

Estimating Expected Calibration Error for Positive-Unlabeled Learning

Anonymous authors

Paper under double-blind review

Abstract

The reliability of probabilistic classifiers hinges on their calibration—the property that their confidence accurately reflect the true class probabilities. The expected calibration error (ECE) is a standard metric for quantifying the calibration of classifiers. However, its estimation presumes access to ground-truth labels. In positive-unlabeled (PU) learning, only positive and unlabeled data are available, which makes the standard ECE estimator inapplicable. Although PU learning has been extensively studied for risk estimation and classifier training, calibration in this setting has received little attention. In this paper, we present PU-ECE, the first ECE estimator for PU data. We provide non-asymptotic bias bounds and prove convergence rates that match those of the fully supervised ECE with an optimal bin size. Furthermore, we develop an information-theoretic generalization error analysis of PU-ECE by formalizing the conditional mutual information (CMI) for a PU setting. Experiments on synthetic and real-world benchmark datasets validate our theoretical analysis and demonstrate that our PU-based ECE estimator achieves performance comparable to that of the fully-labeled ECE estimator.

1 Introduction

Deep learning models achieve state-of-the-art accuracy in various classification tasks (Devlin et al., 2019; Dosovitskiy et al., 2021; Radford et al., 2021). However, their deployment in safety-critical applications requires more than just high prediction accuracy. For instance, in medical diagnosis (Jiang et al., 2012), autonomous systems (Muthali et al., 2023), and financial decision-making (Bahnsen et al., 2014), the confidence of the predictions provided by machine learning models must accurately reflect the true probability of the predicted events. This property, known as *calibration* (Dawid, 1982; Guo et al., 2017), is essential for reliable decision-making. The calibration error, or *true calibration error* (TCE) (Naeini et al., 2015; Guo et al., 2017; Roelofs et al., 2022), is a widely used metric for evaluating the calibration of classifiers. The TCE is typically estimated by partitioning the confidence score space into bins, followed by the calculation of the average difference between the confidence score and the empirical probability of the true label given the confidence score in each bin (Zadrozny & Elkan, 2001; Naeini et al., 2015). The resulting estimator is called the *expected calibration error* (ECE)¹ (Roelofs et al., 2022). Theoretical analysis of the ECE has also been conducted, establishing bias bounds and the consistency of the ECE to the TCE as the number of samples increases, with an appropriate binning strategy Gupta & Ramdas (2021); Futami & Fujisawa (2024). The results show an $\mathcal{O}(n^{-1/3})$ convergence rate of the ECE, where n is the number of test samples. The naive computation of the ECE requires a fully labeled dataset, where each sample has both confidence scores and true labels.

However, in practice, fully labeled data may not be available or may be costly to obtain (Jaskie & Spanias, 2019; Li et al., 2022; Bepler et al., 2019). For instance, when identifying protein particles in microscopy images, annotating negative data is costly due to diverse backgrounds (Bepler et al., 2019). In such cases, positive-unlabeled (PU) learning, which aims to train a classifier from only positive and unlabeled data, is a promising approach (Liu et al., 2002; Elkan & Noto, 2008; Natarajan et al., 2013; du Plessis et al., 2014).

¹As some researchers have pointed out, both TCE and ECE are sometimes called "ECE". For clarity, we follow the convention of Futami & Fujisawa (2024).

Despite extensive research on PU learning methods, the majority of these methods have focused on maximizing classification accuracy (Elkan & Noto, 2008; du Plessis et al., 2015; Kiryo et al., 2017; Chen et al., 2020; Xinrui et al., 2023). On the other hand, the critical issue of calibration in the PU setting has been largely overlooked. As a key motivation for this paper, we first reveal a surprising phenomenon: in PU learning, the standard binary cross-entropy loss, which is strictly proper in supervised learning (Good, 1952; Savage, 1971), can lead to significantly worse calibration than a noise-robust loss that is not strictly proper (Charoenphakdee et al., 2019) but often used in PU learning (Kiryo et al., 2017; Chen et al., 2020; Zhao et al., 2022) (details in Section 5.2). However, evaluating calibration in the PU setting is challenging—the absence of negative labels makes it impossible to apply the standard ECE estimator, which prevents practitioners from evaluating the calibration of PU classifiers.

In this paper, we address the problem of estimating the calibration error in PU learning settings. We introduce *PU-ECE*, a novel method for estimating the ECE using only PU data, without requiring negative labels. The core idea is to reframe the ECE computation by leveraging the statistical properties of PU datasets. By applying Bayes’ rule within each bin, we can estimate from PU data the conditional expectation of the true label given the confidence score, which is necessary for the ECE computation. Our theoretical analysis establishes non-asymptotic bias bounds of PU-ECE, which guarantee the precision of our estimator from the TCE in a finite sample. With these bounds, we prove that with an optimal binning strategy, our estimator achieves a convergence rate that matches the convergence rate of the ECE in fully-supervised settings of $\mathcal{O}(n^{-1/3})$, where n is the number of test samples (Futami & Fujisawa, 2024). An information-theoretic generalization error analysis of PU-ECE is also provided by formalizing the functional conditional mutual information (fCMI) (Steinke & Zakynthinou, 2020; Hellström & Durisi, 2022) for a PU setting. In addition, we demonstrate through extensive experiments on synthetic and benchmark datasets that our approach achieves estimation biases comparable to those of the fully-supervised method.

The rest of this paper is organized as follows. In Section 2, we formally describe the problem setting and define notations. In Section 3, we present the PU-ECE method and its theoretical properties, and in Section 4, we present its information-theoretic analysis. Section 5 provides experimental results to validate our approach. Finally, we conclude the paper in Section 6.

2 Problem Formulation

We denote random variables by capital letters and their realizations by lowercase letters. Let \mathcal{X} be the input space and $\mathcal{Y} = \{0, 1\}$ be the label space. Let $P(X, Y)$ be the probability distribution over $\mathcal{X} \times \mathcal{Y}$, $P(X)$ be the marginal distribution of X , $P(X|Y = y)$ be the conditional distribution of X given $Y = y \in \mathcal{Y}$, and $\pi_P = P(Y = 1)$ be the class prior. The hypothesis space is denoted by $\mathcal{H} \subseteq \mathbb{R}^{\mathcal{X}}$ and a classifier $h \in \mathcal{H}$ outputs a model’s logit score for a given input $x \in \mathcal{X}$. The model’s confidence score is then obtained by applying the sigmoid function $\sigma(z) = (1 + e^{-z})^{-1}$, which we denote by $f(x) := \sigma(h(x))$. Let $\mathcal{F} = \{\sigma(h(\cdot)) \mid h \in \mathcal{H}\}$. We denote the parameterized classifier and its confidence score function by h_w and f_w with $w \in \mathcal{W}$, where \mathcal{W} is the parameter space. If the parameter w is fixed, we simply write h and f . Let $\mathbb{1}_A$ be the indicator function, which returns 1 if the condition A is true and 0 otherwise. We define $[n] = \{1, 2, \dots, n\}$ for $n \in \mathbb{N}$.

2.1 Supervised Learning

In supervised learning, the dataset is a set of independent and identically distributed (i.i.d.) samples drawn from $P(X, Y)$. Let $l : \mathbb{R} \times \mathcal{Y} \rightarrow \mathbb{R}_{\geq 0}$ be the loss function, e.g., the 0-1 loss (Wald, 1945) for the misclassification rate. Let $S_{\text{PN}} := \{(X_m^{\text{PN}}, Y_m^{\text{PN}})\}_{m=1}^{n_{\text{PN}}}$ be a dataset of n_{PN} samples drawn from $P(X, Y)$. The expected risk of h on $P(X, Y)$ is given by $\mathbb{E}_{X,Y} [l(h(X), Y)]$. The empirical risk on S_{PN} is given by $\frac{1}{n} \sum_{m=1}^n l(h(X_m^{\text{PN}}), Y_m^{\text{PN}})$. Classifiers are trained by the empirical risk minimization (ERM) (Vapnik & Chervonenkis, 1971).

