
Binary Classification from Positive-Confidence Data

Takashi Ishida^{1,2} Gang Niu² Masashi Sugiyama^{2,1}

¹ The University of Tokyo, Tokyo, Japan

² RIKEN, Tokyo, Japan

{ishida@ms., sugi@}k.u-tokyo.ac.jp, gang.niu@riken.jp

Abstract

Can we learn a binary classifier from only positive data, without any negative data or unlabeled data? We show that if one can equip positive data with confidence (*positive-confidence*), one can successfully learn a binary classifier, which we name *positive-confidence (Pconf) classification*. Our work is related to *one-class classification* which is aimed at “describing” the positive class by clustering-related methods, but one-class classification does not have the ability to tune hyper-parameters and their aim is not on “discriminating” positive and negative classes. For the Pconf classification problem, we provide a simple empirical risk minimization framework that is model-independent and optimization-independent. We theoretically establish the consistency and an estimation error bound, and demonstrate the usefulness of the proposed method for training deep neural networks through experiments.

1 Introduction

Machine learning with big labeled data has been highly successful in applications such as image recognition, speech recognition, recommendation, and machine translation [14]. However, in many other real-world problems including robotics, disaster resilience, medical diagnosis, and bioinformatics, massive labeled data cannot be easily collected typically. For this reason, machine learning from *weak supervision* has been actively explored recently, including *semi-supervised classification* [6, 30, 40, 53, 23, 36], *one-class classification* [5, 21, 42, 51, 16, 46], *positive-unlabeled (PU) classification* [12, 33, 8, 9, 34, 24, 41], *label-proportion classification* [39, 54], *unlabeled-unlabeled classification* [7, 29, 26], *complementary-label classification* [18, 55, 19], and *similar-unlabeled classification* [1].

In this paper, we consider a novel setting of classification from weak supervision called *positive-confidence (Pconf) classification*, which is aimed at training a binary classifier only from positive data equipped with *confidence*, without negative data. Such a Pconf classification scenario is conceivable in various real-world problems. For example, in purchase prediction, we can easily collect customer data from our own company (positive data), but not from rival companies (negative data). Often times, our customers are asked to answer questionnaires/surveys on how strong their buying intention was over rival products. This may be transformed into a probability between 0 and 1 by pre-processing, and then it can be used as positive-confidence, which is all we need for Pconf classification.

Another example is a common task for app developers, where they need to predict whether app users will continue using the app or unsubscribe in the future. The critical issue is that depending on the privacy/opt-out policy or data regulation, they need to fully discard the unsubscribed user’s data. Hence, developers will not have access to users who quit using their services, but they can associate a positive-confidence score with each remaining user by, e.g., how actively they use the app.

In these applications, as long as positive-confidence data can be collected, Pconf classification allows us to obtain a classifier that discriminates between positive and negative data.




Figure 1: Illustrations of the *Pconf* classification and other related classification settings. Best viewed in color. Red points are positive data, blue points are negative data, and gray points are unlabeled data. The dark/light red colors on the rightmost figure show high/low confidence values for positive data.

Related works *Pconf* classification is related to one-class classification, which is aimed at “describing” the positive class typically from hard-labeled positive data without confidence. To the best of our knowledge, previous one-class methods are motivated geometrically [50, 42], by information theory [49], or by density estimation [4]. However, due to the descriptive nature of all previous methods, there is no systematic way to tune hyper-parameters to “classify” positive and negative data. In the conceptual example in Figure 1, one-class methods (see the second-left illustration) do not have any knowledge of the negative distribution, such that the negative distribution is in the lower right of the positive distribution (see the left-most illustration). Therefore, even if we have an infinite number of training data, one-class methods will still require regularization to have a tight boundary in all directions, wherever the positive posterior becomes low. Note that even if we knew that the negative distribution lies in the lower right of the positive distribution, it is still impossible to find the decision boundary, because we still need to know the degree of overlap between the two distributions and the class prior. One-class methods are designed for and work well for anomaly detection, but have critical limitations if the problem of interest is “classification”.

On the other hand, *Pconf* classification is aimed at constructing a discriminative classifier and thus hyper-parameters can be objectively chosen to discriminate between positive and negative data. We want to emphasize that the key contribution of our paper is to propose a method that is purely based on empirical risk minimization (ERM) [52], which makes it suitable for binary classification.

Pconf classification is also related to positive-unlabeled (PU) classification, which uses hard-labeled positive data and additional unlabeled data for constructing a binary classifier. A practical advantage of our *Pconf* classification method over typical PU classification methods is that our method does not involve estimation of the *class-prior probability*, which is required in standard PU classification methods [8, 9, 24], but is known to be highly challenging in practice [44, 2, 11, 29, 10]. This is enabled by the additional confidence information which indirectly includes the information of the class prior probability, bridging class conditionals and class posteriors.

Organization In this paper, we propose a simple ERM framework for *Pconf* classification and theoretically establish the consistency and an estimation error bound. We then provide an example of implementation to *Pconf* classification by using linear-in-parameter models (such as Gaussian kernel models), which can be implemented easily and can be computationally efficient. Finally, we experimentally demonstrate the practical usefulness of the proposed method for training linear-in-parameter models and deep neural networks.

2 Problem formulation

In this section, we formulate our *Pconf* classification problem. Suppose that a pair of d -dimensional pattern $\mathbf{x} \in \mathbb{R}^d$ and its class label $y \in \{+1, -1\}$ follow an unknown probability distribution with density $p(\mathbf{x}, y)$. Our goal is to train a binary classifier $g(\mathbf{x}) : \mathbb{R}^d \rightarrow \mathbb{R}$ so that the classification risk $R(g)$ is minimized:

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

where $\mathbb{E}_{p(\mathbf{x}, y)}$ denotes the expectation over $p(\mathbf{x}, y)$, and $\ell(z)$ is a loss function. When margin z is small, $\ell(z)$ typically takes a large value. Since $p(\mathbf{x}, y)$ is unknown, the ordinary ERM approach [52] replaces the expectation with the average over training data drawn independently from $p(\mathbf{x}, y)$.

However, in the Pconf classification scenario, we are only given positive data equipped with *confidence* $\mathcal{X} := \{(\mathbf{x}_i, r_i)\}_{i=1}^n$, where \mathbf{x}_i is a positive pattern drawn independently from $p(\mathbf{x}|y = +1)$ and r_i is the positive confidence given by $r_i = p(y = +1|\mathbf{x}_i)$. Note that this equality does not have to strictly hold as later shown in Section 4. Since we have no access to negative data in the Pconf classification scenario, we cannot directly employ the standard ERM approach. In the next section, we show how the classification risk can be estimated only from Pconf data.

