

---

# Self-PU: Self Boosted and Calibrated Positive-Unlabeled Training

---

Xuxi Chen\*<sup>1</sup> Wuyang Chen\*<sup>2</sup> Tianlong Chen<sup>2</sup> Ye Yuan<sup>2</sup> Chen Gong<sup>3</sup> Kewei Chen<sup>4</sup> Zhangyang Wang<sup>2</sup>

## Abstract

Many real-world applications have to tackle the Positive-Unlabeled (PU) learning problem, *i.e.*, learning binary classifiers from a large amount of unlabeled data and a few labeled positive examples. While current state-of-the-art methods employ importance reweighting to design various risk estimators, they ignored the learning capability of the model itself, which could have provided reliable supervision. This motivates us to propose a novel **Self-PU**<sup>1</sup> learning framework, which seamlessly integrates PU learning and self-training. Self-PU highlights three “self”-oriented building blocks: a *self-paced* training algorithm that adaptively discovers and augments confident positive/negative examples as the training proceeds; a *self-calibrated* instance-aware loss; and a *self-distillation* scheme that introduces teacher-students learning as an effective regularization for PU learning. We demonstrate the state-of-the-art performance of Self-PU on common PU learning benchmarks (MNIST and CIFAR-10), which compare favorably against the latest competitors. Moreover, we study a real-world application of PU learning, *i.e.*, classifying brain images of Alzheimer’s Disease. Self-PU obtains significantly improved results on the renowned Alzheimer’s Disease Neuroimaging Initiative (ADNI) database over existing methods.

## 1. Introduction

For standard supervised learning of binary classifiers, both positive and negative classes need to be collected for train-

---

\*Equal contribution The work was done when Xuxi Chen was mentored by Zhangyang Wang. <sup>1</sup>University of Science and Technology of China <sup>2</sup>Texas A&M University <sup>3</sup>Nanjing University of Science and Technology <sup>4</sup>Banner Alzheimer’s Institute. Correspondence to: Zhangyang Wang <atlaswang@tamu.edu>.

*Proceedings of the 37<sup>th</sup> International Conference on Machine Learning*, Online, PMLR 119, 2020. Copyright 2020 by the author(s).

<sup>1</sup>The code is publicly available at: <https://github.com/TAMU-VITA/Self-PU>.

ing purposes. However, this is not always a realistic setting in many applications, where one certain class of data could be difficult to be collected or annotated. For example, in chronic disease diagnosis, while we might safely consider a diagnosed patient to be “positive”, the much larger population of “undiagnosed” individuals are practically mixed with both “positive” (patient) and “negative” (healthy) examples, since people might be undergoing the disease’s incubation period (Armenian & Lilienfeld, 1974) or might just have not seen doctors. Roughly labeling the “undiagnosed” examples all as negative will hence lead to biased classifiers that inevitably underestimate the risk of chronic disease.

Given those practical demands, Positive-Unlabeled (PU) Learning has been increasingly studied in recent years, where a binary classifier is to be learned from a part of positive examples, plus an unlabeled sample pool of mixed and unspecified positive and negative examples. Because of this weak supervision, PU learning is more challenging than standard supervised or semi-supervised classification problems. Early works tried to identify reliable negative examples from the unlabeled data by hand-crafted heuristics or standard semi-supervised learning methods (Liu et al., 2002; Li & Liu, 2003). Recently, importance reweighting methods such as unbiased PU (uPU) (Du Plessis et al., 2014; 2015) and non-negative PU (nnPU) (Kiryo et al., 2017) treat unlabeled data as weighted negative ones.

Despite these successes, self-supervision via auxiliary or surrogate tasks was never considered, which could potentially supply another means of reliable supervision. This motivates us to explore the learning capability of the model itself. Our proposed **Self-PU** learning framework exploits three aspects of such “self-boosts”: (a) we design a self-paced training strategy to progressively select unlabeled examples and update the “trust” set of confident examples; (b) we explore a fine-grained calibration of the functions for unconfident examples in a meta-learning fashion; and (c) we construct a collaborative self-supervision between teacher and student models, and enforce their consistency as a new regularization, against the weak supervision in PU learning. Our main contributions are outlined as follows:

- A novel **self-paced** learning pipeline is first introduced to adaptively mine confident examples from unlabeled data, that will be labeled into trusted positive/negative classes. A hybrid loss is applied to both the augmented

“labeled” examples and remaining unlabeled data for supervision. The procedure is repeated progressively, with more unlabeled examples selected each time.

- A **self-calibration** strategy is leveraged to further explore the fine-grained treatment of loss functions over unconfident examples, in a meta-learning fashion.
- A **self-distillation** scheme is designed via the collaborative training between several teacher networks and student networks, providing a consistency regularization as another fold of self-supervision.
- In addition to standard benchmarks (MNIST, CIFAR-10), a **new real-world testbed** of PU learning, *i.e.*, Alzheimer’s Disease neuroimage classification, is evaluated for the first time. On the Alzheimer’s Disease Neuroimaging Initiative (ADNI) database, Self-PU achieves superior results over existing solutions.

## 2. Related Work

### 2.1. PU Learning

Let  $X \in \mathbb{R}^d$  and  $Y \in \{+1, -1\}$  ( $d \in \mathbb{N}$ ) be the input and output random variables. In PU learning, the training dataset  $D$  is composed of a positive set  $D_P$  and an unlabeled set  $D_U$ , where we have  $D = D_P \cup D_U$ .  $D_P$  contains  $n_p$  positive examples  $x^p$  sampled from  $P(x|Y = +1)$  and  $D_U$  contains  $n_u$  unlabeled examples  $x^u$  sampled from  $P(x)$ . Denote the class prior probability  $\pi_p = P(Y = +1)$  and  $\pi_n = P(Y = -1)$ , where we follow the convention (Kiryo et al., 2017) to assume  $\pi_p$  as known throughout the paper. Let  $g : \mathbb{R}^d \rightarrow \mathbb{R}$  be the binary classifier and  $\theta$  be its parameter, and the  $L : \mathbb{R} \times \{+1, -1\} \rightarrow \mathbb{R}$  be loss function. The risk of classifier  $g$ ,  $\hat{R}_{PU}(g)$  can be approximated by

$$\hat{R}_{PU}(g) = \frac{\pi_p}{n_p} \sum_{i=1}^{n_p} L(g(x_i^p), 1) + \frac{1}{n_u} \sum_{i=1}^{n_u} L(g(x_i^u), -1) - \frac{\pi_p}{n_p} \sum_{i=1}^{n_p} L(g(x_i^p), -1), \quad (1)$$

which has been known as the unbiased risk estimator for uPU (Du Plessis et al., 2014; 2015; Xu et al., 2017; Elkan & Noto, 2008; Xu et al., 2019b). It was later pointed out that the second line in Eq. (1) would become negative due to overfitting complex models (Kiryo et al., 2017).

