

---

# On the Path to an Ideal ROC Curve: Considering Cost Asymmetry in Learning Classifiers

---

**Francis R. Bach\***

Computer Science Division  
University of California  
Berkeley, CA 94720  
*fbach@cs.berkeley.edu*

**David Heckerman & Eric Horvitz**

Microsoft Research  
Redmond, WA 98052  
*{heckerman,horvitz}@microsoft.com*

## Abstract

Receiver Operating Characteristic (ROC) curves are a standard way to display the performance of a set of binary classifiers for all feasible ratios of the costs associated with false positives and false negatives. For linear classifiers, the set of classifiers is typically obtained by training once, holding constant the estimated slope and then varying the intercept to obtain a parameterized set of classifiers whose performances can be plotted in the ROC plane. In this paper, we consider the alternative of varying the asymmetry of the cost function used for training. We show that the ROC curve obtained by varying the intercept and the asymmetry—and hence the slope—always outperforms the ROC curve obtained by varying only the intercept. In addition, we present a path-following algorithm for the support vector machine (SVM) that can compute efficiently the entire ROC curve, that has the same computational properties as training a single classifier. Finally, we provide a theoretical analysis of the relationship between the asymmetric cost model assumed when training a classifier and the cost model assumed in applying the classifier. In particular, we show that the mismatch between the step function used for testing and its convex upper bounds usually used for training leads to a provable and quantifiable difference around extreme asymmetries.

## 1 INTRODUCTION

Receiver Operating Characteristic (ROC) analysis has seen increasing attention in the recent statistics and

---

\*This work was done during a summer internship at Microsoft Research.

machine-learning literature (Pepe, 2000, Provost and Fawcett, 2001, Flach, 2003). The ROC is a representation of choice for displaying the performance of a classifier when the costs assigned by end users to false positives and false negatives are not known at the time of training. For example, when training a classifier for identifying cases of undesirable unsolicited email, end users may have different preferences about the likelihood of a false negative and false positive. The ROC curve for such a classifier reveals the ratio of false negatives and positives at different probability thresholds for classifying an email message as unsolicited or normal email.

In this paper, we consider linear binary classification of points in an Euclidean space—noting that it can be extended in a straightforward manner to non-linear classification problems by using Mercer kernels (Schölkopf and Smola, 2002). That is, given data  $x \in \mathbb{R}^d$ ,  $d \geq 1$ , we consider classifiers of the form  $f(x) = \text{sign}(w^\top x + b)$ , where  $w \in \mathbb{R}^d$  and  $b \in \mathbb{R}$  are referred to as the *slope* and the *intercept*. To date, ROC curves have been usually constructed by training once, holding constant the estimated slope and varying the intercept to obtain the curve. In this paper, we show that, while the latter procedure appears to be the most practical thing to do, it may lead to classifiers with poor performance in some parts of the ROC curve.

The crux of our approach is that we allow the asymmetry of the cost function to vary—i.e., we vary the ratio of the cost of a false positive and the cost of a false negative. For each value of the ratio, we obtain a different slope and intercept, each optimized for this ratio. In a naive implementation, varying the asymmetry would require a retraining of the classifier for each point of the ROC curve, which would be computationally expensive. In Section 3.1, we present an algorithm that can compute the solution of an SVM (Schölkopf and Smola, 2002) for all possible costs of false positives and false negatives, with the same computational complexity as obtaining the solution for only one cost

function. The algorithm extends to asymmetric costs the algorithm of Hastie et al. (2005) and is based on path-following techniques that take advantage of the piecewise linearity of the path of optimal solutions. In Section 3.2, we show how the path-following algorithm can be used to obtain the best possible ROC curve (in expectation). In particular, by allowing both the asymmetry and the intercept to vary, we can obtain provably better ROC curves than by methods that simply vary the intercept.

In Section 4, we provide a theoretical analysis of the link between the asymmetry of costs assumed in training a classifier and the asymmetry desired in its application. In particular, we show that—even in the population (*i.e.*, infinite sample) case—the use of a training loss function which is a convex upper bound on the true or testing loss function (a step function) creates classifiers with sub-optimal accuracy. We quantify this problem around extreme asymmetries for several classical convex-upper-bound loss functions—the square loss and the *erf* loss, an approximation of the logistic loss based on normal cumulative distribution functions (also referred to as the “error function”, and usually abbreviated as erf). The analysis is carried through for Gaussian and mixture of Gaussian class-conditional distributions (see Section 4 for more details). As we shall see, the consequences of the potential mismatch between the cost functions assumed in testing versus training underscore the value of using the algorithm that we introduce in Section 4.3. Even when costs are known (*i.e.*, when only one point on the ROC curve is needed), the classifier resulting from our approach which builds the entire ROC curve is never less accurate and can be more accurate than one trained with the known costs using a convex-upper-bound loss function.

