

A Generalized Unbiased Risk Estimator for Learning with Augmented Classes

Senlin Shu¹, Shuo He², Haobo Wang³, Hongxin Wei⁴, Tao Xiang¹, Lei Feng^{1*}

¹College of Computer Science, Chongqing University, China

²School of Computer Science and Engineering, University of Electronic Science and Technology of China, China

³College of Computer Science and Technology, Zhejiang University, China

⁴School of Computer Science and Engineering, Nanyang Technological University, Singapore

senlinshu@stu.cqu.edu.cn, shuohe@std.uestc.edu.cn, wanghaobo@zju.edu.cn,

hongxin001@e.ntu.edu.sg, {txiang, lfeng}@cqu.edu.cn

Abstract

In contrast to the standard learning paradigm where all classes can be observed in training data, *learning with augmented classes* (LAC) tackles the problem where augmented classes unobserved in the training data may emerge in the test phase. Previous research showed that given unlabeled data, an *unbiased risk estimator* (URE) can be derived, which can be minimized for LAC with theoretical guarantees. However, this URE is only restricted to the specific type of one-versus-rest loss functions for multi-class classification, making it not flexible enough when the loss needs to be changed with the dataset in practice. In this paper, we propose a *generalized* URE that can be equipped with arbitrary loss functions while maintaining the theoretical guarantees, given unlabeled data for LAC. To alleviate the issue of negative empirical risk commonly encountered by previous studies, we further propose a novel risk-penalty regularization term. Experiments demonstrate the effectiveness of our proposed method.

Introduction

Machine learning approaches have achieved great performance on a variety of tasks, and most of them focus on the stationary learning environment. However, the learning environment in many real-world scenarios could be open and change gradually, which requires the learning approaches to have the ability of handling the distribution change in the non-stationary environment (Sugiyama and Kawanabe 2012; Gama et al. 2014; Zhao et al. 2019; Zhang et al. 2020).

This paper considers a specific problem where the class distribution changes from the training phase to the test phase, called *learning with augmented classes* (LAC). In LAC, some augmented classes unobserved in the training phase might emerge in the test phase. In order to make accurate and reliable predictions, the learning model is required to distinguish augmented classes and keep good generalization performance over the test distribution.

The major difficulty in LAC is how to exploit the relationships between known and augmented classes. To overcome this difficulty, various learning methods have been proposed. For example, by learning a compact geometric description of known classes to distinguish augmented classes that are

far away from the description, the anomaly detection or novelty detection methods can be used (e.g., iForest (Liu, Ting, and Zhou 2008), one-class SVM (Schölkopf et al. 2001; Tax and Duin 2004), and kernel density estimation (Parzen 1962; Kim and Scott 2012)). By exploiting unlabeled data with the low-density separation assumption to adjust the classification decision boundary (Da, Yu, and Zhou 2014), the performance of LAC can be empirically improved.

A recent study (Zhang et al. 2020) showed that by exploiting unlabeled data, an *unbiased risk estimator* (URE) for LAC over the test distribution can be constructed under the *class shift condition*. The motivation of this study stems from that although the instances from augmented classes cannot be observed in labeled data, their distribution information may be contained in unlabeled data and estimated by differentiating the distribution of known classes from unlabeled data. Such an URE is advantageous for the learning task, since it can lead to a theoretically grounded method based on empirical risk minimization. However, the URE for LAC derived by the previous study (Zhang et al. 2020) is only restricted to the specific type of one-versus-rest loss functions, making it not flexible enough when the loss needs to be changed with the dataset in practice.

To address the above issue, we make the following contributions in this paper:

- we propose a *generalized* unbiased risk estimator that can be equipped with *arbitrary* loss functions while maintaining the theoretical guarantees, given unlabeled data for learning with augmented classes.
- We provide a theoretical analysis on the estimation error bound, which guarantees that the convergence of the empirical risk minimizer to the true risk minimizer.
- We propose a novel risk-penalty regularization term that can be used to alleviate the issue of negative empirical risk commonly encountered by previous studies.
- We conduct extensive experiments with different base models on widely used benchmark datasets to demonstrate the effectiveness of our proposed method.

Preliminaries

In this section, we introduce some preliminary knowledge.

*Corresponding author: Lei Feng <lfeng@cqu.edu.cn>.

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Learning with augmented classes. In the training phase of LAC, we are given a labeled training set $\mathcal{D}_L = \{\mathbf{x}_i, y_i\}_{i=1}^n$ independently sampled from a distribution of known classes P_{kc} over $\mathcal{X} \times \mathcal{Y}'$, where \mathcal{X} is the feature space and $\mathcal{Y}' = \{1, 2, \dots, k\}$ is the label space of k known classes. However, in the test phase, we need to predict unlabeled data sampled from the test distribution P_{te} , where augmented classes unobserved in the training phase may emerge. Since the specific partition of augmented classes is unknown, they are normally predicted as a single class. In this way, the label space of the test distribution P_{te} is $\mathcal{Y} = \{1, 2, \dots, k, \text{ac}\}$ where ac denotes the class grouped by augmented classes. The *class shift condition* (Zhang et al. 2020) was introduced to describe the relationship between the distributions of known and augmented classes:

$$P_{te} = \theta \cdot P_{kc} + (1 - \theta) \cdot P_{ac}, \quad (1)$$

where $\theta \in [0, 1]$ indicates the mixture proportion. In addition, a set of unlabeled data $\mathcal{D}_U = \{\mathbf{x}_i\}_{i=1}^m$ sampled from the test distribution is provided and will be used for model training. This learning scenario of available unlabeled data is conceivable since we can easily collect a large amount of unlabeled data from current environments. Such unlabeled data can be used to enrich the information of unobservable classes in the training phase, despite no revealed supervision information on the unlabeled data.

