
Cost Sensitive Learning in the Presence of Symmetric Label Noise

Sandhya Tripathi and N. Hemachandra

IEOR, Indian Institute of Technology Bombay, Mumbai 400076, India
{sandhya.tripathi,nh}@iitb.ac.in

Abstract. In binary classification framework, we are interested in making cost sensitive label predictions in the presence of uniform/symmetric label noise. We first observe that 0-1 Bayes classifiers are not (uniform) noise robust in cost sensitive setting. To circumvent this impossibility result, we present two schemes; unlike the existing methods, our schemes do not require noise rate. The first one uses α -weighted γ -uneven margin squared loss function, $l_{\alpha,usq}$, which can handle cost sensitivity arising due to domain requirement (using user given α) or class imbalance (by tuning γ) or both. However, we observe that $l_{\alpha,usq}$ Bayes classifiers are also not cost sensitive and noise robust. We show that regularized ERM of this loss function over the class of linear classifiers yields a cost sensitive uniform noise robust classifier as a solution of a system of linear equations. We also provide a performance bound for this classifier. The second scheme that we propose is a re-sampling based scheme that exploits the special structure of the uniform noise models and uses in-class probability η estimates. Our computational experiments on some UCI datasets with class imbalance show that classifiers of our two schemes are on par with the existing methods and in fact better in some cases w.r.t. Accuracy and Arithmetic Mean, without using/tuning noise rate. We also consider other cost sensitive performance measures viz., F measure and Weighted Cost for evaluation.

If the noise is adversarial then, one can interpret not requiring the noise rates as a dominant strategy and cross-validation over noise rates as a search for a dominant strategy. Also, the required restriction of the hypothesis class to the linear class can be interpreted as an implicit form of regularization. As our re-sampling scheme requires estimates of η , we provide a detailed comparative study of various η estimation methods on synthetic datasets, w.r.t. half a dozen evaluation criterion. Also, we provide understanding on the interpretation of cost parameters α and γ using different synthetic data experiments.

1 Introduction

We are interested in cost sensitive label predictions when only noise corrupted labels are available. The labels might be corrupted when the data has been collected by crowd scouring with not so high labeling expertise. We consider the basic case when the induced label noise is independent of the class or example, i.e., symmetric/uniform noise model. In real world, there are various applications

requiring differential misclassification costs due to class imbalance or domain requirement or both; we elaborate on these below.

In *case 1*, there is no explicit need for different penalization of classes but the data has imbalance. Example: Predicting whether the age of an applicant for a vocational training course is above or below 15 years. As the general tendency is to apply after high school, data is imbalanced but there is no need for asymmetric misclassification cost. Here, asymmetric cost should be learnt from data.

In *case 2*, there is no imbalance in data but one class’s misclassification cost is higher than that of other. Example: Product recommendation by paid web advertisements. Even though the product is liked or disliked by approximately equal proportions of population, losing a potential customer by not recommending the product is more costlier than showing paid advertisement to a customer who doesn’t like it. Here, the misclassification cost has to come from the domain.

In *case 3*, there is both imbalance and need for differential costing. Example: Rare (imbalance) disease diagnosis where the cost of missing a patient with disease is higher than cost of wrongly diagnosing a person with disease. The model should incorporate both the cost from domain and the cost due to imbalance.

In Section 1.1, we provide a summary of how these 3 cases are handled. For cost and uniform noise, we have considered real datasets belonging to cases 2 and 3.

Contributions

- Show that, unlike 0-1 loss, weighted 0-1 loss is not cost sensitive uniform noise robust.
- Show α -weighted γ -uneven margin squared loss $l_{\alpha,usq}$ with linear classifiers is both uniform noise robust and handles cost sensitivity. Present a performance bound of a classifier obtained from $l_{\alpha,usq}$ based regularized ERM.
- Propose a re-sampling based scheme for cost sensitive label prediction in the presence of uniform noise using in-class probability estimates.
- Unlike existing work, both the proposed schemes do not need true noise rate.
- Using a balanced dataset (Bupa) which requires domain cost too, we demonstrate that tuning γ on such corrupted datasets can be beneficial.
- Provide an extensive comparative study on synthetic datasets of in-class probability η estimation methods like the link function based approach ([18]), LSPC ([21]), KLIEP ([22]), k -NN based estimation. Also, a variety of metrics like MSE, MAD, Accuracy, KL divergence, Difference between maximum (minimum) value, etc., of the estimated and actual η values are computed.
- Identify the role of the two cost parameters for different scenarios of differential costing, viz., cost from domain or cost due to class imbalance or both. The interpretation is supported by synthetic data experiments w.r.t. Accuracy, AM and F measure.

Remark 1. Let us consider the case when there is an adversary who adds uniform noise to the data. Different levels of uniform noise is the strategy set of the adversary. The two schemes we propose are our strategies against the adversary. As we do not need true noise rates and hence indifferent to the adversary’s strategy, our strategies would be always a dominant strategy. However, if one

cross validates over the noise rates, like in some existing work, the cross validation can be interpreted as a search for the dominant strategy or best response. Also, we consider differential costing (class based as well as margin based) of the misclassifications.

Remark 2. Restricting the search space for the classifiers from hypothesis class of all measurable functions to the class of linear functions can be viewed as a classical form of regularization. The intuition behind this is as follows: In the space of all measurable functions, the classifier is probably over-fitting and learning the noisy data. Whereas, when the hypothesis class is restricted to the class of linear functions, the generalization of the classifier learnt on noisy data is better.

Related work For classification problems with label noise, particularly in Empirical Risk Minimization framework, the most recent work ([11, 15, 23, 16, 6]) aims to make the loss function noise robust and then develop algorithms. A major roadblock one has to get around in label noise algorithms is the non-robustness of linear classifiers from convex potentials as given in [10]. We circumvent this problem by using squared loss function which is not a convex potential and hence a candidate for label noise robustness. We indeed show that α -weighted γ -uneven margin squared loss with linear classifiers is SLN robust. Cost sensitive learning has been widely studied by [5, 12, 19] and many more. An extensive empirical study on the effect of label noise on cost sensitive learning is presented in [25]. The problem of cost sensitive uniform noise robustness is considered in [14] where asymmetric misclassification cost α is tuned and class dependent noise rates are cross validated over corrupted data. However, our work incorporates cost due to both imbalance (γ) and domain requirement (α) with the added benefit that there is no need to know the true noise rate.

Organization Section 1.1 has some details about weighted uneven margin loss functions and in-class probability estimates. In Section 2, weighted 0-1 loss is shown to be non cost sensitive uniform noise robust. Sections 3 and 4 present two different schemes that make cost sensitive predictions in the presence of uniform label noise. Section 5 has empirical evidence of the performance of proposed methods. Some discussion and future directions are presented in Section 6.

Notations Let D be the joint distribution over $\mathbf{X} \times Y$ with $\mathbf{X} \in \mathcal{X} \subseteq \mathbb{R}^n$ and $Y \in \mathcal{Y} = \{-1, 1\}$. Let the in-class probability and class marginal on D be denoted by $\eta(\mathbf{x}) := P(Y = 1|\mathbf{x})$ and $\pi := P(Y = 1)$. Let the decision function be $f : \mathbf{X} \mapsto \mathbb{R}$, hypothesis class of all measurable functions be \mathcal{H} and class of linear hypothesis be $\mathcal{H}_{lin} = \{(\mathbf{w}, b), \mathbf{w} \in \mathbb{R}^n, b \in \mathbf{R} : \|\mathbf{w}\|_2 \leq W\}$. Let \tilde{D} denote the distribution on $\mathbf{X} \times \tilde{Y}$ obtained by inducing noise to D with $\tilde{Y} \in \{-1, 1\}$. The corrupted sample is $\tilde{S} = \{(\mathbf{x}_1, \tilde{y}_1), \dots, (\mathbf{x}_m, \tilde{y}_m)\} \sim \tilde{D}^m$. The noise rate $\rho := P(\tilde{Y} = -y|Y = y, \mathbf{X} = \mathbf{x})$ is constant across classes and the model is referred to as Symmetric Label Noise (SLN) model. In such cases, the corrupted in-class probability is $\tilde{\eta}(\mathbf{x}) := P(\tilde{Y} = 1|\mathbf{x}) = (1 - 2\rho)\eta(\mathbf{x}) + \rho$ and the corrupted class marginal is $\tilde{\pi} := P(\tilde{Y} = 1) = (1 - 2\rho)\pi + \rho$. Symmetric and uniform noise are synonymous in this work.

1.1 Some relevant background

The first choice of loss function for cost sensitive learning is that of α -weighted 0-1 loss defined as follows:

$$l_{0-1,\alpha}(f(\mathbf{x}), y) = (1 - \alpha)\mathbf{1}_{\{Y=1, f(\mathbf{x}) \leq 0\}} + \alpha\mathbf{1}_{\{Y=-1, f(\mathbf{x}) > 0\}}, \quad \forall \alpha \in (0, 1) \quad (1)$$

Let the α -weighted 0-1 risk be $R_{D,\alpha} := E_D[l_{0-1,\alpha}(f(\mathbf{x}), y)]$. The minimizer of this risk is $f_\alpha^*(\mathbf{x}) = \text{sign}(\eta(\mathbf{x}) - \alpha)$ and referred to as cost-sensitive Bayes classifier. The corresponding surrogate l_α based risk and the minimizer is defined as $R_{D,l_\alpha} := E_D[l_\alpha(f(\mathbf{x}), y)]$ and $f_{l_\alpha}^* \in \mathcal{H}$ respectively. Consider the following notion of α -classification calibration.

Definition 1 (α -Classification Calibration [19]) For $\alpha \in (0, 1)$ and a loss function l , define the α -weighted loss:

$$l_\alpha(f(\mathbf{x}), y) = (1 - \alpha)l_1(f(\mathbf{x})) + \alpha l_{-1}(f(\mathbf{x})), \quad (2)$$

where $l_1(\cdot) := l(\cdot, 1)$ and $l_{-1}(\cdot) := l(\cdot, -1)$. l_α is α -classification calibrated (α -CC) iff there exists a convex, non-decreasing and invertible transformation ψ_{l_α} , with $\psi_{l_\alpha}^{-1}(0) = 0$, such that

$$\psi_{l_\alpha}(R_{D,\alpha}(f) - R_{D,\alpha}(f_\alpha^*)) \leq R_{D,l_\alpha}(f) - R_{D,l_\alpha}(f_{l_\alpha}^*). \quad (3)$$

If the classifiers obtained from α -CC losses are consistent w.r.t l_α -risk then they are also consistent w.r.t α -weighted 0-1 risk. We consider the α -weighted uneven margin squared loss [19] which is by construction α -CC and defined as follows:

$$l_{\alpha,usq}(f(\mathbf{x}), y) = (1 - \alpha)\mathbf{1}_{\{y=1\}}(1 - f(\mathbf{x}))^2 + \alpha\mathbf{1}_{\{y=-1\}}\frac{1}{\gamma}(1 + \gamma f(\mathbf{x}))^2, \quad \gamma > 0 \quad (4)$$

Interpretation of α and γ The role of α and γ can be related to the three cases of differential costing described at the start of this paper. In *case 1*, there are 3 options: fix $\alpha = 0.5$ and let tuned γ pick up the imbalance; fix $\gamma = 1$ and tune α ; tune both α and γ . Our experimental results suggest that latter two perform equally good. For *case 2*, α is given and γ can be fixed at 1. However, we observe that even in this case tuning γ can be more informative. For *case 3*, γ is tuned and α is given a priori. There would be a trade-off between α and γ , i.e., for a given α , there would be an optimal γ in a suitable sense. Above observations are based on the experiments described in Supplementary material Section D.

Choice of η estimation method As η estimates are required in re-sampling scheme, we investigated the performance of 4 methods: *Lk-fun* [18], uses classifier to get the estimate $\hat{\eta}$; interpreting η as a conditional expectation and obtaining it by a suitable squared deviation minimization; LSPC [21], density ratio estimation by l_2 norm minimization; KLIEP [22], density ratio estimation by KL divergence minimization; k -nearest neighbour based estimation method. We chose to use *Lk-fun* with logistic loss l_{log} and squared loss l_{sq} , LSPC, and

$\hat{\eta}_{norm}$, a normalized version of KLIEP because in re-sampling algorithm we are concerned with label prediction and these estimators performed equally well on Accuracy measure. Another method to estimate η , based on k -nearest neighbours, is available in [2]. This is a simple and computationally cheap method. However, it might be affected by the curse of dimensionality and mislead when not used with a carefully chosen value of k . A detailed study is available in Supplementary material Section F.

2 Cost sensitive Bayes classifiers using $l_{0-1,\alpha}$ & $l_{\alpha,usq}$ need not be uniform noise robust

The robustness notion for risk minimization in cost insensitive scenarios was introduced by [11]. They also proved that cost insensitive 0-1 loss based risk minimization is SLN robust. We extend this definition to cost-sensitive learning.

Definition 2 (Cost sensitive noise robust) *Let $f_{\alpha,A}^*$ and $\tilde{f}_{\alpha,A}^*$ be obtained from clean and corrupted distribution D and \tilde{D} using any arbitrary scheme A , then the scheme A is said to be cost sensitive noise robust if*

$$R_{D,\alpha}(\tilde{f}_{\alpha,A}^*) = R_{D,\alpha}(f_{\alpha,A}^*).$$

If $f_{\alpha,A}^$ and $\tilde{f}_{\alpha,A}^*$ are obtained from a cost sensitive loss function l and noise induced is symmetric, then l is said to be cost sensitive uniform noise robust.*

Let the $l_{0-1,\alpha}$ risk on \tilde{D} be denoted by $R_{\tilde{D},\alpha}(f)$. If one is interested in cost sensitive learning with noisy labels, then the sufficient condition of [6] becomes $(1 - \alpha)\mathbf{1}_{[f(\mathbf{x}) \leq 0]} + \alpha\mathbf{1}_{[f(\mathbf{x}) > 0]} = K$. This condition is satisfied if and only if $(1 - \alpha) = K = \alpha$ implying that it cannot be a sufficient condition for SLN robustness if there is a differential costing of $(1 - \alpha, \alpha)$, $\alpha \neq 0.5$.

Let f_{α}^* and \tilde{f}_{α}^* be the minimizers of $R_{D,\alpha}(f)$ and $R_{\tilde{D},\alpha}(f)$. Then, it is known that they have the following form:

$$f_{\alpha}^*(\mathbf{x}) = \text{sign}(\eta(\mathbf{x}) - \alpha) \quad (5) \quad \tilde{f}_{\alpha}^*(\mathbf{x}) = \text{sign}(\tilde{\eta}(\mathbf{x}) - \alpha) = \text{sign}\left(\eta(\mathbf{x}) - \frac{\alpha - \rho}{(1 - 2\rho)}\right) \quad (6)$$

The last equality in (6) follows from the fact that $\tilde{\eta} = (1 - 2\rho)\eta + \rho$. In Example 1 below, we show that $l_{0-1,\alpha}$ is not cost sensitive uniform noise robust with \mathcal{H} .

Example 1. Let Y has a Bernoulli distribution with parameter $p = 0.2$. Let $\mathbf{X} \subset \mathbb{R}$ be such that $\mathbf{X}|Y = 1 \sim \text{Uniform}(0, p)$ and $\mathbf{X}|Y = -1 \sim \text{Uniform}(1 - p, 1)$. Then, the in-class probability $\eta(\mathbf{x})$ is given as follows:

$$\eta(\mathbf{x}) = P(Y = 1|\mathbf{X} = \mathbf{x}) = p = 0.2$$

Suppose $\rho = 0.3$. Then, $\tilde{\eta}(\mathbf{x}) = (1 - 2\rho)\eta(\mathbf{x}) + \rho = 0.38$. If $\alpha = 0.25$, $f_{\alpha}^*(\mathbf{x}) = -1$ and $\tilde{f}_{\alpha}^*(\mathbf{x}) = 1, \forall \mathbf{x} \in \mathbf{X}$. Consider the α -weighted 0-1 risk of $f_{\alpha}^*(\mathbf{x})$ and $\tilde{f}_{\alpha}^*(\mathbf{x})$:

$$\begin{aligned} R_{D,\alpha}(f_{\alpha}^*) &= E_D[l_{0-1,\alpha}(f_{\alpha}^*(\mathbf{x}), y)] = (1 - \alpha)p, & \text{since } f_{\alpha}^*(\mathbf{x}) \leq 0 \quad \forall \mathbf{x} \\ R_{D,\alpha}(\tilde{f}_{\alpha}^*) &= E_D[l_{0-1,\alpha}(\tilde{f}_{\alpha}^*(\mathbf{x}), y)] = \alpha(1 - p), & \text{since } \tilde{f}_{\alpha}^*(\mathbf{x}) > 0 \quad \forall \mathbf{x} \end{aligned}$$

Therefore, $R_{D,\alpha}(f_\alpha^*) \neq R_{D,\alpha}(\tilde{f}_\alpha^*)$ implying that the α -weighted 0-1 loss function $l_{0-1,\alpha}$ is not uniform noise robust with \mathcal{H} . Details are in Supplementary material Section B.1. Note that due to $p < 0.5$, D is linearly separable; a linearly inseparable variant can be obtained by $p > 0.5$. Another linearly inseparable distribution based counter-example is available in Supplementary material Section B.4.

In view of the above example, one can try to use the principle of inductive bias, i.e., consider a strict subset of the above set of classifiers; however, Example 2 below says that the set of linear class of classifiers need not be cost sensitive uniform noise robust.

Example 2. Consider the training set $\{(3, -1), (8, -1), (12, 1)\}$ with uniform probability distribution. Let the linear classifier be of the form $fl = \text{sign}(\mathbf{x} + b)$. Let $\alpha = 0.3$ and the uniform noise be $\rho = 0.42$. Then,

$$\begin{aligned} fl_\alpha^* &= \arg \min_{fl} E_D[l_{0-1,\alpha}(y, fl)] = b^* \in (-8, -12) \quad \text{with} \quad R_{\alpha,D}(fl_\alpha^*) = 0. \\ \tilde{fl}_\alpha^* &= \arg \min_{\tilde{fl}} E_{\tilde{D}}[l_{0-1,\alpha}(\tilde{y}, \tilde{fl})] = \tilde{b}^* \in (-3, \infty) \quad \text{with} \quad R_{\alpha,D}(\tilde{fl}_\alpha^*) = 0.2. \end{aligned}$$

Details of Example 2 are available in Supplementary material Section B.2 To avoid above counter-examples, we resort to convex surrogate loss function and a type of regularization which restricts the hypothesis class. Consider an α -weighted uneven margin loss functions $l_{\alpha,un}$ [19] with its optimal classifiers on D and \tilde{D} denoted by $f_{l_{\alpha,un}}^*$ and $\tilde{f}_{l_{\alpha,un}}^*$ respectively. Regularized risk minimization defined below is known to avoid over-fitting.

$$R_{D,l_{\alpha,un}}^r(f) = E_D[l_{\alpha,un}(f(\mathbf{x}), y)] + \lambda \|f\|_2^2, \quad \text{where} \quad \lambda > 0 \quad (7)$$

Let the regularized risk of $l_{\alpha,un}$ on \tilde{D} be $R_{\tilde{D},l_{\alpha,un}}^r(f)$. Also, let the minimizers of clean and corrupted $l_{\alpha,un}$ -regularized risks be $f_{r,l_{\alpha,un}}^*$ and $\tilde{f}_{r,l_{\alpha,un}}^*$. Now, Definition 2 can be specialized to $l_{\alpha,un}$ to assure cost sensitivity, classification calibration and uniform noise robustness as follows:

Definition 3 ((α, γ, ρ)-robustness of risk minimization) For a loss function $l_{\alpha,un}$ and classifiers $\tilde{f}_{l_{\alpha,un}}^*$ and $f_{l_{\alpha,un}}^*$, risk minimization is said to be (α, γ, ρ) -robust if

$$R_{D,\alpha}(\tilde{f}_{l_{\alpha,un}}^*) = R_{D,\alpha}(f_{l_{\alpha,un}}^*). \quad (8)$$

Further, if the classifiers in equation (8) are $f_{r,l_{\alpha,un}}^*$ and $\tilde{f}_{r,l_{\alpha,un}}^*$ then, we say that regularized risk minimization under $l_{\alpha,un}$ is (α, γ, ρ) -robust.

Due to squared loss's SLN robustness property [11], we check whether $l_{\alpha,usq}$ is (α, γ, ρ) robust or not. It is not with \mathcal{H} as shown in Example 3; details of Example 3 are available in Supplementary material Section B.3.

Example 3. Consider the settings as in Example 1. Let $\alpha = 0.25$ and $\gamma = 0.4$. Then, for all \mathbf{x} ,

$$f_{l_{\alpha,usq}}^*(\mathbf{x}) = \frac{\eta(\mathbf{x}) - \alpha}{\eta(\mathbf{x})(1 - \alpha) + \gamma\alpha(1 - \eta(\mathbf{x}))} = -0.21, \quad \tilde{f}_{l_{\alpha,usq}}^*(\mathbf{x}) = \frac{\tilde{\eta}(\mathbf{x}) - \alpha}{\tilde{\eta}(\mathbf{x})(1 - \alpha) + \gamma\alpha(1 - \tilde{\eta}(\mathbf{x}))} = 0.37.$$

$$\text{And, } R_{D,\alpha}(f_{l_{\alpha,usq}}^*) = (1 - \alpha)p = 0.15, \quad R_{D,\alpha}(\tilde{f}_{l_{\alpha,usq}}^*) = \alpha(1 - p) = 0.2.$$

Hence, $R_{D,\alpha}(f_{l_{\alpha,usq}}^*) \neq R_{D,\alpha}(\tilde{f}_{l_{\alpha,usq}}^*)$. implying that $l_{\alpha,usq}$ based ERM may not be cost sensitive uniform noise robust.

We again have a negative result with $l_{\alpha,usq}$ when we consider hypothesis class \mathcal{H} . In next section, we present a positive result and show that regularized risk minimization under loss function $l_{\alpha,usq}$ is (α, γ, ρ) -robust if the hypothesis class is restricted to \mathcal{H}_{lin} .

