RR-PU: A Synergistic Two-Stage Positive and Unlabeled Learning Framework for Robust Tax Evasion Detection

Shuzhi Cao[*1,2], **Jianfei Ruan**[*1,2,†], **Bo Dong** [2,3], **Bin Shi** [1,2], **Qinghua Zheng** [1,2]

[1] School of Computer Science and Technology, Xi'an Jiaotong University, China
[2] Shaanxi Provincial Key Laboratory of Big Data Knowledge Engineering, Xi'an Jiaotong University, China
[3] School of Distance Education, Xi'an Jiaotong University, China
cao309615@gmail.com, jianfei.ruan@hotmail.com, {dong.bo,shibin,qhzheng}@xjtu.edu.cn

Abstract

Tax evasion, an unlawful practice in which taxpayers deliberately conceal information to avoid paying tax liabilities, poses significant challenges for tax authorities. Effective tax evasion detection is critical for assisting tax authorities in mitigating tax revenue loss. Recently, machine-learning-based methods, particularly those employing positive and unlabeled (PU) learning, have been adopted for tax evasion detection, achieving notable success. However, these methods exhibit two major practical limitations. First, their success heavily relies on the strong assumption that the label frequency (the fraction of identified taxpayers among tax evaders) is known in advance. Second, although some methods attempt to estimate label frequency using approaches like Mixture Proportion Estimation (MPE) without making any assumptions, they subsequently construct a classifier based on the error-prone label frequency obtained from the previous estimation. This two-stage approach may not be optimal, as it neglects error accumulation in classifier training resulting from the estimation bias in the first stage. To address these limitations, we propose a novel PU learning-based tax evasion detection framework called RR-PU, which can revise the bias in a two-stage synergistic manner. Specifically, RR-PU refines the label frequency initialization by leveraging a regrouping technique to fortify the MPE perspective. Subsequently, we integrate a trainable slack variable to fine-tune the initial label frequency, concurrently optimizing this variable and the classifier to eliminate latent bias in the initial stage. Experimental results on three real-world tax datasets demonstrate that RR-PU outperforms state-of-the-art methods in tax evasion detection tasks.

Introduction

Taxation, as the primary and indispensable source of national fiscal revenue, plays a crucial role in fostering national economic development. Compliance with legal provisions and fulfillment of tax obligations constitute the fundamental responsibilities of taxpayers. Unfortunately, a marked increase in corporate tax evasion has been observed recently, leading to significant adverse impacts on overall national fiscal revenue. Tax evasion, characterized by the intentional concealment and deceptive practices employed by taxpayers to circumvent their tax liabilities, represents an illicit behavior that necessitates immediate attention. Recent empirical research indicates that the extent of tax revenue loss in China is approximately 22% (Tian et al. 2016), while governments worldwide face an annual loss of nearly 500 billion US dollars (De Roux et al. 2018). Hence, it is of paramount importance to develop effective methodologies for the detection of tax evasion practices adopted by taxpayers.

Presently, conventional techniques for detecting tax evasion can be classified into three primary categories: manual case selection, whistle-blowing-based selection, and rule-based methods (Krivko 2010; Baesens, Van Vlasselaer, and Verbeke 2015; Wu et al. 2012; Zheng et al. 2023). However, these approaches suffer from inherent limitations. The first two methods demonstrate limited coverage, as they are unable to encompass all taxpayers, relying substantially on auditors' expertise, and consequently leading to time-consuming and inefficient processes (Ruan et al. 2019). On the other hand, rule-based methods utilize predefined rules to identify suspicious taxpayers. Regrettably, these rules gradually become outdated, rendering the approach less adaptable over time.

To mitigate the limitations of traditional techniques, recent advancements have incorporated machine learning-based approaches (Shi et al. 2023; Ruan et al. 2019; Wang et al. 2020; Zhang et al. 2020; Wu et al. 2019; Hemberg et al. 2016; Junqué de Fortuny et al. 2014; Abe et al. 2010). These approaches capitalize on fully labeled tax data to extract pertinent features and train tax evasion detection models utilizing machine learning techniques. In contrast to traditional methods, these approaches demonstrate notable performance improvements while diminishing reliance on human effort. Nonetheless, considering the extensive scale of tax data, obtaining fully labeled datasets for training becomes infeasible in real-world situations (Gao et al. 2021; Mi et al. 2020; Zhang et al. 2020). Consequently, only a small subset of instances are identified as tax evaders, while the majority of instances remain unidentified.

To adapt to real-world scenarios, various semi-supervised techniques have been proposed. Among these methods, positive and unlabeled learning (PU learning) (Wu et al. 2019; Zhang et al. 2020; Mi et al. 2020; Gao et al. 2021) has emerged as the most prevalent approach in the tax evasion

[*] These authors contributed equally.
[†] Corresponding author.

domain. PU learning techniques capitalize on the available labeled and unlabeled instances directly, obviating the need for additional manual annotation. As a special PU learning problem, tax evasion detection has its unique challenge, that is, the label frequency, representing the proportion of identified instances among tax evaders, is unknown. Nevertheless, contemporary PU learning algorithms rely largely on the assumption that the label frequency is known a priori. As a result, these methods cannot be directly applied to detect tax evasion behaviors. Therefore, developing a proprietary PU learning algorithm capable of automatically estimating label frequency to effectively detect tax evasion has become an urgent and critical problem.

