
Estimating the class prior and posterior from noisy positives and unlabeled data

Shantanu Jain, Martha White, Predrag Radivojac
Department of Computer Science
Indiana University, Bloomington, Indiana, USA
{shajain, martha, predrag}@indiana.edu

Abstract

We develop a classification algorithm for estimating posterior distributions from positive-unlabeled data, that is robust to noise in the positive labels and effective for high-dimensional data. In recent years, several algorithms have been proposed to learn from positive-unlabeled data; however, many of these contributions remain theoretical, performing poorly on real high-dimensional data that is typically contaminated with noise. We build on this previous work to develop two practical classification algorithms that explicitly model the noise in the positive labels and utilize univariate transforms built on discriminative classifiers. We prove that these univariate transforms preserve the class prior, enabling estimation in the univariate space and avoiding kernel density estimation for high-dimensional data. The theoretical development and parametric and nonparametric algorithms proposed here constitute an important step towards wide-spread use of robust classification algorithms for positive-unlabeled data.

1 Introduction

Access to positive, negative and unlabeled examples is a standard assumption for most semi-supervised binary classification techniques. In many domains, however, a sample from one of the classes (say, negatives) may not be available, leading to the setting of learning from positive and unlabeled data (Denis et al., 2005). Positive-unlabeled learning often emerges in sciences and commerce where an observation of a positive example (say, that a protein catalyzes reactions or that a social network user likes a particular product) is usually reliable. Here, however, the absence of a positive observation cannot be interpreted as a negative example. In molecular biology, for example, an attempt to label a data point as positive (say, that a protein is an enzyme) may be unsuccessful for a variety of experimental and biological reasons, whereas in social networks an explicit dislike of a product may not be possible. Both scenarios lead to a situation where negative examples cannot be actively collected.

Fortunately, the absence of negatively labeled examples can be tackled by incorporating unlabeled examples as negatives, leading to the development of non-traditional classifiers. Here we follow the terminology by Elkan and Noto (2008) that a traditional classifier predicts whether an example is positive or negative, whereas a non-traditional classifier predicts whether the example is positive or unlabeled. Positive vs. unlabeled (non-traditional) training is reasonable because the class posterior — and also the optimum scoring function for composite losses (Reid and Williamson, 2010) — in the traditional setting is monotonically related to the posterior in the non-traditional setting. However, the true posterior can be fully recovered from the non-traditional posterior only if we know the class prior; i.e., the proportion of positives in unlabeled data. The knowledge of the class prior is also necessary for estimation of the performance criteria such as the error rate, balanced error rate or F-measure, and also for finding the right threshold for the non-traditional scoring function that leads to an optimal classifier with respect to some criteria (Menon et al., 2015).

Class prior estimation in a nonparametric setting has been actively researched in the past decade offering an extensive theory of identifiability (Ward et al., 2009; Blanchard et al., 2010; Scott et al., 2013; Jain et al., 2016) and a few practical solutions (Elkan and Noto, 2008; Ward et al., 2009; du Plessis and Sugiyama, 2014; Sanderson and Scott, 2014; Jain et al., 2016; Ramaswamy et al., 2016). Application of these algorithms to real data, however, is limited in that none of the proposed algorithms simultaneously deals with noise in the labels and practical estimation for high-dimensional data.

Much of the theory on learning class priors relies on the assumption that either the distribution of positives is known or that the positive sample is clean. In practice, however, labeled data sets contain class-label noise, where an unspecified amount of negative examples contaminates the positive sample. This is a realistic scenario in experimental sciences where technological advances enabled generation of high-throughput data at a cost of occasional errors. One example for this comes from the studies of proteins using analytical chemistry technology; i.e., mass spectrometry. For example, in the process of peptide identification (Steen and Mann, 2004), bioinformatics methods are usually set to report results with specified false discovery rate thresholds (e.g., 1%). Unfortunately, statistical assumptions in these experiments are sometimes violated thereby leading to substantial noise in reported results, as in the case of identifying protein post-translational modifications. Similar amounts of noise might appear in social networks such as Facebook, where some users select ‘like’, even when they do not actually like a particular post. Further, the only approach that does consider similar such noise (Scott et al., 2013) requires density estimation, which is known to be problematic for high-dimensional data.

In this work, we propose the first classification algorithm, with class prior estimation, designed particularly for high-dimensional data with noise in the labeling of positive data. We first formalize the problem of class prior estimation from noisy positive and unlabeled data. We extend the existing identifiability theory for class prior estimation from positive-unlabeled data to this noise setting. We then show that we can practically estimate class priors and the posterior distributions by first transforming the input space to a univariate space, where density estimation is reliable. We prove that these transformations preserve class priors and show that they correspond to training a non-traditional classifier. We derive a parametric algorithm and a nonparametric algorithm to learn the class priors. Finally, we carry out experiments on synthetic and real-life data and provide evidence that the new approaches are sound and effective.

2 Problem formulation

Consider a binary classification problem of mapping an input space \mathcal{X} to an output space $\mathcal{Y} = \{0, 1\}$. Let f be the true distribution of inputs. It can be represented as the following mixture

$$f(x) = \alpha f_1(x) + (1 - \alpha) f_0(x), \tag{1}$$

where $x \in \mathcal{X}$, $y \in \mathcal{Y}$, f_y are distributions over \mathcal{X} for the positive ($y = 1$) and negative ($y = 0$) class, respectively; and $\alpha \in [0, 1]$ is the class prior or the proportion of the positive examples in f . We will refer to a sample from f as unlabeled data.

Let g be the distribution of inputs for the labeled data. Because the labeled sample contains some mislabeled examples, the corresponding distribution is also a mixture of f_1 and a small proportion, say $1 - \beta$, of f_0 . That is,

$$g(x) = \beta f_1(x) + (1 - \beta) f_0(x), \tag{2}$$

where $\beta \in (0, 1]$. Observe that both mixtures have the same components but different mixing proportions. The simplest scenario is that the mixing components f_0 and f_1 correspond to the class-conditional distributions $p(x|Y = 0)$ and $p(x|Y = 1)$, respectively. However, our approach also permits transformations of the input space \mathcal{X} , thus resulting in a more general setup.

The objective of this work is to study the estimation of the class prior $\alpha = p(Y = 1)$ and propose practical algorithms for estimating α . The efficacy of this estimation is clearly tied to β , where as β gets smaller, the noise in the positive labels becomes larger. We will discuss identifiability of α and β and give a practical algorithm for estimating α (and β). We will then use these results to estimate the posterior distribution of the class variable, $p(y|x)$, despite the fact that the labeled set does not contain any negative examples.

3 Identifiability

The class prior is identifiable if there is a unique class prior for a given pair (f, g) . Much of the identifiability characterization in this section has already been considered as the case of asymmetric noise (Scott et al., 2013); see Section 7 on related work. We recreate these results here, with the aim to introduce required notation, to highlight several important results for later algorithm development and to include a few missing results needed for our approach. Though the proof techniques are themselves quite different and could be of interest, we include them in the appendix due to space.

There are typically two aspects to address with identifiability. First, one needs to determine if a problem is identifiable, and, second, if it is not, propose a canonical form that is identifiable. In this section we will see that class prior is not identifiable in general because f_0 can be a mixture containing f_1 and vice versa. To ensure identifiability, it is necessary to choose a canonical form that prefers a class prior that makes the two components as different as possible; this canonical form was introduced as the mutual irreducibility condition (Scott et al., 2013) and is related to the proper novelty distribution (Blanchard et al., 2010) and the max-canonical form (Jain et al., 2016).

We discuss identifiability in terms of measures. Let μ, ν, μ_0 and μ_1 be probability measures defined on some σ -algebra \mathcal{A} on \mathcal{X} , corresponding to f, g, f_0 and f_1 , respectively. It follows that

$$\mu = \alpha\mu_1 + (1 - \alpha)\mu_0 \quad (3)$$

$$\nu = \beta\mu_1 + (1 - \beta)\mu_0. \quad (4)$$

Consider a family of pairs of mixtures having the same components

$$\mathcal{F}(\Pi) = \{(\mu, \nu) : \mu = \alpha\mu_1 + (1 - \alpha)\mu_0, \nu = \beta\mu_1 + (1 - \beta)\mu_0, (\mu_0, \mu_1) \in \Pi, 0 \leq \alpha < \beta \leq 1\},$$

where Π is some set of pairs of probability measures defined on \mathcal{A} . The family is parametrized by the quadruple $(\alpha, \beta, \mu_0, \mu_1)$. The condition $\beta > \alpha$ means that ν has a greater proportion of μ_1 compared to μ . This is consistent with our assumption that the labeled sample mainly contains positives. The most general choice for Π is

$$\Pi^{\text{all}} = \mathcal{P}^{\text{all}} \times \mathcal{P}^{\text{all}} \setminus \{(\mu, \mu) : \mu \in \mathcal{P}^{\text{all}}\},$$

where \mathcal{P}^{all} is the set of all probability measures defined on \mathcal{A} and $\{(\mu, \mu) : \mu \in \mathcal{P}^{\text{all}}\}$ is the set of pairs with equal distributions. Removing equal pairs prevents μ and ν from being identical.

We now define the maximum proportion of a component λ_1 in a mixture λ , which is used in the results below and to specify the criterion that enables identifiability; more specifically,

$$a_{\lambda}^{\lambda_1} = \max\{\alpha \in [0, 1] : \lambda = \alpha\lambda_1 + (1 - \alpha)\lambda_0, \lambda_0 \in \mathcal{P}^{\text{all}}\}. \quad (5)$$

Of particular interest is the case when $a_{\lambda}^{\lambda_1} = 0$, which should be read as “ λ is not a mixture containing λ_1 ”. We finally define the set all possible (α, β) that generate μ and ν when (μ_0, μ_1) varies in Π :

$$A^+(\mu, \nu, \Pi) = \{(\alpha, \beta) : \mu = \alpha\mu_1 + (1 - \alpha)\mu_0, \nu = \beta\mu_1 + (1 - \beta)\mu_0, (\mu_0, \mu_1) \in \Pi, 0 \leq \alpha < \beta \leq 1\}.$$

If $A^+(\mu, \nu, \Pi)$ is a singleton set for all $(\mu, \nu) \in \mathcal{F}(\Pi)$, then $\mathcal{F}(\Pi)$ is identifiable in (α, β) .

First, we show that the most general choice for Π , Π^{all} , leads to unidentifiability (Lemma 1). Fortunately, however, by choosing a restricted set

$$\Pi^{\text{res}} = \{(\mu_0, \mu_1) \in \Pi^{\text{all}} : a_{\mu_0}^{\mu_1} = 0, a_{\mu_1}^{\mu_0} = 0\}$$

as Π , we do obtain identifiability (Theorem 1). In words, Π^{res} contains pairs of distributions, where each distribution in a pair cannot be expressed as a mixture containing the other. The proofs of the results below are in the Appendix.

Lemma 1 (Unidentifiability) *Given a pair of mixtures $(\mu, \nu) \in \mathcal{F}(\Pi^{\text{all}})$, let parameters $(\alpha, \beta, \mu_0, \mu_1)$ generate (μ, ν) and $\alpha^+ = a_{\mu}^{\nu}, \beta^+ = a_{\nu}^{\mu}$. It follows that*

1. *There is a one-to-one relation between (μ_0, μ_1) and (α, β) and*

$$\mu_0 = \frac{\beta\mu - \alpha\nu}{\beta - \alpha}, \quad \mu_1 = \frac{(1 - \alpha)\nu - (1 - \beta)\mu}{\beta - \alpha}. \quad (6)$$

2. Both expressions on the right-hand side of Equation 6 are well defined probability measures if and only if $\alpha/\beta \leq \alpha^+$ and $(1-\beta)/(1-\alpha) \leq \beta^+$.
3. $A^+(\mu, \nu, \Pi^{all}) = \{(\alpha, \beta) : \alpha/\beta \leq \alpha^+, (1-\beta)/(1-\alpha) \leq \beta^+\}$.
4. $\mathcal{F}(\Pi^{all})$ is unidentifiable in (α, β) ; i.e., (α, β) is not uniquely determined from (μ, ν) .
5. $\mathcal{F}(\Pi^{all})$ is unidentifiable in α and β , individually; i.e., neither α nor β is uniquely determined from (μ, ν) .