Unbiased risk estimator for LAC. Under the class shift condition in LAC, the expected classification risk over the test distribution can be formulated as

$$\begin{aligned} R(f) &= \mathbb{E}_{(\mathbf{x}, y) \sim P_{te}} [\mathcal{L}(f(\mathbf{x}), y)] \\ &= \theta \mathbb{E}_{(\mathbf{x}, y) \sim P_{kc}} [\mathcal{L}(f(\mathbf{x}), y)] \\ &\quad + (1 - \theta) \mathbb{E}_{(\mathbf{x}, y = \text{ac}) \sim P_{ac}} [\mathcal{L}(f(\mathbf{x}), \text{ac})]. \end{aligned} \quad (2)$$

The recent study (Zhang et al. 2020) only considered the *one-versus-rest* (OVR) loss function as the classification loss in Eq. (2), which takes the following form for k -class classification:

$$\mathcal{L}_{\text{OVR}}(f(\mathbf{x}), y) = \psi(f_y(\mathbf{x})) + \sum_{i=1, i \neq y}^k \psi(-f_i(\mathbf{x})),$$

where $\psi(\cdot) : \mathbb{R} \mapsto \mathbb{R}_+$ denotes a binary loss function. By substituting the OVR loss \mathcal{L}_{OVR} into the classification risk, we have $R^{\text{OVR}}(f) = \mathbb{E}_{(\mathbf{x}, y) \sim P_{te}} [\mathcal{L}_{\text{OVR}}(f(\mathbf{x}), y)]$. The previous study (Zhang et al. 2020) showed that an equivalent expression of $R^{\text{OVR}}(f)$ in the setting of LAC with unlabeled data drawn from the test distribution can be derived:

$$\begin{aligned} R_{\text{LAC}}^{\text{OVR}}(f) &= \theta \cdot \mathbb{E}_{P_{kc}(\mathbf{x}, y)} [\psi(f_y(\mathbf{x})) - \psi(-f_y(\mathbf{x})) \\ &\quad + \psi(-f_{\text{ac}}(\mathbf{x})) - \psi(f_{\text{ac}}(\mathbf{x}))] \\ &\quad + \mathbb{E}_{\mathbf{x} \sim P_{te}} [\psi(f_{\text{ac}}(\mathbf{x})) + \sum_{i=1}^k \psi(-f_k(\mathbf{x}))]. \end{aligned}$$

As can be verified, $R_{\text{LAC}}^{\text{OVR}}(f) = R^{\text{OVR}}(f)$, and thus its empirical version $\widehat{R}_{\text{LAC}}^{\text{OVR}}(f)$ is an *unbiased risk estimator* (URE). As we can see, this URE is only restricted to the OVR loss for multi-class classification, making it not flexible enough when the loss needs to be changed with the dataset in practice.

The Proposed Method

In this section, we first propose a generalized URE and provide a theoretical analysis for the derived URE. Then, we propose a novel risk-penalty regularization term that can be used to alleviate the issue of negative empirical risk.

A Generalized Unbiased Risk Estimator

Theorem 1. *Under the class shift condition in Eq. (1), the classification risk $R(f)$ can be equivalently expressed as:*

$$\begin{aligned} R_{\text{LAC}}(f) &= \theta \mathbb{E}_{(\mathbf{x}, y) \sim P_{kc}} [\mathcal{L}(f(\mathbf{x}), y) - \mathcal{L}(f(\mathbf{x}), \text{ac})] \\ &\quad + \mathbb{E}_{\mathbf{x} \sim P_{te}} [\mathcal{L}(f(\mathbf{x}), \text{ac})] = R(f). \end{aligned} \quad (3)$$

Theorem 1 shows that we are able to learn a classifier from labeled data from the distribution of known classes and unlabeled data from the test distribution, without any restrictions on the loss function \mathcal{L} for multi-class classification.

Corollary 1. *If the one-versus-rest (OVR) loss \mathcal{L}_{OVR} is used as the classification loss in our derived risk $R_{\text{LAC}}(f)$, then we can exactly recover the risk $R_{\text{LAC}}^{\text{OVR}}(f)$ derived by the previous study (Zhang et al. 2020).*

The proof of Corollary 1 is omitted here, since it is quite straightforward to verify that by directly inserting \mathcal{L}_{OVR} into $R_{\text{LAC}}(f)$. Corollary 1 shows that our proposed URE is a generalization of the URE proposed by Zhang et al. (2020) and can be compatible with arbitrary loss functions.

Given a set of labeled data $\mathcal{D}_L = \{\mathbf{x}_i, y_i\}_{i=1}^n$ drawn from the distribution of known classes P_{kc} and a set of unlabeled data $\mathcal{D}_U = \{\mathbf{x}_i\}_{i=1}^m$ drawn from the test distribution P_{te} , we can obtain the following URE, which is the empirical approximation of R_{LAC} :

$$\begin{aligned} \widehat{R}_{\text{LAC}}(f) &= \frac{\theta}{n} \sum_{i=1}^n \left(\mathcal{L}(f(\mathbf{x}_i), y_i) - \mathcal{L}(f(\mathbf{x}_i), \text{ac}) \right) \\ &\quad + \frac{1}{m} \sum_{j=1}^m \mathcal{L}(f(\mathbf{x}_j), \text{ac}). \end{aligned} \quad (4)$$

In this way, we are able to learn an effective multi-class classifier by directly minimizing $\widehat{R}_{\text{LAC}}(f)$. Since there are no restrictions on the loss function \mathcal{L} and the model f , we can use any loss and any model for LAC with unlabeled data.

Theoretical Analysis

Here, we establish an estimation error bound for our proposed URE based on the widely used *Rademacher complexity* (Bartlett and Mendelson 2002).