3 $(l_{\alpha,usq}, \mathcal{H}_{lin})$ is (α, γ, ρ) robust

In this section, we consider the weighted uneven margin squared loss function $l_{\alpha,usq}$ from equation (4) with restricted hypothesis class \mathcal{H}_{lin} and show a positive result that regularized cost sensitive risk minimization under loss function $l_{\alpha,usq}$ is (α, γ, ρ) -robust. A proof is available in Supplementary material Section A.1.

Theorem 1. $(l_{\alpha,usq}, \mathcal{H}_{lin})$ is (α, γ, ρ) -robust, i.e., linear classifiers obtained from α -weighted γ -uneven margin squared loss $l_{\alpha,usq}$ based regularized risk minimization are SLN robust.

Remark 3. The above results relating to counter-examples and Theorem 1 about cost sensitive uniform noise robustness can be summarized as follows: There are two loss functions, $l_{0-1,\alpha}$ and $l_{\alpha,usq}$ and two hypothesis classes, \mathcal{H}_{lin} and \mathcal{H} . Out of the four combinations of loss functions and hypothesis classes only $l_{\alpha,usq}$ with \mathcal{H}_{lin} is cost sensitive uniform noise robust, others are not.

Next, we provide a closed form expression for the classifier learnt on corrupted data by minimizing empirical $l_{\alpha,usq}$ -regularized risk. We also provide a performance bound on the clean risk of this classifier.

3.1 $l_{\alpha,usq}$ based classifier from corrupted data & its performance

In this subsection, we present a descriptive scheme to learn a cost-sensitive linear classifier in the presence of noisy labels, by minimizing $l_{\alpha,usq}$ based regularized empirical risk, i.e.,

$$\hat{f}_r := \arg \min_{f \in \mathcal{H}_{lin}} \hat{R}_{D,l_{\alpha,usq}}^r(f), \quad (9)$$

where $\hat{R}_{D,l_{\alpha,usq}}^r(f) := \frac{1}{m} \sum_{i=1}^m l_{\alpha,usq}(f(\mathbf{x}_i), \tilde{y}_i) + \lambda \|f\|_2^2$, α is user given, γ and regularization parameter $\lambda > 0$ are to be tuned by cross validation. A proof is available in Supplementary material Section A.2.

Proposition 1. Consider corrupted regularized empirical risk $\hat{R}_{D,l_{\alpha,usq}}^r(f)$ of $l_{\alpha,usq}$. Then, the optimal (α, γ, ρ) -robust linear classifier $\hat{f}_r = (\mathbf{w}, b) \in \mathcal{H}_{lin}$ with $\mathbf{w} \in \mathbb{R}^n$ has the following form:

$$\hat{f}_r = \bar{\mathbf{w}}^* = (A + \lambda \mathbf{I})^{-1} \mathbf{c}, \quad \lambda > 0 \quad (10)$$

where $\bar{\mathbf{w}} = [w_1, w_2, \dots, w_n, b]^T$, a $n + 1$ dimensional vector of variables; A , a $(n + 1 \times n + 1)$ dimensional known symmetric matrix and \mathbf{c} , a $n + 1$ dimensional known vector are as follows:

$$A + \lambda I = \begin{bmatrix} \sum_{i=1}^m x_{i1}^2 a_i + \lambda & \sum_{i=1}^m x_{i1} x_{i2} a_i & \dots & \sum_{i=1}^m x_{i1} x_{in} a_i & \sum_{i=1}^m x_{i1} a_i \\ \sum_{i=1}^m x_{i2} x_{i1} a_i & \sum_{i=1}^m x_{i2}^2 a_i + \lambda & \dots & \sum_{i=1}^m x_{i2} x_{in} a_i & \sum_{i=1}^m x_{i2} a_i \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ \sum_{i=1}^m x_{in} x_{i1} a_i & \sum_{i=1}^m x_{in} x_{i2} a_i & \dots & \sum_{i=1}^m x_{in}^2 a_i + \lambda & \sum_{i=1}^m x_{in} a_i \\ \sum_{i=1}^m x_{i1} a_i & \sum_{i=1}^m x_{i2} a_i & \dots & \sum_{i=1}^m x_{in} a_i & \sum_{i=1}^m a_i + \lambda \end{bmatrix}, \quad \mathbf{c} = \begin{bmatrix} \sum_{i=1}^m x_{i1} c_i \\ \sum_{i=1}^m x_{i2} c_i \\ \vdots \\ \sum_{i=1}^m x_{in} c_i \\ \sum_{i=1}^m c_i \end{bmatrix}$$

with $a_i = (\mathbf{1}_{[\tilde{y}_i=1]}(1 - \alpha) + \gamma\alpha\mathbf{1}_{[\tilde{y}_i=-1]})$ and $c_i = (\mathbf{1}_{[\tilde{y}_i=1]}(1 - \alpha) - \alpha\mathbf{1}_{[\tilde{y}_i=-1]})$.

Next, we provide a result on the performance of \hat{f}_r in terms of the Rademacher complexity of the function class \mathcal{H}_{lin} . For this, we need Lemma 1 and 2 whose proofs are available in Supplementary material Section A.3 and A.4 respectively.

Lemma 1. Consider the α -weighted uneven margin squared loss $l_{\alpha, usq}(f(\mathbf{x}), y)$ which is locally L -Lipschitz with $L = 2a + 2$ where $|f(\mathbf{x})| \leq a$, for $a \geq 0$. Then, with probability at least $1 - \delta$,

$$\max_{f \in \mathcal{H}_{lin}} |\hat{R}_{\bar{D}, l_{\alpha, usq}}(f) - R_{\bar{D}, l_{\alpha, usq}}(f)| \leq 2L\mathfrak{R}(\mathcal{H}_{lin}) + \sqrt{\frac{\log(1/\delta)}{2m}},$$

where $\mathfrak{R}(\mathcal{H}_{lin}) := E_{\mathbf{X}, \sigma} \left[\sup_{f \in \mathcal{H}_{lin}} \frac{1}{m} \sum_{i=1}^m \sigma_i f(\mathbf{x}_i) \right]$ is the Rademacher complexity of the function class \mathcal{H}_{lin} with σ_i 's as independent uniform random variables taking values in $\{-1, 1\}$.

Lemma 2. For a classifier $f \in \mathcal{H}$ and user given $\alpha \in (0, 1)$, the $l_{\alpha, usq}$ risk on clean and corrupted distribution satisfy the following equation:

$$R_{\bar{D}, l_{\alpha, usq}}(f) = R_{D, l_{\alpha, usq}}(f) + 4\rho E_D[yf(\mathbf{x})[(1 - \alpha)\mathbf{1}_{[y=1]} + \alpha\mathbf{1}_{[y=-1]}]]. \quad (11)$$

Theorem 2. Under the settings of Lemma 1, with probability at least $1 - \delta$,

$$R_{D, l_{\alpha, usq}}(\hat{f}_r) \leq \min_{f \in \mathcal{H}_{lin}} R_{D, l_{\alpha, usq}}(f) + 4L\mathfrak{R}(\mathcal{H}_{lin}) + 2\sqrt{\frac{\log(1/\delta)}{2m}} + 2\lambda W^2 + \frac{4\rho}{(1 - 2\rho)} E_{\mathbf{X}}[(\tilde{f}l_{\alpha, usq}^*(\mathbf{x}) - (1 - 2\rho)\hat{f}_r(\mathbf{x}))(\eta(\mathbf{x}) - \alpha)]$$

where $\tilde{f}l_{\alpha, usq}^*$ is the linear minimizer of $R_{\bar{D}, l_{\alpha, usq}}$ and $\eta(\mathbf{x})$ is the in-class probability for \mathbf{x} . Furthermore, as $l_{\alpha, usq}$ is α -CC, there exists a non-decreasing and

invertible function $\psi_{l_{\alpha, usq}}$ with $\psi_{l_{\alpha, usq}}^{-1}(0) = 0$ such that,

$$\begin{aligned} R_{D, \alpha}(\hat{f}_r) - R_{D, \alpha}(f_{\alpha}^*) &\leq \psi_{l_{\alpha, usq}}^{-1} \left(\min_{f \in \mathcal{H}_{lin}} R_{D, l_{\alpha, usq}}(f) - \min_{f \in \mathcal{H}} R_{D, l_{\alpha, usq}}(f) + 4L\mathfrak{R}(\mathcal{H}_{lin}) \right. \\ &\quad \left. + 2\sqrt{\frac{\log(1/\delta)}{2m}} + \frac{4\rho}{(1-2\rho)} E_{\mathbf{x}}[(\tilde{f}_{l_{\alpha, usq}}^*(\mathbf{x}) - \hat{f}_r(\mathbf{x}))(\eta(\mathbf{x}) - \alpha)] \right. \\ &\quad \left. + \frac{8\rho^2}{(1-2\rho)} E_{\mathbf{x}}[\hat{f}_r(\mathbf{x})(\eta(\mathbf{x}) - \alpha)] + 2\lambda W^2 \right). \end{aligned} \quad (12)$$

A proof of Theorem 2 is available in Supplementary material Section A.5. The first two terms (involving the difference) in the right hand side of equation (12) denotes the approximation error which is small if \mathcal{H}_{lin} is large and the third term involving the Rademacher complexity denotes the estimation error which is small if \mathcal{H}_{lin} is small. The fourth term denotes the sample complexity which vanishes as the sample size increases. The bound in (12) can be used to show consistency of $l_{\alpha, usq}$ based regularized ERM if the argument of $\psi_{l_{\alpha, usq}}^{-1}$ tends to zero as sample size increases. However, in this case, it is not obvious because the last two terms involving noise rates may not vanish with increasing sample size. In spite of this, our empirical experience with this algorithm is very good.

4 A re-sampling based Algorithm $(\tilde{\eta}, \alpha)$

In this section, we present a cost sensitive label prediction algorithm based on re-balancing (which is guided by the costs) the noisy training set given to the learning algorithm. Let us consider uneven margin version of α -weighted 0-1 loss from equation (1) defined as follows:

$$l_{0-1, \alpha, \gamma}(f(x), y) = (1 - \alpha)\mathbf{1}_{\{Y=1, f(x) \leq 0\}} + \frac{\alpha}{\gamma}\mathbf{1}_{\{Y=-1, \gamma f(x) > 0\}}, \quad \forall \alpha \in (0, 1)$$

where α is user given cost and γ , tunable cost handles the class imbalance. This definition is along the lines of the uneven margin losses defined in [19]. Let $l_{0-1, \alpha, \gamma}$ -risk on D be $R_{D, \alpha, \gamma}(f)$ and corresponding optimal classifier be $f_{0-1, \alpha, \gamma}^*$:

$$f_{0-1, \alpha, \gamma}^* = \arg \min_{f \in \mathcal{H}} R_{D, \alpha, \gamma}(f) = \text{sign} \left(\eta(\mathbf{x}) - \frac{\alpha}{\gamma + (1 - \gamma)\alpha} \right). \quad (13)$$

Also, let $l_{0-1, \alpha, \gamma}$ -risk on \tilde{D} be $R_{\tilde{D}, \alpha, \gamma}(f)$ and the corresponding optimal classifier be $\tilde{f}_{0-1, \alpha, \gamma}^*$ as given below:

$$\tilde{f}_{0-1, \alpha, \gamma}^* = \text{sign} \left(\tilde{\eta}(\mathbf{x}) - \frac{\alpha}{\gamma + (1 - \gamma)\alpha} \right). \quad (14)$$

We propose Algorithm $(\tilde{\eta}, \alpha)$ which is mainly based on two ideas: (i) predictions based on a certain threshold (p^*) can correspond to predictions based on threshold (p_0) if the number of negative examples in the training set is multiplied by $r^* = \frac{p^*}{1-p^*} \frac{1-p_0}{p_0}$ (Theorem 1 of [5]) (ii) for a given \mathbf{x} , $\tilde{\eta}(\mathbf{x})$ and $\eta(\mathbf{x})$ lie on the same side of threshold 0.5 when noise rate is ρ . We first formalize the latter idea in terms of a general result. A proof is available in Supplementary material Section A.6.

Lemma 3. *In SLN models, for a given noise rate $\rho < 0.5$, the clean and corrupted class marginals π and $\tilde{\pi}$ satisfy the following condition:*

$$\pi \underset{\leq}{\geq} 0.5 \Rightarrow \tilde{\pi} \underset{\leq}{\geq} 0.5.$$

Further, the above monotonicity holds for $\eta(\mathbf{x})$ and $\tilde{\eta}(\mathbf{x})$ too.

In our case, the cost sensitive label prediction requires the desired threshold to be $\frac{\alpha}{\gamma + (1-\gamma)\alpha}$ ($= p^*$) but the threshold which we can use is 0.5 ($= p_0$) implying that for us $r^* = \frac{\alpha}{\gamma(1-\alpha)}$. If m_+ and m_- are number of positive and negative examples in m_{tr} , then we should re-sample such that the size of balanced dataset is $m_{tr,b} = m_+ + \left\lfloor \frac{\alpha m_-}{\gamma(1-\alpha)} \right\rfloor$. As we have access to only corrupted data, the learning scheme is: re-balance the corrupted data using r^* and then threshold $\tilde{\eta}$ at 0.5. Since, for SLN model, predictions made by thresholding $\tilde{\eta}$ at 0.5 are same as the predictions made by thresholding η at 0.5, for a test point \mathbf{x}_0 from D , predicted label is $sign(\tilde{\eta}(\mathbf{x}_0) - 0.5)$. The main advantage of this algorithm is that it doesn't require the knowledge of true noise rates. Also, unlike Section 3's scheme involving $l_{\alpha, usq}$ based regularized ERM, this algorithm uses $\tilde{\eta}$ estimates and hence is a generative learning scheme.

Since, we do not want to lose any minority (rare) class examples, we reassign positive labels to the minority class WLOG, if needed, implying that negative class examples are always under-sampled. The performance of Algorithm $(\tilde{\eta}, \alpha)$ is majorly dependent on sampling procedure and $\tilde{\eta}$ estimation methods used.

Algorithm $(\tilde{\eta}, \alpha)$ $\tilde{\eta}$ based scheme to make cost sensitive label predictions from uniform noise corrupted data

Input: Training data $\tilde{S}_{tr} = \{(x_1, \tilde{y}_i)\}_{i=1}^{m_{tr}}$, test data $S_{te} = \{(x_i, y_i)\}_{i=1}^{m_{te}}$, cost α , performance measure $PM \in \{Acc, AM, F, WC\}$

Output: Predicted labels and $\tilde{\eta}$ estimate on test data S_{te}

1: $\gamma_0 = \frac{\alpha}{1-\alpha} + 0.001$, since for under-sampling $r^* = \frac{\alpha}{\gamma(1-\alpha)} < 1 \Rightarrow \gamma_0 > \frac{\alpha}{1-\alpha}$

2: $\Gamma = \left\{ \frac{\alpha}{1-\alpha} + i \times 0.05 \mid i = 1, \dots, 12 \right\}$

3: **for** $\gamma \in \Gamma$ **do**

4: Under-sample the *-ve* class to get $\tilde{S}_{tr,b}$ such that $|\tilde{S}_{tr,b}| = m_+ + \left\lfloor \frac{\alpha m_-}{\gamma(1-\alpha)} \right\rfloor$.

5: Use 5-fold CV to estimate $\tilde{\eta}$ from $\tilde{S}_{tr,b}$ via *Lk-fun* or LSPC or KLIEP.

6: Compute 5-fold cross validated PM from the partitioned data.

7: **end for**

8: $\gamma^* = \arg \max_{\gamma \in \Gamma} PM$ if $PM = \{Acc, AM, F\}$ otherwise $\gamma^* = \arg \min_{\gamma \in \Gamma} WC$.

9: Under-sample the *-ve* class to get $\tilde{S}_{tr,b}$ such that $|\tilde{S}_{tr,b}| = m_+ + \left\lfloor \frac{\alpha m_-}{\gamma^*(1-\alpha)} \right\rfloor$.

10: Estimate $\tilde{\eta}$ from $\tilde{S}_{tr,b}$ using *Lk-fun* method or LSPC or KLIEP.

11: **for** $i = 1, 2, \dots, m_{te}$ **do**

12: Compute $\tilde{\eta}(\mathbf{x}_i)$

13: $\hat{y}_i = sign(\tilde{\eta}(\mathbf{x}_i) - 0.5)$

14: **end for**

Remark 4. Algorithm $(\tilde{\eta}, \alpha)$ exploits the fact that $\text{sign}(\tilde{\eta} - \frac{\alpha}{\gamma + (1-\gamma)\alpha}) = \text{sign}(\tilde{\eta}_b - 0.5)$ where $\tilde{\eta}_b$ is learnt on re-sampled data. This implies $R_{D,\alpha}(\text{sign}(\tilde{\eta}_b - 0.5)) = R_{D,\alpha}(\text{sign}(\tilde{\eta} - \frac{\alpha}{\gamma + (1-\gamma)\alpha}))$ but due to counter-examples in Section 2, these risks may not be equal to $R_{D,\alpha}(\text{sign}(\eta - \frac{\alpha}{\gamma + (1-\gamma)\alpha}))$. Hence, this scheme is not in contradiction to Section 2. However, as $\tilde{\eta}$ estimation methods use a subset of \mathcal{H} (e.g., LSPC and KLIEP use linear combinations of finite Gaussian kernels as basis functions), these risks may be equal to $R_{D,\alpha}(\text{sign}(\tilde{\eta}_e - \frac{\alpha}{\gamma + (1-\gamma)\alpha}))$ where $\tilde{\eta}_e$ is estimate of $\tilde{\eta}$ obtained from strict subset of hypothesis class. Also, based on very good empirical performance of the scheme, we believe that $R_{D,\alpha}(\text{sign}(\tilde{\eta}_e - \frac{\alpha}{\gamma + (1-\gamma)\alpha})) = R_{D,\alpha}(\text{sign}(\eta_e - \frac{\alpha}{\gamma + (1-\gamma)\alpha}))$ where η_e is an estimate of η .

5 Comparison of $l_{\alpha,usq}$ based regularized ERM and Algorithm $(\tilde{\eta}, \alpha)$ to existing methods on UCI datasets

In this section, we consider some UCI datasets [3] and demonstrate that $l_{\alpha,usq}$ is (α, γ, ρ) -robust. Also, we demonstrate the performance of Algorithm $(\tilde{\eta}, \alpha)$ with $\tilde{\eta}$ estimated using *Lk-fun*, LSPC and KLIEP. In addition to Accuracy (Acc), Arithmetic mean (AM) of True positive rate (TPR) and True negative rate (TNR), we also consider two measures suited for evaluating classifiers learnt on imbalanced data, viz., F measure and Weighted cost (WC) defined as below:

$$F = \frac{2TP}{2TP + FP + FN} \quad \text{and} \quad WC = (1 - \alpha)FN + \frac{\alpha}{\gamma}FP$$

where TP, TN, FP, FN are number of true positives, true negatives, false positives and false negatives for a classifier.

To account for randomness in the flips to simulate a given noise rate, we repeat each experiment 10 times, with independent corruptions of the data set for same noise (ρ) setting. In every trial, the data is partitioned into train and test with 80-20 split. Uniform noise induced 80% data is used for training and validation (if there are any parameters to be tuned like γ). Finally, 20% clean test data is used for evaluation. Regularization parameter, λ is tuned over the set $\Lambda = \{0.01, 0.1, 1, 10\}$. On a synthetic dataset, we observed that the performance of our methods and cost sensitive Bayes classifier on clean data w.r.t. Accuracy, AM, F and Weighted Cost measure is comparable for moderate noise rates; details in Supplementary material Section E.3. In all the tables, values within 1% of the best across a row are in bold.

Class imbalance and domain requirement of cost We report the accuracy and AM values of logistic loss based unbiased estimator (MUB) approach and approach of surrogates for weighted 0-1 loss (S-W0-1), as it is from the work of [14] and compare them to Accuracy and AM for our cost sensitive learning schemes. It is to be noted that [14] assumes that the true noise rate ρ is known and cost α is tuned. We are more flexible and user friendly as we don't need the noise rate ρ and allow for user given misclassification cost α and tune γ .

It can be observed in Table 1 that as far as Accuracy is concerned Algorithm $(\tilde{\eta}, \alpha)$ and $l_{\alpha,usq}$ based regularized ERM have comparable values to that

from MUB and S-W0-1 on all datasets. As depicted in Table 2, the proposed algorithms have marginally better values of AM measure than that of MUB and S-W0-1 method. Due to lack of a benchmark w.r.t. F and WC, on these measures, we compared our schemes to the SVMs trained on clean data and observed that our schemes fare well w.r.t to these measures too. However, due to space constraint the details are presented in Supplementary material Section E.1.

Dataset	Cost	ρ	Algorithm ($\tilde{\eta}, \alpha$): $\tilde{\eta}$ estimate from				$l_{\alpha, usq}$ RegERM	MUB	S-W0-1
			<i>Lk-fun</i> with l_{sq}	<i>Lk-fun</i> with l_{log}	LSPC	KLIEP			
(n, m_+, m_-)	α						l_{log}	l_{log}	
Breast Cancer (9,77,186)	0.2	0.2 0.4	0.71±0.06 0.45±0.15	0.72±0.04 0.46 ±0.16	0.72±0.04 0.52± 0.11	0.61±0.09 0.54±0.13	0.72±0.04 0.66±0.06	0.70 0.67	0.66 0.56
Pima Diabetes (8,268,500)	0.16	0.2 0.4	0.74±0.02 0.70±0.06	0.74±0.03 0.71±0.04	0.71±0.04 0.51±0.09	0.63±0.03 0.57±0.07	0.77±0.03 0.71± 0.04	0.76 0.65	0.73 0.66
German (20,300,700)	0.3	0.2 0.4	0.74±0.02 0.64±0.05	0.74±0.02 0.64±0.05	0.72±0.01 0.58±0.04	0.70±0.03 0.66±0.02	0.74±0.03 0.64±0.03	0.66 0.55	0.71 0.67
Thyroid (5,65,150)	0.25	0.2 0.4	0.84±0.04 0.70±0.15	0.85±0.05 0.67±0.15	0.89±0.03 0.61±0.07	0.76±0.13 0.70±0.12	0.84±0.03 0.82±0.06	0.87 0.83	0.82 0.76

Table 1. Averaged Acc ($\pm s.d.$) of the cost sensitive predictions made by Algorithm ($\tilde{\eta}, \alpha$) and $l_{\alpha, usq}$ based regularized ERM on UCI datasets corrupted by uniform noise.