Observe that the definition of α_λ^+ and $\mu \neq \nu$ imply $\alpha^+ < 1$ and, consequently, any $(\alpha, \beta) \in A^+(\mu, \nu, \Pi^{all})$ satisfies $\alpha < \beta$, as expected.

Theorem 1 (Identifiability) Given $(\mu, \nu) \in \mathcal{F}(\Pi^{all})$, let $\alpha^+ = \alpha_\mu^\nu$ and $\beta^+ = \alpha_\nu^\mu$. Let $\mu_0^* = (\mu - \alpha^+ \nu) / (1 - \alpha^+)$, $\mu_1^* = (\nu - \beta^+ \mu) / (1 - \beta^+)$ and

$$\alpha^* = \alpha^+(1 - \beta^+) / (1 - \alpha^+ \beta^+), \quad \beta^* = (1 - \beta^+) / (1 - \alpha^+ \beta^+). \quad (7)$$

It follows that

1. $(\alpha^*, \beta^*, \mu_0^*, \mu_1^*)$ generate (μ, ν)
2. $(\mu_0^*, \mu_1^*) \in \Pi^{res}$ and, consequently, $\alpha^* = \alpha_{\mu_0^*}^{\mu_1^*}$, $\beta^* = \alpha_{\mu_1^*}^{\mu_0^*}$.
3. $\mathcal{F}(\Pi^{res})$ contains all pairs of mixtures in $\mathcal{F}(\Pi^{all})$.
4. $A^+(\mu, \nu, \Pi^{res}) = \{(\alpha^*, \beta^*)\}$.
5. $\mathcal{F}(\Pi^{res})$ is identifiable in (α, β) ; i.e., (α, β) is uniquely determined from (μ, ν) .

We refer to the expressions of μ and ν as mixtures of components μ_0 and μ_1 as a max-canonical form when (μ_0, μ_1) is picked from Π^{res} . This form enforces that μ_1 is not a mixture containing μ_0 and vice versa, which leads to μ_0 and μ_1 having maximum separation, while still generating μ and ν . Each pair of distributions in $\mathcal{F}(\Pi^{res})$ is represented in this form. Identifiability of $\mathcal{F}(\Pi^{res})$ in (α, β) occurs precisely when $A^+(\mu, \nu, \Pi^{res}) = \{(\alpha^*, \beta^*)\}$, i.e., (α^*, β^*) is the only pair of mixing proportions that can appear in a max-canonical form of μ and ν . Moreover, Statement 1 in Theorem 1 and Statement 1 in Lemma 1 imply that the max-canonical form is unique and completely specified by $(\alpha^*, \beta^*, \mu_0^*, \mu_1^*)$, with $\alpha^* < \beta^*$ following from Equation 7. Thus, using $\mathcal{F}(\Pi^{res})$ to model the unlabeled and labeled data distributions makes estimation of not only α , the class prior, but also β, μ_0, μ_1 a well-posed problem. Moreover, due to Statement 3 in Theorem 1, there is no loss in the modeling capability by using $\mathcal{F}(\Pi^{res})$ instead of $\mathcal{F}(\Pi^{all})$. Overall, identifiability, absence of loss of modeling capability and maximum separation between μ_0 and μ_1 combine to justify estimating α^* as the class prior.

4 Univariate Transformation

The theory and algorithms for class prior estimation are agnostic to the dimensionality of the data; in practice, however, this dimensionality can have important consequences. Parametric Gaussian mixture models trained via expectation-maximization (EM) are known to strongly suffer from collinearity in high-dimensional data. Nonparametric (kernel) density estimation is also known to have curse-of-dimensionality issues, both in theory (Liu et al., 2007) and in practice (Scott, 2008).

We address the curse of dimensionality by transforming the data to a single dimension. The transformation $\tau : \mathcal{X} \rightarrow \mathbb{R}$, surprisingly, is simply an output of a non-traditional classifier trained to separate labeled sample, L , from unlabeled sample, U . The transform is similar to that in (Jain et al., 2016), except that it is not required to be calibrated like a posterior distribution; as shown below, a good ranking function is sufficient. First, however, we introduce notation and formalize the data generation steps (Figure 1).

Let X be a random variable taking values in \mathcal{X} , capturing the true distribution of inputs, μ , and Y be an unobserved random variable taking values in \mathcal{Y} , giving the true class of the inputs. It follows that $X|Y = 0$ and $X|Y = 1$ are distributed according to μ_0 and μ_1 , respectively. Let S be a selection random variable, whose value in $\mathcal{S} = \{0, 1, 2\}$ determines the sample to which an input x is added (Figure 1). When $S = 1$, x is added to the noisy labeled sample; when $S = 0$, x is added to the unlabeled sample; and when $S = 2$, x is not added to either of the samples. It follows that

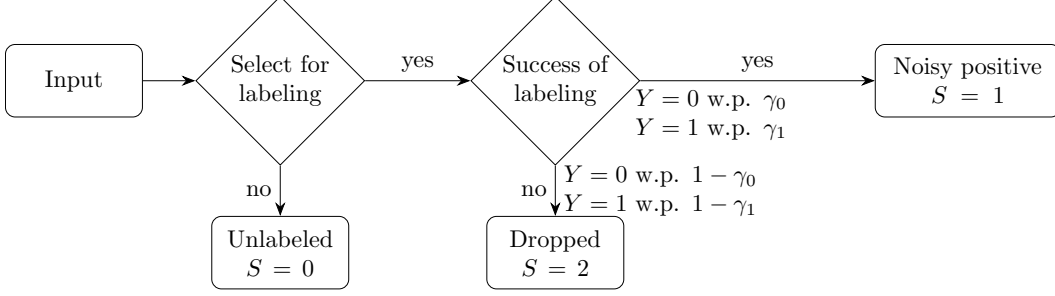


Figure 1: The labeling procedure, with S taking values from $\mathcal{S} = \{0, 1, 2\}$. In the first step, the sample is randomly selected to attempt labeling, with some probability independent of X or Y . If it is not selected, it is added to the “Unlabeled” set. If it is selected, then labeling is attempted. If the true label is $Y = 1$, then with probability $\gamma_1 \in (0, 1)$, the labeling will succeed and it is added to “Noisy positives”. Otherwise, it is added to the “Dropped” set. If the true label is $Y = 0$, then the attempted labeling is much more likely to fail, but because of noise, could succeed. The attempted label of $Y = 0$ succeeds with probability γ_0 , and is added to “Noisy positives”, even though it is actually a negative instance. $\gamma_0 = 0$ leads to the no noise case and the noise increases as γ_0 increases. $\beta = \gamma_1 \alpha / (\gamma_1 \alpha + \gamma_0 (1 - \alpha))$, gives the proportion of positives in the “Noisy positives”.

$X^u = X|S = 0$ and $X^l = X|S = 1$ are distributed according to μ and ν , respectively. We make the following assumptions which are consistent with the statements above:

$$p(y|S = 0) = p(y), \quad (8)$$

$$p(y = 1|S = 1) = \beta, \quad (9)$$

$$p(x|s, y) = p(x|y). \quad (10)$$

Assumptions 8 and 9 states that the proportion of positives in the unlabeled sample and the labeled sample matches the true proportion in μ and ν , respectively. Assumption 10 states that the distribution of the positive inputs (and the negative inputs) in both the unlabeled and the labeled samples is equal and unbiased. Lemma 2 gives the implications of these assumptions. Statement 3 in Lemma 2 is particularly interesting and perhaps counter-intuitive as it states that with non-zero probability some inputs need to be dropped.

Lemma 2 *Let X, Y and S be random variables taking values in \mathcal{X}, \mathcal{Y} and \mathcal{S} , respectively, and $X^u = X|S = 0$ and $X^l = X|S = 1$. For measures μ, ν, μ_0, μ_1 , satisfying Equations 3 and 4 and $\mu_1 \neq \mu_0$, let μ, μ_0, μ_1 give the distribution of $X, X|Y = 0$ and $X|Y = 1$, respectively. If X, Y and S satisfy assumptions 8, 9 and 10, then*

1. X is independent of $S = 0$; i.e., $p(x|S = 0) = p(x)$
2. X^u and X^l are distributed according to μ and ν , respectively.
3. $p(S = 2) \neq 0$.

The proof is in the Appendix. Next, we highlight the conditions under which the score function τ preserves α^* . Observing that S serves as the pseudo class label for labeled vs. unlabeled classification as well, we first give an expression for the posterior:

$$\tau_p(x) = p(S = 1|x, S \in \{0, 1\}), \quad \forall x \in \mathcal{X}. \quad (11)$$

Theorem 2 (α^* -preserving transform) *Let random variables X, Y, S, X^u, X^l and measures μ, ν, μ_0, μ_1 be as defined in Lemma 2. Let τ_p be the posterior as defined in Equation 11 and $\tau = H \circ \tau_p$, where H is a 1-to-1 function on $[0, 1]$ and \circ is the composition operator. Assume*

1. $(\mu_0, \mu_1) \in \Pi^{res}$,
2. X^u and X^l are continuous with densities f and g , respectively,
3. $\mu_\tau, \nu_\tau, \mu_{\tau 1}$ are the measures corresponding to $\tau(X^u), \tau(X^l), \tau(X_1)$, respectively,
4. $(\alpha^+, \beta^+, \alpha^*, \beta^*) = (a_\mu^\nu, a_\nu^\mu, a_\mu^{\mu_1}, a_\nu^{\mu_1})$ and $(\alpha_\tau^+, \beta_\tau^+, \alpha_\tau^*, \beta_\tau^*) = (a_{\mu_\tau}^\nu, a_{\nu_\tau}^\mu, a_{\mu_\tau}^{\mu_{\tau 1}}, a_{\nu_\tau}^{\mu_{\tau 1}})$.

Then

$$(\alpha_\tau^+, \beta_\tau^+, \alpha_\tau^*, \beta_\tau^*) = (\alpha^+, \beta^+, \alpha^*, \beta^*)$$

and so τ is an α^* -preserving transformation.

Moreover, τ_p can also be used to compute the true posterior probability:

$$p(Y = 1|x) = \frac{\alpha^*(1 - \alpha^*)}{\beta^* - \alpha^*} \left(\frac{p(S = 0)}{p(S = 1)} \frac{\tau_p(x)}{1 - \tau_p(x)} - \frac{1 - \beta^*}{1 - \alpha^*} \right). \quad (12)$$

The proof is in the Appendix. Theorem 2 shows that the α^* is the same for the original data and the transformed data, if the transformation function τ can be expressed as a composition of τ_p and a one-to-one function, H , defined on $[0, 1]$. Trivially, τ_p itself is one such function. We emphasize, however, that α^* -preservation is not limited by the efficacy of the calibration algorithm; uncalibrated scoring that ranks inputs as $\tau_p(x)$ also preserves α^* . Theorem 2 further demonstrates how the true posterior, $p(Y = 1|x)$, can be recovered from τ_p by plugging in estimates of τ_p , $p^{(S=0)}/p^{(S=1)}$, α^* and β^* in Equation 12. The posterior probability τ_p can be estimated directly by using a probabilistic classifier or by calibrating a classifier’s score (Platt, 1999; Niculescu-Mizil and Caruana, 2005); $|U|/|L|$ serves as an estimate of $p^{(S=0)}/p^{(S=1)}$; section 5 gives parametric and nonparametric approaches for estimation of α^* and β^* .

