
Classification from Positive, Unlabeled and Biased Negative Data

Yu-Guan Hsieh^{*1} Gang Niu² Masashi Sugiyama^{2,3}

Abstract

In binary classification, there are situations where negative (N) data are too diverse to be fully labeled and we often resort to positive-unlabeled (PU) learning in these scenarios. However, collecting a non-representative N set that contains only a small portion of all possible N data can often be much easier in practice. This paper studies a novel classification framework which incorporates such biased N (bN) data in PU learning. We provide a method based on empirical risk minimization to address this PUBN classification problem. Our approach can be regarded as a novel example-weighting algorithm, with the weight of each example computed through a preliminary step that draws inspiration from PU learning. We also derive an estimation error bound for the proposed method. Experimental results demonstrate the effectiveness of our algorithm in not only PUBN learning scenarios but also ordinary PU learning scenarios on several benchmark datasets.

1. Introduction

In conventional binary classification, examples are labeled as either positive (P) or negative (N), and we train a classifier on these labeled examples. On the contrary, positive-unlabeled (PU) learning addresses the problem of learning a classifier from P and unlabeled (U) data, without the need of explicitly identifying N data (Elkan & Noto, 2008; Ward et al., 2009).

PU learning finds its usefulness in many real-world problems. For example, in one-class remote sensing classification (Li et al., 2011), we seek to extract a specific land-cover class from an image. While it is easy to label examples of this specific land-cover class of interest, examples not

belonging to this class are too diverse to be exhaustively annotated. The same problem arises in text classification, as it is difficult or even impossible to compile a set of N samples that provides a comprehensive characterization of everything that is not in the P class (Liu et al., 2003; Fung et al., 2006). Besides, PU learning has also been applied to other domains such as outlier detection (Hido et al., 2008; Scott & Blanchard, 2009), medical diagnosis (Zuluaga et al., 2011), or time series classification (Nguyen et al., 2011).

By carefully examining the above examples, we find out that the most difficult step is often to collect a fully representative N set, whereas only labeling a small portion of all possible N data is relatively easy. Therefore, in this paper, we propose to study the problem of learning from P, U and biased N (bN) data, which we name PUBN learning hereinafter. We suppose that in addition to P and U data, we also gather a set of bN samples, governed by a distribution distinct from the true N distribution. As described previously, this can be viewed as an extension of PU learning, but such bias may also occur naturally in some real-world scenarios. For instance, let us presume that we would like to judge whether a subject is affected by a particular disease based on the result of a physical examination. While the data collected from the patients represent rather well the P distribution, healthy subjects that request the examination are in general biased with respect to the whole healthy subject population.

We are not the first to be interested in learning with bN data. In fact, both Li et al. (2010) and Fei & Liu (2015) attempted to solve similar problems in the context of text classification. Li et al. (2010) simply discarded N samples and performed ordinary PU classification. It was also mentioned in the paper that bN data could be harmful. Fei & Liu (2015) adopted another strategy. The authors considered even gathering unbiased U data is difficult and learned the classifier from only P and bN data. However, their method is specific to text classification because it relies on the use of effective similarity measures to evaluate similarity between documents (refer to Supplementary Material D.5 for a deeper discussion and an empirical comparison with our method). Therefore, our work differs from these two in that the classifier is trained simultaneously on P, U and bN data, without resorting to domain-specific knowledge. The presence of U data allows us to address the problem from a statistical viewpoint, and thus the proposed method can be

¹École Normale Supérieure, Paris, France ²RIKEN, Tokyo, Japan ³The University of Tokyo, Tokyo, Japan. Correspondence to: Yu-Guan Hsieh <yu-guan.hsieh@ens.fr>.

* Most of the work done during internship at RIKEN.

applied to any PUBN learning problem in principle.

In this paper, we develop an empirical risk minimization-based algorithm that combines both PU learning and importance weighting to solve the PUBN classification problem. We first estimate the probability that an example is sampled into the P or the bN set. Based on this estimate, we regard bN and U data as N examples with instance-dependent weights. In particular, we assign larger weights to U examples that we believe to appear less often in the P and bN sets. P data are treated as P examples with unity weight but also as N examples with usually small or zero weight whose actual value depends on the same estimate.

The contributions of the paper are three-fold:

1. We formulate the PUBN learning problem as an extension of PU learning and propose an empirical risk minimization-based method to address the problem. We also theoretically establish an estimation error bound for the proposed method.
2. We experimentally demonstrate that the classification performance can be effectively improved thanks to the use of bN data during training. In other words, PUBN learning yields better performance than PU learning.
3. Our method can be easily adapted to ordinary PU learning. Experimentally we show that the resulting algorithm allows us to obtain new state-of-the-art results on several PU learning tasks.

Relation with Semi-supervised Learning. With P, N and U data available for training, our problem setup may seem similar to that of semi-supervised learning (Chapelle et al., 2010; Oliver et al., 2018). Nonetheless, in our case, N data are biased and often represent only a small portion of the whole N distribution. Therefore, most of the existing methods designed for the latter cannot be directly applied to the PUBN classification problem. Furthermore, our focus is on deducing a risk estimator using the three sets of data, with U data in particular used to compensate the sampling bias in N data. On the other hand, in semi-supervised learning the main concern is often how U data can be utilized for regularization (Grandvalet & Bengio, 2005; Belkin et al., 2006; Miyato et al., 2016; Laine & Aila, 2017). The two should be compatible and we believe adding such regularization to our algorithm can be beneficial in many cases.

Relation with Dataset Shift. PUBN learning can also be viewed as a special case of dataset shift¹ (Quionero-Candela

¹ Dataset shift refers to any case where training and test distributions differ. The term sample selection bias (Heckman, 1979; Zadrozny, 2004) is sometimes used to describe the same thing. However, strictly speaking, sample selection bias actually refers to the case where training instances are first drawn from the test distributions and then a subset of these data is systematically discarded due to a particular mechanism.

et al., 2009) if we consider that P and bN data are drawn from the training distribution while U data are drawn from the test distribution. Covariate shift (Shimodaira, 2000; Sugiyama & Kawanabe, 2012) is another special case of dataset shift that has been studied intensively. In the covariate shift problem setting, training and test distributions have the same class conditional distribution and only differ in the marginal distribution of the independent variable. One popular approach to tackle this problem is to reweight each training example according to the ratio of the test density to the training density (Huang et al., 2007; Sugiyama et al., 2008). Nevertheless, simply training a classifier on a reweighted version of the labeled set is not sufficient in our case since there may be examples with zero probability to be labeled, and it is therefore essential to involve U samples in the second step of the proposed algorithm. It is also important to notice that the problem of PUBN learning is intrinsically different from that of covariate shift and neither of the two is a special case of the other.

Finally, source component shift (Quionero-Candela et al., 2009) is also related. It assumes that data are generated from several different sources and the proportions of these sources may vary between training and test times. In many practical situations, this is indeed what causes our collected N data to be biased. However, its definition is so general that we are not aware of any universal method which addresses this problem without explicit model assumptions on data distribution.

2. Problem Setting

In this section, we briefly review the formulations of PN, PU and PNU classification and introduce the problem of learning from P, U and bN data.