2.2 Calibration Error Estimation

A classifier is said to be perfectly calibrated if its confidence score matches the true class probability (Murphy & Epstein, 1967; Gupta & Ramdas, 2021). Mathematically, this means that for all confidence scores $p \in [0, 1]$, we have $P(Y = 1 \mid f(X) = p) = p$ almost surely (Gupta & Ramdas, 2021). However, this ideal state is

rarely achieved in practice (Guo et al., 2017). To quantify the deviation from this perfect calibration, the true calibration error (TCE) (Naeini et al., 2015; Gupta & Ramdas, 2021; Roelofs et al., 2022; Futami & Fujisawa, 2024) is used as a calibration metric.

2.2.1 True Calibration Error (TCE)

The TCE of f on $P(X, Y)$ is defined as

$$\text{TCE}(f) = \mathbb{E}_X [|\mathbb{E}[Y | f(X)] - f(X)|].$$

This measures the average absolute difference between the true conditional expectation of Y given $f(X)$ and the confidence score $f(X)$. It quantifies how well the confidence scores align with the true probabilities of the positive class. A lower TCE indicates better calibration, meaning that the confidence scores are closer to the true probabilities. However, calculation of $\text{TCE}(f)$ requires access to the conditional distribution $P(Y | f(X))$, which is usually unavailable.

2.2.2 Expected Calibration Error (ECE)

Instead, to estimate $\text{TCE}(f)$, we often use a binning method (Zadrozny & Elkan, 2001; Guo et al., 2017; Futami & Fujisawa, 2024). This method involves a set of $B \in \mathbb{N}$ bins (intervals), denoted by $\mathcal{I} = \{I_b\}_{b=1}^B$, that covers the entire range of f . We define the true binned function as

$$f_{\mathcal{I}}(X) = \sum_{b=1}^B \mathbb{E}[f(X) | f(X) \in I_b] \mathbb{1}_{f(X) \in I_b}.$$

Then its TCE is calculated as follows (Futami & Fujisawa, 2024):

$$\text{TCE}(f_{\mathcal{I}}) := \sum_{b=1}^B P(f(X) \in I_b) |\mathbb{E}[Y | f(X) \in I_b] - \mathbb{E}[f(X) | f(X) \in I_b]|.$$

$\text{TCE}(f_{\mathcal{I}})$ can be estimated from a supervised evaluation dataset $S_e := \{(X_m^e, Y_m^e)\}_{m=1}^{n_e}$, which can be a validation dataset or a test dataset, by the plug-in method. The resulting empirical estimate is called the expected calibration error (ECE):

$$\text{ECE}(f, S_e) := \sum_{b=1}^B p_b |\bar{y}_{b, S_e} - \bar{f}_{b, S_e}|, \quad (1)$$

where $p_b := \frac{\sum_{m=1}^{n_e} \mathbb{1}_{f(X_m^e) \in I_b}}{n_e}$ is the fraction of samples in bin b , $\bar{y}_{b, S_e} := \frac{\sum_{m=1}^{n_e} \mathbb{1}_{f(X_m^e) \in I_b} Y_m^e}{\sum_{m=1}^{n_e} \mathbb{1}_{f(X_m^e) \in I_b}}$ is the average true label in bin b , and $\bar{f}_{b, S_e} := \frac{\sum_{m=1}^{n_e} \mathbb{1}_{f(X_m^e) \in I_b} f(X_m^e)}{\sum_{m=1}^{n_e} \mathbb{1}_{f(X_m^e) \in I_b}}$ is the average confidence score in bin b . We assume $n \geq 2B$.

For \mathcal{I} , there are two common choices: uniform width binning (UWB) (Guo et al., 2017) and uniform mass binning (UMB) (Zadrozny & Elkan, 2001). In UWB, the bins are defined as $I_b = (\frac{b-1}{B}, \frac{b}{B}]$ for $b \in [B]$. In UMB, the bins are defined as $I_b = (u_{b-1}, u_b]$ for $b \in [B]$, where $u_0 = 0$, $u_b = f_{(k_b)}$ for $b \in [B-1]$, and $u_B = 1$. Here, $k_b = \lfloor \frac{n_e b}{B} \rfloor$, where $\lfloor \cdot \rfloor$ denotes the floor function that returns the greatest integer less than or equal to the input value, and $f_{(k)}$ denotes the k -th order statistic of the set $\{f(X_m^e)\}_{m=1}^{n_e}$.

There are extensions of the TCE and ECE, such as L_p -TCE and L_p -ECE (Kumar et al., 2019; Wenger et al., 2020; Popordanoska et al., 2022), which generalize the TCE and ECE by using the p -th power of the absolute difference. In this paper, we focus on the case of $p = 1$, which is the most commonly used case.

2.3 Positive-Unlabeled Learning

Positive-unlabeled (PU) learning is a classification problem where only positive and unlabeled data are observed, but no negative data. Among various settings of PU learning (Elkan & Noto, 2008; Natarajan

et al., 2013; du Plessis et al., 2014; Kato et al., 2019; Bekker et al., 2020), we focus on the selected-completely-at-random (SCAR) and case-control scenario. Concretely, we observe two independent sample sets: a set of unlabeled samples $S_{\text{tr},\text{U}} := \{X_m^{\text{tr},\text{U}}\}_{m=1}^{n_{\text{tr},\text{U}}} \stackrel{\text{i.i.d.}}{\sim} P(X)$ and a set of positive samples $S_{\text{tr},\text{P}} := \{X_m^{\text{tr},\text{P}}\}_{m=1}^{n_{\text{tr},\text{P}}} \stackrel{\text{i.i.d.}}{\sim} P(X | Y = 1)$, with the full training data denoted by $S_{\text{tr}} = (S_{\text{tr},\text{U}}, S_{\text{tr},\text{P}})$ (du Plessis et al., 2014). This SCAR and case-control scenario is commonly used in the PU learning literature (Natarajan et al., 2013; du Plessis et al., 2014; Kiryo et al., 2017; Yao et al., 2021). Throughout the paper, we assume that the class prior $\pi_{\text{P}} = P(Y = 1)$ is known—in practice, it can be estimated from data (Blanchard et al., 2010; Garg et al., 2021; Zhu et al., 2023).

In PU learning, an unbiased risk estimator of h on S_{tr} is given as follows (Natarajan et al., 2013; du Plessis et al., 2014; 2015):

$$\mathcal{L}_{\text{uPU}}(h) = \frac{\pi_{\text{P}}}{n_{\text{tr},\text{P}}} \sum_{m=1}^{n_{\text{tr},\text{P}}} \{l(h(X_m^{\text{tr},\text{P}}), 1) - l(h(X_m^{\text{tr},\text{P}}), 0)\} + \frac{1}{n_{\text{tr},\text{U}}} \sum_{m=1}^{n_{\text{tr},\text{U}}} l(h(X_m^{\text{tr},\text{U}}), 0). \quad (2)$$

The estimator is suitable for estimating the expected risk $\mathbb{E}[l(h(X), Y)]$ or training simple classifiers (du Plessis et al., 2014; 2015), but it does not work well for training more complex classifiers such as deep neural networks due to the *negative risk* issue (Kiryo et al., 2017). To overcome this issue, non-negative risk estimators were proposed (Kiryo et al., 2017; Lu et al., 2020). The non-negative risk estimator is given by

$$\mathcal{L}_{\text{nnPU}}(h) = \frac{\pi_{\text{P}}}{n_{\text{tr},\text{P}}} \sum_{m=1}^{n_{\text{tr},\text{P}}} l(h(X_m^{\text{tr},\text{P}}), 1) + l_{\text{cc}} \left(\frac{1}{n_{\text{tr},\text{U}}} \sum_{m=1}^{n_{\text{tr},\text{U}}} l(h(X_m^{\text{tr},\text{U}}), 0) - \frac{\pi_{\text{P}}}{n_{\text{tr},\text{P}}} \sum_{m=1}^{n_{\text{tr},\text{P}}} l(h(X_m^{\text{tr},\text{P}}), 0) \right). \quad (3)$$

where $l_{\text{cc}} : \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}$ is called a consistent-correction function. Several functions have been proposed for use as l_{cc} , e.g., $\max\{x, 0\}$ (Kiryo et al., 2017; Lu et al., 2020).

