



Principled analytic classifier for positive-unlabeled learning via weighted integral probability metric

Yongchan Kwon¹ · Wonyoung Kim¹ · Masashi Sugiyama^{2,3} · Myunghee Cho Paik¹

Received: 3 May 2019 / Revised: 7 July 2019 / Accepted: 6 September 2019 / Published online: 4 November 2019
© The Author(s), under exclusive licence to Springer Science+Business Media LLC, part of Springer Nature 2019

Abstract

We consider the problem of learning a binary classifier from only positive and unlabeled observations (called PU learning). Recent studies in PU learning have shown superior performance theoretically and empirically. However, most existing algorithms may not be suitable for large-scale datasets because they face repeated computations of a large Gram matrix or require massive hyperparameter optimization. In this paper, we propose a computationally efficient and theoretically grounded PU learning algorithm. The proposed PU learning algorithm produces a closed-form classifier when the hypothesis space is a closed ball in reproducing kernel Hilbert space. In addition, we establish upper bounds of the estimation error and the excess risk. The obtained estimation error bound is sharper than existing results and the derived excess risk bound has an explicit form, which vanishes as sample sizes increase. Finally, we conduct extensive numerical experiments using both synthetic and real datasets, demonstrating improved accuracy, scalability, and robustness of the proposed algorithm.

Keywords Positive and unlabeled learning · Integral probability metric · Excess risk bound · Approximation error · Reproducing kernel Hilbert space

Editors: Kee-Eung Kim and Jun Zhu.

✉ Myunghee Cho Paik
myungheechopaik@snu.ac.kr

Yongchan Kwon
ykwon0407@snu.ac.kr

Wonyoung Kim
eraser347@snu.ac.kr

Masashi Sugiyama
sugi@k.u-tokyo.ac.jp

¹ Department of Statistics, Seoul National University, Seoul, South Korea

² Center for Advanced Intelligence Project, RIKEN, Tokyo, Japan

³ Department of Complexity Science and Engineering, Graduate School of Frontier Sciences, The University of Tokyo, Tokyo, Japan

1 Introduction

Supervised binary classification assumes that all the training data are labeled as either being positive or negative. However, in many practical scenarios, collecting a large number of labeled samples from the two categories is often costly, difficult, or not even possible. In contrast, unlabeled data are relatively cheap and abundant. As a consequence, semi-supervised learning is used for partially labeled data (Chapelle et al. 2006). In this paper, as a special case of semi-supervised learning, we consider Positive-Unlabeled (PU) learning, the problem of building a binary classifier from only positive and unlabeled samples (Denis et al. 2005; Li and Liu 2005). PU learning provides a powerful framework when negative labels are impossible or very expensive to obtain, and thus has frequently appeared in many real-world applications. Examples include document classification (Elkan and Noto 2008; Xiao et al. 2011), image classification (Zuluaga et al. 2011; Gong et al. 2018), gene identification (Yang et al. 2012, 2014), and novelty detection (Blanchard et al. 2010; Zhang et al. 2017).

Several PU learning algorithms have been developed over the last 2 decades. Liu et al. (2002) and Li and Liu (2003) considered a two-step learning scheme: in Step 1, assigning negative labels to some unlabeled observations believed to be negative, and in Step 2, learning a binary classifier with existing positive samples and the negatively labeled samples from Step 1. Liu et al. (2003) pointed out that the two-step learning scheme is based on heuristics, and suggested fitting a biased support vector machine by regarding all the unlabeled observations as being negative.

Scott and Blanchard (2009) and Blanchard et al. (2010) suggested a modification of supervised Neyman–Pearson classification, whose goal is to find a classifier minimizing the false positive rate keeping the false negative rate low. To circumvent the problem of lack of negative samples, they tried to build a classifier minimizing the marginal probability of being classified as positive while keeping the false negative rate low. Solving the empirical version of this constrained optimization problem is challenging, but the authors did not present an explicit algorithm.

Recently, many PU learning algorithms based on the empirical risk minimization principle have been studied. Du Plessis et al. (2014) proposed the use of the ramp loss and provided an algorithm that requires solving a non-convex optimization problem. Du Plessis et al. (2015) formulated a convex optimization problem by using the logistic loss or double hinge loss. However, all the aforementioned approaches involve solving a non-linear programming problem. This causes massive computational burdens for calculating the large Gram matrix when the sample size is large. Kiryo et al. (2017) suggested a stochastic algorithm for large-scale datasets with a non-negative risk estimator. However, to execute the algorithm, several hyperparameters are required, and choosing the optimal hyperparameter may demand substantial trials of running the algorithm (Oh et al. 2018), causing heavy computation costs.

In supervised binary classification, Sriperumbudur et al. (2012) proposed a computationally efficient algorithm building a closed-form binary discriminant function. The authors showed that their function estimator obtained by evaluating the negative of the empirical integral probability metric (IPM) is the minimizer of the empirical risk using the specific loss defined in Sect. 3.1. They further showed that a closed form can be derived as the result of restricting a hypothesis space to a closed unit ball in reproducing kernel Hilbert space (RKHS).

In this paper, capitalizing on the properties shown in the supervised learning method by Sriperumbudur et al. (2012), we extend it to PU learning settings. In addition, we derive new theoretical results on excess risk bounds. We first define a weighted version of IPM between

two probability measures and call it the weighted integral probability metric (WIPM). We show that computing the negative of WIPM between the unlabeled data distribution and the positive data distribution is equivalent to minimizing the hinge risk. Based on this finding, we propose a binary discriminant function estimator that computes the negative of the empirical WIPM, and then derive associated upper bounds of the estimation error and the excess risk. Under a mild condition, our obtained upper bounds are shown to be sharper than the existing ones because of using Talagrand's inequality over McDiarmid's inequality (Kiryo et al. 2017). Moreover, we pay special attention to the case where the hypothesis space is a closed ball in RKHS and propose a closed-form classifier. We show that the associated excess risk bound has an explicit form that converges to zero as the sample sizes increase. To the best of our knowledge, this is the first result to explicitly show the excess risk bound in PU learning.

As a summary, our main contributions are:

- We formally define WIPM and establish a link with the infimum of the hinge risk (Theorem 1). We derive an estimation error bound and show that it is sharper than existing results (Theorem 2 and Proposition 1).
- The proposed algorithm produces a closed-form classifier when the underlying hypothesis space is a closed ball in RKHS (Proposition 2). Furthermore, we obtain a novel excess risk bound that converges to zero as sample sizes increase (Theorem 3).
- Numerical experiments using both synthetic and real datasets show that our method is comparable to or better than existing PU learning algorithms in terms of accuracy, scalability, and robustness in the case of unknown class-priors.

2 Preliminaries

In this section, we describe the L -risk for binary classification and present its PU representation. We briefly review several PU learning algorithms based on the L -risk minimization principle. We first introduce problem settings and notations.

2.1 Problem settings of PU learning

Let X and Y be random variables for input data and class labels, respectively, whose range is the product space $\mathcal{X} \times \{\pm 1\} \subseteq \mathbb{R}^d \times \{\pm 1\}$. The d is a positive integer. We denote the joint distribution of (X, Y) by $P_{X,Y}$ and the marginal distribution of X by P_X . The distributions of positive and negative data are defined by conditional distributions, $P_{X|Y=1}$ and $P_{X|Y=-1}$, respectively. Let $\pi_+ := P_{X,Y}(Y = 1)$ be the marginal probability of being positive and set $\pi_- = 1 - \pi_+$. We follow the *two samples of data* scheme (Ward et al. 2009; Niu et al. 2016). That is, let $\mathcal{X}_p = \{x_i^p\}_{i=1}^{n_p}$ and $\mathcal{X}_u = \{x_i^u\}_{i=1}^{n_u}$ be observed sets of independently identically distributed samples from the positive data distribution $P_{X|Y=1}$ and the marginal distribution P_X , respectively. Here, the n_p and n_u are the number of positive and unlabeled data points, respectively. Note that the unlabeled data distribution is the marginal distribution.

Let \mathcal{U} be a class of real-valued measurable functions defined on \mathcal{X} . A function $f \in \mathcal{U}$, often called a hypothesis, can be understood as a binary discriminant function and we classify an input x with the sign of a discriminant function, $\text{sign}(f(x))$. Define $\mathcal{M} = \{f : \mathcal{X} \rightarrow \mathbb{R} \mid \|f\|_\infty \leq 1\} \subseteq \mathcal{U}$, where $\|f\|_\infty = \sup_{x \in \mathcal{X}} |f(x)|$ is the supremum norm. We restrict our attention to a class $\mathcal{F} \subseteq \mathcal{M}$ and call \mathcal{F} a hypothesis space. Throughout this paper, we assume that the hypothesis space is symmetric, *i.e.*, $f \in \mathcal{F}$ implies $-f \in \mathcal{F}$. In PU learning,

the main goal is to construct a classifier $\text{sign}(f(x))$ only from the positive dataset \mathcal{X}_p and the unlabeled dataset \mathcal{X}_u with $f \in \mathcal{F}$.

In this paper, the quantity π_+ , often called the class-prior, is assumed to be known as in the literature (Kiryo et al. 2017; Kato et al. 2019) to focus on theoretical and practical benefits of our proposed algorithm. We examine the performance when π_+ is unknown in Experiment 3 of Sects. 6.1 and 6.2.