3 Pconf classification

In this section, we propose an ERM framework for Pconf classification and derive an estimation error bound for the proposed method. Finally we give examples of practical implementations.

3.1 Empirical risk minimization (ERM) framework

Let $\pi_+ = p(y = +1)$ and $r(\mathbf{x}) = p(y = +1|\mathbf{x})$, and let \mathbb{E}_+ denote the expectation over $p(\mathbf{x}|y = +1)$. Then the following theorem holds, which forms the basis of our approach:

Theorem 1. *The classification risk (1) can be expressed as*

$$R(g) = \pi_+ \mathbb{E}_+ \left[\ell(g(\mathbf{x})) + \frac{1 - r(\mathbf{x})}{r(\mathbf{x})} \ell(-g(\mathbf{x})) \right], \quad (2)$$

if we have $p(y = +1|\mathbf{x}) \neq 0$ for all \mathbf{x} sampled from $p(\mathbf{x})$.

A proof is given in Appendix A.1 in the supplementary material. Equation (2) does not include the expectation over negative data, but only includes the expectation over positive data and their confidence values. Furthermore, when (2) is minimized with respect to g , unknown π_+ is a proportional constant and thus can be safely ignored. Conceptually, the assumption of $p(y = +1|\mathbf{x}) \neq 0$ is implying that the support of the negative distribution is the same or is included in the support of the positive distribution.

Based on this, we propose the following ERM framework for Pconf classification:

$$\min_g \sum_{i=1}^n \left[\ell(g(\mathbf{x}_i)) + \frac{1 - r_i}{r_i} \ell(-g(\mathbf{x}_i)) \right]. \quad (3)$$

It might be tempting to consider a similar empirical formulation as follows:

$$\min_g \sum_{i=1}^n \left[r_i \ell(g(\mathbf{x}_i)) + (1 - r_i) \ell(-g(\mathbf{x}_i)) \right]. \quad (4)$$

Equation (4) means that we weigh the positive loss with positive-confidence r_i and the negative loss with negative-confidence $1 - r_i$. This is quite natural and may look straightforward at a glance. However, if we simply consider the population version of the objective function of (4), we have

$$\begin{aligned} & \mathbb{E}_+ \left[r(\mathbf{x}) \ell(g(\mathbf{x})) + (1 - r(\mathbf{x})) \ell(-g(\mathbf{x})) \right] \\ &= \mathbb{E}_+ \left[p(y = +1|\mathbf{x}) \ell(g(\mathbf{x})) + p(y = -1|\mathbf{x}) \ell(-g(\mathbf{x})) \right] \\ &= \mathbb{E}_+ \left[\sum_{y \in \{\pm 1\}} p(y|\mathbf{x}) \ell(yg(\mathbf{x})) \right] = \mathbb{E}_+ \left[\mathbb{E}_{p(y|\mathbf{x})} [\ell(yg(\mathbf{x}))] \right], \end{aligned} \quad (5)$$

which is *not* equivalent to the classification risk $R(g)$ defined by (1). If the outer expectation was over $p(\mathbf{x})$ instead of $p(\mathbf{x}|y = +1)$ in (5), then it would be equal to (1). This implies that if we had a different problem setting of having positive confidence equipped for \mathbf{x} sampled from $p(\mathbf{x})$, this would be trivially solved by a naive weighting idea.

From this viewpoint, (3) can be regarded as an application of *importance sampling* [13, 48] to (4) to cope with the distribution difference between $p(\mathbf{x})$ and $p(\mathbf{x}|y = +1)$, but with the advantage of *not* requiring training data from the test distribution $p(\mathbf{x})$.

In summary, our ERM formulation of (3) is different from naive confidence-weighted classification of (4). We further show in Section 3.2 that the minimizer of (3) converges to the true risk minimizer, while the minimizer of (4) converges to a different quantity and hence learning based on (4) is inconsistent.

3.2 Theoretical analysis

Here we derive an estimation error bound for the proposed method. To begin with, let \mathcal{G} be our function class for ERM. Assume there exists $C_g > 0$ such that $\sup_{g \in \mathcal{G}} \|g\|_\infty \leq C_g$ as well as $C_\ell > 0$ such that $\sup_{|z| \leq C_g} \ell(z) \leq C_\ell$. The existence of C_ℓ may be guaranteed for all reasonable ℓ given a reasonable \mathcal{G} in the sense that C_g exists. As usual [31], assume $\ell(z)$ is Lipschitz continuous for all $|z| \leq C_g$ with a (not necessarily optimal) Lipschitz constant L_ℓ .

Denote by $\widehat{R}(g)$ the objective function of (3) times π_+ , which is unbiased in estimating $R(g)$ in (1) according to Theorem 1. Subsequently, let $g^* = \arg \min_{g \in \mathcal{G}} R(g)$ be the true risk minimizer, and $\hat{g} = \arg \min_{g \in \mathcal{G}} \widehat{R}(g)$ be the empirical risk minimizer, respectively. The estimation error is defined as $R(\hat{g}) - R(g^*)$, and we are going to bound it from above.

In Theorem 1, $(1 - r(\mathbf{x}))/r(\mathbf{x})$ is playing a role inside the expectation, for the fact that

$$r(\mathbf{x}) = p(y = +1 | \mathbf{x}) > 0 \text{ for } \mathbf{x} \sim p(\mathbf{x} | y = +1).$$

In order to derive any error bound based on statistical learning theory, we should ensure that $r(\mathbf{x})$ could never be too close to zero. To this end, assume there is $C_r > 0$ such that $r(\mathbf{x}) \geq C_r$ almost surely. We may trim $r(\mathbf{x})$ and then analyze the bounded but biased version of $\widehat{R}(g)$ alternatively. For simplicity, only the unbiased version is involved after assuming C_r exists.