To surmount this problem, we integrate the mathematical Mixture proportion estimation (MPE) technique with the PU learning method, which is typically used for anomaly detection, to estimate label frequency and finally achieve tax evasion detection. To be specific, we introduce a novel tax evasion detection framework called RR-PU (**R**egrouping-MPE and **R**evision-Based Synergistic **PU** Learning Framework), which capitalizes on the PU learning paradigm. RR-PU is devised as a synergistic two-stage framework that aims to address the latent estimation error of label frequency and concurrently facilitate classifier learning. The framework comprises the following stages: (1) Label frequency initialization stage: In this stage, we augment the label frequency initialization process by utilizing an enhanced MPE method predicated on the regrouping methodology; (2) Revision stage: In this stage, we introduce a slack variable to revise the less accurate initial label frequency estimation. By optimizing both the slack variable and the classifier concurrently, we aim to eliminate the latent estimation bias encountered in the first stage. This synergistic learning approach enables the refinement of the label frequency estimation while improving the overall performance of the classifier. RR-PU, as a potent tax auditing instrument, independent of fully labeled data or established label frequency, is proficient in surveilling tax evaders. While conserving labor, RR-PU amplifies governmental fiscal influx and underpins market regulation integrity. Our paper presents several significant contributions, which can be summarized as follows:

- RR-PU: We introduce a novel tax evasion detection framework called RR-PU, which shows its effectiveness in detecting tax evasion without requiring extra manual annotation. RR-PU capitalizes on a small proportion of positive instances along with a substantial amount of unlabeled instances, offering a labor-saving and efficient approach.
- Automatic label frequency estimation: In contrast to existing PU learning-based tax evasion detection methods, RR-PU automatically estimates the label frequency using tax data, obviating the need for extra assumptions. This feature enhances the adaptability of RR-PU in real-world scenarios
- Synergistic two-stage framework: RR-PU is designed as a synergistic two-stage positive and unlabeled learning framework. It effectively addresses the latent estimation error of label frequency while concurrently training a classifier. Compared to existing methods, RR-PU avoids the accumulation of errors, resulting in a superior classifier.
- Extensive experimental validation: We conduct experiments on three real-world tax datasets. The results validate the effectiveness of RR-PU in both estimating label frequency and detecting tax evasion behaviors. Moreover, our method outperforms the state-of-the-art (SOTA) tax evasion detection algorithms, establishing its superiority.

Related Work

Tax evasion detection Tax evasion detection methods can be categorized into four primary groups: manual case selection, whistle-blowing-based selection, rule-based methods (Tian et al. 2016; Krivko 2010; Baesens, Van Vlasselaer, and Verbeke 2015), and machine-learning-based methods (Shi et al. 2023; Ruan et al. 2019; Wang et al. 2020; Zhang et al. 2020; Mi et al. 2020; Gao et al. 2021). The first two methods involve the random selection of taxpayers for subsequent auditing (Tian et al. 2016). However, they do not encompass all taxpayers, and their effectiveness is heavily dependent on auditors' skills, rendering them time-consuming and inefficient (Ruan et al. 2019). As for rule-based methods, experts typically define rules or derive them from historical cases. When a taxpayer's behavior aligns with a defined rule, the rule-based system issues an alert. Nevertheless, maintaining and updating these rules can be challenging, causing the rule-based methods to lack adaptability. To address these issues, machine-learning-based methods have been recently introduced. In contrast to predefining specific rules, these methods automatically generalize tax evasion behavioral patterns from historical tax data. Consequently, they are not restricted to particular tax evasion behaviors and exhibit superior generalization capabilities. While these methods showcase strong performance, they necessitate a substantial volume of fully labeled instances for training, which can be challenging to acquire in reality. To mitigate this problem, some recent PU learning-based methods have been proposed, which train tax evasion detection models based on limited identified taxpayers and a large number of unidentified taxpayers.

PU learning Given limited positive instances and a large quantity of unlabeled instances as training data, the objective of PU learning is to train a classifier capable of distinguishing between positive and negative instances in the test data. Existing PU learning methods can be classified into three categories: two-step techniques (Liu et al. 2002; Li and Liu 2003), biased learning (Northcutt, Wu, and Chuang 2017; Patrini et al. 2016), and reweighting methods (Elkan and Noto 2008; Du Plessis, Niu, and Sugiyama 2015; Kiryo et al. 2017; Zhao et al. 2022). Two-step techniques typically identify reliable negative instances among unlabeled instances and then perform ordinary supervised learning. However, such algorithms rely heavily on heuristics, which can introduce additional errors. Biased learning methods treat unlabeled instances as negative instances with class label noise (Li et al. 2021). To mitigate interference from label noise, these methods often impose higher penalties on incorrectly classified positive instances. Nonetheless, penalties are closely associated with the positive class prior or label frequency, which is generally unknown in practice, rendering these methods challenging to implement. Reweighting methods assign varying weights to different instances, calibrating the inaccurate data distribution to a potentially correct one. However, these algorithms also perform poorly when the positive class prior or label frequency is unknown.

MPE The MPE problem is a statistical inference problem wherein, given data from a mixture and one of its two components, the objective is to identify the proportion of each component. Notable works in this area include the following: Blanchard et al. (Blanchard, Lee, and Scott 2010) first conducted a systematic study on the MPE problem, identifying it as an ill-posed problem without any assumptions. They introduced the irreducible assumption (Blanchard, Lee, and Scott 2010) to ensure a unique solution for the MPE problem. However, under this assumption, the convergence rate is exceedingly slow (Scott 2015). Subsequently, Ramaswamy et al. (Ramaswamy, Scott, and Tewari 2016) proposed the first computationally feasible algorithms, KM1 and KM2, to enhance the convergence rate by embedding the distributions into a reproducing kernel Hilbert space (RKHS) (Berlinet and Thomas-Agnan 2011). Regrettably, their methods falter when applied to high-dimensional data. To counteract this issue, Jain et al. (Jain et al. 2016) and Ivanov (Ivanov 2020) proposed AlphaMax and DEDPUL, respectively, both of which explore dimensional reduction techniques for handling high-dimensional data problems. Nevertheless, all these MPE methods rely on the irreducible assumption or its variants. In reality, the distribution of tax data is more complex and may not conform to these assumptions. Consequently, directly employing these conventional MPE methods for estimating label frequency is unreliable. Investigating an assumption-free MPE method for estimating label frequency in taxation scenarios is essential.

Definition and Problem Formulation

In this section, we first present key definitions pertinent to the tax scenario, followed by the systematic formulations of tax evasion detection, PU learning, and the MPE problem.

Definition 1: Compliant Taxpayer. This term refers to a taxpayer who adheres to tax laws and regulations. In the context of this study, Compliant Taxpayers are designated as negative instances.

Definition 2: Tax Evader. A taxpayer who deliberately infringes upon tax laws by engaging in fraudulent practices, concealment, or other illicit activities to evade tax payments is classified as a Tax Evader. For the purpose of this study, Tax Evaders are identified as positive instances.