2.2 L-risk minimization in PU learning

In supervised binary classification, the L -risk is defined by

$$\begin{aligned}
 R_L(f) &:= \int_{\mathcal{X} \times \{\pm 1\}} L(y, f(x)) dP_{X,Y}(x, y) \\
 &= \pi_+ \int_{\mathcal{X}} L(1, f(x)) dP_{X|Y=1}(x) + \pi_- \int_{\mathcal{X}} L(-1, f(x)) dP_{X|Y=-1}(x), \quad (1)
 \end{aligned}$$

for a loss function $L : \{\pm 1\} \times \mathbb{R} \rightarrow \mathbb{R}$ (Steinwart and Christmann 2008, Section 2.1). We denote the margin-based loss function by $\ell(yt) := L(y, t)$ if a loss function $L(y, t)$ can be represented as a function of margin yt , the product of a label y and a score t for all possible $y \in \{\pm 1\}$ and $t \in \mathbb{R}$.

Under the PU learning framework, however, the right-hand side of Eq. (1) cannot be directly estimated due to lack of negatively labeled observations. To circumvent this problem, many studies in the field of PU learning exploited the relationship $P_X = \pi_+ P_{X|Y=1} + \pi_- P_{X|Y=-1}$ and replaced $P_{X|Y=-1}$ in Eq. (1) with $(P_X - \pi_+ P_{X|Y=1})/\pi_-$ (Du Plessis et al. 2014; Sakai et al. 2017). That is, the L -risk can be alternatively expressed as:

$$R_L(f) = \int_{\mathcal{X}} L(-1, f(x)) dP_X(x) + \pi_+ \int_{\mathcal{X}} L(1, f(x)) - L(-1, f(x)) dP_{X|Y=1}(x). \quad (2)$$

Now the right-hand side of Eq. (2) can be empirically estimated by the positive dataset \mathcal{X}_p and the unlabeled dataset \mathcal{X}_u . However, the L -risk $R_L(f)$ is not convex with respect to f in general, and minimizing an empirical estimator for $R_L(f)$ is often formulated as a complicated non-convex optimization problem.

There have been several approaches to resolving the computational difficulty by modifying loss functions. Du Plessis et al. (2014) proposed to use non-convex loss functions satisfying the symmetric condition, $L(1, f(x)) + L(-1, f(x)) = 1$. They proposed to optimize the empirical risk based on the ramp loss $\ell_{\text{ramp}}(yt) = 0.5 \times \max(0, \min(2, 1 - yt))$ via the concave-convex procedure (Collobert et al. 2006). Du Plessis et al. (2015) converted the problem to convex optimization through the linear-odd condition, $L(1, f(x)) - L(-1, f(x)) = -f(x)$. They showed that the logistic loss $\ell_{\log}(yt) = \log(1 + \exp(-yt))$ and the double hinge loss $\ell_{\text{dh}}(yt) = \max(0, \max(-yt, (1 - yt)/2))$ satisfy the linear-odd condition. However, all the aforementioned methods utilized a weighted sum of $n_p + n_u$ predefined basis functions as a binary discriminant function, which triggered calculating the $(n_p + n_u) \times (n_p + n_u)$ Gram matrix. Hence, executing algorithms is not scalable and can be intractable when n_p and n_u are large (Sansone et al. 2019). Our first goal in this paper is to overcome this computational problem by providing a computationally efficient method.

3 Weighted integral probability metric and L -risk

In this section, we formally define WIPM, a key tool for constructing the proposed algorithm, and build a link with the L -risk in Theorem 1 below. Based on the link, we propose a new binary discriminant function estimator and present its theoretical properties in Theorem 2. We first introduce the earlier work by Sriperumbudur et al. (2012) that provided a closed-form classifier in supervised binary classification.

3.1 Relation between IPM and L -risk in supervised binary classification

Müller (1997) introduced an IPM for any two probability measures P and Q defined on \mathcal{X} and a class \mathcal{F} of bounded measurable functions, given by

$$\text{IPM}(P, Q; \mathcal{F}) := \sup_{f \in \mathcal{F}} \left| \int_{\mathcal{X}} f(x) dP(x) - \int_{\mathcal{X}} f(x) dQ(x) \right|.$$

IPM has been studied as either a metric between two probability measures (Sriperumbudur et al. 2010a; Arjovsky et al. 2017; Tolstikhin et al. 2018) or a hypothesis testing tool (Gretton et al. 2012).

Under the supervised binary classification setting, Sriperumbudur et al. (2012) showed that calculating IPM between $P_{X|Y=1}$ and $P_{X|Y=-1}$ is negatively related to minimizing the risk with a loss function, *i.e.*, $\text{IPM}(P_{X|Y=1}, P_{X|Y=-1}; \mathcal{F}) = -\inf_{f \in \mathcal{F}} R_{L_c}(f)$, where $L_c(1, t) = -t/\pi_+$ and $L_c(-1, t) = t/\pi_-$ for all $t \in \mathbb{R}$. They further showed that a discriminant function minimizing the L_c -risk can be obtained analytically when \mathcal{F} is a closed unit ball in RKHS. This result cannot be directly extended to PU learning due to absence of negatively labeled observations. In the next subsection, we define a generalized version of IPM and extend the previous results for supervised binary classification to PU learning.

3.2 Extension to WIPM and L -risk in PU learning

Let \mathcal{F} be a given class of bounded measurable functions and let $\tilde{w} : \mathcal{X} \rightarrow \mathbb{R}$ be a weight function such that $\|\tilde{w}\|_{\infty} < \infty$. We define WIPM¹ between two probability measures P and Q with a function class \mathcal{F} and a weight function \tilde{w} by

$$\text{WIPM}(P, Q; \tilde{w}, \mathcal{F}) := \sup_{f \in \mathcal{F}} \left| \int_{\mathcal{X}} f(x) dP(x) - \int_{\mathcal{X}} \tilde{w}(x) f(x) dQ(x) \right|. \quad (3)$$

Note that WIPM reduces to IPM if $\tilde{w}(x) = 1$ for all $x \in \mathcal{X}$. Other special cases of Eq. (3) have been discussed in many applications. In the covariate shift problem, Huang et al. (2007) and Gretton et al. (2009) proposed to minimize WIPM with respect to \tilde{w} when \mathcal{F} is the unit ball in RKHS and P, Q are empirical distributions of test and training data, respectively. In unsupervised domain adaptation, Yan et al. (2017) regarded P, Q as empirical distributions of target and source data, respectively, where in this case, \tilde{w} is a ratio of two class-prior distributions.

¹ Although WIPM is not a metric in general, we keep saying the name WIPM to emphasize that it is a weighted version of IPM.

We pay special attention to the case where $\tilde{w}(x)$ is constant, $w \in \mathbb{R}$, for every input value and denote WIPM by $\text{WIPM}(P, Q; w, \mathcal{F})$,

$$\text{WIPM}(P, Q; w, \mathcal{F}) := \sup_{f \in \mathcal{F}} \left| \int_{\mathcal{X}} f(x) dP(x) - w \int_{\mathcal{X}} f(x) dQ(x) \right|.$$

In the following theorem, we establish a link between $\text{WIPM}(P_X, P_{X|Y=1}; 2\pi_+, \mathcal{F})$ and the infimum of the ℓ_h -risk over \mathcal{F} for the hinge loss $\ell_h(yt) = \max(0, 1 - yt)$.

Theorem 1 (Relationship between ℓ_h -risk and WIPM) *Let \mathcal{F} be a symmetric hypothesis space in \mathcal{M} and $\ell_h(yt) = \max(0, 1 - yt)$ be the hinge loss. Then, we have*

$$\inf_{f \in \mathcal{F}} R_{\ell_h}(f) = 1 - \text{WIPM}(P_X, P_{X|Y=1}; 2\pi_+, \mathcal{F}).$$

Moreover, if $g_{\mathcal{F}}$ satisfies

$$\text{WIPM}(P_X, P_{X|Y=1}; 2\pi_+, \mathcal{F}) = \int_{\mathcal{X}} g_{\mathcal{F}}(x) dP_X(x) - 2\pi_+ \int_{\mathcal{X}} g_{\mathcal{F}}(x) dP_{X|Y=1}(x),$$

then $\inf_{f \in \mathcal{F}} R_{\ell_h}(f) = R_{\ell_h}(-g_{\mathcal{F}})$.

Theorem 1 shows that the infimum of the ℓ_h -risk over a hypothesis space \mathcal{F} equals the negative WIPM between the unlabeled data distribution P_X and the positive data distribution $P_{X|Y=1}$ with the same hypothesis space \mathcal{F} and the weight $2\pi_+$ up to addition by constant. Furthermore, by negating the WIPM optimizer $g_{\mathcal{F}}$, we obtain the minimizer of the ℓ_h -risk over the hypothesis space \mathcal{F} . Here, we define a WIPM optimizer $g_{\mathcal{F}}$ as a function that attains the supremum, i.e., $\text{WIPM}(P_X, P_{X|Y=1}; 2\pi_+, \mathcal{F}) = \int_{\mathcal{X}} g_{\mathcal{F}}(x) dP_X(x) - 2\pi_+ \int_{\mathcal{X}} g_{\mathcal{F}}(x) dP_{X|Y=1}(x)$ and we set $f_{\mathcal{F}} = -g_{\mathcal{F}}$ for later notational convenience. Sriperumbudur et al. (2012) derived a similar result to Theorem 1 by showing $\text{IPM}(P_{X|Y=1}, P_{X|Y=-1}; \mathcal{F}) = -\inf_{f \in \mathcal{F}} R_{L_c}(f)$ in supervised binary classification. However, as we mentioned in Sect. 3.1, their method is only applicable to supervised binary classification settings.