Lemma 2. *For any $\delta > 0$, the following uniform deviation bound holds with probability at least $1 - \delta$ (over repeated sampling of data for evaluating $\widehat{R}(g)$):*

$$\sup_{g \in \mathcal{G}} |\widehat{R}(g) - R(g)| \leq 2\pi_+ \left(L_\ell + \frac{L_\ell}{C_r} \right) \mathfrak{R}_n(\mathcal{G}) + \pi_+ \left(C_\ell + \frac{C_\ell}{C_r} \right) \sqrt{\frac{\ln(2/\delta)}{2n}}, \quad (6)$$

where $\mathfrak{R}_n(\mathcal{G})$ is the Rademacher complexity of \mathcal{G} for \mathcal{X} of size n drawn from $p(\mathbf{x} | y = +1)$.¹

Lemma 2 guarantees that with high probability $\widehat{R}(g)$ concentrates around $R(g)$ for all $g \in \mathcal{G}$, and the degree of such concentration is controlled by $\mathfrak{R}_n(\mathcal{G})$. Based on this lemma, we are able to establish an estimation error bound, as follows:

Theorem 3. *For any $\delta > 0$, with probability at least $1 - \delta$ (over repeated sampling of data for training \hat{g}), we have*

$$R(\hat{g}) - R(g^*) \leq 4\pi_+ \left(L_\ell + \frac{L_\ell}{C_r} \right) \mathfrak{R}_n(\mathcal{G}) + 2\pi_+ \left(C_\ell + \frac{C_\ell}{C_r} \right) \sqrt{\frac{\ln(2/\delta)}{2n}}. \quad (7)$$

Theorem 3 guarantees learning with (3) is consistent [25]: $n \rightarrow \infty$ always means $R(\hat{g}) \rightarrow R(g^*)$. Consider linear-in-parameter models defined by

$$\mathcal{G} = \{g(\mathbf{x}) = \langle w, \phi(\mathbf{x}) \rangle_{\mathcal{H}} \mid \|w\|_{\mathcal{H}} \leq C_w, \|\phi(\mathbf{x})\|_{\mathcal{H}} \leq C_\phi\},$$

where \mathcal{H} is a Hilbert space, $\langle \cdot, \cdot \rangle_{\mathcal{H}}$ is the inner product in \mathcal{H} , $w \in \mathcal{H}$ is the normal, $\phi : \mathbb{R}^d \rightarrow \mathcal{H}$ is a feature map, and $C_w > 0$ and $C_\phi > 0$ are constants [43]. It is known that $\mathfrak{R}_n(\mathcal{G}) \leq C_w C_\phi / \sqrt{n}$ [31] and thus $R(\hat{g}) \rightarrow R(g^*)$ in $\mathcal{O}_p(1/\sqrt{n})$, where \mathcal{O}_p denotes the order in probability. This order is already the optimal parametric rate and cannot be improved without additional strong assumptions

¹ $\mathfrak{R}_n(\mathcal{G}) = \mathbb{E}_{\mathcal{X}} \mathbb{E}_{\sigma_1, \dots, \sigma_n} [\sup_{g \in \mathcal{G}} \frac{1}{n} \sum_{\mathbf{x}_i \in \mathcal{X}} \sigma_i g(\mathbf{x}_i)]$ where $\sigma_1, \dots, \sigma_n$ are n Rademacher variables following [31].

on $p(\mathbf{x}, y)$, ℓ and \mathcal{G} jointly [28]. Additionally, if ℓ is strictly convex we have $\hat{g} \rightarrow g^*$, and if the aforementioned \mathcal{G} is used $\hat{g} \rightarrow g^*$ in $\mathcal{O}_p(1/\sqrt{n})$ [3].

At first glance, learning with (4) is numerically more stable; however, it is generally inconsistent, especially when g is linear in parameters and ℓ is strictly convex. Denote by $\widehat{R}'(g)$ the objective function of (4) times π_+ , which is unbiased to $R'(g) = \pi_+ \mathbb{E}_+ \mathbb{E}_{p(y|\mathbf{x})}[\ell(yg(\mathbf{x}))]$ rather than $R(g)$. By the same technique for proving (6) and (7), it is not difficult to show that with probability at least $1 - \delta$,

$$\sup_{g \in \mathcal{G}} |\widehat{R}'(g) - R'(g)| \leq 4\pi_+ L_\ell \mathfrak{R}_n(\mathcal{G}) + 2\pi_+ C_\ell \sqrt{\frac{\ln(2/\delta)}{2n}},$$

and hence

$$R'(\hat{g}') - R'(g'^*) \leq 8\pi_+ L_\ell \mathfrak{R}_n(\mathcal{G}) + 4\pi_+ C_\ell \sqrt{\frac{\ln(2/\delta)}{2n}},$$

where

$$g'^* = \arg \min_{g \in \mathcal{G}} R'(g) \quad \text{and} \quad \hat{g}' = \arg \min_{g \in \mathcal{G}} \widehat{R}'(g).$$

As a result, when the strict convexity of $R'(g)$ and $\widehat{R}'(g)$ is also met, we have $\hat{g}' \rightarrow g'^*$. This demonstrates the inconsistency of learning with (4), since $R'(g) \neq R(g)$ which leads to $g'^* \neq g^*$ given any reasonable \mathcal{G} .

3.3 Implementation

Finally we give examples of implementations. As a classifier g , let us consider a linear-in-parameter model $g(\mathbf{x}) = \boldsymbol{\alpha}^\top \boldsymbol{\phi}(\mathbf{x})$, where $^\top$ denotes the transpose, $\boldsymbol{\phi}(\mathbf{x})$ is a vector of basis functions, and $\boldsymbol{\alpha}$ is a parameter vector. Then from (3), the ℓ_2 -regularized ERM is formulated as

$$\min_{\boldsymbol{\alpha}} \sum_{i=1}^n \left[\ell(\boldsymbol{\alpha}^\top \boldsymbol{\phi}(\mathbf{x}_i)) + \frac{1-r_i}{r_i} \ell(-\boldsymbol{\alpha}^\top \boldsymbol{\phi}(\mathbf{x}_i)) \right] + \frac{\lambda}{2} \boldsymbol{\alpha}^\top \mathbf{R} \boldsymbol{\alpha},$$

where λ is a non-negative constant and \mathbf{R} is a positive semi-definite matrix. In practice, we can use any loss functions such as squared loss $\ell_S(z) = (z-1)^2$, hinge loss $\ell_H(z) = \max(0, 1-z)$, and ramp loss $\ell_R(z) = \min(1, \max(0, 1-z))$. In the experiments in Section 4, we use the logistic loss $\ell_L(z) = \log(1 + e^{-z})$, which yields,

$$\min_{\boldsymbol{\alpha}} \sum_{i=1}^n \left[\log(1 + e^{-\boldsymbol{\alpha}^\top \boldsymbol{\phi}(\mathbf{x}_i)}) + \frac{1-r_i}{r_i} \log(1 + e^{\boldsymbol{\alpha}^\top \boldsymbol{\phi}(\mathbf{x}_i)}) \right] + \frac{\lambda}{2} \boldsymbol{\alpha}^\top \mathbf{R} \boldsymbol{\alpha}. \quad (8)$$

The above objective function is continuous and differentiable, and therefore optimization can be efficiently performed, for example, by quasi-Newton [35] or stochastic gradient methods [45].