2.1. Standard Binary Classification

Let $\mathbf{x} \in \mathbb{R}^d$ and $y \in \{+1, -1\}$ be random variables following an unknown probability distribution with density $p(\mathbf{x}, y)$. Let $g : \mathbb{R}^d \rightarrow \mathbb{R}$ be an arbitrary decision function for binary classification and $\ell : \mathbb{R} \rightarrow \mathbb{R}_+$ be a loss function of margin $yg(\mathbf{x})$ that usually takes a small value for a large margin. The goal of binary classification is to find g that minimizes the classification risk:

$$R(g) = \mathbb{E}_{(\mathbf{x}, y) \sim p(\mathbf{x}, y)}[\ell(yg(\mathbf{x}))], \quad (1)$$

where $\mathbb{E}_{(\mathbf{x}, y) \sim p(\mathbf{x}, y)}[\cdot]$ denotes the expectation over the joint distribution $p(\mathbf{x}, y)$. When we care about classification accuracy, ℓ is the zero-one loss $\ell_{01}(z) = (1 - \text{sign}(z))/2$. However, for ease of optimization, ℓ_{01} is often substituted with a surrogate loss such as the sigmoid loss $\ell_{\text{sig}}(z) = 1/(1 + \exp(z))$ or the logistic loss $\ell_{\text{log}}(z) = \ln(1 + \exp(-z))$ during learning.

In standard supervised learning scenarios (PN classification), we are given P and N data that are sampled independently from $p_P(\mathbf{x}) = p(\mathbf{x} \mid y = +1)$ and $p_N(\mathbf{x}) = p(\mathbf{x} \mid y = -1)$ as $\mathcal{X}_P = \{\mathbf{x}_i^P\}_{i=1}^{n_P}$ and $\mathcal{X}_N = \{\mathbf{x}_i^N\}_{i=1}^{n_N}$. Let us denote by $R_P^+(g) = \mathbb{E}_{\mathbf{x} \sim p_P(\mathbf{x})}[\ell(g(\mathbf{x}))]$, $R_N^-(g) = \mathbb{E}_{\mathbf{x} \sim p_N(\mathbf{x})}[\ell(-g(\mathbf{x}))]$ partial risks and $\pi = p(y = 1)$ the P prior. We have the equality $R(g) = \pi R_P^+(g) + (1 - \pi)R_N^-(g)$. The classification risk (1) can then be empirically approximated from data by

$$\hat{R}_{PN}(g) = \pi \hat{R}_P^+(g) + (1 - \pi) \hat{R}_N^-(g),$$

where $\hat{R}_P^+(g) = \frac{1}{n_P} \sum_{i=1}^{n_P} \ell(g(\mathbf{x}_i^P))$ and $\hat{R}_N^-(g) = \frac{1}{n_N} \sum_{i=1}^{n_N} \ell(-g(\mathbf{x}_i^N))$. By minimizing $\hat{R}_{PN}(g)$ we obtain the ordinary empirical risk minimizer \hat{g}_{PN} .

2.2. PU Classification

In PU classification, instead of N data \mathcal{X}_N we have only access to $\mathcal{X}_U = \{\mathbf{x}_i^U\}_{i=1}^{n_U} \sim p(\mathbf{x})$ a set of U samples drawn from the marginal density $p(\mathbf{x})$. Several effective algorithms have been designed to address this problem. Liu et al. (2002) proposed the S-EM approach that first identifies reliable N data in the U set and then run the Expectation-Maximization (EM) algorithm to build the final classifier. The biased support vector machine (Biased SVM) introduced by Liu et al. (2003) regards U samples as N samples with smaller weights. Mordelet & Vert (2014) solved the PU problem by aggregating classifiers trained to discriminate P data from a small random subsample of U data. An ad hoc algorithm designed for linear classifiers, treating the U set as an N set influenced by label noise, was proposed in (Shi et al., 2018).

Recently, attention has also been paid on the unbiased risk estimator proposed by du Plessis et al. (2014; 2015). The key idea is to use the following equality:

$$(1 - \pi)R_N^-(g) = R_U^-(g) - \pi R_P^-(g),$$

where $R_U^-(g) = \mathbb{E}_{\mathbf{x} \sim p(\mathbf{x})}[\ell(-g(\mathbf{x}))]$ and $R_P^-(g) = \mathbb{E}_{\mathbf{x} \sim p_P(\mathbf{x})}[\ell(-g(\mathbf{x}))]$. This equality is acquired by exploiting the fact $p(\mathbf{x}) = \pi p_P(\mathbf{x}) + (1 - \pi)p_N(\mathbf{x})$. As a result, we can approximate the classification risk (1) by

$$\hat{R}_{PU}(g) = \pi \hat{R}_P^+(g) - \pi \hat{R}_P^-(g) + \hat{R}_U^-(g), \quad (2)$$

where $\hat{R}_P^-(g) = \frac{1}{n_P} \sum_{i=1}^{n_P} \ell(-g(\mathbf{x}_i^P))$ and $\hat{R}_U^-(g) = \frac{1}{n_U} \sum_{i=1}^{n_U} \ell(-g(\mathbf{x}_i^U))$. We then minimize $\hat{R}_{PU}(g)$ to obtain another empirical risk minimizer \hat{g}_{PU} . Note that as the loss is always positive, the classification risk (1) that $\hat{R}_{PU}(g)$ approximates is also positive. However, Kiryo et al. (2017) pointed out that when the model of g is too flexible, that is, when the function class \mathcal{G} is too large, $\hat{R}_{PU}(\hat{g}_{PU})$ indeed goes negative and the model severely overfits the training

data. To alleviate overfitting, the authors observed that $R_U^-(g) - \pi R_P^-(g) = (1 - \pi)R_N^-(g) \geq 0$ and proposed the non-negative risk estimator for PU learning:

$$\hat{R}_{PU}(g) = \pi \hat{R}_P^+(g) + \max\{0, \hat{R}_U^-(g) - \pi \hat{R}_P^-(g)\}. \quad (3)$$

In terms of implementation, stochastic optimization was used and when $r = \hat{R}_U^-(g) - \pi \hat{R}_P^-(g)$ becomes smaller than some threshold value $-\beta$ for a mini-batch, they performed a step of gradient ascent along ∇r to make the mini-batch less overfitted.

2.3. PNU Classification

In semi-supervised learning (PNU classification), P, N and U data are all available. An abundance of works have been dedicated to solving this problem. Here we in particular introduce the PNU risk estimator proposed by Sakai et al. (2017). By directly leveraging U data for risk estimation, it is the most comparable to our method. The PNU risk is simply defined as a linear combination of PN and PU/NU risks. Let us just consider the case where PN and PU risks are combined, then for some $\gamma \in [0, 1]$, the PNU risk estimator is expressed as

$$\begin{aligned} \hat{R}_{PNU}^\gamma(g) &= \gamma \hat{R}_{PN}(g) + (1 - \gamma) \hat{R}_{PU}(g) \\ &= \pi \hat{R}_P^+(g) + \gamma(1 - \pi) \hat{R}_N^-(g) \\ &\quad + (1 - \gamma)(\hat{R}_U^-(g) - \pi \hat{R}_P^-(g)). \end{aligned} \quad (4)$$

We can again consider the non-negative correction by forcing the term $\gamma(1 - \pi) \hat{R}_N^-(g) + (1 - \gamma)(\hat{R}_U^-(g) - \pi \hat{R}_P^-(g))$ to be non-negative. In the rest of the paper, we refer to the resulting algorithm as non-negative PNU (nnPNU) learning (see Supplementary Material D.4 for an alternative definition of nnPNU and the corresponding results).

