Estimating the accuracies of multiple classifiers without labeled data

Ariel Jaffe
Weizmann Institute of Science

Boaz Nadler
Weizmann Institute of Science

Yuval Kluger
Yale University School of Medicine

Abstract

In various situations one is given only the predictions of multiple classifiers over a large unlabeled test data. This scenario raises the following questions: Without any labeled data and without any a-priori knowledge about the reliability of these different classifiers, is it possible to consistently and computationally efficiently estimate their accuracies? Furthermore, also in a completely unsupervised manner, can one construct a more accurate unsupervised ensemble classifier? In this paper, focusing on the binary case, we present simple, computationally efficient algorithms to solve these questions. Furthermore, under standard classifier independence assumptions, we prove our methods are consistent and study their asymptotic error. Our approach is spectral, based on the fact that the off-diagonal entries of the classifiers' covariance matrix and 3-d tensor are rank-one. We illustrate the competitive performance of our algorithms via extensive experiments on both artificial and real datasets.

1 Introduction

Consider a classification problem from an instance space $\mathcal{X}$ to an output label set $\mathcal{Y} = \{1, \ldots, K\}$. In contrast to the classical supervised setting, in various contemporary applications, one has access only to the predictions of multiple experts or classifiers over a large number of unlabeled instances. Moreover, the reliability of these experts may be unknown, and at test time there is no labeled data to assess it. This occurs for example when due to privacy considerations each classifier is trained with its own possibly proprietary labeled data, unavailable to us. Another scenario is crowdsourcing, where an annotation task over many instances is distributed to many annotators whose reliability is a-priori unknown, see for example Welinder et al. [2010], Whitehill et al. [2009], Sheshadri and Lease [2013]. This setup, denoted as unsupervised-supervised learning in Donmez et al. [2010], appears in several other application domains, including decision science, economics and medicine, see Snow et al. [2008], Raykar et al. [2010], Parisi et al. [2014].

Given only the $m \times n$ matrix Z, or a significant part of it, with $Z_{ij} = f_i(x_j)$ holding the predictions of the given m classifiers over n instances, and without any labeled data, two fundamental questions arise: (i) Under the assumption that different classifiers make independent errors, is it possible to consistently estimate the accuracies of the m classifiers in a computationally efficient way; and (ii) is it possible to construct, again by some computationally efficient procedure, an unsupervised ensemble learner, more accurate than most if not all of the original m classifiers.

The first question is important in cases where obtaining the predictions of these m classifiers is by itself an expensive task, and after collecting a certain number of instances and their predictions, we wish to pick only a few of the most accurate ones, see Rokach [2009]. The second question, also known as *offline consensus*, is of utmost importance in improving the quality of automatic decision making systems based on multiple sources of information.

Beyond the simplest approach of majority voting, perhaps the first to define and address these questions were Dawid and Skene [1979]. With the increasing popularity of crowdsourcing and large scale expert opinion systems, the last years have seen a surge of interest in these problems, see Sheng et al. [2008], Whitehill et al. [2009], Raykar et al. [2010], Platanios et al. [2014] and references therein. Yet, the most common methods to address questions (i) and (ii) above are based on the expectation maximization (EM) algorithm, already proposed in this context by Dawid and Skene, and whose only guarantee is convergence to a local maxima.

Appearing in Proceedings of the 18^{th} International Conference on Artificial Intelligence and Statistics (AISTATS) 2015, San Diego, CA, USA. JMLR: W&CP volume 38.

Two recent exceptions, proposing spectral (and thus computationally efficient) methods with strong consistency guarantees are Karger et al. [2011] and Parisi et al. [2014]. Karger et al. [2011] assume a spammer-hammer model, where each classifier is either perfectly correct or totally random and develop a spectral method to detect which one is which. Parisi et al. [2014] derive a spectral approach to address questions (i) and (ii) above in the context of binary classification. Their approach, however, has several limitations. First, they do not actually estimate each classifier sensitivity and specificity, but only show how to consistently rank them according to their balanced accuracies. Second, their unsupervised learner assumes that all classifiers have balanced accuracies close to 1/2 (random). Hence, their ensemble learner may be suboptimal, for example, when few classifiers are significantly more accurate than all others.

In this paper we extend and generalize the results of Parisi et al. [2014] in several directions and make the following contributions: In Section 3, focusing on the binary case, we present a simple spectral method to estimate the sensitivity and specificity of each classifier, assuming the class imbalance is known. Hence, the problem boils down to estimating a single-scalar – the class imbalance. In Section 4 we present two different methods to do so. In Sec. 4.1, we prove that the off-diagonal elements of the $m \times m$ covariance matrix and the $m \times m \times m$ joint covariance tensor of the set of classifiers are both rank 1. Moreover the covariance matrix and tensor share the same eigenvector but with different eigenvalues, from which the class imbalance can be extracted by a simple least-squares procedure. In Sec. 4.2, we devise a second algorithm to estimate the class imbalance by a restricted likelihood approach, that requires optimizing a function of a single variable. Both algorithms are computationally efficient, and under the assumption that classifiers make independent errors, are also proven to be consistent. For the first method, we also prove it is rate optimal with asymptotic error $\mathcal{O}_P(1/\sqrt{n})$, where n is the number of unlabeled samples. Our work thus provides a simple and elegant solution to the long-standing problem originally posed by Dawid and Skene [2], whose previous solutions were mostly based on expectation maximization approaches to the full likelihood function.

In Sec. 5 we consider the multiclass case. Building upon standard reductions from multiclass to binary, we devise a method to estimate the class probabilities and the diagonal entries of the confusion matrices of all classifiers. We also prove that in the multiclass case, using only the first and second moments of these binary reductions, it is in general not possible to estimate all entries of the confusion matrices of all classifiers. This motivates the development of tensor or higher order methods to solve the multi-class case, as for example in Zhang et al. [2014]. In Sec. 6 we illustrate our methods on both real and artificial data. The results on real data show that our proposed ensemble learner achieves a competitive performance even in practical scenarios where the assumption of independent classifier errors does not hold precisely.