Dataset	Cost	ρ	Algorithm ($\tilde{\eta}, \alpha$): $\tilde{\eta}$ estimate from				$l_{\alpha, usq}$ RegERM	MUB	S-W0-1
			<i>Lk-fun</i> with l_{sq}	<i>Lk-fun</i> with l_{log}	LSPC	KLIEP			
(n, m_+, m_-)	α						l_{log}	l_{log}	
Breast Cancer (9,77,186)	0.2	0.2 0.4	0.62±0.06 0.49±0.04	0.64±0.05 0.50±0.03	0.60±0.07 0.54±0.05	0.56±0.09 0.52±0.05	0.59±0.04 0.56±0.05	0.59 0.51	0.63 0.56
Pima Diabetes (8,268,500)	0.16	0.2 0.4	0.71±0.02 0.67±0.06	0.72±0.02 0.68±0.07	0.71±0.02 0.54±0.03	0.65±0.01 0.57±0.06	0.73±0.03 0.68±0.05	0.63 0.56	0.74 0.67
German (20,300,700)	0.3	0.2 0.4	0.67±0.02 0.57±0.06	0.66±0.01 0.56±0.06	0.63±0.03 0.57±0.03	0.59±0.04 0.53±0.04	0.65±0.04 0.60±0.03	0.67 0.51	0.69 0.56
Thyroid (5,65,150)	0.25	0.2 0.4	0.78±0.06 0.63±0.14	0.80±0.05 0.62±0.11	0.85±0.06 0.64±0.06	0.72±0.18 0.63±0.19	0.75±0.05 0.74±0.08	0.82 0.53	0.78 0.72

Table 2. Averaged AM ($\pm s.d.$) of the cost sensitive predictions made by Algorithm ($\tilde{\eta}, \alpha$) and $l_{\alpha, usq}$ based regularized ERM on UCI datasets corrupted by uniform noise.

Class balanced and domain requirement of cost We consider the Bupa dataset [3] with 6 features where the label is +1 if the value of feature 6 is greater than 3 otherwise -1 ($m_+ = 176, m_- = 169$). We learn a cost sensitive classifier by implementing Algorithm ($\tilde{\eta}, \alpha$) with $\alpha = 0.25$ and $\tilde{\eta}$ estimated using *Lk-fun* with logistic loss. When comparing tuned γ case (from the set $\Gamma = \{0.5, 0.8, 1, 1.2, 1.5\}$) and $\gamma = 1$ case, we observed in Table 3 that tuning γ is favorable for all measures. Increase in noise rates from (0.1/0.2/0.3) to 0.4 leads to change in γ values from 0.5 to 1 or 1.5 due to imbalance induced by noise. Results for $\tilde{\eta}$ estimated using *Lk-fun* with l_{sq} are available in Supplementary material E.2.

6 Discussion

We considered the binary classification problem of cost sensitive learning in the presence of uniform label noise. We are interested in the scenarios where there

Algorithm $(\tilde{\eta}, \alpha)$: $\tilde{\eta}$ estimate from <i>Lk-fun</i> with l_{log}								
ρ	γ tuned				$\gamma = 1$			
	Acc	AM	F	WC	Acc	AM	F	WC
0.0	1.0±0.0	1.0±0.0	1.0±0.0	0.0±0.0	0.99±0.00	0.99±0.005	0.99±0.005	0.05±0.1
0.1	0.96±0.02	0.96±0.03	0.96±0.04	2.36±1.01	0.88±0.88	0.88±0.07	0.89±0.05	2.22±1.29
0.2	0.91±0.03	0.90±0.03	0.91±0.02	4.09±1.33	0.68±0.11	0.68±0.11	0.77±0.06	5.95±1.88
0.3	0.88±0.03	0.88±0.03	0.82±0.1	5.47±0.32	0.55±0.05	0.55±0.05	0.69±0.02	8.02±0.96
0.4	0.61±0.11	0.58±0.09	0.67±0.00	5.64±0.29	0.51±0.01	0.51±0.01	0.67±0.01	8.3±0.475

Table 3. Dataset: BUPA Liver Disorder (6, 176, 169). The above table depicts the performance of Algorithm $(\tilde{\eta}, \alpha)$ when $\tilde{\eta}$ estimate is from *Lk-fun* with l_{log} based regularized ERM (λ tuned via CV). Here, $\alpha = 0.25$ and γ is tuned.

can be two separate costs: α , the one fixed due to domain requirement and γ that can be tuned to capture class imbalance. We first show that weighted 0-1 loss is neither uniform noise robust with measurable class of classifiers nor with linear classifiers. In spite of this, we propose two schemes to address the problem in consideration, without requiring the knowledge of true noise rates.

For the first scheme, we show that linear classifiers obtained using weighted uneven margin squared loss $l_{\alpha, usq}$ is uniform noise robust and incorporates cost sensitivity. This classifier is obtained by solving $l_{\alpha, usq}$ based regularized ERM, that only requires a matrix inversion. Also, a performance bound with respect to clean distribution for such a classifier is provided. One possible issue here could be weighted uneven margin squared loss function’s susceptibility to outliers. However, in our experiments, this scheme performed well. The second scheme is a re-sampling based algorithm using the corrupted in-class probability estimates. It handles class imbalance in the presence of uniform label noise but it can be heavily influenced by the quality of $\tilde{\eta}$ estimates. However, we chose the $\tilde{\eta}$ estimation method based on their label prediction ability (as discussed in Supplementary material Section F.6) and obtained good empirical results.

We provide empirical evidence for the performance of our schemes. We observed that the Accuracy and AM values for our schemes are comparable to those given by [14]. Our schemes have other measures like F and Weighted Cost comparable to that of SVMs trained on clean data; details are in Supplementary material Section E.1. The two proposed schemes have comparable performance measures among themselves.

An interesting direction to explore would be cost sensitive learning in Class Conditional Noise (CCN) models without using/tuning noise rates ρ_+ and ρ_- . We could invoke Elkan’s result [5] in the re-sampling scheme due to the uniform nature of noise. However, the estimation of $\tilde{\eta}$ induces an iota of approximation bias in the scheme. A theoretical proof for correctness of re-sampling scheme could provide directions in dealing with CCN models, including perhaps a bound on the error in SLN models; see Remark 4. Even though the constant term appears in the upper bound of Theorem 2, it would be interesting to understand whether it is due to the affect of noise on the systemic component of risk or something else. Also, using synthetic data, one can comment on the tightness of the bound. A more fundamental question would be “Is lack of guarantee on

consistency a price that one has to always pay while learning to get away without requiring the knowledge of noise rate?”

References

1. Shivani Agarwal. Surrogate regret bounds for the area under the roc curve via strongly proper losses. In *Conference on Learning Theory*, pages 338–353, 2013.
2. George H. Chen and Devavrat Shah. Explaining the success of nearest neighbor methods in prediction. *Foundations and Trends® in Machine Learning*, 10(5-6):337–588, 2018.
3. Dua Dheeru and Efi Karra Taniskidou. UCI Machine Learning Repository, 2017.
4. Bradley Efron and Robert Tibshirani. Improvements on cross-validation: the 632+ bootstrap method. *Journal of the American Statistical Association*, 92(438):548–560, 1997.
5. Charles Elkan. The foundations of cost-sensitive learning. In *International Joint Conference on Artificial Intelligence*, volume 17, pages 973–978, 2001.
6. Aritra Ghosh, Naresh Manwani, and PS Sastry. Making risk minimization tolerant to label noise. *Neurocomputing*, 160:93–107, 2015.
7. Takafumi Kanamori, Shohei Hido, and Masashi Sugiyama. A least-squares approach to direct importance estimation. *Journal of Machine Learning Research*, 10(Jul):1391–1445, 2009.
8. CX Ling and VS Sheng. Cost-sensitive learning and the class imbalance problem. 2011. *Encyclopedia of Machine Learning: Springer*.
9. Tongliang Liu and Dacheng Tao. Classification with noisy labels by importance reweighting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(3):447–461, 2016.
10. Philip M Long and Rocco A Servedio. Random classification noise defeats all convex potential boosters. *Machine learning*, 78(3):287–304, 2010.
11. Naresh Manwani and PS Sastry. Noise tolerance under risk minimization. *IEEE Transactions on Cybernetics*, 43(3):1146–1151, 2013.
12. Hamed Masnadi-Shirazi and Nuno Vasconcelos. Risk minimization, probability elicitation, and cost-sensitive SVMs. In *International Conference on Machine Learning*, pages 759–766, 2010.
13. Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. *Foundations of Machine Learning*. The MIT Press, 2012.
14. Nagarajan Natarajan, Inderjit S. Dhillon, Pradeep Ravikumar, and Ambuj Tewari. Cost-sensitive learning with noisy labels. *Journal of Machine Learning Research*, 18(155):1–33, 2018.
15. Nagarajan Natarajan, Inderjit S Dhillon, Pradeep K Ravikumar, and Ambuj Tewari. Learning with noisy labels. In *Advances in Neural Information Processing Systems*, pages 1196–1204, 2013.
16. Giorgio Patrini, Frank Nielsen, Richard Nock, and Marcello Carioni. Loss factorization, weakly supervised learning and label noise robustness. In *International Conference on Machine Learning*, pages 708–717, 2016.
17. J. C. Platt. Probabilities for SV machines. *Advances in Large Margin Classifiers*, pages 61–74, 2000.
18. Mark D Reid and Robert C Williamson. Composite binary losses. *Journal of Machine Learning Research*, 11(Sep):2387–2422, 2010.
19. Clayton Scott. Calibrated asymmetric surrogate losses. *Electronic Journal of Statistics*, 6:958–992, 2012.

20. Ingo Steinwart. How to compare different loss functions and their risks. *Constructive Approximation*, 26(2):225–287, 2007.
21. Masashi Sugiyama. Superfast-trainable multi-class probabilistic classifier by least-squares posterior fitting. *IEICE Transactions on Information and Systems*, 93(10):2690–2701, 2010.
22. Masashi Sugiyama, Shinichi Nakajima, Hisashi Kashima, Paul V Buenau, and Motoaki Kawanabe. Direct importance estimation with model selection and its application to covariate shift adaptation. In *Advances in Neural Information Processing Systems*, pages 1433–1440, 2008.
23. Brendan Van Rooyen, Aditya Menon, and Robert C Williamson. Learning with symmetric label noise: The importance of being unhinged. In *Advances in Neural Information Processing Systems*, pages 10–18, 2015.
24. Tong Zhang. Statistical behavior and consistency of classification methods based on convex risk minimization. *Annals of Statistics*, pages 56–85, 2004.
25. Xingquan Zhu, Xindong Wu, Taghi M Khoshgoftaar, and Yong Shi. An empirical study of the noise impact on cost-sensitive learning. In *International Joint Conference on Artificial Intelligence*, pages 1168–1174, 2007.

Supplementary Material

Cost Sensitive Learning in the Presence of Symmetric Label Noise

A Proofs

A.1 Proof of Theorem 1

Proof. Let the linear classifier be of the form $\mathbf{x}^T \mathbf{w}_\alpha + b = \tilde{\mathbf{w}}_\alpha^T \tilde{\mathbf{x}}$ where $\tilde{\mathbf{w}}_\alpha = [\mathbf{w}_\alpha \ b]^T$ and $\tilde{\mathbf{x}} = [\mathbf{x} \ 1]^T$ trained on clean data D and $\tilde{\mathbf{w}}_{\rho, \alpha}$ trained on \tilde{D} with noise rate ρ . For $\lambda > 0$, the weighted clean regularized risk is

$$\begin{aligned} R_{D, l_{\alpha, usq}}^r(\tilde{\mathbf{w}}_\alpha) &= E_D[(1 - \alpha)\mathbf{1}_{\{y=1\}}(f(\mathbf{x}) - y)^2 + \alpha\mathbf{1}_{\{y=-1\}}\frac{1}{\gamma}(\gamma f(\mathbf{x}) - y)^2] + \lambda\|f\|_2^2 \\ &= E_D[(1 - \alpha)\mathbf{1}_{\{y=1\}}(\tilde{\mathbf{x}}^T \tilde{\mathbf{w}}_\alpha - y)^2 + \alpha\mathbf{1}_{\{y=-1\}}\frac{1}{\gamma}(\gamma \tilde{\mathbf{x}}^T \tilde{\mathbf{w}}_\alpha - y)^2] + \lambda\|\tilde{\mathbf{w}}_\alpha\|_2^2 \end{aligned} \quad (15)$$

Differentiating (15) w.r.t. $\tilde{\mathbf{w}}_\alpha$ and equating to 0, we get

$$\begin{aligned} E_D[2\tilde{\mathbf{x}}(1 - \alpha)\mathbf{1}_{\{y=1\}}(\tilde{\mathbf{x}}^T \tilde{\mathbf{w}}_\alpha - y) + 2\tilde{\mathbf{x}}\alpha\mathbf{1}_{\{y=-1\}}(\gamma \tilde{\mathbf{x}}^T \tilde{\mathbf{w}}_\alpha - y)] + 2\lambda\tilde{\mathbf{w}}_\alpha &= 0 \\ E_D[(1 - \alpha)\mathbf{1}_{\{y=1\}}(\tilde{\mathbf{x}}\tilde{\mathbf{x}}^T \tilde{\mathbf{w}}_\alpha - \tilde{\mathbf{x}}y) + \alpha\mathbf{1}_{\{y=-1\}}(\gamma \tilde{\mathbf{x}}\tilde{\mathbf{x}}^T \tilde{\mathbf{w}}_\alpha - \tilde{\mathbf{x}}y)] + \lambda\tilde{\mathbf{w}}_\alpha &= 0 \end{aligned}$$

$$\begin{aligned} \tilde{\mathbf{w}}_\alpha^* &= [E_D[(1 - \alpha)\mathbf{1}_{\{y=1\}}\tilde{\mathbf{x}}\tilde{\mathbf{x}}^T + \gamma\alpha\mathbf{1}_{\{y=-1\}}\tilde{\mathbf{x}}\tilde{\mathbf{x}}^T] \\ &\quad + \lambda\mathbf{I}]^{-1} E_D[(1 - \alpha)\mathbf{1}_{\{y=1\}}\tilde{\mathbf{x}}y + \alpha\mathbf{1}_{\{y=-1\}}\tilde{\mathbf{x}}y] \end{aligned} \quad (16)$$

where \mathbf{I} is $n + 1 \times n + 1$ dimensional identity matrix. Now, consider the weighted corrupted regularized risk as:

$$\begin{aligned} R_{\tilde{D}, l_{\alpha, usq}}^r(\tilde{\mathbf{w}}_{\rho, \alpha}) &= E_{\tilde{D}}[(1 - \alpha)\mathbf{1}_{\{\tilde{y}=1\}}(\tilde{\mathbf{x}}^T \tilde{\mathbf{w}}_{\rho, \alpha} - \tilde{y})^2 + \alpha\mathbf{1}_{\{\tilde{y}=-1\}}\frac{1}{\gamma}(\gamma \tilde{\mathbf{x}}^T \tilde{\mathbf{w}}_{\rho, \alpha} - \tilde{y})^2] + \lambda\|\tilde{\mathbf{w}}_{\rho, \alpha}\|_2^2 \\ &= (1 - \rho)E_D[(1 - \alpha)\mathbf{1}_{\{y=1\}}(\tilde{\mathbf{x}}^T \tilde{\mathbf{w}}_{\rho, \alpha} - y)^2 + \alpha\mathbf{1}_{\{y=-1\}}\frac{1}{\gamma}(\gamma \tilde{\mathbf{x}}^T \tilde{\mathbf{w}}_{\rho, \alpha} - y)^2] \\ &\quad + \rho E_D[(1 - \alpha)\mathbf{1}_{\{y=1\}}(\tilde{\mathbf{x}}^T \tilde{\mathbf{w}}_{\rho, \alpha} + y)^2 + \alpha\mathbf{1}_{\{y=-1\}}\frac{1}{\gamma}(\gamma \tilde{\mathbf{x}}^T \tilde{\mathbf{w}}_{\rho, \alpha} + y)^2] + \lambda\|\tilde{\mathbf{w}}_{\rho, \alpha}\|_2^2 \end{aligned}$$

Differentiating (17) w.r.t. $\tilde{\mathbf{w}}_{\rho, \alpha}$, we get

$$\begin{aligned} (1 - \rho)E_D[2\tilde{\mathbf{x}}(1 - \alpha)\mathbf{1}_{\{y=1\}}(\tilde{\mathbf{x}}^T \tilde{\mathbf{w}}_{\rho, \alpha} - y) + 2\tilde{\mathbf{x}}\alpha\mathbf{1}_{\{y=-1\}}(\gamma \tilde{\mathbf{x}}^T \tilde{\mathbf{w}}_{\rho, \alpha} - y)] \\ + \rho E_D[2\tilde{\mathbf{x}}(1 - \alpha)\mathbf{1}_{\{y=1\}}(\tilde{\mathbf{x}}^T \tilde{\mathbf{w}}_{\rho, \alpha} + y) + 2\tilde{\mathbf{x}}\alpha\mathbf{1}_{\{y=-1\}}(\gamma \tilde{\mathbf{x}}^T \tilde{\mathbf{w}}_{\rho, \alpha} + y)] + 2\lambda\|\tilde{\mathbf{w}}_{\rho, \alpha}\| &= 0 \\ (1 - \rho)E_D[(1 - \alpha)\mathbf{1}_{\{y=1\}}\tilde{\mathbf{x}}\tilde{\mathbf{x}}^T \tilde{\mathbf{w}} + \gamma\alpha\mathbf{1}_{\{y=-1\}}\tilde{\mathbf{x}}\tilde{\mathbf{x}}^T \tilde{\mathbf{w}}] + \rho E_D[(1 - \alpha)\mathbf{1}_{\{y=1\}}\tilde{\mathbf{x}}\tilde{\mathbf{x}}^T \tilde{\mathbf{w}} + \gamma\alpha\mathbf{1}_{\{y=-1\}}\tilde{\mathbf{x}}\tilde{\mathbf{x}}^T \tilde{\mathbf{w}}] \\ + \lambda\|\tilde{\mathbf{w}}_{\rho, \alpha}\| &= (1 - \rho)E_D[(1 - \alpha)\mathbf{1}_{\{y=1\}}\tilde{\mathbf{x}}y + \alpha\mathbf{1}_{\{y=-1\}}\tilde{\mathbf{x}}y] - \rho E_D[(1 - \alpha)\mathbf{1}_{\{y=1\}}\tilde{\mathbf{x}}y + \alpha\mathbf{1}_{\{y=-1\}}\tilde{\mathbf{x}}y] \end{aligned}$$

Therefore, we get the optimal noisy weighted classifier as

$$\begin{aligned} \tilde{\mathbf{w}}_{\rho, \alpha}^* &= (1 - 2\rho)[E_D[(1 - \alpha)\mathbf{1}_{\{y=1\}}\tilde{\mathbf{x}}\tilde{\mathbf{x}}^T + \gamma\alpha\mathbf{1}_{\{y=-1\}}\tilde{\mathbf{x}}\tilde{\mathbf{x}}^T] + \lambda\mathbf{I}]^{-1} E_D[(1 - \alpha)\mathbf{1}_{\{y=1\}}\tilde{\mathbf{x}}y + \alpha\mathbf{1}_{\{y=-1\}}\tilde{\mathbf{x}}y] \\ &= (1 - 2\rho)\mathbf{w}_\alpha^* \end{aligned} \quad (18)$$

Taking the transpose in above equation and post multiplying by $\tilde{\mathbf{x}}$, we get $\tilde{\mathbf{w}}_{\rho, \alpha}^{*T} \tilde{\mathbf{x}} = (1 - 2\rho) \mathbf{w}_{\alpha}^{*T} \tilde{\mathbf{x}}$ and hence,

$$\tilde{f}l_{r, l_{\alpha, usq}}^* = (1 - 2\rho) fl_{r, l_{\alpha, usq}}^*, \quad (19)$$

where $fl_{r, l_{\alpha, usq}}^*$ and $\tilde{f}l_{r, l_{\alpha, usq}}^*$ are the optimal liner classifiers learnt using $l_{\alpha, usq}$ on D and \tilde{D} respectively. Since the noise rate $\rho < 0.5$, $sign(\tilde{f}l_{r, l_{\alpha, usq}}^*) = sign(fl_{r, l_{\alpha, usq}}^*)$. This implies that (8) holds for $(l_{\alpha, usq}, \mathcal{H}_{lin})$.