2.4. PUBN Classification

In this paper, we study the problem of PUBN learning. It differs from usual semi-supervised learning in the fact that labeled N data are not fully representative of the underlying N distribution $p_N(\mathbf{x})$. To take this point into account, we introduce a latent random variable s and consider the joint distribution $p(\mathbf{x}, y, s)$ with constraint $p(s = +1 \mid \mathbf{x}, y = +1) = 1$. Equivalently, $p(y = -1 \mid \mathbf{x}, s = -1) = 1$. Let $\rho = p(y = -1, s = +1)$. Both π and ρ are assumed known throughout the paper. In practice they often need to be estimated from data (Jain et al., 2016; Ramaswamy et al., 2016; du Plessis et al., 2017). In place of ordinary N data we collect a set of bN samples

$$\mathcal{X}_{bN} = \{\mathbf{x}_i^{bN}\}_{i=1}^{n_{bN}} \sim p_{bN}(\mathbf{x}) = p(\mathbf{x} \mid y = -1, s = +1).$$

For instance, in text classification, if our bN data is composed of a small set of all possible N topics, $s = +1$ means

that a sample is either from these topics that make up the bN set or in the P class. The goal remains the same: we would like to minimize the classification risk (1).

3. Method

In this section, we propose a risk estimator for PUbN classification and establish an estimation error bound for the proposed method. Finally we show how our method can be applied to PU learning as a special case when no bN data are available.

3.1. Risk Estimator

Let $R_{\text{bN}}^-(g) = \mathbb{E}_{\mathbf{x} \sim p_{\text{bN}}(\mathbf{x})}[\ell(-g(\mathbf{x}))]$ and $R_{s=-1}^-(g) = \mathbb{E}_{\mathbf{x} \sim p(\mathbf{x} | s=-1)}[\ell(-g(\mathbf{x}))]$. Since $p(\mathbf{x}, y = -1) = p(\mathbf{x}, y = -1, s = +1) + p(\mathbf{x}, s = -1)$, we have

$$R(g) = \pi R_{\text{P}}^+(g) + \rho R_{\text{bN}}^-(g) + (1 - \pi - \rho) R_{s=-1}^-(g). \quad (5)$$

The first two terms on the right-hand side of the equation can be approximated directly from data by $\hat{R}_{\text{P}}^+(g)$ and $\hat{R}_{\text{bN}}^-(g) = \frac{1}{n_{\text{bN}}} \sum_{i=1}^{n_{\text{bN}}} \ell(-g(\mathbf{x}_i^{\text{bN}}))$. We therefore focus on the third term $\bar{R}_{s=-1}^-(g) = (1 - \pi - \rho) R_{s=-1}^-(g)$. Our approach is mainly based on the following theorem. We relegate all proofs to the Supplementary Material.

Theorem 1. *Let $\sigma(\mathbf{x}) = p(s = +1 | \mathbf{x})$. For all $\eta \in [0, 1]$ and $h : \mathbb{R}^d \rightarrow [0, 1]$ satisfying the condition $h(\mathbf{x}) > \eta \Rightarrow \sigma(\mathbf{x}) > 0$, the risk $\bar{R}_{s=-1}^-(g)$ can be expressed as*

$$\begin{aligned} \bar{R}_{s=-1}^-(g) &= \mathbb{E}_{\mathbf{x} \sim p(\mathbf{x})} [\mathbb{1}_{h(\mathbf{x}) \leq \eta} \ell(-g(\mathbf{x})) (1 - \sigma(\mathbf{x}))] \\ &+ \pi \mathbb{E}_{\mathbf{x} \sim p_{\text{P}}(\mathbf{x})} \left[\mathbb{1}_{h(\mathbf{x}) > \eta} \ell(-g(\mathbf{x})) \frac{1 - \sigma(\mathbf{x})}{\sigma(\mathbf{x})} \right] \\ &+ \rho \mathbb{E}_{\mathbf{x} \sim p_{\text{bN}}(\mathbf{x})} \left[\mathbb{1}_{h(\mathbf{x}) > \eta} \ell(-g(\mathbf{x})) \frac{1 - \sigma(\mathbf{x})}{\sigma(\mathbf{x})} \right]. \end{aligned} \quad (6)$$

In the theorem, $\bar{R}_{s=-1}^-(g)$ is decomposed into three terms, and when the expectation is substituted with the average over training samples, these three terms are approximated respectively using data from \mathcal{X}_{U} , \mathcal{X}_{P} and \mathcal{X}_{bN} . The choice of h and η is thus very crucial because it determines what each of the three terms tries to capture in practice. Ideally, we would like h to be an approximation of σ . Then, for \mathbf{x} such that $h(\mathbf{x})$ is close to 1, $\sigma(\mathbf{x})$ is close to 1, so the last two terms on the right-hand side of the equation can be reasonably evaluated using \mathcal{X}_{P} and \mathcal{X}_{bN} (i.e., samples drawn from $p(\mathbf{x} | s = +1)$). On the contrary, if $h(\mathbf{x})$ is small, $\sigma(\mathbf{x})$ is small and such samples can be hardly found in \mathcal{X}_{P} or \mathcal{X}_{bN} . Consequently the first term appeared in the decomposition is approximated with the help of \mathcal{X}_{U} . Finally, in the empirical risk minimization paradigm, η becomes a hyperparameter that controls how important U data is

against P and bN data when we evaluate $\bar{R}_{s=-1}^-(g)$. The larger η is, the more attention we would pay to U data.

One may be curious about why we do not simply approximate the whole risk using only U samples, that is, set η to 1. There are two main reasons. On one hand, if we have a very small U set, which means $n_{\text{U}} \ll n_{\text{P}}$ and $n_{\text{U}} \ll n_{\text{bN}}$, approximating a part of the risk with labeled samples should help us reduce the estimation error. This may seem unrealistic but sometimes unbiased U samples can also be difficult to collect (Ishida et al., 2018). On the other hand, more importantly, we have empirically observed that when the model of g is highly flexible, even a sample regarded as N with small weight gets classified as N in the latter stage of training and performance of the resulting classifier can thus be severely degraded. Introducing η alleviates this problem by avoiding treating all U data as N samples.

As σ is not available in reality, we propose to replace σ by its estimate $\hat{\sigma}$ in (6). We further substitute h with the same estimate and obtain the following expression:

$$\begin{aligned} \bar{R}_{s=-1, \eta, \hat{\sigma}}^-(g) &= \mathbb{E}_{\mathbf{x} \sim p(\mathbf{x})} [\mathbb{1}_{\hat{\sigma}(\mathbf{x}) \leq \eta} \ell(-g(\mathbf{x})) (1 - \hat{\sigma}(\mathbf{x}))] \\ &+ \pi \mathbb{E}_{\mathbf{x} \sim p_{\text{P}}(\mathbf{x})} \left[\mathbb{1}_{\hat{\sigma}(\mathbf{x}) > \eta} \ell(-g(\mathbf{x})) \frac{1 - \hat{\sigma}(\mathbf{x})}{\hat{\sigma}(\mathbf{x})} \right] \\ &+ \rho \mathbb{E}_{\mathbf{x} \sim p_{\text{bN}}(\mathbf{x})} \left[\mathbb{1}_{\hat{\sigma}(\mathbf{x}) > \eta} \ell(-g(\mathbf{x})) \frac{1 - \hat{\sigma}(\mathbf{x})}{\hat{\sigma}(\mathbf{x})} \right]. \end{aligned}$$