4 Experiments

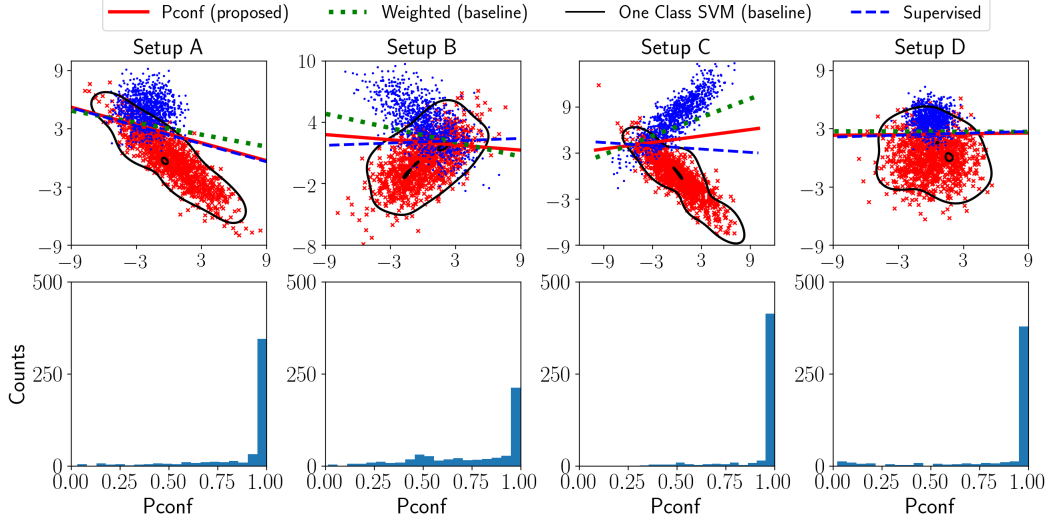
In this section, we numerically illustrate the behavior of the proposed method on synthetic datasets for linear models. We further demonstrate the usefulness of the proposed method on benchmark datasets for deep neural networks that are highly nonlinear models. The implementation is based on PyTorch [37], Sklearn [38], and mpmath [20]. Our code will be available on <http://github.com/takashiishida/pconf>.

4.1 Synthetic experiments with linear models

Setup: We used two-dimensional Gaussian distributions with means $\boldsymbol{\mu}_+$ and $\boldsymbol{\mu}_-$ and covariance matrices $\boldsymbol{\Sigma}_+$ and $\boldsymbol{\Sigma}_-$, for $p(\mathbf{x}|y = +1)$ and $p(\mathbf{x}|y = -1)$, respectively. For these parameters, we tried various combinations visually shown in Figure 2. The specific parameters used for each setup are:

- Setup A: $\boldsymbol{\mu}_+ = [0, 0]^\top$, $\boldsymbol{\mu}_- = [-2, 5]^\top$, $\boldsymbol{\Sigma}_+ = \begin{bmatrix} 7 & -6 \\ -6 & 7 \end{bmatrix}$, $\boldsymbol{\Sigma}_- = \begin{bmatrix} 2 & 0 \\ 0 & 2 \end{bmatrix}$.

Figure 2: Illustrations based on a single trail of the four setups used in experiments with various Gaussian distributions. The red and green lines are decision boundaries obtained by Pconf and Weighted classification, respectively, where only positive data with confidence are used (no negative data). The black boundary is obtained by O-SVM, which uses only hard-labeled positive data. The blue boundary is obtained by the fully-supervised method using data from both classes. Histograms of confidence of positive data are shown below.



- Setup B: $\mu_+ = [0, 0]^\top, \mu_- = [0, 4]^\top, \Sigma_+ = \begin{bmatrix} 5 & 3 \\ 3 & 5 \end{bmatrix}, \Sigma_- = \begin{bmatrix} 5 & -3 \\ -3 & 5 \end{bmatrix}$.
- Setup C: $\mu_+ = [0, 0]^\top, \mu_- = [0, 8]^\top, \Sigma_+ = \begin{bmatrix} 7 & -6 \\ -6 & 7 \end{bmatrix}, \Sigma_- = \begin{bmatrix} 7 & 6 \\ 6 & 7 \end{bmatrix}$.
- Setup D: $\mu_+ = [0, 0]^\top, \mu_- = [0, 4]^\top, \Sigma_+ = \begin{bmatrix} 4 & 0 \\ 0 & 4 \end{bmatrix}, \Sigma_- = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$.

In the case of using two Gaussian distributions, $p(y = +1|\mathbf{x}) > 0$ is satisfied for any \mathbf{x} sampled from $p(\mathbf{x})$, which is a necessary condition for applying Theorem 1. 500 positive data and 500 negative data were generated independently from each distribution for training.² Similarly, 1,000 positive and 1,000 negative data were generated for testing. We compared our proposed method (3) with the weighted classification method (4), a regression based method (predict the confidence value itself and post-process output to a binary signal by comparing it to 0.5), one-class support vector machine (O-SVM, [42]) with the Gaussian kernel, and a fully-supervised method based on the empirical version of (1). Note that the proposed method, weighted method, and regression based method only use Pconf data, O-SVM only uses (hard-labeled) positive data, and the fully-supervised method uses both positive and negative data.

In the proposed, weighted, fully-supervised methods, linear-in-input model $g(\mathbf{x}) = \alpha^\top \mathbf{x} + b$ and the logistic loss were commonly used and *vanilla gradient descent* with 5,000 epochs (full-batch size) and learning rate 0.001 was used for optimization. For the regression-based method, we used the squared loss and analytical solution [15]. For the purpose of clear comparison of the risk, we did not use regularization in this toy experiment. An exception was O-SVM, where the user is required to subjectively pre-specify regularization parameter ν and Gaussian bandwidth γ . We set them at $\nu = 0.05$ and $\gamma = 0.1$.³

Analysis with true positive-confidence: Our first experiments were conducted when true positive-confidence was known. The positive-confidence $r(\mathbf{x})$ was analytically computed from the two

²Negative training data are used only in the fully-supervised method that is tested for performance comparison.

³If we naively use default parameters in Sklearn [38] instead, which is the usual case in the real world without negative data for validation, the classification accuracy of O-SVM is worse for all setups except D in Table 1, which demonstrates the difficulty of using O-SVM.

Table 1: Comparison of the proposed Pconf classification with other methods, with varying degrees of overlap between the positive and negative distributions. We report the mean and standard deviation of the classification accuracy over 20 trials. We show the best and equivalent methods based on the 5% t-test in bold, excluding the fully-supervised method and O-SVM whose settings are different from the others.