Related Work Under the assumption that all classifiers make independent errors, our problem is equivalent to learning a mixture of discrete product distributions. This problem was studied, among others, by Freund and Mansour [1999] for the case of $k = 2$ distributions, and by Feldman et al. [2008] for $k > 2$. Important observations regarding the low-rank spectral structure of the second and third moments of such distributions were made by Anandkumar et al. [2014, 2012]. Building upon these results, recently Jain and Oh [2014] and Zhang et al. [2014], devised computationally efficient algorithms to estimate the parameters of the mixture of product distributions, which are equivalent to the confusion matrices and class probabilities in our problem.

Our first method to estimate the class imbalance in the binary case using the mean-centered 3-d tensor is closely related to these works, with some notable differences. One key difference is that the above works study non-centered tensors of classifiers' outputs, and hence for a k-class problem, need to resolve the structure of rank-k tensors. In contrast, we work with centered matrices and tensors. In the binary case with $k = 2$, we thus obtain a simpler rank-1 tensor, which we do not even need to decompose, but only extract a single scalar from it. A second difference is that the above methods require stronger assumptions on the classifiers. For example, Zhang et al. [2014] divide the classifiers into groups and assume that within each group, on average classifiers are better than random. Due to these differences, our resulting algorithm is significantly simpler.

Our second algorithm for estimating the class imbalance, based on a restricted likelihood approach is totally different from these tensor-based works, as it requires only a spectral decomposition of the classifiers' covariance matrix, and then optimizes a 1-d function of the full likelihood of the data. On both simulated and real data, this second approach had at least as good as, and in some cases better accuracy compared to the tensor based method. Finally, while we focus on classification, our algorithms may also be of interest to learning a mixture of discrete product distributions.

2 Problem Setup

We consider the following binary classification problem, as also studied in several works (Dawid and Skene [1979], Raykar et al. [2010], Parisi et al. [2014]). Let $\mathcal{X}$ be an instance space with an output space $\mathcal{Y} = \{-1, 1\}$. A labeled instance $(x, y) \in \mathcal{X} \times \mathcal{Y}$ is a realization of the random variable (X, Y), which has an unknown probability density $p(x, y)$, and X and Y marginals $p_X(x)$ and $p_Y(y)$, respectively. We further denote by b the class imbalance of Y,

$$b = \Pr(Y = 1) - \Pr(Y = -1) = p_Y(1) - p_Y(-1).$$

Let $\{f_i\}_{i=1}^m$ be $m \geq 3$ classifiers operating on $\mathcal{X}$. In this binary setting, the accuracy of the i-th classifier is fully specified by its sensitivity ψ_i and specificity η_i,

$$\psi_i = \Pr\left(f_i(X)=1|Y=1\right), \; \eta_i = \Pr\left(f_i(X)=-1|Y=-1\right)$$

For future use, we denote by π_i its balanced accuracy,

$$\pi_i = (\psi_i + \eta_i)/2.$$

In this paper we consider the following totally *unsupervised* scenario. Let Z be a $m \times n$ matrix with entries $Z_{ij} = f_i(x_j), i = 1, \ldots, m, j = 1, \ldots, n$, where $f_i(x_j)$ is the label predicted at instance x_j by classifier f_i. In particular, we assume no prior knowledge about the m classifiers, so their accuracies (sensitivities ψ_i and specificities η_i) are all unknown.

Given only the matrix Z of binary predictions[1], we consider the following two problems: (i) consistently and computationally efficiently estimate the sensitivity and specificity of each classifier, and (ii) construct a more accurate ensemble classifier. As discussed below, under certain assumptions, a solution to the first problem readily yields a solution to the second one.

To tackle these problems, we make the following three assumptions: (i) The n instances x_j are i.i.d. realizations from the marginal $p_X(x)$. (ii) The m classifiers are conditionally independent. That is, for all f_i, f_j with $i \neq j$ and for all labels $a_i, a_j \in \{-1, 1\}$,

$$\Pr(f_i = a_i, f_j = a_j|Y = y) = \Pr(f_i = a_i|Y = y) \Pr(f_j = a_j|Y = y). \quad (1)$$

(iii) Most classifiers are better than random, in the sense that for more than half of all classifiers, $\pi_i > 0.5$. Note that (i)-(ii) are standard assumptions in both the supervised and unsupervised settings, see Dietterich [2000], Dawid and Skene [1979], Raykar et al. [2010], Parisi et al. [2014]. Assumption (iii) or a variant thereof is needed, given an inherent ± 1 sign ambiguity in this fully unsupervised problem.

[1]For simplicity of exposition, we assume the matrix is fully observed. While beyond the scope of this paper, our proposed methods and theory continue to hold if few entries are missing (at random), such that accurate estimates of various means, covariances and tensors, as detailed in Sections 3-4 are still possible.

3 Estimating ψ and η with a known class imbalance.

For some classification problems, the class imbalance b is known. One example is in epidemiology, where the overall prevalence of a certain disease in the population is known, and the classification problem is to predict the presence or future onset of the disease in individuals given their observed features (such as blood results, height, weight, age, genetic profile, etc).

Assuming b is known, Donmez et al. [2010] presented a simple method to estimate the error rates of all classifiers under a symmetric noise model, where $\psi_i = \eta_i$ for all i. They further proposed EM methods in the general case, see also Raykar et al. [2010]. We instead build upon the spectral approach in Parisi et al. [2014], and present a computationally efficient method to consistently estimate the sensitivities and specificities of all m classifiers. To motivate our approach, it is instructive to study the limit of an infinite unlabeled set size, $n \to \infty$, where the mean values of the classifiers $\mu_i = \mathbb{E}[f_i(X)]$, and their $m \times m$ population covariance matrix $R = \mathbb{E}\left[(f_i(x) - \mu_i)(f_j(x) - \mu_j)\right]$, are all perfectly known.