Definition 3: Identified Taxpayer. In real-world tax auditing scenarios, only a small proportion of Tax Evaders are distinctly identified as such, owing to the substantial cost of labeling. In the framework of PU learning, these individuals are regarded as positive instances.

Definition 4: Unidentified Taxpayer. This categorization refers to taxpayers who have not been labeled by tax auditors. An Unidentified Taxpayer could either be a Compliant Taxpayer or a Tax Evader. In PU learning, Unidentified Taxpayers are denoted as unlabeled instances.

Formulation of Tax Evasion Detection Problem

In the context of tax evasion, a small subset of tax evaders is identified, with a considerable majority of taxpayers remaining unidentified. We designate $X \in \mathbb{R}^d$ as the variable for instances and $Y \in \mathbb{R}$ as the variable for labels. Moreover, $\mathcal{X}$ is defined as the feature space, and $\mathcal{Y} = \{0,1\}$ is the label space. An instance $x \in \mathcal{X}$ signifies a taxpayer in real-world taxation scenarios. If $y = 1$, taxpayer x is classified as a tax evader (positive instance), and if $y = 0$, as a compliant taxpayer (negative instances). S is introduced as a binary variable indicating whether an instance x is identified; in other words, a taxpayer x is identified if $s = 1$ and unidentified if $s = 0$. Given the general unavailability of the true label y, the aim of tax evasion detection is to identify tax evaders based on a limited number of identified taxpayers and a large pool of unidentified taxpayers.

Formulation of PU Learning Problem

In this context, we designate $\mathcal{S}_P = \{x_i^p, s_i^p = 1\}_{i=1}^{n_p}$ as the ensemble of identified taxpayers and $\mathcal{S}_U = \{x_i^u, s_i^u = 0\}_{i=1}^{n_u}$ as the collective of unidentified taxpayers. Here, n_p and n_u represent the quantities of identified and unidentified taxpayers, respectively. Within the framework of PU learning, $\mathcal{S}_P$ and $\mathcal{S}_U$ are correspondingly defined as the positive and unlabeled data. The primary objective of PU learning is to develop a binary classifier, which is proficient in predicting the posterior probability $P(Y|X = x) = [P(Y = 0|X = x), P(Y = 1|X = x)]^\top$, based on the union of $\mathcal{S}_P$ and $\mathcal{S}_U$. In essence, this probability indicates whether a given taxpayer x is a tax evader. PU learning typically utilizes one of two problem settings, each dependent on the data sampling methodology employed: the single-training-set setting (Gong et al. 2019) and the case-control setting (Niu et al. 2016; Bekker and Davis 2020). Under the single-training-set setting, it is postulated that all unlabeled training instances are sampled from the marginal density $p(x)$. If a given instance x is positive, its positive label is discerned with a probability c, leaving x unlabeled with a probability of $1 - c$. Herein, $c = P(S = 1|Y = 1)$ symbolizes the label frequency. In contrast, should x be negative, the negative label remains unobserved, thereby consistently leaving x unlabeled. Contrarily, the case-control setting postulates that positive instances and unlabeled instances are independently sampled from the marginal densities $p(x|Y = 1)$ and $p(x)$, respectively,signified by $S_P \sim p(x|Y = 1)$ and $S_U \sim p(x)$. In real-world tax scenarios, an exhaustive collection of taxpayer information is initially gathered, after which a portion of taxpayers is designated as tax evaders by auditors. This process of data generation mirrors the methodology of the single-training-set setting, thereby justifying its selection for use in this paper.

Formulation of MPE Problem

Problem definition Given two distributions G and H over a metric space $\mathcal{X}$, and a parameter $\kappa \in (0, 1)$, let F be a convex combination of G and H, i.e., $F = (1 - \kappa)G + \kappa H$. The MPE problem involves determining κ from instances S_F and S_H i.i.d. drawn from the mixture distribution F and the distribution H, respectively. In the tax evasion framework, the known distributions F and H correspond to the probability density functions of unidentified taxpayers and tax evaders, respectively, while the unknown distribution G characterizes the probability density function of compliant taxpayers. The parameter κ represents the fraction of tax evaders among unidentified taxpayers. Upon determination of κ, the label frequency c is readily computable, i.e., $c = |P| / (\kappa |U| + |P|)$, where $|P|$ and $|U|$ denote the number of identified and unidentified taxpayers, respectively. We denote $\kappa(F|H) \triangleq sup \{\kappa | F = (1-\kappa)G + \kappa H\}$ as the maximum proportion of H in F. Consequently, any $\kappa \in (0, \kappa(F|H))$ could serve as a feasible solution to the MPE problem (Yao et al. 2020). Without introducing certain assumptions, the MPE problem is ill-posed. To make κ identifiable, a number of propositions concerning distribution G have been presented. Presently, the irreducible assumption (Yao et al. 2021), which is defined as follows, stands as the least constraining of these propositions:

Definition 5: Irreducible Assumption. The distribution G is deemed irreducible with respect to the distribution H if G does not contain H in its mixture. Formally, this condition implies that the decomposition $G = (1-\beta)Q + \beta H$ does not exist, where Q denotes a distribution over the metric space $\mathcal{X}$ and $0 < \beta \leq 1$.

Assuming that the unknown distribution G is irreducible with respect to H, the parameter κ will converge to its supremum, represented as $\kappa(F|H)$. This condition enables the identifiability of the MPE problem under the irreducible assumption. Extant MPE methodologies, encompassing EN (Elkan and Noto 2008), KM (Ramaswamy, Scott, and Tewari 2016), ROC (Scott 2015), AlphaMax (Jain et al. 2016), and DEDPUL (Ivanov 2020), aim to estimate $\kappa(F|H)$ to approximate the true κ. The irreducible assumption, when valid, results in $\kappa(F|H)$ being a robust approximation. However, any violation of this assumption could lead to significant estimation errors.