A non-negative version (nnPU) of Eq. (1) was therefore suggested:

$$\hat{R}_{PU}(g) = \frac{\pi_p}{n_p} \sum_{i=1}^{n_p} L(g(x_i^p), 1) + \max(0, \frac{1}{n_u} \sum_{i=1}^{n_u} L(g(x_i^u), -1) - \frac{\pi_p}{n_p} \sum_{i=1}^{n_p} L(g(x_i^p), -1)) \quad (2)$$

Importance reweighting methods (e.g. uPU, nnPU) achieve the state of the arts, although treating unlabeled data as “weighted” negative examples still brings in unreliable supervision. Generative adversarial networks were introduced by (Hou et al., 2018), where the conditional generator produced both negative and positive examples resembling the unlabeled real data. DAN (Liu et al., 2019) tried to recover the positive and negative distributions from the unlabeled data without requiring the class prior.

### 2.2. Self-Paced Learning

Self-paced learning (Kumar et al., 2010) was presented as a special case of curriculum learning (Bengio et al., 2009), where the feed of training examples was dynamically generated by the model based on its learning history, aiming to simulate the learning principle of starting by learning easy instances and then gradually taking more challenging cases (Khan et al., 2011). Early PU works designed heuristics for sample selection. In (Xu et al., 2019a), positive instead of negative examples are permanently selected by analyzing the distribution of sample loss. Unlike previous PU learning works which rely on crafted sample selection heuristics, we are the first to leverage the data-driven self-paced learning to progressively turn unlabeled data into labeled ones.

### 2.3. Self-Supervised Learning

In many supervision-starved fields, it is generally difficult to obtain accurate annotations, despite the vast number of unlabeled data available. Self-supervised learning aims to form “pseudo” supervision for learning informative/discriminative features from the data, where models are required to predict on “proxy” tasks formed to be relevant to the target goal. It is known to benefit data-efficient learning (Trinh et al., 2019; Jing & Tian, 2020), adversarial robustness (Chen et al., 2020), and outlier detection (Mohseni et al., 2020).

For example, (Laine & Aila, 2016) augmented each unlabeled example with random noises, and forced consistency between the two predictions. In (Tarvainen & Valpola, 2017), two identical models were used during training: the student learned as usual while the teacher model generated labels and updated its weights through a moving average of the student, forcing consistency between two models. (Zhang et al., 2018) further suggested that, instead of exchanging examples, mutual feature distillation between peer networks can form another strong source of supervision, and can enable the collaborative learning of an ensemble of students. To our best knowledge, we are the first to consider such self-supervision to improve PU learning.

### 3. The Self-PU Framework

Our proposed Self-PU framework exploits the learning capability of the model itself (Figure 1). We first design a self-paced learning pipeline to progressively select and label confident examples from unlabeled data for supervised learning. On top of that, we calibrate the loss functions over the unconfident examples via meta-learning. Moreover, a consistency loss is introduced between peer networks with different learning paces, which collaboratively teach each other. We further extend our consistency from peer networks to their moving-averages (Tarvainen & Valpola, 2017; Laine & Aila, 2016), as another form of supervision.

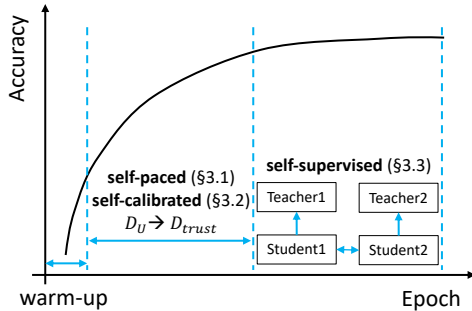


Figure 1. Illustration of the proposed “Self-PU” framework. After a short warm-up period, the classifier is first trained with self-paced learning, where confident examples in  $D_U$  are progressively selected and labeled (positive/negative) into a trusted  $D_{trust}$  subset for supervised learning, with the loss functions over unconfident examples carefully calibrated. After collecting enough confident examples, we start the self-supervised learning via distillation between two collaborative students and their teacher networks.

#### 3.1. Self-Paced PU Learning

Despite the success of unbiased PU risk estimators, they still rely on the estimated class prior and reduced weights on unlabeled data. As shown in (Arpit et al., 2017), during gradient descent, deep neural networks tend not to memorize all training data at the same time but tend to memorize frequent or easy patterns first and later irregular patterns. If we could first find out easy examples and label them with confidence, and then augmenting this labeled pool for the training progress, then we can enjoy “progressively increased” confident full supervision along with the training, in addition to the weak supervision from the PU risk estimators.

Given the model  $g$ , an input example  $x$  and the corresponding label  $y$ , we may compute the output  $g(x)$  and then calculate the probability of  $x$  being positive as  $p(x) = P(Y = +1|x) = f(g(x))$ , where  $f$  is a monotonic function of mapping  $\mathbb{R} \rightarrow [0, 1]$  (e.g. sigmoid function). A greater  $p(x)$  suggests higher confidence that  $x$  belongs to positive class as predicted by  $g$ , and vice versa. By descending sort of  $p(x)$  each time, we can select  $n$  most confident positive and

$n$  most confident negative examples from the current unlabeled data pool  $D_U$ . They will be removed from  $D_U$  and added to our trusted subset  $D_{trust}$ , considered as labeled training examples hereinafter.

Let  $L_{CE}(x, y)$  be the cross entropy loss:

$$L_{CE}(x, y) = \log f(g(x))\mathbb{I}_{y=1} + \log(1 - f(g(x))\mathbb{I}_{y=-1}$$

,  $L_{nnPU}(x, y)$  be the nnPU risk with Sigmoid loss, and together with the given positive subset  $D_P$ , our hybrid loss for self-paced learning becomes:

$$L_{SP} = \sum_{(x,y) \in D_{trust}} L_{CE}(x, y) + \sum_{x \in D - D_{trust}} L_{nnPU}(x) \quad (3)$$

Note that previous works select either only confident positive examples (Xu et al., 2019a) or negative examples (Li & Liu, 2003), while our self-paced learning selects both. Since the cross entropy is used as our supervised loss, one advantage is that the trusted sets of positive/negative samples are balanced in size at each sampling step, avoiding the potential pitfall of extreme class imbalance caused by incrementally sampling only one class.