The following two lemmas show that knowing R and $\{\mu_i\}_{i=1}^m$ suffice to extract the specificities and sensitivities of the m classifiers. Lemma 1 appeared in Parisi et al. [2014], and implies that given the value of b one may compute the balanced accuracies of all classifiers. Lemma 2, proven in the appendix, is new and shows how to extract their sensitivities and specificities.

Lemma 1. *Under assumptions (i)-(iii) of Section 2, the off diagonal elements of the matrix R are identical to those of a rank one matrix $\mathbf{v}\mathbf{v}^T$, whose vector $\mathbf{v}$, up to a ± 1 sign ambiguity, is equal to*

$$\mathbf{v} = \sqrt{1 - b^2}(2\boldsymbol{\pi} - 1), \quad (2)$$

where the vector $\boldsymbol{\pi} = (\pi_1, \ldots, \pi_m)$ contains the balanced accuracies of the m classifiers.

Lemma 2. *Given the class imbalance b, the mean values $\boldsymbol{\mu} = (\mu_1, \ldots, \mu_m)$ of the m classifiers and the vector $\mathbf{v}$ of Eq. (2), the values of $\boldsymbol{\psi} = (\psi_1, \ldots, \psi_m)$ and $\boldsymbol{\eta} = (\eta_1, \ldots, \eta_m)$ with the specificities and sensitivities of the m classifiers are given by*

$$\boldsymbol{\psi} = \tfrac{1}{2}\left(1 + \boldsymbol{\mu} + \mathbf{v}\sqrt{\tfrac{1-b}{1+b}}\right), \boldsymbol{\eta} = \tfrac{1}{2}\left(1 - \boldsymbol{\mu} + \mathbf{v}\sqrt{\tfrac{1+b}{1-b}}\right). \quad (3)$$

To uniquely recover $\mathbf{v}$ from the off-diagonal entries of R, we further assume that at least three classifiers have balanced accuracies different from 1/2 (so $2\pi_i - 1 \neq 0$). In practice, the quantities $\{\mu_i\}_{i=1}^m$, R and consequently the eigenvector $\mathbf{v}$ are all unknown. We thus estimate them from the given data, and plug into Eq. (3). Let us denote by $\hat{\boldsymbol{\mu}}$ and $\hat{R}$ the sample mean and covariance matrix of all classifiers, whose entries are given by

$$\begin{aligned} \hat{\mu}_i &= \frac{1}{n}\sum_{k=1}^{n} f_i(x_k), \qquad (4) \\ \hat{r}_{ij} &= \frac{1}{n-1}\sum_{k=1}^{n}(f_i(x_k)-\hat{\mu}_i)(f_j(x_k)-\hat{\mu}_j). \end{aligned}$$

Estimating the vector $\mathbf{v}$ from the noisy matrix $\hat{R}$ can be cast as a low-rank matrix completion problem. Parisi et al. [2014] present several methods to construct such an estimate $\hat{\mathbf{v}}$, and resolve its inherent ± 1 sign ambiguity, via assumption (iii). Inserting $\hat{\boldsymbol{\mu}}$ and $\hat{\mathbf{v}}$ into (3), gives the following estimates for $\boldsymbol{\psi}$ and $\boldsymbol{\eta}$,

$$\hat{\boldsymbol{\psi}} = \tfrac{1}{2}\Big(1+\hat{\boldsymbol{\mu}}+\hat{\mathbf{v}}\sqrt{\tfrac{1-b}{1+b}}\Big), \hat{\boldsymbol{\eta}} = \tfrac{1}{2}\Big(1-\hat{\boldsymbol{\mu}}+\hat{\mathbf{v}}\sqrt{\tfrac{1+b}{1-b}}\Big). \quad (5)$$

The following lemma, proven in the appendix, presents some statistical properties of $\hat{\boldsymbol{\psi}}$ and $\hat{\boldsymbol{\eta}}$.

Lemma 3. *Under assumptions (i)-(iii) of Section 2, $\hat{\boldsymbol{\psi}}$ and $\hat{\boldsymbol{\eta}}$ are consistent estimators of $\boldsymbol{\psi}$ and $\boldsymbol{\eta}$. Furthermore, as $n \to \infty$,*

$$\hat{\psi}_i = \psi_i + \mathcal{O}_P\left(\frac{1}{\sqrt{n}}\right), \quad \hat{\eta}_i = \eta_i + \mathcal{O}_P\left(\frac{1}{\sqrt{n}}\right). \quad (6)$$

In summary, assuming the class imbalance b is known, Eq. (5) gives a computationally efficient way to estimate the sensitivities and specificities of all classifiers. Lemma 3 ensures that the resulting estimates are consistent. In the next section we show that the assumption of explicit knowledge of b can be removed, whereas in Section 5 we show that a similar approach can also (partly) handle the multiclass case.

3.1 Unsupervised Ensemble Learning

We now consider the second problem discussed in Section 2, the construction of an unsupervised ensemble learner. To this end, note that under the stronger assumption that all classifiers make independent errors, the likelihood of a label y at an instance x with predicted labels $f_1(x), \ldots, f_m(x)$ is

$$\mathcal{L}(f_1(x), \ldots, f_m(x)) \,|\, y) = \prod_{i=1}^{m} \Pr(f_i(x) \,|\, y). \quad (7)$$

In Eq. (7), the i-th term $\Pr(f_i(x)|y)$ depends on the specificity and sensitivity ψ_i and η_i of the i-th classifier. While the likelihood is non-convex in ψ_i, η_i and y, if the former are known, there is a closed form solution for the maximum-likelihood value of the class label,

$$\hat{y}^{(\mathrm{ML})} = \mathrm{sign}\left(\textstyle\sum_i f_i(x) \ln \alpha_i + \ln \beta_i\right) \quad (8)$$

where

$$\alpha_i = \frac{\psi_i \eta_i}{(1-\psi_i)(1-\eta_i)}, \quad \beta_i = \frac{\psi_i(1-\psi_i)}{\eta_i(1-\eta_i)}. \quad (9)$$