Definition 1 (Rademacher Complexity). *Let n be a positive integer, $\mathbf{x}_1, \dots, \mathbf{x}_n$ be independent and identically distributed random variables drawn from a probability distribution with density μ , $\mathcal{G} = g : \mathcal{X} \mapsto \mathbb{R}$ be a class of measurable functions, and $\boldsymbol{\sigma} = (\sigma_1, \dots, \sigma_n)$ be Rademacher variables that take value from $\{+1, -1\}$ with even probabilities. Then, the Rademacher complexity of \mathcal{G} is defined as*

$$\mathfrak{R}_n(\mathcal{G}) := \mathbb{E}_{\mathbf{x}_1, \dots, \mathbf{x}_n \stackrel{\text{i.i.d.}}{\sim} \mu} \mathbb{E}_{\boldsymbol{\sigma}} \left[\sup_{g \in \mathcal{G}} \frac{1}{n} \sum_{i=1}^n \sigma_i g(\mathbf{x}_i) \right].$$

Let us represent \mathcal{F} by $\mathcal{F} = \mathcal{F}_1 \times \mathcal{F}_2 \cdots \mathcal{F}_{k+1}$ where $\mathcal{F}_i = \{\mathbf{x} \mapsto f_i(\mathbf{x}) \mid f \in \mathcal{F}\}$ and \mathcal{F}_{k+1} is specially for the augmented class ac. Thus we can denote

by $\mathfrak{R}_n(\mathcal{F}_y)$ the Rademacher complexity of \mathcal{F}_y for the y -th class, given the sample size n over the test distribution P_{te} . It is commonly assumed that $\mathfrak{R}_n(\mathcal{F}_y) \leq C_{\mathcal{F}}/\sqrt{n}$ for all $y \in \mathcal{Y}$ (Bao, Niu, and Sugiyama 2018), where $C_{\mathcal{F}}$ is a positive constant. Let the learned classifier by our URE be $\hat{f} = \arg \min_{f \in \mathcal{F}} \hat{R}_{LAC}(f)$ and the optimal classifier learned by minimizing the true classification risk be $f^* = \arg \min_{f \in \mathcal{F}} R(f)$. Then we have the following theorem.

Theorem 2. Assume the multi-class loss $\mathcal{L}(f(\mathbf{x}), y)$ is ρ -Lipschitz ($0 < \rho < \infty$) with respect to $f(\mathbf{x})$ for all $y \in \mathcal{Y}$ and upper bounded by a positive constant $C_{\mathcal{L}}$, i.e., $C_{\mathcal{L}} = \sup_{\mathbf{x} \in \mathcal{X}, y \in \mathcal{Y}, f \in \mathcal{F}} \mathcal{L}(f(\mathbf{x}), y)$. Then, for any $\delta > 0$, with probability at least $1 - \delta$,

$$R(\hat{f}) - R(f^*) \leq C_{k,\rho,\delta} \left(\frac{2\theta}{\sqrt{n}} + \frac{1}{\sqrt{m}} \right),$$

where $C_{k,\rho,\delta} = (4\sqrt{2}\rho(k+1)C_{\mathcal{F}} + 2C_{\mathcal{L}}\sqrt{\frac{\log \frac{1}{\delta}}{2}})$.

Theorem 2 shows that our proposed method is consistent, i.e., $R(\hat{f}) \rightarrow R(f^*)$ as $n \rightarrow \infty$ and $m \rightarrow \infty$. The convergence rate is $\mathcal{O}_p(1/\sqrt{n} + 1/\sqrt{m})$, where \mathcal{O}_p denotes the order in probability. This order is the optimal parametric rate for the empirical risk minimization without additional assumptions (Mendelson, 2008).

Risk-Penalty Regularization

As shown in Eq. (4), there exist negative terms in the URE $\hat{R}_{LAC}(f)$. Since the widely used classification loss (i.e., the cross entropy loss) is unbounded above, $\hat{R}_{LAC}(f)$ could be unbounded below, which causes the overfitting issue, and this issue becomes especially terrible when deep models are employed. Such an issue was encountered and demonstrated by many previous studies (Kiryo et al. 2017; Lu et al. 2020; Zhang et al. 2020; Feng et al. 2021). To alleviate this issue, these studies normally wrapped the terms that could make empirical risk negative by certain correction functions, such as the *rectified linear unit* (ReLU) function $g(z) = \max(0, z)$ and the *absolute value* (ABS) function $g(z) = |z|$. Despite the effectiveness of the correction functions used by previous studies (Lu et al. 2020), using a correction function could make the empirical risk not unbiased anymore. Motivated by this, we aim to keep the original URE unchanged and impose a regularization term to alleviate the issue of negative empirical risk.

We note that the above issue in our work comes from the use of equality: $(1 - \theta)\mathbb{E}_{(\mathbf{x}, y=\text{ac}) \sim P_{ac}}[\mathcal{L}(f(\mathbf{x}), \text{ac})] = \mathbb{E}_{(\mathbf{x}, y) \sim P_{te}}[\mathcal{L}(f(\mathbf{x}), \text{ac})] - \theta\mathbb{E}_{(\mathbf{x}, y) \sim P_{kc}}[\mathcal{L}(f(\mathbf{x}), \text{ac})]$. Let us denote by \hat{R}_{PAC} the empirical version of the expected risk on the left hand side of the above equality:

$$\hat{R}_{PAC}(f) = \frac{1}{m} \sum_{j=1}^m \mathcal{L}(f(\mathbf{x}_j), \text{ac}) - \frac{\theta}{n} \sum_{i=1}^n \mathcal{L}(f(\mathbf{x}_i), \text{ac}).$$

Our goal is to prevent the whole training objective from going to be unbounded below in the training process. Therefore, we choose to add an regularization term on the training objective, which could incur a cost when $\hat{R}_{PAC}(f)$ goes to




Figure 1: The function curve of our proposed risk-penalty regularization.