A.2 Proof of Proposition 1

Proof. Consider the empirical noisy regularized risk $\hat{R}_{\tilde{D}, l_{\alpha, usq}}^r(f)$ as follows:

$$\begin{aligned} \hat{R}_{\tilde{D}, l_{\alpha, usq}}^r(f) &= \frac{1}{m} \sum_{i=1}^m [\mathbf{1}_{[\tilde{y}_i=1]} (\mathbf{w}^T \mathbf{x}_i + b - y_i)^2 (1 - \alpha) + \frac{\alpha}{\gamma} \mathbf{1}_{[\tilde{y}_i=-1]} (\gamma \mathbf{w}^T \mathbf{x}_i + \gamma b - y_i)^2] + \lambda \|(\mathbf{w}, b)\|_2^2 \\ &= \frac{1}{m} \sum_{i=1}^m [\mathbf{1}_{[\tilde{y}_i=1]} (\mathbf{w}^T \mathbf{x}_i + b - 1)^2 (1 - \alpha) + \frac{\alpha}{\gamma} \mathbf{1}_{[\tilde{y}_i=-1]} (\gamma \mathbf{w}^T \mathbf{x}_i + \gamma b + 1)^2] + \lambda \|(\mathbf{w}, b)\|_2^2 \end{aligned} \quad (20)$$

Differentiating equation (20) w.r.t. w_j and equating to 0, we get

$$\frac{1}{m} \sum_{i=1}^m x_{ij} \left[\left(\sum_{k=1}^n w_k x_{ik} + b \right) (\mathbf{1}_{[\tilde{y}_i=1]} (1 - \alpha) + \gamma \alpha \mathbf{1}_{[\tilde{y}_i=-1]}) - (\mathbf{1}_{[\tilde{y}_i=1]} (1 - \alpha) - \alpha \mathbf{1}_{[\tilde{y}_i=-1]}) \right] + \lambda w_j = 0 \quad (21)$$

Differentiating equation (20) w.r.t. b and equating to 0, we get

$$\frac{1}{m} \sum_{i=1}^m \left[\left(\sum_{k=1}^n w_k x_{ik} + b \right) (\mathbf{1}_{[\tilde{y}_i=1]} (1 - \alpha) + \gamma \alpha \mathbf{1}_{[\tilde{y}_i=-1]}) - (\mathbf{1}_{[\tilde{y}_i=1]} (1 - \alpha) - \alpha \mathbf{1}_{[\tilde{y}_i=-1]}) \right] + \lambda b = 0 \quad (22)$$

Let $a_i = (\mathbf{1}_{[\tilde{y}_i=1]} (1 - \alpha) + \gamma \alpha \mathbf{1}_{[\tilde{y}_i=-1]})$ and $c_i = (\mathbf{1}_{[\tilde{y}_i=1]} (1 - \alpha) - \alpha \mathbf{1}_{[\tilde{y}_i=-1]})$. Writing equation (21) and (22) in matrix form, we get the following system of linear equations:

$$\underbrace{\begin{bmatrix} \sum_{i=1}^m x_{i1}^2 a_i + \lambda & \sum_{i=1}^m x_{i1} x_{i2} a_i & \dots & \sum_{i=1}^m x_{i1} x_{in} a_i & \sum_{i=1}^m x_{i1} a_i \\ \sum_{i=1}^m x_{i2} x_{i1} a_i & \sum_{i=1}^m x_{i2}^2 a_i + \lambda & \dots & \sum_{i=1}^m x_{i2} x_{in} a_i & \sum_{i=1}^m x_{i2} a_i \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ \sum_{i=1}^m x_{in} x_{i1} a_i & \sum_{i=1}^m x_{in} x_{i2} a_i & \dots & \sum_{i=1}^m x_{in}^2 a_i + \lambda & \sum_{i=1}^m x_{in} a_i \\ \sum_{i=1}^m x_{i1} a_i & \sum_{i=1}^m x_{i2} a_i & \dots & \sum_{i=1}^m x_{in} a_i & \sum_{i=1}^m a_i + \lambda \end{bmatrix}}_{A + \lambda \mathbf{I}} \underbrace{\begin{bmatrix} w_1 \\ w_2 \\ \vdots \\ w_n \\ b \end{bmatrix}}_{\bar{\mathbf{w}}} = \underbrace{\begin{bmatrix} \sum_{i=1}^m x_{i1} c_i \\ \sum_{i=1}^m x_{i2} c_i \\ \vdots \\ \sum_{i=1}^m x_{in} c_i \\ \sum_{i=1}^m c_i \end{bmatrix}}_{\mathbf{c}} \quad (23)$$

where A is a $(n + 1) \times (n + 1)$ dimensional known symmetric matrix, $\bar{\mathbf{w}}$ is $n + 1$ dimensional vector of variables and \mathbf{c} is $n + 1$ dimensional known vector. The above linear system of equations can be written as $(A + \lambda \mathbf{I}) \bar{\mathbf{w}} = \mathbf{c}$ and the cost-sensitive and (α, γ, ρ) -robust classifier \hat{f}_r is as follows:

$$\hat{f}_r = \bar{\mathbf{w}}^* = (A + \lambda \mathbf{I})^{-1} \mathbf{c} \quad (24)$$

A.3 Proof of Lemma 1

Proof. Consider the difference between corrupted true risk and corrupted empirical risk, i.e., $R_{\tilde{D},l_{\alpha,usq}}(f) - \hat{R}_{\tilde{D},l_{\alpha,usq}}(f)$. Now, using the Rademacher bound ([13]) on the maximal deviation between risks and empirical risks over $f \in \mathcal{H}_{lin}$, we get with probability at least $1 - \delta$

$$\max_{f \in \mathcal{H}_{lin}} |\hat{R}_{\tilde{D},l_{\alpha,usq}}(f) - R_{\tilde{D},l_{\alpha,usq}}(f)| \leq 2\mathfrak{R}(l_{\alpha,usq} \circ \mathcal{H}_{lin}) + \sqrt{\frac{\log(1/\delta)}{2m}}, \quad (25)$$

where $\mathfrak{R}(l_{\alpha,usq} \circ \mathcal{H}_{lin}) := E_{\tilde{D},\sigma} \left[\sup_{f \in \mathcal{H}_{lin}} \frac{1}{m} \sum_{i=1}^m \sigma_i l_{\alpha,usq}(f(\mathbf{x}_i, \tilde{y}_i)) \right]$. Now, using the Lipschitz composition property of the Rademacher averages, we have $\mathfrak{R}(l_{\alpha,usq} \circ \mathcal{H}_{lin}) \leq L\mathfrak{R}(\mathcal{H}_{lin})$. Hence, the inequality in lemma statement follows.

A.4 Proof of Lemma 2

Proof. Consider the corrupted $l_{\alpha,usq}$ risk of a classifier $f \in \mathcal{H}$ as given in equation (17).

$$\begin{aligned} R_{\tilde{D},l_{\alpha,usq}}(f) &= (1 - \rho)E_D[(1 - \alpha)\mathbf{1}_{\{y=1\}}(f(\mathbf{x}) - y)^2 + \alpha\mathbf{1}_{\{y=-1\}}\frac{1}{\gamma}(\gamma f(\mathbf{x}) - y)^2] \\ &\quad + \rho E_D[(1 - \alpha)\mathbf{1}_{\{y=1\}}(f(\mathbf{x}) + y)^2 + \alpha\mathbf{1}_{\{y=-1\}}\frac{1}{\gamma}(\gamma f(\mathbf{x}) + y)^2] \\ &= E_D[(1 - \alpha)\mathbf{1}_{\{y=1\}}(f(\mathbf{x}) - y)^2 + \alpha\mathbf{1}_{\{y=-1\}}\frac{1}{\gamma}(\gamma f(\mathbf{x}) - y)^2] \\ &\quad + \rho E_D[(1 - \alpha)\mathbf{1}_{\{y=1\}}(f(\mathbf{x}) + y)^2 + \alpha\mathbf{1}_{\{y=-1\}}\frac{1}{\gamma}(\gamma f(\mathbf{x}) + y)^2] \\ &\quad - (1 - \alpha)\mathbf{1}_{\{y=1\}}(f(\mathbf{x}) - y)^2 - \alpha\mathbf{1}_{\{y=-1\}}\frac{1}{\gamma}(\gamma f(\mathbf{x}) - y)^2] \\ &= R_{D,l_{\alpha,usq}}(f) + 4\rho E_D[yf(\mathbf{x})[(1 - \alpha)\mathbf{1}_{\{y=1\}} + \alpha\mathbf{1}_{\{y=-1\}}]] \end{aligned}$$

A.5 Proof of Theorem 2

Proof. Let $fl_{\alpha,usq}^* \in \mathcal{H}_{lin}$ be the minimizer of $R_{D,l_{\alpha,usq}}$. Then, the excess D risk of \hat{f}_r is

$$\begin{aligned} &R_{D,l_{\alpha,usq}}(\hat{f}_r) - R_{D,l_{\alpha,usq}}(fl_{\alpha,usq}^*) \\ &= R_{D,l_{\alpha,usq}}(\hat{f}_r) - \hat{R}_{\tilde{D},l_{\alpha,usq}}^r(\hat{f}_r) + \hat{R}_{\tilde{D},l_{\alpha,usq}}^r(\hat{f}_r) - \hat{R}_{\tilde{D},l_{\alpha,usq}}^r(fl_{\alpha,usq}^*) \\ &\quad + \hat{R}_{\tilde{D},l_{\alpha,usq}}^r(fl_{\alpha,usq}^*) - R_{D,l_{\alpha,usq}}(fl_{\alpha,usq}^*) \\ &\leq 0 + R_{D,l_{\alpha,usq}}(\hat{f}_r) - \hat{R}_{\tilde{D},l_{\alpha,usq}}^r(\hat{f}_r) + \hat{R}_{\tilde{D},l_{\alpha,usq}}^r(fl_{\alpha,usq}^*) - R_{D,l_{\alpha,usq}}(fl_{\alpha,usq}^*) \end{aligned}$$

The last inequality holds because \hat{f}_r is the minimizer of $\hat{R}_{\bar{D}, l_{\alpha, usq}}^r(f)$. Using Lemma 2, we get

$$\begin{aligned}
& R_{D, l_{\alpha, usq}}(\hat{f}_r) - R_{D, l_{\alpha, usq}}(f l_{\alpha, usq}^*) \\
& \leq R_{\bar{D}, l_{\alpha, usq}}(\hat{f}_r) - 4\rho E_D[\hat{f}_r(\mathbf{x})[(1-\alpha)\mathbf{1}_{[y=1]} - \alpha\mathbf{1}_{[y=-1]}]] - \hat{R}_{\bar{D}, l_{\alpha, usq}}^r(\hat{f}_r) \\
& \quad - R_{\bar{D}, l_{\alpha, usq}}(f l_{\alpha, usq}^*) + 4\rho E_D[f l_{\alpha, usq}^*(\mathbf{x})[(1-\alpha)\mathbf{1}_{[y=1]} - \alpha\mathbf{1}_{[y=-1]}]] + \hat{R}_{\bar{D}, l_{\alpha, usq}}^r(f l_{\alpha, usq}^*) \\
& \leq 2 \max_{f \in \mathcal{H}_{lin}} |\hat{R}_{\bar{D}, l_{\alpha, usq}}^r(f) - R_{\bar{D}, l_{\alpha, usq}}(f)| + 4\rho E_D[((1-\alpha)\mathbf{1}_{[y=1]} + \alpha\mathbf{1}_{[y=-1]})y(f l_{\alpha, usq}^* - \hat{f}_r)] \\
& = 2 \max_{f \in \mathcal{H}_{lin}} |\hat{R}_{\bar{D}, l_{\alpha, usq}}(f) - R_{\bar{D}, l_{\alpha, usq}}(f) + \lambda \|f\|_2^2| + 4\rho E_D[((1-\alpha)\mathbf{1}_{[y=1]} + \alpha\mathbf{1}_{[y=-1]})y(f l_{\alpha, usq}^* - \hat{f}_r)] \\
& \leq 2 \max_{f \in \mathcal{H}_{lin}} |\hat{R}_{\bar{D}, l_{\alpha, usq}}(f) - R_{\bar{D}, l_{\alpha, usq}}(f)| + 2\lambda \max_{f \in \mathcal{H}_{lin}} \|f\|_2^2 \\
& \quad + 4\rho E_D[((1-\alpha)\mathbf{1}_{[y=1]} + \alpha\mathbf{1}_{[y=-1]})y(f l_{\alpha, usq}^* - \hat{f}_r)] \\
& \leq 4L\mathfrak{R}(\mathcal{H}_{lin}) + 2\sqrt{\frac{\log(1/\delta)}{2m}} + 2\lambda W^2 + 4\rho E_D \left[((1-\alpha)\mathbf{1}_{[y=1]} + \alpha\mathbf{1}_{[y=-1]})y \left(\frac{\tilde{f} l_{\alpha, usq}^*}{1-2\rho} - \hat{f}_r \right) \right]
\end{aligned}$$

Last second inequality follows by using Triangle's inequality for absolute deviations. The last inequality follows from Lemma 1 and equation (19). Next, we simplify the term including expectation

$$\begin{aligned}
& E_D \left[((1-\alpha)\mathbf{1}_{[y=1]} + \alpha\mathbf{1}_{[y=-1]})y \left(\frac{\tilde{f} l_{\alpha, usq}^*}{1-2\rho} - \hat{f}_r \right) \right] \\
& = E_{\mathbf{X}} \left[\left(\frac{\tilde{f} l_{\alpha, usq}^*(\mathbf{x})}{1-2\rho} - \hat{f}_r(\mathbf{x}) \right) E_{Y|\mathbf{X}} [y((1-\alpha)\mathbf{1}_{[y=1]} + \alpha\mathbf{1}_{[y=-1]})] \right] \\
& = E_{\mathbf{X}} \left[\left(\frac{\tilde{f} l_{\alpha, usq}^*(\mathbf{x})}{1-2\rho} - \hat{f}_r(\mathbf{x}) \right) [(1-\alpha)\eta(\mathbf{x}) - \alpha(1-\eta(\mathbf{x}))] \right] \\
& = E_{\mathbf{X}} \left[\left(\frac{\tilde{f} l_{\alpha, usq}^*(\mathbf{x})}{1-2\rho} - \hat{f}_r(\mathbf{x}) \right) (\eta(\mathbf{x}) - \alpha) \right]
\end{aligned}$$

Substituting it back in the upper bound we get,

$$\begin{aligned}
R_{D, l_{\alpha, usq}}(\hat{f}_r) & \leq \min_{f \in \mathcal{H}_{lin}} R_{D, l_{\alpha, usq}}(f) + 4L\mathfrak{R}(\mathcal{H}_{lin}) + 2\sqrt{\frac{\log(1/\delta)}{2m}} + 2\lambda W^2 \\
& \quad + \frac{4\rho}{(1-2\rho)} E_{\mathbf{X}}[(\tilde{f} l_{\alpha, usq}^*(\mathbf{x}) - \hat{f}_r(\mathbf{x}))(\eta(\mathbf{x}) - \alpha)] \\
& \quad + \frac{8\rho^2}{(1-2\rho)} E_{\mathbf{X}}[\hat{f}_r(\mathbf{x})(\eta(\mathbf{x}) - \alpha)].
\end{aligned}$$

Subtracting $\min_{f \in \mathcal{H}} R_{D, l_{\alpha, usq}}(f)$ from both sides of above inequality and invoking equation (3) we get the second result in the theorem statement.

A.6 Proof of Lemma 3

Proof. Consider the case when $\pi < 0.5$ and $\rho < 0.5$ as follows:

$$\begin{aligned}
\pi < 0.5 & \Rightarrow (1-2\rho)\pi + \rho < (1-2\rho)0.5 + \rho \quad \text{since } \rho < 0.5 \\
\pi < 0.5 & \Rightarrow \tilde{\pi} < 0.5
\end{aligned}$$

The case of $\pi > 0.5 \Rightarrow \tilde{\pi} > 0.5$ can be proved in similar lines as above by reversing the inequality.

Since, $\tilde{\eta}(\mathbf{x}) = (1 - 2\rho)\eta(\mathbf{x}) + \rho$ has the same coefficients as those relating π and $\tilde{\pi}$, monotonicity w.r.t. 0.5 holds for $\eta(\mathbf{x})$ too.

B Details about various counter-examples

B.1 Details of the counter-example 1

Let Y has a Bernoulli distribution with parameter $p = 0.2$. Let $\mathbf{X} \subset \mathbb{R}$ be such that such that $\mathbf{X}|Y = 1 \sim Uniform(0, p)$ and $\mathbf{X}|Y = -1 \sim Uniform(1 - p, 1)$. Then, the in-class probability $\eta(\mathbf{x})$ is given as follows:

$$\begin{aligned} \eta(\mathbf{x}) = P(Y = 1|\mathbf{X} = \mathbf{x}) &= \frac{P(Y = 1)P(\mathbf{X} = \mathbf{x}|Y = 1)}{P(\mathbf{X} = \mathbf{x})} \\ &= \frac{P(Y = 1)P(\mathbf{X} = \mathbf{x}|Y = 1)}{P(Y = 1)P(\mathbf{X} = \mathbf{x}|Y = 1) + P(Y = -1)P(\mathbf{X} = \mathbf{x}|Y = -1)} \\ &= \frac{p^{\frac{1}{p}}}{p^{\frac{1}{p}} + (1 - p)^{\frac{1}{p}}} = p = 0.2 \end{aligned}$$

Suppose uniform noise rate be $\rho = 0.3$. Then, $\tilde{\eta}(\mathbf{x}) = (1 - 2\rho)\eta(\mathbf{x}) + \rho = 0.38$. Now, if we are given that positive class is more important with misclassification costs with $\alpha = 0.25$, then the corresponding $l_{0-1, \alpha}$ based optimal classifier on D and \tilde{D} are as follows:

$$f_{\alpha}^*(\mathbf{x}) = -1 \text{ and } \tilde{f}_{\alpha}^*(\mathbf{x}) = 1 \quad \forall \mathbf{x} \in \mathbf{X}.$$

Consider the α weighted 0-1 risk of the above classifiers as below :

$$\begin{aligned} R_{D, \alpha}(f_{\alpha}^*) &= E_D[l_{0-1, \alpha}(f_{\alpha}^*(\mathbf{x}), y)] \\ &= \sum_y \int_{\mathbf{x}} (\mathbf{1}_{\{f_{\alpha}^*(\mathbf{x}) > 0, y = -1\}} \alpha + \mathbf{1}_{\{f_{\alpha}^*(\mathbf{x}) \leq 0, y = 1\}} (1 - \alpha)) P(\mathbf{X} = \mathbf{x}|Y = y) P(Y = y) d\mathbf{x} \\ &= \int_{1-p}^1 (\mathbf{1}_{\{f_{\alpha}^*(\mathbf{x}) > 0\}} \alpha) P(\mathbf{X} = \mathbf{x}|Y = -1) (1 - p) d\mathbf{x} + \int_0^p (\mathbf{1}_{\{f_{\alpha}^*(\mathbf{x}) \leq 0\}} (1 - \alpha)) P(\mathbf{X} = \mathbf{x}|Y = 1) p d\mathbf{x} \\ &= (1 - \alpha)p \quad \text{since } f_{\alpha}^*(\mathbf{x}) \leq 0 \quad \forall \mathbf{x} \end{aligned}$$

$$\begin{aligned} R_{D, \alpha}(\tilde{f}_{\alpha}^*) &= E_D[l_{0-1, \alpha}(\tilde{f}_{\alpha}^*(\mathbf{x}), y)] \\ &= \sum_y \int_{\mathbf{x}} (\mathbf{1}_{\{\tilde{f}_{\alpha}^*(\mathbf{x}) > 0, y = -1\}} \alpha + \mathbf{1}_{\{\tilde{f}_{\alpha}^*(\mathbf{x}) \leq 0, y = 1\}} (1 - \alpha)) P(\mathbf{X} = \mathbf{x}|Y = y) P(Y = y) d\mathbf{x} \\ &= \int_{1-p}^1 (\mathbf{1}_{\{\tilde{f}_{\alpha}^*(\mathbf{x}) > 0\}} \alpha) P(\mathbf{X} = \mathbf{x}|Y = -1) (1 - p) d\mathbf{x} + \int_0^p (\mathbf{1}_{\{\tilde{f}_{\alpha}^*(\mathbf{x}) \leq 0\}} (1 - \alpha)) P(\mathbf{X} = \mathbf{x}|Y = 1) p d\mathbf{x} \\ &= \alpha(1 - p) \quad \text{since } \tilde{f}_{\alpha}^*(\mathbf{x}) > 0 \quad \forall \mathbf{x} \end{aligned}$$

Therefore, $R_{D, \alpha}(f_{\alpha}^*) \neq R_{D, \alpha}(\tilde{f}_{\alpha}^*)$ implying that the α weighted 0-1 loss function $l_{0-1, \alpha}$ is not uniform noise robust. Note that in this example D is linearly

separable because of $p < 0.5$. The counter-example suggests that even in the easiest case when the clean distribution is linearly separable, $l_{0-1,\alpha}$ is not uniform noise tolerant. If $p > 0.5$, D will be linearly inseparable and by changing α value another counter-example can be generated.

B.2 Details of counter-example 2

Consider the training set $\{(3, -1), (8, -1), (12, 1)\}$. Let the probability distribution on the feature space be uniformly concentrated on the training dataset. Let the linear classifier be of the form $fl = \text{sign}(\mathbf{x} + b)$. Let $\alpha = 0.3$ and the uniform noise be $\rho = 0.42$. Then,

$$\begin{aligned} fl_\alpha^* &= \arg \min_{fl} E_D[l_{0-1,\alpha}(y, fl)] \\ &= \arg \min_b \frac{1}{3} [\alpha \mathbf{1}_{\{3+b>0, y=-1\}} + \alpha \mathbf{1}_{\{8+b>0, y=-1\}} + (1-\alpha) \mathbf{1}_{\{12+b \leq 0, y=1\}}] \end{aligned}$$

Therefore, $fl_\alpha^* = b^* \in (-8, -12)$ with $R_{\alpha,D}(fl_\alpha^*) = 0$. Also,

$$\begin{aligned} \tilde{fl}_\alpha^* &= \arg \min_{\tilde{fl}} E_{\tilde{D}}[l_{0-1,\alpha}(\tilde{y}, \tilde{fl})] \\ &= \arg \min_{\tilde{fl}} (1-\rho) E_D[l_{0-1,\alpha}(y, \tilde{fl})] + \rho E_D[l_{0-1,\alpha}(-y, \tilde{fl})] \\ &= \arg \min_{\tilde{b}} \frac{1-\rho}{3} [\alpha \mathbf{1}_{\{3+\tilde{b}>0, y=-1\}} + \alpha \mathbf{1}_{\{8+\tilde{b}>0, y=-1\}} + (1-\alpha) \mathbf{1}_{\{12+\tilde{b} \leq 0, y=1\}}] \\ &\quad + \frac{\rho}{3} [(1-\alpha) \mathbf{1}_{\{3+\tilde{b} \leq 0, y=1\}} + (1-\alpha) \mathbf{1}_{\{8+\tilde{b} \leq 0, y=1\}} + \alpha \mathbf{1}_{\{12+\tilde{b}>0, y=-1\}}] \\ &= \arg \min_{\tilde{b}} \frac{\alpha}{3} [(1-\rho)(\mathbf{1}_{\{3+\tilde{b}>0, y=-1\}} + \mathbf{1}_{\{8+\tilde{b}>0, y=-1\}}) + \rho \mathbf{1}_{\{12+\tilde{b}>0, y=-1\}}] \\ &\quad + \frac{1-\alpha}{3} [\rho(\mathbf{1}_{\{3+\tilde{b} \leq 0, y=1\}} + \mathbf{1}_{\{8+\tilde{b} \leq 0, y=1\}}) + (1-\rho) \mathbf{1}_{\{12+\tilde{b} \leq 0, y=1\}}] \\ &= \arg \min_{\tilde{b}} \begin{cases} 0.474 & \text{when } \tilde{b} \in (-3, \infty) \\ 0.994 & \text{when } \tilde{b} \in (-\infty, -12) \end{cases} \end{aligned}$$

Therefore, $\tilde{fl}_\alpha^* = \tilde{b}^* \in (-3, \infty)$ with $R_{\alpha,D}(\tilde{fl}_\alpha^*) = 0.2$. Hence, $R_{\alpha,D}(fl_\alpha^*) \neq R_{\alpha,D}(\tilde{fl}_\alpha^*)$ implying that weighted 0-1 loss is not uniform noise robust even with a classifier in \mathcal{H}_{lin} .