3 PU-ECE: An Estimator for Calibration Error in PU Learning

Conventional ECE estimation, as shown in Eq. (1), requires labeled data (both positive and negative), but in the PU environment, negative data are unavailable, as introduced in Section 2.3. The main challenges are (i) estimating the true positive rate within bins (ii) the data are not from a single distribution $P(X, Y)$ but from $P(X)$ and $P(X|Y = 1)$. In this section, we propose a method to estimate the ECE in the PU environment. The proposed method is based on the idea of using positive and unlabeled data to estimate $\mathbb{E}[Y | f(X) \in I_b]$. We denote the PU evaluation data as $S_{\text{U}} := \{X_m^{\text{U}}\}_{m=1}^{n_{\text{U}}}$, $S_{\text{P}} := \{X_m^{\text{P}}\}_{m=1}^{n_{\text{P}}}$, and $S_{\text{PU}} := (S_{\text{U}}, S_{\text{P}})$. The evaluation data are independent of the training data and f , and can be a validation dataset or a test dataset in practice.

3.1 Proposed Estimator

Since $Y \in \{0, 1\}$, the conditional expectation can be expressed as

$$\mathbb{E}[Y | f(X) \in I_b] = P(Y = 1 | f(X) \in I_b) = \frac{P(Y = 1, f(X) \in I_b)}{P(f(X) \in I_b)} = \frac{\pi_{\text{P}} P(f(X) \in I_b | Y = 1)}{P(f(X) \in I_b)}. \quad (4)$$

In the PU learning setting, we have access to positive samples $S_{\text{P}} = \{X_m^{\text{P}}\}_{m=1}^{n_{\text{P}}}$ drawn i.i.d. from $P(X|Y = 1)$ and unlabeled samples $S_{\text{U}} = \{X_m^{\text{U}}\}_{m=1}^{n_{\text{U}}}$ drawn i.i.d. from $P(X)$. The term $P(f(X) \in I_b | Y = 1)$ in Eq. (4) can be estimated by the empirical frequency of positive samples falling into bin I_b : $p_{b,\text{P}} := \frac{1}{n_{\text{P}}} \sum_{m=1}^{n_{\text{P}}} \mathbb{1}_{f(X_m^{\text{P}}) \in I_b}$. Similarly, $P(f(X) \in I_b)$ can be estimated by the empirical frequency of unlabeled samples falling into bin I_b : $p_{b,\text{U}} := \frac{1}{n_{\text{U}}} \sum_{m=1}^{n_{\text{U}}} \mathbb{1}_{f(X_m^{\text{U}}) \in I_b}$. These estimators are unbiased, that is, $\mathbb{E}[p_{b,\text{P}}] = P(f(X) \in I_b | Y = 1)$ and $\mathbb{E}[p_{b,\text{U}}] = P(f(X) \in I_b)$. Let $|I_{b,\text{P}}| = \sum_{m=1}^{n_{\text{P}}} \mathbb{1}_{f(X_m^{\text{P}}) \in I_b}$ and $|I_{b,\text{U}}| = \sum_{m=1}^{n_{\text{U}}} \mathbb{1}_{f(X_m^{\text{U}}) \in I_b}$ be the number of positive and unlabeled samples in bin I_b , respectively. Then $p_{b,\text{P}} = |I_{b,\text{P}}|/n_{\text{P}}$ and $p_{b,\text{U}} = |I_{b,\text{U}}|/n_{\text{U}}$. Using these estimates, an estimator for the conditional probability

$P(Y = 1 \mid f(X) \in I_b)$ is naively obtained as $\hat{y}_{b, S_{PU}} := \frac{\pi_P p_{b, P}}{p_{b, U}}$. The ECE for PU learning (PU-ECE) can then be estimated as

$$\sum_{b=1}^B p_{b, U} |\hat{y}_{b, S_{PU}} - \bar{f}_{b, S_{PU}}|,$$

where $\bar{f}_{b, S_{PU}} := \frac{\sum_{m=1}^{n_U} \mathbb{1}_{f(X_m^U) \in I_b} f(X_m^U)}{|I_{b, U}|}$ is the average confidence score for unlabeled samples in bin I_b . If $|I_{b, U}| = 0$, then $p_{b, U} = 0$. In this case, $\hat{y}_{b, S_{PU}}$ would involve division by zero. However, this is a removable singularity. For all cases where $p_{b, U} > 0$, we can algebraically reformulate the estimator as follows:

$$\text{ECE}_{PU}(f, S_{PU}) = \sum_{b=1}^B \left| \frac{\pi_P}{n_P} \sum_{m=1}^{n_P} \mathbb{1}_{f(X_m^P) \in I_b} - \frac{1}{n_U} \sum_{m=1}^{n_U} \mathbb{1}_{f(X_m^U) \in I_b} f(X_m^U) \right|.$$

This new form is mathematically equivalent for all $p_{b, U} > 0$, and remains well-defined at $p_{b, U} = 0$, and its value at $p_{b, U} = 0$ coincides with the limit of the estimator as $p_{b, U} \rightarrow 0$.

We define UWB and UMB for the PU setting as follows: in UWB, the bins are defined in the same way as in supervised learning, while in UMB, the bins are defined by *the unlabeled data*; $I_b = (u_{b-1}, u_b]$ for $b \in [B]$, where $u_0 = 0$, $u_b = f_{(k_b)}$ for $b \in [B-1]$, and $u_B = 1$. Here, $k_b = \lfloor \frac{n_U b}{B} \rfloor$, and $f_{(k)}$ denotes the k -th order statistic of the set $\{f(X_m^U)\}_{m=1}^{n_U}$, i.e., the k -th smallest value.

3.2 Bias Analysis of PU-ECE

We analyze the bias of the proposed estimator given the evaluation data. In the analysis, we assume the following:

Assumption 1. *The number of positive samples n_P satisfies $n_P \geq 1$. The number of unlabeled samples n_U satisfies $n_U \geq 2B$ for UMB and $n_U \geq 1$ for UWB.*

Assumption 2. *$f(X)$ is absolutely continuous with respect to the Lebesgue measure, i.e., $f(X)$ has a probability density function.*

Assumption 3. *$\mathbb{E}[Y \mid f(X)]$ satisfies L -Lipschitz continuity with respect to $f(X)$, i.e., $|\mathbb{E}[Y \mid f(X) = v] - \mathbb{E}[Y \mid f(X) = v']| \leq L|v - v'|$ for all $v, v' \in [0, 1]$.*

The (total) bias of the PU-ECE estimator is defined as

$$\text{Bias}_{\text{tot}}(f, S_{PU}) := |\text{ECE}_{PU}(f, S_{PU}) - \text{TCE}(f)|.$$

Then, we have the following theorem, which provides non-asymptotic bias bounds for the PU-ECE estimator.

Theorem 1. *Under Assumptions 1, 2 and 3, we have*

$$\mathbb{E}_{S_{PU}} [\text{Bias}_{\text{tot}}(f, S_{PU})] \leq \begin{cases} \frac{1+L}{B} + \sqrt{2 \left(\frac{\pi_P^2}{n_P} + \frac{1}{n_U} \right) B \log 2} & (\text{for UWB}), \\ \frac{1+L}{B} + \frac{(3+2L)B}{n_U - B} + \sqrt{2B \log 2} \left(\frac{1+L}{\sqrt{n_U - B}} + \sqrt{\frac{\pi_P^2}{n_P} + \frac{1}{n_U - B}} \right) & (\text{for UMB}). \end{cases}$$

Furthermore, for any $\delta \in (0, 1)$, with probability at least $1 - \delta$ over S_{PU} , we have

$$\text{Bias}_{\text{tot}}(f, S_{PU}) \leq \begin{cases} \frac{1+L}{B} + \sqrt{2 \left(\frac{\pi_P^2}{n_P} + \frac{1}{n_U} \right) (B \log 2 + \log(1/\delta))} & (\text{for UWB}), \\ \frac{1+L}{B} + \frac{(3+2L)B}{n_U - B} + \sqrt{2(B \log 2 + \log(1/\delta))} \left(\frac{1+L}{\sqrt{n_U - B}} + \sqrt{\frac{\pi_P^2}{n_P} + \frac{1}{n_U - B}} \right) & (\text{for UMB}). \end{cases}$$

The proof of Theorem 1 is provided in Appendix A.1. Briefly, the proof is based on decomposing the total bias into two parts: the bias due to binning and the bias due to the estimation by the empirical frequency, and then bounding each part separately by concentration inequalities.