be negative. In this way, the whole training objective would not be unbounded below, because there is always a positive cost incurred if $\hat{R}_{PAC}(f)$ goes to be negative. Our proposed regularization term is presented as follows:

$$\Omega(f) = \begin{cases} (-\hat{R}_{PAC}(f))^t, & \text{if } \hat{R}_{PAC}(f) < 0, \\ 0, & \text{otherwise,} \end{cases} \quad (5)$$

where $t \geq 0$ is a hyper-parameter. We plot the above regularization function in Figure 1. Using this regularization, our final training objective becomes $\hat{R}_{LAC}(f) + \lambda\Omega(f)$, where $\hat{R}_{LAC}(f)$ is our derived URE defined in Eq. (4) and λ is a hyper-parameter that controls the importance of the regularization term. As can be seen from Figure 1, when the empirical risk becomes negative, an extra loss will be generated, thereby alleviating the issue of the negative empirical risk during training. Additionally, the widely used correction functions ReLU and ABS can be considered as special cases of our proposed negative risk penalty regularization. Specifically, when $t = 1$ and $\lambda = 1$, we obtain a formulation by using ReLU as the correction function. When $t = 1$ and $\lambda = 2$, we obtain a formulation by using ABS as the correction function. Therefore, our proposed risk-penalty regularization could be considered as a generalized solution to the issue of negative empirical risk.

Handling Class Prior Shift

The class shift condition in Eq. (1) describes the relationship between the distributions of known classes and augmented classes, with the assumption that the distribution of known classes keeps identical in both the training phase and the test phase. However, the learning environment could dynamically change in real-world scenarios, and thus the prior of each class could also vary from the training phase and the test phase. As shown by Zhang et al. (2020), the following generalized class shift condition can be obtained:

$$P_{te} = \sum_{i=1}^k \theta_{te}^i \cdot P_i + (1 - \sum_{i=1}^k \theta_{te}^i) \cdot P_{ac}, \quad (6)$$

where P_i denotes the data distribution of the i -th known class and θ_{te}^i denotes the class prior of the i -th known class in the test phase. Given the generalized class shift condition in Eq. (6), we can still obtain an equivalent representation of the classification risk $R(f)$.

Theorem 3. Under the generalized class shift condition in Eq. (6), the classification risk $R(f)$ can be equivalently expressed as follows:

$$R_{\text{LAC}}^{\text{shift}}(f) = \sum_{i=1}^k \theta_{\text{te}}^i \cdot \mathbb{E}_{(\mathbf{x}, y) \sim P_i} [\mathcal{L}(f(\mathbf{x}), y) - \mathcal{L}(f(\mathbf{x}), \text{ac})] + \mathbb{E}_{\mathbf{x} \sim P_{\text{te}}} [\mathcal{L}(f(\mathbf{x}), \text{ac})] = R(f). \quad (7)$$

Proof. We omitted the proof of Theorem 3, since it is quite straightforward by following the proof of Theorem 1. \square

Theorem 3 shows that we can learn a classifier for LAC with unlabeled data under the class prior shift condition of known classes by empirically minimizing $R_{\text{LAC}}^{\text{shift}}(f)$.

Experiments

In this section, we conduct extensive experiments to evaluate the performance of our proposed method on various datasets using different models.

Experimental Setup

Datasets. We use six regular-scale datasets downloaded from the UCI Machine Learning Repository¹ (Dua, Graff et al. 2017), including Har, Msplice, Normal, Optdigits, Texture, and Usps. Since they are not large-scale datasets, we train a linear model on these datasets. We also use four widely large-scale benchmark datasets, including MNIST² (LeCun et al. 1998), Fashion-MNIST³ (Xiao, Rasul, and Vollgraf 2017), Kuzushiji-MNIST⁴ (Clanuwat et al. 2018), and SVHN⁵ (Netzer et al. 2011). For MNIST, Fashion-MNIST, and Kuzushiji-MNIST, we train a multilayer perceptron (MLP) model with three layers ($d - 500 - k$) and the ReLU activation function is used. For the SVHN dataset, we train a VGG16 model (Simonyan and Zisserman 2015). The brief characteristics of all the used datasets are reported in Table 1. For each regular-scale dataset, half of the classes are selected as augmented classes and the remaining classes are considered as known classes. Besides, the number of labeled, unlabeled, and test examples set to 500, 1000, and 1000, respectively. For large-scale datasets, we select six classes as known classes and other classes are regarded as augmented classes. For MNIST, Fashion-MNIST, and Kuzushiji-MNIST, the number of the labeled, unlabeled, and test examples is set to 24000 (4000 per known class), 10000 (1000 per class), and 1000 (100 per class), respectively. For SVHN, the number of the labeled, unlabeled, and test examples is set to 24000 (4000 per known class), 25000 (2500 per class), and 1000 (100 per class), respectively. To comprehensively evaluate the performance of our proposed method, we report the mean value with standard deviation over 10 trials, in terms of the evaluation metrics including accuracy, Macro-F1, and AUC.

¹<https://archive.ics.uci.edu>

²<http://yann.lecun.com/exdb/mnist/>

³<https://github.com/zalandoresearch/fashion-mnist>

⁴<http://codh.rois.ac.jp/kmnist/>

⁵<http://ufldl.stanford.edu/housenumbers/>

Dataset	# Examples	# Features	# Classes
Har	10,699	561	6
Msplice	3,175	240	3
Normal	7,000	2	7
Optdigits	5,620	62	10
Texture	5,500	40	11
Usps	9,298	256	10
MNIST	70,000	784	10
Fashion	70,000	784	10
Kuzushiji	70,000	784	10
SVHN	99,268	3,072	10

Table 1: Brief characteristics of the used datasets.