B.3 Details of counter-example 3

Consider the minimizer of $R_{D,l_{\alpha,usq}}(f)$, the clean $l_{\alpha,usq}$ -risk of f given below:

$$\begin{aligned} fl_{\alpha,usq}^* &= \arg \min_{f \in \mathcal{H}} R_{D,l_{\alpha,usq}}(f) \\ &= \arg \min_{f \in \mathcal{H}} E_D[(1-\alpha) \mathbf{1}_{\{y=1\}} (1-f(\mathbf{x}))^2 + \alpha \mathbf{1}_{\{y=-1\}} \frac{1}{\gamma} (1+\gamma f(\mathbf{x}))^2] \\ &= \arg \min_{f \in \mathcal{H}} E_{\mathbf{X}}[E_{Y|\mathbf{X}}[(1-\alpha)\eta(\mathbf{x})(f(\mathbf{x})-y)^2 + \frac{\alpha}{\gamma}(1-\eta(\mathbf{x}))(1+\gamma f(\mathbf{x}))^2]] \\ &= \frac{\eta(\mathbf{x}) - \alpha}{\eta(\mathbf{x})(1-\alpha) + \gamma\alpha(1-\eta(\mathbf{x}))} \end{aligned} \tag{26}$$

Given \mathbf{x} , the optimal decision would be $\text{sign}(f_{l_{\alpha, usq}}^*) = \text{sign}(\eta(\mathbf{x}) - \alpha)$ since the denominator in the RHS of equation (26) is always positive. In similar lines as above, one can show that the minimizer of $R_{\tilde{D}, l_{\alpha, usq}}(f)$ is as follows:

$$\tilde{f}_{l_{\alpha, usq}}^* = \frac{\tilde{\eta}(\mathbf{x}) - \alpha}{\tilde{\eta}(\mathbf{x})(1 - \alpha) + \gamma\alpha(1 - \tilde{\eta}(\mathbf{x}))} \quad (27)$$

We show using following counter-example that clean weighted risks of $f_{l_{\alpha, usq}}^*$ and $\tilde{f}_{l_{\alpha, usq}}^*$ need not be equal.

Consider the settings as in Example 1. Let $\alpha = 0.25$ and $\gamma = 0.4$. Then, from equation (26) and (27), we get

$$f_{l_{\alpha, usq}}^* = -0.21 < 0 \quad \tilde{f}_{l_{\alpha, usq}}^* = 0.37 > 0 \quad \forall x.$$

Now, we need to check for the clean $l_{0-1, \alpha}$ risk of above two classifiers. This can be computed in the same way as in Example 1. Therefore, we get

$$R_{D, \alpha}(f_{l_{\alpha, usq}}^*) = (1 - \alpha)p = 0.15, \quad R_{D, \alpha}(\tilde{f}_{l_{\alpha, usq}}^*) = \alpha(1 - p) = 0.2$$

Hence, $R_{D, \alpha}(f_{l_{\alpha, usq}}^*) \neq R_{D, \alpha}(\tilde{f}_{l_{\alpha, usq}}^*)$ implying that the classifiers $f \in \mathcal{H}$ learnt from $l_{\alpha, usq}$ based ERM may not be both cost-sensitive and uniform noise robust.

B.4 Another counter-example to show that weighted 0-1 loss, $l_{0-1, \alpha}$ is not cost sensitive uniform noise robust

Following example shows that cost sensitive risk minimization under $l_{0-1, \alpha}$ is not uniform noise tolerant.

Example 4. Let Y has a Bernoulli distribution with parameter $p = 0.8$. Let $\mathbf{X} \subset \mathbb{R}$ be such that $\mathbf{X}|Y = 1 \sim N(2, 1)$ and $\mathbf{X}|Y = -1 \sim N(-2.5, 1)$. Then, using Lemma 4, the in-class probability $\eta(\mathbf{x})$ is given as follows:

$$\eta(\mathbf{x}) = \frac{1}{1 + 0.25 \exp(-1.125 - 4.5\mathbf{x})}. \quad (28)$$

Suppose we are given that the negative class is more important with misclassification cost $\alpha = 0.65$. Also, let the uniform noise rate to corrupt the data be $\rho = 0.3$. Now, using equation (5) and (6), the thresholds corresponding to $\eta(\mathbf{x})$ are 0.65 and 0.875. Then, the optimal classifiers are as follows:

$$\begin{aligned} f_{\alpha}^*(\mathbf{x}) &= \begin{cases} +1 & \text{if } \eta(\mathbf{x}) > 0.65 \\ -1 & \text{if } \eta(\mathbf{x}) \leq 0.65 \end{cases} & \tilde{f}_{\alpha}^*(\mathbf{x}) &= \begin{cases} +1 & \text{if } \eta(\mathbf{x}) > 0.875 \\ -1 & \text{if } \eta(\mathbf{x}) \leq 0.875 \end{cases} \\ &= \begin{cases} +1 & \text{if } \mathbf{x} > 0.195629 \\ -1 & \text{if } \mathbf{x} \leq 0.195629 \end{cases} & & = \begin{cases} +1 & \text{if } \mathbf{x} > 0.490489 \\ -1 & \text{if } \mathbf{x} \leq 0.490489 \end{cases} \end{aligned} \quad (29) \quad (30)$$

The last equality is obtained by substituting the value of $\eta(\mathbf{x})$ in (28) and solving for \mathbf{x} . Now we compute the weighted 0-1 risks of these classifiers show

that they are different.

$$\begin{aligned}
R_{D,\alpha}(f_\alpha^*) &= E_D[l_{0-1,\alpha}(f_\alpha^*(\mathbf{x}), y)] \\
&= \sum_y \int_{\mathbf{x}} (\mathbf{1}_{\{f_\alpha^*(\mathbf{x}) > 0, y = -1\}} \alpha + \mathbf{1}_{\{f_\alpha^*(\mathbf{x}) \leq 0, y = 1\}} (1 - \alpha)) P(\mathbf{X} = \mathbf{x} | Y = y) P(Y = y) d\mathbf{x} \\
&= \int_{-\infty}^{\infty} (\mathbf{1}_{\{f_\alpha^*(\mathbf{x}) > 0\}} \alpha) P(\mathbf{X} = \mathbf{x} | Y = -1) P(Y = -1) d\mathbf{x} \\
&\quad + \int_{-\infty}^{\infty} (\mathbf{1}_{\{f_\alpha^*(\mathbf{x}) \leq 0\}} (1 - \alpha)) P(\mathbf{X} = \mathbf{x} | Y = 1) P(Y = 1) d\mathbf{x} \\
&= (1 - p)\alpha \int_{0.1956}^{\infty} \frac{1}{\sqrt{2\pi}} \exp(-\frac{1}{2}(\mathbf{x} + 2.5)^2) d\mathbf{x} + p(1 - \alpha) \int_{-\infty}^{0.1956} \frac{1}{\sqrt{2\pi}} \exp(-\frac{1}{2}(\mathbf{x} - 2)^2) d\mathbf{x} \\
&= 0.13(1 - \Phi(2.69)) + 0.28\Phi(-1.81) \quad \text{where } \Phi(\mathbf{x}) \text{ is the c.d.f of standard normal dist at } \mathbf{x}. \\
&= 0.0103
\end{aligned}$$

$$\begin{aligned}
R_{D,\alpha}(\tilde{f}_\alpha^*) &= E_D[l_{0-1,\alpha}(\tilde{f}_\alpha^*(\mathbf{x}), y)] \\
&= \sum_y \int_{\mathbf{x}} (\mathbf{1}_{\{\tilde{f}_\alpha^*(\mathbf{x}) > 0, y = -1\}} \alpha + \mathbf{1}_{\{\tilde{f}_\alpha^*(\mathbf{x}) \leq 0, y = 1\}} (1 - \alpha)) P(\mathbf{X} = \mathbf{x} | Y = y) P(Y = y) d\mathbf{x} \\
&= \int_{-\infty}^{\infty} (\mathbf{1}_{\{\tilde{f}_\alpha^*(\mathbf{x}) > 0\}} \alpha) P(\mathbf{X} = \mathbf{x} | Y = -1) P(Y = -1) d\mathbf{x} \\
&\quad + \int_{-\infty}^{\infty} (\mathbf{1}_{\{\tilde{f}_\alpha^*(\mathbf{x}) \leq 0\}} (1 - \alpha)) P(\mathbf{X} = \mathbf{x} | Y = 1) P(Y = 1) d\mathbf{x} \\
&= (1 - p)\alpha \int_{0.4908}^{\infty} \frac{1}{\sqrt{2\pi}} \exp(-\frac{1}{2}(\mathbf{x} + 2.5)^2) d\mathbf{x} + p(1 - \alpha) \int_{-\infty}^{0.4908} \frac{1}{\sqrt{2\pi}} \exp(-\frac{1}{2}(\mathbf{x} - 2)^2) d\mathbf{x} \\
&= 0.13(1 - \Phi(2.99)) + 0.28\Phi(-1.51) \quad \text{where } \Phi(\mathbf{x}) \text{ is the c.d.f of standard normal dist at } \mathbf{x}. \\
&= 0.0185
\end{aligned}$$

Therefore, $R_{D,\alpha}(f_\alpha^*) \neq R_{D,\alpha}(\tilde{f}_\alpha^*)$ implying that the risk minimization under weighted 0-1 loss $l_{0-1,\alpha}$ is not uniform robust.

C Explicit form for η when conditional distributions are normally distributed

Here, we present a result which provides the explicit form of in-class probability η when the conditional distributions $\mathbf{X}|Y$ are normally distributed.

Lemma 4. *Let Y has Bernoulli distribution with parameter p . Let $\mathbf{X} \subset \mathbb{R}^n$ be such that $\mathbf{X}|Y = 1 \sim N(\boldsymbol{\mu}_+, \Sigma_+)$ and $\mathbf{X}|Y = -1 \sim N(\boldsymbol{\mu}_-, \Sigma_-)$. Then, the*

in-class probability $\eta(\mathbf{x}) = P(Y = 1|\mathbf{X} = \mathbf{x})$ is given as follows:

$$\eta(\mathbf{x}) = \frac{1}{1 + \frac{1-p}{p} \sqrt{\frac{|\Sigma_+|}{|\Sigma_-|}} \exp(-\frac{1}{2}(\mathbf{x}^T(\Sigma_- - \Sigma_+)\mathbf{x} - 2\mathbf{x}^T(\Sigma_- \boldsymbol{\mu}_- - \Sigma_+ \boldsymbol{\mu}_+) + \boldsymbol{\mu}_-^T \Sigma_- \boldsymbol{\mu}_- - \boldsymbol{\mu}_+^T \Sigma_+ \boldsymbol{\mu}_+))} \quad (31)$$

where Σ_+^- and Σ_-^- are the inverse of Σ_+ and Σ_- respectively.

Proof. Consider the in-class probability $\eta(\mathbf{x})$ given below:

$$\begin{aligned} P(Y = 1|\mathbf{X} = \mathbf{x}) &= \frac{P(Y = 1, \mathbf{X} = \mathbf{x})}{P(\mathbf{X} = \mathbf{x})} \\ &= \frac{P(\mathbf{X} = \mathbf{x}|Y = 1)P(Y = 1)}{\sum_{y \in \{-1, 1\}} P(\mathbf{X} = \mathbf{x}|Y = y)P(Y = y)} \\ &= \frac{\frac{1}{\sqrt{2\pi|\Sigma_+|}} \exp(-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu}_+)^T \Sigma_+^- (\mathbf{x} - \boldsymbol{\mu}_+)) \times p}{\frac{1}{\sqrt{2\pi|\Sigma_+|}} \exp(-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu}_+)^T \Sigma_+^- (\mathbf{x} - \boldsymbol{\mu}_+)) \times p + \frac{1}{\sqrt{2\pi|\Sigma_-|}} \exp(-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu}_-)^T \Sigma_-^- (\mathbf{x} - \boldsymbol{\mu}_-)) \times (1-p)} \\ &= \frac{1}{1 + \frac{1-p}{p} \sqrt{\frac{|\Sigma_+|}{|\Sigma_-|}} \exp(-\frac{1}{2}[(\mathbf{x} - \boldsymbol{\mu}_-)^T \Sigma_-^- (\mathbf{x} - \boldsymbol{\mu}_-) - (\mathbf{x} - \boldsymbol{\mu}_+)^T \Sigma_+^- (\mathbf{x} - \boldsymbol{\mu}_+)])} \end{aligned}$$

The last second equality uses the fact that $P(Y = 1) = p$ and $\mathbf{X}|Y$ is normally distributed. Now, expanding the product terms in exponent of the last equality and collecting the common terms, we get equation (31).

This result was used in synthetic data experiments to get the Bayes classifiers using η .

D Experiments demonstrating the role of α and γ in cost sensitive learning on clean data

In this section, our goal is to understand the role of α and γ in $l_{\alpha, un}(f(\mathbf{x}, y))$. Specifically, we work with α -weighted uneven margin squared loss defined below:

$$l_{\alpha, usq}(f(\mathbf{x}), y) = (1-\alpha)\mathbf{1}_{\{y=1\}}(1-f(\mathbf{x}))^2 + \alpha\mathbf{1}_{\{y=-1\}}\frac{1}{\gamma}(1+\gamma f(\mathbf{x}))^2, \quad \gamma > 0 \quad (32)$$

We consider 3 performance measures, viz., Accuracy (Acc), Arithmetic mean (AM) and F measure.

Every dataset is partitioned into training and test set with 80-20 split. Further, for tuning the parameters like γ , we split the training data into 65-15 and use 15% as validation set for tuning the parameters. To account for the randomness while partitioning of the data, this process is repeated 10 times. The reported values are performance measures averaged over 10 trials along with their standard deviation. To choose the value of the regularization parameter λ , we perform a CV over the set $\Lambda = \{.01, .1, 1, 10\}$. Next, we separately provide examples for the cases where there is a need for differential costing of misclassifications.

D.1 Data has only class imbalance and no user given cost

This is the case when there is only imbalance in the dataset and there is no explicit requirement of costing by the domain. In literature, researchers use only α to account for the imbalance, i.e., use surrogate loss function $l_\alpha(f(\mathbf{x}), y)$ given below:

$$l_\alpha(f(\mathbf{x}), y) = (1 - \alpha)l_1(f(\mathbf{x})) + \alpha l_{-1}(f(\mathbf{x})) \quad (33)$$

where $\alpha > 0$ is tuned. However, as stated in [19], uneven margins $\gamma \neq 1$ have been found to yield improved empirical performance in classification problems involving imbalanced data. With this motivation, we demonstrate the performance of cost sensitive linear classifier from $l_{\alpha, usq}$ based regularized ERM on two synthetic datasets across different level of imbalances given in the set $Im = \{0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8\}$ with 3 approaches given below:

1. Ap1 : α is tuned and $\gamma = 1$
2. Ap2 : γ is tuned and $\alpha = 0.5$
3. Ap3 : α and γ both are tuned

Results on the following two synthetic datasets w.r.t. Accuracy, AM, and F measure are presented in Table 4, 5 and 6 respectively.

Syn-dataset1 We first generate 1000 binary class labels Y from Bernoulli distribution with parameter $\pi \in Im$ where $\pi = P(Y = 1)$. Then a 2-dimensional feature vector \mathbf{X} for each label is drawn from two different Gaussian distributions: $\mathbf{X}|Y = 1 \sim N([1, 0], \Sigma)$ & $\mathbf{X}|Y = -1 \sim N([-1, 0], \Sigma)$ where $\Sigma = [[1, -0.25]; [-0.25, 1]]$.

Syn-dataset2 We first generate 1000 binary class labels Y from Bernoulli distribution with parameter $\pi \in Im$ where $\pi = P(Y = 1)$. Then, a 2-dimensional feature vector \mathbf{X} for each label is drawn from two different Gaussian distributions: $\mathbf{X}|Y = 1 \sim N([0.9, -0.5], \Sigma)$ & $\mathbf{X}|Y = -1 \sim N([-0.9, -0.8], \Sigma)$ where $\Sigma = [[1, 0.5]; [0.5, 1]]$.

π	Syn-dataset1			Syn-dataset2		
	α tuned; $\gamma = 1$	γ tuned; $\alpha = 0.5$	α & γ tuned	α tuned; $\gamma = 1$	γ tuned; $\alpha = 0.5$	α & γ tuned
0.2	0.877 \pm 0.017	0.873 \pm 0.017	0.877 \pm 0.018	0.877 \pm 0.018	0.872 \pm 0.016	0.871 \pm 0.014
0.3	0.853 \pm 0.024	0.852 \pm 0.025	0.853 \pm 0.025	0.852 \pm 0.022	0.856 \pm 0.025	0.854 \pm 0.024
0.4	0.850 \pm 0.016	0.852 \pm 0.016	0.852 \pm 0.016	0.841 \pm 0.018	0.843 \pm 0.014	0.841 \pm 0.012
0.5	0.851 \pm 0.016	0.849 \pm 0.0160	0.852 \pm 0.016	0.825 \pm 0.018	0.834 \pm 0.018	0.8325 \pm 0.0172
0.6	0.859 \pm 0.018	0.857 \pm 0.019	0.857 \pm 0.019	0.852 \pm 0.022	0.85 \pm 0.021	0.852 \pm 0.022
0.7	0.885 \pm 0.020	0.884 \pm 0.016	0.882 \pm 0.017	0.869 \pm 0.015	0.872 \pm 0.023	0.874 \pm 0.018
0.8	0.895 \pm 0.009	0.900 \pm 0.010	0.9015 \pm 0.0105	0.892 \pm 0.0120	0.890 \pm 0.014	0.895 \pm 0.013

Table 4. Averaged accuracy over 10 trials of the linear classifiers obtained from $l_{\alpha, usq}$ based regularized ERM on two synthetic datasets. The 3 columns for each dataset correspond to the three approaches mentioned when there is no user given cost but imbalance in the data.

π	Syn-dataset1			Syn-dataset2		
	α tuned; $\gamma = 1$	γ tuned; $\alpha = 0.5$	α & γ tuned	α tuned; $\gamma = 1$	γ tuned; $\alpha = 0.5$	α & γ tuned
0.2	0.812 \pm 0.033	0.8195 \pm 0.03	0.782 \pm 0.034	0.8136 \pm 0.028	0.78 \pm 0.041	0.816 \pm 0.024
0.3	0.831 \pm 0.0270	0.839 \pm 0.0281	0.828 \pm 0.032	0.830 \pm 0.036	0.829 \pm 0.031	0.840 \pm 0.034
0.4	0.842 \pm 0.021	0.847 \pm 0.018	0.849 \pm 0.016	0.821 \pm 0.017	0.827 \pm 0.020	0.827 \pm 0.019
0.5	0.850 \pm 0.0170	0.848 \pm 0.016	0.848 \pm 0.014	0.831 \pm 0.019	0.837 \pm 0.0186	0.837 \pm 0.018
0.6	0.841 \pm 0.023	0.850 \pm 0.024	0.849 \pm 0.019	0.846 \pm 0.023	0.840 \pm 0.026	0.842 \pm 0.023
0.7	0.859 \pm 0.025	0.862 \pm 0.019	0.867 \pm 0.018	0.838 \pm 0.018	0.843 \pm 0.031	0.844 \pm 0.027
0.8	0.845 \pm 0.031	0.846 \pm 0.032	0.850 \pm 0.029	0.844 \pm 0.024	0.841 \pm 0.030	0.843 \pm 0.024

Table 5. Averaged AM over 10 trials of the linear classifiers obtained from $l_{\alpha,usq}$ based regularized ERM on two synthetic datasets. The 3 columns for each dataset correspond to the three approaches mentioned when there is no user given cost but imbalance in the data.

π	Syn-dataset1			Syn-dataset2		
	α tuned; $\gamma = 1$	γ tuned; $\alpha = 0.5$	α & γ tuned	α tuned; $\gamma = 1$	γ tuned; $\alpha = 0.5$	α & γ tuned
0.2	0.671 \pm 0.054	0.670 \pm 0.054	0.677 \pm 0.057	0.661 \pm 0.034	0.659 \pm 0.057	0.670 \pm 0.052
0.3	0.744 \pm 0.039	0.753 \pm 0.040	0.751 \pm 0.043	0.748 \pm 0.035	0.746 \pm 0.044	0.745 \pm 0.041
0.4	0.807 \pm 0.023	0.812 \pm 0.021	0.815 \pm 0.020	0.782 \pm 0.020	0.787 \pm 0.023	0.788 \pm 0.021
0.5	0.847 \pm 0.011	0.846 \pm 0.013	0.849 \pm 0.010	0.833 \pm 0.018	0.837 \pm 0.017	0.837 \pm 0.017
0.6	0.885 \pm 0.015	0.886 \pm 0.013	0.885 \pm 0.015	0.886 \pm 0.015	0.882 \pm 0.016	0.882 \pm 0.015
0.7	0.919 \pm 0.015	0.919 \pm 0.013	0.919 \pm 0.013	0.909 \pm 0.013	0.909 \pm 0.010	0.909 \pm 0.01
0.8	0.936 \pm 0.005	0.939 \pm 0.006	0.939 \pm 0.006	0.935 \pm 0.007	0.934 \pm 0.007	0.936 \pm 0.008

Table 6. Averaged F measure over 10 trials of the linear classifiers obtained from $l_{\alpha,usq}$ based regularized ERM on two synthetic datasets. The 3 columns for each dataset correspond to the three approaches mentioned when there is no user given cost but imbalance in the data.

Observations In Table 4, we observe that when there is only imbalance in the data and no user given cost, all three approaches, i.e., Ap1, Ap2 and Ap3 have almost equal accuracy on both the synthetic datasets considered. When AM measure is used for evaluation, one can observe in Table 5 that Ap2 has marginally better performance for some values of π . Again with F measure, deciding on one approach to handle only class imbalance is not easy. We believe that one can consider either tuning α with $\gamma = 1$ or tuning γ with $\alpha = 0.5$. Also, as can be seen in Table 4, Accuracy values exhibits convexity w.r.t the imbalance π . This provides us with an evidence that Accuracy is not the correct measure when there is imbalance in the data. It will lead to high values of Accuracy in the presence of imbalance.