Parisi et al. [2014], assumed all classifiers are close to random, and via a Taylor expansion near $\psi = \eta = 1/2$, showed that β is approximately zero, and $\alpha_i \approx 1 + 4(2\pi_i - 1)$. Plugging these into Eq. (8), they derived the following spectral meta-learner (SML),

$$\hat{y}^{(\mathrm{SML})} = \mathrm{sign}\left(\textstyle\sum_i f_i(x)\hat{v}_i\right). \quad (10)$$

Their motivation was that they only had estimates of the vector $\mathbf{v}$, which according to Eq. (2) is proportional to $(2\boldsymbol{\pi} - 1)$. Since we consistently estimate the individual specificities and sensitivities of the m classifiers, we suggest to plug in these estimates directly into Eqs. (9) and (8). Our improved spectral approach, denoted i-SML, yields a more accurate ensemble learner when few classifiers are significantly better than random, so the linearization around $\psi = \eta = 1/2$ is inaccurate. We present such examples in Sec. 6. Finally, we note that as in Parisi et al. [2014] and Zhang et al. [2014], we may use our i-SML as a starting guess for EM methods that maximize the full likelihood.

4 Estimation of the class imbalance

We now consider the problem of estimating $\boldsymbol{\psi}$ and $\boldsymbol{\eta}$ when the class imbalance b is unknown. Our proposed approach is to first estimate b, and then plug this estimate into Eq. (5). We present two different methods to estimate the class imbalance. The first uses the covariance matrix and the 3-dimensional covariance tensor of all m classifiers. The second method exploits properties of the likelihood function. As detailed below, both methods are computationally efficient, but require stronger assumptions than Eq.(1) on independence of classifier errors to prove their consistency.

4.1 Estimation via the 3-D covariance tensor

For the method derived in this subsection, we assume that the classifiers are conditionally independent in triplets. That is, for every f_i, f_j, f_k with $i \neq j \neq k$ and for all labels $a_i, a_j, a_k \in \{-1, 1\}$,

$$\begin{aligned} &\Pr(f_i = a_i, f_j = a_j, f_k = a_k | y) = \\ &\quad \Pr(f_i = a_i|y)\Pr(f_j = a_j|y)\Pr(f_k = a_k|y). \quad (11) \end{aligned}$$

Let $T = (T_{ijk})$ denote the 3-dimensional covariance tensor of the m classifiers $\{f_i(X)\}_{i=1}^m$,

$$T_{ijk} = \mathbb{E}\left[(f_i(X) - \mu_i)(f_j(X) - \mu_j)(f_k(X) - \mu_k)\right]. \quad (12)$$

The following lemma, proven in the appendix, provides the relation between the tensor T, the class imbalance b and the balanced accuracies of the m classifiers.

Lemma 4. *Under assumption* (11)*, the following holds for all $i \neq j \neq k$,*

$$T_{ijk} = -2b(1-b^2)(2\pi_i - 1)(2\pi_j - 1)(2\pi_k - 1). \quad (13)$$

According to (13), the off diagonal elements of T (with $i \neq j \neq k$) correspond to a *rank one tensor*,

$$T = \mathbf{w} \otimes \mathbf{w} \otimes \mathbf{w}, \quad (14)$$

where $\otimes$ denotes the outer product and the vector $\mathbf{w} \in \mathbb{R}^m$ is equal to

$$\mathbf{w} = \left(-2b(1-b^2)\right)^{\frac{1}{3}} \cdot (2\boldsymbol{\pi} - 1). \quad (15)$$

Note that unlike the vector $\mathbf{v}$ of the covariance matrix R, there is no sign ambiguity in the vector $\mathbf{w}$. Moreover, comparing Eqs. (2) and (15), the vectors $\mathbf{v}$ of R and $\mathbf{w}$ of T are both proportional to $(2\boldsymbol{\pi} - 1)$, where the proportionality factor depends on the class imbalance b. Hence, $\mathbf{w} = \alpha(b)^{1/3}\,\mathbf{v}$, and

$$T = \alpha(b)\,\mathbf{v} \otimes \mathbf{v} \otimes \mathbf{v} \quad (16)$$

where $\alpha(b) = (-2b)/\sqrt{1-b^2}$. Inverting this expression yields the following relation,

$$b = -\alpha/\sqrt{4+\alpha^2}. \quad (17)$$

Eq. (17) thus shows, that in our setup, as $n \to \infty$, the first three moments of the data $(\boldsymbol{\mu}, R, T)$ are sufficient to determine both the class imbalance and the sensitivities and specificities of all m classifiers. In practice, the tensor T is unknown, though it can be estimated from the observed data by

$$\hat{T}_{ijk} = \frac{1}{n}\sum_{l=1}^{n}(f_i(x_l) - \hat{\mu}_i)(f_j(x_l) - \hat{\mu}_j)(f_k(x_l) - \hat{\mu}_k). \quad (18)$$

Given an estimate $\hat{\mathbf{v}}$ from the matrix $\hat{R}$, the scalar α of Eq. (16) is estimated by least squares,

$$\hat{\alpha} = \underset{\alpha}{\operatorname{argmin}} \sum_{i<j<k}\left(\hat{T}_{ijk} - \alpha\,\hat{v}_i\hat{v}_j\hat{v}_k\right)^2. \quad (19)$$

The steps to estimate the class imbalance with the 3 dimensional tensor appear in Algorithm 1. In terms of complexity the heaviest step is the estimation of T which requires $\mathcal{O}(m^3 n)$ operations. The following lemma shows that this method yields an asymptotic error of $\mathcal{O}_P(1/\sqrt{n})$. This error rate is optimal since it is equal to the rate achieved when estimating b with the ground truth labels y_i.