Regrouping-MPE Despite its significance, the irreducible assumption remains challenging to validate due to the unobservable nature of distribution G. This influences the effectiveness of MPE algorithms. To address this issue, the Regrouping-MPE method (Yao et al. 2020) has recently been introduced. This method ingeniously addresses the violation of the irreducible assumption by formulating an entirely new MPE problem, in which a new component distribution H'—generated through a regrouping technique—complies with the irreducible assumption. In this restructured context, the solution to the new MPE problem expressed as κ', aligns with the original solution, κ. This alignment allows traditional MPE methodologies to resolve the new MPE problem with diminished bias, regardless of whether the original MPE problem adheres to the irreducible assumption or not. While the Regrouping-MPE method is successful in reducing the estimation bias of κ, it continues to construct a classifier based on the label frequency, which is potentially error-prone due to it being derived from the estimated κ. This two-stage procedure may not yield optimal results, as it doesn't account for error accumulation during classifier training that stems from the estimation bias in the initial stage. Hence, the direct application of this algorithm in tax evasion may provide unreliable results. These challenges emphasize the need for a more robust algorithm.

Proposed Method

This section provides an in-depth examination of the RR-PU method, which fundamentally incorporates two principal stages: 1) Label Frequency Initialization, and 2) Revision.

Label Frequency Initialization Stage

In this context, we denote the probability density function of positive (tax evaders), negative (compliant taxpayers), and unlabeled (unidentified taxpayers) instance spaces as $f_P(x)$, $f_N(x)$, and $f_U(x)$, respectively. Given the single-training-set configuration , we observe the following equation:

$$f_U(x) = \theta_+ f_P(x) + (1-\theta_+) f_N(x), \tag{1}$$

Herein, θ_+ represents the ratio of positive instances (tax evaders) within the unlabeled instance (unidentified taxpayers). Following the division of the PU training data, denoted S_{tr}, into positive data S_P and unlabeled data S_U, the task of estimating θ_+ from Eq. (1) becomes an MPE problem. To solve this, we employ the Regrouping-MPE method, using S_P and S_U as inputs to the algorithm. Notably, by duplicating $p \times |S_U|$ instances with low negative class-posterior probability from S_U to S_P, we generate a novel regrouped distribution, P', which is irreducible relative to distribution N. A standard MPE solver is subsequently utilized to estimate $\kappa(U|P')$, which approximates θ_+ with $S_{P'}$ (i.i.d drawn from distribution P') and S_U. Finally, we derive the estimator $\hat{\theta}_+$, which is defined as:

$$\hat{\theta}_+ = \kappa(U|P') = \underset{f_{P'}(x)>0}{inf} \frac{f_U(x)}{f_{P'}(x)}, \tag{2}$$

Upon establishing θ_+, the positive class prior—denoted as $\pi_+ = P(Y=1)$ and indicating the ratio of tax evaders amongst all taxpayers—can be formulated as $(\theta_+ |S_U| + |S_P|)/(|S_U| + |S_P|)$. In this context, $|S_P|$ and $|S_U|$ correspond to the volumes of the positive and unlabeled data, respectively. Further, in the framework of PU learning, the label frequency—denoted as c—can be expressed as:

$$c = P(S=1|Y=1) = \frac{P(S=1, Y=1)}{P(Y=1)} = \frac{P(S=1)}{\pi_+}, \tag{3}$$

The relationship encapsulated in Eq. (3) is valid under the premise that all labeled instances (identified taxpayers) are inherently positive instances (tax evaders), i.e., $P(S = 1, Y = 1) = P(S = 1)$. Here, $P(S = 1) = |S_P| / (|S_P| + |S_U|)$ represents the fraction of labeled instances among the total instances. As a result, the initial label frequency, denoted as $\hat{c}$, can be further defined as:

$$\hat{c} = \frac{|S_P|}{|S_P| + \hat{\theta}_+ |S_U|}. \tag{4}$$

Revision Stage

Consider $g(x)$ as the predictive function, producing the actual class posterior probability $P(Y|X = x) = [P(Y = 0|X = x), P(Y = 1|X = x)]^\top$, thereby determining if a taxpayer x is a tax evader. Let $q(x)$ be another predictive function, approximating the posterior probability $P(S|X = x) = [P(S = 0|X = x), P(S = 1|X = x)]^\top$, ascertaining whether a taxpayer x is labeled. Define $f(x) = \arg\max_{j\in\{0,1\}} g_j(x)$ as the decision function, where $g_j(x)$ represents an estimation of $P(Y = j|X = x)$, and $h(x) = \arg\max_{j\in\{0,1\}} q_j(x)$ as the decision function, where $q_j(x)$ estimates $P(S = j|X = x)$. Allow $T(X = x)$ to symbolize the transition matrix where $T_{ij}(X = x) = P(S = j|Y = i, X = x)$, from which $q(x)$ can be deduced, i.e., $T(X = x)^\top g(x) = q(x)$. In real-world taxation scenarios, auditors typically select a random subset for further auditing, making the labeling of a tax evader entirely arbitrary. Thus, we adopt the Selected Completely At Random (SCAR) assumption (Li et al. 2019a), which asserts that the label frequency is independent of instance features, i.e., $P(S = j|Y = i, X = x) = P(S = j|Y = i), \forall i, j \in \{0, 1\}$. Under this premise, the transition matrix estimator, denoted as $\hat{T}$, can be simplified accordingly as follows:

$$\hat{T} = \begin{bmatrix} 1 & 0 \\ 1-\hat{c} & \hat{c} \end{bmatrix}. \tag{5}$$

Following Eq. (5), we infer that the determination of the initial label frequency $\hat{c}$ precedes the identification of the transition matrix $\hat{T}$. Consequently, we utilize $\hat{T}^\top g(x)$ as an approximation of $q(x)$, denoted as $q(x) \approx \hat{q}(x) = \hat{T}^\top g(x)$. Keeping the transition matrix $\hat{T}$ constant, the backbone network $g(x)$ is updated by minimizing the unweighted risk $\mathcal{R}_{unweighted}(g)$, defined as:

$$\mathcal{R}_{unweighted}(g) = \mathbb{E}[\ell(\hat{q}(x), s)] = \mathbb{E}[\ell(\hat{T}^\top g(x), s)], \tag{6}$$