3.3 Theoretical properties of empirical WIPM optimizer

We denote the empirical distributions of $P_{X|Y=1}$ and P_X by $P_{X|Y=1, n_p}$ and P_{X, n_u} , respectively. Let $P_{X|Y=1, n_p} = n_p^{-1} \sum_{i=1}^{n_p} \delta_{x_i^p}$ and $P_{X, n_u} = n_u^{-1} \sum_{i=1}^{n_u} \delta_{x_i^u}$, where $\delta(\cdot)$ defined on \mathcal{X} is the Dirac delta function and $\delta_x(\cdot) := \delta(\cdot - x)$ for $x \in \mathcal{X}$. The empirical Rademacher complexity of \mathcal{F} given a set $S = \{z_1, \dots, z_m\}$ is defined by $\mathfrak{R}_S(\mathcal{F}) := \mathbb{E}_{\sigma} \left(\frac{1}{m} \sup_{f \in \mathcal{F}} \left| \sum_{i=1}^m \sigma_i f(z_i) \right| \right)$. Here, $\{\sigma_i\}_{i=1}^m$ is a set of independent Rademacher random variables taking 1 or -1 with probability 0.5 each and $\mathbb{E}_{\sigma}(\cdot)$ is the expectation operator over the Rademacher random variables (Bartlett and Mendelson 2002). Denote a maximum by $a \vee b := \max(a, b)$, a minimum by $a \wedge b := \min(a, b)$. For a probability measure Q defined on \mathcal{X} , denote the expectation of a discriminant function f by $\mathbb{E}_Q(f) := \int_{\mathcal{X}} f(x) dQ(x)$ and the variance by $\text{Var}_Q(f) := \mathbb{E}_Q(f^2) - (\mathbb{E}_Q(f))^2$.

The empirical estimator for $\text{WIPM}(P_X, P_{X|Y=1}; w, \mathcal{F})$ is given by plugging in the empirical distributions,

$$\text{WIPM}(P_{X, n_u}, P_{X|Y=1, n_p}; w, \mathcal{F}) = \sup_{f \in \mathcal{F}} \left| \frac{1}{n_u} \sum_{i=1}^{n_u} f(x_i^u) - \frac{w}{n_p} \sum_{i=1}^{n_p} f(x_i^p) \right|,$$

and we define an empirical WIPM optimizer $\hat{g}_{\mathcal{F}} \in \mathcal{F}$ that satisfies the following equation,

$$\text{WIPM}(P_{X, n_u}, P_{X|Y=1, n_p}; w, \mathcal{F}) = \frac{1}{n_u} \sum_{i=1}^{n_u} \hat{g}_{\mathcal{F}}(x_i^u) - \frac{w}{n_p} \sum_{i=1}^{n_p} \hat{g}_{\mathcal{F}}(x_i^p). \tag{4}$$

We set $\hat{f}_{\mathcal{F}} = -\hat{g}_{\mathcal{F}}$ for notational convenience as in Sect. 3.2.

We analyze the estimation error $R_{\ell_h}(\hat{f}_{\mathcal{F}}) - \inf_{f \in \mathcal{F}} R_{\ell_h}(f)$ in the following theorem. To begin, let $\chi_{n_p, n_u}^{(1)}(w) = w/\sqrt{n_p} + 1/\sqrt{n_u}$ and $\chi_{n_p, n_u}^{(2)}(w) = 2(w/n_p + 1/n_u)$.

Theorem 2 (Estimation error bound for general function space) *Let $\hat{g}_{\mathcal{F}}$ be an empirical WIPM optimizer defined in Eq. (4) and set $\hat{f}_{\mathcal{F}} = -\hat{g}_{\mathcal{F}}$. Let \mathcal{F} be a symmetric hypothesis space such that $\|f\|_{\infty} \leq \nu \leq 1$, $\text{Var}_{P_{X|Y=1}}(f) \leq \sigma_{X|Y=1}^2$, and $\text{Var}_{P_X}(f) \leq \sigma_X^2$. Denote $\rho^2 = \sigma_{X|Y=1}^2 \vee \sigma_X^2$. Then, for all $\alpha, \tau > 0$, the following holds with probability at least $1 - e^{-\tau}$,*

$$\begin{aligned} R_{\ell_h}(\hat{f}_{\mathcal{F}}) - \inf_{f \in \mathcal{F}} R_{\ell_h}(f) &\leq C_{\alpha}(\mathbb{E}_{P_X^{n_u}}(\mathfrak{R}_{\mathcal{X}_u}(\mathcal{F})) + 2\pi_+ \mathbb{E}_{P_{X|Y=1}^{n_p}}(\mathfrak{R}_{\mathcal{X}_p}(\mathcal{F}))) \\ &\quad + C_{\tau, \rho^2}^{(1)} \chi_{n_p, n_u}^{(1)}(2\pi_+) + C_{\tau, \nu, \alpha}^{(2)} \chi_{n_p, n_u}^{(2)}(2\pi_+), \end{aligned} \tag{5}$$

where $C_{\alpha} = 4(1 + \alpha)$, $C_{\tau, \rho^2}^{(1)} = 2\sqrt{2\tau\rho^2}$, $C_{\tau, \nu, \alpha}^{(2)} = 2\tau\nu(\frac{2}{3} + \frac{1}{\alpha})$.

Due to Talagrand’s inequality, Theorem 2 provides a sharper bound than the existing result based on McDiarmid’s inequality. Specifically, Kiryo et al. (2017, Theorem 4) utilized McDiarmid’s inequality and showed that for $\tau > 0$ and some $\Delta > 0$ the following holds with probability at least $1 - e^{-\tau}$,

$$\begin{aligned} R_{\ell_h}(\hat{f}) - \inf_{f \in \mathcal{F}} R_{\ell_h}(f) &\leq 8(\mathbb{E}_{P_X^{n_u}}(\mathfrak{R}_{\mathcal{X}_u}(\mathcal{F})) + 2\pi_+ \mathbb{E}_{P_{X|Y=1}^{n_p}}(\mathfrak{R}_{\mathcal{X}_p}(\mathcal{F}))) \\ &\quad + \chi_{n_p, n_u}^{(1)}(2\pi_+)(1 + \nu)\sqrt{2\tau} + \Delta. \end{aligned} \tag{6}$$

The following proposition shows that the proposed upper bound (5) is sharper than the upper bound (6) under a certain condition.

Proposition 1 *With the notations defined in Theorem 2, suppose that the following holds,*

$$\frac{1 + \nu}{2} - \frac{5\sqrt{2\tau} \chi_{n_p, n_u}^{(2)}(2\pi_+) \nu}{6\chi_{n_p, n_u}^{(1)}(2\pi_+)} \geq \rho. \tag{7}$$

Then, the proposed upper bound (5) is sharper than the previous result (6) proposed by Kiryo et al. (2017).

It is noteworthy that the second term in the left-hand side of (7) converges to zero as n_p and n_u increase because $\chi_{n_p, n_u}^{(1)}(2\pi_+) = O_{P_{X|Y=1}, P_X}((n_p \wedge n_u)^{-1/2})$ and $\chi_{n_p, n_u}^{(2)}(2\pi_+) = O_{P_{X|Y=1}, P_X}((n_p \wedge n_u)^{-1})$. Due to $(1 + \nu)/2 \geq \nu \geq \rho$, the condition (7) is quite reasonable if the upper bounds of the variances, σ_X^2 and $\sigma_{X|Y=1}^2$, are sufficiently small.

In binary classification, one ultimate goal is to find a classifier minimizing the misclassification error, or equivalently, minimizing the excess risk. Bartlett et al. (2006) showed that there is an invertible function $\psi : [-1, 1] \rightarrow [0, \infty)$ such that the excess risk $R_{\ell_{01}}(\hat{f}_{\mathcal{F}}) - \inf_{f \in \mathcal{U}} R_{\ell_{01}}(f)$ is bounded above by $\psi^{-1}(R_{\ell}(\hat{f}_{\mathcal{F}}) - \inf_{f \in \mathcal{U}} R_{\ell}(f))$ if the margin-based loss ℓ is classification-calibrated. In particular, Zhang (2004) showed that the excess risk is bounded above by the excess ℓ_h -risk, i.e., $R_{\ell_{01}}(\hat{f}_{\mathcal{F}}) - \inf_{f \in \mathcal{U}} R_{\ell_{01}}(f) \leq$

$R_{\ell_h}(\hat{f}_{\mathcal{F}}) - \inf_{f \in \mathcal{U}} R_{\ell_h}(f)$. This implies that an excess risk bound can be obtained by analyzing the excess ℓ_h -risk bound with Theorem 2. The following corollary provides the excess risk bound.

Corollary 1 (Excess risk bound for general function space) *With the notations defined in Theorem 2, for all $\alpha, \tau > 0$, the following holds with probability at least $1 - e^{-\tau}$,*

$$\begin{aligned} R_{\ell_{01}}(\hat{f}_{\mathcal{F}}) - \inf_{f \in \mathcal{U}} R_{\ell_{01}}(f) &\leq \inf_{f \in \mathcal{F}} R_{\ell_h}(f) - \inf_{f \in \mathcal{U}} R_{\ell_h}(f) \\ &\quad + C_{\alpha}(\mathbb{E}_{P_X^{n_u}}(\mathfrak{R}_{\mathcal{X}_u}(\mathcal{F})) + 2\pi_+ \mathbb{E}_{P_{X|Y=1}^{n_p}}(\mathfrak{R}_{\mathcal{X}_p}(\mathcal{F}))) \\ &\quad + C_{\tau, \rho^2}^{(1)} \chi_{n_p, n_u}^{(1)}(2\pi_+) + C_{\tau, \nu, \alpha}^{(2)} \chi_{n_p, n_u}^{(2)}(2\pi_+). \end{aligned}$$

4 WIPM optimizer with reproducing kernel Hilbert space

In this section, we provide a computationally efficient PU learning algorithm which builds an analytic classifier when a hypothesis space is a closed ball in RKHS. In addition, unlike the excess risk bound in Corollary 1, we explicitly derive the bound that converges to zero when the sample sizes n_p and n_u increase.