Methods. For the experiments on regular-scale datasets, we compare with the following five methods:

- **OVR** (Rifkin and Klautau 2004), which uses the one-versus-reset loss for training a multi-class classifier. For the LAC problem, OVR predicts the augmented class if $\max_{y \in \mathcal{Y}'} f_y(\mathbf{x}) < 0$, otherwise it makes normal predictions $\arg \max_{y \in \mathcal{Y}} f_y$ in known classes.
- **EVM** (Rudd et al. 2017), which uses the extreme value machine to perform nonlinear kernel-free variable bandwidth incremental learning.
- **LACU** (Da, Yu, and Zhou 2014), which exploits unlabeled data with the low-density separation assumption to adjust the classification decision boundary.
- **iForest** (Liu, Ting, and Zhou 2008), which first uses iForest to detect augmented classes and then uses the OVR method (Rifkin and Klautau 2004) to make normal predictions for known classes.
- **EULAC** (Zhang et al. 2020), which derives an unbiased risk estimator for LAC with unlabeled data.

For the experiments on large-scale datasets, we compare with OVR, EULAC, and following three methods:

- **Softmax**, which directly trains a multi-class classifier by using the softmax cross entropy loss and only makes normal predictions for known classes.
- **Softmax-T**, which denotes the Softmax method with a threshold set for predicting the augmented class. It predicts the augmented class if $\max_{y \in \mathcal{Y}'} g_y(\mathbf{x})$ is less than a given threshold, where $g_y(\mathbf{x})$ is a softmax score of the class y for the instance \mathbf{x} . The threshold is set to 0.95.
- **Openmax** (Bendale and Boult 2016), which can be considered as a calibrated version of the Softmax-T method. In Openmax, Weibull calibration is implemented as an augment to the SoftMax method while replacing the Softmax layer with a new OpenMax layer.

For the compared methods, we adopt the hyper-parameter settings suggested by respective papers. For our proposed method, we utilize the *generalized cross entropy* (GCE) loss (Zhang and Sabuncu 2018) as the multi-class loss function, since it is a generalized version of the widely used cross entropy loss. We use the Adam optimization method (Kingma and Ba 2015), with the number of training epochs set to 1500 on regular-scale datasets and 200

Datasets	Accuracy					
	OVR	EVM	LACU	iForest	EULAC	NRPR
Har	0.452±0.012	0.413±0.007	0.542±0.077	0.958±0.005	0.967±0.009	0.986±0.007
Msplice	0.725±0.008	0.658±0.005	0.664±0.015	0.716±0.007	0.781±0.032	0.785±0.020
Normal	0.570±0.001	0.567±0.003	0.567±0.050	0.775±0.036	0.755±0.051	0.851±0.009
Optdigits	0.533±0.014	0.460±0.009	0.724±0.068	0.728±0.020	0.791±0.047	0.863±0.037
Texture	0.570±0.010	0.474±0.018	0.520±0.062	0.601±0.029	0.732±0.109	0.930±0.009
Usps	0.631±0.026	0.547±0.007	0.572±0.057	0.551±0.005	0.874±0.023	0.878±0.012
Datasets	Macro-F1					
	OVR	EVM	LACU	iForest	EULAC	NRPR
Har	0.526±0.014	0.530±0.008	0.277±0.020	0.950±0.006	0.950±0.011	0.980±0.009
Msplice	0.549±0.007	0.490±0.005	0.537±0.049	0.537±0.007	0.759±0.034	0.767±0.021
Normal	0.679±0.001	0.658±0.028	0.405±0.221	0.747±0.050	0.586±0.103	0.762±0.013
Optdigits	0.606±0.012	0.572±0.014	0.745±0.069	0.708±0.025	0.690±0.103	0.859±0.026
Texture	0.709±0.008	0.561±0.026	0.294±0.194	0.696±0.019	0.540±0.192	0.927±0.011
Usps	0.667±0.026	0.579±0.008	0.370±0.111	0.590±0.010	0.853±0.038	0.867±0.012
Datasets	AUC					
	OVR	EVM	LACU	iForest	EULAC	NRPR
Har	0.789±0.006	0.754±0.007	0.695±0.051	0.953±0.005	0.998±0.001	0.999±0.001
Msplice	0.737±0.007	0.652±0.006	0.748±0.011	0.733±0.006	0.924±0.013	0.917±0.010
Normal	0.849±0.001	0.847±0.002	0.748±0.029	0.803±0.031	0.976±0.026	0.999±0.000
Optdigits	0.866±0.005	0.836±0.008	0.834±0.041	0.798±0.021	0.907±0.053	0.987±0.002
Texture	0.894±0.004	0.831±0.016	0.720±0.036	0.828±0.012	0.756±0.126	0.995±0.001
Usps	0.870±0.007	0.846±0.007	0.743±0.034	0.843±0.004	0.979±0.005	0.981±0.002

Table 2: Test performance (mean±std) of each method on UCI datasets. The best performance is highlighted in bold.