## 2 PROBLEM OVERVIEW

Given data  $x \in \mathbb{R}^d$  and labels  $y \in \{-1, 1\}$ , we consider linear classifiers of the form  $f(x) = \text{sign}(w^\top x + b)$ , where  $w$  is the *slope* of the classifier and  $b$  the *intercept*. A classifier is determined by the parameters  $(w, b) \in \mathbb{R}^{d+1}$ . In Section 2.1, we introduce notation and definitions; in Section 2.2, we lay out the necessary concepts of ROC analysis. In Section 2.3, we describe how these classifiers and ROC curves are typically obtained from data.

### 2.1 ASYMMETRIC COST AND LOSS FUNCTIONS

Positive (resp. negative) examples are those for which  $y = 1$  (resp.  $y = -1$ ). The two types of misclassification, false positives and false negatives, are assigned

two different costs, and the total *expected* cost is equal to

$$R(C_+, C_-, w, b) = C_+ P\{w^\top x + b < 0, y = 1\} + C_- P\{w^\top x + b \geq 0, y = -1\}$$

If we let  $\phi_{0-1}(u) = 1_{u < 0}$  be the 0-1 loss, we can write the expected cost as

$$R(C_+, C_-, w, b) = C_+ E\{1_{y=1} \phi_{0-1}(w^\top x + b)\} + C_- E\{1_{y=-1} \phi_{0-1}(-w^\top x - b)\}$$

where  $E$  denotes the expectation with respect to the joint distribution of  $(x, y)$ . The expected cost defined using the 0-1 loss is the cost that end users are usually interested in during the use of the classifier, while the other cost functions that we define below are used solely for training purposes. The convexity of these cost functions makes learning algorithms convergent without local minima, and leads to attractive asymptotic properties (Bartlett et al., 2004).

A traditional set-up for learning linear classifiers from labeled data is to consider a convex upper bound  $\phi$  on the 0-1 loss  $\phi_{0-1}$ , and use the expected  $\phi$ -cost :

$$R_\phi(C_+, C_-, w, b) = C_+ E\{1_{y=1} \phi(w^\top x + b)\} + C_- E\{1_{y=-1} \phi(-w^\top x - b)\}$$

We refer to the ratio  $C_+/(C_- + C_+)$  as the *asymmetry*. We shall use *training asymmetry* to refer to the asymmetry used for training a classifier using a  $\phi$ -cost, and the *testing asymmetry* to refer to the asymmetric cost characterizing the testing situation (reflecting end user preferences) with the actual cost based on the 0-1 loss. In Section 4, we will show that these may be different in the general case.

We shall consider several common loss functions. Some of the loss functions (square loss, hinge loss) lead to attractive computational properties, while others (square loss, erf loss) are more amenable to theoretical manipulations (see Figure 1 for the plot of the loss functions, as they are commonly used and defined below<sup>1</sup>):

- **square loss** :  $\phi_{sq}(u) = \frac{1}{2}(u - 1)^2$ ; the classifier is equivalent to linear regression,
- **hinge loss** :  $\phi_{hi}(u) = \max\{1 - u, 0\}$ ; the classifier is the support vector machine (Schölkopf and Smola, 2002),
- **erf loss** :  $\phi_{erf}(u) = 2 \left[ \frac{u}{2} \psi\left(\frac{u}{2}\right) - \frac{u}{2} + \psi'\left(\frac{u}{2}\right) \right]$ , where  $\psi$  is the cumulative distribution of the standard normal distribution, *i.e.* :  $\psi(v) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^v e^{-t^2/2} dt$ , and  $\psi'(v) = \frac{1}{\sqrt{2\pi}} e^{-v^2/2}$ .

<sup>1</sup>Note that by rescaling, all of these loss functions can be made to be an upper bound on the 0-1 loss which is tight at zero.