4.1 An analytic classifier via WMMD optimizer

To this end, we assume that $\mathcal{X} \subseteq [0, 1]^d$ is compact. Let $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ be a reproducing kernel defined on \mathcal{X} and \mathcal{H}_k be the associated RKHS with the inner product $\langle \cdot, \cdot \rangle_{\mathcal{H}_k} : \mathcal{H}_k \times \mathcal{H}_k \rightarrow \mathbb{R}$. We denote the induced norm by $\|\cdot\|_{\mathcal{H}_k}$. Denote a closed ball in RKHS \mathcal{H}_k with a radius $r > 0$, by $\mathcal{H}_{k,r} = \{f : \|f\|_{\mathcal{H}_k} \leq r\}$. We define the weighted maximum mean discrepancy (WMMD) between two probability measures P and Q with a weight w and a closed ball $\mathcal{H}_{k,r}$ by $\text{WMMD}_k(P, Q; w, r) := \text{WIPM}(P, Q; w, \mathcal{H}_{k,r})$. The name of WMMD comes from the maximum mean discrepancy (MMD), a popular example of the IPM whose function space is the unit ball $\mathcal{H}_{k,1}$, i.e., $\text{MMD}_k(P, Q) := \text{IPM}(P, Q; \mathcal{H}_{k,1})$ (Sriperumbudur et al. 2010a, b). As defined in Eq. (4), let $\hat{g}_{\mathcal{H}_{k,r}} \in \mathcal{H}_{k,r}$ be the empirical WMMD optimizer such that

$$\text{WMMD}_k(P_{X, n_u}, P_{X|Y=1, n_p}; w, r) = \frac{1}{n_u} \sum_{i=1}^{n_u} \hat{g}_{\mathcal{H}_{k,r}}(x_i^u) - \frac{w}{n_p} \sum_{i=1}^{n_p} \hat{g}_{\mathcal{H}_{k,r}}(x_i^p).$$

In addition, we set $\hat{f}_{\mathcal{H}_{k,r}} = -\hat{g}_{\mathcal{H}_{k,r}}$, which leads the corresponding classification rule to $\text{sign}(\hat{f}_{\mathcal{H}_{k,r}}(z))$. In the following proposition, we show that this classification rule has an analytic expression by exploiting the reproducing property $f(x) = \langle f, k(\cdot, x) \rangle_{\mathcal{H}_k}$ and the Cauchy-Schwarz inequality.

Proposition 2 *Let $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ be a bounded reproducing kernel. Then, the classification rule has a closed-form expression given by*

$$\text{sign}(\hat{f}_{\mathcal{H}_{k,r}}(z)) = \begin{cases} +1 & \text{if } (2\pi_+)^{-1} < \hat{\lambda}_{n_p, n_u}(z), \\ -1 & \text{otherwise,} \end{cases} \tag{8}$$

where

$$\hat{\lambda}_{n_p, n_u}(z) = \frac{n_p^{-1} \sum_{i=1}^{n_p} k(z, x_i^p)}{n_u^{-1} \sum_{i=1}^{n_u} k(z, x_i^u)}.$$

We call the classifier defined in Eq. (8) the *WMMD classifier* and the score $\hat{\lambda}_{n_p, n_u}(z)$ the *WMMD score* for z . One strength of the WMMD classifier is that the classification rule has a closed-form expression, resulting in computational efficiency. Furthermore, the WMMD score $\hat{\lambda}_{n_p, n_u}$ is independent of the class-prior π_+ , and thus we can obtain the score function without prior knowledge of the class-prior.

4.2 Explicit excess risk bound of WMMD classifier

Since the empirical WMMD optimizer $\hat{g}_{\mathcal{H}_{k,r}}$ is a special case of the empirical WIPM optimizer, we have an excess risk bound from the result of Corollary 1. However, without knowing convergence rates of the Rademacher complexities, $\mathbb{E}_{P_X^{n_u}}(\mathfrak{R}_{\mathcal{X}_u}(\mathcal{F}))$ and $\mathbb{E}_{P_{X|Y=1}^{n_p}}(\mathfrak{R}_{\mathcal{X}_p}(\mathcal{F}))$, and the approximation error, the consistency of the classifier remains unclear. In this subsection, we establish an explicit excess risk bound that vanishes. We first derive an explicit estimation error bound in the following lemma.

Lemma 1 (Explicit estimation error bound) *With the notations defined in Theorem 2, assume that a reproducing kernel k defined on a compact space \mathcal{X} is bounded. Let $r_1^{-1} = \sup_{x \in \mathcal{X}} \sqrt{k(x, x)}$. Then, we have $\mathcal{H}_{k,r_1} \subseteq \mathcal{M}$. Moreover, for all $\alpha, \tau > 0$, the following holds with probability at least $1 - e^{-\tau}$,*

$$R_{\ell_h}(\hat{f}_{\mathcal{H}_{k,r_1}}) - \inf_{f \in \mathcal{H}_{k,r_1}} R_{\ell_h}(f) \leq (C_\alpha + C_{\tau, \rho^2}^{(1)}) \chi_{n_p, n_u}^{(1)}(2\pi_+) + C_{\tau, v, \alpha}^{(2)} \chi_{n_p, n_u}^{(2)}(2\pi_+).$$

While the bound in Theorem 2 is expressed in terms $\mathbb{E}_{P_X^{n_u}}(\mathfrak{R}_{\mathcal{X}_u}(\mathcal{F}))$ and $\mathbb{E}_{P_{X|Y=1}^{n_p}}(\mathfrak{R}_{\mathcal{X}_p}(\mathcal{F}))$, these are evaluated in terms of n_p and n_u in the upper bound in Lemma 1, giving an explicit estimation error bound with $O((n_p \wedge n_u)^{-1/2})$ convergence rate. The key idea is to use reproducing property $f(x) = \langle f, k(\cdot, x) \rangle_{\mathcal{H}_k}$ and the Cauchy-Schwarz inequality to obtain an upper bound for the Rademacher complexity.

In the following lemma, we elaborate on the approximation error bound. To begin, for any $0 < \beta \leq 1$, let $\beta\mathcal{M} := \{\beta f : f \in \mathcal{M}\}$. Set $f_1^*(x) = \text{sign}(P(Y = 1 | X = x) - \frac{1}{2})$.

Lemma 2 (Approximation error bound over uniformly bounded hypothesis space) *With the notations defined in Lemma 1, we have*

$$\inf_{f \in \mathcal{H}_{k,r_1}} R_{\ell_h}(f) - \inf_{f \in \beta\mathcal{M}} R_{\ell_h}(f) \leq \beta \inf_{g \in \mathcal{H}_{k,r_1/\beta}} \|g - f_1^*\|_{L_2(P_X)},$$

for any $0 < \beta \leq 1$.

When $\beta = 1$, Lemma 2 implies that the approximation error $\inf_{f \in \mathcal{H}_{k,r_1}} R_{\ell_h}(f) - \inf_{f \in \mathcal{U}} R_{\ell_h}(f)$ is bounded above by $\inf_{g \in \mathcal{H}_{k,r_1}} \|g - f_1^*\|_{L_2(P_X)}$ due to $\inf_{f \in \mathcal{U}} R_{\ell_h}(f) = \inf_{f \in \mathcal{M}} R_{\ell_h}(f)$ (Lin 2002). Hence, a naive substitution to Corollary 1 will give a sub-optimal bound because $\inf_{g \in \mathcal{H}_{k,r_1}} \|g - f_1^*\|_{L_2(P_X)}$ is non-zero in general. In the following theorem, we rigorously establish the explicit excess risk bound which vanishes as n_p and n_u increase. We provide the pointer to a proof and the conditions (C1) and (C2) in ‘‘Appendix A’’.

Theorem 3 (Explicit excess risk bound) *With the notations defined in Theorem 2 and in Lemma 1, assume that the Gaussian kernel $k(x, y) = \exp(-\|x - y\|_2^2/2)$ is used. Then, $r_1 = 1$. Furthermore, under the conditions (C1) and (C2), for all $0 < s < 1$, $\alpha > 0$, and $\tau > 0$, the following holds with probability at least $1 - e^{-\tau}$:*

$$\begin{aligned} & R_{\ell_{01}}(\hat{f}_{\mathcal{H}_{k,1}}) - \inf_{f \in \mathcal{U}} R_{\ell_{01}}(f) \\ & \leq (C_\alpha + C_{\tau, \rho^2}^{(1)}) \chi_{n_p, n_u}^{(1)} (2\pi_+)^{1-s} + C_{\tau, \nu, \alpha}^{(2)} \chi_{n_p, n_u}^{(1)} (2\pi_+)^{-s} \chi_{n_p, n_u}^{(2)} (2\pi_+) \\ & \quad + C_d \left\| \frac{dP_X}{d\lambda} \right\|_{L_2(\lambda)} \|f_1^*\|_{1/4, 2} \{-s \ln \chi_{n_p, n_u}^{(1)} (2\pi_+)\}^{-1/16}, \end{aligned}$$

where

$$C_d = \left\{ \left(\frac{\sqrt{2\pi}}{\log 2} \sqrt{d} \right)^{1/8} + 64\sqrt{d} \left(\frac{8}{\pi} \right)^d \right\}.$$

Compared to Corollary 1, the excess risk bound in Theorem 3 has an explicit form, thus obviates the approximation error term. In addition, this bound converges to zero when n_p and n_u increase and the convergence rate is $O(\{\ln(n_p \wedge n_u)\}^{-1/16})$. To derive the excess risk bound, we first note that the misclassification error is determined by the sign of a discriminant function alone, *i.e.*, $\inf_{f \in \mathcal{U}} R_{\ell_{01}}(f) = \inf_{f \in \beta \mathcal{M}} R_{\ell_{01}}(f)$ for any $0 < \beta \leq 1$. Next, we show that $R_{\ell_{01}}(g) - \inf_{f \in \beta \mathcal{M}} R_{\ell_{01}}(f) \leq \{R_{\ell_h}(g) - \inf_{f \in \beta \mathcal{M}} R_{\ell_h}(f)\}/\beta$ by modifying Bartlett et al. (2006). We next obtain bounds for $\{R_{\ell_h}(\hat{f}_{\mathcal{H}_{k,r_1}}) - \inf_{f \in \mathcal{H}_{k,r_1}} R_{\ell_h}(f)\}/\beta$ using Lemma 1 and $\{\inf_{f \in \mathcal{H}_{k,r_1}} R_{\ell_h}(f) - \inf_{f \in \beta \mathcal{M}} R_{\ell_h}(f)\}/\beta$ using Lemma 2 and the previous result by Smale and Zhou (2003, Proposition 1.1) in terms of β . A carefully chosen β provides the explicit excess risk bound.