Datasets	Accuracy					
	OVR	Softmax	Softmax-T	Openmax	EULAC	NRPR
MNIST	0.743±0.018	0.592±0.002	0.883±0.009	0.882±0.007	0.952±0.006	0.960±0.006
Fashion	0.587±0.009	0.559±0.007	0.595±0.011	0.601±0.011	0.840±0.034	0.875±0.010
Kuzushiji	0.696±0.017	0.578±0.004	0.828±0.009	0.835±0.012	0.881±0.016	0.927±0.009
SVHN	0.595±0.022	0.564±0.006	0.720±0.037	0.786±0.013	0.835±0.021	0.873±0.012
Datasets	Macro-F1					
	OVR	Softmax	Softmax-T	Openmax	EULAC	NRPR
MNIST	0.794±0.015	0.653±0.004	0.898±0.007	0.898±0.006	0.951±0.007	0.960±0.006
Fashion	0.667±0.010	0.637±0.007	0.685±0.009	0.688±0.008	0.813±0.070	0.872±0.010
Kuzushiji	0.746±0.014	0.630±0.005	0.847±0.009	0.851±0.011	0.877±0.021	0.926±0.008
SVHN	0.640±0.020	0.614±0.007	0.760±0.029	0.801±0.014	0.831±0.021	0.876±0.011
Datasets	AUC					
	OVR	Softmax	Softmax-T	Openmax	EULAC	NRPR
MNIST	0.929±0.005	0.890±0.002	0.939±0.005	0.939±0.005	0.996±0.002	0.998±0.001
Fashion	0.868±0.004	0.864±0.005	0.822±0.005	0.820±0.004	0.953±0.039	0.985±0.002
Kuzushiji	0.908±0.005	0.880±0.003	0.909±0.008	0.905±0.008	0.981±0.009	0.994±0.001
SVHN	0.835±0.011	0.868±0.005	0.889±0.012	0.865±0.018	0.931±0.019	0.984±0.003

Table 3: Test performance (mean±std) of each method on benchmark datasets. The best performance is highlighted in bold.

on large-scale datasets respectively. We set the mini-batch size to 500 on large-scale datasets and use the full batch size on regular-scale datasets. For regular-scale datasets, learning rate and weight decay are selected in $\{10^{-2}, 10^{-3}, 10^{-4}\}$. For large-scale datasets, learning rate and weight decay are selected in $\{10^{-3}, 10^{-4}, 10^{-5}\}$. For our method, t and λ are selected in $\{1, 2, 3\}$ and $\{0.2, 0.4, \dots, 1.8, 2\}$, respectively. Since our method uses a URE with negave-risk penalty regularization, we name it **NRPR**. All the experiments are con-

ducted on GeForce GTX 3090 GPUs.

Experimental Results

Results on regular-scale datasets. Table 2 records the accuracy, Macro-F1, and AUC of each method on the six regular-scale UCI datasets trained with a linear model. From Table 2, we can observe that our proposed method achieves the best performance in all the cases, which validates the effectiveness of our proposed method. Besides, the superi-

Datasets	Accuracy		
	ReLU	ABS	NRPR
MNIST	0.865±0.034	0.776±0.014	0.960±0.006
Fashion	0.809±0.021	0.747±0.019	0.870±0.009
Kuzushiji	0.853±0.029	0.757±0.027	0.927±0.009
SVHN	0.734±0.018	0.695±0.007	0.865±0.010
Datasets	Macro-F1		
	ReLU	ABS	NRPR
MNIST	0.883±0.027	0.813±0.011	0.960±0.006
Fashion	0.825±0.017	0.782±0.015	0.868±0.009
Kuzushiji	0.868±0.024	0.790±0.023	0.926±0.008
SVHN	0.767±0.015	0.736±0.006	0.867±0.011
Datasets	AUC		
	ReLU	ABS	NRPR
MNIST	0.997±0.001	0.996±0.001	0.998±0.001
Fashion	0.974±0.001	0.966±0.002	0.985±0.002
Kuzushiji	0.992±0.002	0.984±0.004	0.994±0.001
SVHN	0.983±0.003	0.980±0.002	0.981±0.003

Table 4: Performance comparison between our proposed risk-penalty regularization and risk-correction functions.

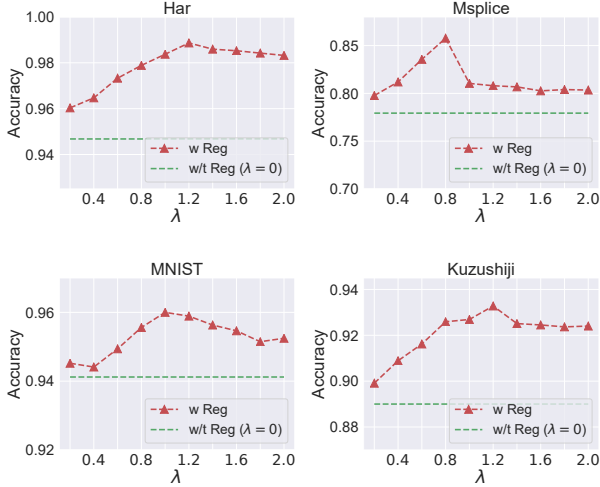


Figure 2: Classification accuracy with different values of the regularization parameter λ .

ority of our method (using a generalized URE) over the EULAC method (using a URE relying on the OVR loss) is also clearly demonstrated, which indicates that our method is not only more flexible on the choice of used loss functions but also is more effective by substituting other loss functions rather than the OVR loss.

Results on large-scale datasets. Table 3 records the accuracy, Macro-F1, and AUC of each method on the four large-scale benchmark datasets. As can be observed from Table 3, our proposed method also achieves the best performance in all the cases. In addition, the performance gap between the EULAC method and our proposed method is also considerably evident. These positive experimental results also support our proposed method.




Figure 3: Influence of the mixture proportion θ (\diamond denotes the empirically estimated mixture proportion $\hat{\theta}$).




Figure 4: Ability to handle class prior shift.

Further Analysis

Effectiveness of the risk-penalty regularization. In order to verify the effectiveness of our proposed risk-penalty regularization, we further compare it with the widely used risk-correction functions (Kiryo et al. 2017; Lu et al. 2020) including the *rectified linear unit* (ReLU) function $g(z) = \max(0, z)$ and the *absolute value* (ABS) function $g(z) = |z|$. We conduct experiments on the four large-scale benchmark datasets and report the accuracy, Macro-F1, and AUC of each method in Table 4. As can be seen from Table 4, our proposed risk-penalty regularization significantly outperforms ReLU and ABS, which demonstrates the superiority of our proposed risk-penalty regularization over the risk correction scheme.