5 Algorithms

In this section, we derive a parametric and a nonparametric algorithm to estimate α^* and β^* from the unlabeled sample, $U = \{X_i^u\}$, and the noisy positive sample, $L = \{X_i^l\}$. In theory, both approaches can handle multivariate samples; in practice, however, to circumvent the curse of dimensionality, we exploit the theory of α^* -preserving univariate transforms to transform the samples.

Parametric approach. The parametric approach is derived by modeling each sample as a two component Gaussian mixture, sharing the same components but having different mixing proportions:

$$\begin{aligned} X_i^u &\sim \alpha \mathcal{N}(u_1, \Sigma_1) + (1 - \alpha) \mathcal{N}(u_0, \Sigma_0) \\ X_i^l &\sim \beta \mathcal{N}(u_1, \Sigma_1) + (1 - \beta) \mathcal{N}(u_0, \Sigma_0) \end{aligned}$$

where $u_1, u_0 \in \mathbb{R}^d$ and $\Sigma_1, \Sigma_0 \in \mathbb{S}_{++}^d$, the set of all $d \times d$ positive definite matrices. The algorithm is an extension to the EM approach for Gaussian mixture models (GMMs) where, instead of estimating the parameters of a single mixture, the parameters of both mixtures ($\alpha, \beta, u_0, u_1, \Sigma_0, \Sigma_1$) are estimated simultaneously by maximizing the combined likelihood over both U and L . This approach, which we refer to as a multi-sample GMM (MSGMM), exploits the constraint that the two mixtures share the same components. The update rules and their derivation are given in the Appendix.

Nonparametric approach. Our nonparametric strategy directly exploits the results of Lemma 1 and Theorem 1, which give a direct connection between $(\alpha^+ = a_\mu^+, \beta^+ = a_\mu^+)$ and (α^*, β^*) . Therefore, for a two-component mixture sample, M , and a sample from one of the components, C , it only requires an algorithm to estimate the maximum proportion of C in M . For this purpose, we use the AlphaMax algorithm (Jain et al., 2016), briefly summarized in the Appendix. Specifically, our two-step approach for estimating α^* and β^* is as follows: (i) Estimate α^+ and β^+ as outputs of AlphaMax(U, L) and AlphaMax(L, U), respectively; (ii) Estimate (α^*, β^*) from the estimates of (α^+, β^+) by applying Equation 7. We refer to our nonparametric algorithm as AlphaMax-N.

6 Empirical investigation

In this section we systematically evaluate the new algorithms in a controlled, synthetic setting as well as on a variety of data sets from the UCI Machine Learning Repository (Lichman, 2013).

Experiments on synthetic data: We start by evaluating all algorithms in a univariate setting where both mixing proportions, α and β , are known. We generate unit-variance Gaussian and unit-scale Laplace-distributed i.i.d. samples and explore the impact of mixing proportions, the size of the component sample, and the separation and overlap between the mixing components on the accuracy of estimation. The class prior α was varied from $\{0.05, 0.25, 0.50\}$ and the noise component β from $\{1.00, 0.95, 0.75\}$. The size of the labeled sample L was varied from $\{100, 1000\}$, whereas the size of the unlabeled sample U was fixed at 10000.

Experiments on real-life data: We considered twelve real-life data sets from the UCI Machine Learning Repository. To adjust these data to our problems, categorical features were transformed into numerical using sparse binary representation, the regression data sets were transformed into classification based on mean of the target variable, and the multi-class classification problems were converted into binary problems by combining classes. In each data set, a subset of positive and negative examples was randomly selected to provide a labeled sample while the remaining data (without class labels) were used as unlabeled data. The size of the labeled sample was kept at 1000 (or 100 for small data sets) and the maximum size of unlabeled data was set 10000.

Algorithms: We compare the AlphaMax-N and MSGMM algorithms to the Elkan-Noto algorithm (Elkan and Noto, 2008) as well as the noiseless version of AlphaMax (Jain et al., 2016). There are several versions of the Elkan-Noto estimator and each can use any underlying classifier. We used the e_1 alternative estimator combined with the ensembles of 100 two-layer feed-forward neural networks, each with five hidden units. The out-of-bag scores of the same classifier were used as a class-prior preserving transformation that created an input to the AlphaMax algorithms. It is important to mention that neither Elkan-Noto nor AlphaMax algorithm was developed to handle noisy labeled data. In addition, the theory behind the Elkan-Noto estimator restricts its use to class-conditional distributions with non-overlapping supports. The algorithm by du Plessis and Sugiyama (2014) minimizes the same objective as the e_1 Elkan-Noto estimator and, thus, was not implemented.

Evaluation: All experiments were repeated 50 times to be able to draw conclusions with statistical significance. In real-life data, the labeled sample was created randomly by choosing an appropriate number of positive and negative examples to satisfy the condition for β and the size of the labeled sample, while the remaining data was used as the unlabeled sample. Therefore, the class prior in the unlabeled data varies with the selection of the noise parameter β . The mean absolute difference between the true and estimated class priors was used as a performance measure. The best performing algorithm on each data set was determined by multiple hypothesis testing using the P-value of 0.05 and Bonferroni correction.

Results: The comprehensive results for synthetic data drawn from univariate Gaussian and Laplace distributions are shown in Appendix (Table 2). In these experiments no transformation was applied prior to running any of the algorithms. As expected, the results show excellent performance of the MSGMM model on the Gaussian data. These results significantly degrade on Laplace-distributed data, suggesting sensitivity to the underlying assumptions. On the other hand, AlphaMax-N was accurate over all data sets and also robust to noise. These results suggest that new parametric and nonparametric algorithms perform well in these controlled settings.

Table 1 shows the results on twelve real data sets. Here, AlphaMax and AlphaMax-N algorithms demonstrate significant robustness to noise, although the parametric version MSGMM was competitive in some cases. On the other hand, the Elkan-Noto algorithm expectedly degrades with noise. Finally, we investigated the practical usefulness of the α^* -preserving transform. Table 3 (Appendix) shows the results of AlphaMax-N and MSGMM on the real data sets, with and without using the transform. Because of computational and numerical issues, we reduced the dimensionality by using principal component analysis (the original data caused matrix singularity issues for MSGMM and density estimation issues for AlphaMax-N). MSGMM deteriorates significantly without the transform, whereas AlphaMax-N preserves some signal for the class prior. AlphaMax-N with the transform, however, shows superior performance on most data sets.

7 Related work

Class prior estimation in a semi-supervised setting, including positive-unlabeled learning, has been extensively discussed previously; see Saerens et al. (2002); Cortes et al. (2008); Elkan and Noto (2008); Blanchard et al. (2010); Scott et al. (2013); Jain et al. (2016) and references therein. Recently, a general setting for label noise has also been introduced, called the mutual contamination model. The aim under this model is to estimate multiple unknown base distributions, using multiple random samples that are composed of different convex combinations of those base distributions (Katz-Samuels and Scott, 2016). The setting of asymmetric label noise is a subset of this more general setting, treated under general conditions by Scott et al. (2013), and previously investigated under a more restrictive setting as co-training (Blum and Mitchell, 1998). A natural approach is to use robust estimation to learn in the presence of class noise; this strategy, however, has been shown to be ineffective, both theoretically (Long and Servedio, 2010; Manwani and Sastry, 2013) and empirically (Hawkins and McLachlan, 1997; Bashir and Carter, 2005), indicating the need to explicitly model the noise. Generative mixture model approaches have also been developed, which explicitly model the noise (Lawrence and Scholkopf, 2001; Bouveyron and Girard, 2009); these algorithms, however, assume labeled data for each class. As the most related work, though Scott et al. (2013) did not explicitly treat the positive-unlabeled learning with noisy positives, their formulation can incorporate this setting by using $\pi_0 = \alpha$ and $\beta = 1 - \pi_1$. The theoretical and algorithmic treatment, however, is very different. Their focus is on identifiability and analyzing convergence rates and statistical properties, assuming access to some κ^* function which can obtain proportions

Table 1: Mean absolute difference between estimated and true mixing proportion over twelve data sets from the UCI Machine Learning Repository. Statistical significance was evaluated by comparing the Elkan-Noto algorithm, AlphaMax, AlphaMax-N, and the multi-sample GMM after applying a multivariate-to-univariate transform (MSGMM-T). The bold font type indicates the winner and the asterisk indicates statistical significance. For each data set, shown are the true mixing proportion (α), true proportion of the positives in the labeled sample (β), sample dimensionality (d), the number of positive examples (n_1), the total number of examples (n), and the area under the ROC curve (AUC) for a model trained between labeled and unlabeled data.

Data	α	β	AUC	d	n_1	n	Elkan-Noto	AlphaMax	AlphaMax-N	MSGMM-T
Bank	0.095	1.00	0.842	13	5188	45000	0.241	0.070	0.037*	0.163
	0.096	0.95	0.819	13	5188	45000	0.284	0.079	0.036*	0.155
	0.101	0.75	0.744	13	5188	45000	0.443	0.124	0.040*	0.127
Concrete	0.419	1.00	0.685	8	490	1030	0.329	0.141	0.181	0.077*
	0.425	0.95	0.662	8	490	1030	0.363	0.174	0.231	0.095*
	0.446	0.75	0.567	8	490	1030	0.531	0.212	0.272	0.233
Gas	0.342	1.00	0.825	127	2565	5574	0.017	0.011	0.017	0.008*
	0.353	0.95	0.795	127	2565	5574	0.078	0.016	0.006	0.006
	0.397	0.75	0.672	127	2565	5574	0.396	0.137	0.009	0.006*
Housing	0.268	1.00	0.810	13	209	506	0.159	0.087	0.094	0.209
	0.281	0.95	0.777	13	209	506	0.226	0.094	0.110	0.204
	0.330	0.75	0.651	13	209	506	0.501	0.125	0.134	0.172
Landsat	0.093	1.00	0.933	36	1508	6435	0.074	0.009	0.007*	0.157
	0.103	0.95	0.904	36	1508	6435	0.110	0.015	0.008*	0.152
	0.139	0.75	0.788	36	1508	6435	0.302	0.063	0.012*	0.143
Mushroom	0.409	1.00	0.792	126	3916	8124	0.029	0.015*	0.022	0.037
	0.416	0.95	0.766	126	3916	8124	0.087	0.015	0.008*	0.037
	0.444	0.75	0.648	126	3916	8124	0.370	0.140	0.020	0.024
Pageblock	0.086	1.00	0.885	10	560	5473	0.116	0.026*	0.044	0.129
	0.087	0.95	0.858	10	560	5473	0.137	0.031*	0.052	0.125
	0.090	0.75	0.768	10	560	5473	0.256	0.041*	0.064	0.111
Pendigit	0.243	1.00	0.875	16	3430	10992	0.030	0.006*	0.009	0.081
	0.248	0.95	0.847	16	3430	10992	0.071	0.011	0.005*	0.074
	0.268	0.75	0.738	16	3430	10992	0.281	0.093	0.007*	0.062
Pima	0.251	1.00	0.735	8	268	768	0.351	0.120	0.111	0.171
	0.259	0.95	0.710	8	268	768	0.408	0.118	0.110	0.168
	0.289	0.75	0.623	8	268	768	0.586	0.144	0.156	0.175
Shuttle	0.139	1.00	0.929	9	8903	58000	0.024*	0.027	0.029	0.157
	0.140	0.95	0.903	9	8903	58000	0.052	0.004*	0.007	0.157
	0.143	0.75	0.802	9	8903	58000	0.199	0.047	0.004*	0.148
Spambase	0.226	1.00	0.842	57	1813	4601	0.184	0.046	0.041	0.059
	0.240	0.95	0.812	57	1813	4601	0.246	0.059	0.042*	0.063
	0.295	0.75	0.695	57	1813	4601	0.515	0.155	0.044*	0.059
Wine	0.566	1.00	0.626	11	4113	6497	0.290	0.083	0.060	0.070
	0.575	0.95	0.610	11	4113	6497	0.322	0.113	0.063	0.076
	0.612	0.75	0.531	11	4113	6497	0.420	0.322	0.353	0.293

between samples. They do not explicitly address issues with high-dimensional data nor focus on algorithms to obtain κ^* . In contrast, we focus primarily on the univariate transformation to handle high-dimensional data and practical algorithms for estimating α^* . Supervised learning used for class prior-preserving transformation provides a rich set of techniques to address high-dimensional data.