Algorithm 1 Estimating class imbalance with the 3dimensional covariance tensor

1: Estimate covariance matrix R by Eq. (5).
2: Estimate $\mathbf{v}$ from the off diagonal entries of $\hat{R}$ (see appendix).
3: Estimate the 3 dimensional tensor T by Eq. (18).
4: Estimate α via Eq. (19) and b via Eq. (17).

Lemma 5. *Let $\hat{\alpha}$ be given by Eq. (19) and let $\hat{b}_n$ be the plug-in estimator from Eq. (17). Then,*

$$\hat{b}_n = b + \mathcal{O}_P\left(1/\sqrt{n}\right). \quad (20)$$

Consequently the plug-in estimators $\hat{\psi}_i, \hat{\eta}_i$ in Eq. (5) also have the same asymptotic error $\mathcal{O}_P(1/\sqrt{n})$.

The proof of Lemma 5 appears in the appendix. Following it are some remarks regarding the accuracy of various estimates as a function of the number of classifiers and their accuracies. A detailed study of this issue is beyond the scope of this paper.

4.2 A restricted-likelihood approach

The algorithm in Section 4.1 relied only on the first three moments of the data. We now present a second method to estimate the class imbalance, based on a restricted likelihood function of all the data. This method is potentially more accurate, however it requires the following stronger assumption of joint conditional independence of all m classifiers,

$$\Pr(f_1 = a_1, \ldots, f_m = a_m | y) = \prod_{i=1}^{m} \Pr(f_i = a_i | y). \quad (21)$$

It is important to note that under this assumption, the problem at hand is equivalent to learning a mixture of two product distributions, addressed in Freund and Mansour [1999]. For this problem, several recent works suggested spectral tensor decomposition approaches, see Anandkumar et al. [2014], Jain and Oh [2014], Zhang et al. [2014].

In contrast, we now present a totally different approach, not based on tensor decompositions. Our starting point is Eq. (5) which provides consistent estimates of $\boldsymbol{\psi}$ and $\boldsymbol{\eta}$ given the class imbalance b. In particular, any guess $\tilde{b}$ of the class imbalance, yields corresponding guesses for the sensitivities and specificities of all m classifiers, $\hat{\boldsymbol{\psi}}(\tilde{b})$ and $\hat{\boldsymbol{\eta}}(\tilde{b})$. As described below, our approach is to construct a suitable functional $\hat{G}_n(Z|\tilde{b})$, that depends on both $\tilde{b}$ and on the observed data Z, whose maxima as a function of $\tilde{b}$, as $n \to \infty$ is attained at the true class imbalance b.

To this end, let $\mathbf{f}(x) = (f_1(x), \ldots, f_m(x))$ denote the vector of labels predicted by the m classifiers at an

instance x. We define the following approximate log-likelihood, assuming class imbalance $\tilde{b}$

$$\hat{g}_n(\mathbf{f}(x)|\tilde{b}) = \log \Pr\left(\mathbf{f}(x)|\hat{\boldsymbol{\psi}}(\tilde{b}), \hat{\boldsymbol{\eta}}(\tilde{b}), \tilde{b}\right) \tag{22}$$

where $\hat{\boldsymbol{\psi}}$ and $\hat{\boldsymbol{\eta}}$ are given by Eq. (5), and an expression for the above probability is given in Eq. (38) in the appendix. Our functional $\hat{G}_n(Z|\tilde{b})$ is the average of $\hat{g}_n(\mathbf{f}(x)|\tilde{b})$ over all instances x_j,

$$\hat{G}_n(Z|\tilde{b}) = \frac{1}{n}\sum_{j=1}^{n} \hat{g}_n(\mathbf{f}(x_j)|\tilde{b}). \tag{23}$$

Note that the estimates of $\boldsymbol{\psi}, \boldsymbol{\eta}$ in Eq. (5) become numerically unstable for b close to ± 1. Hence, in what follows we assume there is an a-priori known $\delta > 0$, such that the true class imbalance $b \in [-1+\delta, 1-\delta]$. The estimate of the class imbalance is then defined as

$$\hat{b}_n = \underset{\tilde{b}\in[-1+\delta,1-\delta]}{\operatorname{argmax}} \hat{G}_n(Z|\tilde{b}). \tag{24}$$

To justify Eq. (24), it is again constructive to consider the limit $n \to \infty$. First, for any $\tilde{b} \in [-1+\delta, 1-\delta]$, the convergence of $\hat{\boldsymbol{\psi}}(\tilde{b})$ and $\hat{\boldsymbol{\eta}}(\tilde{b})$ to $\boldsymbol{\psi}(\tilde{b})$ and $\boldsymbol{\eta}(\tilde{b})$, respectively, implies that at any instance x,

$$\lim_{n\to\infty} \hat{g}_n(\mathbf{f}(x)|\tilde{b}) = g(\mathbf{f}(x)|\tilde{b}) \equiv \log \Pr(\mathbf{f}(x)|\boldsymbol{\psi}(\tilde{b}), \boldsymbol{\eta}(\tilde{b}), \tilde{b}).$$

Next, since the n instances x_j are i.i.d, by the law of large numbers, combined with the delta method

$$\lim_{n\to\infty} \hat{G}_n(Z|\tilde{b}) = G(\tilde{b}) \equiv \mathbb{E}_{(X,Y)}\left[g(\mathbf{f}(X)|\tilde{b})\right]. \tag{25}$$

The following theorem, proven in the appendix, shows that the maxima of $G(\tilde{b})$ is obtained at the true class imbalance $\tilde{b} = b$, and that $\hat{b}_n \to b$ in probability.

Theorem 1. *Assume all classifier errors are independent, so Eq. (21) holds. Let $\epsilon, \delta > 0$ be a-priori known, such that classifiers sensitivities and specificities satisfy $\epsilon < \psi_i, \eta_i < 1-\epsilon$, and $b \in [-1+\delta, 1-\delta]$. Then,*

$$b = \underset{\tilde{b}\in[-1+\delta,1-\delta]}{\operatorname{argmax}} \mathbb{E}_{(X,Y)}\left[g(\mathbf{f}(X)|\tilde{b})\right] \tag{26}$$

and as $n \to \infty$ the estimate $\hat{b}_n$ of Eq. (24) converges to b in probability.