Regularization parameter analysis. To show the concrete effect of our risk-penalty regularization, we further conduct parameter sensitivity analysis of the regularization parameter λ on four datasets. Specifically, λ is varied in $\{0, 0.2, 0.4, \dots, 1.8, 2\}$ ($\lambda = 0$ means that the risk-penalty regularization is not used). From Figure 2, we can find that the risk-penalty regularization indeed improves the learning performance and the best performance is achieved at some intermediate value of λ , which suggests that the regularization plays an important role.

Influence of the mixture proportion. For the mixture parameter θ , we follow Zhang et al. (2020) to empirically estimate θ by using the kernel mean embedding method (Ramswamy, Scott, and Tewari 2016) on regular datasets. To show the influence of the mixture proportion θ , we conduct experiments on the Usps and Optdigits datasets by varying the preset mixture proportion $\hat{\theta}$ from 0.1 to 1 under different values of the true mixture proportion θ . As shown in




Figure 5: Test performance on four regular-scale datasets when the number of unlabeled examples increases.

Figure 3, the true mixture proportion can be accurately estimated in most cases.

Handling class prior shift. In order to verify the ability of our proposed method to handle more complex learning environments, we further conduct experiments on Usps and Optdigits with class prior shift. Specifically, we select $\{1, 2, 3, 4, 5\}$ as known classes with prior 0.2 per class in labeled data and $\{6, 7, 8, 9, 10\}$ as augmented classes with prior 0.1 per class in test and unlabeled data. As for the prior shift in the test distribution, we reset the prior of five known classes to $\{1 - \alpha, 1 - \frac{\alpha}{2}, 1, 1 + \frac{\alpha}{2}, 1 + \alpha\} \times 0.1$, where α controls shift intensity and is selected in $\{0, 0.3, 0.5, 0.7, 0.9\}$. Figure 4 reports the accuracy for NRPR (ignoring class prior shift) and NRPR-shift (considering class prior shift in model training according to our derived Theorem 3). As shown in Figure 4, when the shift intensity increases, performance of NRPR rapidly decreases while NRPR-shift retains satisfactory performance even with a high shift intensity. This observation suggests that our proposed NRPR-shift method could well handle changing learning environments.

Performance of increasing unlabeled examples. As shown in Theorem 2, the performance of our NRPR method is expected to be improved if more unlabeled examples are provided. To empirically validate such a theoretical finding, we conduct experiments on four regular-scale datasets by changing the number of unlabeled examples from 100 to 1200. As shown in Figure 5, the accuracy, Macro-F1, and AUC of our NRPR method generally increases given more unlabeled examples. This observation clearly supports the derived estimation error bound in Theorem 2.

Related Work

In this section, we discuss some research topics and studies that are relevant to our work.

Class-incremental learning. As many practical machine learning systems require the ability to continually learn new knowledge, *class-incremental learning* (CIL) (Zhou and Chen 2002; Tao et al. 2020) was proposed to handling the increasing number of classes. *Learning with augmented classes* (LAC) (Da, Yu, and Zhou 2014; Zhang et al. 2020; Ding, Liu, and Zhang 2018), the problem studied in this paper, is a main task in CIL. The previous research on LAC (Da, Yu, and Zhou 2014) showed that unlabeled data can be exploited to enhance the performance of LAC. Furthermore,

Zhang et al. (2020) showed that given unlabeled data, an unbiased risk estimator (URE) can be derived, which can be minimized for LAC with theoretical guarantees.

Open-set recognition. In *open-set recognition* (OSR) (Geng, Huang, and Chen 2020), new classes unseen in training the phase could appear in the test phase, and the learned classifier is required to not only accurately classify known classes but also effectively deal with unknown classes. Therefore, OSR can be seen as a cousin of the LAC problem in the computer vision and pattern recognition communities. Most studies on OSR aim to explore the relationship between known and augmented classes by using specific techniques, such as the open space risk (Scheirer, Jain, and Boulton 2014), nearest neighbor approach (Mendes Júnior et al. 2017), and extreme value theory (Rudd et al. 2017). It is noteworthy that many OSR methods utilized the domain knowledge, while our work focuses on a general setting without additional information.

Multi-class positive-unlabeled learning. Learning from multi-class positive-unlabeled data (Xu et al. 2017; Shu et al. 2020) is similar to the LAC problem with unlabeled data, by taking the known classes as the multi-class positive. Although our work shares similarity with the two studies (Xu et al. 2017; Shu et al. 2020), our LAC risk derived in a different context brings novel understandings for the LAC and can handle more complex environments (i.e., class prior shift), while they assumed the class priors to be known.

Conclusion

In this paper, we studied the problem of *learning with augmented classes* (LAC). To solve this problem, we proposed a generalized *unbiased risk estimator* (URE) that can be equipped with arbitrary loss functions while maintaining the theoretical guarantees, given unlabeled data for LAC. We showed that our generalized URE can recover the URE proposed by the previous study (Zhang et al. 2020). We also established an estimation error bound, which achieves the optimal parametric convergence rate. To alleviate the issue of negative empirical risk commonly encountered by previous studies, we further proposed a novel risk-penalty regularization term. Comprehensive experimental results clearly verified the effectiveness of our proposed method. In future work, it would be interesting to develop more powerful regularization techniques for further improving the performance of our proposed method.

Acknowledgments

This work is supported by the National Natural Science Foundation of China (Grant No. 62106028), Chongqing Overseas Chinese Entrepreneurship and Innovation Support Program, and CAAI-Huawei MindSpore Open Fund.