8 Conclusion

In this paper, we developed a practical algorithm for classification of positive-unlabeled data with noise in the labeled data set. In particular, we focused on a strategy for high-dimensional data, providing a univariate transform that reduces the dimension of the data, preserves the class prior so that estimation in this reduced space remains valid and is then further useful for classification. This approach provides a simple algorithm that simultaneously improves estimation of the class prior and provides a resulting classifier. We derived a parametric and a nonparametric version of the algorithm and then evaluated its performance on a wide variety of learning scenarios and data sets. To the best of our knowledge, this algorithm represents one of the first practical and easy-to-use approaches to learning with high-dimensional positive-unlabeled data with noise in the labels.

Acknowledgements

We thank Prof. Michael W. Trosset for helpful comments. Grant support: NSF DBI-1458477, NIH R01MH105524, NIH R01GM103725, and the Indiana University Precision Health Initiative.

References

- S. Bashir and E. M. Carter. High breakdown mixture discriminant analysis. *J Multivar Anal*, 93(1):102–111, 2005.
- G. Blanchard, G. Lee, and C. Scott. Semi-supervised novelty detection. *J Mach Learn Res*, 11:2973–3009, 2010.
- A. Blum and T. Mitchell. Combining labeled and unlabeled data with co-training. COLT 1998, pages 92–100, 1998.
- C. Bouveyron and S. Girard. Robust supervised classification with mixture models: learning from data with uncertain labels. *Pattern Recognit*, 42(11):2649–2658, 2009.
- C. Cortes, M. Mohri, M. Riley, and A. Rostamizadeh. Sample selection bias correction theory. ALT 2008, pages 38–53, 2008.
- F. Denis, R. Gilleron, and F. Letouzey. Learning from positive and unlabeled examples. *Theor Comput Sci*, 348(16):70–83, 2005.
- M. C. du Plessis and M. Sugiyama. Class prior estimation from positive and unlabeled data. *IEICE Trans Inf & Syst*, E97-D(5):1358–1362, 2014.
- C. Elkan and K. Noto. Learning classifiers from only positive and unlabeled data. KDD 2008, pages 213–220, 2008.
- D. M. Hawkins and G. J. McLachlan. High-breakdown linear discriminant analysis. *J Am Stat Assoc*, 92(437):136–143, 1997.
- S. Jain, M. White, M. W. Trosset, and P. Radivojac. Nonparametric semi-supervised learning of class proportions. *arXiv preprint arXiv:1601.01944*, 2016. URL <http://arxiv.org/abs/1601.01944>.
- J. Katz-Samuels and C. Scott. A mutual contamination analysis of mixed membership and partial label models. *arXiv preprint arXiv:1602.06235*, 2016. URL <http://arxiv.org/abs/1602.06235>.
- N. D. Lawrence and B. Scholkopf. Estimating a kernel Fisher discriminant in the presence of label noise. ICML 2001, pages 306–313, 2001.
- M. Lichman. UCI Machine Learning Repository, 2013. URL <http://archive.ics.uci.edu/ml>.
- H. Liu, J. D. Lafferty, and L. A. Wasserman. Sparse nonparametric density estimation in high dimensions using the rodeo. AISTATS 2007, pages 283–290, 2007.
- P. M. Long and R. A. Servedio. Random classification noise defeats all convex potential boosters. *Mach Learn*, 78(3):287–304, 2010.
- N. Manwani and P. S. Sastry. Noise tolerance under risk minimization. *IEEE T Cybern*, 43(3):1146–1151, 2013.
- A. K. Menon, B. van Rooyen, C. S. Ong, and R. C. Williamson. Learning from corrupted binary labels via class-probability estimation. ICML 2015, pages 125–134, 2015.
- A. Niculescu-Mizil and R. Caruana. Obtaining calibrated probabilities from boosting. UAI 2005, pages 413–420, 2005.
- J. C. Platt. *Probabilistic outputs for support vector machines and comparison to regularized likelihood methods*, pages 61–74. MIT Press, 1999.
- H. G. Ramaswamy, C. Scott, and A. Tewari. Mixture proportion estimation via kernel embedding of distributions. *arXiv preprint arXiv:1603.02501*, 2016. URL <https://arxiv.org/abs/1603.02501>.
- M. D. Reid and R. C. Williamson. Composite binary losses. *J Mach Learn Res*, 11:2387–2422, 2010.
- M. Saerens, P. Latinne, and C. Decaestecker. Adjusting the outputs of a classifier to new a priori probabilities: a simple procedure. *Neural Comput*, 14:21–41, 2002.
- T. Sanderson and C. Scott. Class proportion estimation with application to multiclass anomaly rejection. AISTATS 2014, pages 850–858, 2014.
- C. Scott, G. Blanchard, and G. Handy. Classification with asymmetric label noise: consistency and maximal denoising. *J Mach Learn Res W&CP*, 30:489–511, 2013.
- D. W. Scott. The curse of dimensionality and dimension reduction. *Multivariate Density Estimation: Theory, Practice, and Visualization*, pages 195–217, 2008.
- H. Steen and M. Mann. The ABC’s (and XYZ’s) of peptide sequencing. *Nat Rev Mol Cell Biol*, 5(9):699–711, 2004.
- G. Ward, T. Hastie, S. Barry, J. Elith, and J.R. Leathwick. Presence-only data and the EM algorithm. *Biometrics*, 65(2):554–563, 2009.

Appendix

A Identifiability proofs

We will need the following Lemma for the proofs.

Lemma A.1 *Let μ, μ_1 and μ_0 be three measures defined on \mathcal{A} such that $\mu = \alpha\mu_1 + (1 - \alpha)\mu_0$ for some $\alpha \in [0, 1]$. Define*

$$R(\mu, \mu_1) = \{\mu(A)/\mu_1(A) : A \in \mathcal{A}, \mu_1(A) > 0\}.$$

It follows that

1. $a_\mu^{\mu_1} = \inf R(\mu, \mu_1)$.
2. If $a_{\mu_0}^{\mu_1} = 0$, then $\alpha = a_\mu^{\mu_1}$.

Proof: The proof follows from Lemma 4 and Theorem 3 of Jain et al. (2016). ■

As a consequence of Statement 1 in Lemma A.1, we use an alternate characterization of $a_\lambda^{\lambda_1}$

$$a_\lambda^{\lambda_1} = \inf R(\lambda, \lambda_1) \tag{13}$$

where $R(\lambda, \lambda_1) = \{\lambda(A)/\lambda_1(A) : A \in \mathcal{A}, \lambda_1(A) > 0\}$.

Lemma 1 (Unidentifiability) *Given a pair of mixtures $(\mu, \nu) \in \mathcal{F}(\Pi^{\text{all}})$, let parameters $(\alpha, \beta, \mu_0, \mu_1)$ generate (μ, ν) and $\alpha^+ = a_\mu^\nu, \beta^+ = a_\nu^\mu$. It follows that*

1. *There is a one-to-one relation between (μ_0, μ_1) and (α, β) and*

$$\mu_0 = \frac{\beta\mu - \alpha\nu}{\beta - \alpha}, \quad \mu_1 = \frac{(1 - \alpha)\nu - (1 - \beta)\mu}{\beta - \alpha}. \tag{6}$$

2. *Both expressions on the right-hand side of Equation 6 are well defined probability measures if and only if $\alpha/\beta \leq \alpha^+$ and $(1-\beta)/(1-\alpha) \leq \beta^+$.*
3. $A^+(\mu, \nu, \Pi^{\text{all}}) = \{(\alpha, \beta) : \alpha/\beta \leq \alpha^+, (1-\beta)/(1-\alpha) \leq \beta^+\}$.
4. $\mathcal{F}(\Pi^{\text{all}})$ is unidentifiable in (α, β) ; i.e., (α, β) is not uniquely determined from (μ, ν) .
5. $\mathcal{F}(\Pi^{\text{all}})$ is unidentifiable in α and β , individually; i.e., neither α nor β is uniquely determined from (μ, ν) .

Proof:

Statement 1: Because $(\alpha, \beta, \mu_0, \mu_1)$ generate μ and ν , $\beta > \alpha$ and Equations 3 and 4 hold. μ_0 can be expressed in terms of μ, ν, α, β by eliminating μ_1 ($\beta \times$ (Equation 3) $- \alpha \times$ (Equation 4)). Similarly μ_1 can be expressed in terms of μ, ν, α, β and Equation 6 follows. Thus, given μ and ν , μ_0, μ_1 can be uniquely determined from α, β . Equation 3 implies that for all $A \in \mathcal{A}$ such that $\mu_1(A) - \mu_0(A) > 0$, $\alpha = (\mu(A) - \mu_1(A))/(\mu_1(A) - \mu_0(A))$. Existence of A is guaranteed because $\mu_1 \neq \mu_0$ from the definition of Π^{all} . Thus, given μ , α is uniquely determined from μ_0, μ_1 . Similarly, Equation 4 implies that, given ν , β is uniquely determined from μ_0, μ_1 .

Statement 2: It is easy to observe that $(\beta\mu - \alpha\nu)/(\beta - \alpha)$ is a valid probability measure, if and only if $\beta\mu(A) - \alpha\nu(A) \geq 0$ for all $A \in \mathcal{A}$. Because the inequality is trivially true when $\nu(A) = 0$, the necessary and sufficient condition can be reduced to $\beta\mu(A) - \alpha\nu(A) \geq 0$ for all $A \in \mathcal{A}$ with $\nu(A) > 0$. This can be rewritten as $\alpha/\beta \leq \mu(A)/\nu(A)$ for all $A \in \mathcal{A}$ such that $\nu(A) > 0$. Equivalently, α/β is a lower bound to $R(\mu, \nu)$; in other words, $\alpha/\beta \leq \alpha^+$ due to the alternate characterization given by equation 13. Similarly, $((1-\alpha)\nu - (1-\beta)\mu)/(\beta - \alpha)$ is a valid probability measure provided $(1-\beta)/(1-\alpha) \leq \beta^+$.

Statement 3: Because system of equations in 6 is essentially equivalent to equations 3 and 4, $(\alpha, \beta) \in A^+(\mu, \nu, \Pi^{\text{all}})$ if and only if, μ_0 and μ_1 as defined in equations 6 come from Π^{all} . When the conditions of statement 2 are not satisfied, μ_0, μ_1 are not well-defined probability measures and

consequently $(\mu_0, \mu_1) \notin \Pi^{\text{all}}$. On the other hand, when they are satisfied, μ_0, μ_1 are well-defined probability measures. Moreover, $\mu \neq \nu$ implies $\mu_0 \neq \mu_1$ and consequently $(\mu_0, \mu_1) \in \Pi^{\text{all}}$.