We notice that $\bar{R}_{s=-1, \eta, \hat{\sigma}}^-$ depends both on η and $\hat{\sigma}$. It can be directly approximated from data by

$$\begin{aligned} \hat{R}_{s=-1, \eta, \hat{\sigma}}(g) &= \frac{1}{n_{\text{U}}} \sum_{i=1}^{n_{\text{U}}} \left[\mathbb{1}_{\hat{\sigma}(\mathbf{x}_i^{\text{U}}) \leq \eta} \ell(-g(\mathbf{x}_i^{\text{U}})) (1 - \hat{\sigma}(\mathbf{x}_i^{\text{U}})) \right] \\ &+ \frac{\pi}{n_{\text{P}}} \sum_{i=1}^{n_{\text{P}}} \left[\mathbb{1}_{\hat{\sigma}(\mathbf{x}_i^{\text{P}}) > \eta} \ell(-g(\mathbf{x}_i^{\text{P}})) \frac{1 - \hat{\sigma}(\mathbf{x}_i^{\text{P}})}{\hat{\sigma}(\mathbf{x}_i^{\text{P}})} \right] \\ &+ \frac{\rho}{n_{\text{bN}}} \sum_{i=1}^{n_{\text{bN}}} \left[\mathbb{1}_{\hat{\sigma}(\mathbf{x}_i^{\text{bN}}) > \eta} \ell(-g(\mathbf{x}_i^{\text{bN}})) \frac{1 - \hat{\sigma}(\mathbf{x}_i^{\text{bN}})}{\hat{\sigma}(\mathbf{x}_i^{\text{bN}})} \right]. \end{aligned}$$

We are now able to derive the empirical version of Equation (5) as

$$\hat{R}_{\text{PubN}, \eta, \hat{\sigma}}(g) = \pi \hat{R}_{\text{P}}^+(g) + \rho \hat{R}_{\text{bN}}^-(g) + \hat{R}_{s=-1, \eta, \hat{\sigma}}^-(g). \quad (7)$$

3.2. Practical Implementation

To complete our algorithm, we need to be able to estimate σ and find appropriate η . Given that the value of η can be hard to tune, we introduce another intermediate hyperparameter τ and choose η such that $\#\{\mathbf{x} \in \mathcal{X}_{\text{U}} | \hat{\sigma}(\mathbf{x}) \leq \eta\} = \lfloor \tau(1 - \pi - \rho)n_{\text{U}} \rfloor$, where $\lfloor \cdot \rfloor$ is the floor function. The number $\tau(1 - \pi - \rho)$ is then the portion of unlabeled samples that are involved in the second step of our algorithm. Intuitively, we can set a higher τ and include more U samples in the minimization of (7) when we have a good

Algorithm 1 PUBN Classification

- 1: **Input:** data $(\mathcal{X}_P, \mathcal{X}_{bN}, \mathcal{X}_U)$, hyperparameter τ
- 2: **Step 1:**
- 3: Compute $\hat{\sigma}$ by minimizing an nnPU risk involving \mathcal{X}_P , \mathcal{X}_{bN} as P data and \mathcal{X}_U as U data
- 4: **Step 2:**
- 5: Initialize model parameter θ of g
- 6: Choose \mathcal{A} a SGD-like stochastic optimization algorithm
- 7: Set η such that

$$\#\{x \in \mathcal{X}_U \mid \hat{\sigma}(x) \leq \eta\} = \lfloor \tau(1 - \pi - \rho)n_U \rfloor$$
- 8: **for** $i = 1 \dots \mathbf{do}$
- 9: Shuffle $(\mathcal{X}_P, \mathcal{X}_{bN}, \mathcal{X}_U)$ into M mini-batches
- 10: **for** each mini-batch $(\mathcal{X}_P^j, \mathcal{X}_{bN}^j, \mathcal{X}_U^j)$ **do**
- 11: Compute the corresponding $\hat{R}_{\text{PUBN}, \eta, \hat{\sigma}}(g)$
- 12: Use \mathcal{A} to update θ with the gradient information $\nabla_{\theta} \hat{R}_{\text{PUBN}, \eta, \hat{\sigma}}(g)$
- 13: **end for**
- 14: **end for**
- 15: **Return:** θ minimizing the validation loss

estimate $\hat{\sigma}$ and otherwise we should prefer a smaller τ to reduce the negative effect that can be caused by the use of $\hat{\sigma}$ of poor quality. The use of validation data to select the final τ should also be prioritized as what we do in the experimental part.

Estimating σ . If we regard s as a class label, the problem of estimating σ is then equivalent to training a probabilistic classifier separating the classes with $s = +1$ and $s = -1$. Upon noting that $(\pi + \rho)\mathbb{E}_{\mathbf{x} \sim p(\mathbf{x}|s=+1)}[\ell(\epsilon g(\mathbf{x}))] = \pi\mathbb{E}_{\mathbf{x} \sim p_P(\mathbf{x})}[\ell(\epsilon g(\mathbf{x}))] + \rho\mathbb{E}_{\mathbf{x} \sim p_{bN}(\mathbf{x})}[\ell(\epsilon g(\mathbf{x}))]$ for $\epsilon \in \{+1, -1\}$, it is straightforward to apply nnPU learning with availability of \mathcal{X}_P , \mathcal{X}_{bN} and \mathcal{X}_U to minimize $\mathbb{E}_{(\mathbf{x}, s) \sim p(\mathbf{x}, s)}[\ell(sg(\mathbf{x}))]$. In other words, here we regard \mathcal{X}_P and \mathcal{X}_{bN} as P and \mathcal{X}_U as U, and attempt to solve a PU learning problem by applying nnPU. Since we are interested in the class-posterior probabilities, we minimize the risk with respect to the logistic loss and apply the sigmoid function to the output of the model to get $\hat{\sigma}(\mathbf{x})$. However, the above risk estimator accepts any reasonable $\hat{\sigma}$ and we are not limited to using nnPU for computing $\hat{\sigma}$. For example, the least-squares fitting approach proposed by Kanamori et al. (2009) for direct density ratio estimation can also be adapted to solving the problem.

To handle large datasets, it is preferable to adopt stochastic optimization algorithms to minimize $\hat{R}_{\text{PUBN}, \eta, \hat{\sigma}}(g)$.

3.3. Estimation Error Bound

Here we establish an estimation error bound for the proposed method. Let \mathcal{G} be the function class from which we find a function. The Rademacher complexity of \mathcal{G} for the samples of size n drawn from $q(\mathbf{x})$ is defined as

$$\mathfrak{R}_{n, q}(\mathcal{G}) = \mathbb{E}_{\mathcal{X} \sim q^n} \mathbb{E}_{\xi} \left[\sup_{g \in \mathcal{G}} \frac{1}{n} \sum_{x_i \in \mathcal{X}} \xi_i g(\mathbf{x}_i) \right],$$

where $\mathcal{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ and $\xi = \{\xi_1, \dots, \xi_n\}$ with each \mathbf{x}_i drawn from $q(\mathbf{x})$ and ξ_i as a Rademacher variable (Mohri et al., 2012). In the following we will assume that $\mathfrak{R}_{n, q}(\mathcal{G})$ vanishes asymptotically as $n \rightarrow \infty$. This holds for most of the common choices of \mathcal{G} if proper regularization is considered (Bartlett & Mendelson, 2002; Golowich et al., 2018). Assume additionally the existence of $C_g > 0$ such that $\sup_{g \in \mathcal{G}} \|g\|_{\infty} \leq C_g$ as well as $C_{\ell} > 0$ such that $\sup_{|z| \leq C_g} \ell(z) \leq C_{\ell}$. We also assume that ℓ is Lipschitz continuous on the interval $[-C_g, C_g]$ with a Lipschitz constant L_{ℓ} .