References

- Bao, H.; Niu, G.; and Sugiyama, M. 2018. Classification from Pairwise Similarity and Unlabeled Data. In *ICML*, 452–461.
- Bartlett, P. L.; and Mendelson, S. 2002. Rademacher and Gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3(11): 463–482.
- Bendale, A.; and Boulton, T. E. 2016. Towards open set deep networks. In *CVPR*, 1563–1572.
- Clanuwat, T.; Bober-Irizar, M.; Kitamoto, A.; Lamb, A.; Yamamoto, K.; and Ha, D. 2018. Deep learning for classical Japanese literature. *arXiv preprint arXiv:1812.01718*.
- Da, Q.; Yu, Y.; and Zhou, Z.-H. 2014. Learning with augmented class by exploiting unlabeled data. In *AAAI*.
- Ding, S.-Y.; Liu, X.-Y.; and Zhang, M.-L. 2018. Imbalanced augmented class learning with unlabeled data by label confidence propagation. In *ICDM*, 79–88. IEEE.
- Dua, D.; Graff, C.; et al. 2017. UCI machine learning repository. URL <http://archive.ics.uci.edu/ml>, 7(1).
- Feng, L.; Shu, S.; Lu, N.; Han, B.; Xu, M.; Niu, G.; An, B.; and Sugiyama, M. 2021. Pointwise Binary Classification with Pairwise Confidence Comparisons. In *ICML*.
- Gama, J.; Žliobaitė, I.; Bifet, A.; Pechenizkiy, M.; and Bouchachia, A. 2014. A survey on concept drift adaptation. *ACM Computing Surveys*, 46(4): 1–37.
- Geng, C.; Huang, S.-j.; and Chen, S. 2020. Recent advances in open set recognition: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(10): 3614–3631.
- Kim, J.; and Scott, C. D. 2012. Robust kernel density estimation. *Journal of Machine Learning Research*, 13(1): 2529–2565.
- Kingma, D. P.; and Ba, J. 2015. Adam: A method for stochastic optimization. In *ICLR*.
- Kiryo, R.; Niu, G.; Du Plessis, M. C.; and Sugiyama, M. 2017. Positive-unlabeled learning with non-negative risk estimator. In *NeurIPS*, 1675–1685.
- LeCun, Y.; Bottou, L.; Bengio, Y.; and Haffner, P. 1998. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11): 2278–2324.
- Liu, F. T.; Ting, K. M.; and Zhou, Z.-H. 2008. Isolation forest. In *ICDM*, 413–422.
- Lu, N.; Zhang, T.; Niu, G.; and Sugiyama, M. 2020. Mitigating overfitting in supervised classification from two unlabeled datasets: A consistent risk correction approach. In *AISTATS*, 1115–1125.
- Mendes Júnior, P. R.; De Souza, R. M.; Werneck, R. d. O.; Stein, B. V.; Pazinato, D. V.; de Almeida, W. R.; Penatti, O. A.; Torres, R. d. S.; and Rocha, A. 2017. Nearest neighbors distance ratio open-set classifier. *Machine Learning*, 106(3): 359–386.
- Netzer, Y.; Wang, T.; Coates, A.; Bissacco, A.; Wu, B.; and Ng, A. Y. 2011. Reading digits in natural images with unsupervised feature learning. In *NeurIPS Workshop*.
- Parzen, E. 1962. On estimation of a probability density function and mode. *The Annals of Mathematical Statistics*, 33(3): 1065–1076.
- Ramaswamy, H.; Scott, C.; and Tewari, A. 2016. Mixture proportion estimation via kernel embeddings of distributions. In *ICML*, 2052–2060.
- Rifkin, R.; and Klautau, A. 2004. In defense of one-vs-all classification. *Journal of Machine Learning Research*, 5: 101–141.
- Rudd, E. M.; Jain, L. P.; Scheirer, W. J.; and Boulton, T. E. 2017. The extreme value machine. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(3): 762–768.
- Scheirer, W. J.; Jain, L. P.; and Boulton, T. E. 2014. Probability models for open set recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(11): 2317–2324.
- Schölkopf, B.; Platt, J. C.; Shawe-Taylor, J.; Smola, A. J.; and Williamson, R. C. 2001. Estimating the support of a high-dimensional distribution. *Neural Computation*, 13(7): 1443–1471.
- Shu, S.; Lin, Z.; Yan, Y.; and Li, L. 2020. Learning from multi-class positive and unlabeled data. In *ICDM*, 1256–1261.
- Simonyan, K.; and Zisserman, A. 2015. Very Deep Convolutional Networks for Large-Scale Image Recognition. In *ICLR*.
- Sugiyama, M.; and Kawanabe, M. 2012. *Machine learning in non-stationary environments: Introduction to covariate shift adaptation*. MIT press.
- Tao, X.; Hong, X.; Chang, X.; Dong, S.; Wei, X.; and Gong, Y. 2020. Few-Shot Class-Incremental Learning. In *CVPR*.
- Tax, D. M.; and Duin, R. P. 2004. Support vector data description. *Machine Learning*, 54(1): 45–66.
- Xiao, H.; Rasul, K.; and Vollgraf, R. 2017. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*.
- Xu, Y.; Xu, C.; Xu, C.; and Tao, D. 2017. Multi-Positive and Unlabeled Learning. In *IJCAI*, 3182–3188.
- Zhang, Y.-J.; Zhao, P.; Ma, L.; and Zhou, Z.-H. 2020. An unbiased risk estimator for learning with augmented classes. In *NeurIPS*, 10247–10258.
- Zhang, Z.; and Sabuncu, M. 2018. Generalized cross entropy loss for training deep neural networks with noisy labels. In *NeurIPS*.
- Zhao, P.; Wang, X.; Xie, S.; Guo, L.; and Zhou, Z.-H. 2019. Distribution-free one-pass learning. *IEEE Transactions on Knowledge and Data Engineering*, 33(3): 951–963.
- Zhou, Z.-H.; and Chen, Z.-Q. 2002. Hybrid decision tree. *Knowledge-based systems*, 15(8): 515–528.