Minimizing the upper bound of the total bias with respect to B , we have the following upper bound orders:

Corollary 1. *Under Assumption 1, the scale of the optimal B that minimizes the upper bound in Theorem 1 for both UWB and UMB is given by*

$$B = \Theta \left(\left(\frac{\pi_P^2}{n_P} + \frac{1}{n_U} \right)^{-\frac{1}{3}} \right),$$

and the corresponding expected total bias convergence rate is given by

$$\mathbb{E}_{S_{PU}} [\text{Bias}_{\text{tot}}(f, S_{PU})] = \mathcal{O} \left(\left(\frac{\pi_P^2}{n_P} + \frac{1}{n_U} \right)^{\frac{1}{3}} \right).$$

The proof can be found in Appendix A.2. In practice, this means that to minimize the bias of the PU-ECE estimator, the number of bins B should be chosen according to the formula above, balancing the sample sizes of the positive and unlabeled data. The result corresponds to the convergence rate of the supervised setting; the optimal B is $\mathcal{O}(n^{\frac{1}{3}})$ and the expected total bias is $\mathcal{O}(n^{-\frac{1}{3}})$, where n is the size of the fully-labeled evaluation data (Futami & Fujisawa, 2024). The term π_P^2 in these expressions is a constant factor that does not depend on the sample size, but it indicates that a smaller class prior π_P leads to a smaller expected bias.

In practical PU learning scenarios, the true class prior π_P is often unknown and must be estimated from data (Bepler et al., 2019; Ito & Sugiyama, 2023), which introduces additional estimation error into the calibration process. It is therefore important to understand how inaccuracies in class prior estimation affect PU-ECE. We also analyze how the class prior estimation error affects the PU-ECE estimation.

Theorem 2 (Error bound due to class prior estimation error). *Let $\hat{\pi}_P$ be an estimate of the class prior π_P , and define the PU-ECE estimator using $\hat{\pi}_P$ as*

$$\widehat{\text{ECE}}_{PU}(f, S_{PU}) = \sum_{b=1}^B \left| \frac{\hat{\pi}_P}{n_P} \sum_{m=1}^{n_P} \mathbb{1}_{f(X_m^P) \in I_b} - \frac{1}{n_U} \sum_{m=1}^{n_U} f(X_m^U) \mathbb{1}_{f(X_m^U) \in I_b} \right|.$$

Then, we have the following bound on the difference between the PU-ECE with the true prior and that with the estimated prior:

$$\left| \text{ECE}_{PU}(f, S_{PU}) - \widehat{\text{ECE}}_{PU}(f, S_{PU}) \right| \leq |\pi_P - \hat{\pi}_P|.$$

The proof is provided in Appendix A.3. This inequality does not require expectation and holds for any f and S_{PU} . This theorem shows that when the class prior is unknown and must be estimated, the calibration error estimation incurs an additional error of at most the class prior estimation error.

4 Information-Theoretic Generalization Analysis of PU-ECE

The previous analysis focused on the bias of the ECE estimation using evaluation data, which is independent of the training data used for learning f_w . However, in practice, we may recalibrate classifiers using a recalibration dataset and then evaluate the PU-ECE on the same recalibration dataset (Kumar et al., 2019), or we may calibrate classifiers using the training dataset (Futami & Fujisawa, 2024). In these cases, the data used for PU-ECE evaluation and the classifier recalibration/training are no longer independent. Therefore, analyzing the generalization error of the PU-ECE is crucial to ensure reliable model selection and evaluation in practice. In this section, we extend the conditional mutual information (CMI) framework (Steinke & Zakynthinou, 2020; Haghifam et al., 2021) used in supervised learning to the PU setting and analyze the generalization error of PU-ECE using it.

4.1 Information-Theoretic Generalization Error Analysis for Supervised Learning

A rigorous understanding of the generalization error is crucial for evaluating the reliability of machine learning algorithms. Classical paradigms for assessing generalization include the concepts of Vapnik-Chervonenkis (VC) dimension (Vapnik & Chervonenkis, 1968; Blumer et al., 1989) and metric entropy (Kolmogorov &

(Tikhomirov, 1959; Dudley, 1974). The VC dimension quantifies how many points a hypothesis class can shatter, while metric entropy measures how finely a function class can be approximated using a finite set of representative functions at a given precision. Despite the established significance of these tools, recent advances in information-theoretic perspectives, particularly those centered on CMI (Steinke & Zakynthinou, 2020; Haghifam et al., 2021), have provided new insights into generalization. Intuitively, CMI quantifies the amount of information that a model retains about the training data after observing a test sample. This offers an alternative view for analyzing the generalization of a learned model. By bounding this data-model dependency, CMI-based analysis methods provide new insights into why certain learning algorithms exhibit strong generalization performance. For instance, Futami & Fujisawa (2024) applied this approach to derive non-asymptotic generalization error bounds for the ECE in the supervised setting.

4.2 Conditional Mutual Information (CMI) for PU Learning

To analyze the generalization error, we extend the conditional mutual information (CMI) framework to the PU setting. We define the functional conditional mutual information (fCMI) for PU data, denoted by fCMI , which measures the information between the model’s confidence score and the PU training data. We also define a variant, fCMI_U , which considers only the unlabeled training data. The formal definitions of these quantities are technical and are detailed in Appendix B. By the data processing inequality (Cover & Thomas, 2006), we have $\text{fCMI}_U \leq \text{fCMI}$. The PU CMI metrics can be bounded by data-independent model complexity measures. For instance, if the model has a VC dimension d_{VC} , $\text{fCMI} \leq \max\{(d_{VC} + 1) \log 2, d_{VC} \log(2e(n_{tr,U} + n_{tr,P})/d_{VC})\}$ (Steinke & Zakynthinou, 2020; Harutyunyan et al., 2021). If $f_{\mathcal{A}(\tilde{X}_M, R)}$ is a d -dimensional parametric function that is L -Lipschitz continuous on $\|\cdot\|_\infty$ and whose domain is $[0, 1]^d$ and the function class has an ϵ -cover, then $\text{fCMI} \lesssim d \log(L/\epsilon)$ and $\text{fCMI}_U \lesssim d \log(L/\epsilon)$ (Wainwright, 2019; Futami & Fujisawa, 2024).

4.3 Generalization Bound for PU-ECE

This subsection analyzes the generalization error of PU-ECE using the fCMIs defined above. We follow Futami & Fujisawa (2024) and use $\text{Bias}_{\text{tot}}(f_W, S_{tr}) := |\text{TCE}(f_W) - \text{ECE}_{PU}(f_W, S_{tr})|$ as the generalization error of $\text{ECE}_{PU}(f_W, S_{tr})$, where we denote by $W = \mathcal{A}(S_{tr}, R) \in \mathcal{W}$ the parameters of f_W learned by a randomized algorithm \mathcal{A} on the training data S_{tr} and the source of randomness R .

Theorem 3 (Expected generalization error bound of PU-ECE). *Under Assumptions 2 and 3, we have*

$$\begin{aligned} & \mathbb{E}_{R, S_{tr}} [\text{Bias}_{\text{tot}}(f_W, S_{tr})] \\ & \leq \begin{cases} \frac{1+L}{B} + \sqrt{8 \left(\frac{\pi_P^2}{n_{tr,P}} + \frac{1}{n_{tr,U}} \right) \{\text{fCMI} + B \log 2\}} & (\text{for UWB}), \\ \frac{1+L}{B} + \sqrt{8 \left(\frac{\pi_P^2}{n_{tr,P}} + \frac{1}{n_{tr,U}} \right) \{\text{fCMI} + B \log 2\}} + (1+L) \sqrt{\frac{8}{n_{tr,U}} \{\text{fCMI}_U + B \log 2\}} & (\text{for UMB}). \end{cases} \end{aligned}$$

The proof of Theorem 3 is provided in Appendix C.1. Briefly, the proof is similar to the proof of Theorem 1, except the dependency on the training data is handled by fCMI.