Here $\mathbb{E}$ represents the expectation over the joint density $p(x, s)$, and $\ell : \mathbb{R} \times \{0, 1\} \longrightarrow \mathbb{R}$ is a specific loss function. We represent the distributions for PU data and clean data (the fully labeled data) as $\mathcal{D}_{PU}$ and $\mathcal{D}$, respectively. Leveraging the importance reweighting technique (Liu and Tao 2015), we reformulate the expected risk associated with distribution $\mathcal{D}$ as:

$$\begin{aligned} R_{\ell,\mathcal{D}}(f) &= \mathbb{E}_{(X,Y)\sim\mathcal{D}}[\ell(f(X), Y)] \\ &= \mathbb{E}_{(X,S)\sim\mathcal{D}_{PU}}[\frac{P_{\mathcal{D}}(X,Y)}{P_{\mathcal{D}_{PU}}(X,S)}\ell(f(X), S)] \\ &= \mathbb{E}_{(X,S)\sim\mathcal{D}_{PU}}[\frac{P_{\mathcal{D}}(Y|X)P_{\mathcal{D}}(X)}{P_{\mathcal{D}_{PU}}(S|X)P_{\mathcal{D}_{PU}}(X)}\ell(f(X), S)] \\ &= \mathbb{E}_{(X,S)\sim\mathcal{D}_{PU}}[\frac{P_{\mathcal{D}}(Y|X)}{P_{\mathcal{D}_{PU}}(S|X)}\ell(f(X), S)]. \end{aligned} \tag{7}$$

Eq. (7) establishes that the expected risk corresponding to clean data and the loss $\ell(f(X), Y)$ is comparable to an expected risk linked to PU data and a weighted loss. Once the transition matrix T is identified, we can express the weighted risk in Eq. (7) as:

$$\mathcal{R}_{weighted}(T, f) = \mathbb{E}_{(X,S)\sim\mathcal{D}_{PU}}[\frac{g_{\mathcal{D}_{PU}}(X)}{(T^\top g)_{\mathcal{D}_{PU}}(X)}\ell(f(X), S)]. \tag{8}$$

It is critical to recognize the potential estimation error between the initial label frequency and its true value.

Algorithm 1: RR-PU

Input: PU training data S_{tr} and PU validation data S_v.
Output: Binary classifier f and the estimator of label frequency $\hat{c}$.

1: Split the PU training data S_{tr} into positive data S_P and unlabeled data S_U.
2: Take S_U, S_P as inputs S_F and S_H in Regrouping-MPE to estimate label frequency $\hat{c}$.
3: Initialize the transition matrix $\hat{T}$ and minimize the $\mathcal{R}_{unweighted}(g)$ to optimize the backbone network g while keeping the transition matrix $\hat{T}$ fixed.
4: Minimize $R_{weighted}(\hat{T} + \Delta T, f)$ to learn ΔT and f simultaneously. // Stopping criterion: when the $\hat{P}(S|X = x)$ yields the minimum classification error on validation set S_v.
5: Update the transition matrix $\hat{T} \leftarrow \hat{T} + \Delta T$ and then perform row normalization.
6: Update the label frequency by $\hat{c} \leftarrow \hat{T}_{11}$.

To enhance the precision of the estimated label frequency, we introduce a slack variable ΔT, substituting the transition matrix $\hat{T}$ with $(\hat{T} + \Delta T)$ in Eq. (8). By minimizing $R_{weighted}(\hat{T} + \Delta T, f)$, we achieve synergistic optimization of the backbone network g and the slack variable ΔT. This methodology proves effective as it minimizes the weighted risk, which is asymptotically identical to the expected risk on clean data, leading to a more robust classifier $\hat{P}(Y|X)$. Concurrently, we validate the slack variable on the validation set, ensuring $\hat{P}(S|X)$ aligns with the validation set. More precise $\hat{P}(Y|X)$ and $\hat{P}(S|X)$ facilitate the estimation of ΔT, further mitigating bias in the initial stage. We detail the implementation of the RR-PU in Algorithm 1, with a visual representation provided in Fig 1.

Experiments

Datasets and evaluation metrics In light of the absence of a standard public tax dataset for method evaluation, we procured raw tax data from value-added invoices and taxpayer registration details gathered by tax bureaus spanning diverse regions in China. Basic features such as tax amount, long-term debt, and profit ratio were extracted from the taxpayer registration data, and transaction features were derived from the value-added invoices via PnCGCN (Gao et al. 2021). Consequently, we assembled three real-world tax datasets - TaxS, TaxH, and TaxZ, with all instances being fully labeled. The datasets were divided into training, validation, and test sets following a 3:1:1 ratio. To simulate actual taxation scenarios, we randomly picked 50 percent of the original positive instances, combining them with all negative instances to create unlabeled instances in the training and validation sets. For the test set, instances were assigned their ground truth labels. Comprehensive details about the datasets are further discussed in the Appendix. The adopted evaluation metrics for our experiments include Accuracy, F1-score, and AUC, with their respective definitions outlined in the Appendix.

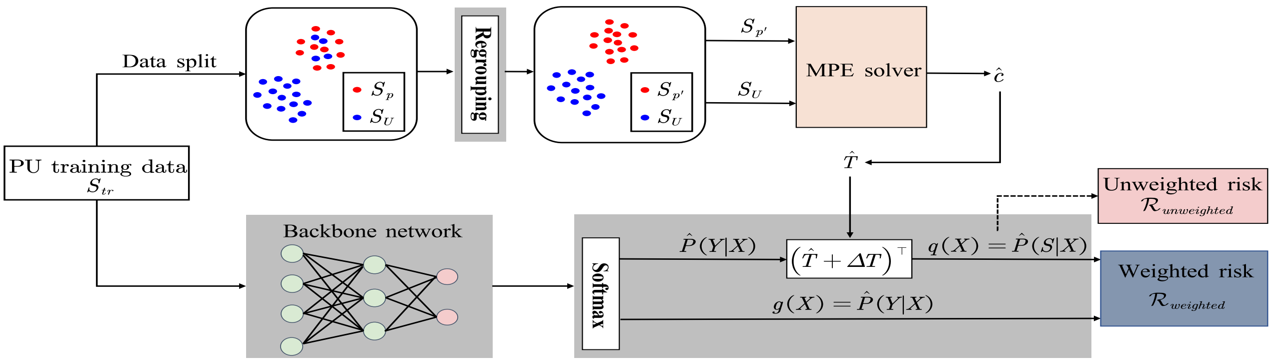

Figure 1: The overview of the proposed algorithm.