Besides, previous sample selection (Xu et al., 2019a) often sticks to a pre-fixed learning schedule. In contrast, we unleash more flexibility for the model to automatically and adaptively adjust its own learning pace, via the following techniques. Later on, we will experimentally verify their effectiveness via a step-by-step ablation study.

##### 3.1.1. DYNAMIC RATE SAMPLING

As the learning progresses, training examples with easy/frequent patterns and those with harder/irregular patterns are memorized in different training stages (Arpit et al., 2017). It is important to make our self-paced learning compatible with the memorizing process of the model. A small number of easy examples should be selected first, and then intermediate to hard examples can be labeled after the model is well-trained. Instead of fixing the number of selected confident examples, we propose to dynamically choose the number of confident examples during the self-paced learning. Specifically, as the self-paced learning proceeds, we linearly increase the size of  $D_{trust}$  from 0 to  $r|D_U|$ , where the sampling ratio  $r$  could range from 10% to 40% in our experiments (see section 4.3.2). Empirically, we first “warm-up” the model by training 10 epochs before starting the self-paced learning, in order to keep the selected confident examples as accurate as possible.

##### 3.1.2. “IN-AND-OUT” TRUSTED SET

In previous sample selection approaches, once selected, a trusted example will never be deprived of its label during the subsequent training. In contrast, we allow our training to “regret” on earlier selections in  $D_{trust}$ . Especially at the

early training stage, the intermediate model might not be well-trained enough and not always reliable for predicting labels, which could mislead training if continuing acting as supervision. To this end, we adaptively update  $D_{\text{trust}}$  by also re-examining its current examples each time when we augment new confident ones. The previously selected examples will be removed from  $D_{\text{trust}}$  if their predictions by the current model become of low-confidence, and will be treated as unlabeled again.

### 3.1.3. SOFT LABELS

Instead of giving the selected confident examples hard labels, we directly use the prediction  $f(g(x))$  as soft labels:  $[1 - f(g(x)), f(g(x))]$  as  $[P(Y = -1|x), P(Y = +1|x)]$ , as the practice of label smoothing (Szegedy et al., 2016) appears to benefit learning robustness against label noise.

## 3.2. Self-Calibrated Loss Reweighting

Only leveraging nnPU risk on  $D_U - D_{\text{trust}}$  may not be optimal, as some examples in this set can still provide meaningful supervision. To exploit more supervision from this noisy sets, we introduce a *learning-to-reweight* paradigm (Ren et al., 2018) to the PU learning field for the first time. Letting  $L_{\text{CES}}(x) = f(g(x)) \log f(g(x)) + (1 - f(g(x))) \log(1 - f(g(x)))$  be the cross entropy loss function with soft labels in Sec. 3.1.3, we adaptively combine  $L_{\text{CES}}$  and  $L_{\text{nnPU}}$  for each example  $x_i$  in a batch from  $D_U - D_{\text{trust}}$ , namely:

$$l(x_i, \mathbf{w}_i) = w_{i,1} L_{\text{CES}}(x_i) + w_{i,2} L_{\text{nnPU}}(x_i)$$

Let  $n$  be the mini-batch size. To learn the optimal  $\mathbf{w} = [\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_n]^T$  together in training, we update the model  $g$  for a single gradient descent step on  $l$  with  $\mathbf{w}_i$  very small (*i.e.* a perturbation) of a mini-batch of training examples w.r.t. parameters of models  $\theta$ , followed by a gradient descent step on the cross entropy loss of a mini-batch of validation examples w.r.t  $\mathbf{w}$ , and then rectify the output to be non-negative. The procedure is described as follows:

$$\theta^* = \theta - \delta \nabla \sum_{i=1}^n l(x_i, y_i, \mathbf{w}_i) \quad (4)$$

$$\mathbf{u}_i = -\frac{\partial}{\partial \mathbf{w}_i} \frac{1}{m} \sum_{j=1}^m L_{\text{CE}}^*(x_j^v, y_j^v) |_{\mathbf{w}_i=0} \quad (5)$$

$$\tilde{\mathbf{w}}_i = \max(\mathbf{u}_i, \mathbf{0}), w_{i,1} = \frac{\tilde{w}_{i,1}}{\sum_i \tilde{w}_{i,1}}, w_{i,2} = \frac{\tilde{w}_{i,2}}{\sum_i \tilde{w}_{i,2}} \quad (6)$$

where  $\delta$  denotes the step size,  $m$  the mini-batch size on the validation set which contains clean positive and negative examples, and  $(x_j^v, y_j^v)$ ,  $j = 1, 2, \dots, m$  an example from the validation set with the ground-truth label.  $L_{\text{CE}}^*(x, y)$  calculates loss using the updated parameters  $\theta^*$ .

Meanwhile, on  $D_U - D_{\text{trust}}$ , weighting the cross-entropy

loss too much might not be beneficial to the classifier, especially when the soft labels are not accurate enough.

Therefore, we restrict the total weights of the cross-entropy loss via a balancing factor  $\gamma$ :

$$T = \sup \{k : \sum_{i=1}^k w_{2,i} < \gamma n\} \quad (7)$$

$$w_{i,1}^* = w_{i,1} \mathbf{1}_{\{i < T\}}, w_{i,2}^* = w_{i,2} \mathbf{1}_{\{i < T\}} + \mathbf{1}_{\{i \geq T\}} \quad (8)$$

The corresponding hybrid loss becomes:

$$L_{\text{SP+Reweight}} = \sum_{(x,y) \in D_{\text{trust}}} L_{\text{CE}}(x, y) + \quad (9)$$

$$\sum_{\mathbf{x} \in (D_U - D_{\text{trust}})} l^*(\mathbf{x}) + \sum_{x \in D_P} L_{\text{nnPU}}(x), \quad (10)$$

where  $l^*(\mathbf{x}) = \frac{\sum_{i=1}^n l(x_i, \mathbf{w}_i^*)}{n}$ .

## 3.3. Self-Supervised Consistency via Distillation

To explore additional sources of supervision, we encourage *two forms of self-supervised consistency*: among different learning paces of the model, and along the model’s own moving averaged trajectory. The two goals are altogether achieved by an innovative distillation scheme, with a pair of collaborative student models and their teacher models.

