haddow, and Birch 2016) to split words into sub-word units on the dataset. We use a joint dictionary for the two languages and the vocabulary size is 10k.

To create the noisy IWSLT’14 De→En dataset, we add noise to each sentence pair with probability p . If a sentence is selected to add noise, each word in the source language sentence and the target language sentence is uniformly replaced with a special token (MASK) with probability q . In our experiment, we set p and q as 0.1 and 0.15 respectively.

For all translation tasks, we reuse the settings in image classification to train the teacher model, and use Transformer (Vaswani et al. 2017) as the student model. We derive most settings from the fairseq official implementation¹ and use the `transformer_small` configuration for the student model, where the embedding dimension, hidden dimension of feed-forward layers and number of layers are 512, 1024 and 6 respectively. We first use Adam algorithm with learning 5×10^{-4} to train an NMT model until convergence, which takes about one day. Then we use pre-trained models to initialize the student model and use momentum SGD for finetuning with learning rate 10^{-3} , which takes about three hours. The batchsize is 4096 tokens per GPU. We implement the data teaching (Fan et al. 2018) as a baseline, which is denoted as L2T-data. We use BLEU score (Papineni et al. 2002) as the evaluation metric, which is calculated by `multi-bleu.perl`².

Results

The BLEU score results of neural machine translation tasks are reported in Table 7. We can see that our proposed method can achieve more than 1 points gain on all translation tasks compared with the baseline, and also outperforms the previous approach of L2T-data. For noisy IWSLT’14 De→En task, our approach can improve the baseline by 1.88 points, which indicates that our proposed method is more competitive on noisy datasets.

Task	Baseline	L2T-data	Ours
clean	34.95	35.61	36.00
noisy	33.68	34.42	35.56

Table 7: BLEU scores on IWSLT De→En NMT tasks.

Conclusion and Future Work

We propose a new data teaching paradigm, where the teacher and student model have deep interactions. The internal states are fed into the teacher model to calculate the weights of the data, and we propose an algorithm to jointly optimize the two models. Experiments on CIFAR-10/100 and neural machine translation tasks with clean and noisy labels demonstrate the effectiveness of our approach. Rich ablation studies are conducted in this work.

For future work, the first is to study how to apply deeper interaction to the learning to teach framework (preliminary results in Section 6). Second, we want that the teacher model could be transferred across different tasks, which is lacked for the current teacher (see Appendix B (Fan et al. 2020) for the exploration). Third, we will carry out theoretical analysis on the convergence of the optimization algorithm.

¹<https://github.com/pytorch/fairseq/blob/master/fairseq/models/transformer.py>

²<https://github.com/moses-smt/mosesdecoder/blob/master/scripts/generic/multi-bleu.perl>