Niu et al. (2016) provided the excess risk bound expressed as a function of n_p, n_u . However, their bound included combined terms of the approximation error and the Rademacher complexity, as in Corollary 1. To the best of our knowledge, we are the first to explicitly derive the excess risk bound with convergence rate in terms of a function of n_p, n_u in PU learning.

5 Related work

Excess risk bound in noisy label literature PU learning can be considered as a special case of classification with asymmetric label noise, and many studies in this literature have shown consistency results similar to Theorem 3 (Natarajan et al. 2013). Patrini et al. (2016) derived an explicit estimation error when \mathcal{F} is a set of linear hypotheses and Blanchard et al. (2016) showed a consistency result of the excess risk bound when the hypothesis space is RKHS with universal kernels. While the two studies assumed the *one sample of data* scheme, the proposed bound is based on the *two samples of data* scheme. Therefore, our proposed excess risk bound is expressed in n_p and n_u , giving a new consistency theory.

Closed-form classifier Blanchard et al. (2010) suggested a score function similar to the WMMD score by using different bandwidth hyperparameters for the denominator and the numerator. However, with these differences, our method gains theoretical justification while their score function does not. Du Plessis et al. (2015) derived a closed-form classifier based on the squared loss. They estimated $P(Y = 1 | X) - P(Y = -1 | X)$ and showed the consistency of the estimation error bound in the *two samples of data* scheme. However, the

classifier is not scalable because it requires to compute the inverse of a $(n_p + n_u) \times (n_p + n_u)$ matrix.

6 Numerical experiments

In this section, we empirically analyze the proposed algorithm to demonstrate its practical efficacy using synthetic and real datasets. Pytorch implementation for the experiments is available at https://github.com/eraser347/WMMD_PU.

6.1 Synthetic data analysis

We first visualize the effect of increasing the sample sizes n_p and n_u on the discriminant ability of the proposed algorithm (Experiment 1). Then we compare performance with (i) the logistic loss ℓ_{\log} , denoted by LOG, (ii) the double hinge loss ℓ_{dh} , denoted by DH, both proposed by Du Plessis et al. (2015), (iii) the non-negative risk estimator method, denoted by NNPU, proposed by Kiryo et al. (2017), (iv) the threshold adjustment method, denoted by tADJ, proposed by Elkan and Noto (2008), and (v) the proposed algorithm, denoted by WMMD (Experiments 2, 3, and 4).

Experiment 1 In this case, we used the `two_moons` dataset whose underlying distributions are

$$X|Y = y, U \sim N \left(\begin{bmatrix} 2(1+y) - 4y \cos(\pi U) \\ (1+y) - 4y \sin(\pi U) \end{bmatrix}, \begin{bmatrix} 0.4^2 & 0 \\ 0 & 0.4^2 \end{bmatrix} \right),$$

where U refers to a uniform random variable ranges from 0 to 1 and $N(\mu, \Sigma)$ is the normal distribution with mean μ and covariance Σ . We used the 'make_moons' function in the Python module 'sklearn.datasets' (Pedregosa et al. 2011) to generate the datasets.

Figure 1 illustrates the decision boundaries of WMMD using the `two_moons` dataset. The first row displays the case where the unlabeled sample size is small, $n_u = 50$, and the second row displays the case where the unlabeled sample size is large, $n_u = 400$. The first and second columns display the case where the positive sample sizes are $n_p = 5$ and $n_p = 10$, respectively. The class-prior is fixed to $\pi_+ = 0.5$, and we assumed that the class-prior is known. We visualize the true mean function of the positive and negative data distributions with blue and red lines, respectively. The positive data are represented by blue diamond points, and the unlabeled data are represented by gray points. The decision boundaries of the WMMD classifier tend to correctly separate the two clusters as n_p and n_u increase.

In Experiments 2, 3, and 4, we evaluate: (i) the accuracy and area under the receiver operating characteristic curve (AUC) as n_u and π_+ change when the class-prior is known (Experiment 2) and unknown (Experiment 3); (ii) the elapsed training time (Experiment 4). In these experiments, we set up the underlying joint distribution as follows:

$$X | Y = y \sim N \left(y \frac{\mathbf{1}_2}{\sqrt{2}}, I_2 \right), Y \sim 2 \times \text{Bern}(\pi_+) - 1, \quad (9)$$

where $\text{Bern}(p)$ is the Bernoulli distribution with mean p , $\mathbf{1}_2 = (1, 1)^T$ is the 2 dimensional vector of all ones and I_2 is the identity matrix of size 2.

Experiment 2 In this experiment, we compare the accuracy and AUC of the five PU learning algorithms when the true class-prior π_+ is known. Figure 2a, c show the accuracy and AUC on various n_u . The training sample size for the positive data is $n_p = 100$ and the

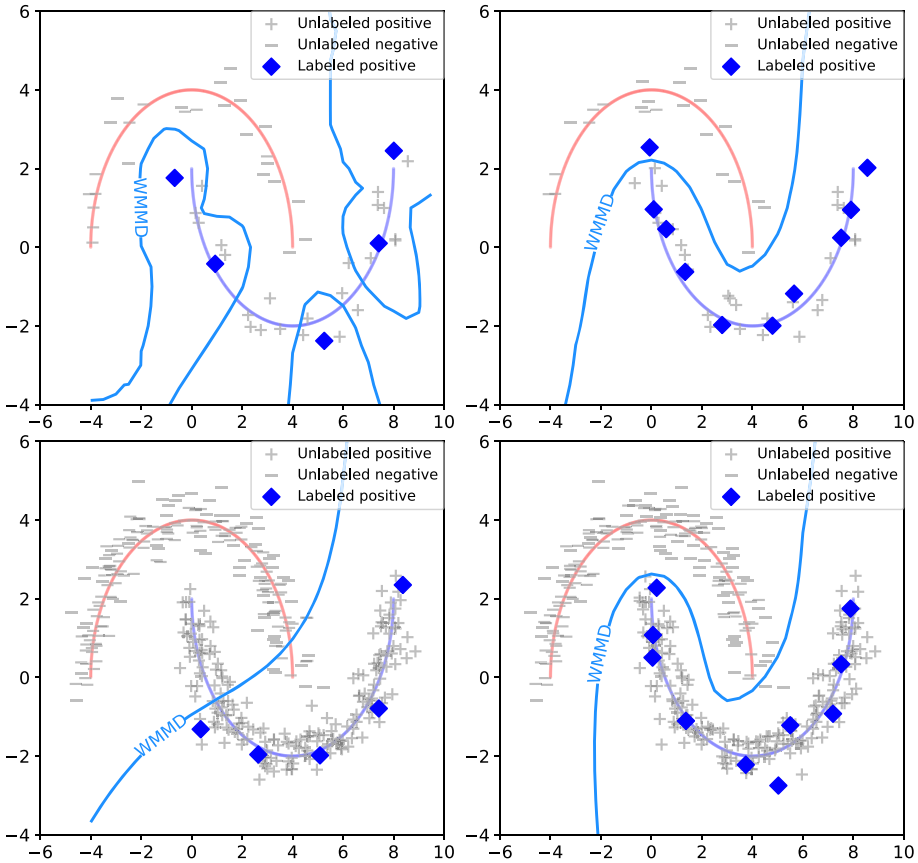


Fig. 1 The illustration of the decision boundaries of the WMMD classifier using the `two_moons` dataset with the increases in the size of positive and unlabeled samples. The true means of the positive and negative data distributions are plotted by blue and red lines respectively. The gray ‘+’ points and the gray ‘-’ points refer to the unlabeled positive and unlabeled negative training data, respectively (Color figure online)

class prior is $\pi_+ = 0.5$. The unlabeled sample size changes from 40 to 500 by 20. We repeat a random generation of training and test data 100 times. For comparison purposes, we add the 1-Bayes risk for each unlabeled sample size. In terms of accuracy, the proposed WMMD tends to be closer to the 1-Bayes risk as the n_u increases. Compared with other PU learning algorithms, WMMD achieves higher accuracy in every n_u and achieves comparable to or better AUC.