Theorem 2. Let $g^* = \arg \min_{g \in \mathcal{G}} R(g)$ be the true risk minimizer and $\hat{g}_{\text{PUBN}, \eta, \hat{\sigma}} = \arg \min_{g \in \mathcal{G}} \hat{R}_{\text{PUBN}, \eta, \hat{\sigma}}(g)$ be the PUBN empirical risk minimizer. We suppose that $\hat{\sigma}$ is a fixed function independent of data used to compute $\hat{R}_{\text{PUBN}, \eta, \hat{\sigma}}(g)$ and $\eta \in (0, 1]$. Let $\zeta = p(\hat{\sigma}(\mathbf{x}) \leq \eta)$ and $\epsilon = \mathbb{E}_{\mathbf{x} \sim p(\mathbf{x})} [|\hat{\sigma}(\mathbf{x}) - \sigma(\mathbf{x})|^2]$. Then for any $\delta > 0$, with probability at least $1 - \delta$,

$$\begin{aligned} & R(\hat{g}_{\text{PUBN}, \eta, \hat{\sigma}}) - R(g^*) \\ & \leq 4L_{\ell} \mathfrak{R}_{n_U, p}(\mathcal{G}) + \frac{4\pi L_{\ell}}{\eta} \mathfrak{R}_{n_P, p_P}(\mathcal{G}) + \frac{4\rho L_{\ell}}{\eta} \mathfrak{R}_{n_{bN}, p_{bN}}(\mathcal{G}) \\ & \quad + 2C_{\ell} \sqrt{\frac{\ln(6/\delta)}{2n_U}} + \frac{2\pi C_{\ell}}{\eta} \sqrt{\frac{\ln(6/\delta)}{2n_P}} + \frac{2\rho C_{\ell}}{\eta} \sqrt{\frac{\ln(6/\delta)}{2n_{bN}}} \\ & \quad + 2C_{\ell} \sqrt{\zeta \epsilon} + \frac{2C_{\ell}}{\eta} \sqrt{(1 - \zeta)\epsilon}. \end{aligned}$$

Theorem 2 combined with the Borel-Cantelli lemma implies that as $n_P \rightarrow \infty$, $n_{bN} \rightarrow \infty$ and $n_U \rightarrow \infty$, the inequality $\limsup R(\hat{g}_{\text{PUBN}, \eta, \hat{\sigma}}) - R(g^*) \leq 2C_{\ell} \sqrt{\zeta \epsilon} + 2(C_{\ell}/\eta) \sqrt{(1 - \zeta)\epsilon}$ holds almost surely. Furthermore, if there is $C_{\mathcal{G}} > 0$ such that $\mathfrak{R}_{n, q}(\mathcal{G}) \leq C_{\mathcal{G}}/\sqrt{n}$ ², the convergence of $[(R(\hat{g}_{\text{PUBN}, \eta, \hat{\sigma}}) - R(g^*)) - (2C_{\ell} \sqrt{\zeta \epsilon} + 2(C_{\ell}/\eta) \sqrt{(1 - \zeta)\epsilon})]^+$ to 0 is in $\mathcal{O}_p(1/\sqrt{n_P} + 1/\sqrt{n_{bN}} + 1/\sqrt{n_U})$, where \mathcal{O}_p denotes the order in probability and $[\cdot]^+ = \max\{0, \cdot\}$. As for ϵ , knowing that $\hat{\sigma}$ is also estimated from data in practice³, apparently its value depends on both the estimation algorithm and the number of samples that are involved in the estimation process. For example, in our approach we applied nnPU with the logistic loss to obtain $\hat{\sigma}$, so the excess risk can be written as $\mathbb{E}_{\mathbf{x} \sim p(\mathbf{x})} \text{KL}(\sigma(\mathbf{x}) \parallel \hat{\sigma}(\mathbf{x}))$, where by abuse of notation $\text{KL}(p \parallel q) = p \ln(p/q) + (1 -$

² For instance, this holds for linear-in-parameter model class $\mathcal{F} = \{f(\mathbf{x}) = \mathbf{w}^T \phi(\mathbf{x}) \mid \|\mathbf{w}\| \leq C_w, \|\phi\|_{\infty} \leq C_{\phi}\}$, where C_w and C_{ϕ} are positive constants (Mohri et al., 2012).

³ These data, according to theorem 2, must be different from those used to evaluate $\hat{R}_{\text{PUBN}, \eta, \hat{\sigma}}(g)$. This condition is however violated in most of our experiments. See Supplementary Material D.3 for more discussion.

$p) \ln((1-p)/(1-q))$ denotes the KL divergence between two Bernoulli distributions with parameters respectively p and q . It is known that $\epsilon = \mathbb{E}_{\mathbf{x} \sim p(\mathbf{x})} [|\hat{\sigma}(\mathbf{x}) - \sigma(\mathbf{x})|^2] \leq (1/2) \mathbb{E}_{\mathbf{x} \sim p(\mathbf{x})} \text{KL}(\sigma(\mathbf{x}) \parallel \hat{\sigma}(\mathbf{x}))$ (Zhang, 2004). The excess risk itself can be decomposed into the sum of the estimation error and the approximation error. Kiryo et al. (2017) showed that under mild assumptions the estimation error part converges to zero when the sample size increases to infinity in nnPU learning. It is however impossible to get rid of the approximation error part which is fixed once we fix the function class \mathcal{G} . To circumvent this problem, we can either resort to kernel-based methods with universal kernels (Zhang, 2004) or simply enlarge the function class when we get more samples.

3.4. PU Learning Revisited

In PU learning scenarios, we only have P and U data and bN data are not available. Nevertheless, if we let y play the role of s and ignore all the terms related to bN data, our algorithm is naturally applicable to PU learning. Let us name the resulting algorithm PUBN\N, then

$$\hat{R}_{\text{PUBN}\backslash\text{N}, \eta, \hat{\sigma}}(g) = \pi \hat{R}_{\text{P}}^+(g) + \hat{R}_{y=-1, \eta, \hat{\sigma}}^-(g),$$

where $\hat{\sigma}$ is an estimate of $p(y = +1 \mid \mathbf{x})$ and

$$\begin{aligned} \hat{R}_{y=-1, \eta, \hat{\sigma}}^-(g) &= \mathbb{E}_{\mathbf{x} \sim p(\mathbf{x})} [\mathbb{1}_{\hat{\sigma}(\mathbf{x}) \leq \eta} \ell(-g(\mathbf{x})) (1 - \hat{\sigma}(\mathbf{x}))] \\ &+ \pi \mathbb{E}_{\mathbf{x} \sim p(\mathbf{x})} \left[\mathbb{1}_{\hat{\sigma}(\mathbf{x}) > \eta} \ell(-g(\mathbf{x})) \frac{1 - \hat{\sigma}(\mathbf{x})}{\hat{\sigma}(\mathbf{x})} \right]. \end{aligned}$$

PUBN\N can be viewed as a variant of the traditional two-step approach in PU learning which first identifies possible N data in U data and then perform ordinary PN classification to distinguish P data from the identified N data. However, being based on state-of-the-art nnPU learning, our method is more promising than other similar algorithms. Moreover, by explicitly considering the posterior $p(y = +1 \mid \mathbf{x})$, we attempt to correct the bias induced by the fact of only taking into account confident negative samples. The benefit of using an unbiased risk estimator is that the resulting algorithm is always statistically consistent, i.e., the estimation error converges in probability to zero as the number of samples grows to infinity.