D.2 Data has no class imbalance but there is a user given cost

This is the case where there a user given cost α and γ can be set to 1. However, using Bupa dataset [3] we show that, in this case also tuning γ can lead to improvement in the performance of the classifier.

	γ tuned & $\alpha = 0.25$	$\gamma = 1$ & $\alpha = 0.25$
Acc	0.971 ± 0.019	0.915 ± 0.035
AM	0.971 ± 0.019	0.914 ± 0.035
F	0.966 ± 0.017	0.923 ± 0.029

Table 7. Averaged accuracy, AM and F measures of the linear classifiers obtained from $l_{\alpha,usq}$ based regularized ERM on Bupa dataset. The dataset has 6 features and is balanced. This is the case when there is a user given cost but the data has class imbalance. The above results show that in this case too tuning γ improves the performance w.r.t all 3 measures.

D.3 Data has class imbalance and there is a user given cost too

This is the scenario when there is a domain requirement of differential costing and the dataset available is imbalanced too. In this case, we fix the value of α and let the data decide the value of γ . Here, we present the performance of linear classifiers learnt from $l_{\alpha,usq}$ based regularized ERM on Syn-dataset 2. The results are presented across different level of imbalances from the set Im with the corresponding cost set as $C = \{0.25, 0.35, 0.45, 0.4, 0.55, 0.65, 0.75\}$, i.e., $\alpha \in C$.

p	α	Acc	AM	F
0.2	0.25	0.870 ± 0.014	0.816 ± 0.030	0.678 ± 0.040
0.3	0.35	0.853 ± 0.023	0.841 ± 0.033	0.746 ± 0.039
0.4	0.45	0.836 ± 0.019	0.825 ± 0.018	0.786 ± 0.020
0.5	0.4	0.829 ± 0.015	0.834 ± 0.017	0.833 ± 0.019
0.6	0.55	0.853 ± 0.020	0.842 ± 0.024	0.879 ± 0.014
0.7	0.65	0.870 ± 0.019	0.846 ± 0.025	0.909 ± 0.013
0.8	0.75	0.882 ± 0.018	0.842 ± 0.027	0.926 ± 0.012

Table 8. Averaged accuracy, AM and F measures of the classifiers obtained from $l_{\alpha,usq}$ based regularized ERM on Syn-dataset2. This is the case when there is a user given cost α and there is imbalance also in the dataset as given by p .

Observations In this case, the user given cost is α and the value of γ is tuned to account for the class imbalance in data. We observe that the values of γ picked up when evaluation measure was Accuracy are almost complementary to the value of γ picked up when evaluation measure is AM. Specifically, with Accuracy measure, for $p < 0.5$ the γ value is less than 1 and for $p > 0.5$, it is more than 1. However, with AM measure, for $p < 0.5$ the γ value is greater than 1 and for $p > 0.5$, it is less than 1. Based on the functional form of $l_{\alpha,un}$, the correct thing would to choose the value of γ as given by AM measure. The F measure based evaluation leads to γ values lying somewhere between the values obtained from Accuracy and AM measure. For $p < 0.5$, there is a mix of both γ less than 1 and more than 1 but for $p > 0.5$, all values of γ are more than 1.

E Additional experimental results when there is a need for differential costing and data has uniform label noise

E.1 F and Weighted Cost measure of $l_{\alpha, usq}$ -regularized ERM based classifiers and Re-sampling based algorithm on UCI datasets

In this section, we present the performance of Algorithm $(\tilde{\eta}, \alpha)$ and $l_{\alpha, usq}$ based regularized ERM on UCI datasets w.r.t. F measure and Weighted Cost measure. Due to lack of a benchmark for F and WC measure of our cost sensitive label predictions, we make a comparison with the F and WC measure of linear classifiers learnt on clean data using SVM. From Table 9 and 10, we conclude that for Breast Cancer and German dataset, at least one of the proposed methods using noisy data is better than SVM classifier trained on clean data. For other two datasets, the predictions made using our schemes fare well keeping in mind that they have been learnt from noisy data..

Dataset	Cost	SVM	ρ	Algorithm $(\tilde{\eta}, \alpha)$: $\tilde{\eta}$ estimate from				$l_{\alpha, usq}$ RegERM
				<i>Lk-fun</i> with l_{sq}	<i>Lk-fun</i> with l_{log}	LSPC	KLIEP	
(n, m_+, m_-)	α	$\rho = 0$						
Breast Cancer (9,77,186)	0.2	0.28±0.07	0.2 0.4	0.49±0.02 0.44±0.01	0.50±0.04 0.43±0.01	0.46±0.07 0.44±0.03	0.45±0.06 0.43±0.02	0.35±0.10 0.32±0.14
Pima Diabetes (8,268,500)	0.16	0.63±0.05	0.2 0.4	0.60±0.03 0.51±0.01	0.62±0.04 0.52±0.61	0.58±0.03 0.51±0.04	0.58±0.02 0.52±0.37	0.64±0.04 0.58±0.07
German (20,300,700)	0.3	0.50±0.05	0.2 0.4	0.52±0.01 0.46±0.1	0.53±0.02 0.46±0.01	0.52±0.03 0.45±0.01	0.48±0.02 0.46±0.01	0.49±0.72 0.45±0.04
Thyroid (5,65,150)	0.25	0.79±0.07	0.2 0.4	0.63±0.12 0.44±0.04	0.63±0.10 0.46±0.02	0.76±0.06 0.46±0.0	0.56±0.26 0.43±0.13	0.65±0.09 0.63±0.14

Table 9. Averaged F measure of the cost sensitive predictions made by Algorithm $(\tilde{\eta}, \alpha)$ and $l_{\alpha, usq}$ based regularized ERM on UCI datasets corrupted by uniform noise.

Dataset	Cost	SVM	ρ	Algorithm $(\tilde{\eta}, \alpha)$: $\tilde{\eta}$ estimate from				$l_{\alpha, usq}$ RegERM
				<i>Lk-fun</i> with l_{sq}	<i>Lk-fun</i> with l_{log}	LSPC	KLIEP	
(n, m_+, m_-)	α	$\rho = 0$						
Breast Cancer (9,77,186)	0.2	10.67±1.93	0.2 0.4	8.40±1.16 9.93±0.09	8.95±0.91 9.98±0.0	9.88± 1.41 9.96± 0.07	7.89±0.27 8.66±1.13	12.30±1.75 15.09 ±2.62
Pima Diabetes (8,268,500)	0.16	20.37±2.76	0.2 0.4	22.34±2.84 28.19±0.51	22.01±2.76 28.07±0.61	28.21±0.11 28.26±3.55	21.83±1.62 24.09±0.89	26.22 ±3.36 33.02±5.09
German (20,300,700)	0.3	31.20±2.83	0.2 0.4	27.57±1.11 30.03±0.45	26.98±1.07 30.20±0.14	29.51±0.97 30.16±0.08	24.56±0.50 24.65±0.37	43.60±5.56 59.54±5.10
Thyroid (5,65,150)	0.25	3.0±0.98	0.2 0.4	5.63±2.36 9.57±0.78	6.29±2.79 9.34±0.58	3.07±0.69 9.18±0.0	7.54±2.17 9.90 ± 2.22	4.68±0.99 5.38±1.87

Table 10. Averaged Weighted Cost of the cost sensitive predictions made by Algorithm $(\tilde{\eta}, \alpha)$ and $l_{\alpha, usq}$ based regularized ERM on UCI datasets corrupted by uniform noise.

E.2 Result of Algorithm $(\tilde{\eta}, \alpha)$ on Bupa dataset when $\tilde{\eta}$ is estimated using $Lk\text{-fun}$ with squared loss

In this section, we obtained the classifiers from Algorithm $(\tilde{\eta}, \alpha)$ when $\tilde{\eta}$ is estimated using $Lk\text{-fun}$ with l_{sq} on Bupa datasets. We compared the cases when the classifiers are learnt by tuning γ and when $\gamma = 1$. We observed that tuning γ is always better than fixing γ at 1 because all the classifiers have improved value of all performance measures with the former.

Algorithm $(\tilde{\eta}, \alpha)$: $\tilde{\eta}$ estimate from $Lk\text{-fun}$ with l_{sq}								
ρ	γ tuned				$\gamma = 1$			
	Acc	AM	F	WC	Acc	AM	F	WC
0.0	0.96±0.02	0.97±0.02	0.96±0.02	1.49±1.02	0.90±0.03	0.90±0.03	0.91±0.02	1.67±0.84
0.1	0.96±0.02	0.96±0.03	0.97±0.02	3.70±2.05	0.71±0.88	0.70±0.09	0.77±0.05	6.12±1.61
0.2	0.9±0.03	0.89±0.03	0.90±0.03	5.65±0.05	0.56±0.05	0.55±0.06	0.69±0.02	7.95±0.79
0.3	0.85±0.04	0.85±0.04	0.81±0.1	5.4±1.74	0.51±0.01	0.50±0.01	0.67±0.01	8.5±0.23
0.4	0.59±0.09	0.58±0.09	0.67±0.00	5.66±0.00	0.51±0.0	0.50±0.01	0.67±0.01	8.32±0.37

Table 11. Dataset: BUPA Liver Disorder. The above table depicts the performance of Algorithm $(\tilde{\eta}, \alpha)$ when $\tilde{\eta}$ estimate is from $Lk\text{-fun}$ with l_{sq} based regularized ERM (λ tuned via CV). Here, $\alpha = 0.25$ and γ is tuned.

E.3 Performance of $l_{\alpha, usq}$ -regularized ERM based classifier and Re-sampling based algorithm on Synthetic dataset

Syn3D_{imb} (adapted from [4]): Generate 4000 binary class labels $Y \sim \text{Bern}(0.35)$. Then, a 3-dimensional feature vector \mathbf{X} for each label is drawn from two different Gaussian distributions: $\mathbf{X}|Y = 1 \sim N([1.5; 0; -1], \Sigma)$ & $\mathbf{X}|Y = -1 \sim N([-1.5; 1; 1], \Sigma)$ where $\Sigma = [[2, -0.3, -0.3]; [-0.3, 2, -0.3]; [-0.3, -0.3, 2]]$. For $\alpha = 0.3$ and γ tuned suitably for the measure in consideration, we compare the quality of cost sensitive predictions from Algorithm $(\tilde{\eta}, \alpha)$ and $l_{\alpha, usq}$ based regularized ERM implemented on corrupted data with cost sensitive Bayes classifier $f_{0-1, \alpha, \gamma}^*(\mathbf{x}) = \text{sign}(\eta(\mathbf{x}) - \frac{\alpha}{\gamma + \alpha(1-\gamma)})$ learnt on clean data.

It can be observed in Table 13 that Accuracy and AM values for Algorithm $(\tilde{\eta}, \alpha)$ and $l_{\alpha, usq}$ based classifiers are comparable to that of $f_{0-1, \alpha, \gamma}^*$ till the noise rate ρ does not cross 0.4 after which it starts deteriorating. However, as shown in Table 14, for F measure and Weighted Cost, this threshold is reached at close to $\rho = 0.3$ implying that these measures are more sensitive to noise.

$f_{0-1, \alpha, \gamma}^*$ performance				
	AC	AM	F	WC
$\rho = 0$	0.90±0.01	0.88±0.01	0.85±0.01	42.91±4.01

Table 12. Performance of cost weighted Bayes classifier on $Syn3D_{imb}$ w.r.t. various performance measures. The misclassification cost α is taken as 0.3 and γ is tuned.

ρ	Accuracy			AM		
	Algorithm $(\tilde{\eta}, \alpha)$: $\tilde{\eta}$ from			Algorithm $(\tilde{\eta}, \alpha)$: $\tilde{\eta}$ from		
	Link-func with l_{log}	LSPC	$l_{\alpha, usq}$ RegERM	Link-func with l_{log}	LSPC	$l_{\alpha, usq}$ RegERM
0.1	0.90±0.01	0.89±0.01	0.90±0.01	0.89±0.01	0.89±0.01	0.88±0.01
0.2	0.89±0.01	0.89±0.00	0.90±0.01	0.88±0.01	0.88±0.01	0.89±0.01
0.3	0.89±0.01	0.89±0.01	0.89±0.01	0.88±0.01	0.88±0.01	0.88±0.01
0.4	0.88±0.01	0.71±0.04	0.89±0.01	0.87±0.01	0.69±0.03	0.88±0.01
0.45	0.83±0.03	0.60±0.06	0.80±0.06	0.81±0.07	0.57±0.05	0.81±0.07
0.49	0.46±0.16	0.38±0.02	0.48±0.16	0.53±0.14	0.51±0.01	0.54±0.15

Table 13. Comparison of Algorithm $(\tilde{\eta}, \alpha)$ and $l_{\alpha, usq}$ based regularized ERM’s linear classifier on Accuracy and AM measure. The value of $\alpha = 0.3$ and γ is tuned.

ρ	F measure			Weighted Cost		
	Algorithm $(\tilde{\eta}, \alpha)$: $\tilde{\eta}$ from			Algorithm $(\tilde{\eta}, \alpha)$: $\tilde{\eta}$ from		
	Link-func with l_{log}	LSPC	$l_{\alpha, usq}$ RegERM	Link-func with l_{log}	LSPC	$l_{\alpha, usq}$ RegERM
0.1	0.85±0.02	0.85±0.01	0.85±0.02	45.62±3.37	40.16±3.61	44.06±3.58
0.2	0.84±0.01	0.83±0.01	0.84±0.01	60.40±4.14	58.01±7.14	58.09±4.75
0.3	0.69±0.03	0.61±0.01	0.68±0.01	106.9±7.35	109.77±11.24	107.47±6.84
0.4	0.53±0.01	0.52±0.01	0.53±0.01	187.41±6.45	158.76±1.76	190.10±6.20
0.45	0.51±0.00	0.51±0.00	0.52±0.00	200.37±0.24	159.55±0.0	200.49±0.00
0.49	0.51±0.00	0.52±0.02	0.52±0.00	200.49±2.84	159.55±0.0	200.49±0.00

Table 14. Comparison of Algorithm $(\tilde{\eta}, \alpha)$ and $l_{\alpha, usq}$ based regularized ERM’s linear classifier on F and Weighted Cost measure. The value of $\alpha = 0.3$ and γ is tuned.

F Various in-class probability estimation methods

In this section, we present a comparative study between 3 existing in-class probability η estimation methods and a minimum squared deviation method. The later method is based on our idea that η is a conditional expectation and hence a minimizer of a suitable expected squared loss. We first present the details of 3 existing η estimation methods viz., Link function based method, LSPC, KLIEP and k -Nearest Neighbour (k -NN). Then, we present our method and various estimators based on KLIEP.

F.1 Link function based approach [18]

This method is the most commonly used approach for getting the in-class probabilities from a given dataset. In this method, first a classifier is learnt from the data using a special class of loss functions and then a suitable transformation is used to get the η estimates. This special class of loss functions is the class of strongly proper loss functions defined below:

Definition 4 (Strongly proper loss [1]) Let $c : \{\pm 1\} \times [0, 1] \mapsto \mathbb{R}_+$ be a binary Class probability estimation (CPE) loss and let $\lambda > 0$. Then, c is λ -strongly proper if $\forall \eta, \hat{\eta} \in [0, 1]$

$$C_c(\eta, \hat{\eta}) - H_c(\eta) \geq \frac{\lambda}{2}(\eta - \hat{\eta})^2$$

where $H_c(\eta) = \inf_{\eta' \in [0,1]} C_c(\eta, \eta')$.

Therefore, a loss $l : \{\pm 1\} \times \mathbb{R} \mapsto \mathbb{R}_+$ is said to be strongly proper composite if it can be written as

$$l(y, \hat{y}) = c(y, \psi^{-1}(\hat{y}))$$

for some strongly proper loss c and a strictly increasing link function $\psi : [0, 1] \mapsto \mathbb{R}$.

Given a strongly proper loss function l , one can obtain an optimal classifier $f^*(\mathbf{x})$ by minimizing l based empirical risk minimization (ERM). Then, the in-class probability can be computed as $\eta_l(\mathbf{x}) = \psi_l^{-1}(f(\mathbf{x}))$. Some examples of strongly proper composite loss functions along with their corresponding inverse link function $\psi^{-1}(\cdot)$ [1, 24] are as follows:

1. Logistic loss: $l_{log}(f(\mathbf{x}), y) = \log(1 + e^{-yf(\mathbf{x})}) / \log(2)$, $\psi_{log}^{-1}(f(\mathbf{x})) = \frac{1}{1+e^{-f(\mathbf{x})}}$
2. Squared loss: $l_{sq}(f(\mathbf{x}), y) = (1 - yf(\mathbf{x}))^2$, $\psi_{sq}^{-1}(f(\mathbf{x})) = \frac{1+f(\mathbf{x})}{2}$
3. Mod-squared loss: $l_{msq}(f(\mathbf{x}), y) = (\max\{0, (1 - yf(\mathbf{x}))\})$, $\psi_{msq}^{-1}(f(\mathbf{x})) = \frac{1+T(f(\mathbf{x}))}{2}$ where $T(p) = \min\{\max\{p, -1\}, 1\}$
4. Exponential loss: $l_{exp}(f(\mathbf{x}), y) = e^{-yf(\mathbf{x})}$, $\psi_{log}^{-1}(f(\mathbf{x})) = \frac{1}{1+e^{-2f(\mathbf{x})}}$

It is shown in [1] that if l is a strongly proper loss then the L_2 distance between $\psi_l^{-1}(f(\mathbf{X}))$ and $\eta_l(\mathbf{X})$ can be upper bounded in terms of l -regret of f . Some special cases of this result are also presented in [24], [20] (for logistic loss). In this method, for some loss functions like squared loss and mod-squared loss, one has to suitably truncate the inverse link function ψ_l^{-1} so that it is a valid probability.

The quality of η estimate will depend upon the quality of the classifier too. So, whether one uses linear or nonlinear (kernel) classifiers is expected to make a difference. Also, in case of non-linear classifier, regularization parameter tuning is required.

As hinge loss is not a strongly proper loss, one cannot use the above approach to get the $\eta(\mathbf{x})$. It was shown in [17] that one can get the in-class probability by using a sigmoid transformation and solving an optimization problem in 2 variables. However, in [24], it is shown that SVM formulation doesn't lead to reliable probability estimates.

For the sake of convenience, a compact form of this scheme is available in Algorithm ($\mathbf{f}, \hat{\boldsymbol{\eta}}$).

Algorithm ($f, \hat{\eta}$) $\hat{\eta}$ via classifier scheme

Input: Training data $\tilde{S}_{tr} = \{(x_1, y_i)\}_{i=1}^{m_{tr}}$, test data $S_{te} = \{(x_i, y_i)\}_{i=1}^{m_{te}}$, strongly proper composite loss $l(f(x), y)$, its inverse link $\psi_l^{-1}(\cdot)$.

Output: $\hat{\eta}$ estimate of η on test data S_{te}

- 1: Compute $\hat{f}_{D,l}^* = \arg \min_{f \in \mathcal{H}_{lin}} \sum_{i=1}^{m_{tr}} l(f(x_i), y_i)$.
 - 2: **for** $i = 1, 2, \dots, m_{te}$ **do**
 - 3: Compute the estimate $\hat{\eta}(x_i) = \psi_l^{-1}(\hat{f}_{D,l}^*(x_i))$.
 - 4: **end for**
-

F.2 Least Squared Probabilistic Classifier (LSPC) [21]

Least Squared Probabilistic Classifier (LSPC) [21] is a classifier which makes predictions based on the in-class probability η . This is a direct η estimation method. Suppose the in-class probability is modeled as follows :

$$q(y|\mathbf{x}, \boldsymbol{\alpha}) := \sum_{l=1}^b \alpha_l \phi_l(\mathbf{x}, y) = \boldsymbol{\alpha}^T \boldsymbol{\phi}(\mathbf{x}, y)$$

where $\boldsymbol{\alpha}$, a b -dimensional vector is to be learned from the sample and $\boldsymbol{\phi}(\mathbf{x}, y)$ is non-negative b -dimensional vector of basis functions. The parameter $\boldsymbol{\alpha}$ is obtained so that the following squared error J is minimized:

$$J = \frac{1}{2} \sum_{y=1}^c \int (q(y|\mathbf{x}, \boldsymbol{\alpha}) - p(y|\mathbf{x}))^2 p(\mathbf{x}) d\mathbf{x} \quad (34)$$

Here, c is the number of classes. Solving a regularized empirical version of the above optimization problem leads to the following closed form expression for the optimal value of $\boldsymbol{\alpha}$,

$$\hat{\boldsymbol{\alpha}} = (\hat{H} + \lambda I_b)^{-1} \hat{h}$$

where $\hat{H} = \frac{1}{m} \sum_{y=1}^c \sum_{i=1}^m \phi(\mathbf{x}_i, y_i) \phi(\mathbf{x}_i, y_i)^T$ and $\hat{h} = \frac{1}{m} \sum_{i=1}^m \phi(\mathbf{x}_i, y_i)$ and λ is the regularization parameter. The final solution is obtained by normalizing the solution to get,

$$\hat{p}(y|\mathbf{x}) = \frac{\max(0, \hat{\boldsymbol{\alpha}}^T \boldsymbol{\phi}(\mathbf{x}, y))}{\sum_{y'=1}^c \max(0, \hat{\boldsymbol{\alpha}}^T \boldsymbol{\phi}(\mathbf{x}, y'))} \quad (35)$$

In particular, LSPC uses the Gaussian kernels centered at training points. Use of kernel functions, allows to learn the parameters in a class-wise manner. The kernel parameters are chosen based on the dataset and the regularization parameter is chosen by CV. It is claimed that LSPC can be viewed as an application of density ratio estimation method called unconstrained Least Squared

Importance Fitting method (uLSIF [7]) and hence the theoretical guarantees like consistency and stability will follow.

However, the consistency results of [7] are under the following assumption. If n_{nu} and n_{de} are the data points from the numerator and denominator density respectively, then $n_{nu} = \omega(n_{de}^2)$. This means that there exists c such that $|cn_{de}^2| < |n_{nu}|$. Now, the density ratio in case of in-class probability is $\eta(\mathbf{x}) = \frac{P(\mathbf{x}, y)}{P(\mathbf{x})}$. Therefore, the number of sample points in numerator and denominator would be equal i.e., $n_{nu} = n_{de}$ and the assumption of uLSIF is not satisfied.