Setup	Pconf	Weighted	Regression	O-SVM	Supervised
A	89.7 ± 0.6	88.7 ± 1.2	68.4 ± 6.5	76.0 ± 3.5	89.8 ± 0.7
B	81.2 ± 1.1	78.1 ± 1.8	73.2 ± 3.2	71.3 ± 2.3	81.4 ± 1.0
C	90.2 ± 9.1	82.7 ± 13.1	50.5 ± 1.7	90.8 ± 1.2	93.6 ± 0.5
D	91.5 ± 0.5	90.8 ± 0.7	64.6 ± 5.3	57.1 ± 4.8	91.4 ± 0.5

Table 2: Mean and standard deviation of the classification accuracy with noisy positive confidence. The experimental setup is the same as Table 1, except that positive confidence scores for positive data are noisy. Std. is the standard deviation of Gaussian noise.

Setup A			Setup C		
Std.	Pconf	Weighted	Std.	Pconf	Weighted
0.01	89.8 ± 0.6	88.8 ± 0.9	0.01	92.4 ± 1.7	84.0 ± 8.2
0.05	89.7 ± 0.6	88.3 ± 1.1	0.05	92.2 ± 3.3	78.5 ± 11.3
0.10	89.2 ± 0.7	87.6 ± 1.4	0.10	90.8 ± 9.5	72.6 ± 12.9
0.20	85.9 ± 2.5	85.8 ± 2.5	0.20	88.0 ± 9.5	65.5 ± 13.1
Setup B			Setup D		
Std.	Pconf	Weighted	Std.	Pconf	Weighted
0.01	81.2 ± 0.9	78.2 ± 1.4	0.01	91.6 ± 0.5	90.6 ± 0.9
0.05	80.7 ± 2.3	78.1 ± 1.4	0.05	91.5 ± 0.5	89.9 ± 1.2
0.10	80.8 ± 1.2	77.8 ± 1.5	0.10	90.8 ± 0.7	88.7 ± 1.8
0.20	77.8 ± 1.4	77.2 ± 1.9	0.20	87.7 ± 0.8	85.5 ± 3.7

Gaussian densities and given to each positive data. The results in Table 1 show that the proposed Pconf method is significantly better than the baselines in all cases. In most cases, the proposed Pconf method has similar accuracy compared with the fully supervised case, excluding Setup C where there is a few percent loss. Note that the naive weighted method is consistent if the model is correctly specified, but becomes inconsistent if misspecified [48].⁴

Analysis with noisy positive-confidence: In the above toy experiments, we assumed that true positive confidence $r(\mathbf{x}) = p(y = +1|\mathbf{x})$ is exactly accessible, but this can be unrealistic in practice. To investigate the influence of noise in positive-confidence, we conducted experiments with noisy positive-confidence.

As noisy positive confidence, we added zero-mean Gaussian noise with standard deviation chosen from $\{0.01, 0.05, 0.1, 0.2\}$. As the standard deviation gets larger, more noise will be incorporated into positive-confidence. When the modified positive-confidence was over 1 or below 0.01, we clipped it to 1 or rounded up to 0.01 respectively.

The results are shown in Table 2. As expected, the performance starts to deteriorate as the confidence becomes more noisy (i.e., as the standard deviation of Gaussian noise is larger), but the proposed method still works reasonably well in almost all cases.

4.2 Benchmark experiments with neural network models

Here, we use more realistic benchmark datasets and more flexible neural network models for experiments.

⁴Since our proposed method has coefficient $\frac{1-r_i}{r_i}$ in the 2nd term of (3), it may suffer numerical problems, e.g., when r_i is extremely small. To investigate this, we used the mpmath package [20] to compute the gradient with arbitrary precision. The experimental results were actually not that much different from the ones obtained with single precision, implying that the numerical problems are not much troublesome.

Fashion-MNIST: The *Fashion-MNIST dataset*⁵ consists of 70,000 examples where each sample is a 28×28 gray-scale image (input dimension is 784), associated with a label from 10 fashion item classes. We standardized the data to have zero mean and unit variance.

First, we chose “T-shirt/top” as the positive class, and another item for the negative class. The binary dataset was then divided into four sub-datasets: a training set, a validation set, a test set, and a dataset for learning a probabilistic classifier to estimate positive-confidence. Note that we ask labelers for positive-confidence values in real-world Pconf classification, but we obtained positive-confidence values through a probabilistic classifier here.

We used logistic regression with the same network architecture as a probabilistic classifier to generate confidence.⁶ However, instead of weight decay, we used *dropout* [47] with rate 50% after each fully-connected layer, and early-stopping with 20 epochs, since softmax output of flexible neural networks tends to be extremely close to 0 or 1 [14], which is not suitable as a representation of confidence. Furthermore, we rounded up positive confidence less than 1% to 1% to stabilize the optimization process.

We compared Pconf classification (3) with weighted classification (4) and fully-supervised classification based on the empirical version of (1). We used the logistic loss for these methods. We also compared our method with Auto-Encoder [17] as a one-class classification method.

Except Auto-Encoder, we used a fully-connected neural network of three hidden layers (d -100-100-100-1) with *rectified linear units (ReLU)* [32] as the activation functions, and weight decay candidates were chosen from $\{10^{-7}, 10^{-4}, 10^{-1}\}$. *Adam* [22] was again used for optimization with 200 epochs and mini-batch size 100.

To select hyper-parameters with validation data, we used the zero-one loss versions of (3) and (4) for Pconf classification and weighted classification, respectively, since no negative data were available in the validation process and thus we could not directly use the classification accuracy. On the other hand, the classification accuracy was directly used for hyper-parameter tuning of the fully-supervised method, which is extremely advantageous. We reported the test accuracy of the model with the best validation score out of all epochs.

Auto-Encoder was trained with (hard-labeled) positive data, and we classified test data into positive class if the mean squared error (MSE) is below a threshold of 70% quantile, and into negative class otherwise. Since we have no negative data for validating hyper-parameters, we sort the MSEs of training positive data in ascending order. We set the weight decay to 10^{-4} . The architecture is d -100-100-100-100 for encoding and the reversed version for decoding, with *ReLU* after hidden layers and *Tanh* after the final layer.

CIFAR-10: The *CIFAR-10 dataset*⁷ consists of 10 classes, with 5,000 images in each class. Each image is given in a $32 \times 32 \times 3$ format. We chose “airplane” as the positive class and one of the other classes as the negative class in order to construct a dataset for binary classification. We used the neural network architecture specified in Appendix B.1.