4. Experiments

In this section, we experimentally investigate the proposed method and compare its performance against several baseline methods.

4.1. Basic Setup

We focus on training neural networks with stochastic optimization. For simplicity, in an experiment, $\hat{\sigma}$ and g always

use the same model and are trained for the same number of epochs. All models are learned using AMSGrad (Reddi et al., 2018) as the optimizer and the logistic loss as the surrogate loss unless otherwise specified. In all the experiments, an additional validation set, equally composed of P, U and bN data, is sampled for both hyperparameter tuning and choosing the model parameters with the lowest validation loss among those obtained after every epoch. Regarding the computation of the validation loss, we use the PU risk estimator (2) with the sigmoid loss for g and an empirical approximation of $\mathbb{E}_{\mathbf{x} \sim p(\mathbf{x})} [|\hat{\sigma}(\mathbf{x}) - \sigma(\mathbf{x})|^2] - \mathbb{E}_{\mathbf{x} \sim p(\mathbf{x})} [\sigma(\mathbf{x})^2]$ for $\hat{\sigma}$ (see Supplementary Material B).

4.2. Effectiveness of the Algorithm

We assess the performance of the proposed method on three benchmark datasets: MNIST, CIFAR-10 and 20 Newsgroups. Experimental details are given in Supplementary Material C. To recapitulate, for the three datasets we respectively use a 4-layer ConvNet, PreAct ResNet-18 (He et al., 2016) and a 3-layer fully connected neural network. On 20 Newsgroups text features are generated thanks to the use of ELMo word embedding (Peters et al., 2018). Since all the three datasets are originally designed for multiclass classification, we group different categories together to form a binary classification problem.

Baselines. When \mathcal{X}_{bN} is given, two baseline methods are considered. The first one is nnPNU adapted from (4). In the second method, named as PU \rightarrow PN, we train two binary classifiers: one is learned with nnPU while we regard s as the class label, and the other is learned from \mathcal{X}_{P} and \mathcal{X}_{bN} to separate P samples from bN samples. A sample is classified in the P class only if it is so classified by the two classifiers. When \mathcal{X}_{bN} is not available, nnPU is compared with the proposed PUBN\N.

Sampling bN Data. To sample \mathcal{X}_{bN} , we suppose that the bias of N data is caused by a latent prior probability change (Sugiyama & Storkey, 2007; Hu et al., 2018) in the N class. Let $z \in \mathcal{Z} = \{1, \dots, S\}$ be some latent variable which we call a latent category, where S is a constant. It is assumed

$$\begin{aligned} p(\mathbf{x} \mid z, y = -1) &= p(\mathbf{x} \mid z, y = -1, s = +1), \\ p(z \mid y = -1) &\neq p(z \mid y = -1, s = +1). \end{aligned}$$

In the experiments, the latent categories are the original class labels of the datasets. Concrete definitions of \mathcal{X}_{bN} with experimental results are summarized in Table 1.

Results. Overall, our proposed method consistently achieves the best or comparable performance in all the scenarios, including those of standard PU learning. Additionally, using bN data can effectively help improving the classification performance. However, the choice of algorithm is essential. Both nnPNU and the naive PU \rightarrow PN are

Table 1. Mean and standard deviation of misclassification rates over 10 trials for MNIST, CIFAR-10 and 20 Newsgroups under different choices of P class and bN data sampling strategies. For a same learning task, different methods are compared using the same 10 random samplings. Underlines denote that with the use of bN data the method leads to an improvement of performance according to the 5% t-test. Boldface indicates the best method in each task.

[†] Biased N data uniformly sampled from the indicated latent categories.

* Probabilities that a sample of \mathcal{X}_{bN} belongs to the latent categories [1, 3, 5, 7, 9] / [bird, cat, deer, dog, frog, horse] / [sci., soc., talk.] are [0.03, 0.15, 0.3, 0.02, 0.5] / [0.1, 0.02, 0.2, 0.08, 0.2, 0.4] / [0.1, 0.5, 0.4].

Dataset	P	biased N	ρ	nnPU/nnPNU	PUBN(\N)	PU→PN
MNIST	0, 2, 4, 6, 8	Not given	NA	5.76 ± 1.04	4.64 ± 0.62	NA
		1, 3, 5 [†]	0.3	5.33 ± 0.97	<u>4.05 ± 0.27</u>	<u>4.00 ± 0.30</u>
		9 > 5 > others [*]	0.2	4.60 ± 0.65	<u>3.91 ± 0.66</u>	<u>3.77 ± 0.31</u>
CIFAR-10	Airplane, automobile, ship, truck	Not given	NA	12.02 ± 0.65	10.70 ± 0.57	NA
		Cat, dog, horse [†]	0.3	10.25 ± 0.38	<u>9.71 ± 0.51</u>	10.37 ± 0.65
		Horse > deer = frog > others [*]	0.25	<u>9.98 ± 0.53</u>	<u>9.92 ± 0.42</u>	<u>10.17 ± 0.35</u>
CIFAR-10	Cat, deer, dog, horse	Not given	NA	23.78 ± 1.04	21.13 ± 0.90	NA
		Bird, frog [†]	0.2	22.00 ± 0.53	<u>18.83 ± 0.71</u>	<u>19.88 ± 0.62</u>
		Car, truck [†]	0.2	22.00 ± 0.74	<u>20.19 ± 1.06</u>	21.83 ± 1.36
20 Newsgroups	alt., comp., misc., rec.	Not given	NA	14.67 ± 0.87	13.30 ± 0.53	NA
		sci. [†]	0.21	14.69 ± 0.46	13.10 ± 0.90	13.58 ± 0.97
		talk. [†]	0.17	14.38 ± 0.74	<u>12.61 ± 0.75</u>	13.76 ± 0.66
		soc. > talk. > sci. [*]	0.1	14.41 ± 0.76	<u>12.18 ± 0.59</u>	12.92 ± 0.51

able to leverage bN data to enhance classification accuracy in only relatively few tasks. In the contrast, the proposed PUBN successfully reduce the misclassification error most of the time.

Clearly, the performance gain that we can benefit from the availability of bN data is case-dependent. On CIFAR-10, the greatest improvement is achieved when we regard mammals (i.e. cat, deer, dog and horse) as P class and drawn samples from latent categories bird and frog as labeled negative data. This is not surprising because birds and frogs are more similar to mammals than vehicles, which makes the classification harder specifically for samples from these two latent categories. By explicitly labeling these samples as N data, we allow the classifier to make better predictions for these difficult samples.