Statement 5: We construct mixtures $\ddot{\mu}$ and $\ddot{\nu}$ such that there are multiple values for α that generate the mixtures. For $0 < \ddot{\alpha} < \ddot{\beta} < 1$ and probability measures $\ddot{\mu}_0, \ddot{\mu}_1$ on \mathcal{A} , let $(\ddot{\alpha}, \ddot{\beta}, \ddot{\mu}_0, \ddot{\mu}_1)$ generate mixtures $\ddot{\mu}$ and $\ddot{\nu}$. Thus, $(\ddot{\alpha}, \ddot{\beta}) \in A^+(\ddot{\mu}, \ddot{\nu}, \Pi^{\text{all}})$. It follows that $\ddot{\alpha}^+ (= a_{\ddot{\mu}}^{\ddot{\nu}})$ and $\ddot{\beta}^+ (= a_{\ddot{\nu}}^{\ddot{\mu}})$ are both strictly greater than 0 because $\ddot{\alpha}^+ = 0$ implies $\ddot{\alpha} = 0$ and $\ddot{\beta}^+ = 0$ implies $\ddot{\beta} = 1$ (using statement 3), which contradicts our assumption. Now, $(0, 1) \in A^+(\ddot{\mu}, \ddot{\nu}, \Pi^{\text{all}})$ follows trivially from statement 3. It is easy to observe that $(\alpha, 1) \in A^+(\ddot{\mu}, \ddot{\nu}, \Pi^{\text{all}})$ for any $\alpha \in [0, \ddot{\alpha}^+]$. Thus, there are multiple values for α that generate $\ddot{\mu}$ and $\ddot{\nu}$. Similarly, $(0, \beta) \in A^+(\ddot{\mu}, \ddot{\nu}, \Pi^{\text{all}})$ for any $\beta \in [1 - \ddot{\beta}^+, 1]$ and there are multiple values for β that generate $\ddot{\mu}$ and $\ddot{\nu}$.

Statement 4: Statement 4 follows trivially from statement 5 and also by observing that $A^+(\ddot{\mu}, \ddot{\nu}, \Pi^{\text{all}})$ cannot be a singleton set because $\ddot{\alpha}^+, \ddot{\beta}^+ > 0$. ■

Theorem 1 (Identifiability) Given $(\mu, \nu) \in \mathcal{F}(\Pi^{\text{all}})$, let $\alpha^+ = a_{\mu}^{\nu}$ and $\beta^+ = a_{\nu}^{\mu}$. Let $\mu_0^* = (\mu - \alpha^+ \nu) / (1 - \alpha^+)$, $\mu_1^* = (\nu - \beta^+ \mu) / (1 - \beta^+)$ and

$$\alpha^* = \alpha^+ (1 - \beta^+) / (1 - \alpha^+ \beta^+), \quad \beta^* = (1 - \beta^+) / (1 - \alpha^+ \beta^+). \quad (7)$$

It follows that

1. $(\alpha^*, \beta^*, \mu_0^*, \mu_1^*)$ generate (μ, ν)
2. $(\mu_0^*, \mu_1^*) \in \Pi^{\text{res}}$ and, consequently, $\alpha^* = a_{\mu_0^*}^{\mu_1^*}$, $\beta^* = a_{\nu_1^*}^{\mu_0^*}$.
3. $\mathcal{F}(\Pi^{\text{res}})$ contains all pairs of mixtures in $\mathcal{F}(\Pi^{\text{all}})$.
4. $A^+(\mu, \nu, \Pi^{\text{res}}) = \{(\alpha^*, \beta^*)\}$.
5. $\mathcal{F}(\Pi^{\text{res}})$ is identifiable in (α, β) ; i.e., (α, β) is uniquely determined from (μ, ν) .

Proof:

Statement 1: First, we show that μ_0^*, μ_1^* are well defined probability measures.

$\alpha^+ \neq 1, \beta^+ \neq 1$: Suppose $\alpha^+ = 1$. It follows that $\mu = \nu$ and, consequently, from Equation 6, $\mu_0 = \mu_1 = \nu$. However, $\mu_0 \neq \mu_1$ because they are picked from Π^{all} . Thus, $\alpha^+ \neq 1$ by contradiction. Similarly $\beta^+ \neq 1$. Thus, the denominator in the R.H.S. of μ_0^* and μ_1^* is not 0.

By definition $\alpha^+ = \inf R(\mu, \nu)$. Thus, $\alpha^+ \leq \mu(A) / \nu(A)$ when $\nu(A) > 0$. Consequently, $\mu(A) - \alpha^+ \nu(A) \geq 0$. The inequality is trivially true when $\nu(A) = 0$. Thus, $\mu(A) - \alpha^+ \nu(A) \geq 0$ for all $A \in \mathcal{A}$. Hence, μ_0^* is a probability measure. Similarly, μ_1^* is also a probability measure.

Second, we show that $(\alpha^*, \beta^*, \mu_0^*, \mu_1^*)$ generate μ, ν .

$(\alpha^*, \beta^*, \mu_0^*, \mu_1^*) \rightarrow (\mu, \nu)$: Observe that μ_0^*, μ_1^* can also be expressed as $\mu_0^* = \frac{\beta^* \mu - \alpha^* \nu}{\beta^* - \alpha^*}$ and $\mu_1^* = \frac{(1 - \alpha^*) \nu - (1 - \beta^*) \mu}{\beta^* - \alpha^*}$. Moreover, after some algebraic manipulation of Equations labeled 7, $\alpha^* / \beta^* = \alpha^+$ and $(1 - \beta^*) / (1 - \alpha^*) = \beta^+$ can be derived. Thus, from Lemma 1 (statements 1 and 3), $(\alpha^*, \beta^*, \mu_0^*, \mu_1^*)$ generate μ, ν .

Statement 2: We show that $a_{\mu_1^*}^{\mu_0^*} = 0$ and $a_{\mu_0^*}^{\mu_1^*} = 0$, giving $(\mu_0^*, \mu_1^*) \in \Pi^{\text{res}}$.

We start by showing that $a_{\mu_1^*}^{\mu_0^*} = 0$. Suppose, for some $\epsilon > 0$ and

$$\begin{aligned} \nu(A) - \beta^+ \mu(A) &\geq \epsilon (\mu(A) - \alpha^+ \nu(A)) && \text{(for all } A \in \mathcal{A}) \\ \Rightarrow (1 + \epsilon \alpha^+) \nu(A) &\geq (\epsilon + \beta^+) \mu(A) \\ \Rightarrow \nu(A) / \mu(A) &\geq (\epsilon + \beta^+) / (1 + \epsilon \alpha^+) && \text{(when } \mu(A) > 0) \end{aligned}$$

Thus, $(\epsilon + \beta^+) / (1 + \epsilon \alpha^+)$ is a lower bound to $R(\nu, \mu)$ and, consequently, $\beta^+ \geq (\epsilon + \beta^+) / (1 + \epsilon \alpha^+)$. However, because $\alpha^+ \beta^+ < 1$, $(\epsilon + \beta^+) / (1 + \epsilon \alpha^+) > (\epsilon \alpha^+ \beta^+ + \beta^+) / (1 + \epsilon \alpha^+) = \beta^+$. This is a contradiction. Thus, for all $\epsilon > 0$ there exists some $A_\epsilon \in \mathcal{A}$ such that $\nu(A_\epsilon) - \beta^+ \mu(A_\epsilon) < \epsilon (\mu(A_\epsilon) - \alpha^+ \nu(A_\epsilon))$. We now divide the inequality on both sides by $\mu(A_\epsilon) - \alpha^+ \nu(A_\epsilon)$, observing that the divisor is strictly greater than 0 because $\nu(A_\epsilon) - \beta^+ \mu(A_\epsilon) \geq 0$ as μ_1^* is a probability measure. Thus,

$$\begin{aligned} (\nu(A_\epsilon) - \beta^+ \mu(A_\epsilon)) / (\mu(A_\epsilon) - \alpha^+ \nu(A_\epsilon)) &< \epsilon \\ \Rightarrow \mu_1^*(A_\epsilon) / \mu_0^*(A_\epsilon) &< \epsilon (1 - \alpha^+) / (1 - \beta^+). \end{aligned}$$

Because the choice of ϵ is arbitrary and $\beta^+ \neq 1$, any lower bound of $R(\mu_1^*, \mu_0^*)$ cannot be greater than 0. Thus, $a_{\mu_1^*}^{\mu_0^*} = 0$.

A similar argument shows that $a_{\mu_0^*}^{\mu_1^*} = 0$. Therefore, $a_{\mu_1^*}^{\mu_0^*} = 0$ and $a_{\mu_0^*}^{\mu_1^*} = 0$; consequently, $(\mu_0^*, \mu_1^*) \in \Pi^{\text{res}}$. Because μ_0^* is not a mixture containing μ_1^* and because $\mu = \alpha^* \mu_1^* + (1 - \alpha^*) \mu_0^*$, from Statement 2 in Lemma A.1 we get that $\alpha^* = a_{\mu_0^*}^{\mu_1^*}$. Similarly, $\beta^* = a_{\mu_1^*}^{\mu_0^*}$.

Statement 3: Because statements 1 and 2 are true for any $(\mu, \nu) \in \mathcal{F}(\Pi^{\text{all}})$, statement 3 is true.

Statement 4: It follows from statements 1 and 2, that $(\alpha^*, \beta^*) \in A^+(\mu, \nu, \Pi^{\text{res}})$. To show that $A^+(\mu, \nu, \Pi^{\text{res}})$ contains no other element, we give a proof by contradiction. Suppose $(\alpha, \beta) \in A^+(\mu, \nu, \Pi^{\text{res}})$ and $(\alpha, \beta) \neq (\alpha^*, \beta^*)$. Let $(\alpha, \beta, \mu_0, \mu_1)$ generate (μ, ν) , for some $(\mu_0, \mu_1) \in \Pi^{\text{res}}$. First, we show that $\alpha^+ = \alpha/\beta$ and $\beta^+ = (1-\beta)/(1-\alpha)$:

$\alpha^+ = \alpha/\beta$: Suppose for some $0 < \epsilon < (\beta - \alpha)/\beta(1 - \beta)$ and all $A \in \mathcal{A}$ where $\nu(A) > 0$,

$$\begin{aligned} \mu(A)/\nu(A) &\geq \alpha/\beta + \epsilon \\ \Rightarrow \frac{\alpha + (1 - \alpha)\mu_0(A)/\mu_1(A)}{\beta + (1 - \beta)\mu_0(A)/\mu_1(A)} &\geq \frac{\alpha + \beta\epsilon}{\beta} \\ \Rightarrow \mu_0(A)/\mu_1(A) &\geq \beta^2\epsilon/(\beta - \alpha - \epsilon\beta(1 - \beta)) \\ \Rightarrow a_{\mu_0}^{\mu_1} &\geq \beta^2\epsilon/(\beta - \alpha - \epsilon\beta(1 - \beta)) > 0 \end{aligned}$$

However, this is a contradiction because $(\mu_0, \mu_1) \in \Pi^{\text{res}}$. Thus, for every $0 < \epsilon < (\beta - \alpha)/\beta(1 - \beta)$ there exists some $A_\epsilon \in \mathcal{A}$ with $\nu(A_\epsilon) > 0$ such that $\mu(A_\epsilon)/\nu(A_\epsilon) < \alpha/\beta + \epsilon$. Thus, $\alpha^+ < \alpha/\beta + \epsilon$, because ϵ can be made arbitrarily small, $\alpha^+ \leq \alpha/\beta$. However, because $A^+(\mu, \nu, \Pi^{\text{res}}) \subseteq A^+(\mu, \nu, \Pi^{\text{all}})$, (α, β) also belongs to $A^+(\mu, \nu, \Pi^{\text{all}})$ and consequently $\alpha^+ \geq \alpha/\beta$ from Lemma 1 (statement 3). Thus, $\alpha^+ = \alpha/\beta$.