For the probabilistic classifier, the same architecture as that for Fashion-MNIST was used except *dropout* with rate 50% was added after the first two fully-connected layers. For Auto-Encoder, the MSE threshold was set to 80% quantile, and we used the architecture specified in Appendix B.2. Other details such as the loss function and weight-decay follow the same setup as the Fashion-MNIST experiments.

Results: The results in Table 3 and Table 4 show that in most cases, Pconf classification either outperforms or is comparable to the weighted classification baseline, outperforms Auto-Encoder, and is even comparable to the fully-supervised method in some cases.

⁵<https://github.com/zalando-research/fashion-mnist>

⁶Both positive and negative data are used to train the probabilistic classifier to estimate confidence, and this data is separated from any other process of experiments.

⁷<https://www.cs.toronto.edu/~kriz/cifar.html>

Table 3: Mean and standard deviation of the classification accuracy over 20 trials for the Fashion-MNIST dataset with fully-connected three hidden-layer neural networks. Pconf classification was compared with the baseline Weighted classification method, Auto-Encoder method and fully-supervised method, with *T-shirt* as the positive class and different choices for the negative class. The best and equivalent methods are shown in bold based on the 5% t-test, excluding the Auto-Encoder method and fully-supervised method.

P / N	Pconf	Weighted	Auto-Encoder	Supervised
T-shirt / trouser	92.14 ± 4.06	85.30 ± 9.07	71.06 ± 1.00	98.98 ± 0.16
T-shirt / pullover	96.00 ± 0.29	96.08 ± 1.05	70.27 ± 1.22	96.17 ± 0.34
T-shirt / dress	91.52 ± 1.14	89.31 ± 1.08	53.82 ± 0.93	96.56 ± 0.34
T-shirt / coat	98.12 ± 0.33	98.13 ± 1.12	68.74 ± 0.98	98.44 ± 0.13
T-shirt / sandal	99.55 ± 0.22	87.83 ± 18.79	82.02 ± 0.49	99.93 ± 0.09
T-shirt / shirt	83.70 ± 0.46	83.60 ± 0.65	57.76 ± 0.55	85.57 ± 0.69
T-shirt / sneaker	89.86 ± 13.32	58.26 ± 14.27	83.70 ± 0.26	100.00 ± 0.00
T-shirt / bag	97.56 ± 0.99	95.34 ± 1.00	82.79 ± 0.70	99.02 ± 0.29
T-shirt / ankle boot	98.84 ± 1.43	88.87 ± 7.86	85.07 ± 0.37	99.76 ± 0.07

Table 4: Mean and standard deviation of the classification accuracy over 20 trials for the CIFAR-10 dataset with convolutional neural networks. Pconf classification was compared with the baseline Weighted classification method, Auto-Encoder method and fully-supervised method, with *airplane* as the positive class and different choices for the negative class. The best and equivalent methods are shown in bold based on the 5% t-test, excluding the Auto-Encoder method and fully-supervised method.

P / N	Pconf	Weighted	Auto-Encoder	Supervised
airplane / automobile	82.68 ± 1.89	76.21 ± 2.43	75.13 ± 0.42	93.96 ± 0.58
airplane / bird	82.23 ± 1.21	80.66 ± 1.60	54.83 ± 0.39	87.76 ± 4.97
airplane / cat	85.18 ± 1.35	89.60 ± 0.92	61.03 ± 0.59	92.90 ± 0.58
airplane / deer	87.68 ± 1.36	87.24 ± 1.58	55.60 ± 0.53	93.35 ± 0.77
airplane / dog	89.91 ± 0.85	89.08 ± 1.95	62.64 ± 0.63	94.61 ± 0.45
airplane / frog	90.80 ± 0.98	81.84 ± 3.92	62.52 ± 0.68	95.95 ± 0.40
airplane / horse	89.82 ± 1.07	85.10 ± 2.61	67.55 ± 0.73	95.65 ± 0.37
airplane / ship	69.71 ± 2.37	70.68 ± 1.45	52.09 ± 0.42	81.45 ± 8.87
airplane / truck	81.76 ± 2.09	86.74 ± 0.85	73.74 ± 0.38	92.10 ± 0.82

5 Conclusion

We proposed a novel problem setting and algorithm for binary classification from positive data equipped with confidence. Our key contribution was to show that an unbiased estimator of the classification risk can be obtained for positive-confidence data, without negative data or even unlabeled data. This was achieved by reformulating the classification risk based on both positive and negative data, to an equivalent expression that only requires positive-confidence data. Theoretically, we established an estimation error bound, and experimentally demonstrated the usefulness of our algorithm.

Acknowledgments

TI was supported by Sumitomo Mitsui Asset Management. MS was supported by JST CREST JPMJCR1403. We thank Ikko Yamane and Tomoya Sakai for the helpful discussions. We also thank anonymous reviewers for pointing out numerical issues in our experiments, and for pointing out the necessary condition in Theorem 1 in our earlier work of this paper.

References

- [1] H. Bao, G. Niu, and M. Sugiyama. Classification from pairwise similarity and unlabeled data. In *ICML*, 2018.
- [2] G. Blanchard, G. Lee, and C. Scott. Semi-supervised novelty detection. *Journal of Machine Learning Research*, 11:2973–3009, 2010.
- [3] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- [4] M. M. Breunig, H. P. Kriegel, R. T. Ng, and J. Sander. Lof: Identifying density-based local outliers. In *ACM SIGMOD*, 2000.
- [5] V. Chandola, A. Banerjee, and V. Kumar. Anomaly detection: A survey. *ACM Computing Surveys*, 41(3), 2009.
- [6] O. Chapelle, B. Schölkopf, and A. Zien, editors. *Semi-Supervised Learning*. MIT Press, 2006.
- [7] M. C. du Plessis, G. Niu, and M. Sugiyama. Clustering unclustered data: Unsupervised binary labeling of two datasets having different class balances. In *TAAI*, 2013.
- [8] M. C. du Plessis, G. Niu, and M. Sugiyama. Analysis of learning from positive and unlabeled data. In *NIPS*, 2014.
- [9] M. C. du Plessis, G. Niu, and M. Sugiyama. Convex formulation for learning from positive and unlabeled data. In *ICML*, 2015.
- [10] M. C. du Plessis, G. Niu, and M. Sugiyama. Class-prior estimation for learning from positive and unlabeled data. *Machine Learning*, 106(4):463–492, 2017.
- [11] M. C. du Plessis and M. Sugiyama. Class prior estimation from positive and unlabeled data. *IEICE Transactions on Information and Systems*, E97-D(5):1358–1362, 2014.
- [12] C. Elkan and K. Noto. Learning classifiers from only positive and unlabeled data. In *KDD*, 2008.
- [13] G. S. Fishman. *Monte Carlo: Concepts, Algorithms, and Applications*. Springer-Verlag, 1996.
- [14] I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*. MIT Press, 2016.
- [15] T. Hastie, R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer, 2009.
- [16] S. Hido, Y. Tsuboi, H. Kashima, M. Sugiyama, and T. Kanamori. Inlier-based outlier detection via direct density ratio estimation. In *ICDM*, 2008.
- [17] G. E. Hinton and R. R. Salakhutdinov. Reducing the dimensionality of data with neural networks. *Science*, 313(5786):504–507, 2006.
- [18] T. Ishida, G. Niu, W. Hu, and M. Sugiyama. Learning from complementary labels. In *NIPS*, 2017.
- [19] T. Ishida, G. Niu, A. K. Menon, and M. Sugiyama. Complementary-label learning for arbitrary losses and models. *arXiv preprint arXiv:1810.04327*, 2018.
- [20] F. Johansson et al. *mpmath: a Python library for arbitrary-precision floating-point arithmetic (version 0.18)*, December 2013. <http://mpmath.org/>.
- [21] S. S. Khan and M. G. Madden. A survey of recent trends in one class classification. In *Irish Conference on Artificial Intelligence and Cognitive Science*, 2009.
- [22] D. P. Kingma and J. L. Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015.
- [23] T. N. Kipf and M. Welling. Semi-supervised classification with graph convolutional networks. In *ICLR*, 2017.