4.3. Illustration on How the Presence of bN Data Help

Through experiments we have demonstrated that the presence of bN data effectively helps learning a better classifier. Here we would like to provide some intuition for the reason behind this. Let us consider the MNIST learning task where \mathcal{X}_{bN} is uniformly sampled from the latent categories 1, 3 and 5. We project the representations learned by the classifier (i.e., the activation values of the last hidden layer of the neural network) into a 2D plane using PCA for both nnPU

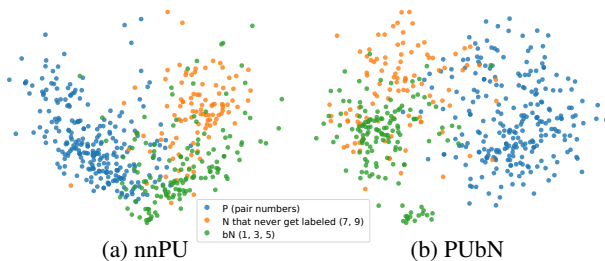


Figure 1. PCA embeddings of the representations learned by the nnPU and PUBN classifiers for 500 samples from the test set in the MNIST learning task where \mathcal{X}_{bN} is uniformly sampled from latent categories 1, 3 and 5.

and PUBN algorithms, as shown in Figure 1.

For both nnPU and PUBN classifiers, the first two principal components account around 90% of variance. We can therefore presume that the figure depicts fairly well the learned representations. Thanks to the use of bN data, in the high-level feature space 1, 3, 5 and P data are further pushed away when we employ the proposed PUBN learning algorithm, and we are always able to separate 7, 9 from P to some extent. This explains the better performance which is achieved by PUBN learning and the benefit of incorporating bN data into the learning process.

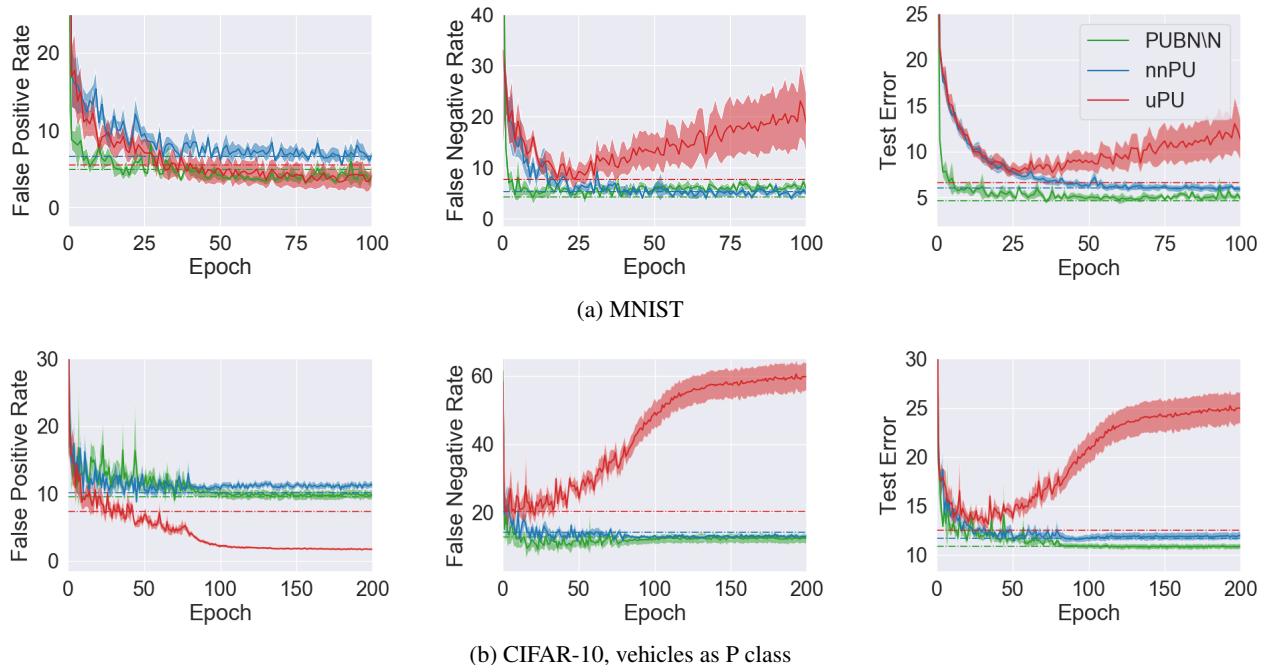


Figure 2. Comparison of uPU, nnPU and PUBN\N over two of the four PU learning tasks. For each task, means and standard deviations are computed based on the same 10 random samplings. Dashed lines indicate the corresponding values of the final classifiers (recall that at the end we select the model with the lowest validation loss out of all epochs).

4.4. Why Does PUBN\N Outperform nnPU ?

Our method, specifically designed for PUBN learning, naturally outperforms other baseline methods in this problem. Nonetheless, Table 1 equally shows that the proposed method when applied to PU learning, achieves significantly better performance than the state-of-the-art nnPU algorithm. Here we numerically investigate the reason behind this phenomenon with help of the first two PU tasks of the table.

Besides nnPU and PUBN\N, we compare with unbiased PU (uPU) learning (2). Both uPU and nnPU are learned with the sigmoid loss, learning rate 10^{-3} for MNIST and initial learning rate 10^{-4} for CIFAR-10, as uPU learning is unstable with the logistic loss. The other parts of the experiments remain unchanged. On the test sets we compute the false positive rates, false negative rates and misclassification errors for the three methods and plot them in Figure 2. We first notice that PUBN\N still outperforms nnPU trained with the sigmoid loss. In fact, the final performance of the nnPU classifier does not change much when we replace the logistic loss with the sigmoid loss.

In (Kiryu et al., 2017), the authors observed that uPU overfits training data with the risk going to negative. In other words, a large portion of U samples are classified as N. This is confirmed in our experiments by an increase of false negative rate and decrease of false positive rate. nnPU remedies the problem by introducing the non-negative risk estimator

(3). While the non-negative correction successfully prevents false negative rate from going up, it also causes more N samples to be classified as P compared to uPU. However, since the gain in terms of false negative rate is enormous, at the end nnPU achieves a lower misclassification error. By further identifying possible N samples after nnPU learning, we expect that our algorithm can yield lower false positive rate than nnPU without misclassifying too many P samples as N as in the case of uPU. Figure 2 suggests that this is effectively the case. In particular, we observe that on MNIST, our method achieves the same false positive rate as uPU whereas its false negative rate is comparable to nnPU.

5. Conclusion

This paper studies the PUBN classification problem, where a binary classifier is trained on P, U and bN data. The proposed method is a two-step approach inspired from both PU learning and importance weighting. The key idea is to attribute appropriate weights to each example for evaluation of the classification risk using the three sets of data. We theoretically established an estimation error bound for the proposed risk estimator and experimentally showed that our approach successfully leveraged bN data to improve the classification performance on several real-world datasets. A variant of our algorithm was able to achieve state-of-the-art results in PU learning.

Acknowledgements

MS was supported by JST CREST Grant Number JP-MJCR1403.