Comparative methods and experimental setup RR-PU was compared with contemporary state-of-the-art approaches. Specifically, for label frequency estimation, RR-PU was juxtaposed with MPE methods such as EN (Elkan and Noto 2008), KM1 and KM2 (Ramaswamy, Scott, and Tewari 2016), AlphaMax (AM) (Jain et al. 2016), and DEDPUL (DP) (Ivanov 2020). Furthermore, to evaluate our algorithm's efficacy in detecting tax evasion behaviors, we incorporated the leading tax evasion detection methods: FBNE-PU (Gao et al. 2021), Eagle (Shi et al. 2023) and related PU-learning algorithms including Biased PU (Li et al. 2019b), nnPU (Kiryo et al. 2017), uPU (Du Plessis, Niu, and Sugiyama 2015), RankPruning (Northcutt, Wu, and Chuang 2017), WSVM (Elkan and Noto 2008), VPU (Chen et al. 2020) and Dist-PU (Zhao et al. 2022) for comparison. To guarantee an equitable comparison, recommended hyper-parameters were employed for all the comparison methods. Alongside, we set the hyper-parameter p at 10% in the Regrouping-MPE, a value confirmed as optimal for minimizing estimation error (Yao et al. 2020). A five-layer multi-layer perceptron (MLP) served as the backbone network for all the methods. In the revision stage, the network was trained utilizing the Adam method with a learning rate of $5e-5$, weight decay of $3e-4$, and batch size of 128. All the experiments were executed using Pytorch on two GPUs (NVIDIA RTX 3090) functioning in parallel.

Experimental Results

Label frequency estimation RR-PU leverages the regrouping technique to fortify the MPE method and introduces a revision stage to bolster the estimation of label frequency. To corroborate the effectiveness of this two-stage PU learning framework, five SOTA MPE methods are utilized as baseline procedures (base). Simultaneously, experimental analysis is conducted on the baseline, augmented with the regrouping technique (with-R), and baseline supplemented with both the regrouping technique and the revision stage (with-RR). These methods are analyzed concerning the estimation error of label frequency. Fig. 2 encapsulates the outcomes of different methods on TaxS, TaxH, and TaxZ, where the X-axis and Y-axis denote the baselines and estimation error, respectively. Compared with baselines, the regrouping technique yields a diminished estimation error, which further contracts upon the integration of the revision stage. Thus, both the regrouping technique and the revision stage are integral to the enhancement of estimation accuracy.

Tax evasion detection To evaluate the performance of RR-PU in identifying tax evaders, one tax evasion detection method and six PU learning methods mentioned above, are selected as comparison baselines. Considering that the comparison methods such as FBNE-PU, uPU, and nnPU require a known positive class prior and are all designed under a case-control setting, the KM1 estimator is employed in advance to estimate the positive class prior, ensuring a fair comparison. Subsequently, the positive data is reinserted into the unlabeled set to comply with the sampling requirements of these methods. The experimental results are exhibited in Table 1, with the best results emphasized in bold. The results indicate that RR-PU outperforms the baselines in detecting tax evasion behaviors across all datasets. More-

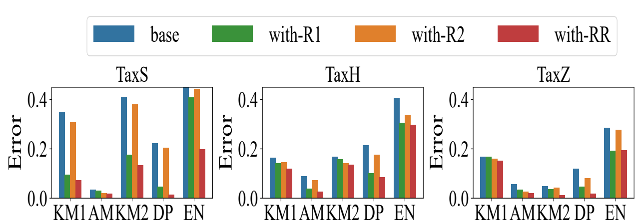

Figure 2: Label frequency estimation error with different methods on TaxS, TaxH, and TaxZ.

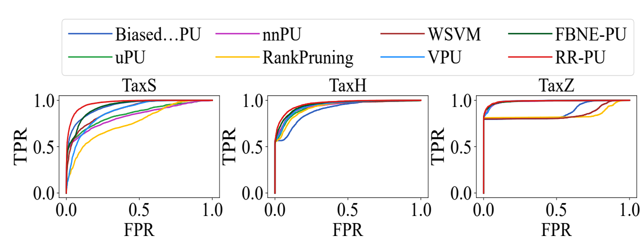

Figure 3: The ROC curves of different methods on TaxS, TaxH, and TaxZ.

Method	Accuracy (%)			F1-score			AUC		
	TaxS	TaxH	TaxZ	TaxS	TaxH	TaxZ	TaxS	TaxH	TaxZ
Biased PU	74.081	92.662	95.520	0.5612	0.7014	0.8763	0.9411	0.9027	0.8754
nnPU	70.357	92.569	95.751	0.5804	0.7182	0.8816	0.8309	0.9170	0.9831
uPU	70.644	92.730	95.231	0.4881	0.7090	0.8795	0.8424	0.9272	0.9827
RankPruning	71.867	92.327	95.095	0.6184	0.7065	0.8871	0.7624	0.9310	0.9454
WSVM	77.232	92.415	93.968	0.7668	0.7056	0.8037	0.8946	0.9426	0.8523
VPU	79.365	92.721	95.412	0.7679	0.7192	0.8941	0.8757	0.9448	0.9890
FBNE-PU	80.403	93.327	95.456	0.7325	0.6921	0.8850	0.9334	0.9547	0.9891
Eagle	84.573	92.377	95.642	0.7883	0.7197	0.8966	0.9376	0.9582	0.9843
Dist-PU	86.719	92.851	95.813	0.8159	0.7337	0.8965	0.9511	0.9514	0.9735
RR-PU	**90.286**	**93.658**	**96.726**	**0.9088**	**0.7834**	**0.9182**	**0.9742**	**0.9643**	**0.9903**

Table 1: Tax evasion detection results on the real-world datasets

Method	Accuracy (%)			F1-score		
	TaxS	TaxH	TaxZ	TaxS	TaxH	TaxZ
base	87.611	92.471	95.615	0.8536	0.7155	0.8893
with-R	89.451	92.597	96.161	0.8749	0.7268	0.9067
with-RR	**90.847**	**93.517**	**96.734**	**0.8938**	**0.7645**	**0.9219**

Table 2: The results of ablation experiments

over, receiver operating characteristic (ROC) curves of different methods are illustrated in Fig. 3, wherein the X-axis and Y-axis represent the False Positive Rate (FPR) and True Positive Rate (TPR) respectively. As depicted in Fig. 3, the ROC curve corresponding to RR-PU covers the largest area. These results show the superiority of RR-PU in tax evasion detection.