Note that since $\hat{b}_n$ is the maximizer of a restricted likelihood, its convergence to b is not a direct consequence of the consistency of ML estimators. Instead, what is needed is uniform convergence in probability of $\hat{G}_n(\tilde{b})$ to $G(\tilde{b})$, see Newey [1991] and appendix. Also note that even though $\hat{G}_n(\tilde{b})$ is not necessarily concave, finding its global maxima requires optimization of a smooth function of only one variable.

Algorithm 2 Estimating the class imbalance using the restricted likelihood functional

1: Estimate the mean values $\{\hat{\mu}_i\}_{i=1}^m$, the covariance matrix $\hat{R}$, and the vector $\hat{\mathbf{v}}$.
2: **for** $\tilde{b} \in (-1+\delta, 1-\delta)$ **do**
3: Estimate $\hat{\boldsymbol{\psi}}(\tilde{b})$ and $\hat{\boldsymbol{\eta}}(\tilde{b})$ via Eq. (5).
4: Calculate $\hat{G}_n(Z|\tilde{b})$ by Eqs. (22) and (23).
5: **end for**
6: Estimate b by Eq. (24).

Algorithm 2 summarizes the method to estimate b by the restricted-likelihood method. This algorithm scans possible values of $\tilde{b}$, where each evaluation of $\hat{G}_n$ requires $O(mn)$ operations. Since $\hat{g}_n$ and consequently $\hat{G}_n$ are smooth functions of $\tilde{b}$ in $(-1+\delta, 1-\delta)$, the finite grid of values of $\tilde{b}$ can be of size polynomial in n and the method is computationally efficient.

5 The multi-class case

We now consider the multi-class case, with $K > 2$ classes. Here we are given the predictions of m classifiers, $f_i : \mathcal{X} \to \mathcal{Y}$, where $\mathcal{Y} = \{1, \ldots, K\}$. Instead of the class imbalance b, we now have a vector of K class probabilities $p_k = \Pr(Y = k)$. Similarly, instead of specificity and sensitivity, now each classifier is characterized by a $K \times K$ confusion matrix ψ^i

$$\psi^i_{kk'} = \Pr(f_i(X) = k | Y = k') \qquad k, k' \in \mathcal{Y}.$$

In analogy to Section 2, given only an $m \times n$ matrix of predictions, with elements $f_i(x_j) \in \{1 \ldots K\}$, the problem is to estimate the confusion matrices ψ^i of all classifiers and the class probabilities p_k.

As in the binary case, we make an assumption regarding the mutual independence of errors made by different classifiers. The precise independence assumption (pairs, triplets or the full set of classifiers) depends on the method employed.

By a simple reduction to the binary case, we now present a partial solution to this problem. We develop a method to consistently estimate the class probabilities p_k and the diagonals of the confusion matrices, namely the probabilities $\Pr(f_i(X) = k | Y = k)$. However, we prove that even if the class probabilities are a-priori known, estimating all entries of the m confusion matrices is not possible via this binary reduction.

To this end, we build upon the methods developed in Sections 3 and 4 for binary problems. Consider a split of the group $\mathcal{Y} = \{1 \ldots K\}$ into two non-empty disjoint subsets, $\mathcal{Y} = \mathcal{A} \cup (\mathcal{Y} \setminus \mathcal{A})$, where $\mathcal{A} \subset \mathcal{Y}$ is a non trivial subset of $\mathcal{Y}$, with $0 < |\mathcal{A}| < K$. Next,

define the binary classifiers $\{f_i^{\mathcal{A}}\}_{i=1}^m$:

$$f_i^{\mathcal{A}}(X) = \begin{cases} 1 & f_i(X) \in \mathcal{A} \\ -1 & f_i(X) \notin \mathcal{A} \end{cases}$$

Using one of the algorithms described in Section 4, we estimate the probability of the group $\mathcal{A}$

$$p^{\mathcal{A}} = \Pr(Y \in \mathcal{A}) = \sum_{k \in \mathcal{A}} p_k$$

and the sensitivity of each classifier $f_i^{\mathcal{A}}$ by Eq. (5).

In particular, when $\mathcal{A} = \{k\}$, $p^{\mathcal{A}} = p_k$ and $\psi_i^{\mathcal{A}} = \psi_{kk}^i$. Hence, by considering all 1-vs.-all splits, we consistently and computationally efficiently estimate all class probabilities p_k, and all diagonal entries ψ_{kk}^i.

The following theorem, proven in the appendix, states a negative result, that estimating the full confusion matrix is not possible by this binary reduction method.

Theorem 2. *Let $\mu_{\mathcal{A}}^i = \mathbb{E}[f_i^{\mathcal{A}}]$ and let $R_{\mathcal{A}}$ be the covariance matrix of the classifiers $\{f_i^{\mathcal{A}}\}_{i=1}^m$. The inverse problem of estimating the m confusions matrices ψ^i, from the values of $\{\mu_{\mathcal{A}}^i\}_{i=1}^m$ and $R_{\mathcal{A}}$ for all possible subsets $\mathcal{A}$ of $\mathcal{Y} = \{1 \dots K\}$, is in general ill posed with multiple solutions.*

Theorem 2 implies that in order to completely estimate the confusion matrices in a multiclass problem, it is necessary to use higher-order dependencies such as tensors or even the full likelihood. Indeed, both Zhang et al. [2014] and Jain and Oh [2014] derived such methods based on three-dimensional tensors.

While beyond the scope of this paper, we remark that combining our simpler method with these tensor-based approaches might produce more accurate algorithms for the multiclass case.