### 3.3.1. CONSISTENCY FOR DIFFERENT LEARNING PACES: MAKING A PAIR OF STUDENTS

The consistency between two self-paced models trained with different paces (*i.e.* sampling ratio in self-paced learning) makes the trained model more resilient to perturbations caused by training stochasticity. To form this self-supervision, we simultaneously train two networks that share the identical architecture, with the same  $D_P$  and  $D_U$  to start on. However, they are set with different confidence thresholds and select different amounts from  $D_U$  to  $D_{\text{trust}}$  each time, making their learning paces un-synchronized and resulting in two different “trusted sets”  $D_{\text{trust}1}$  and  $D_{\text{trust}2}$ . Since two students’ estimations of the probabilities of classes may differ, we force the consistency between two students as a source of distillation, via a mean-square-error (MSE) loss on two models’ predictions.

Denote two networks  $g_1$  and  $g_2$ , the MSE from  $g_1$  to  $g_2$  is defined over  $D - D_{\text{trust}1}$ :

$$L_{\text{MSE}}(g_1, g_2, x) = \|f(g_1(x)) - f(g_2(x))\|^2, x \in D - D_{\text{trust}1} \quad (11)$$

The MSE from  $g_2$  to  $g_1$  is defined over  $D - D_{\text{trust}2}$ :

$$L_{\text{MSE}}(g_2, g_1, x) = \|f(g_2(x)) - f(g_1(x))\|^2, x \in D - D_{\text{trust}2} \quad (12)$$

A *hard sample mining* strategy is further adopted on top, where we only calculate MSE over those “challenging” unlabeled examples whose nnPU risks (2) are large.



Later on, the pair of networks will become two student models for distillation. We therefore have our *first part* of self-supervised consistency loss as:

$$L_{\text{students}} = \sum_{x \in D - D_{\text{trust1}}} l_{\text{stu}}(g_1, g_2, x) + \sum_{x \in D - D_{\text{trust2}}} l_{\text{stu}}(g_2, g_1, x), \quad (13)$$

where

$$l_{\text{stu}}(g_1, g_2, x) = \begin{cases} L_{\text{MSE}}(g_1, g_2, x), & L_{\text{nnPU}}(x) > \alpha L_{\text{MSE}}(g_1, g_2, x) \\ 0, & L_{\text{nnPU}}(x) \leq \alpha L_{\text{MSE}}(g_1, g_2, x) \end{cases} \quad (14)$$

We study the effect of choosing  $\alpha$  in section 4.3.3. The mean squared error between the two students is only calculated on  $D - D_{\text{trust1}}$  and  $D - D_{\text{trust2}}$ . One reason why we choose such design is that the accuracy on  $D_{\text{trust1}}$  and  $D_{\text{trust2}}$  discovered by the self-paced learning is much higher than the accuracy on  $D$ . In addition, on the  $D_{\text{trust1}}$  and  $D_{\text{trust2}}$  the prediction entropy is 0.005, while on the unlabeled set it is 0.074, which indicates much lower confidence.

### 3.3.2. CONSISTENCY FOR MOVING AVERAGED WEIGHTS: ADDING TEACHERS TO DISTILL

Inspired by (Tarvainen & Valpola, 2017), in addition to the consistency between two students, we also encourage them to be consistent with their moving averaged trajectory of weights. Assume that  $g_1$  and  $g_2$  are parameterized by  $\theta_1$  and  $\theta_2$ . For each student we introduce a new teacher model,  $G_1$  and  $G_2$ , parameterized by  $\Theta_1$  and  $\Theta_2$  with the same structure as  $g_1$  and  $g_2$ . The weights of  $G_1, G_2$  are updated via the following moving average:

$$\begin{aligned} \Theta_{1,t} &= \beta \theta_{1,t-1} + (1 - \beta) \theta_{1,t} \\ \Theta_{2,t} &= \beta \theta_{2,t-1} + (1 - \beta) \theta_{2,t} \end{aligned} \quad (15)$$

where  $\theta_{1,t}$  denotes the instance of  $\theta_1$  at time  $t$ , and similarly for others. We study the effect of  $\beta$  in section 4.3.4.

An MSE loss is next enforced for  $G_1$  and  $G_2$  to distill from  $g_1$  and  $g_2$ , namely:

$$L_{\text{teachers}} = \sum_{x \in D} \|f(G_1(x)) - f(g_1(x))\|^2 + \sum_{x \in D} \|f(G_2(x)) - f(g_2(x))\|^2 \quad (16)$$

The above constitutes the *second part* of our self-supervised consistency cost.

In summary, the benefits of self-supervised learning for PU learning come from two folds: 1) the enlarged labeled examples ( $D_{\text{trust}}$ ) introduces stronger supervision into PU learning and brings high accuracy; 2) the consistency cost

between diverse student and teacher models introduces the learning stability (low variance). Eventually, our overall loss function<sup>2</sup> of Self-PU is:

$$L = L_{\text{SP+Reweight}} + L_{\text{students}} + L_{\text{teachers}}. \quad (17)$$

In all experiments and as shown in Figure 1, we first apply self-paced learning and self-calibrated loss reweighting from the 10th epoch to the 50th epoch, followed by a self-distillation period from 50th to 200th epoch. That allows for the models to learn sufficient meaningful information before being distilled. After training, we compare the validation accuracy of two teacher models and select the better performer to be applied on the testing set<sup>3</sup>.

## 4. Experiments

### 4.1. Datasets

In order to evaluate our proposed ‘‘Self-PU’’ learning framework, we conducted experiments on two common testbeds for PU learning: MNIST, and CIFAR-10; plus **a new real-world benchmark**, *i.e.* ADNI (Jack Jr et al., 2008), for the application of Alzheimer’s Disease diagnosis.

#### 4.1.1. INTRODUCTION TO THE ADNI DATABASE

The Alzheimer’s Disease Neuroimaging Initiative (ADNI) database<sup>4</sup> was constructed to test whether brain scans, *e.g.* magnetic resonance imaging (MRI) and other biological markers, can be utilized to predict the early-stage Alzheimer’s disease (AD), in order for more timely prevention and treatment. The dataset, especially its MRI image collection, has been widely adopted and studied for the classification of Alzheimer’s disease (Khvostikov et al., 2018; Li et al., 2015). Fig. 2 shows visual examples.

Traditionally, the machine learning community considers the AD diagnosis task as a binary, fully supervised classification task, between the *patient* and the *healthy* classes. It has **never** been connected to PU learning. Yet, we advocate that this task could become a new **suitable, realistic and challenging** application benchmark for PU learning.