Figure 2b, d show a comparison of accuracy and AUC as π_+ changes. The training sample size for the positive and unlabeled data are $n_p = 100$ and $n_u = 400$, respectively. The class-prior π_+ changes from 0.05 to 0.95 by 0.05. The test sample size is 10^3 . Training and test data are repeatedly generated 100 times with different random seeds. In terms of accuracy, the proposed WMMD performs comparably with LOG and NNPU, showing advantages over DH and tADJ. When the true class-prior is less than equal to 0.8, WMMD performs better in terms of AUC, except for tADJ. The tADJ achieves the highest AUC because $P(Y = 1 | X = x)$ is proportional to $P(\{x \text{ is from the positive dataset} \} | X = x)$. This empirically shows that

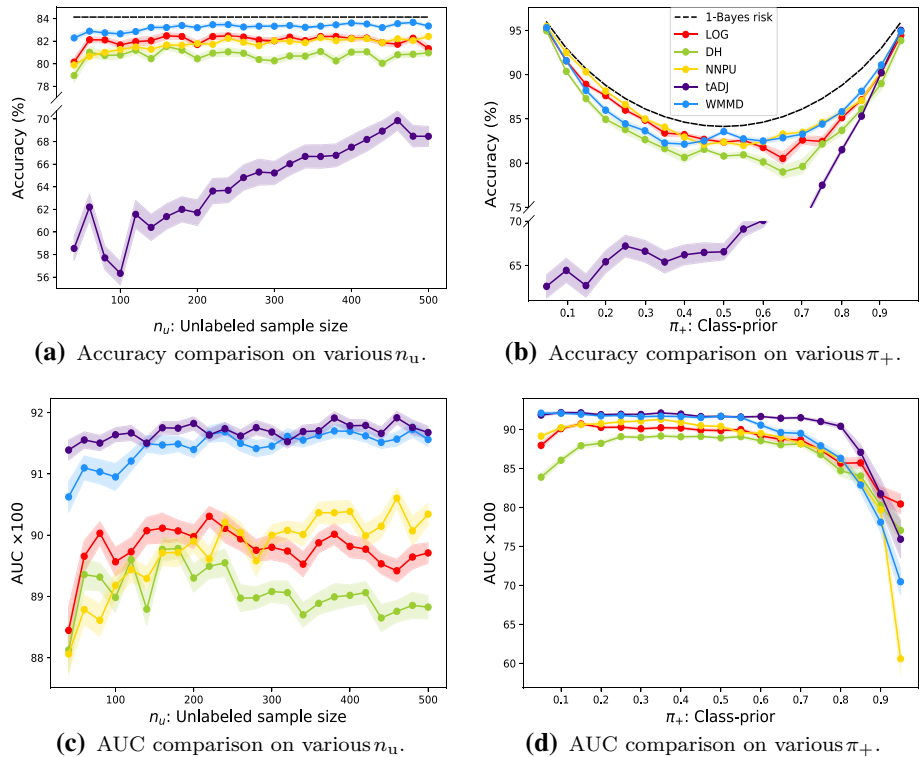



Fig. 2 The comparison of the accuracy and AUC of the five PU learning algorithms when each of n_u and π_+ changes. The dashed curve represents the 1-Bayes risk. The curve and the shaded region represent the average and the standard error, respectively, based on 100 replications

WMMD has a comparable discriminant ability to the other algorithms for a wide range of class-priors.

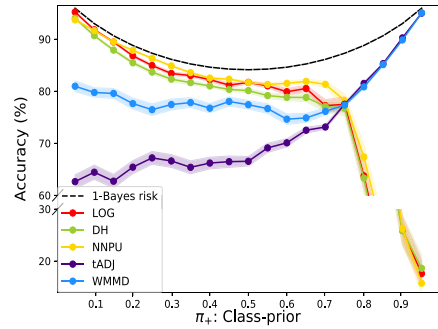
Experiment 3 The main goal of this subsection is to show the robustness of the proposed classifier in the case of unknown class-prior π_+ . In PU learning literature, π_+ has been frequently assumed to be known (Du Plessis et al. 2015; Niu et al. 2016; Kiryo et al. 2017; Kato et al. 2019). However, this assumption can be considered to be strong in real-world applications, and to correctly execute existing PU learning algorithms, an accurate estimate of π_+ is necessary. In this experiment, we compare the accuracy and AUC when the class-prior π_+ is unknown. For the WMMD classifier, we used a density-based method for the class-prior estimation which can be obtained as a byproduct of the proposed algorithm. We provide the pointer to a description in “Appendix A”. The results of LOG, DH, and NNPU are given for completeness sake using the ‘KM1’ method² by Ramaswamy et al. (2016). We take these estimates as true values and repeat the same comparative numerical experiments in Experiment 2.

Since the objective functions of the LOG, DH, and NNPU algorithms depend on the estimate $\hat{\pi}_+$, we anticipate that both the accuracy and AUC rely on the quality of the estimation.

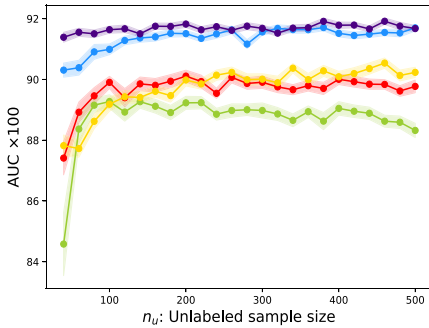
² While the ‘KM2’ method by Ramaswamy et al. (2016) is often considered to be a state-of-the-art method for estimating π_+ , in our experiments, estimates based on the ‘KM2’ method have a larger estimation error than that of the ‘KM1’ method and thus we omitted it.




(a) Accuracy comparison on various n_u when the class-prior is unknown.



(b) Accuracy comparison on various π_+ when the class-prior is unknown.



(c) AUC comparison on various n_u when the class-prior is unknown.



(d) AUC comparison on various π_+ when the class-prior is unknown.

Fig. 3 The comparison of the accuracy and AUC of the five PU learning algorithms when each of n_u and π_+ changes under the situation where π_+ is unknown. The dashed curve represents the 1-Bayes risk. The curve and the shaded region represent the average and the standard error, respectively, based on 100 replications. LOG, DH, and NNPU use the estimate of the class-prior from the ‘KMI’ method

On the other hand, the tADJ algorithm does not depend on the class-prior, so the performance is not affected. Also, as the proposed score function does not depend on the class-prior π_+ , and since π_+ is used only to determine a cutoff, the AUC of the proposed algorithm is less affected by the estimation of π_+ .

Figure 3a, c compare the accuracy and AUC as a function of n_u . WMMD performs worse than LOG, DH, and NNPU, while AUC is higher. Though tADJ shows poor accuracy in a wide range, it achieves high AUC comparable to WMMD. As we anticipated, WMMD is more robust than LOG, DH, and NNPU in AUC. This is possibly because our score function $\hat{\lambda}_{n_p, n_u}$ does not depend on π_+ . A similar trend can be found in Figure 3b, d. We note that the ‘KMI’ method is not scalable and thus may not be used for large-scale datasets.

Experiment 4 In this experiment, we compare the elapsed training time, including hyperparameter optimization, of the five PU learning algorithms. The data are generated from the distributions described in Eq. (9), and we set $n_p = 100, n_u = 400$, and $\pi_+ = 0.5$. The elapsed time is measured with 20 Intel®Xeon®E5-2630 v4@2.20GHz CPU processors.

Table 1 compares the elapsed training time and its ratio relative to that of WMMD. WMMD takes the shortest time among the five baseline methods. In particular, the training time for WMMD is at least about 300 times shorter than that of the LOG and DH methods. This is

Table 1 A summary of elapsed training time and its ratio for the five PU learning algorithms based on 100 replications

	LOG	DH	NNPU	tADJ	WMMD
In seconds $\times 10$	90.0 \pm 4.7	96.1 \pm 6.1	6.0 \pm 0.1	0.4 \pm 0.0	0.2 \pm 0.0
In ratio	347.9 \pm 23.4	371.0 \pm 28.5	23.2 \pm 1.1	1.8 \pm 0.0	1.0 \pm 0.0

We set $n_p = 100$, $n_u = 400$, and $\pi_+ = 0.5$. Average and standard error are denoted by ‘average \pm standard error’

Table 2 A summary of the eleven binary classification datasets

Dataset	d	# of samples	n_p	n_u	n_{te}	π_+	Scale
heart_scale	12	122	10	60	60	0.62	Small
sonar_scale	60	207	10	100	100	0.47	Small
australian_scale	12	449	20	220	220	0.51	Small
australian_scale2	12	449	10	130	130	0.15	Small
breast-cancer_scale	10	683	20	340	340	0.35	Small
breast-cancer_scale2	10	683	40	340	340	0.65	Small
diabetes_scale	8	759	50	380	370	0.65	Small
skin_nonskin	3	245,057	10^3	10^5	10^5	0.79	Large
skin_nonskin2	3	245,057	10^3	10^5	10^5	0.21	Large
epsilon_normalized	2000	500,000	10^3	4×10^5	10^5	0.50	Large
HIGGS	26	8,786,441	10^3	10^6	10^5	0.50	Large

‘# of samples’ denotes the number of total samples after removing incomplete observations. We denote the number of positive, unlabeled, and test samples, by n_p , n_u , and n_{te} after the random sampling, respectively. We categorize the eleven datasets into two groups: the first seven datasets as small-scale and the last four datasets as large-scale

because the WMMD classifier has an analytic form while the LOG and DH methods require solving a non-linear programming problem.