6 Experiments

6.1 Artificial Data

First, we demonstrate the performance of the two class imbalance estimators on artificial binary data. In the following we constructed an ensemble of $m = 10$ classifiers that make independent errors and thus satisfy Eq. (21). Their sensitivities and specificities were chosen uniformly at random from the interval $[0.5, 0.8]$. Thus, assumption (iii) on the balanced accuracies $\boldsymbol{\pi}$ holds. The vector of true labels $\mathbf{y} \in \{\pm 1\}^n$ was randomly generated according to the class imbalance b, and the data matrix Z was randomly generated according to $\mathbf{y}, \boldsymbol{\psi}$, and $\boldsymbol{\eta}$.

Fig. 1 presents the accuracy (mean and standard deviation) of the estimates $\hat{b}$ of the class imbalance,

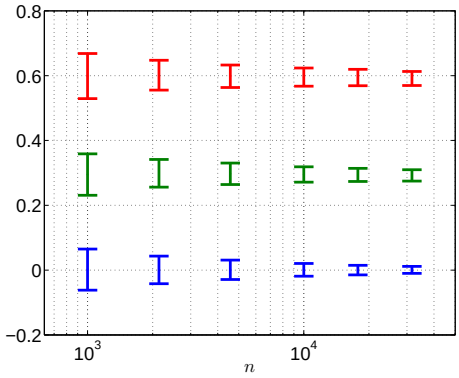

(a) Estimating b via the 3-D tensor T.

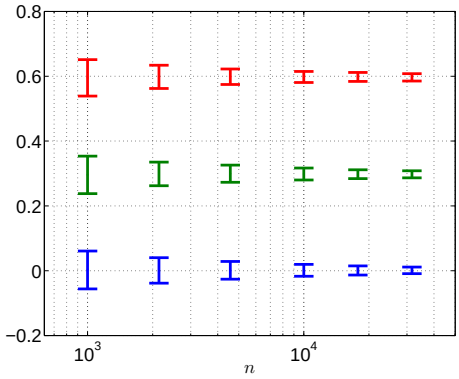

(b) Estimating b via the restricted likelihood $\hat{G}_n$

Fig. 1: Mean and variance of the tensor-based and likelihood-based class imbalance estimators vs. number of instances n, for several values of b.

achieved by the two different algorithms of Sections 4.1 and 4.2, vs. the number of unlabeled instances n, for several values of the class imbalance, $b = 0, 0.3, 0.6$. As expected, the accuracy of both methods improves with the number of instances. Fig. 2 shows the mean squared error (MSE) $\mathbb{E}[(\hat{b}-b)^2]$ vs. the number of samples n, on a log-log scale. The linear line with slope ≈ -1 shows that empirically $\hat{b}_n = b + \mathcal{O}_P(1/\sqrt{n})$, in accordance to Lemma 5. In addition, on simulated data, the restricted likelihood estimator is more accurate than the tensor-based estimator.

6.2 Real data

We applied our algorithms on various binary and multi-class problems using a total of 5 datasets: 4 datasets from the UCI repository Bache and Lichman [2013] and the MNIST data. Our ensemble consisted of $m = 10$ classification methods implemented in the software package Weka Hall et al. [2009]. Due to page limits, we present here results only on the 'magic' dataset. Further details on the different datasets, classifiers and additional results appear in the appendix.

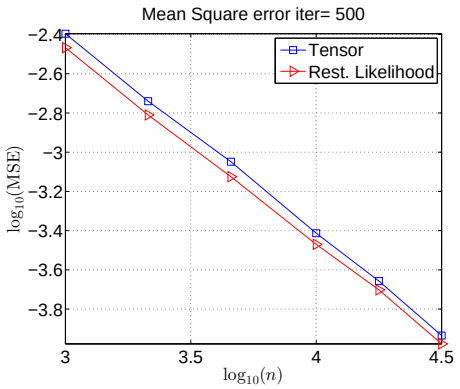

Fig. 2: The MSE of the two class imbalance estimators vs. number of samples on a log-log scale.

The magic data contains 19,000 instances with 11 attributes. The task is to distinguish each instance as either background or high energy gamma rays. Each of the $m = 10$ classifiers was trained on its own randomly chosen set of 200 instances. The classifiers were then applied to the whole dataset, thus providing the $m \times n$ prediction matrix. We compared the results of 4 different unsupervised ensemble methods: (i) Majority voting; (ii) SML of Parisi et al. [2014]; (iii) i-SML as described in section 4; and (iv) Oracle ML: the MLE formula (8) with the values of $\boldsymbol{\psi}$ and $\boldsymbol{\eta}$, estimated from the full dataset with its labels.

To assess the stability of the different methods, for each dataset we repeated the above simulation 30 times, each realization with different randomly chosen training sets. Fig. 3a shows the mean and standard deviation of the balanced accuracy π achieved by the four methods on the 'magic' dataset. It shows that on average, i-SML improves upon the SML by approximately 2%, and both are significantly better than majority voting. Fig. 3b displays the error rates $1-\pi_{\text{i-SML}}$ vs. $1-\pi_{\text{SML}}$ for all 30 realizations. As all points are below the diagonal, the improvement over SML was consistent in all 30 simulation runs. As shown in the appendix, similar results, and in particular the improvement of i-SML over SML, were observed also in all 4 other datasets.

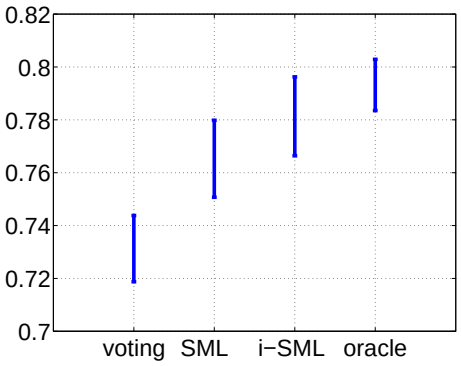

(a) The balanced accuracies of 4 unsupervised ensemble methods on the magic dataset.