The early-stage AD prediction/diagnosis is highly nontrivial for multi-fold, field-specific reasons. **First**, many nuance factors can heavily affect the feature effectiveness, ranging from individual patient variability to (mechanical/optical) equipment functional fluctuations, to manual operation and sensor/environment noise. **Second**, within the whole-brain scans, only some (not fully-specified) local

<sup>2</sup>Since here we have two students of different learning paces, our  $L_{\text{SP+Reweight}}$  is also extended to both  $D_{\text{trust1}}$  and  $D_{\text{trust2}}$ .

<sup>3</sup>Note that, here we only select one and discard the other, only for simplicity purpose. Other approaches, such as average or weighted-fusion of the two teachers models, are applicable too.

<sup>4</sup><http://adni.loni.usc.edu>

Table 1. Specification of benchmark datasets and models.

Dataset	#Train	#Test	Input Size	$\pi_p$	Positive/Negative	Model
MNIST	60,000	10,000	28×28	0.49	Odd/Even	6-layer MLP
CIFAR-10	50,000	10,000	3×32×32	0.40	Traffic tools/Animals	13-layer CNN
ADNI	822	113	104×128×112	0.43	AD Positive/Negative	3-branch 2-layer CNN




Figure 2. Cross-sectional imaging of a  $104 \times 128 \times 112$  MRI example from the ADNI dataset. Images are from the  $52^{nd}$ ,  $64^{th}$ ,  $56^{th}$  slice of sagittal, coronal, and axial plane, respectively. An MRI image is of gray scale with a value from 0 to 255 for each voxel, and was processed by the intensity inhomogeneity correction, skull-stripping and cerebellum removing.

brain regions are found to be indicative of AD symptoms. Third and most importantly, in contrast to the diagnosed patients, the remaining population, who are not yet clinically diagnosed with AD, *cannot* be simply treated as all *healthy*: on one hand, the above challenges of AD diagnosis inevitably lead to incorrectly missed patient cases; on the other hand, and **more notably**, the AD patients go through a stage called *mild cognitive impairment* (MCI) (Larson et al., 2004; Duyckaerts & Hauw, 1997), a critical transition period between the expected cognitive decline of normal aging, and the severe decline of true dementia. During the MCI stage, those people were “clinically” considered as AD patients already (if diagnosed with more intrusive biochemical means); however, no symptom is known to be observable in current MRI images or other bio-markers.

In other words, the MCI examples have definitely been included in the currently “healthy”-labeled samples in ADNI, while they should have belonged to the “patient” class. In training, we label patients as **positive** class, the “healthy”-labeled examples can then be considered as **unlabeled** class, which mixes the true healthy people (*i.e.*, from the actual **negative** class) and the MCI people (*i.e.*, from the **positive** class). We communicated with several seasoned medical doctors practicing in AD fields, and **they unanimously agreed** that AD diagnosis should be described as a PU learning problem rather than (the traditional treatment as) a binary classification problem. In this paper, we study the specific setting of MRI image classification task on the ADNI dataset, while other bio-marker classification can be

similarly studied in PU settings.

#### 4.1.2. DATASET SETTING

We report our dataset protocols towards PU learning. More metadata are summarized in Table 1.

- **MNIST**: odd numbers 1, 3, 5, 7, 9 form the positive class while even numbers 0, 2, 4, 6, 8 form the negative.
- **CIFAR-10**: four vehicles classes (“airplane”, “automobile”, “ship”, “truck”) constitute the positive class, and six animal classes (“bird”, “cat”, “deer”, “dog”, “frog”, “horse”) constitute the negative.
- **ADNI**: We utilized the public ADNI data set as (Li et al., 2015; Yuan et al., 2018) suggested: The T1-weighted MRI images were processed by first correcting the intensity inhomogeneity, followed by skull-stripping and cerebellum removing. We consider the subjects as positive class if they: 1) have positive clinical diagnosis records on file; or 2) have their standardized uptake value ratio (SUVR) values<sup>5</sup> no less than 1.08 (Villeneuve et al., 2015; Ott et al., 2017; Yuan et al., 2018). While this estimate can be treated as “golden rule” in clinical practice and is shown to work well in our experiments (Table 8), it can be further adjusted flexibly and used in our framework with ease.

Following the convention of nnPU (Kiryo et al., 2017), we use  $n_p = |D_P| = 1000$  in MNIST and CIFAR-10. In ADNI, we end up with  $n_p = 113$ .  $n_u = |D_U|$  equals the size of remaining training data on all three datasets.  $\pi_p$  is the proportion of true positive examples in the dataset.

## 4.2. Baselines and Implementations

Following nnPU (Kiryo et al., 2017), we used a 6-layer *multilayer perceptron* (MLP) with ReLU on MNIST. On CIFAR-10, we use a 13-layer CNN with ReLU. We design a multi-scale network for ADNI, which is used as the backbone for all compared baselines: please see supplementary materials for details. We use Adam optimizer with a cosine annealing learning rate scheduler for training. The batch size is 256 for MNIST and CIFAR-10, and 64 for ADNI.

<sup>5</sup>SUVR is a therapy monitoring or response, considered as an important indicator of Alzheimer’s Disease.

The  $\gamma$  is set to  $\frac{1}{16}$ . The batch size of validation examples equals to the batch size of the training examples.

For a fair comparison, each experiment runs five times, and the mean and standard deviations of accuracy are reported.

### 4.3. Ablation Study

In this section, we carry out a thorough ablation study, on the key components introduced during the self-paced learning stage (*i.e.*, the selection of trusted set) and the self-supervised distillation stage (*i.e.*, the diversity of students, the effect of hard sample mining when training students, and the effect of weight averaging further by teachers). All experiments are conducted on the **CIFAR-10** dataset<sup>6</sup>. We will study the effect of  $\gamma$  in the supplementary materials.

#### 4.3.1. SELECTION OF “TRUSTED SET” $D_{\text{TRUST}}$

Since self-paced learning aims to mine more confident positive/negative examples, it is crucial to ensure the trustworthiness of the selected  $D_{\text{trust}}$ . We therefore calculate the accuracy of assigned labels for  $D_{\text{trust}}$  along the self-paced training, as an indicator of the sampling strategy reliability.