Similar to Corollary 1, the optimal bin size to minimize the upper bound in Theorem 3 and the order of the upper bound can also be derived if a learning algorithm has a sufficiently small fCMI. For instance, such an assumption holds when \mathcal{H} is a finite VC class since \mathcal{F} is also a finite VC class, and hence fCMI is bounded by $\mathcal{O}(\log(n_{tr,P} + n_{tr,U}))$ (Steinke & Zakynthinou, 2020). For such a case, we can derive the following result.

Corollary 2 (Total bias upper bound order). *Under Assumption 2 and 3 and assuming $\text{fCMI} = \mathcal{O}\left(\left(\frac{\pi_P^2}{n_{tr,P}} + \frac{1}{n_{tr,U}}\right)^{-1/3}\right)$, the scale of the optimal B that minimizes the upper bound in Theorem 3 is given by*

$$B = \Theta \left(\left(\frac{\pi_P^2}{n_{tr,P}} + \frac{1}{n_{tr,U}} \right)^{-1/3} \right),$$

Dataset	$n_{\text{tr},U}$	π_P	Model	Learning rate
MNIST (LeCun et al., 1998)	60,000	0.49	2-layer MLP (784-100-100-1)	1.0×10^{-3}
CIFAR-10 (Krizhevsky, 2009)	50,000	0.40	ResNet-18	1.0×10^{-5}

Table 1: Specification of benchmark datasets, models, and learning rates. $n_{\text{tr},U}$ denotes the number of unlabeled samples for training PU classifiers.

and the corresponding expected generalization error convergence rate is given by

$$\mathbb{E}_{R,S_{\text{tr}}} [\text{Bias}_{\text{tot}}(f_W, S_{\text{tr}})] = \mathcal{O} \left(\left(\frac{\pi_P^2}{n_{\text{tr},P}} + \frac{1}{n_{\text{tr},U}} \right)^{1/3} \right).$$

The proof is provided in Appendix A.2.

By Corollaries 1 and 2, $\mathbb{E}_{R,S_{\text{tr}}} [\text{Bias}_{\text{tot}}(f_W, S_{\text{tr}})] \ll \mathbb{E}_{S_{\text{PU}}} [\text{Bias}_{\text{tot}}(f_W, S_{\text{PU}})]$ if $n_{\text{tr},P} \gg n_P$ and $n_{\text{tr},U} \gg n_U$ and the classifier generalizes well. This would actually happen in practice since the training dataset is often much larger than the evaluation dataset. Such a property is also observed in the supervised setting (Futami & Fujisawa, 2024). This represents an important future direction for PU learning to develop a method that jointly optimizes the TCE and accuracy, although this is beyond the scope of this paper.

5 Experiments

In this section, we empirically evaluate the theoretical properties of the proposed PU-ECE estimator. We first present a motivational study that highlights the unique challenges of calibration in PU learning (Section 5.2). We then empirically validate the theoretical bias and convergence properties of our proposed PU-ECE estimator on both synthetic and real-world benchmark datasets (Section 5.3).

5.1 Common Setup

For all experiments, we used UMB since UMB has demonstrated smaller biases than UWB (Roelofs et al., 2022). The code for all experiments was implemented in Python and is available at <https://github.com/will-be-released>.

Synthetic data Following previous work (Vaicenavicius et al., 2019; Zhang et al., 2020; Futami & Fujisawa, 2024), we generated data from a logistic model where $P(Y = 1 \mid X = x) = \sigma(2x)$, with the conditional probability densities $p(X \mid Y = 1) = \mathcal{N}(x; 1, 1)$ and $p(X \mid Y = 0) = \mathcal{N}(x; -1, 1)$. We evaluated a classifier $f(x) = \sigma(\beta_0 + \beta_1 x)$ under two settings: a less-calibrated case ($\beta_0 = -0.5, \beta_1 = 1.5$) and a better-calibrated case ($\beta_0 = -0.2, \beta_1 = 1.9$). The ground-truth TCE(f) was computed by the trapezoidal rule. Further details are provided in Appendix D.

Benchmark data We used the MNIST and CIFAR-10 datasets, with specifications detailed in Table 1. Since these datasets are multiclass, we converted them into binary classification tasks by selecting a positive class and treating all other classes as negative. For MNIST, we used the digits 5-9 as the positive class, while for CIFAR-10, we selected the vehicle class (airplane, automobile, ship, and truck) as the positive class. Classifiers were trained using the standard ERM for supervised learning and the nnPU learning method (Kiryo et al., 2017) for PU learning, and the outputs of the classifiers were used as logit scores. The logistic loss (a.k.a. binary cross-entropy loss) $\log(\sigma(-(2y - 1)f(x)))$ or sigmoid loss $\sigma(-(2y - 1)f(x))$ was used as the loss function, depending on the experiment. The logistic loss was used as a strictly proper loss function in supervised learning (Zadrozny & Elkan, 2001). The sigmoid loss was used as a noise-robust loss function (Ghosh et al., 2017) and classification calibrated (Charoenphakdee et al., 2019), and hence are widely used in weakly-supervised learning including PU learning (Kiryo et al., 2017; Chen et al., 2020; Zhao et al., 2022; Ye et al., 2023). Although the sigmoid loss is not strictly proper (Charoenphakdee et al., 2019), its confidence scores are conventionally interpreted as class-posterior probabilities (Platt, 1999; Chen et al., 2020; Ye et al., 2023). The other details are provided in Appendix F.

5.2 Motivational Study

We first conducted a controlled study to compare the ECE of classifiers trained from fully-supervised data and PU data to see how calibration evolves with the number of training samples. MNIST was used as the dataset. We evaluated the classifiers trained under four conditions: supervised learning with logistic loss (PN-logistic), supervised learning with sigmoid loss (PN-sigmoid), PU learning trained with the logistic loss (PU-logistic), and PU learning with the sigmoid loss (PU-sigmoid). Results are averaged over 100 trials, each with a different random seed. The number of unlabeled samples is fixed at 60,000.

As shown in Figure 1, supervised methods improve calibration steadily with more data regardless of the loss function. As expected, the logistic loss leads to better calibration than the sigmoid loss in supervised learning. In contrast, PU learning behaves very differently: PU-sigmoid is better calibrated than PU-logistic.

This collapse is likely due to the non-negative correction in the nnPU risk estimator (Eq. (3)). When the logistic loss is used, the nnPU risk estimator tends to cram the logit scores of positive samples into a small constant value so that the non-negative correction has little effect for positive samples, which leads to poor calibration. As shown in Figure 2, when trained with the logistic loss, the logit scores for positive examples collapse toward a constant, yielding high ECE. These findings demonstrate that, in practice, PU learning is considerably more prone to miscalibration than standard supervised learning, especially when using the logistic loss, which is strictly proper in supervised learning.

5.3 Validation of Bias and Convergence

Let N denote an integer ranging from 10 to 10,000. We evaluated the convergence behavior by varying N , comparing PU-ECE in settings with $(n_P = N, n_U = \infty)$, $(n_P = \infty, n_U = N)$, $(n_P = N, n_U = 10N)$, and standard ECE with $n = N$. For settings with either $n_U = \infty$ or $n_P = \infty$, the corresponding components (e.g., $P(f(X) \in I_b)$ when $n_U = \infty$) were computed via numerical integration using the trapezoidal rule with up to 10,000 points.

We also investigated the effect of fixing n_P by comparing PU-ECE with $(n_P = 100, n_U = N)$, $(n_P = 1,000, n_U = N)$, $(n_P = 10,000, n_U = N)$, and $(n_P = N, n_U = \infty)$. The number of bins B was set to $\lceil (\pi_P^2/n_P + 1/n_U)^{-1/3} \rceil$ for PU-ECE and $n^{1/3}$ or $(n/\pi_P^2)^{1/3}$ for ECE, following the optimal order derived in our analysis and in Futami & Fujisawa (2024).