$\beta^+ = (1-\beta)/(1-\alpha)$: The proof is similar to $\alpha^+ = \alpha/\beta$ —supposing $\nu(A)/\mu(A) \geq (1-\beta)/(1-\alpha) + \epsilon$ for some $\epsilon > 0$ and all $A \in \mathcal{A}$ with $\mu(A) > 0$, reaching a contradiction and following the subsequent steps.

$\alpha^+ = \alpha/\beta$, $\beta^+ = (1-\beta)/(1-\alpha)$ and Equation 7 implies $\alpha = \alpha^*$ and $\beta = \beta^*$, which contradicts our assumption. Hence, (α^*, β^*) is the only element in $A^+(\mu, \nu, \Pi^{\text{res}})$, which proves statement 4.

Statement 5: This statement follows by observing that $A^+(\mu, \nu, \Pi^{\text{res}})$ is a singleton set. ■

B Proofs for properties of the univariate transform

We first prove Lemma 2, which we can then use to construct a suitable univariate transform.

Lemma 2 *Let X, Y and S be random variables taking values in \mathcal{X}, \mathcal{Y} and \mathcal{S} , respectively, and $X^u = X|S = 0$ and $X^l = X|S = 1$. For measures μ, ν, μ_0, μ_1 , satisfying Equations 3 and 4 and $\mu_1 \neq \mu_0$, let μ, μ_0, μ_1 give the distribution of $X, X|Y = 0$ and $X|Y = 1$, respectively. If X, Y and S satisfy assumptions 8, 9 and 10, then*

1. X is independent of $S = 0$; i.e., $p(x|S = 0) = p(x)$
2. X^u and X^l are distributed according to μ and ν , respectively.
3. $p(S = 2) \neq 0$.

Proof:

Statement 1: Observe that

$$\begin{aligned} p(x|S = 0) &= p(x, Y = 1|S = 0) + p(x, Y = 0|S = 0) \\ &= p(Y = 1|S = 0)p(x|Y = 1, S = 0) + p(Y = 0|S = 0)p(x|Y = 0, S = 0) \\ &= p(Y = 1)p(x|Y = 1) + p(Y = 0)p(x|Y = 0) \quad (\text{from assumptions 8 and 10}) \\ &= p(x). \end{aligned}$$

Thus, X is independent of $S = 0$.

Statement 2: From statement 1, X^u has the same distribution as X , which is μ . Now,

$$\begin{aligned} p(x|S=1) &= p(x, Y=1|S=1) + p(x, Y=0|S=1) \\ &= p(Y=1|S=1)p(x|Y=1, S=1) + p(Y=0|S=1)p(x|Y=0, S=1) \\ &= \beta p(x|Y=1) + (1-\beta)p(x|Y=0). \end{aligned} \quad (\text{from assumptions 9 and 10})$$

Thus, the distribution of $X|S=1$ is $\beta\mu_1 + (1-\beta)\mu_0$, which is ν .

Statement 3: Now,

$$\begin{aligned} p(S=2|x) &= 1 - p(S=0|x) - p(S=1|x) \\ &= 1 - p(S=0) - \frac{p(x|S=1)}{p(x)}p(S=1) \quad (\text{because } S=0 \text{ and } X \text{ are independent}) \end{aligned}$$

The probability $p(S=2|x)$ is independent of x only if $p(x|S=1)/p(x)$ is a constant with respect to x . Let $p(x|S=1)/p(x) = c$, where c is some constant. Integrating over x on both sides gives $\int_{\mathcal{X}} p(x|S=1)dx = c \int_{\mathcal{X}} p(x)dx$. Since both integrals are 1, it follows that $c = 1$. Thus, $p(x|S=1)/p(x) = 1$, which implies $\mu = \nu$; i.e., the labeled and unlabeled samples have the same distribution. However, this implies $\mu_1 = \mu_0$, which contradicts the assumption. Therefore, $S=2$ is not independent of X . \blacksquare

To prove the main result about the preservation properties of the univariate transform, we will make use of the following theorem.

Lemma B.1 (Restatement of Theorem 9 in Jain et al. (2016)) *Let X and X_1 be random variables with densities f and f_1 and measures μ and μ_1 respectively. For $\overline{\mathbb{R}}^+ = \mathbb{R}^+ \cup \{0, \infty\}$ and an abstract space \mathcal{X}_τ , given any one-to-one function $G : \overline{\mathbb{R}}^+ \rightarrow \mathcal{X}_\tau$, define function $\tau : \mathcal{X} \rightarrow \mathcal{X}_\tau$*

$$\tau = G \circ \tau_d, \quad \left| \begin{array}{l} \text{where} \\ \tau_d(x) = \begin{cases} f(x)/f_1(x) & \text{if } f_1(x) > 0 \\ \infty & \text{if } f_1(x) = 0. \end{cases} \end{array} \right.$$

Let μ_τ and μ_{τ_1} be the measures for the random variables $\tau(X)$, $\tau(X_1)$ respectively for σ -algebra \mathcal{A}_τ on \mathcal{X}_τ . Then $a_\mu^{\mu_1} = a_{\mu_\tau}^{\mu_{\tau_1}}$.

Theorem 2 (α^* -preserving transform) *Let random variables X, Y, S, X^u, X^l and measures μ, ν, μ_0, μ_1 be as defined in Lemma 2. Let τ_p be the posterior as defined in Equation 11 and $\tau = H \circ \tau_p$, where H is a 1-to-1 function on $[0, 1]$ and \circ is the composition operator. Assume*

1. $(\mu_0, \mu_1) \in \Pi^{res}$,
2. X^u and X^l are continuous with densities f and g , respectively,
3. $\mu_\tau, \nu_\tau, \mu_{\tau_1}$ are the measures corresponding to $\tau(X^u), \tau(X^l), \tau(X_1)$, respectively,
4. $(\alpha^+, \beta^+, \alpha^*, \beta^*) = (a_\mu^\nu, a_\nu^\mu, a_\mu^{\mu_1}, a_\nu^{\mu_1})$ and $(\alpha_\tau^+, \beta_\tau^+, \alpha_\tau^*, \beta_\tau^*) = (a_{\mu_\tau}^{\nu_\tau}, a_{\nu_\tau}^{\mu_\tau}, a_{\mu_\tau}^{\mu_{\tau_1}}, a_{\nu_\tau}^{\mu_{\tau_1}})$.

Then

$$(\alpha_\tau^+, \beta_\tau^+, \alpha_\tau^*, \beta_\tau^*) = (\alpha^+, \beta^+, \alpha^*, \beta^*)$$

and so τ is an α^ -preserving transformation.*

Moreover, τ_p can also be used to compute the true posterior probability:

$$p(Y=1|x) = \frac{\alpha^*(1-\alpha^*)}{\beta^* - \alpha^*} \left(\frac{p(S=0)}{p(S=1)} \frac{\tau_p(x)}{1-\tau_p(x)} - \frac{1-\beta^*}{1-\alpha^*} \right). \quad (12)$$

Proof: First we prove that $(\alpha^+, \beta^+) = (\alpha_\tau^+, \beta_\tau^+)$. To this end, we expand $\tau_p(x)$ as follows

$$\begin{aligned}\tau_p(x) &= \frac{p(S=1, x, S \in \{0, 1\})}{p(x, S \in \{0, 1\})} \\ &= \frac{p(S=1, x)}{p(x, S=0) + p(x, S=1)} \\ &= \frac{p(x|S=1)p(S=1)}{p(x|S=0)p(S=0) + p(x|S=1)p(S=1)}\end{aligned}\tag{14}$$

$$\begin{aligned}&= \frac{p(S=1)}{\frac{p(x|S=0)}{p(x|S=1)}p(S=0) + p(S=1)} \\ &= \frac{p(S=1)}{\frac{f(x)}{g(x)}p(S=0) + p(S=1)}\end{aligned}\tag{15}$$

the last step is justified because $f(x) = p(X=x|S=0)$ and $g(x) = p(X=x|S=1)$.

Consider one-to-one functions $G_1(t) = \frac{p(S=1)}{tp(S=0)+p(S=1)}$ and $G_2(t) = \frac{p(S=1)}{\frac{1}{t}p(S=0)+p(S=1)}$ defined on $\mathbb{R}^+ \cup \{0, \infty\}$. We can apply Lemma B.1 to G_1 to get $\alpha^+ = \alpha_\tau^+$ and to G_2 to get $\beta^+ = \beta_\tau^+$. To satisfy the condition of Lemma B.1 for G_1 , let X^u , X^l and $H \circ G_1$ play the role of X , X_1 and G , respectively. Now,

$$\begin{aligned}\tau &= H \circ \tau_p \\ &= H \circ G_1 \circ \tau_d\end{aligned}\tag{from Equation 15}$$

Because H and G_1 are both one-to-one functions, $H \circ G_1$ is one-to-one as well. Thus, all the conditions of Lemma B.1 are satisfied and, consequently, $\alpha^+ = \alpha_\tau^+$.

Similarly, we can use Lemma B.1 with X^u , X^l and $H \circ G_2$ playing the role of X_1 , X and G , respectively, now giving that $\beta^+ = \beta_\tau^+$.

From Statement 2 in Theorem 1 and Equation 7

$$\begin{aligned}\alpha^* &= \alpha^+(1-\beta^+)/ (1-\alpha^+\beta^+) \\ \beta^* &= (1-\beta^+)/ (1-\alpha^+\beta^+) \\ \alpha_\tau^* &= \alpha_\tau^+(1-\beta_\tau^+)/ (1-\alpha_\tau^+\beta_\tau^+) \\ \beta_\tau^* &= (1-\beta_\tau^+)/ (1-\alpha_\tau^+\beta_\tau^+).\end{aligned}$$

and, thus, $(\alpha_\tau^*, \beta_\tau^*) = (\alpha^*, \beta^*)$.

Next we prove Equation 12. Let $c = p(S=0)/p(S=1)$. Rearranging the terms of Equation 14,

$$\begin{aligned}\frac{\tau(x)}{1-\tau(x)} &= \frac{p(x|S=1)p(S=1)}{p(x|S=0)p(S=0)} \\ &= \frac{p(S=1)}{p(S=0)} \left(\frac{p(x, Y=1|S=1) + p(x, Y=1|Y=0)}{p(x)} \right) \tag{from Lemma 2 statement 1} \\ &= \frac{p(S=1)}{p(S=0)} \left(p(Y=1|S=1) \frac{p(x|Y=1, S=1)}{p(x)} + p(Y=0|S=1) \frac{p(x|Y=0, S=1)}{p(x)} \right) \\ &= \frac{p(S=1)}{p(S=0)} \left(p(Y=1|S=1) \frac{p(x|Y=1)}{p(x)} + p(Y=0|S=1) \frac{p(x|Y=0)}{p(x)} \right) \\ &\tag{from assumption 10} \\ &= \frac{1}{c} \left(\frac{p(Y=1|S=1)}{p(Y=1)} p(Y=1|x) + \frac{p(Y=0|S=1)}{p(Y=0)} p(Y=0|x) \right) \\ &= \frac{1}{c} \left(\frac{\beta^*}{\alpha^*} p(Y=1|x) + \frac{1-\beta^*}{1-\alpha^*} (1-p(Y=1|x)) \right) \\ &= \frac{1}{c} \left(\frac{1-\beta^*}{1-\alpha^*} + \left(\frac{\beta^*}{\alpha^*} - \frac{1-\beta^*}{1-\alpha^*} \right) p(Y=1|x) \right).\end{aligned}$$

Rearranging the terms,

$$p(Y = 1|x) = \frac{\alpha^*(1 - \alpha^*)}{\beta^* - \alpha^*} \left(\frac{c\tau(x)}{1 - \tau(x)} - \frac{1 - \beta^*}{1 - \alpha^*} \right). \quad \blacksquare$$

C MSGMM

Let $U = \{X_i^u\}$ and $L = \{X_i^l\}$ be the unlabeled sample and the noisy positive sample, respectively. The parametric approach is derived by modeling each sample as a two component Gaussian mixture, sharing the same components but having different mixing proportions:

$$\begin{aligned} X_i^u &\sim \alpha \mathcal{N}(u_1, \Sigma_1) + (1 - \alpha) \mathcal{N}(u_0, \Sigma_0) \\ X_i^l &\sim \beta \mathcal{N}(u_1, \Sigma_1) + (1 - \beta) \mathcal{N}(u_0, \Sigma_0) \end{aligned}$$

where $u_1, u_0 \in \mathbb{R}^d$ and $\Sigma_1, \Sigma_0 \in \mathbb{S}_{++}^d$, the set of all $d \times d$ positive definite matrices. The algorithm is an extension to the EM approach for Gaussian mixture models (GMMs) where, instead of estimating the parameters of a single mixture, the parameters of both mixtures $(\alpha, \beta, u_0, u_1, \Sigma_0, \Sigma_1)$ are estimated simultaneously by maximizing the combined likelihood over both U and L . This approach, that we refer to as a multi-sample GMM (MSGMM), exploits the constraint that the two mixtures share the same components.