We compare three settings: 1) “*Fixed sampling size*”: each time, the model selects a fixed number of samples (*e.g.* 25% of examples in  $D_U$ ), assigning soft labels and adding them into  $D_{\text{trust}}$ . Meanwhile, low-confidence samples in  $D_{\text{trust}}$  will also be removed in next round of selection. 2) “*Sampling without replacement*”: each example selected by model will permanently reside in  $D_{\text{trust}}$ . Here the size of  $D_{\text{trust}}$  is linearly increased along the training progress. 3) Our *default approach* in Self-PU: both “Dynamic Rate Sampling” and “In-and-Out Trusted Set” are enabled. All three settings end up with  $|D_{\text{trust}}| = 0.25|D_U|$ .

From Figure 3, we clearly see that sampling either with a fixed size or without replacement results in a less reliable selection of  $D_{\text{trust}}$ , compared to our strategy. Moreover, the inaccurately selected examples in  $D_{\text{trust}}$  will further cause much unstable training (dash line). We demonstrate that both “Dynamic Rate Sampling” and “In-and-Out Trusted Set” are vital to achieving an accurate and stable self-paced learning (solid line). Table 2 shows the final test accuracy of three settings, where our proposed self-paced learning pipeline ( $L_{\text{SP}}$ ) significantly outperforms the other two settings ( $L_{\text{SP}}$  of fixed sampling size and sampling without replacement). The better accuracy and lower variance show the advantage of our strategy.

<sup>6</sup>To conduct controlled experiments we disable the self-calibration strategy in Table 3, 4, 5

Table 2. Classification comparison on CIFAR-10: we report both **means** and **standard deviations** (in parentheses) from five runs.  $L_{\text{SP}}$ : self-paced training.  $L_{\text{SPS}}$ : self-paced training with soft label in Sec. 3.1.3.  $L_{\text{SP}} + \text{Reweighting}$ : self-paced training with self-calibrated loss reweighting in section 3.2.  $L_{\text{SPS}} + \text{Reweighting}$ : self-paced training with soft label and self-calibrated loss reweighting in section 3.2.  $L_{\text{students}}$ : self-distillation from a pair of students.  $L_{\text{teacher}}$ : self-distillation from teacher networks. Self-PU:  $L_{\text{SPS}} + \text{Reweighting} + L_{\text{students}} + L_{\text{teacher}}$

Method	CIFAR-10 %
nnPU (baseline)	88.60 (0.40)
$L_{\text{SP}}$ (fixed size)	88.05 (0.59)
$L_{\text{SP}}$ (w.o. replacement)	88.27 (0.43)
$L_{\text{SP}}$	88.66 (0.40)
$L_{\text{SPS}}$	88.75 (0.27)
$L_{\text{SP}} + \text{Reweighting}$	89.25 (0.42)
$L_{\text{SPS}} + \text{Reweighting}$	89.39 (0.36)
$L_{\text{SP}} + L_{\text{students}}$	88.84 (0.36)
$L_{\text{SPS}} + L_{\text{students}}$	88.93 (0.28)
$L_{\text{SP}} + L_{\text{students}} + L_{\text{teacher}}$	89.43 (0.42)
$L_{\text{SPS}} + L_{\text{students}} + L_{\text{teacher}}$	89.65 (0.33)
<b>Self-PU</b>	<b>89.68 (0.22)</b>

Table 3. Study of student diversity (learning paces) for two-student distillation on CIFAR-10 dataset. Pace1/Pace2 denotes the final ratio of  $|D_{\text{trust}}|$  over  $|D_U|$ .

Pace1	Pace2	Test Accuracy %
10%	40%	89.32 (0.36)
15%	35%	89.55 (0.46)
<b>20%</b>	<b>30%</b>	<b>89.65 (0.33)</b>
25%	25%	89.64 (0.47)

Table 4. Study of hard sample mining threshold  $\alpha$  for two-student distillation on CIFAR-10 dataset. Smaller  $\alpha$  indicates stronger distillation (Eq. (14))

$\alpha$	Test Accuracy %
5	89.59 (0.39)
<b>10</b>	<b>89.65 (0.33)</b>
20	89.38 (0.51)

#### 4.3.2. EFFECTS OF STUDENT DIVERSITY

Different learning paces enable the diversity of two students and thus make the collaborative teaching between two students effective. Therefore we study how the student diversity, *i.e.* combination of their different learning paces, can affect the final results. Table 3 considers three different pace pairs. For example, Pace1 “10%” means that the

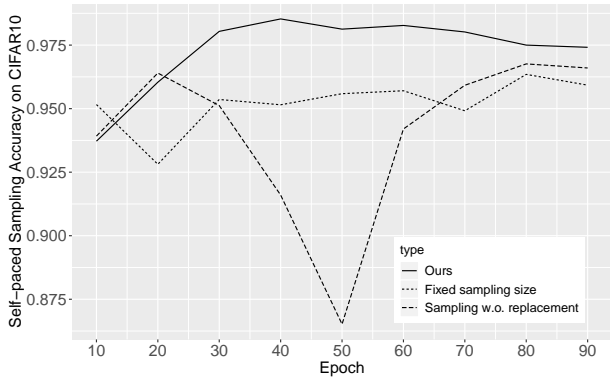


Figure 3. Accuracy of selected confident examples during self-paced learning. We compared self-paced learning with three different sampling settings: fixed sampling size (dot line), sampling without replacement (dash line), and our proposed dynamic “in-and-out” sampling (solid line). It is clear that self-paced learning with fixed sampling size or without replacement suffers from low sampling accuracy, and no-replacement is even jeopardized by the inaccurate examples remain in the  $D_{\text{trust}}$ .

Table 5. Study of smoothing coefficient  $\beta$  for teacher networks on CIFAR-10 dataset. Greater  $\beta$  indicates slower updates of teachers from the students (Eq. (16))

$\beta$	Test Accuracy %
0.2	89.37 (0.39)
<b>0.3</b>	<b>89.65 (0.33)</b>
0.4	89.47 (0.41)

Table 6. Study of  $\gamma$  for self-calibrated loss reweighting on CIFAR-10. Greater  $\gamma$  indicates larger weight on cross-entropy (Eq. (10))

$\gamma$	Accuracy
0.125	89.29%
0.100	89.42%
0.075	89.55%
<b>0.063</b>	<b>89.68%</b>
0.050	89.67%
0.000	89.65%

self-paced learning of the first student model will end up with  $|D_{\text{trust}}| = 0.1|D_U|$ , and all students will complete the sampling for self-paced learning within the same number of training epochs. Table 3 shows that, while student diversity helps (“20%” + “30%” > “25%” + “25%”), too large student pace discrepancy will hurt the learning too (“20%” + “30%” > “10%” + “40%”). Students with very different paces are harmful because a large gap in two learning paces results in a smaller intersection-over-union of  $D_{\text{trust}1}$  and  $D_{\text{trust}2}$ . It is difficult to keep consistency between two mod-

els trained with different amounts of labeled data. Therefore, it is important to keep diversity, while not too extreme.