- [24] R. Kiryo, G. Niu, M. C. du Plessis, and M. Sugiyama. Positive-unlabeled learning with non-negative risk estimator. In *NIPS*, 2017.
- [25] M. Ledoux and M. Talagrand. *Probability in Banach Spaces: Isoperimetry and Processes*. Springer, 1991.
- [26] N. Lu, G. Niu, A. K. Menon, and M. Sugiyama. On the minimal supervision for training any binary classifier from only unlabeled data. *arXiv preprint arXiv:1808.10585v2*, 2018.
- [27] C. McDiarmid. On the method of bounded differences. In J. Siemons, editor, *Surveys in Combinatorics*, pages 148–188. Cambridge University Press, 1989.
- [28] S. Mendelson. Lower bounds for the empirical minimization algorithm. *IEEE Transactions on Information Theory*, 54(8):3797–3803, 2008.
- [29] A. Menon, B. Van Rooyen, C. S. Ong, and B. Williamson. Learning from corrupted binary labels via class-probability estimation. In *ICML*, 2015.
- [30] T. Miyato, S. Maeda, M. Koyama, K. Nakae, and S. Ishii. Distributional smoothing with virtual adversarial training. In *ICLR*, 2016.
- [31] M. Mohri, A. Rostamizadeh, and A. Talwalkar. *Foundations of Machine Learning*. MIT Press, 2012.
- [32] V. Nair and G.E. Hinton. Rectified linear units improve restricted boltzmann machines. In *ICML*, 2010.
- [33] N. Natarajan, I. S. Dhillon, P. Ravikumar, and A. Tewari. Learning with noisy labels. In *NIPS*, 2013.
- [34] G. Niu, M. C. du Plessis, T. Sakai, Y. Ma, and M. Sugiyama. Theoretical comparisons of positive-unlabeled learning against positive-negative learning. In *NIPS*, 2016.
- [35] J. Nocedal and S. Wright. *Numerical Optimization*. Springer, 2006.
- [36] A. Oliver, A. Odena, C. Raffel, E. D. Cubuk, and I. J. Goodfellow. Realistic evaluation of deep semi-supervised learning algorithms. In *NeurIPS*, 2018.
- [37] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer. Automatic differentiation in pytorch. In *Autodiff Workshop in NIPS*, 2017.
- [38] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [39] N. Quadrianto, A. Smola, T. Caetano, and Q. Le. Estimating labels from label proportions. In *ICML*, 2008.
- [40] T. Sakai, M. C. du Plessis, G. Niu, and M. Sugiyama. Semi-supervised classification based on classification from positive and unlabeled data. In *ICML*, 2017.
- [41] T. Sakai, G. Niu, and M. Sugiyama. Semi-supervised AUC optimization based on positive-unlabeled learning. *Machine Learning*, 107(4):767–794, 2018.
- [42] B. Schölkopf, J. C. Platt, J. Shawe-Taylor, A. J. Smola, and R. C. Williamson. Estimating the support of a high-dimensional distribution. *Neural Computation*, 13:1443–1471, 2001.
- [43] B. Schölkopf and A. Smola. *Learning with Kernels*. MIT Press, 2001.
- [44] C. Scott and G. Blanchard. Novelty detection: Unlabeled data definitely help. In *AISTATS*, 2009.
- [45] S. Shalev-Shwartz and S. Ben-David. *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press, 2014.
- [46] A. Smola, L. Song, and C. H. Teo. Relative novelty detection. In *AISTATS*, 2009.
- [47] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15:1929–1958, 2014.
- [48] M. Sugiyama and M. Kawanabe. *Machine Learning in Non-Stationary Environments: Introduction to Covariate Shift Adaptation*. MIT Press, 2012.

- [49] M. Sugiyama, G. Niu, M. Yamada, M. Kimura, and H. Hachiya. Information-maximization clustering based on squared-loss mutual information. *Neural Computation*, 26(1):84–131, 2014.
- [50] D. Tax and R. Duin. Support vextor domain description. In *Pattern Recognition Letters*, 1999.
- [51] D. M. J. Tax and R. P. W. Duin. Support vector data description. *Machine Learning*, 54(1):45–66, 2004.
- [52] V. N. Vapnik. *Statistical Learning Theory*. John Wiley and Sons, 1998.
- [53] Z. Yang, W. W. Cohen, and R. Salakhutdinov. Revisiting semi-supervised learning with graph embeddings. In *ICML*, 2016.
- [54] F. X. Yu, D. Liu, S. Kumar, T. Jebara, and S.-F. Chang. α svm for learning with label proportions. In *ICML*, 2013.
- [55] X. Yu, T. Liu, M. Gong, and D. Tao. Learning with biased complementary labels. In *ECCV*, 2018.