To derive the update equations, we introduce missing variables W_i^u, W_j^l that give the true class of the i^{th} and j^{th} example in U and L , respectively. The variables W_i^u, W_j^l are Bernoulli distributed; i.e., $W_i^u \sim \text{Bernoulli}(\alpha)$ and $W_j^l \sim \text{Bernoulli}(\beta)$. For

$$W = \{W_i^u\}_{i=1}^{|U|}, V = \{W_j^l\}_{j=1}^{|L|}$$

the quartet (U, L, W, V) forms the observed and unobserved variables in the EM framework. The complete data log-likelihood, ll_C is given by,

$$\begin{aligned} ll_C &= \sum_{i=1}^{|U|} W_i^u \log[\alpha \phi_1(x_i^u)] + (1 - W_i^u) \log[(1 - \alpha) \phi_0(x_i^u)] \\ &\quad + \sum_{i=1}^{|L|} W_i^l \log[\beta \phi_1(x_i^l)] + (1 - W_i^l) \log[(1 - \beta) \phi_0(x_i^l)], \end{aligned}$$

where ϕ_i is the density of $\mathcal{N}(u_i, \Sigma_i)$. Our goal is to maximize $E[ll_C]$. To do so we take the conditional expectation of ll_C with respect to W and V given U and L . For

$$\begin{aligned} \bar{w}_i^u &= E[W_i^u | X_i^u = x_i^u] = \frac{\alpha \phi_1(x_i^u)}{\alpha \phi_1(x_i^u) + (1 - \alpha) \phi_0(x_i^u)}, \\ \bar{w}_i^l &= E[W_i^l | X_i^l = x_i^l] = \frac{\beta \phi_1(x_i^l)}{\beta \phi_1(x_i^l) + (1 - \beta) \phi_0(x_i^l)}, \end{aligned}$$

we obtain

$$\begin{aligned} E[ll_C] &= \sum_{i=1}^{|U|} \bar{w}_i^u \log[\alpha \phi_1(x_i^u)] + (1 - \bar{w}_i^u) \log[(1 - \alpha) \phi_0(x_i^u)] \\ &\quad + \sum_{i=1}^{|L|} \bar{w}_i^l \log[\beta \phi_1(x_i^l)] + (1 - \bar{w}_i^l) \log[(1 - \beta) \phi_0(x_i^l)] \end{aligned}$$

which up to constants, that are ignored in the optimization, can be explicitly written as

$$\begin{aligned} E[ll_C] &= \sum_{i=1}^{|U|} \bar{w}_i^u \left[\log \alpha - \frac{1}{2} \log |\Sigma_1| - (x_i^u - u_1)^T \Sigma_1^{-1} (x_i^u - u_1) \right] \\ &\quad + \sum_{i=1}^{|U|} (1 - \bar{w}_i^u) \left[\log(1 - \alpha) - \frac{1}{2} \log |\Sigma_0| - (x_i^u - u_0)^T \Sigma_0^{-1} (x_i^u - u_0) \right] \\ &\quad + \sum_{i=1}^{|L|} \bar{w}_i^l \left[\log \beta - \frac{1}{2} \log |\Sigma_1| - (x_i^l - u_1)^T \Sigma_1^{-1} (x_i^l - u_1) \right] \\ &\quad + \sum_{i=1}^{|L|} (1 - \bar{w}_i^l) \left[\log(1 - \beta) - \frac{1}{2} \log |\Sigma_0| - (x_i^l - u_0)^T \Sigma_0^{-1} (x_i^l - u_0) \right]. \end{aligned}$$

Finally, we obtain the parameter update equations by maximizing $E[ll_C]$ with respect to $(\alpha, \beta, u_0, u_1, \Sigma_0, \Sigma_1)$:

$$\begin{aligned}\alpha &\leftarrow 1/|U| \sum_{i=1}^{|U|} \bar{w}_i^u \\ \beta &\leftarrow 1/|L| \sum_{j=1}^{|L|} \bar{w}_j^l \\ u_1 &\leftarrow \frac{\sum_{i=1}^{|U|} \bar{w}_i^u x_i^u + \sum_{j=1}^{|L|} \bar{w}_j^l x_j^l}{\sum_{i=1}^{|U|} \bar{w}_i^u + \sum_{j=1}^{|L|} \bar{w}_j^l} \\ u_0 &\leftarrow \frac{\sum_{i=1}^{|U|} (1 - \bar{w}_i^u) x_i^u + \sum_{j=1}^{|L|} (1 - \bar{w}_j^l) x_j^l}{\sum_{i=1}^{|U|} (1 - \bar{w}_i^u) + \sum_{j=1}^{|L|} (1 - \bar{w}_j^l)} \\ \Sigma_1 &\leftarrow \frac{\sum_{i=1}^{|U|} \bar{w}_i^u (x_i^u - u_1)(x_i^u - u_1)^T + \sum_{j=1}^{|L|} \bar{w}_j^l (x_j^l - u_1)(x_j^l - u_1)^T}{\sum_{i=1}^{|U|} \bar{w}_i^u + \sum_{j=1}^{|L|} \bar{w}_j^l} \\ \Sigma_0 &\leftarrow \frac{\sum_{i=1}^{|U|} (1 - \bar{w}_i^u) (x_i^u - u_0)(x_i^u - u_0)^T + \sum_{j=1}^{|L|} (1 - \bar{w}_j^l) (x_j^l - u_0)(x_j^l - u_0)^T}{\sum_{i=1}^{|U|} (1 - \bar{w}_i^u) + \sum_{j=1}^{|L|} (1 - \bar{w}_j^l)}\end{aligned}$$

The update rules reduce to the standard GMM when the labeled sample is not provided. Further generalization to more than two samples and/or mixing components is straightforward.

D AlphaMax

For a mixture sample M and a component sample C , AlphaMax(M, C) estimates the maximum proportion of C in M (Jain et al., 2016). AlphaMax is based on the constrained maximization of the log likelihood of samples M and C , derived using nonparametric estimates of their densities m and c , respectively. We list the main steps of AlphaMax below.

1. Estimate c nonparametrically as \hat{c} using sample C . Obtain the weights, v_i , and components κ_i from nonparametric density estimation of m as a k -component mixture, $\hat{m}(x) = \sum_{i=1}^k v_i \kappa_i(x)$, using M .
2. Construct two density functions $\tilde{c}(\cdot|\omega)$ and $\tilde{m}(\cdot|\omega)$ from v_i, κ_i and \hat{c} parameterized by a k -dimensional weight vector $\omega = [\omega_i], 0 \leq \omega_i \leq 1$, which re-weights components κ_i :

$$\begin{aligned}\tilde{c}(x|\omega) &= \frac{\sum_{i=1}^k \omega_i v_i \kappa_i(x)}{\sum_{i=1}^k \omega_i v_i}, \\ \tilde{m}(x|\omega) &= \left(\sum_{i=1}^k \omega_i v_i \right) \hat{c}(x) + \left(1 - \sum_{i=1}^k \omega_i v_i \right) \tilde{c}(x|1 - \omega); \end{aligned}$$

3. Maximize the log likelihood of M and C constructed with \tilde{m} and \tilde{c} under the constraint $\sum_{i=1}^k \omega_i v_i = r$ for many values of r equispaced in $[0, 1]$.

$$\begin{aligned}ll_r &= \max_{w.r.t. \omega} \sum_{x \in M} \log \tilde{m}(x|\omega) + \sum_{x \in C} \log \tilde{c}(x|\omega), \\ \text{subject to} & \quad \sum_{i=1}^k \omega_i v_i = r, \\ & \quad 0 \leq \omega_i \leq 1, i = 1, \dots, k.\end{aligned}$$

4. Estimate the maximum proportion of c in m , a_m^c (minor abuse of notation¹), as the x -coordinate of the elbow in the ll_r versus r graph.

The densities $\tilde{m}(\cdot|\omega)$ and $\tilde{c}(\cdot|\omega)$ are constructed to approximate m and c . The efficacy of the approximation depends on the value of ω ; there exists ω such that $\tilde{m}(\cdot|\omega)$ and $\tilde{c}(\cdot|\omega)$ are good approximations provided $\sum_{i=1}^k \omega_i v_i \leq a_m^c$, however, the approximation deteriorates progressively, even

¹ a_m^c is a minor abuse of notation; replacing the densities with the underlying measures makes the notation correct.

with the optimum ω , as $\sum_{i=1}^k \omega_i v_i$ moves beyond a_m^c . This suggests that the graph of ll_r versus r should be approximately a flat line from 0 to a_m^c and decrease progressively beyond a_m^c exposing an elbow at a_m^c , which is detected in the last step. The pseudo code for elbow detection is provided in (Jain et al., 2016).

The practical implementation, backed by the α^* -preservation theory, reduces the dimension of the data to a single dimension by using the scoring function of a non-traditional classifier and employs histograms as the nonparametric method to obtain \hat{n} and \hat{c} . The bin-width is chosen to cover the component sample’s (after the transformation) range and reveal the shape of its distribution, using the default option in Matlab’s `histogram` function. More bins with the same bin-width are subsequently added to cover the mixture sample’s range.

The total computation includes the time to (a) train a classifier (b) perform density estimation in 1D and (c) perform optimization in AlphaMax. The complexity for (a) varies; for neural networks it is $O(n)$ per epoch. For (b) we used k bin histograms, where k can be $O(\log(n))$ to $O(\sqrt{n})$ depending on the bandwidth selection rule, giving $O(nk)$ complexity. For (c), (size k optimization) the computation of the objective and gradient is $O(nk)$ per step; e.g., LBFGS is $O(nk)$ per step. An execution of AlphaMax takes about 10 minutes on a laptop computer for the Shuttle data set (for results reported in Table 1; i.e., 1000 labeled and 10000 unlabeled examples); 12 hours on the entire data (8903 labeled and 49097 unlabeled examples).

E Empirical results

Results for univariate synthetic data. The results for synthetic data, with scalar inputs, are summarized in Table 2.