We trained the classifiers on the PU training data with 10,000 positive and $n_{tr,U}$ unlabeled samples specified in Table 1. Results are averaged over 100 trials, each with a different random seed.

To compute $TCE(f)$ for these datasets, we first need to assume parametric models for the true conditional probability densities. Fortunately, the TCE can be computed using the distribution of the logit score $h(X)$ or the confidence score $f(X)$, without modeling the distribution of X , which would otherwise suffer from the curse of dimensionality (see Appendix E for details). We observed that the confidence scores were heavily saturated near 0 or 1, while the logit scores were spread out (Figure 2 and 3). Therefore, we modeled the logit score $h(X)$ instead of the confidence score $f(X)$, specifically modeling the distributions of the logit score, $P(h(X) | Y = 1)$ and $P(h(X) | Y = 0)$. The empirical logit histograms were multi-modal, so a single-mode parametric family, such as the logistic-beta distribution used by Roelofs et al. (2022), is inadequate. For instance, the MNIST logits (Fig. 2) exhibited two clear peaks. Instead, we fitted Gaussian Mixture Models

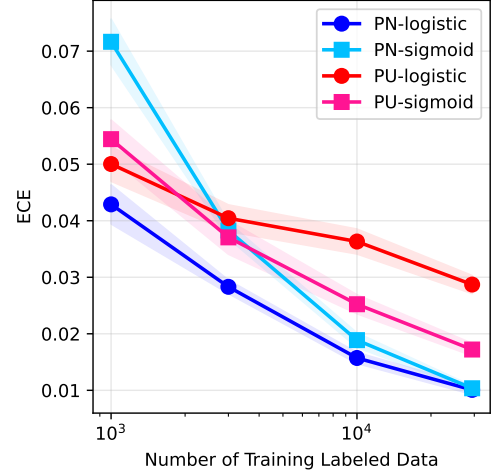


Figure 1: Comparison of ECE for supervised and PU learning on MNIST. The x-axis represents the number of labeled samples used for training. For PU learning, this corresponds to the number of positive samples. The y-axis is the ECE on the test set. The lines represent the ECE values and the shaded areas represent the standard error.

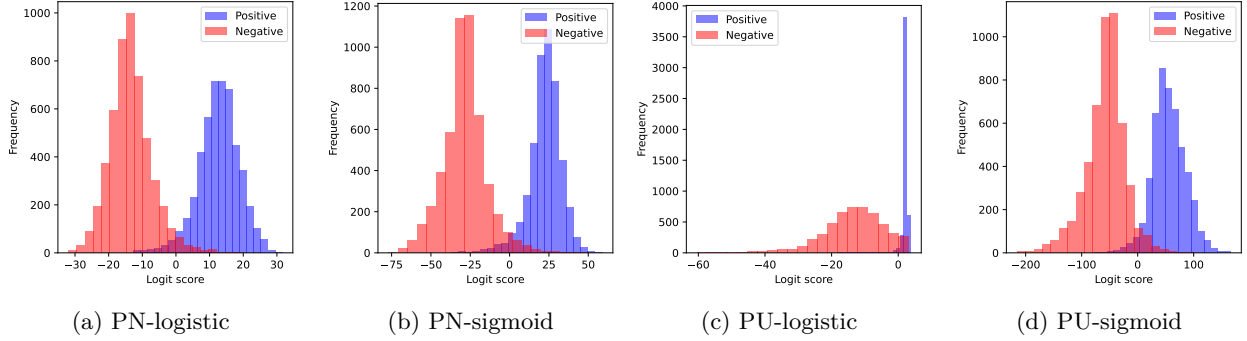


Figure 2: Histogram of logit scores for different training settings on MNIST test data. The x-axis is the logit and the y-axis is the frequency.

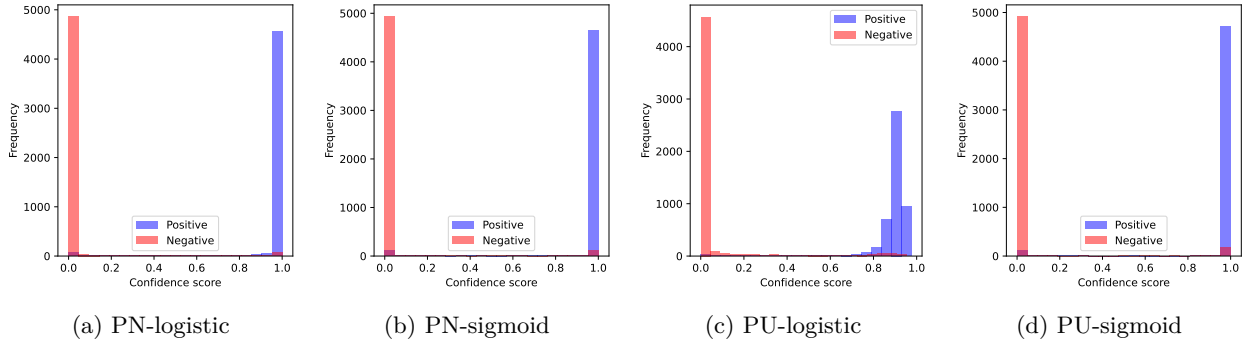


Figure 3: Histogram of the confidence scores for different training settings on MNIST test data. The x-axis is the confidence score and the y-axis is the frequency.

(GMMs) to $P(h(X) | Y = 1)$ and $P(h(X) | Y = 0)$, selecting the number of components for each GMM from 1 to 10 via the Bayesian Information Criterion (BIC) (Schwarz, 1978). In MNIST, the GMMs with 4 positive class components and 3 negative class components were selected, while in CIFAR-10, the GMMs with 2 positive class components and 2 negative class components were selected. With these estimated densities, we sampled the logit scores from the GMMs. $\text{TCE}(f)$ was calculated using the trapezoidal rule at 10,000 points.

Figure 4 shows the total bias of PU-ECE and standard ECE (absolute bias compared to the TCE) on both synthetic and benchmark datasets. This experiment is designed to validate the theoretical convergence rate of our estimator. According to our analysis, the total bias of PU-ECE is expected to decrease at a rate of $\mathcal{O}((\pi_P^2/n_P + 1/n_U)^{-1/3})$. When both n_P and n_U grow proportionally with N (or are set to ∞), this simplifies to $\mathcal{O}(N^{-1/3})$.

The empirical results clearly exhibit this $\mathcal{O}(N^{-1/3})$ convergence, providing strong evidence in support of our theoretical claims. Notably, this rate matches the known optimal convergence rate for standard ECE (Futami & Fujisawa, 2024), demonstrating that PU-ECE can achieve competitive calibration performance even without access to fully labeled data.

Figure 5 investigates the effect of fixing n_P while increasing n_U . We observe that PU-ECE converges to the limiting case of PU-ECE with $n_P = N, n_U = \infty$ but not to the true TCE. This behavior is consistent with our theoretical prediction: even as $n_U \rightarrow \infty$, the bias remains bounded below by the term depending on n_P . This residual error creates an error floor that prevents convergence to the TCE when n_P is finite.

These results highlight a key practical insight: sufficiently large sets of both positive and unlabeled samples are therefore essential for effectively estimating the TCE.

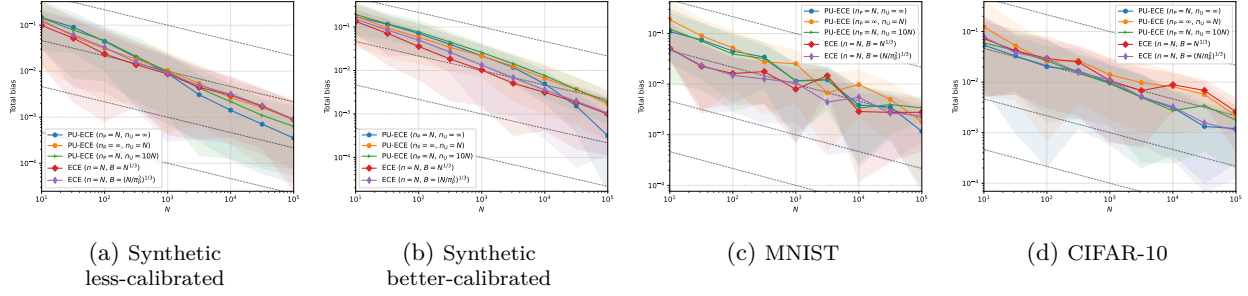


Figure 4: Total bias of ECE and PU-ECE, as a function of sample size N . Both axes are log-scale. The lines represent the total bias, and the shaded areas represent the 90% percentile interval. The dashed lines represent the theoretical convergence rate of $\mathcal{O}(N^{-1/3})$.