Centiles from the data are known to be good candidates for Gaussian kernel width. Pseudo code given in Pseudocode Centile method is borrowed from the code made available by [21]. It generates the set of kernel width σ to be chosen from for tuning Gaussian kernel's parameter σ .

Pseudocode Centile method Method for generating a set of σ values for Gaussian kernels.

Input: Sample points $\mathbf{x}_i, i = 1, \dots, m$, Centile Vector C of size k e.g. $C = \{10, 20, \dots, 90\}$ implies $k=9$.

Output: Σ_{set} , a vector of size k containing possible σ values.

- 1: Set number of pairs for which the distance is computed as $np = \min(2000, m)$ and $n_{cen} = k$.
 - 2: $randOrder1 = \pi(1 : m)$ and $randOrder2 = \pi(1 : m)$ where $\pi(\cdot)$ is a permutation of $(1, \dots, m)$.
 - 3: **for** $i = 1, 2, \dots, np$ **do**
 - 4: $dist[i] = \|\mathbf{x}_{randOrder1[i]} - \mathbf{x}_{randOrder2[i]}\|_2^2$
 - 5: **end for**
 - 6: Sort the vector $dist$.
 - 7: **for** $i = 1, \dots, n_{cen}$ **do**
 - 8: $out[i] = dist \left[np \times \frac{C[i]}{100} \right]$
 - 9: **end for**
 - 10: Return $\Sigma_{set} = \sqrt{out}$.
-

F.3 Kullback-Leibler Importance Estimation Procedure (KLIEP) [22]

The in-class probability for a class y can be written as

$$P(Y = y|\mathbf{x}) = \frac{P(\mathbf{x}, y)}{P(\mathbf{x})} = \frac{P(\mathbf{x}|Y = y)}{P(\mathbf{x})}P(Y = y) \quad (36)$$

The authors in [22] deal with the problem of estimating importance $w(\mathbf{x}) := \frac{p_{nu}(\mathbf{x})}{p_{de}(\mathbf{x})}$ using $\{\mathbf{x}_i^{de}\}_{i=1}^{n_{de}}$ and $\{\mathbf{x}_j^{nu}\}_{j=1}^{n_{nu}}$ from the probability density p_{de} and p_{nu} respectively. The estimated ratio is given a linear parametric form of $\hat{w}(\mathbf{x}) = \sum_{l=1}^b \alpha_l \phi_l(\mathbf{x})$. Here, ϕ_l are known non negative basis functions and α_l are the parameters to be estimated. The numerator density can be estimated as $\hat{p}_{nu}(\mathbf{x}) =$

$\hat{w}(\mathbf{x})p_{de}(\mathbf{x})$. The parameters are determined by minimizing the KL divergence between the true and estimated numerator density as follows:

$$\begin{aligned} KL[p_{nu}(\mathbf{x})\|\hat{p}_{nu}(\mathbf{x})] &= \int p_{nu}(\mathbf{x}) \log \frac{p_{nu}(\mathbf{x})}{\hat{w}(\mathbf{x})p_{de}(\mathbf{x})} d\mathbf{x} \\ &= \int p_{nu}(\mathbf{x}) \log \frac{p_{nu}(\mathbf{x})}{p_{de}(\mathbf{x})} d\mathbf{x} - \int p_{nu}(\mathbf{x}) \log \hat{w}(\mathbf{x}) d\mathbf{x} \end{aligned} \quad (37)$$

The first term in equation (37) is independent of the $\{\alpha_l\}_{l=1}^b$, and incorporating the fact that $\hat{p}_{nu}(\mathbf{x})$ is probability density, the following optimization problem is obtained:

$$\max_{\{\alpha_l\}_{l=1}^b} \left[\sum_{j=1}^{n_{nu}} \log \left(\sum_{l=1}^b \alpha_l \phi_l(\mathbf{x}_j^{nu}) \right) \right] \quad (38)$$

$$\text{subject to } \sum_{i=1}^{n_{de}} \sum_{l=1}^b \alpha_l \phi_l(\mathbf{x}_i^{de}) = n_{de} \quad \text{and} \quad \alpha_1, \dots, \alpha_b \geq 0 \quad (39)$$

This is a convex optimization problem and the authors have provided an algorithm based on gradient ascent to solve this problem. The solution is claimed to be sparse and hence leads to a reduction in computational complexity.

F.4 k -Nearest Neighbour method for η estimation [2]

In classical binary setup, k -NN works based on the labels of the neighbours. Firstly, an estimate of in-class probability is computed and then it is thresholded to make the label prediction.

$$\hat{\eta}_{k-NN}(\mathbf{x}) = \frac{1}{k} \sum_{i=1}^k Y_{(i)}(\mathbf{x})$$

Here, $Y_{(i)}$'s are ordered w.r.t. some distance measure ρ . Choice of k and ρ are the two levers available while learning. Also, there is a trade-off between bias and variance.

F.5 Minimization of Squared Deviation method

This is a direct η estimation method, i.e., it doesn't involve learning a classifier. It is known that the expected squared loss is minimized by a suitable conditional expectation. In particular, given two r.v.s P, Q such that Q is integrable, $Z^* = E[Q|P]$ is the minimizer of $E[(Q - Z)^2]$ a.s. among all random variables Z , i.e.,

$$\forall h(P), \quad E[(Q - Z^*)^2] \leq E[(Q - h(P))^2]$$

where h belongs to the class of all measurable functions. Given (\mathbf{X}, Y) , $\mathbf{X} \in \mathbb{R}^n$, $Y \in \{-1, 1\}$ and interpretation of $\eta(\mathbf{x}) = P(Y = 1|\mathbf{X})$ as the conditional

expectation, we are interested in knowing what is the corresponding P, Q in the binary classification framework.

$$\begin{aligned} E[Q|P] &= P(Y = 1|\mathbf{X}) \quad a.s. \\ &= E[\mathbf{1}_{\{Y=1\}}|\mathbf{X}] \quad a.s. \end{aligned}$$

Therefore, we have $P = \mathbf{X}$ and $Q = \mathbf{1}_{\{Y=1\}}$ and the minimization problem is

$$\min_Z E_{\mathbf{X} \times Y} [(\mathbf{1}_{\{Y=1\}} - Z)^2] \quad (40)$$

From above, $Z^* = \eta = P[Y = 1|\mathbf{X}]$ *a.s.* and hence $Z^* \in [0, 1]$ *a.s.* as a function of \mathbf{X} . This has some resemblance to the squared loss for CPE which is $l_1 = (1 - \hat{\eta})^2$, $l_{-1} = \hat{\eta}^2$. The advantage here is that the minimizer is the conditional expectation which in our case is the conditional probability. This need not be true for the log-loss ($l_1(\hat{\eta}) = \log(\hat{\eta})$, $l_{-1}(\hat{\eta}) = \log(1 - \hat{\eta})$), etc.

Consider the class of positive valued functions with an element denoted by $\phi_l(\cdot, \cdot)$. Let ϕ_l be called a basis function. Later, we will be specializing this class to be the class of Gaussian kernels. Now, consider the following functional form of Z in equation (40),

$$Z = \sum_{l=1}^b \alpha_l \phi_l(\mathbf{X}, Y) = \boldsymbol{\alpha}^T \boldsymbol{\phi}(\mathbf{X}, Y) \quad (41)$$

where $\boldsymbol{\alpha} \in \mathbb{R}^b$ is the variable to be optimized over and $\boldsymbol{\phi}(\mathbf{X}, Y) \geq 0$ is the vector of basis functions ϕ_l , $l = 1, \dots, b$, i.e., Z is a weighted combination of the positive basis functions. An L_2 regularized version of equation (40) can be rewritten as follows:

$$\min_{\boldsymbol{\alpha} \in \mathbb{R}^b} E_{\mathbf{X} \times Y} [(\mathbf{1}_{\{Y=1\}} - \boldsymbol{\alpha}^T \boldsymbol{\phi}(\mathbf{X}, Y))^2] + \lambda \boldsymbol{\alpha}^T \boldsymbol{\alpha} \quad (42)$$

Here, $\lambda > 0$ is the regularization parameter which is usually determined using CV. If one considers the multi-class classification problem with c classes, then the in-class probability for each class is defined as follows:

$$\eta_y(\mathbf{x}) = P(Y = y|\mathbf{X} = \mathbf{x}), \quad y \in \{1, 2, \dots, c\} \quad (43)$$

The idea that expected squared loss is minimized by a suitable conditional expectation can be extended to multi class problems too by formulating following optimization problems, one for each class $y \in \{1, 2, \dots, c\}$, to get the $\eta_y(\mathbf{x})$'s.

$$\min_{\boldsymbol{\alpha}^{(y)}} E_{\mathbf{X} \times Y} [(\mathbf{1}_{\{Y=y\}} - \boldsymbol{\alpha}^{(y)T} \boldsymbol{\phi}(\mathbf{X}, Y))^2] + \lambda \boldsymbol{\alpha}^{(y)T} \boldsymbol{\alpha}^{(y)}, \quad \lambda > 0 \quad (44)$$

Solving for $\boldsymbol{\alpha}^{(y)}$ will lead to the following system of equations,

$$\boldsymbol{\alpha}_{reg}^{(y)*} = [E[\boldsymbol{\phi}(\mathbf{X}, Y)\boldsymbol{\phi}^T(\mathbf{X}, Y)] + \lambda I_b]^{-1} E[\boldsymbol{\phi}(\mathbf{X}, Y)\mathbf{1}_{\{Y=y\}}] \quad (45)$$

Given the multi-class data (\mathbf{x}_i, y_i) , $i \in \{1, 2, \dots, m\}$, $\mathbf{x}_i \in \mathbf{R}^n$, $y \in \{1, \dots, c\}$, we have the sample versions of the solution as follows:

$$\hat{\boldsymbol{\alpha}}_{reg}^{(y)} = \left[\underbrace{\frac{1}{m} \sum_{i=1}^m \phi(\mathbf{x}_i, y_i) \phi^T(\mathbf{x}_i, y_i)}_{\hat{A}} + \lambda \mathbf{1}_b \right]^{-1} \left[\underbrace{\frac{1}{m} \sum_{i \in m_y} \phi(\mathbf{x}_i, y_i)}_{\hat{\mathbf{a}}^{(y)}} \mathbf{1}_{\{y_i=y\}} \right] \quad (46)$$

$$= [\hat{A} + \lambda \mathbf{1}_b]^{-1} \hat{\mathbf{a}}^{(y)} \quad (47)$$

where m_y is the index of examples with label y . One important question here is to decide about the nature of basis functions ϕ . Kernel function $K(\mathbf{x}, \mathbf{x}')$ comes as natural choice for the basis functions due to the information they carry. Using these kernels, the elements of \hat{A} and $\hat{\mathbf{a}}^{(y)}$ can be written as follows:

$$\hat{A}_{l,l'} = \frac{1}{m} \sum_{i=1}^m K(\mathbf{x}_i, \mathbf{x}_l) K(\mathbf{x}_i, \mathbf{x}_{l'})$$

$$\hat{\mathbf{a}}^{(y)} = \frac{1}{m} \sum_{i \in m_y} K(\mathbf{x}_i, \mathbf{x}_l)$$

Therefore, here we would get c sets of $\hat{\boldsymbol{\alpha}}^{(y)}$, one from each system of equation. To get the final estimates of the in-class probabilities, one has to do the max and the normalization as follows:

$$\hat{\eta}_y(x) = \frac{\max\{0, \sum_{l=1}^m \hat{\boldsymbol{\alpha}}^{(y)} K(\mathbf{x}, \mathbf{x}_l)\}}{\sum_{y'=1}^c \max\{0, \sum_{l=1}^m \hat{\boldsymbol{\alpha}}^{(y')} K(\mathbf{x}, \mathbf{x}_l)\}} \quad \forall y \in \{1, 2, \dots, c\} \quad (48)$$

In-class probability for a binary classification problem can be written as a special case of equation (48).

Remark 5. It is not necessary that the regularized expected squared loss is minimized by conditional expectation. Mathematically,

$$\min_{\boldsymbol{\alpha} \in \mathbf{R}^b} E[(\mathbf{1}_{\{Y=1\}} - Z)^2] + \lambda \boldsymbol{\alpha}^T \boldsymbol{\alpha} \quad (49)$$

where $Z = \boldsymbol{\alpha}^T \phi(\mathbf{X}, Y)$. Here, it is not necessary that $Z^* = \eta = E[\mathbf{1}_{\{Y=1\}} | \mathbf{X}]$. So, there would be an approximation error if one tries to use minimizer of regularized expected squared loss to estimate in-class probability. Let us denote the estimates of in-class probability from regularized and unregularized version by (40) as $\hat{\eta}_{reg}$ and $\hat{\eta}$ respectively. We expect $\hat{\eta}$ to have model explaining properties and $\hat{\eta}_{reg}$ to have good prediction properties due to regularization and avoidance of overfitting.

Remark 6. It can be observed that the solution in equation (48) and (35) are same even though the initial objective is quite different. Therefore, in the experiments, we denote the results by the name of LSPC.

F.6 Various estimators based on KLIEP

We use KLIEP to estimate $\frac{P(\mathbf{x}|Y=y)}{P(\mathbf{x})}$, ratio of two continuous densities. One requirement in this method is that the ratio is such that its product with estimated class marginal, $\hat{\pi} = \hat{P}(Y = y)$ is a valid probability i.e., $\hat{P}(Y = y|\mathbf{x}) < 1$. This condition need not be true in general as we observed in our experiments.

For choosing the number of basis functions b , the authors in [22] suggest to use $b = \min(100, n_{nu})$. In our case, n_{nu} would be number of points from class y . As seen in equation (36), one can either use KLIEP with positive class points in the numerator and get $\hat{\eta}$ or use KLIEP twice, one with positive class in the numerator and other with negative class in the numerator and get $\hat{\eta}$.

Based on these observations, we have come up with following estimates of η using KLIEP.

1. $\hat{\eta}_{\text{pos}}$: This estimator is obtained by implementing KLIEP with positive class points in the numerator and obtaining the ratio \hat{w}_{pos} . That is, the ratio of density to be estimated is $\frac{P(\mathbf{x}|Y=1)}{P(\mathbf{X})}$.

$$\hat{\eta}_{\text{pos}}(\mathbf{x}) = \hat{w}_{\text{pos}}(\mathbf{x})\hat{\pi}$$

This estimator of η integrates to 1 but there is no guarantee that it will always be less than 1.

2. $\hat{\eta}_{\text{neg}}$: This estimator is obtained by implementing KLIEP with negative class points in the numerator and obtaining the ratio \hat{w}_{neg} . That is, the ratio of density to be estimated is $\frac{P(\mathbf{x}|Y=-1)}{P(\mathbf{X})}$.

$$\hat{\eta}_{\text{neg}}(\mathbf{x}) = 1 - \hat{w}_{\text{neg}}(\mathbf{x})(1 - \hat{\pi})$$

This estimator of η integrates to 1 but it is prone to being less than 0.

3. $\hat{\eta}_{\text{norm}}$: This estimator is obtained by implementing two KLIEPs, one with positive class in the numerator and other with negative class in the numerator. Then, the estimates are combined as follows:

$$\hat{\eta}_{\text{norm}} = \frac{\hat{\eta}_{\text{pos}}}{\hat{\eta}_{\text{pos}} + \hat{\eta}_{\text{neg}}}$$

This η estimate is a valid density as it integrates to 1 and always lies in the interval $[0, 1]$.

4. $\hat{\eta}_{\text{pos},s}$: This is a scaled version of $\hat{\eta}_{\text{pos}}$. Let $\hat{\eta}_{\text{max},p}$ be the maximum value of $\hat{\eta}_{\text{pos}}$ in the training data. Then,

$$\hat{\eta}_{\text{pos},s} = \frac{\hat{\eta}_{\text{pos}}}{\hat{\eta}_{\text{max},p}}$$

This estimator will integrate to 1 and will be in the interval $[0, 1]$ for training points but for test points it can cross 1.

5. $\hat{\eta}_{\text{neg},s}$: This is a scaled version of $\hat{\eta}_{\text{neg}}$. Let $\hat{\eta}_{\text{max},n}$ be the maximum value of $\hat{\eta}_{\text{neg}}$ in the training data. Then,

$$\hat{\eta}_{\text{neg},s} = \frac{\hat{\eta}_{\text{neg}}}{\hat{\eta}_{\text{max},n}}$$

This estimator will integrate to 1 and will be in the interval $[0, 1]$ for training points but for test points it can take values less than 0.

6. $\hat{\eta}_{\text{norm},s}$: This is obtained by normalizing the estimators $\hat{\eta}_{\text{pos},s}$ and $\hat{\eta}_{\text{neg},s}$ as follows:

$$\hat{\eta}_{\text{norm},s} = \frac{\hat{\eta}_{\text{pos},s}}{\hat{\eta}_{\text{pos},s} + \hat{\eta}_{\text{neg},s}}$$

This is a valid probability density estimate as it integrates to 1 and lies in $[0, 1]$.

7. $\hat{\eta}_{\text{Cnorm},s}$: This is also a normalized estimator but the normalizing is different from $\hat{\eta}_{\text{norm},s}$. It is given as below:

$$\hat{\eta}_{\text{Cnorm},s} = \frac{\hat{\eta}_{\text{neg},s}}{\hat{\eta}_{\text{neg},s} + \hat{\eta}_{\text{pos},s}}$$

This is a valid probability density estimate as it integrates to 1 and lies in $[0, 1]$.

8. $\hat{\eta}_{\text{avg},s}$: This is the averaged version of scaled estimators $\hat{\eta}_{\text{pos},s}$ and $\hat{\eta}_{\text{neg},s}$ and given as follows:

$$\hat{\eta}_{\text{avg},n} = \frac{\hat{\eta}_{\text{pos},s} + \hat{\eta}_{\text{neg},s}}{2}$$

This estimator integrates to 1 but for test data points it is prone to lying outside the interval $[0, 1]$.

The estimators which does not lie inside the interval $[0, 1]$ can be truncated to lie inside the interval. However, there performance might be affected. A compact form of the above estimators is available in Table 15.

$\hat{\eta}$	$\hat{\eta}(\mathbf{x}) \in [0, 1]$ for \mathbf{x} in test set
$\hat{\eta}_{\text{pos}}(\mathbf{x}) = \hat{w}_{\text{pos}}(\mathbf{x})\hat{\pi}$	May not
$\hat{\eta}_{\text{neg}}(\mathbf{x}) = 1 - \hat{w}_{\text{neg}}(\mathbf{x})(1 - \hat{\pi})$	May not
$\hat{\eta}_{\text{norm}} = \frac{\hat{\eta}_{\text{pos}}}{\hat{\eta}_{\text{pos}} + \hat{\eta}_{\text{neg}}}$	Yes
$\hat{\eta}_{\text{pos},s} = \frac{\hat{\eta}_{\text{pos}}}{\hat{\eta}_{\text{max},p}}$	May not
$\hat{\eta}_{\text{neg},s} = \frac{\hat{\eta}_{\text{neg}}}{\hat{\eta}_{\text{max},n}}$	May not
$\hat{\eta}_{\text{norm},s} = \frac{\hat{\eta}_{\text{pos},s}}{\hat{\eta}_{\text{pos},s} + \hat{\eta}_{\text{neg},s}}$	Yes
$\hat{\eta}_{\text{Cnorm},s} = \frac{\hat{\eta}_{\text{neg},s}}{\hat{\eta}_{\text{neg},s} + \hat{\eta}_{\text{pos},s}}$	Yes
$\hat{\eta}_{\text{avg},n} = \frac{\hat{\eta}_{\text{pos},s} + \hat{\eta}_{\text{neg},s}}{2}$	May not

Table 15. Details about various KLIEP based η estimates and whether they are valid probability or not. We only check for $\hat{\eta} \in [0, 1]$ as all the estimates integrate to 1.

The comparison between these estimators w.r.t. various measures is done in Section F.7. In our implementation, we have used Gaussian kernels as the basis functions with centres as the sample points from the numerator density (positive/negative class). The kernel width σ is tuned by KLIEP’s inbuilt CV procedure where the criterion is maximum objective value of the KL based optimization problem. The possible candidates for kernel width σ to be used in KLIEP’s CV procedure were generated through Centile method described in Pseudocode Centile method.