Results for multivariate AlphaMax-N and MSGMM. To demonstrate the efficacy of the class prior preserving transform, we implemented the multivariate versions of AlphaMax-N and MSGMM and evaluated them on the twelve real data sets without applying the transform. There were significant stability and computational issues related to the high-dimensional nature of the data sets. MSGMM was numerically unstable because of singular/nearly-singular covariance matrix, whereas AlphaMax-N became computationally demanding because the number of bins (for histogram based density estimation) grow exponentially with the dimension, resulting in a large parameter vector ω and, consequently, a large optimization problem, even after removing the zero-count bins. This is expected, as density estimation for multivariate data is known to be problematic, which is one of the main reasons for introducing our transform. To make estimation feasible under these stability and computational issues, we used dimensionality reduction. Though not all data sets posed the same level of difficulty, to have a standard approach and permit effective density estimation, we used the top three principal components, obtained via principal component analysis on the z-score normalized data (mixture and component samples combined), as input to the two algorithms. We also attempted using top k principal components that preserve 75 percent of the total variance, however, for some of the data sets, the dimension was still too high.

In the same manner as in the univariate case, we used histograms in the multivariate implementation of AlphaMax-N. The bin-width for a dimension was selected to minimize the asymptotic mean integrated squared error (AMISE) with a normal reference rule, using the component sample, C . The formula for the bin-width of dimension k is given by:

$$b_k = 3.5\sigma_k|C|^{-1/(2+d)},$$

where d is the total number of dimensions, σ_k is the standard deviation of the k^{th} dimension and $|C|$ is the size of the component sample. Bins were added to cover the range of the entire data, mixture and component combined, and empty bins were removed to reduce the size of the optimization problem.

Table 3 contains the results of AlphaMax-N and MSGMM on the real-life data sets, using the top three principal components under column headings AlphaMax-NM (M for multi-dimensional) and MSGMM, respectively. The results of AlphaMax-N and MSGMM-T with the class-prior preserving transform are also provided for comparison. Notice that, though AlphaMax-NM (without transform) performs well, AlphaMax-N (with transform) is significantly better in terms of estimation error, despite having a lower computational cost. Also notice the deterioration in the performance of MSGMM (without transform) compared to MSGMM-T (with transform).

Table 2: Mean absolute difference between estimated and true mixing proportion over a selection of true mixing proportions and the following data sets: \mathcal{N} = Gaussian with $\Delta\mu \in \{1, 2, 4\}$, \mathcal{L} = Laplace with $\Delta\mu \in \{1, 2, 4\}$. Statistical significance was evaluated by comparing the Elkan-Noto algorithm, AlphaMax, AlphaMax-N, and the multi-sample Gaussian Mixture Model (MSGMM). The bold font type indicates the winner and the asterisk indicates statistical significance.

Data	α	β	Elkan-Noto		AlphaMax		AlphaMax-N		MSGMM	
			100	1000	100	1000	100	1000	100	1000
\mathcal{N} $\Delta\mu = 1$	0.05	1.00	0.473	0.460	0.079	0.102	0.064	0.085	0.034*	0.020*
	0.25	1.00	0.484	0.435	0.132	0.160	0.123	0.122	0.063*	0.044*
	0.50	1.00	0.395	0.347	0.155	0.125	0.173	0.096	0.073*	0.040*
	0.05	0.95	0.484	0.496	0.099	0.124	0.076	0.099	0.039*	0.022*
	0.25	0.95	0.510	0.469	0.127	0.167	0.115	0.111	0.074*	0.037*
	0.50	0.95	0.433	0.378	0.180	0.152	0.186	0.087	0.068*	0.048*
	0.05	0.75	0.663	0.630	0.124	0.135	0.089	0.076	0.056*	0.024*
	0.25	0.75	0.641	0.608	0.152	0.209	0.141	0.120	0.141	0.060*
	0.50	0.75	0.548	0.485	0.219	0.218	0.244	0.137	0.040*	0.068*
\mathcal{N} $\Delta\mu = 2$	0.05	1.00	0.112	0.136	0.018	0.016	0.017	0.015	0.006*	0.004*
	0.25	1.00	0.177	0.168	0.050	0.049	0.049	0.042	0.018*	0.010*
	0.50	1.00	0.219	0.153	0.104	0.053	0.109	0.043	0.024*	0.015*
	0.05	0.95	0.125	0.162	0.015	0.019	0.015	0.015	0.009*	0.004*
	0.25	0.95	0.212	0.208	0.063	0.061	0.060	0.043	0.025*	0.011*
	0.50	0.95	0.271	0.211	0.099	0.077	0.106	0.046	0.023*	0.018*
	0.05	0.75	0.215	0.285	0.030	0.036	0.022	0.012	0.034	0.005*
	0.25	0.75	0.405	0.403	0.100	0.139	0.084	0.036	0.025*	0.014*
	0.50	0.75	0.457	0.415	0.156	0.184	0.159	0.048	0.030*	0.020*
\mathcal{N} $\Delta\mu = 4$	0.05	1.00	0.010	0.013	0.022	0.019	0.022	0.019	0.001*	0.001*
	0.25	1.00	0.032	0.018	0.022	0.004	0.027	0.008	0.002*	0.002*
	0.50	1.00	0.070	0.020	0.032	0.005	0.039	0.012	0.002*	0.002*
	0.05	0.95	0.018	0.029	0.018	0.014	0.019	0.015	0.001*	0.001*
	0.25	0.95	0.063	0.056	0.010	0.017	0.013	0.009	0.002*	0.001*
	0.50	0.95	0.115	0.083	0.040	0.032	0.045	0.010	0.002*	0.001*
	0.05	0.75	0.062	0.114	0.011	0.006	0.018	0.015	0.001*	0.001*
	0.25	0.75	0.236	0.249	0.063	0.090	0.026	0.005	0.002*	0.002*
	0.50	0.75	0.380	0.353	0.130	0.168	0.106	0.021	0.002*	0.002*
\mathcal{L} $\Delta\mu = 1$	0.05	1.00	0.410	0.389	0.195	0.256	0.147*	0.190*	0.418	0.390
	0.25	1.00	0.410	0.356	0.151	0.209	0.103*	0.117*	0.148	0.183
	0.50	1.00	0.367	0.299	0.195	0.154	0.190	0.038*	0.243	0.236
	0.05	0.95	0.455	0.430	0.204	0.271	0.140*	0.192*	0.424	0.406
	0.25	0.95	0.455	0.401	0.180	0.230	0.115*	0.112*	0.157	0.187
	0.50	0.95	0.412	0.331	0.222	0.179	0.214	0.047*	0.249	0.241
	0.05	0.75	0.593	0.578	0.225	0.325	0.133*	0.181*	0.415	0.440
	0.25	0.75	0.602	0.562	0.191	0.327	0.119*	0.123*	0.210	0.230
	0.50	0.75	0.520	0.470	0.264	0.278	0.272	0.053*	0.254	0.247
\mathcal{L} $\Delta\mu = 2$	0.05	1.00	0.116	0.123	0.052	0.061	0.045	0.056	0.448	0.014*
	0.25	1.00	0.158	0.132	0.054	0.067	0.041	0.051	0.037	0.027*
	0.50	1.00	0.186	0.125	0.077	0.049	0.074	0.021*	0.130	0.068
	0.05	0.95	0.131	0.146	0.053	0.066	0.042*	0.053	0.455	0.016*
	0.25	0.95	0.186	0.175	0.071	0.082	0.049	0.049	0.055	0.021*
	0.50	0.95	0.239	0.183	0.082	0.082	0.073	0.024*	0.134	0.075
	0.05	0.75	0.230	0.268	0.079	0.097	0.046*	0.052*	0.461	0.261
	0.25	0.75	0.375	0.377	0.123	0.159	0.090*	0.039	0.191	0.050
	0.50	0.75	0.450	0.406	0.201	0.212	0.216	0.045	0.101*	0.067
\mathcal{L} $\Delta\mu = 4$	0.05	1.00	0.015	0.020	0.024	0.020	0.025	0.021	0.015	0.011*
	0.25	1.00	0.040	0.025	0.018	0.005	0.022	0.005	0.009*	0.009
	0.50	1.00	0.094	0.027	0.022	0.005	0.028	0.006	0.002*	0.002*
	0.05	0.95	0.024	0.037	0.019	0.017	0.020	0.018	0.022	0.012*
	0.25	0.95	0.074	0.063	0.013	0.021	0.011	0.011	0.009	0.009*
	0.50	0.95	0.141	0.089	0.028	0.034	0.026	0.008	0.002*	0.002*
	0.05	0.75	0.072	0.124	0.008*	0.005*	0.015	0.010	0.034	0.012
	0.25	0.75	0.244	0.258	0.076	0.095	0.033	0.009	0.009*	0.008
	0.50	0.75	0.389	0.355	0.128	0.173	0.106	0.014	0.002*	0.002*

Table 3: Mean absolute difference between estimated and true mixing proportion over twelve data sets from the UCI Machine Learning Repository. Statistical significance was evaluated by comparing AlphaMax-NM, MSGMM (both using top three principal components of the data as input), AlphaMax-N, MSGMM-T (both using class-prior preserving transform). The bold font type indicates the winner and the asterisk indicates statistical significance. For each data set, shown are the true mixing proportion (α), true proportion of the positives in the labeled sample (β) and the percent of the total variance explained by the top three principal components rounded to the nearest integer (%variance).

Data	α	β	%variance	AlphaMax-N	AlphaMax-NM	MSGMM-T	MSGMM
Bank	0.095	1.00	40	0.037	0.028*	0.163	0.528
	0.096	0.95		0.036	0.032	0.155	0.574
	0.101	0.75		0.040	0.047	0.127	0.580
Concrete	0.419	1.00	60	0.181	0.276	0.077	0.020*
	0.425	0.95		0.231	0.269	0.095	0.028*
	0.446	0.75		0.272	0.320	0.233	0.063*
Gas	0.342	1.00	81	0.017	0.030	0.008*	0.585
	0.353	0.95		0.006	0.021	0.006	0.575
	0.397	0.75		0.009	0.064	0.006*	0.533
Housing	0.268	1.00	69	0.094*	0.132	0.209	0.316
	0.281	0.95		0.110	0.110	0.204	0.308
	0.330	0.75		0.134	0.205	0.172	0.283
Landsat	0.093	1.00	89	0.007	0.008	0.157	0.443
	0.103	0.95		0.008*	0.028	0.152	0.298
	0.139	0.75		0.012*	0.053	0.143	0.270
Mushroom	0.409	1.00	24	0.022*	0.075	0.037	0.432
	0.416	0.95		0.008*	0.021	0.037	0.398
	0.444	0.75		0.020	0.050	0.024	0.375
Pageblock	0.086	1.00	72	0.044	0.046	0.129	0.178
	0.087	0.95		0.052	0.040*	0.125	0.178
	0.090	0.75		0.064	0.031*	0.111	0.188
Pendigit	0.243	1.00	65	0.009*	0.071	0.081	0.289
	0.248	0.95		0.005*	0.070	0.074	0.286
	0.268	0.75		0.007*	0.092	0.062	0.260
Pima	0.251	1.00	60	0.111	0.123	0.171	0.299
	0.259	0.95		0.110*	0.156	0.168	0.292
	0.289	0.75		0.156	0.178	0.175	0.286
Shuttle	0.139	1.00	67	0.029*	0.064	0.157	0.232
	0.140	0.95		0.007*	0.055	0.157	0.227
	0.143	0.75		0.004*	0.015	0.148	0.356
Spambase	0.226	1.00	22	0.041	0.034	0.059	0.487
	0.240	0.95		0.042	0.041	0.063	0.485
	0.295	0.75		0.044*	0.072	0.059	0.434
Wine	0.566	1.00	65	0.060	0.258	0.070	0.134
	0.575	0.95		0.063	0.256	0.076	0.126
	0.612	0.75		0.353	0.302	0.293	0.096*