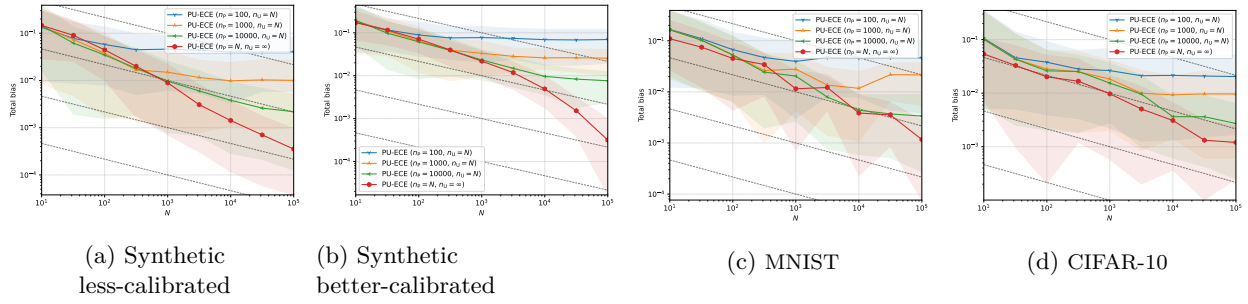


Figure 5: Convergence of total bias of PU-ECE varying n_P . Both axes are log-scale. The lines represent the total bias, and the shaded areas represent the 90% percentile interval. The dashed lines represent the theoretical convergence rate of $\mathcal{O}(N^{-1/3})$.

6 Conclusions

We proposed PU-ECE, the first ECE estimator that can be computed from PU data. Our theoretical analysis provides non-asymptotic bias bounds for PU-ECE and shows that, with an optimal binning strategy, the bias of our estimator converges to the TCE at a rate of $\mathcal{O}((\pi_P^2/n_P + 1/n_U)^{1/3})$. This rate matches the convergence rate of the standard ECE in the supervised setting. We further extended the CMI framework to PU learning, deriving generalization error bounds with the same convergence rate as the bias. Experiments on synthetic and benchmark datasets corroborated the theory: PU-ECE performs comparably to the standard ECE, despite lacking negative labels. They also revealed a practical caveat—if the number of positive samples is fixed and the number of unlabeled samples increases, PU-ECE does not converge to the true TCE, which highlights the importance of having sufficient positive data as well as unlabeled data for accurate estimation.

References

- Kendall Atkinson. *An introduction to numerical analysis*. John Wiley & Sons, New York, 2 edition, January 1989.
- Alejandro Correa Bahnsen, Aleksandar Stojanovic, Djamila Aouada, and Björn Ottersten. Improving credit card fraud detection with calibrated probabilities. In *Proceedings of the 2014 SIAM International Conference on Data Mining*, pp. 677–685, Philadelphia, PA, April 2014. Society for Industrial and Applied Mathematics.
- Jessa Bekker, Pieter Robberechts, and Jesse Davis. Beyond the selected completely at random assumption for learning from positive and unlabeled data. In *Machine Learning and Knowledge Discovery in Databases*, pp. 71–85. Springer International Publishing, 2020.
- Tristan Bepler, Andrew Morin, Micah Rapp, Julia Brasch, Lawrence Shapiro, Alex J Noble, and Bonnie Berger. Positive-unlabeled convolutional neural networks for particle picking in cryo-electron micrographs. *Nature Methods*, 16(11):1153–1160, November 2019.
- Gilles Blanchard, Gyemin Lee, and Clayton Scott. Semi-supervised novelty detection. *Journal of Machine Learning Research*, 11(99):2973–3009, 2010.
- Anselm Blumer, Andrzej Ehrenfeucht, David Haussler, and Manfred K. Warmuth. Learnability and the vapnik-chervonenkis dimension. *Journal of the ACM*, 36(4):929–965, October 1989.
- Nontawat Charoenphakdee, Jongyeong Lee, and Masashi Sugiyama. On symmetric losses for learning from corrupted labels. In Kamalika Chaudhuri and Ruslan Salakhutdinov (eds.), *Proceedings of the 36th International Conference on Machine Learning*, volume 97, pp. 961–970. PMLR, January 2019.
- Xuxi Chen, Wuyang Chen, Tianlong Chen, Ye Yuan, Chen Gong, Kewei Chen, and Zhangyang Wang. Self-PU: Self boosted and calibrated positive-unlabeled training. In Hal Daumé Iii and Aarti Singh (eds.), *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 1510–1519. PMLR, July 2020.
- Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory: Cover/elements of information theory, second edition*. John Wiley & Sons, Nashville, TN, 2 edition, June 2006.
- Alexander P. Dawid. The well-calibrated bayesian. *Journal of the American Statistical Association*, 77(379): 605–610, 1982.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In Jill Burstein, Christy Doran, and Thamar Solorio (eds.), *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 4171–4186, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics.
- Monroe D. Donsker and Srinivasa R. S. Varadhan. On a variational formula for the principal eigenvalue for operators with maximum principle. *Proceedings of the National Academy of Sciences of the United States of America*, 72(3):780–783, 1975.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, May 2021.
- Marthinus C. du Plessis, Gang Niu, and Masashi Sugiyama. Analysis of learning from positive and unlabeled data. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Q. Weinberger (eds.), *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc., 2014.

- Marthinus C. du Plessis, Gang Niu, and Masashi Sugiyama. Convex formulation for learning from positive and unlabeled data. In Francis Bach and David Blei (eds.), *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pp. 1386–1394, Lille, France, 2015. PMLR.
- Richard M. Dudley. Metric entropy of some classes of sets with differentiable boundaries. *Journal of Approximation Theory*, 10(3):227–236, 1974.
- Charles Elkan and Keith Noto. Learning classifiers from only positive and unlabeled data. In *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 213–220, New York, NY, USA, August 2008. ACM.
- Futoshi Futami and Masahiro Fujisawa. Information-theoretic generalization analysis for expected calibration error. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (eds.), *Advances in Neural Information Processing Systems*, volume 37, pp. 84246–84297. Curran Associates, Inc., 2024.
- Saurabh Garg, Yifan Wu, Alex Smola, Sivaraman Balakrishnan, and Zachary Chase Lipton. Mixture proportion estimation and PU learning: a modern approach. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan (eds.), *Advances in Neural Information Processing Systems*, volume 34, pp. 8532–8544. Curran Associates, Inc., 2021.
- Aritra Ghosh, Himanshu Kumar, and P. S. Sastry. Robust loss functions under label noise for deep neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31, February 2017.
- Irving J. Good. Rational decisions. *Journal of the Royal Statistical Society. Series B (Methodological)*, 14(1):107–114, 1952.
- Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q. Weinberger. On calibration of modern neural networks. In Doina Precup and Yee Whye Teh (eds.), *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 1321–1330. PMLR, August 2017.
- Chirag Gupta and Aaditya Ramdas. Distribution-free calibration guarantees for histogram binning without sample splitting. In Marina Meila and Tong Zhang (eds.), *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 3942–3952. PMLR, July 2021.
- Mahdi Haghifam, Gintare Karolina Dziugaite, Shay Moran, and Dan Roy. Towards a unified information-theoretic framework for generalization. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan (eds.), *Advances in Neural Information Processing Systems*, volume 34, pp. 26370–26381. Curran Associates, Inc., 2021.
- Hrayr Harutyunyan, Maxim Raginsky, Greg Ver Steeg, and Aram Galstyan. Information-theoretic generalization bounds for black-box learning algorithms. In Marc’Aurelio Ranzato, Alina Beygelzimer, Yann Dauphin, Percy S. Liang, and Jennifer Wortman Vaughan (eds.), *Advances in Neural Information Processing Systems*, pp. 24670–24682. Curran Associates, Inc., 2021.
- Fredrik Hellström and Giuseppe Durisi. A new family of generalization bounds using samplewise evaluated cmi. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (eds.), *Advances in Neural Information Processing Systems*, volume 35, pp. 10108–10121. Curran Associates, Inc., 2022.
- Nobutaka Ito and Masashi Sugiyama. Audio signal enhancement with learning from positive and unlabeled data. In *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing*, pp. 1–5, 2023.
- Kristen Jaskie and Andreas Spanias. Positive and unlabeled learning algorithms and applications: A survey. In *2019 10th International Conference on Information, Intelligence, Systems and Applications*, pp. 1–8, 2019.