A Proofs

A.1 Proof of Theorem 1

The classification risk (1) can be expressed and decomposed as

$$\begin{aligned}
R(g) &= \sum_{y=\pm 1} \int \ell(yg(\mathbf{x}))p(\mathbf{x}|y)p(y)d\mathbf{x} \\
&= \int \ell(g(\mathbf{x}))p(\mathbf{x}|y=+1)p(y=+1)d\mathbf{x} + \int \ell(-g(\mathbf{x}))p(\mathbf{x}|y=-1)p(y=-1)d\mathbf{x} \\
&= \pi_+ \mathbb{E}_+[\ell(g(\mathbf{x}))] + \pi_- \mathbb{E}_-[\ell(-g(\mathbf{x}))], \tag{9}
\end{aligned}$$

where $\pi_- = p(y = -1)$ and \mathbb{E}_- denotes the expectation over $p(\mathbf{x}|y = -1)$. Since

$$\begin{aligned}
\pi_+ p(\mathbf{x}|y=+1) + \pi_- p(\mathbf{x}|y=-1) &= p(\mathbf{x}, y=+1) + p(\mathbf{x}, y=-1) \\
&= p(\mathbf{x}) \\
&= \frac{p(\mathbf{x}, y=+1)}{p(y=+1|\mathbf{x})} \\
&= \frac{\pi_+ p(\mathbf{x}|y=+1)}{r(\mathbf{x})},
\end{aligned}$$

where the third equality requires the assumption of $p(y = +1|\mathbf{x}) \neq 0$ stated in Theorem 1. we have

$$\pi_- p(\mathbf{x}|y=-1) = \pi_+ p(\mathbf{x}|y=+1) \left(\frac{1-r(\mathbf{x})}{r(\mathbf{x})} \right).$$

Then the second term in (9) can be expressed as

$$\begin{aligned}
\pi_- \mathbb{E}_-[\ell(-g(\mathbf{x}))] &= \int \pi_- p(\mathbf{x}|y=-1) \ell(-g(\mathbf{x})) d\mathbf{x} \\
&= \int \pi_+ p(\mathbf{x}|y=+1) \left(\frac{1-r(\mathbf{x})}{r(\mathbf{x})} \right) \ell(-g(\mathbf{x})) d\mathbf{x} \\
&= \pi_+ \mathbb{E}_+ \left[\frac{1-r(\mathbf{x})}{r(\mathbf{x})} \ell(-g(\mathbf{x})) \right],
\end{aligned}$$

which concludes the proof. \square

A.2 Proof of Lemma 2

By assumption, it holds almost surely that

$$\frac{1-r(\mathbf{x})}{r(\mathbf{x})} \leq \frac{1}{C_r};$$

due to the existence of C_ℓ , the change of $\widehat{R}(g)$ will be no more than $(C_\ell + C_\ell/C_r)/n$ if some \mathbf{x}_i is replaced with \mathbf{x}'_i .

Consider a single direction of the uniform deviation: $\sup_{g \in \mathcal{G}} \widehat{R}(g) - R(g)$. Note that the change of $\sup_{g \in \mathcal{G}} \widehat{R}(g) - R(g)$ shares the same upper bound with the change of $\widehat{R}(g)$, and *McDiarmid's inequality* [27] implies that

$$\Pr \left\{ \sup_{g \in \mathcal{G}} \widehat{R}(g) - R(g) - \mathbb{E}_{\mathcal{X}} \left[\sup_{g \in \mathcal{G}} \widehat{R}(g) - R(g) \right] \geq \epsilon \right\} \leq \exp \left(- \frac{2\epsilon^2 n}{(C_\ell + C_\ell/C_r)^2} \right),$$

or equivalently, with probability at least $1 - \delta/2$,

$$\sup_{g \in \mathcal{G}} \widehat{R}(g) - R(g) \leq \mathbb{E}_{\mathcal{X}} \left[\sup_{g \in \mathcal{G}} \widehat{R}(g) - R(g) \right] + \left(C_\ell + \frac{C_\ell}{C_r} \right) \sqrt{\frac{\ln(2/\delta)}{2n}}.$$

Since $\widehat{R}(g)$ is unbiased, it is routine to show that [31]

$$\begin{aligned}\mathbb{E}_{\mathcal{X}} \left[\sup_{g \in \mathcal{G}} \widehat{R}(g) - R(g) \right] &\leq 2\mathfrak{R}_n \left(\left(1 + \frac{1-r}{r} \right) \circ \ell \circ \mathcal{G} \right) \\ &\leq 2 \left(1 + \frac{1}{C_r} \right) \mathfrak{R}_n(\ell \circ \mathcal{G}) \\ &\leq 2 \left(L_\ell + \frac{L_\ell}{C_r} \right) \mathfrak{R}_n(\mathcal{G}),\end{aligned}$$

which proves this direction.

The other direction $\sup_{g \in \mathcal{G}} R(g) - \widehat{R}(g)$ can be proven similarly. □

A.3 Proof of Theorem 3

Based on Lemma 2, the estimation error bound (7) is proven through

$$\begin{aligned}R(\hat{g}) - R(g^*) &= \left(\widehat{R}(\hat{g}) - \widehat{R}(g^*) \right) + \left(R(\hat{g}) - \widehat{R}(\hat{g}) \right) + \left(\widehat{R}(g^*) - R(g^*) \right) \\ &\leq 0 + 2 \sup_{g \in \mathcal{G}} |\widehat{R}(g) - R(g)| \\ &\leq 4 \left(L_\ell + \frac{L_\ell}{C_r} \right) \mathfrak{R}_n(\mathcal{G}) + 2 \left(C_\ell + \frac{C_\ell}{C_r} \right) \sqrt{\frac{\ln(2/\delta)}{2n}},\end{aligned}$$

where $\widehat{R}(\hat{g}) \leq \widehat{R}(g^*)$ by the definition of \widehat{R} . □

B Neural Network Architectures used in Section 4.2

B.1 CNN architecture

- Convolution (3 in- /18 out-channels, kernel size 5).
- Max-pooling (kernel size 2, stride 2).
- Convolution (18 in- /48 out-channels, kernel size 5).
- Max-pooling (kernel size 2, stride 2).
- Fully-connected (800 units) with ReLU.
- Fully-connected (400 units) with ReLU.
- Fully-connected (1 unit).

B.2 AutoEncoder Architecture

- Convolution (3 in- /18 out-channels, kernel size 5, stride 1) with ReLU.
- Max-pooling (kernel size 2, stride 2).
- Convolutional layer (18 in- /48 out-channels, kernel size 5, stride 1) with ReLU.
- Max-pooling (kernel size 2, stride 2).
- Deconvolution (48 in- /18 out-channels, kernel size 5, stride 2) with ReLU.
- Deconvolution (18 in- /5 out-channels, kernel size 5, stride 2).
- Deconvolution (5 in- /3 out-channels, kernel size 4, stride 1) with Tanh.