6.2 Real data analysis

We demonstrate the practical utility of the proposed algorithm using the eight real binary classification datasets from the LIBSVM³ (Chang and Lin 2011). Since some observations from the raw datasets are not completely recorded, we removed such observations and construct the dataset with fully recorded data. Next, to investigate the effect of varying π_+ , we artificially reconstructed \mathcal{X}_p and \mathcal{X}_u through a random sampling from the fully recorded datasets. For the three datasets *australian_scale*, *breast-cancer_scale* and *skin_nonskin*, we reconstructed the data so that the resulting class-prior π_+ ranges from 0.15 to 0.79. We add the suffix 2 for those datasets. We randomly resampled data 100 times for the seven small datasets and 10 times for the four big datasets: *skin_nonskin*, *skin_nonskin2*, *epsilon_normalized*, and *HIGGS*. Table 2 summarizes statistics for the eleven real datasets. We conduct two comparative numerical experiments when π_+ is known and unknown.

³ <https://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/>.

Table 3 Accuracy and AUC comparison using the real datasets when the class-prior π_+ is known

Dataset	LOG	DH	NNPU	tADJ	WMMD
Accuracy (in %)					
heart_scale	70.5 ± 0.8	68.4 ± 0.9	71.0 ± 0.8	65.1 ± 0.8	71.6 ± 0.8
sonar_scale	55.8 ± 0.6	52.9 ± 0.6	63.2 ± 0.6	60.7 ± 0.6	62.4 ± 0.6
australian_scale	85.4 ± 0.4	84.9 ± 0.6	79.2 ± 0.5	80.0 ± 0.7	84.2 ± 0.6
australian_scale2	85.7 ± 0.2	85.7 ± 0.2	86.7 ± 0.3	75.4 ± 2.0	86.2 ± 0.2
breast-cancer_scale	95.8 ± 0.1	96.0 ± 0.3	90.1 ± 0.3	91.0 ± 0.4	89.3 ± 0.5
breast-cancer_scale2	95.6 ± 0.3	94.4 ± 0.8	95.9 ± 0.1	92.5 ± 0.2	94.2 ± 0.3
diabetes_scale	66.7 ± 0.7	65.5 ± 0.9	69.4 ± 0.4	67.9 ± 0.3	66.4 ± 0.2
skin_nonskin	–	–	98.2 ± 0.1	78.0 ± 0.4	85.3 ± 0.7
skin_nonskin2	–	–	98.6 ± 0.0	93.9 ± 0.1	98.1 ± 0.2
epsilon_normalized	–	–	64.5 ± 0.3	63.1 ± 0.1	56.3 ± 1.3
HIGGS	–	–	56.3 ± 0.2	52.6 ± 0.1	54.0 ± 0.2
AUC × 100					
heart_scale	78.4 ± 1.0	78.3 ± 1.1	73.8 ± 0.9	72.5 ± 1.1	79.0 ± 0.9
sonar_scale	61.2 ± 0.8	60.6 ± 0.9	67.4 ± 0.7	66.2 ± 0.7	68.9 ± 0.8
australian_scale	91.1 ± 0.2	91.3 ± 0.3	87.8 ± 0.4	87.8 ± 0.5	90.4 ± 0.4
australian_scale2	89.2 ± 0.4	87.3 ± 0.4	84.3 ± 0.6	85.9 ± 0.7	88.6 ± 0.6
breast-cancer_scale	99.4 ± 0.0	99.3 ± 0.0	97.8 ± 0.1	95.6 ± 0.4	99.5 ± 0.0
breast-cancer_scale2	99.3 ± 0.0	99.2 ± 0.1	99.3 ± 0.0	97.2 ± 0.2	98.7 ± 0.2
diabetes_scale	74.0 ± 0.6	71.5 ± 1.1	73.5 ± 0.6	74.7 ± 0.5	74.5 ± 0.7
skin_nonskin	–	–	99.5 ± 0.1	94.8 ± 0.1	99.4 ± 0.1
skin_nonskin2	–	–	99.7 ± 0.0	94.6 ± 0.0	99.8 ± 0.0
epsilon_normalized	–	–	70.0 ± 0.4	69.3 ± 0.1	62.2 ± 2.3
HIGGS	–	–	59.6 ± 0.2	65.3 ± 0.1	55.7 ± 0.3

We denote the memory error results for LOG and DH by the hyphen. Average and standard error are denoted by ‘average ± standard error’. Boldface numbers denote the best and equivalent algorithms with respect to a t-test with a significance level of 5%

Table 3 shows the average and the standard error of the accuracy and AUC when the class-prior π_+ is known. LOG and DH fail to compute the $(n_p + n_u) \times (n_p + n_u)$ Gram matrix due to out of memory in the 12 GB GPU memory limit. WMMD achieves comparable to or better accuracy and AUC than LOG, DH, and tADJ on most datasets. Compared to NNPU, WMMD performs comparably on the small datasets. However, NNPU achieves higher accuracy on skin_nonskin, epsilon_normalized, and HIGGS. The neural network used in NNPU fits well to the complicated and high-dimensional structure of data and shows high accuracy.

Table 4 compares the average and the standard error of the accuracy and AUC when the class-prior π_+ is unknown. As in Experiment 3 in Sect. 6.1, we estimate π_+ using the ‘KM1’ method for LOG, DH, and NNPU, and using the density-based method for WMMD. The LOG, DH, and NNPU algorithms are implemented on the seven small-scale datasets alone because the method by Ramaswamy et al. (2016) is not feasible with the large-scale datasets (Bekker and Davis 2018). Overall, WMMD shows comparable to or better performances than other PU learning algorithms on most datasets. Compared to Table 3, WMMD and

Table 4 Accuracy and AUC comparison using the real datasets when the class-prior π_+ is unknown

Dataset	LOG	DH	NNPU	tADJ	WMMD
Accuracy (in %)					
heart_scale	39.5 ± 0.8	39.1 ± 0.7	42.4 ± 0.9	65.1 ± 0.8	70.5 ± 0.7
sonar_scale	53.5 ± 0.6	52.1 ± 0.5	59.9 ± 0.7	60.7 ± 0.6	54.3 ± 0.8
australian_scale	50.0 ± 0.2	50.0 ± 0.2	50.0 ± 0.2	80.0 ± 0.7	79.4 ± 1.0
australian_scale2	84.9 ± 0.2	84.9 ± 0.2	85.6 ± 0.3	85.5 ± 1.0	80.2 ± 1.0
breast-cancer_scale	65.0 ± 0.2	65.2 ± 0.2	65.0 ± 0.2	91.0 ± 0.4	93.0 ± 1.0
breast-cancer_scale2	35.0 ± 0.2	35.0 ± 0.2	35.0 ± 0.2	92.5 ± 0.2	92.6 ± 0.4
diabetes_scale	36.0 ± 0.3	41.1 ± 1.1	37.9 ± 0.5	67.9 ± 0.3	65.1 ± 0.2
skin_nonskin	–	–	–	78.0 ± 0.4	82.2 ± 0.9
skin_nonskin2	–	–	–	93.9 ± 0.1	95.7 ± 0.4
epsilon_normalized	–	–	–	63.1 ± 0.1	49.9 ± 0.1
HIGGS	–	–	–	52.6 ± 0.1	50.9 ± 0.0
AUC × 100					
heart_scale	67.2 ± 1.8	67.2 ± 1.4	71.0 ± 0.9	72.5 ± 1.1	77.1 ± 0.9
sonar_scale	60.5 ± 0.9	62.2 ± 0.9	66.6 ± 0.8	66.2 ± 0.7	69.7 ± 0.8
australian_scale	78.4 ± 1.1	72.5 ± 1.4	80.3 ± 0.6	87.8 ± 0.5	90.3 ± 0.3
australian_scale2	92.9 ± 0.2	89.6 ± 0.6	85.9 ± 0.7	92.4 ± 0.3	93.3 ± 0.2
breast-cancer_scale	98.9 ± 0.1	93.5 ± 1.8	54.8 ± 1.1	95.6 ± 0.4	99.5 ± 0.0
breast-cancer_scale2	14.4 ± 2.0	19.4 ± 3.8	91.5 ± 0.3	97.2 ± 0.2	99.0 ± 0.1
diabetes_scale	64.0 ± 1.3	63.8 ± 1.4	72.6 ± 0.5	74.7 ± 0.5	75.9 ± 0.5
skin_nonskin	–	–	–	94.8 ± 0.1	99.5 ± 0.1
skin_nonskin2	–	–	–	94.6 ± 0.0	99.8 ± 0.0
epsilon_normalized	–	–	–	69.3 ± 0.1	59.7 ± 1.8
HIGGS	–	–	–	65.3 ± 0.1	55.4 ± 0.2

The ‘KM1’ method by Ramaswamy et al. (2016) is used for LOG, DH, and NNPU, and the density-based method is used for WMMD. We denote the infeasible cases due to ‘KM1’ method by the hyphen. Other details are given in Table 3

tADJ show robustness to unknown π_+ in terms of AUC. This is because WMMD and tADJ do not require estimation of π_+ to construct score functions. In contrast, the other methods require an estimate $\hat{\pi}_+$, and we observe a substantial drop in accuracy and AUC when the ‘KM1’ method estimate is used.

7 Concluding remarks

Existing methods use different objective functions and hypothesis spaces, and as a consequence, different optimization algorithms. Hence, there is no reason that one method outperforms uniformly for all scenarios. It is possible that one particular method may outperform in one scenario, for example, NNPU proposed by Kiryo et al. (2017) would perform better in complicated data settings because of the expressive power of neural networks. However, the proposed method has a clear computational advantage due to the closed-form as well as theoretical strength in terms of the explicit excess risk bound. Further, the proposed

method works reasonably well in both cases in which π_+ is known or unknown. In this regard, we believe the proposed method can be used as a principled and easy-to-compute baseline algorithm in PU learning.

Acknowledgements YK, WK, and MCP were supported by the National Research Foundation of Korea under Grant NRF-2017R1A2B4008956. MS was supported by JST CREST JPMJCR1403.