References

- Aguilar, G.; Ling, Y.; Zhang, Y.; Yao, B.; Fan, X.; and Guo, E. 2020. Knowledge Distillation from Internal Representations. *AAAI*.
- Anderson, P.; He, X.; Buehler, C.; Teney, D.; Johnson, M.; Gould, S.; and Zhang, L. 2018. Bottom-up and top-down attention for image captioning and visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 6077–6086.
- Angluin, D.; and Laird, P. 1988. Learning from noisy examples. *Machine Learning* 2(4): 343–370.
- Bengio, Y.; Louradour, J.; Collobert, R.; and Weston, J. 2009. Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning*, 41–48.
- Dong, Q.; Gong, S.; and Zhu, X. 2017. Class rectification hard mining for imbalanced deep learning. In *Proceedings of the IEEE International Conference on Computer Vision*, 1851–1860.
- Fan, Y.; Tian, F.; Qin, T.; Li, X.-Y.; and Liu, T.-Y. 2018. Learning to teach. In *Sixth International Conference on Learning Representations*.
- Fan, Y.; Xia, Y.; Wu, L.; Xie, S.; Liu, W.; Bian, J.; Qin, T.; Li, X.-Y.; and Liu, T.-Y. 2020. Learning to teach with deep interactions. *arXiv preprint arXiv:2007.04649*.
- Freung, Y.; and Shapire, R. 1997. A decision-theoretic generalization of on-line learning and an application to boosting. *J Comput Syst Sci* 55: 119–139.
- Friedman, J.; Hastie, T.; Tibshirani, R.; et al. 2000. Additive logistic regression: a statistical view of boosting (with discussion and a rejoinder by the authors). *The annals of statistics* 28(2): 337–407.
- Gong, C.; Tao, D.; Liu, W.; Liu, L.; and Yang, J. 2016a. Label propagation via teaching-to-learn and learning-to-teach. *IEEE transactions on neural networks and learning systems* 28(6): 1452–1465.
- Gong, C.; Tao, D.; Yang, J.; and Liu, W. 2016b. Teaching-to-learn and learning-to-teach for multi-label propagation. In *Thirtieth AAAI conference on artificial intelligence*.
- Hastie, T.; Rosset, S.; Zhu, J.; and Zou, H. 2009. Multi-class adaboost. *Statistics and its Interface* 2(3): 349–360. URL https://www.intlpress.com/site/pub/files/_fulltext/journals/sii/2009/0002/0003/SII-2009-0002-0003-a008.pdf.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep Residual Learning for Image Recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 770–778. URL <https://arxiv.org/pdf/1512.03385.pdf>.
- Ho, M. K.; Littman, M.; MacGlashan, J.; Cushman, F.; and Austerweil, J. L. 2016. Showing versus doing: Teaching by demonstration. In *Advances in neural information processing systems*, 3027–3035.
- Jiang, L.; Zhou, Z.; Leung, T.; Li, L.-J.; and Fei-Fei, L. 2018. Mentornet: Learning data-driven curriculum for very deep neural networks on corrupted labels. In *Thirty-fifth International Conference on Machine Learning*. URL <https://arxiv.org/pdf/1712.05055>.
- Katharopoulos, A.; and Fleuret, F. 2018. Not all samples are created equal: Deep learning with importance sampling. *arXiv preprint arXiv:1803.00942*.
- Khan, S. H.; Hayat, M.; Bennamoun, M.; Soheli, F. A.; and Togneri, R. 2018. Cost-Sensitive Learning of Deep Feature Representations From Imbalanced Data. *IEEE Transactions on Neural Networks and Learning Systems* 29(8): 3573–3587.
- Koh, P. W.; and Liang, P. 2017. Understanding black-box predictions via influence functions. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, 1885–1894. JMLR. org.
- Krizhevsky, A.; Hinton, G.; et al. 2009. Learning Multiple Layers of Features from Tiny Images. *University of Toronto*.
- Kumar, M. P.; Packer, B.; and Koller, D. 2010. Self-Paced Learning for Latent Variable Models. In Lafferty, J. D.; Williams, C. K. I.; Shawe-Taylor, J.; Zemel, R. S.; and Culotta, A., eds., *Advances in Neural Information Processing Systems 23*, 1189–1197. Curran Associates, Inc. URL <http://papers.nips.cc/paper/3923-self-paced-learning-for-latent-variable-models.pdf>.
- Lessard, L.; Zhang, X.; and Zhu, X. 2019. An optimal control approach to sequential machine teaching. In *The 22nd International Conference on Artificial Intelligence and Statistics*, 2495–2503. PMLR.
- Lin, T.-Y.; Goyal, P.; Girshick, R.; He, K.; and Dollár, P. 2017. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, 2980–2988.
- Liu, J.; and Zhu, X. 2016. The Teaching Dimension of Linear Learners. *Journal of Machine Learning Research* 17(162): 1–25. URL <http://jmlr.org/papers/v17/15-630.html>.
- Liu, R.; Wu, T.; and Mozafari, B. 2020. Adam with Bandit Sampling for Deep Learning. *Advances in Neural Information Processing Systems* 33.
- Liu, W.; Dai, B.; Humayun, A.; Tay, C.; Yu, C.; Smith, L. B.; Rehg, J. M.; and Song, L. 2017. Iterative machine teaching. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, 2149–2158. JMLR. org.
- Liu, W.; Dai, B.; Li, X.; Liu, Z.; Rehg, J.; and Song, L. 2018. Towards black-box iterative machine teaching. In *International Conference on Machine Learning*, 3141–3149.
- Maaten, L. v. d.; and Hinton, G. 2008. Visualizing data using t-SNE. *Journal of machine learning research* 9(Nov): 2579–2605.
- Maclaurin, D.; Duvenaud, D.; and Adams, R. 2015. Gradient-based hyperparameter optimization through reversible learning. In *International Conference on Machine Learning*, 2113–2122.
- Malisiewicz, T.; Gupta, A.; and Efros, A. A. 2011. Ensemble of exemplar-svms for object detection and beyond. In *2011 International conference on computer vision*, 89–96. IEEE.
- Namkoong, H.; Sinha, A.; Yadlowsky, S.; and Duchi, J. C. 2017. Adaptive sampling probabilities for non-smooth optimization. In *International Conference on Machine Learning*, 2574–2583.