F.7 An experimental comparative study of various in-class probability estimation methods

In this section, we compare the performance of the η estimation methods described in Section F.1, F.2 and F.6. We implemented Algorithm $(\mathbf{f}, \hat{\eta})$ with squared loss, logistic loss and modified squared loss and LSPC, KLIEP and k -NN on 3 synthetic datasets so that we can comment on the quality of η estimates. In case of k -NN, the value of number of neighbours k to be used is selected by cross validating w.r.t. the measure to be evaluated on. The dataset generation scheme is adapted from [4]. We consider 9 performance measures out of which first 5 (MSE, RMSE, MAD, MD, KL) are purely for η estimation, next 2 (Acc, BS) are for the scenario when these estimates are used in label prediction and the last 2 measure the algorithms estimation capability at the boundary i.e., maximum and minimum. For an estimate $\hat{\eta}$, when the train data is $\{(\mathbf{x}_i, y_i)\}_{i=1}^{m_{tr}}$ and test data set is $\{(\mathbf{x}_i, y_i)\}_{i=1}^{m_{te}}$, the details about all these measures are given below:

1. Mean squared error (MSE): $MSE = \frac{1}{m_{te}} \sum_{i=1}^{m_{te}} (\hat{\eta}(\mathbf{x}_i) - \eta(\mathbf{x}_i))^2$
2. Root mean squared error (RMSE): $RMSE = \sqrt{MSE}$
3. Mean absolute deviation (MAD): $MAD = \frac{1}{m_{te}} \sum_{i=1}^{m_{te}} |\hat{\eta}(\mathbf{x}_i) - \eta(\mathbf{x}_i)|$
4. Mean deviation (MD): $MD = \frac{1}{m_{te}} \sum_{i=1}^{m_{te}} (\hat{\eta}(\mathbf{x}_i) - \eta(\mathbf{x}_i))$
5. Averaged Kullback-Leibler divergence (KL): Since η is a density, we can compute the KL between its estimate and true value as follows:

$$\begin{aligned} KL &= \frac{1}{m} \sum_{i=1}^{m_{te}} KL[\eta(\mathbf{x}_i) \|\hat{\eta}(\mathbf{x}_i)] \\ &= \frac{1}{m} \sum_{i=1}^{m_{te}} \eta(\mathbf{x}_i) \log \frac{\eta(\mathbf{x}_i)}{\hat{\eta}(\mathbf{x}_i)} + (1 - \eta(\mathbf{x}_i)) \log \frac{1 - \eta(\mathbf{x}_i)}{1 - \hat{\eta}(\mathbf{x}_i)} \end{aligned}$$

6. Accuracy (Acc): $Acc = \frac{1}{m_{te}} \sum_{i=1}^{m_{te}} \mathbf{1}_{[(2\hat{\eta}(\mathbf{x}_i)-1)y_i \leq 0]}$
7. Brier Score (BS): Brier score is a proper score function which measures the accuracy of probabilistic predictions and is defined below:

$$BS = \frac{1}{m_{te}} \sum_{i=1}^{m_{te}} (\hat{\eta}(\mathbf{x}_i) - y_i)^2$$

8. DiffMax: In some applications, one requires how good the estimates are at the boundary. In such cases, we have to compare the maximum value of true η to the maximum value of $\hat{\eta}$ on the training set. This measure plays an important role when one is interested in estimating noise rates as some noise estimating schemes use corrupted η values [9]. Therefore,

$$DiffMax := \max_{i=1, \dots, m_{tr}} \hat{\eta}(\mathbf{x}_i) - \max_{i=1, \dots, m_{tr}} \eta(\mathbf{x}_i)$$

9. DiffMin: Similar to DiffMax, this measure is also a deciding factor for quality of noise rates.

$$DiffMin := \min_{i=1, \dots, m_{tr}} \hat{\eta}(\mathbf{x}_i) - \min_{i=1, \dots, m_{tr}} \eta(\mathbf{x}_i)$$

The synthetic dataset is partitioned into training and test set with 80-20 split. The first 7 performance measures are computed on the test set and the last two on the training set. To make the comparisons more robust, this process is repeated for 10 trials. Finally, the reported values are performance measures averaged over the 10 trials along with the standards deviation.

Synthetic dataset 1 (Syn_{2D}) This is a 2 dimensional dataset. We first generate 1000 binary class labels Y from Bernoulli distribution with parameter $p = 0.5$. Then a 2-dimensional feature vector \mathbf{X} for each label is drawn from two different Gaussian distributions: $\mathbf{X}|Y = 1 \sim N([1.2, 0], \Sigma)$ & $\mathbf{X}|Y = -1 \sim N([-1.2, 0], \Sigma)$ where $\Sigma = [[1, 0.4]; [0.4, 1]]$. Table 16, 17 and 18 present the values of various measures for this dataset.

Performance measures	η estimated from			LSPC	k-NN
	Algorithm ($f, \hat{\eta}$) with l_{sq}	Algorithm ($f, \hat{\eta}$) with l_{log}	Algorithm ($f, \hat{\eta}$) with l_{msq}		
MSE	0.02±0.002	0.0008±0.0007	0.001±0.0003	0.003±0.002	0.013 ± 0.003
RMSE	0.144±0.007	0.025±0.012	0.034±0.004	0.054±0.02	0.114 ± 0.014
MAD	0.123±0.005	0.018±0.010	0.025±0.002	0.037±0.015	0.066 ± 0.008
MD	-0.0032±0.011	-0.0076±0.001	-0.007±0.002	-0.007±0.004	-0.007 ± 0.005
KL	NA	0.004±0.004	NA	NA	NA
Acc	0.911±0.016	0.91±0.017	0.91±0.018	0.91±0.01	0.906 ± 0.016
BS	0.0804±0.004	0.0651±0.010	0.0657±0.007	0.06809±0.0068	0.077 ± 0.010
DiffMax	0.531±0.04	0.0 ±0.0	0.0 ±0.0	0.0 ±0.0	0.0 ±0.0
DiffMin	-0.497±0.042	0.0 ±0.0	0.0 ±0.0	0.0 ±0.0	0.0 ±0.0

Table 16. Performance of estimates obtained by Algorithm ($f, \hat{\eta}$) with squared loss, logistic loss and modified squared loss and LSPC and k -NN on Syn_{2D} dataset. The cells with NA means that some estimate was 0 and the KL couldn't be defined. Cells with 0.0 entry means that the value is less than 1×10^{-5} .

Synthetic dataset 2 (Syn_{3D}) This is a 3 dimensional dataset. We first generate 1000 binary class labels Y from Bernoulli distribution with parameter

	KLIEP based estimator used		
Performance measures	$\hat{\eta}_{pos}$	$\hat{\eta}_{neg}$	$\hat{\eta}_{norm}$
MSE	0.012± 0.003	0.014±0.007	0.001± 0.00
RMSE	0.108± 0.015	0.118±0.027	0.034±0.009
MAD	0.062± 0.006	0.069±0.016	0.021±0.005
MD	-0.006±0.008	-0.006±0.013	-0.003±0.005
KL	NA	NA	0.005± 0.002
Acc	0.89 ± 0.02	0.90 ± 0.01	0.91 ± 0.02
BS	0.075 ± 0.011	0.074± 0.006	0.065±0.006
DiffMax	0.16±0.02	0.0 ±0.0	0.0 ±0.0
DiffMin	0.0 ±0.0	-0.16 ±0.06	0.0 ±0.0

Table 17. Performance of $\hat{\eta}_{pos}$, $\hat{\eta}_{neg}$, $\hat{\eta}_{norm}$, KLIEP based estimates on Syn_{2D} dataset. The cells with NA means that some estimate was 0 and the KL couldn't be defined. Cells with 0.0 entry means that the value is less than 1×10^{-5} .

	KLIEP based estimator used				
Performance measures	$\hat{\eta}_{pos,s}$	$\hat{\eta}_{neg,s}$	$\hat{\eta}_{norm,s}$	$\hat{\eta}_{Cnorm,s}$	$\hat{\eta}_{avg,s}$
MSE	0.009±0.004	0.019±0.01	0.001±0.00	0.178±0.004	0.009±0.002
RMSE	0.096±0.019	0.133±0.035	0.036±0.009	0.133±0.013	0.097±0.01
MAD	0.095±0.008	0.072±0.027	0.021±0.005	0.105±0.014	0.072±0.01
MD	-0.08±0.008	0.06±0.027	-0.003±0.007	-0.019±0.018	-0.012±0.014
KL	0.073 ±0.002	NA	0.005 ± 0.002	NA	0.047±0.008
Acc	0.87 ± 0.02	0.89 ± 0.02	0.90 ± 0.01	0.90 ±0.01	0.90 ± 0.02
BS	0.086±0.010	0.084±0.009	0.066±0.006	0.083±0.006	0.0748±0.005
DiffMax	0.0 ±0.0	0.0 ±0.0	0.0 ±0.0	0.0 ±0.0	-0.003±0.00
DiffMin	0.0 ±0.0	0.0 ±0.0	0.0 ±0.0	0.0 ±0.0	-0.005±0.004

Table 18. Performance of $\hat{\eta}_{pos,s}$, $\hat{\eta}_{neg,s}$, $\hat{\eta}_{norm,s}$, $\hat{\eta}_{Cnorm,s}$, $\hat{\eta}_{avg,s}$, KLIEP based estimates on Syn_{2D} dataset. The cells with NA means that some estimate was 0 and the KL couldn't be defined. Cells with 0.0 entry means that the value is less than 1×10^{-5} .

$p = 0.5$. Then a 3-dimensional feature vector \mathbf{X} for each label is drawn from two different Gaussian distributions: $\mathbf{X}|Y = 1 \sim N([1.3, 0, 0], \Sigma)$ & $\mathbf{X}|Y = -1 \sim N([-1.3, 0, 0], \Sigma)$ where $\Sigma = [[1, 0.1, 0.1]; [0.1, 1, 0.1]; [0.1, 0.1, 1]]$. Table 19, 20 and 21 present the values of various measures for this dataset.

Performance measures	η estimated from			LSPC	k-NN
	Algorithm ($f, \hat{\eta}$) with l_{sq}	Algorithm ($f, \hat{\eta}$) with l_{log}	Algorithm ($f, \hat{\eta}$) with l_{msq}		
MSE	0.02±0.001	0.0008±0.0007	0.001±0.0003	0.003±0.002	0.0108 ± 0.0022
RMSE	0.145±0.007	0.026±0.007	0.035±0.003	0.064±0.003	0.103 ± 0.0102
MAD	0.126±0.003	0.018±0.007	0.025±0.002	0.045±0.015	0.0612 ± 0.005
MD	-0.016±0.008	-0.0078±0.002	-0.006±0.0036	-0.009±0.005	-0.008 ± 0.007
KL	NA	0.004±0.003	NA	NA	NA
Acc	0.92±0.014	0.923±0.015	0.924±0.014	0.923±0.016	0.929 ± 0.0197
BS	0.0820±0.006	0.0638±0.008	0.0663±0.008	0.0692±0.008	0.0667 ± 0.0127
DiffMax	0.628±0.01	0.0 ±0.0	0.0 ±0.0	0.0 ±0.0	0.0 ±0.0
DiffMin	-0.458±0.01	0.0 ±0.0	0.0 ±0.0	0.0 ±0.0	0.0 ±0.0

Table 19. Performance of estimates obtained by Algorithm ($f, \hat{\eta}$) with squared loss, logistic loss and modified squared loss and LSPC and k -NN on Syn_{3D} dataset. The cells with NA means that some estimate was 0 and the KL couldn't be defined. Cells with 0.0 entry means that the value is less than 1×10^{-5} .

Performance measures	KLIEP based estimator used		
	$\hat{\eta}_{pos}$	$\hat{\eta}_{neg}$	$\hat{\eta}_{norm}$
MSE	0.022± 0.003	0.025±0.005	0.002± 0.00
RMSE	0.146± 0.010	0.158±0.015	0.050±0.005
MAD	0.097± 0.007	0.103±0.007	0.040±0.005
MD	-0.0004±0.009	-0.009±0.019	-0.006±0.004
KL	NA	NA	0.014± 0.003
Acc	0.90 ± 0.02	0.89 ± 0.0	0.92 ± 0.016
BS	0.0793 ± 0.0077	0.084± 0.010	0.067±0.007
DiffMax	0.34±0.02	0.0 ±0.0	0.0 ±0.0
DiffMin	0.0002 ±0.0	-0.344 ±0.02	0.006 ±0.002

Table 20. Performance of $\hat{\eta}_{pos}, \hat{\eta}_{neg}, \hat{\eta}_{norm}$, KLIEP based estimates on Syn_{3D} dataset. The cells with NA means that some estimate was 0 and the KL couldn't be defined. Cells with 0.0 entry means that the value is less than 1×10^{-5} .

Synthetic dataset 3 (Syn_{10D}) This is a 10 dimensional dataset. We first generate 1000 binary class labels Y from Bernoulli distribution with parameter $p = 0.5$. Then, a 10-dimensional feature vector \mathbf{X} for each label is drawn from two different Gaussian distributions: $\mathbf{X}|Y = 1 \sim N(\boldsymbol{\mu}_+, \Sigma)$ & $\mathbf{X}|Y = -1 \sim N(\boldsymbol{\mu}_-, \Sigma)$ where Σ is such that all the diagonal elements (variances) are 1 and the non-diagonal elements (covariances) are 10. Here, the means are $\boldsymbol{\mu}_+ =$

KLIEP based estimator used					
Performance measures	$\hat{\eta}_{pos,s}$	$\hat{\eta}_{neg,s}$	$\hat{\eta}_{norm,s}$	$\hat{\eta}_{Cnorm,s}$	$\hat{\eta}_{avg,s}$
MSE	0.014±0.004	0.044±0.001	0.03±0.005	0.041±0.003	0.024±0.002
RMSE	0.118±0.014	0.209±0.021	0.051±0.005	0.203±0.01	0.154±0.009
MAD	0.164±0.008	0.141±0.014	0.042±0.005	0.183±0.010	0.136±0.009
MD	-0.151±0.004	0.122±0.0151	-0.008±0.015	-0.018±0.009	-0.014±0.008
KL	0.150 ± 0.011	NA	0.015 ± 0.003	NA	0.102±0.009
Acc	0.8535 ± 0.01	0.83 ± 0.02	0.92 ± 0.01	0.92 ± 0.01	0.92 ± 0.02
BS	0.1147±0.006	0.1144±0.011	0.0677±0.0.007	0.108±0.005	0.0905±0.005
DiffMax	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	-0.006±0.02
DiffMin	0.00 ± 0.0	0.0 ± 0.0	0.0006 ± 0.0	0.0 ± 0.0	-0.008±0.002

Table 21. Performance of $\hat{\eta}_{pos,s}, \hat{\eta}_{neg,s}, \hat{\eta}_{norm,s}, \hat{\eta}_{Cnorm,s}, \hat{\eta}_{avg,s}$, KLIEP based estimates on Syn_{3D} dataset. The cells with NA means that some estimate was 0 and the KL couldn't be defined. Cells with 0.0 entry means that the value is less than 1×10^{-5} .

[2, 1.2, 2, 2.1, 2, 2, 2, 0.2, 2, 3.6] and $\mu_+ = [2, -1.2, 2, -2.1, 2, 2, 2, -0.2, 2, -3.6]$. Table 22, 23 and 24 present the values of various measures for this dataset.

η estimated from					
Performance measures	Algorithm ($f, \hat{\eta}$) with l_{sq}	Algorithm ($f, \hat{\eta}$) with l_{log}	Algorithm ($f, \hat{\eta}$) with l_{msq}	LSPC	k-NN
MSE	0.026±0.002	0.002±0.0008	0.003±0.0008	0.011±0.002	0.0232 ± 0.002
RMSE	0.161±0.006	0.047±0.009	0.005±0.007	0.106±0.010	0.152 ± 0.007
MAD	0.138±0.005	0.026±0.005	0.031±0.004	0.077±0.011	0.0911 ± 0.00756
MD	0.004±0.006	0.002±0.006	0.003±0.006	0.003±0.003	0.01434 ± 0.005
KL	NA	0.010±0.003	NA	NA	NA
Acc	0.912±0.018	0.911±0.016	0.908±0.016	0.913±0.017	0.9195 ± 0.0181
BS	0.0795±0.008	0.0645±0.012	0.0636±0.012	0.0733±0.0102	0.0834 ± 0.011
DiffMax	0.463±0.03	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0
DiffMin	-0.551±0.024	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0

Table 22. Performance of estimates obtained by Algorithm ($f, \hat{\eta}$) with squared loss, logistic loss and modified squared loss and LSPC and k -NN on Syn_{10D} dataset. The cells with NA means that some estimate was 0 and the KL couldn't be defined. Cells with 0.0 entry means that the value is less than 1×10^{-5} .

Observations We observe that as far as label prediction is concerned Algorithm ($f, \hat{\eta}$) with squared loss, logistic loss and modified squared loss, LSPC, k -NN and KLIEP's $\hat{\eta}_{norm}$ estimators perform equally well as their accuracy is the highest in comparison to others. This trend has been observed across all 3 synthetic datasets.

With respect to DiffMax and DiffMin measures, Algorithm ($f, \hat{\eta}$) with logistic loss and modified squared loss, LSPC, k -NN and KLIEP's $\hat{\eta}_{norm}$ estimators are good.

Performance measures	KLIEP based estimator used		
	$\hat{\eta}_{\text{pos}}$	$\hat{\eta}_{\text{neg}}$	$\hat{\eta}_{\text{norm}}$
MSE	0.06± 0.003	0.062±0.006	0.035± 0.005
RMSE	0.250± 0.006	0.249±0.012	0.189±0.014
MAD	0.208± 0.007	0.207±0.013	0.177±0.014
MD	0.0031±0.01	0.006±0.012	0.008±0.005
KL	NA	NA	0.014± 0.017
Acc	0.848 ± 0.03	0.846 ± 0.017	
BS	0.125 ± 0.010	0.123± 0.009	0.0995±0.006
DiffMax	0.32±0.098	-0.02 ±0.006	-0.04 ±0.001
DiffMin	0.015 ±0.005	-0.422 ±0.12	0.031 ±0.009

Table 23. Performance of $\hat{\eta}_{\text{pos}}, \hat{\eta}_{\text{neg}}, \hat{\eta}_{\text{norm}}$, KLIEP based estimates on Syn_{10D} dataset. The cells with NA means that some estimate was 0 and the KL couldn't be defined. Cells with 0.0 entry means that the value is less than 1×10^{-5} .

Performance measures	KLIEP based estimator used				
	$\hat{\eta}_{\text{pos},s}$	$\hat{\eta}_{\text{neg},s}$	$\hat{\eta}_{\text{norm},s}$	$\hat{\eta}_{C\text{norm},s}$	$\hat{\eta}_{\text{avg},s}$
MSE	0.035±0.005	0.061±0.003	0.104±0.014	0.036±0.005	0.084±0.006
RMSE	0.247±0.006	0.321±0.022	0.191±0.013	0.290±0.01	0.260±0.006
MAD	0.245±0.009	0.263±0.009	0.178±0.014	0.275±0.011	0.247±0.007
MD	-0.181±0.03	0.157±0.035	0.019±0.007	0.018±0.012	-0.019±0.010
KL	0.262 ±0.022	NA	0.143 ± 0.016	NA	0.228±0.009
Acc	0.750 ± 0.044	0.736 ± 0.044	0.913 ± 0.016	0.913 ± 0.016	0.913 ± 0.016
BS	0.158±0.0159	0.165±0.0158	0.099±0.006	0.149±0.008	0.1318±0.006
DiffMax	0.0 ±0.0	-0.01 ±0.006	-0.040 ±0.012	0.0 ±0.0	-0.08±0.018
DiffMin	0.011 ±0.005	0.0 ±0.0	0.033 ±0.009	0.0 ±0.0	0.081±0.02

Table 24. Performance of $\hat{\eta}_{\text{pos},s}, \hat{\eta}_{\text{neg},s}, \hat{\eta}_{\text{norm},s}, \hat{\eta}_{C\text{norm},s}, \hat{\eta}_{\text{avg},s}$, KLIEP based estimates on Syn_{10D} dataset. The cells with NA means that some estimate was 0 and the KL couldn't be defined. Cells with 0.0 entry means that the value is less than 1×10^{-5} .

G Risk minimization under weighted 0-1 loss in the extended cost space

In view of the negative result obtained with the α -weighted 0-1 loss function, we explore a larger cost space. In a similar kind of representation as that of C_1, C_{-1} , one can extend the parameter space from (C_1, C_{-1}) to $(C_1, C_{-1}, c_1, c_{-1})$ given as follows.

- C_1 (or C_{-1}): loss when a positive (or negative) class point is misclassified
- c_1 (or c_{-1}): loss when a positive (or negative) class point is correctly classified

Let the extended cost sensitive 0-1 loss be defined as

$$l_{c,4}(f(\mathbf{x}), y) = \mathbf{1}_{\{f(\mathbf{x}) > 0, y=1\}} c_1 + \mathbf{1}_{\{f(\mathbf{x}) > 0, y=-1\}} C_{-1} + \mathbf{1}_{\{f(\mathbf{x}) \leq 0, y=1\}} C_1 + \mathbf{1}_{\{f(\mathbf{x}) \leq 0, y=-1\}} c_{-1} \quad (50)$$

Then, the corresponding risk on D and \tilde{D} be denoted by $R_{D,l_{c,4}}(f)$ and $R_{\tilde{D},l_{c,4}}(f)$ respectively.

One can reduce the extended $(C_1, c_1, C_{-1}, c_{-1})$ based 4 cost space again to $(C_1 - c_1, C_{-1} - c_{-1})$ based 2 cost space as the threshold of the optimal classifiers does not change by reducing the parameter space. This has also been shown in [8]. However, we still compute the minimizers of the risks $R_{D,l_{c,4}}(f)$ and $R_{\tilde{D},l_{c,4}}(f)$ and show that they need not be equal.

Using the same technique as in (C_1, C_{-1}) cost case, one can show that the minimizers of $R_{D,l_{c,4}}$ and $R_{\tilde{D},l_{c,4}}$ are as follows:

$$f_{4c}^*(\mathbf{x}) = \text{sign} \left(\eta(\mathbf{x}) - \frac{C_{-1} - c_{-1}}{C_1 - c_1 + C_{-1} - c_{-1}} \right) \quad (51)$$

$$\tilde{f}_{4c}^*(\mathbf{x}) = \text{sign} \left(\eta(\mathbf{x}) - \frac{C_{-1} - c_{-1} - \rho(C_1 - c_1 + C_{-1} - c_{-1})}{(1 - 2\rho)(C_1 - c_1 + C_{-1} - c_{-1})} \right) \quad (52)$$

The threshold of η in equation (51) and (52) cannot be same unless $C_1 - c_1 = s = C_{-1} - c_{-1}$ or $\rho = 0$. Therefore, even in the extended cost space, under uniform noise and differential cost, $f_{4c}^* \neq \tilde{f}_{4c}^*$, i.e., learning a classifier from $l_{c,4}$ which is both noise robust and cost sensitive is not always possible. This is independent of the extended sufficient condition (E-SC) that we now introduce; as we will see, this condition doesn't help us for joint uniform noise robustness and cost sensitivity.

A natural extension of sufficient condition (Symmetry condition SC [6]) in cost sensitive case when used with weighted 0-1 loss can be written as follows. For an arbitrary classifier $f(\mathbf{x})$ and $K > 0$,

$$\begin{aligned} \mathbf{E}\text{-SC} \quad & \text{when } f(\mathbf{x}) > 0 \text{ SC becomes } C_{-1} + c_1 = K \\ & \text{when } f(\mathbf{x}) \leq 0 \text{ SC becomes } C_1 + c_{-1} = K \end{aligned} \quad (53)$$

Clearly, $C_{-1} + c_1 = K = C_1 + c_{-1}$ which can be equivalently written as $C_{-1} - c_{-1} = s = C_1 - c_1$. This implies that the effective cost for each class is same, i.e. s and hence no differential costs for the classes. Therefore, E-SC cannot be sufficient condition for SLN robustness if there is differential costing of $(C_1, c_1, C_{-1}, c_{-1})$.