A Proofs and implementation details

All the proofs, implementation details, and additional experiments are provided in the extended version due to the page limit (<https://arxiv.org/abs/1901.09503>). However, we have no objection to the inclusion of the proofs here when the paper get accepted.

References

- Arjovsky, M., Chintala, S., & Bottou, L. (2017). Wasserstein generative adversarial networks. In *International conference on machine learning*, pp. 214–223.
- Bartlett, P. L., Jordan, M. I., & McAuliffe, J. D. (2006). Convexity, classification, and risk bounds. *Journal of the American Statistical Association*, *101*(473), 138–156.
- Bartlett, P. L., & Mendelson, S. (2002). Rademacher and gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, *3*(Nov), 463–482.
- Bekker, J., & Davis, J. (2018). Estimating the class prior in positive and unlabeled data through decision tree induction. In *Proceedings of the 32th AAAI conference on artificial intelligence*.
- Blanchard, G., Flaska, M., Handy, G., Pozzi, S., & Scott, C. (2016). Classification with asymmetric label noise: Consistency and maximal denoising. *Electronic Journal of Statistics*, *10*(2), 2780–2824.
- Blanchard, G., Lee, G., & Scott, C. (2010). Semi-supervised novelty detection. *Journal of Machine Learning Research*, *11*(Nov), 2973–3009.
- Chang, C.-C., & Lin, C.-J. (2011). Libsvm: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology (TIST)*, *2*(3), 27.
- Chapelle, O., Schölkopf, B., & Zien, A. (2006). *Semi-supervised Learning*. Cambridge: MIT Press.
- Collobert, R., Sinz, F., Weston, J., & Bottou, L. (2006). Trading convexity for scalability. In *Proceedings of the 23rd international conference on machine learning*, ACM, pp. 201–208.
- Denis, F., Gilleron, R., & Letouzey, F. (2005). Learning from positive and unlabeled examples. *Theoretical Computer Science*, *348*(1), 70–83.
- Du Plessis, M. C., Niu, G., & Sugiyama, M. (2014). Analysis of learning from positive and unlabeled data. In *Advances in neural information processing systems*, pp. 703–711.
- Du Plessis, M. C., Niu, G., & Sugiyama, M. (2015). Convex formulation for learning from positive and unlabeled data. In *International conference on machine learning*, pp. 1386–1394.
- Elkan, C., & Noto, K. (2008). Learning classifiers from only positive and unlabeled data. In *Proceedings of the 14th ACM SIGKDD international conference on knowledge discovery and data mining*, ACM, pp. 213–220.
- Gong, T., Wang, G., Ye, J., Xu, Z., & Lin, M. (2018). Margin based pu learning. In *AAAI conference on artificial intelligence*.
- Gretton, A., Borgwardt, K. M., Rasch, M. J., Schölkopf, B., & Smola, A. (2012). A kernel two-sample test. *Journal of Machine Learning Research*, *13*(Mar), 723–773.
- Gretton, A., Smola, A. J., Huang, J., Schmittfull, M., Borgwardt, K. M., Schölkopf, B., Candela, Q., Sugiyama, M., Schwaighofer, A., Lawrence, N. D. et al. (2009). Covariate shift by kernel mean matching. In *Dataset shift in machine learning*, MIT Press, pp. 131–160.
- Huang, J., Gretton, A., Borgwardt, K. M., Schölkopf, B., & Smola, A. J. (2007). Correcting sample selection bias by unlabeled data. In *Advances in neural information processing systems*, pp. 601–608.
- Kato, M., Teshima, T., & Honda, J. (2019). Learning from positive and unlabeled data with a selection bias. In *International conference on learning representations*. URL <https://openreview.net/forum?id=rJzLciCqKm>.
- Kiryo, R., Niu, G., Plessis, Du Marthinus C., & Sugiyama, M. (2017). Positive-unlabeled learning with non-negative risk estimator. In *Advances in neural information processing systems*, pp. 1675–1685.

- Li, X., & Liu, B. (2003). Learning to classify texts using positive and unlabeled data. In *Proceedings of the 18th international joint conference on artificial intelligence*, Morgan Kaufmann Publishers Inc., pp. 587–592.
- Li, X.-L., & Liu, B. (2005). Learning from positive and unlabeled examples with different data distributions. In *European conference on machine learning*, pp. 218–229.
- Lin, Y. (2002). Support vector machines and the bayes rule in classification. *Data Mining and Knowledge Discovery*, 6(3), 259–275.
- Liu, B., Dai, Y., Li, X., Lee, W. S., & Yu, P. S. (2003). Building text classifiers using positive and unlabeled examples. In *Third IEEE international conference on data mining, 2003, ICDM 2003*, IEEE, pp. 179–186.
- Liu, B., Lee, W. S., Yu, P. S., & Li, X. (2002). Partially supervised classification of text documents. In *International conference on machine learning*, Vol. 2, Citeseer, pp. 387–394.
- Müller, A. (1997). Integral probability metrics and their generating classes of functions. *Advances in Applied Probability*, 29(2), 429–443.
- Natarajan, N., Dhillon, I. S., Ravikumar, P. K., & Tewari, A. (2013). Learning with noisy labels. In *Advances in neural information processing systems*, pp. 1196–1204.
- Niu, G., Du Plessis, M. C., Sakai, T., Ma, Y., & Sugiyama, M. (2016). Theoretical comparisons of positive-unlabeled learning against positive-negative learning. In *Advances in neural information processing systems*, pp. 1199–1207.
- Oh, C. Y., Gavves, E., & Welling, M. (2018). Bock: Bayesian optimization with cylindrical kernels. *arXiv preprint arXiv:1806.01619*.
- Patrini, G., Nielsen, F., Nock, R., & Carioni, M. (2016). Loss factorization, weakly supervised learning and label noise robustness. In *International conference on machine learning*, pp. 708–717.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., et al. (2011). Scikit-learn: Machine learning in python. *Journal of Machine Learning Research*, 12, 2825–2830.
- Ramaswamy, H., Scott, C., & Tewari, A. (2016). Mixture proportion estimation via kernel embeddings of distributions. In *International conference on machine learning*, pp. 2052–2060.
- Sakai, T., Du Plessis, M. C., Niu, G., & Sugiyama, M. (2017). Semi-supervised classification based on classification from positive and unlabeled data. In *International conference on machine learning*, pp. 2998–3006.
- Sansone, E., De Natale, F. G. B., & Zhou, Z. (2019). Efficient training for positive unlabeled learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(11), 2584–2598.
- Scott, C., & Blanchard, G. (2009). Novelty detection: Unlabeled data definitely help. In *Artificial intelligence and statistics*, pp. 464–471.
- Smale, S., & Zhou, D.-X. (2003). Estimating the approximation error in learning theory. *Analysis and Applications*, 1(01), 17–41.
- Sriperumbudur, B. K., Fukumizu, K., Gretton, A., Schölkopf, B., & Lanckriet, G. R. G. (2012). On the empirical estimation of integral probability metrics. *Electronic Journal of Statistics*, 6, 1550–1599.
- Sriperumbudur, B. K., Fukumizu, K., & Lanckriet, G. (2010a). On the relation between universality, characteristic kernels and rkhs embedding of measures. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pp. 773–780.
- Sriperumbudur, B. K., Gretton, A., Fukumizu, K., Schölkopf, B., & Lanckriet, G. R. G. (2010b). Hilbert space embeddings and metrics on probability measures. *Journal of Machine Learning Research*, 11(Apr), 1517–1561.
- Steinwart, I., & Christmann, A. (2008). *Support vector machines*. Berlin: Springer.
- Tolstikhin, I., Bousquet, O., Gelly, S., & Schoelkopf, B. (2018). Wasserstein auto-encoders. In *International conference on learning representations*.
- Ward, G., Hastie, T., Barry, S., Elith, J., & Leathwick, J. R. (2009). Presence-only data and the EM algorithm. *Biometrics*, 65(2), 554–563.
- Xiao, Y., Liu, B., Yin, J., Cao, L., Zhang, C., & Hao, Z. (2011). Similarity-based approach for positive and unlabelled learning. In *Twenty-second international joint conference on artificial intelligence*.
- Yan, H., Ding, Y., Li, P., Wang, Q., Xu, Y., & Zuo, W. (2017). Mind the class weight bias: Weighted maximum mean discrepancy for unsupervised domain adaptation. In *2017 IEEE conference on computer vision and pattern recognition (CVPR)*, IEEE, pp. 945–954.
- Yang, P., Li, X., Chua, H.-N., Kwok, C.-K., & Ng, S.-K. (2014). Ensemble positive unlabeled learning for disease gene identification. *PLoS ONE*, 9(5), e97079.
- Yang, P., Li, X.-L., Mei, J.-P., Kwok, C.-K., & Ng, S.-K. (2012). Positive-unlabeled learning for disease gene identification. *Bioinformatics*, 28(20), 2640–2647.
- Zhang, J., Wang, Z., Yuan, J., Tan, Y.-P. (2017). Positive and unlabeled learning for anomaly detection with multi-features. In *Proceedings of the 2017 ACM on multimedia conference*, ACM, pp. 854–862.
- Zhang, T. (2004). Statistical behavior and consistency of classification methods based on convex risk minimization. *Annals of Statistics*, 32, 56–85.

Zuluaga, M. A., Hush, D., Delgado, E., Leyton, J. F., Hoyos, M. H., & Orkisz, M. (2011). Learning from only positive and unlabeled data to detect lesions in vascular ct images. In *International conference on medical image computing and computer-assisted intervention*, Springer, pp. 9–16.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.