- Papineni, K.; Roukos, S.; Ward, T.; and Zhu, W.-J. 2002. BLEU: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting on association for computational linguistics*, 311–318. Association for Computational Linguistics.
- Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; Desmaison, A.; Kopf, A.; Yang, E.; DeVito, Z.; Raison, M.; Tejani, A.; Chilamkurthy, S.; Steiner, B.; Fang, L.; Bai, J.; and Chintala, S. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems 32*, 8024–8035. Curran Associates, Inc. URL <http://papers.nips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>.
- Peters, M.; Neumann, M.; Iyyer, M.; Gardner, M.; Clark, C.; Lee, K.; and Zettlemoyer, L. 2018. Deep Contextualized Word Representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, 2227–2237. New Orleans, Louisiana: Association for Computational Linguistics. URL <https://www.aclweb.org/anthology/N18-1202>.
- Polyak, B. T. 1964. Some methods of speeding up the convergence of iteration methods. *USSR Computational Mathematics and Mathematical Physics* 4(5): 1–17.
- Reed, S.; Lee, H.; Anguelov, D.; Szegedy, C.; Erhan, D.; and Rabinovich, A. 2014. Training deep neural networks on noisy labels with bootstrapping. *arXiv preprint arXiv:1412.6596*.
- Ren, M.; Zeng, W.; Yang, B.; and Urtasun, R. 2018. Learning to reweight examples for robust deep learning. In *Thirty-fifth International Conference on Machine Learning*.
- Ren, S.; He, K.; Girshick, R.; and Sun, J. 2015. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. In Cortes, C.; Lawrence, N. D.; Lee, D. D.; Sugiyama, M.; and Garnett, R., eds., *Advances in Neural Information Processing Systems 28*, 91–99. Curran Associates, Inc. URL <http://papers.nips.cc/paper/5638-faster-r-cnn-towards-real-time-object-detection-with-region-proposal-networks.pdf>.
- Romero, A.; Ballas, N.; Kahou, S. E.; Chassang, A.; Gatta, C.; and Bengio, Y. 2015. Fitnets: Hints for thin deep nets. *ICLR* URL <https://arxiv.org/pdf/1412.6550.pdf>.
- Schapiro, R. E.; Freund, Y.; Bartlett, P.; Lee, W. S.; et al. 1998. Boosting the margin: A new explanation for the effectiveness of voting methods. *The annals of statistics* 26(5): 1651–1686.
- Sennrich, R.; Haddow, B.; and Birch, A. 2016. Neural Machine Translation of Rare Words with Subword Units. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, 1715–1725.
- Shafto, P.; Goodman, N. D.; and Griffiths, T. L. 2014. A rational account of pedagogical reasoning: Teaching by, and learning from, examples. *Cognitive psychology* 71: 55–89.
- Shu, J.; Xie, Q.; Yi, L.; Zhao, Q.; Zhou, S.; Xu, Z.; and Meng, D. 2019. Meta-Weight-Net: Learning an Explicit Mapping For Sample Weighting. In *Advances in Neural Information Processing Systems 32*, 1919–1930. Curran Associates, Inc. URL <http://papers.nips.cc/paper/8467-meta-weight-net-learning-an-explicit-mapping-for-sample-weighting.pdf>.
- Sukhbaatar, S.; and Fergus, R. 2014. Learning from noisy labels with deep neural networks. *arXiv preprint arXiv:1406.2080* 2(3): 4.
- Sun, Y.; Kamel, M. S.; Wong, A. K.; and Wang, Y. 2007. Cost-sensitive boosting for classification of imbalanced data. *Pattern Recognition* 40(12): 3358 – 3378. ISSN 0031-3203. doi:<https://doi.org/10.1016/j.patcog.2007.04.009>. URL <http://www.sciencedirect.com/science/article/pii/S0031320307001835>.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*, 5998–6008.
- Wu, L.; Tian, F.; Xia, Y.; Fan, Y.; Qin, T.; Jian-Huang, L.; and Liu, T.-Y. 2018. Learning to Teach with Dynamic Loss Functions. In *Advances in Neural Information Processing Systems 31*, 6466–6477. URL <http://papers.nips.cc/paper/7882-learning-to-teach-with-dynamic-loss-functions.pdf>.
- Xu, K.; Ba, J.; Kiros, R.; Cho, K.; Courville, A.; Salakhudinov, R.; Zemel, R.; and Bengio, Y. 2015. Show, attend and tell: Neural image caption generation with visual attention. In *International conference on machine learning*, 2048–2057.
- Zadrozny, B. 2004. Learning and evaluating classifiers under sample selection bias. In *Proceedings of the twenty-first international conference on Machine learning*, 114.
- Zagoruyko, S.; and Komodakis, N. 2016. Wide Residual Networks. In Richard C. Wilson, E. R. H.; and Smith, W. A. P., eds., *Proceedings of the British Machine Vision Conference (BMVC)*, 87.1–87.12. BMVA Press. ISBN 1-901725-59-6. doi:10.5244/C.30.87. URL <https://dx.doi.org/10.5244/C.30.87>.
- Zhang, C.; Bengio, S.; Hardt, M.; Recht, B.; and Vinyals, O. 2017. Understanding deep learning requires rethinking generalization. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net. URL <https://openreview.net/forum?id=Sy8gdB9xx>.
- Zhang, X.; Bharti, S. K.; Ma, Y.; Singla, A.; and Zhu, X. 2020. The Teaching Dimension of Q-learning. *arXiv preprint arXiv:2006.09324*.
- Zhu, X. 2015. Machine teaching: An inverse problem to machine learning and an approach toward optimal education. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*.