Figure 1: Loss functions: (left) plain: 0-1 loss, dotted: hinge loss, dashed: erf loss, dash-dotted: square loss. (Right) plain: 0-1 loss, dotted: probit loss, dashed: logistic loss.

The erf loss provides a good approximation of the *logistic loss*  $\log(1 + e^{-u})$  as well as its derivative, and is amenable to closed-form computations for Gaussians and mixture of Gaussians (see Section 4 for more details). Note that the erf loss is different from the *probit loss*  $-\log \psi(u)$ , which leads to probit regression (Hastie et al., 2001).

## 2.2 ROC ANALYSIS

The aim of ROC analysis is to display in a single graph the performance of classifiers for all possible costs of misclassification. In this paper, we consider sets of classifiers  $f_\gamma(x)$ , parameterized by a variable  $\gamma \in \mathbb{R}$  ( $\gamma$  can either be the intercept or the training asymmetry).

For a classifier  $f(x)$ , we can define a point  $(u, v)$  in the “ROC plane,” where  $u$  is the proportion of false positives  $u = P(f(x) = 1 | y = -1)$ , and  $v$  is the proportion of true positives  $v = P(f(x) = 1 | y = 1)$ .

When  $\gamma$  is varied, we obtain a curve in the ROC plane, the ROC curve (see Figure 2 for an example). Whether  $\gamma$  is the intercept or the training asymmetry, the ROC curve always passes through the point  $(0, 0)$  and  $(1, 1)$ , which corresponds to classifying all points as negative (resp. positive).

The upper convex envelope of the curve is the set of optimal ROC points that can be achieved by the set of classifiers; indeed, if a point in the envelope is not one of the original points, it must lie in a segment between two points  $(u(\gamma_0), v(\gamma_0))$  and  $(u(\gamma_1), v(\gamma_1))$ , and all points in a segment between two classifiers can always be attained by choosing randomly between the two classifiers (note that this classifier itself is not a linear classifier; this performance can only be achieved by a mixture of two linear classifiers).

Denoting  $p_+ = P(y = 1)$  and  $p_- = P(y = -1)$ , the expected  $(C_+, C_-)$ -cost for a classifier  $(u, v)$  in the ROC space, is simply  $p_+C_+(1 - v) + p_-C_-u$ , and thus optimal classifiers for the  $(C_+, C_-)$ -cost can be found by looking at lines of slope that are normal to  $(p_-C_-, -p_+C_+)$ —and thus proportional to




Figure 2: (Left) ROC curve: (plain) regular ROC curve; (dashed) convex envelope. The points  $a$  and  $c$  are ROC-consistent and the point  $b$  is ROC-inconsistent. (Right) ROC curve and dashed equi-cost lines: All lines have direction  $(p_+C_+, p_-C_-)$ , the plain line is optimal and the point “a” is the optimal classifier.

$(p_+C_+, p_-C_-)$ —and which intersects the ROC curve and are as close as the point  $(0, 1)$  as possible (see Figure 2).

A point  $(u(\gamma), v(\gamma))$  is said to be *ROC-consistent* if it lies on the upper convex envelope; In this case, the tangent direction  $(du/d\gamma, dv/d\gamma)$  defines a cost  $(C_+(\gamma), C_-(\gamma))$  for which the classifier is optimal (for the testing cost, which is defined using the 0-1 loss), by having  $(p_+C_+(\gamma), p_-C_-(\gamma))$  proportional to  $(du/d\gamma, dv/d\gamma)$ . This leads to an *optimal testing asymmetry*  $\beta(\gamma)$ , defined as  $\beta(\gamma) = \frac{C_+(\gamma)}{C_+(\gamma) + C_-(\gamma)} = \frac{1}{1 + \frac{p_+}{p_-} \frac{dv/d\gamma(\gamma)}{du/d\gamma(\gamma)}}$ .

If a point  $(u(\gamma), v(\gamma))$  is ROC-inconsistent, then the quantity  $\beta(\gamma)$  has no meaning, and such a classifier is generally useless, because, for all settings of the misclassification cost, that classifier can be outperformed by the other classifiers or a combination of classifiers.

In Section 4, we relate the optimal asymmetry of cost in the testing or eventual use of a classifier in the real world, to the asymmetry of cost used to train that classifier; in particular, we show that they differ and quantify this difference for extreme asymmetries (*i.e.*, close to the corner points  $(0, 0)$  and  $(1, 1)$ ). This analysis highlights the value of generating the entire ROC curve, even when only one point is needed, as we will present in Section 4.3.