#### 4.3.3. EFFECTS OF SAMPLE MINING THRESHOLD

$L_{\text{students}}$  takes the hard sample mining threshold  $\alpha$  as an important hyperparameter: the smaller  $\alpha$  is, the more examples are counted in computing the mean squared error, which implies stronger self-supervision consistency between two students. Table 4 shows that a moderate  $\alpha = 10$  leads to the best performance. Understandably, either “under-mining” ( $\alpha = 20$ ) or “over-mining” ( $\alpha = 5$ ) hurts the performance: the former is not sufficiently regularized, while the latter starts to eliminate the emphasis over hard examples.

#### 4.3.4. EFFECTS OF SMOOTHING COEFFICIENT $\beta$

The smoothing coefficient  $\beta$  controls how “conservative” we distill the teachers from the students: the larger  $\beta$  is, the more reluctant the teacher models get updated from the students. Table 5 investigate three different  $\beta$  values: similar to the previous experiment on  $\alpha$ ,  $\beta$  also favors a reasonably moderate value, while either overly large or small  $\beta$ , corresponding to “over-smoothing” and “under-smoothing” during the distillation from students to teachers respectively, degrades the final performance.

#### 4.3.5. EFFECTS OF $\gamma$ FOR SELF-CALIBRATED LOSS REWEIGHTING

The balancing factor  $\gamma$  restricts the total weights of the cross-entropy loss. In Table 6, we report the validation accuracy with different  $\gamma$ , and we can see that the optimal choice of  $\gamma$  is 0.063. Table 6 indicates that mining examples with our calibrated loss contribute to better supervision than using only  $L_{\text{nnPU}}$ , while too much weight on  $L_{\text{XE}}$  may lead to worse validation accuracy.

#### 4.3.6. EFFECT OF TEACHER AND STUDENTS

We verify the effect of two types of distillation in our self-supervised learning: mutually between  $L_{\text{students}}$ , and by  $L_{\text{teachers}}$ . In Table 2, distillation from two students with different learning paces ( $L_{\text{students}}$ ) improve the accuracy of nnPU baseline from 88.60% to 88.84% on CIFAR-10. By adding two teachers for self-distillation, the performance is further boosted to 89.43%, which endorses the complementary power of two types of self-distillation.

### 4.4. Comparison to State-of-the-Art Methods

#### 4.4.1. RESULTS ON THE MNIST AND CIFAR-10 BENCHMARKS

We compare the performance of the proposed Self-PU with several popular baselines: the unbiased PU learning (uPU)



Table 7. Classification comparison on MNIST and CIFAR-10. “\*” indicates that 3,000 positive examples were initialized for training, while others used 1,000.

Method	MNIST %	CIFAR-10 %
uPU (Du Plessis et al., 2014)	92.52 (0.39)	88.00 (0.62)
nnPU (Kiryo et al., 2017)	93.41 (0.20)	88.60 (0.40)
DAN* (Liu et al., 2019)	-	89.7 (0.40)
Self-PU	94.21 (0.54)	89.68 (0.22)
Self-PU*	<b>96.00 (0.29)</b>	<b>90.77 (0.21)</b>

(Du Plessis et al., 2014); the non-negative PU learning (nnPU) (Kiryo et al., 2017)<sup>7</sup>, and DAN (a latest GAN-based PU method) (Liu et al., 2019).

Table 7 summarizes the comparison results on MNIST and CIFAR-10. On MNIST, Self-PU outperforms uPU and nnPU by over 0.5%, setting new performance records. On CIFAR-10, Self-PU surpasses nnPU by over 1% (a considerable gap). More importantly, by only leveraging 1000 positive examples, Self-PU achieves comparable performance as DAN where 3000 positive samples were used. Training with 3000 positive examples further boosts our performance which outperforms DAN by 1%.




Figure 4. Validation accuracy during training on CIFAR-10. Our “Self-PU” framework<sup>9</sup> achieved a more stable training compared with uPU and nnPU methods.

Our “Self-PU” achieved not only high accuracy, but more importantly a **much more stable PU learning process** (Figure 4). As noted in (Kiryo et al., 2017), uPU suffers from overfitting complex models. We also empirically found a similar phenomenon in PU learning with nnPU risk estimator, where the validation accuracy remains unstable and even drops in the late training stage. However, the training process of Self-PU is significantly more stable than uPU and nnPU. This training stability benefits from both accurately

<sup>7</sup>We reproduced the uPU and nnPU baselines using the official codebase from: <https://github.com/kiryor/nnPUlearning>

<sup>9</sup>Since in “Self-PU” we use the teacher model  $G$  for the final prediction, the solid line shows the accuracy of  $G$  starts from epoch 50 when the self-paced training ends.

Table 8. Classification accuracy of different methods on ADNI. “naive” means that we treat the entire unlabeled class as negative.

Method	ADNI %
naive	73.27 (1.45)
uPU	73.45 (1.77)
nnPU	75.96 (1.42)
Self-PU	<b>79.50 (1.80)</b>

identified examples in self-paced training and prediction consistency forced by our self-supervised distillations.

#### 4.4.2. RESULTS ON THE NEW ADNI TESTBED

Finally, we demonstrate the promise of our method on the more complex real-world ADNI data in Table 8. We first run a naive fully supervised classification baseline, by treating the entire **unlabeled** class as **negative**. Its accuracy is much inferior to our PU learning results, validating our PU formulation of the ADNI task. Next, Self-PU gains significantly over uPU and nnPU, showing highly promising performance on ADNI and setting new state-of-the-arts. Our sophisticated building blocks seem to add robustness to handling the real-world data variations and challenges.

Furthermore, our above results seems to suggest that conventional PU benchmarks like CIFAR-10 and MNIST may have been saturated (as they already did in image classification). We recommend the community to pay more attention to more challenging and realistic new PU learning testbeds, and suggest ADNI as an effective, illustrative and practically important option.

## 5. Conclusion

We proposed Self-PU, that bridges self-training strategy into PU learning for the first time. It leverages both the self-paced selected set of trusted samples and the consistency supervision via self-distillation and self-calibration. Experiments report state-of-the-art performance of Self-PU on two conventional (and potentially oversimplified) benchmarks, plus our newly introduced real-world PU testbed of ADNI classification. Our future work will explore more realistic PU learning setting, which we believe will motivate new algorithmic findings.