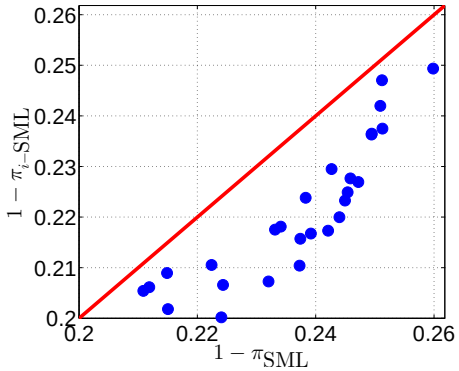

(b) The empirical test error $(1-\pi_{\text{i-SML}})$ vs. $(1-\pi_{\text{SML}})$ for 30 random realizations.

Fig. 3: Comparing 4 unsupervised ensemble learning algorithms, based on $m = 10$ classifiers.

7 Summary and Discussion

In this paper we presented a simple spectral-based approach to estimate, in an unsupervised manner, the accuracies of multiple classifiers, mainly in the binary case. This, in turn, resulted in a novel unsupervised spectral ensemble learner, denoted i-SML. The empirical results on several real data sets attest to its competitive performance in practical situations where clearly the underlying idealized assumptions that all classifiers make independent errors do not hold exactly.

There are several interesting directions to extend this work. One possible direction is to relax the strict assumptions of independence of classifier errors across all instances, for example by introducing the concept of instance difficulty. A second interesting direction is the construction of novel semi-supervised ensemble learners, when one is given not only the predictions of m classifiers on a large unlabeled set of instances, but also their predictions on a small set of labeled ones.

Acknowledgments

This research was funded in part by the Intel Collaborative Research Institute for Computational Intelligence (ICRI-CI). Y.K. is supported by the National Institutes of Health Grant R0-1 CA158167.

References

A. Anandkumar, D. Hsu, and S.M. Kakade. A method of moments for mixture models and hidden Markov models. In *Conference on Learning Theory (COLT)*, 2012.

A. Anandkumar, R. Ge, D. Hsu, S.M. Kakade, and M. Telgarsky. Tensor decompositions for learning latent variable models. *Journal of Machine Learning Research*, 15:2773–2832, Aug. 2014.

K. Bache and M. Lichman. UCI machine learning repository, 2013.

A. P Dawid and A. M Skene. Maximum likelihood estimation of observer error-rates using the em algorith. *Journal of the Royal Statistical Society. Series C*, 28:20–28, 1979.

T.G. Dietterich. Ensemble methods in machine learning. In *Lecture Notes in Computer Science*, volume 1857, pages 1–15. Springer, Berlin, 2000.

P. Donmez, G. Lebanon, and K. Balasubramanian. Unsupervised supervised learning i: Estimating classification and regression errors without labels. *Journal of Machine Learning Research*, pages 1323–1351, 2010.

J. Feldman, R. O'Donnell, and R.A. Servedio. Learning mixtures of product distributions over discrete domains. *SIAM Journal on Computing*, 37(5):1536–1564, 2008.

Y. Freund and Y. Mansour. Estimating a mixture of two product distributions. In *COLT '99 Proceedings of the twelfth annual conference on Computational learning theory*, pages 53–62, 1999.

M. Hall, E. Frank, G. Holmes, G. Pfahringer, P. Reutemann, and I.H Witten. The Weka data mining software: An update. *SIGKDD Explorations*, 11(1), 2009.

P. Jain and S. Oh. Learning mixtures of discrete product distributions using spectral decompositions. *Journal of Machine Learning Research*, 35:1–33, 2014.

D.R. Karger, S. Oh, and D. Shah. Budget-optimal crowdsourcing using low-rank matrix approximations. In *IEEE Alerton Conference on Communication, Control and Computing*, pages 284–291, 2011.

W. K. Newey. Uniform convergence in probability and stochastic equicontinuity. *Econometrica*, 59:1161–1167, 1991.

F. Parisi, F. Strino, B. Nadler, and Y. Kluger. Ranking and combining multiple predictors without labeled data. *Proceedings of the National Academy of Sciences*, 111:1253–1258, 2014.

E.A. Platanios, A. Blum, and T. Mitchell. Estimating accuracy from unlabeled data. In *Uncertainty in Artificial Intelligence*, 2014.

V.C. Raykar, Y. Shipeng, L.H. Zhao, G.H. Valdez, C. Florin, L. Bogoni, and Moy L. Learning from crowds. *J. Machine Learning Research*, 11:1297–1322, 2010.

L. Rokach. Collective-agreement-based pruning of ensembles. *Computational Statistics and Data Analysis*, 53:1015–1026, 2009.

V.S. Sheng, F. Provost, and P.G. Ipeirotis. Get another label? improving data quality and data mining using multiple, noisy labelers. In *14th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 614–622, 2008.

A. Sheshadri and M. Lease. Square: A benchmark for research on computing crowd concensus. In *AAAI conference on human computation and crowdsourcing*, 2013.

R. Snow, B. O'Connor, D. Jurafsky, and A.Y. Ng. Cheap and fast – but is it good? In *Conference on Empirical Methods in Natural Language Processing*, 2008.

P. Welinder, S. Branson, S. Belongie, and P. Perona. The multidimensional wisdom of crowds. In *Advances in Neural Information Processing Systems 23 (NIPS 2010)*, 2010.

J. Whitehill, P. Ruvolo, T. Wu, J Bergsma, and J.R. Movellan. Whose vote should count more: Optimal integration of labels from labelers of unknown expertise. In *Advances in Neural Information Processing Systems 22 (NIPS 2009)*, 2009.

Y. Zhang, X. Chen, D. Zhou, and M.I. Jordan. Spectral methods meet EM: A provably optimal algorithm for crowdsourcing. In *Advances in Neural Information Processing Systems*, volume 27, pages 1260–1268, 2014.