Ablation experiments We conducted ablation experiments to underscore the significance of both the regrouping technique and the revision stage in augmenting classifier performance. Specifically, AlphaMax, which yields the minimum estimation error (as per Fig.2), was chosen as the base MPE method. Adhering to prior definitions, the performance of the classifier was assessed on the baseline, 'with-R', and 'with-RR', in terms of accuracy and F1-score. The results are delineated in Table 2. The findings show that the regrouping technique refines the classifier, and its performance is further amplified with the revision stage.

Analysis of Results

The results reveal that RR-PU surpasses the best extant method by 3.567%, 0.0929, and 0.0231 in terms of accuracy, F1-score, and AUC on TaxS, respectively. These enhancements are markedly more considerable than those observed on TaxZ and TaxH. To comprehend the mechanism driving these results, we employed the t-SNE technique to project the tax data onto a two-dimensional plane, thereby visualizing its distribution (as depicted in the Appendix). The pronounced advancement on the TaxS dataset can be elucidated by the following aspects: **1) Breach of irreducibility:** The TaxS dataset exhibits notable overlap between positive and unlabeled instances, leading to severe violation of the irreducibility assumption by their corresponding density functions. Direct implementation of MPE methods for estimating the positive class prior, devoid of any modifications, results in substantial estimation errors, causing performance deterioration in methods such as nnPU, uPU, and FBNE-PU; **2) Data separability challenges:** The poor separability of data complicates the development of a non-traditional classifier (NTC) that outputs the posterior probability $P(s = 1|x)$. For comparison methods that rely on NTC, such as Biased PU, WSVM, RankPruning, and VPU, their performances are significantly affected. Conversely, RR-PU employs the regrouping technique to lessen the estimation error caused by the irreducibility violation and introduces a revision stage to revise the estimator. Hence, RR-PU is less impacted by these adverse factors and achieves superior performance.

Conclusion

In this work, we propose a novel synergistic two-stage PU learning framework RR-PU for robust detection of tax evasion. The first stage employs an enhanced MPE method, exploiting the regrouping technique to initialize the label frequency. Subsequently, RR-PU introduces a slack variable to revise the initially estimated label frequency, simultaneously optimizing this slack variable and the classifier to mitigate potential bias. Extensive experiments show that RR-PU outperforms a range of comparison methods in tax evasion detection. Despite its advantages, RR-PU faces limitations due to its MPE-based design that inherently relies on kernel density estimation. This reliance becomes a significant challenge when dealing with high-dimensional input, as it can lead to inflated estimation errors that adversely affect RR-PU's effectiveness. In the future, we aim to address these limitations and explore advanced interpretable models for tax evasion detection, striving to provide understandable evidence from the model's findings.

Acknowledgments

This research was partially supported by the Key Research and Development Project in Shaanxi Province No. 2022GXLH-01-03, the National Science Foundation of China under Grant Nos. 62002282, 62250009, and 6219278, and the Major Technological Innovation Project of Hangzhou No. 2022AIZD0113.

References

Abe, N.; Melville, P.; Pendus, C.; Reddy, C. K.; Jensen, D. L.; Thomas, V. P.; Bennett, J. J.; Anderson, G. F.; Cooley, B. R.; Kowalczyk, M.; et al. 2010. Optimizing debt collections using constrained reinforcement learning. In *Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining*, 75–84.

Baesens, B.; Van Vlasselaer, V.; and Verbeke, W. 2015. *Fraud analytics using descriptive, predictive, and social network techniques: a guide to data science for fraud detection*. John Wiley & Sons.

Bekker, J.; and Davis, J. 2020. Learning from positive and unlabeled data: A survey. *Machine Learning*, 109(4): 719–760.

Berlinet, A.; and Thomas-Agnan, C. 2011. *Reproducing kernel Hilbert spaces in probability and statistics*. Springer Science & Business Media.

Blanchard, G.; Lee, G.; and Scott, C. 2010. Semi-supervised novelty detection. *The Journal of Machine Learning Research*, 11: 2973–3009.

Chen, H.; Liu, F.; Wang, Y.; Zhao, L.; and Wu, H. 2020. A variational approach for learning from positive and unlabeled data. *Advances in Neural Information Processing Systems*, 33: 14844–14854.

De Roux, D.; Perez, B.; Moreno, A.; Villamil, M. d. P.; and Figueroa, C. 2018. Tax fraud detection for under-reporting declarations using an unsupervised machine learning approach. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 215–222.

Du Plessis, M.; Niu, G.; and Sugiyama, M. 2015. Convex formulation for learning from positive and unlabeled data. In *International conference on machine learning*, 1386–1394. PMLR.

Elkan, C.; and Noto, K. 2008. Learning classifiers from only positive and unlabeled data. In *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*, 213–220.

Gao, Y.; Shi, B.; Dong, B.; Wang, Y.; Mi, L.; and Zheng, Q. 2021. Tax evasion detection with FBNE-PU algorithm based on PnCGCN and PU learning. *IEEE Transactions on Knowledge and Data Engineering*.

Gong, C.; Shi, H.; Liu, T.; Zhang, C.; Yang, J.; and Tao, D. 2019. Loss decomposition and centroid estimation for positive and unlabeled learning. *IEEE transactions on pattern analysis and machine intelligence*, 43(3): 918–932.

Hemberg, E.; Rosen, J.; Warner, G.; Wijesinghe, S.; and O'Reilly, U.-M. 2016. Detecting tax evasion: a coevolutionary approach. *Artificial Intelligence and Law*, 24(2): 149–182.

Ivanov, D. 2020. Dedpul: Difference-of-estimated-densities-based positive-unlabeled learning. In *2020 19th IEEE International Conference on Machine Learning and Applications (ICMLA)*, 782–790. IEEE.