## References

- Armenian, H. and Lilienfeld, A. The distribution of incubation periods of neoplastic diseases. *American journal of epidemiology*, 1974.
- Arpit, D., Jastrzebski, S., Ballas, N., Krueger, D., Bengio, E., Kanwal, M., Maharaj, T., Fischer, A., Courville, A.,

- and Bengio, Y. A closer look at memorization in deep networks. In *ICML*, pp. 233–242, 2017.
- Bengio, Y., Louradour, J., Collobert, R., and Weston, J. Curriculum learning. In *ICML*, pp. 41–48, 2009.
- Chen, T., Liu, S., Chang, S., Cheng, Y., Amini, L., and Wang, Z. Adversarial robustness: From self-supervised pre-training to fine-tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 699–708, 2020.
- Du Plessis, M., Niu, G., and Sugiyama, M. Convex formulation for learning from positive and unlabeled data. In *ICML*, pp. 1386–1394, 2015.
- Du Plessis, M. C., Niu, G., and Sugiyama, M. Analysis of learning from positive and unlabeled data. In *NeurIPS*, 2014.
- Duyckaerts, C. and Hauw, J.-J. Diagnosis and staging of alzheimer disease. *Neurobiology of aging*, 18(4):S33–S42, 1997.
- Elkan, C. and Noto, K. Learning classifiers from only positive and unlabeled data. In *ACM SIGKDD*, 2008.
- Hou, M., Chaib-draa, B., Li, C., and Zhao, Q. Generative adversarial positive-unlabelled learning. *IJCAI*, Jul 2018. doi: 10.24963/ijcai.2018/312.
- Jack Jr, C., Bernstein, M., Fox, N., Thompson, P., Alexander, G., Harvey, D., Borowski, B., Britson, P., Whitwell, J., and Ward, C. The alzheimer’s disease neuroimaging initiative (adni): Mri methods. *Journal of Magnetic Resonance Imaging*, pp. 685–691, 2008.
- Jing, L. and Tian, Y. Self-supervised visual feature learning with deep neural networks: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.
- Khan, F., Mutlu, B., and Zhu, J. How do humans teach: On curriculum learning and teaching dimension. In *NeurIPS*, pp. 1449–1457, 2011.
- Khvostikov, A., Aderghal, K., Benois-Pineau, J., Krylov, A., and Catheline, G. 3d cnn-based classification using smri and md-dti images for alzheimer disease studies. *arXiv preprint arXiv:1801.05968*, 2018.
- Kiryo, R., Niu, G., du Plessis, M. C., and Sugiyama, M. Positive-unlabeled learning with non-negative risk estimator. In *NeurIPS*, pp. 1675–1685, 2017.
- Kumar, M. P., Packer, B., and Koller, D. Self-paced learning for latent variable models. In *NeurIPS*. 2010.
- Laine, S. and Aila, T. Temporal ensembling for semi-supervised learning. *arXiv:1610.02242*, 2016.
- Larson, E. B., Shadlen, M.-F., Wang, L., McCormick, W. C., Bowen, J. D., Teri, L., and Kukull, W. A. Survival after initial diagnosis of alzheimer disease. *Annals of internal medicine*, 140(7):501–509, 2004.
- Li, F., Tran, L., Thung, K.-H., Ji, S., Shen, D., and Li, J. A robust deep model for improved classification of ad/mci patients. *Biomedical Health Informatics*, 2015.
- Li, X. and Liu, B. Learning to classify texts using positive and unlabeled data. In *IJCAI*, 2003.
- Liu, B., Lee, W. S., Yu, P. S., and Li, X. Partially supervised classification of text documents. In *ICML*, 2002.
- Liu, F., Chen, H., and Wu, H. Discriminative adversarial networks for positive-unlabeled learning. *arXiv:1906.00642*, 2019.
- Mohseni, S., Pitale, M., Yadawa, J., and Wang, Z. Self-supervised learning for generalizable out-of-distribution detection. *AAAI*, 2020.
- Ott, B. R., Jones, R. N., Noto, R. B., Yoo, D. C., Snyder, P. J., Bernier, J. N., Carr, D. B., and Roe, C. M. Brain amyloid in preclinical alzheimer’s disease is associated with increased driving risk. *Alzheimer’s & Dementia: Diagnosis, Assessment & Disease Monitoring*, 6:136–142, 2017.
- Ren, M., Zeng, W., Yang, B., and Urtasun, R. Learning to reweight examples for robust deep learning. *arXiv preprint arXiv:1803.09050*, 2018.
- Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., and Wojna, Z. Rethinking the inception architecture for computer vision. *CVPR*, Jun 2016. doi: 10.1109/cvpr.2016.308.
- Tarvainen, A. and Valpola, H. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *NeurIPS*, 2017.
- Trinh, T. H., Luong, M.-T., and Le, Q. V. Selfie: Self-supervised pretraining for image embedding. *arXiv preprint arXiv:1906.02940*, 2019.
- Villeneuve, S., Rabinovici, G. D., Cohn-Sheehy, B. I., Madison, C., Ayakta, N., Ghosh, P. M., La Joie, R., Arthur-Bentil, S. K., Vogel, J. W., Marks, S. M., et al. Existing pittsburgh compound-b positron emission tomography thresholds are too high: statistical and pathological evaluation. *Brain*, 138(7):2020–2033, 2015.
- Xu, M., Li, B., Niu, G., Han, B., and Sugiyama, M. Revisiting sample selection approach to positive-unlabeled learning: Turning unlabeled data into positive rather than negative. 2019a.

- Xu, Y., Xu, C., Xu, C., and Tao, D. Multi-positive and unlabeled learning. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, pp. 3182–3188, 2017.
- Xu, Y., Wang, Y., Chen, H., Han, K., Chunjing, X., Tao, D., and Xu, C. Positive-unlabeled compression on the cloud. In *Advances in Neural Information Processing Systems*, pp. 2561–2570, 2019b.
- Yuan, Y., Wang, Z., Lee, W., Thiyyagura, P., Reiman, E. M., and Chen, K. Feasibility of quantifying amyloid burden using volumetric mri data: Preliminary findings based on the deep learning 3d convolutional neural network approach. *Alzheimer's & Dementia: The Journal of the Alzheimer's Association*, 14(7):P695, 2018.
- Zhang, Y., Xiang, T., Hospedales, T. M., and Lu, H. Deep mutual learning. In *ICCV*, 2018.