References

- Bartlett, P. L. and Mendelson, S. Rademacher and gaussian complexities: Risk bounds and structural results. *JMLR*, 3(Nov):463–482, 2002.
- Belkin, M., Niyogi, P., and Sindhvani, V. Manifold regularization: A geometric framework for learning from labeled and unlabeled examples. *JMLR*, 7(Nov):2399–2434, 2006.
- Chapelle, O., Schölkopf, B., and Zien, A. *Semi-Supervised Learning*. The MIT Press, 1st edition, 2010.
- 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*, pp. 703–711, 2014.
- du Plessis, M. C., Niu, G., and Sugiyama, M. Class-prior estimation for learning from positive and unlabeled data. *Maching Learning*, 106(4):463–492, April 2017. ISSN 0885-6125.
- Elkan, C. and Noto, K. Learning classifiers from only positive and unlabeled data. In *KDD*, 2008.
- Fei, G. and Liu, B. Social media text classification under negative covariate shift. In *Proc. of EMNLP*, pp. 2347–2356, 2015.
- Fung, G. P. C., Yu, J. X., Lu, H., and Yu, P. S. Text classification without negative examples revisited. *TKDE*, 18(1): 6–20, 2006.
- Golowich, N., Rakhlin, A., and Shamir, O. Size-independent sample complexity of neural networks. In *COLT*, pp. 297–299, 2018.
- Grandvalet, Y. and Bengio, Y. Semi-supervised learning by entropy minimization. In *NeurIPS*, pp. 529–536, 2005.
- He, K., Zhang, X., Ren, S., and Sun, J. Identity mappings in deep residual networks. In *ECCV*, pp. 630–645. Springer, 2016.
- Heckman, J. J. Sample selection bias as a specification error. *Econometrica*, 47(1):153–161, 1979.
- Hido, S., Tsuboi, Y., Kashima, H., Sugiyama, M., and Kanamori, T. Inlier-based outlier detection via direct density ratio estimation. In *ICDM*, pp. 223–232. IEEE, 2008.
- Hu, W., Niu, G., Sato, I., and Sugiyama, M. Does distributionally robust supervised learning give robust classifiers? In *ICML*, pp. 2034–2042, 2018.
- Huang, J., Gretton, A., Borgwardt, K. M., Schölkopf, B., and Smola, A. J. Correcting sample selection bias by unlabeled data. In *NeurIPS*, pp. 601–608, 2007.
- Ishida, T., Niu, G., and Sugiyama, M. Binary classification from positive-confidence data. In *NeurIPS*, pp. 5921–5932, 2018.
- Jain, S., White, M., and Radivojac, P. Estimating the class prior and posterior from noisy positives and unlabeled data. In *NeurIPS*, pp. 2693–2701, 2016.
- Kanamori, T., Hido, S., and Sugiyama, M. A least-squares approach to direct importance estimation. *JMLR*, 10: 1391–1445, 2009.
- Kiryu, R., Niu, G., du Plessis, M. C., and Sugiyama, M. Positive-unlabeled learning with non-negative risk estimator. In *NeurIPS*, pp. 1675–1685, 2017.
- Laine, S. and Aila, T. Temporal ensembling for semi-supervised learning. In *ICLR*, 2017.
- Li, W., Guo, Q., and Elkan, C. A positive and unlabeled learning algorithm for one-class classification of remote-sensing data. *TGRS*, 49(2):717–725, 2011.
- Li, X.-L., Liu, B., and Ng, S.-K. Negative training data can be harmful to text classification. In *Proc. of EMNLP*, pp. 218–228, 2010.
- Liu, B., Lee, W. S., Yu, P. S., and Li, X. Partially supervised classification of text documents. In *ICML*, volume 2, pp. 387–394, 2002.
- Liu, B., Dai, Y., Li, X., Lee, W. S., and Yu, P. S. Building text classifiers using positive and unlabeled examples. In *ICDM*, pp. 179–186. IEEE, 2003.
- Miyato, T., Maeda, S.-i., Koyama, M., Nakae, K., and Ishii, S. Distributional smoothing with virtual adversarial training. In *ICLR*, 2016.
- Mohri, M., Rostamizadeh, A., and Talwalkar, A. *Foundations of machine learning*. The MIT press, 2012.
- Mordelet, F. and Vert, J.-P. A bagging svm to learn from positive and unlabeled examples. *Pattern Recognition Letters*, 37:201–209, 2014.

- Nguyen, M. N., Li, X.-L., and Ng, S.-K. Positive unlabeled learning for time series classification. In *IJCAI*, volume 11, pp. 1421–1426, 2011.
- Oliver, A., Odena, A., Raffel, C. A., Cubuk, E. D., and Goodfellow, I. Realistic evaluation of deep semi-supervised learning algorithms. In *NeurIPS*, pp. 3239–3250, 2018.
- Peters, M. E., Neumann, M., Iyyer, M., Gardner, M., Clark, C., Lee, K., and Zettlemoyer, L. Deep contextualized word representations. In *Proc. of NAACL*, 2018.
- Quionero-Candela, J., Sugiyama, M., Schwaighofer, A., and Lawrence, N. D. *Dataset Shift in Machine Learning*. The MIT Press, 2009.
- Ramaswamy, H., Scott, C., and Tewari, A. Mixture proportion estimation via kernel embeddings of distributions. In *ICML*, pp. 2052–2060, 2016.
- Reddi, S. J., Kale, S., and Kumar, S. On the convergence of adam and beyond. In *ICLR*, 2018.
- Rücklé, A., Eger, S., Peyrard, M., and Gurevych, I. Concatenated power mean word embeddings as universal cross-lingual sentence representations. *arXiv preprint arXiv:1803.01400*, 2018.
- Sakai, T., du Plessis, M. C. d., Niu, G., and Sugiyama, M. Semi-supervised classification based on classification from positive and unlabeled data. In *ICML*, volume 70, pp. 2998–3006, 2017.
- Scott, C. and Blanchard, G. Novelty detection: Unlabeled data definitely help. In *AISTATS*, pp. 464–471, 2009.
- Shalev-Shwartz, S. and Ben-David, S. *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press, 2014.
- Shi, H., Pan, S., Yang, J., and Gong, C. Positive and unlabeled learning via loss decomposition and centroid estimation. In *IJCAI*, pp. 2689–2695, 2018.
- Shimodaira, H. Improving predictive inference under covariate shift by weighting the log-likelihood function. *JMLR*, 90(2):227–244, 2000.
- Sugiyama, M. and Kawanabe, M. *Machine Learning in Non-Stationary Environments: Introduction to Covariate Shift Adaptation*. The MIT Press, Cambridge, Massachusetts, USA, 2012.
- Sugiyama, M. and Storkey, A. J. Mixture regression for covariate shift. In *NeurIPS*, pp. 1337–1344, 2007.
- Sugiyama, M., Nakajima, S., Kashima, H., Buenau, P. V., and Kawanabe, M. Direct importance estimation with model selection and its application to covariate shift adaptation. In *NeurIPS*, pp. 1433–1440, 2008.
- Ward, G. A., Hastie, T. J., Barry, S. T., Elith, J., and Leathwick, J. R. Presence-only data and the em algorithm. *Biometrics*, 65 2:554–63, 2009.
- Zadrozny, B. Learning and evaluating classifiers under sample selection bias. In *ICML*, pp. 903–910, 2004.
- Zhang, T. Statistical behavior and consistency of classification methods based on convex risk minimization. *Annals of Statistics*, pp. 56–85, 2004.
- Zuluaga, M., Hush, D., J F Delgado Leyton, E., Hernandez Hoyos, M., and Orkisz, M. Learning from only positive and unlabeled data to detect lesions in vascular ct images. In *MICCAI*, volume LNCS 6893, pp. 9–16, 2011.