Jain, S.; White, M.; Trosset, M. W.; and Radivojac, P. 2016. Nonparametric semi-supervised learning of class proportions. *arXiv preprint arXiv:1601.01944*.

Junqué de Fortuny, E.; Stankova, M.; Moeyersoms, J.; Minnaert, B.; Provost, F.; and Martens, D. 2014. Corporate residence fraud detection. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, 1650–1659.

Kiryo, R.; Niu, G.; Du Plessis, M. C.; and Sugiyama, M. 2017. Positive-unlabeled learning with non-negative risk estimator. *Advances in neural information processing systems*, 30.

Krivko, M. 2010. A hybrid model for plastic card fraud detection systems. *Expert Systems with Applications*, 37(8): 6070–6076.

Li, T.; Wang, C.-C.; Ma, Y.; Ortal, P.; Zhao, Q.; Stenger, B.; and Hirate, Y. 2019a. Learning classifiers on positive and unlabeled data with policy gradient. In *2019 IEEE International Conference on Data Mining (ICDM)*, 399–408. IEEE.

Li, T.; Wang, C.-C.; Ma, Y.; Ortal, P.; Zhao, Q.; Stenger, B.; and Hirate, Y. 2019b. Learning classifiers on positive and unlabeled data with policy gradient. In *2019 IEEE International Conference on Data Mining (ICDM)*, 399–408. IEEE.

Li, X.; and Liu, B. 2003. Learning to classify texts using positive and unlabeled data. In *IJCAI*, volume 3, 587–592. Citeseer.

Li, X.; Liu, T.; Han, B.; Niu, G.; and Sugiyama, M. 2021. Provably end-to-end label-noise learning without anchor points. In *International Conference on Machine Learning*, 6403–6413. PMLR.

Liu, B.; Lee, W. S.; Yu, P. S.; and Li, X. 2002. Partially supervised classification of text documents. In *ICML*, volume 2, 387–394. Sydney, NSW.

Liu, T.; and Tao, D. 2015. Classification with noisy labels by importance reweighting. *IEEE Transactions on pattern analysis and machine intelligence*, 38(3): 447–461.

Mi, L.; Dong, B.; Shi, B.; and Zheng, Q. 2020. A tax evasion detection method based on positive and unlabeled learning with network embedding features. In *International Conference on Neural Information Processing*, 140–151. Springer.

Niu, G.; du Plessis, M. C.; Sakai, T.; Ma, Y.; and Sugiyama, M. 2016. Theoretical comparisons of positive-unlabeled learning against positive-negative learning. *Advances in neural information processing systems*, 29.

Northcutt, C. G.; Wu, T.; and Chuang, I. L. 2017. Learning with confident examples: Rank pruning for robust classification with noisy labels. *arXiv preprint arXiv:1705.01936*.

Patrini, G.; Nielsen, F.; Nock, R.; and Carioni, M. 2016. Loss factorization, weakly supervised learning and label noise robustness. In *International conference on machine learning*, 708–717. PMLR.

Ramaswamy, H.; Scott, C.; and Tewari, A. 2016. Mixture proportion estimation via kernel embeddings of distributions. In *International conference on machine learning*, 2052–2060. PMLR.

Ruan, J.; Yan, Z.; Dong, B.; Zheng, Q.; and Qian, B. 2019. Identifying suspicious groups of affiliated-transaction-based tax evasion in big data. *Information Sciences*, 477: 508–532.

Scott, C. 2015. A rate of convergence for mixture proportion estimation, with application to learning from noisy labels. In *Artificial Intelligence and Statistics*, 838–846. PMLR.

Shi, B.; Dong, B.; Xu, Y.; Wang, J.; Wang, Y.; and Zheng, Q. 2023. An edge feature aware heterogeneous graph neural network model to support tax evasion detection. *Expert Systems with Applications*, 213: 118903.

Tian, F.; Lan, T.; Chao, K.-M.; Godwin, N.; Zheng, Q.; Shah, N.; and Zhang, F. 2016. Mining suspicious tax evasion groups in big data. *IEEE Transactions on Knowledge and Data Engineering*, 28(10): 2651–2664.

Wang, Y.; Zheng, Q.; Ruan, J.; Gao, Y.; Chen, Y.; Li, X.; and Dong, B. 2020. T-egat: A temporal edge enhanced graph attention network for tax evasion detection. In *2020 IEEE International Conference on Big Data (Big Data)*, 1410–1415. IEEE.

Wu, R.-S.; Ou, C.-S.; Lin, H.-y.; Chang, S.-I.; and Yen, D. C. 2012. Using data mining technique to enhance tax evasion detection performance. *Expert Systems with Applications*, 39(10): 8769–8777.

Wu, Y.; Zheng, Q.; Gao, Y.; Dong, B.; Wei, R.; Zhang, F.; and He, H. 2019. Tedm-pu: A tax evasion detection method based on positive and unlabeled learning. In *2019 IEEE International Conference on Big Data (Big Data)*, 1681–1686. IEEE.

Yao, Y.; Liu, T.; Han, B.; Gong, M.; Niu, G.; Sugiyama, M.; and Tao, D. 2020. Towards mixture proportion estimation without irreducibility. *arXiv preprint arXiv:2002.03673*.

Yao, Y.; Liu, T.; Han, B.; Gong, M.; Niu, G.; Sugiyama, M.; and Tao, D. 2021. Rethinking class-prior estimation for positive-unlabeled learning. In *International Conference on Learning Representations*.

Zhang, F.; Shi, B.; Dong, B.; Zheng, Q.; and Ji, X. 2020. TTED-PU: A Transferable Tax Evasion Detection Method Based on Positive and Unlabeled Learning. In *2020 IEEE 44th Annual Computers, Software, and Applications Conference (COMPSAC)*, 207–216. IEEE.

Zhao, Y.; Xu, Q.; Jiang, Y.; Wen, P.; and Huang, Q. 2022. Dist-pu: Positive-unlabeled learning from a label distribution perspective. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14461–14470.

Zheng, Q.; Xu, Y.; Liu, H.; Shi, B.; Wang, J.; and Dong, B. 2023. A Survey of Tax Risk Detection Using Data Mining Techniques. *Engineering*.