- Xiaoqian Jiang, Melanie Osl, Jihoon Kim, and Lucila Ohno-Machado. Calibrating predictive model estimates to support personalized medicine. *Journal of the American Medical Informatics Association*, 19(2):263–274, March 2012.
- Masahiro Kato, Takeshi Teshima, and Junya Honda. Learning from positive and unlabeled data with a selection bias. In *International Conference on Learning Representations*, May 2019.
- Ryuichi Kiryo, Gang Niu, Marthinus C. du Plessis, and Masashi Sugiyama. Positive-unlabeled learning with non-negative risk estimator. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- Andrey N. Kolmogorov and Vladimir M. Tikhomirov. ε -entropy and ε -capacity of sets in function spaces. *Uspekhi Matematicheskikh Nauk*, 14:3–86, 1959.
- Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009.
- Ananya Kumar, Percy S. Liang, and Tengyu Ma. Verified uncertainty calibration. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- Fuyi Li, Shuangyu Dong, André Leier, Meiya Han, Xudong Guo, Jing Xu, Xiaoyu Wang, Shirui Pan, Cangzhi Jia, Yang Zhang, Geoffrey I Webb, Lachlan J M Coin, Chen Li, and Jiangning Song. Positive-unlabeled learning in bioinformatics and computational biology: a brief review. *Briefings in Bioinformatics*, 23(1):bbab461, January 2022.
- Bing Liu, Wee Sun Lee, Philip S. Yu, and Xiaoli Li. Partially supervised classification of text documents. In *Proceedings of the 19th International Conference on Machine Learning*, pp. 387–394. Morgan Kaufmann Publishers Inc., July 2002.
- Nan Lu, Tianyi Zhang, Gang Niu, and Masashi Sugiyama. Mitigating overfitting in supervised classification from two unlabeled datasets: A consistent risk correction approach. In Silvia Chiappa and Roberto Calandra (eds.), *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pp. 1115–1125. PMLR, August 2020.
- Allan H. Murphy and Edward S. Epstein. Verification of probabilistic predictions: A brief review. *Journal of Applied Meteorology*, 6(5):748–755, October 1967.
- Anish Muthali, Haotian Shen, Sampada Deglurkar, Michael H. Lim, Rebecca Roelofs, Aleksandra Faust, and Claire Tomlin. Multi-agent reachability calibration with conformal prediction. In *2023 62nd IEEE Conference on Decision and Control*, pp. 6596–6603, 2023.
- Mahdi P. Naeini, Gregory Cooper, and Milos Hauskrecht. Obtaining well calibrated probabilities using bayesian binning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 29, February 2015.
- Nagarajan Natarajan, Inderjit S. Dhillon, Pradeep K. Ravikumar, and Ambuj Tewari. Learning with noisy labels. In C.J. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K.Q. Weinberger (eds.), *Advances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc., 2013.
- John C. Platt. Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods. In *Advances in large margin classifiers*. MIT Press, 1999.
- Teodora Popordanoska, Raphael Sayer, and Matthew B. Blaschko. A consistent and differentiable lp canonical calibration error estimator. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (eds.), *Advances in Neural Information Processing Systems*, volume 35, pp. 7933–7946. Curran Associates, Inc., 2022.

- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In Marina Meila and Tong Zhang (eds.), *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 8748–8763. PMLR, July 2021.
- Christian Robert and George Casella. *Monte Carlo Statistical Methods*. Springer texts in statistics. Springer, New York, NY, November 2010.
- Rebecca Roelofs, Nicholas Cain, Jonathon Shlens, and Michael C. Mozer. Mitigating bias in calibration error estimation. In Gustau Camps-Valls, Francisco J R Ruiz, and Isabel Valera (eds.), *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*, volume 151 of *Proceedings of Machine Learning Research*, pp. 4036–4054. PMLR, March 2022.
- Leonard J. Savage. Elicitation of personal probabilities and expectations. *Journal of the American Statistical Association*, 66(336):783–801, 1971.
- Gideon Schwarz. Estimating the dimension of a model. *Annals of Statistics*, 6(2):461–464, March 1978.
- Thomas Steinke and Lydia Zakyntinou. Reasoning about generalization via conditional mutual information. In *Conference on Learning Theory*, pp. 3437–3452. PMLR, July 2020.
- Juozas Vaicenavicius, David Widmann, Carl Andersson, Fredrik Lindsten, Jacob Roll, and Thomas Schön. Evaluating model calibration in classification. In Kamalika Chaudhuri and Masashi Sugiyama (eds.), *Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics*, volume 89 of *Proceedings of Machine Learning Research*, pp. 3459–3467. PMLR, April 2019.
- Vladimir N. Vapnik and Aleksei Y. Chervonenkis. The uniform convergence of frequencies of the appearance of events to their probabilities. *Doklady Akademii Nauk*, 181(4):781–783, 1968.
- Vladimir N. Vapnik and Alexey Y. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability & Its Applications*, 16(2):264–280, 1971.
- Martin J. Wainwright. *High Dimensional Statistics Non Asymptotic Viewpoint*. Cambridge University Press, 2019.
- Abraham Wald. Statistical decision functions which minimize the maximum risk. *Annals of Mathematics*, 46(2):265–280, 1945.
- Jonathan Wenger, Hedvig Kjellström, and Rudolph Triebel. Non-parametric calibration for classification. In Silvia Chiappa and Roberto Calandra (eds.), *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pp. 178–190. PMLR, August 2020.
- Wang Xinrui, Wenhai Wan, Chuanxing Geng, Shao-Yuan Li, and Songcan Chen. Beyond myopia: Learning from positive and unlabeled data through holistic predictive trends. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (eds.), *Advances in Neural Information Processing Systems*, volume 36, pp. 67589–67602. Curran Associates, Inc., 2023.
- Yu Yao, Tongliang Liu, Bo Han, Mingming Gong, Gang Niu, Masashi Sugiyama, and Dacheng Tao. Re-thinking class-prior estimation for positive-unlabeled learning. In *International Conference on Learning Representations*, May 2021.
- Xichen Ye, Xiaoqiang Li, songmin dai, Tong Liu, Yan Sun, and Weiqin Tong. Active negative loss functions for learning with noisy labels. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (eds.), *Advances in Neural Information Processing Systems*, volume 36, pp. 6917–6940. Curran Associates, Inc., 2023.

- Bianca Zadrozny and Charles Elkan. Obtaining calibrated probability estimates from decision trees and naive bayesian classifiers. In *Proceedings of the 18th International Conference on Machine Learning*, pp. 609–616. Morgan Kaufmann Publishers Inc., June 2001.
- Jize Zhang, Bhavya Kailkhura, and T. Yong-Jin Han. Mix-n-match : Ensemble and compositional methods for uncertainty calibration in deep learning. In Hal Daumé III and Aarti Singh (eds.), *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 11117–11128. PMLR, July 2020.
- Yunrui Zhao, Qianqian Xu, Yangbangyan Jiang, Peisong Wen, and Qingming Huang. Dist-PU: Positive-unlabeled learning from a label distribution perspective. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14461–14470, June 2022.
- Yilun Zhu, Aaron Fjeldsted, Darren Holland, George Landon, Azaree Lintereur, and Clayton Scott. Mixture proportion estimation beyond irreducibility. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (eds.), *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 42962–42982. PMLR, 2023.