## 2.3 LEARNING FROM DATA

Given  $n$  labelled data points  $(x_i, y_i)$ , the *empirical cost* is equal to:

$$\widehat{R}(C_+, C_-, w, b) = \frac{C_+}{n} \#\{y_i(w^\top x_i + b) < 0, y_i = 1\} + \frac{C_-}{n} \#\{y_i(w^\top x_i + b) < 0, y_i = -1\}$$

while the *empirical  $\phi$ -cost* is equal to

$$\widehat{R}_\phi(C_+, C_-, w, b) = \frac{C_+}{n} \sum_{i \in \mathcal{I}_+} \phi(y_i(w^\top x_i + b)) + \frac{C_-}{n} \sum_{i \in \mathcal{I}_-} \phi(y_i(w^\top x_i + b)),$$

where  $\mathcal{I}_+ = \{i, y_i = 1\}$  and  $\mathcal{I}_- = \{i, y_i = -1\}$ . When learning a classifier from data, a classical setup is to minimize the sum of the *empirical  $\phi$ -cost* and a regularization term  $\frac{1}{2n}\|w\|^2$ , *i.e.*, to minimize  $\widehat{J}_\phi(C_+, C_-, w, b) = \widehat{R}_\phi(C_+, C_-, w, b) + \frac{1}{2n}\|w^2\|$ .

Note that the objective function is no longer homogeneous in  $(C_+, C_-)$ ; the sum  $C_+ + C_-$  is referred to as the total amount of regularization. Thus, two end-user-defined parameters are needed to train a linear classifier: the *total amount of regularization*  $C_+ + C_- \in \mathbb{R}^+$ , and the *asymmetry*  $\frac{C_+}{C_+ + C_-} \in [0, 1]$ . In Section 3.1, we show how the minimum of  $\widehat{J}_\phi(C_+, C_-, w, b)$ , with respect to  $w$  and  $b$ , can be computed efficiently for the hinge loss, for many values of  $(C_+, C_-)$ , with a computational cost that is within a constant factor of the computational cost of obtaining a solution for one value of  $(C_+, C_-)$ .

**Building an ROC curve from data** If a sufficiently large validation set is available, we can train on the training set and use the empirical distribution of the validation data to plot the ROC curve. If sufficient validation data is not available, then we can use 10 random half splits of the data, train a classifier on one half and use the other half to obtain the ROC scores. Then, for each value of the parameter  $\gamma$  that defines the ROC curve (either the intercept or the training asymmetry), we average the 10 scores. We can also use this approach to obtain confidence intervals (Flach, 2003).

### 3 BUILDING ROC CURVES FOR THE SVM

In this section, we will present an algorithm to compute ROC curves for the SVM that explores the two-dimensional space of cost parameters  $(C_+, C_-)$  efficiently. We first show how to obtain optimal solutions of the SVM without solving the optimization problems many times for each value of  $(C_+, C_-)$ . This method generalizes the results of Hastie et al. (2005) to the case of asymmetric cost functions. We then describe how the space  $(C_+, C_-)$  can be appropriately explored and how ROC curves can be constructed.

#### 3.1 BUILDING PATHS OF CLASSIFIERS

Given  $n$  data points  $x_i, i = 1, \dots, n$  which belong to  $\mathbb{R}^d$ , and  $n$  labels  $y_i \in \{-1, 1\}$ , minimizing the regularized empirical hinge loss is equivalent to solving the following convex optimization problem (Schölkopf and Smola, 2002):

$$\min_{w, b, \xi} \sum_i C_i \xi_i + \frac{1}{2}\|w\|^2 \quad \text{s.t.} \quad \forall i, \xi_i \geq 0, \\ \forall i, \xi_i \geq 1 - y_i(w^\top x_i + b)$$

where  $C_i = C_+$  if  $y_i = 1$  and  $C_i = C_-$  if  $y_i = -1$ .

The dual problem is the following:

$$\max_{\alpha \in \mathbb{R}^n} -\frac{1}{2}\alpha^\top \text{Diag}(y)K \text{Diag}(y)\alpha + 1^\top \alpha \\ \text{s.t.} \quad \alpha^\top y = 0 \text{ and } \forall i, 0 \leq \alpha_i \leq C_i$$

Following Hastie et al. (2005), from the KKT optimality conditions, for an optimal set of primal-dual variables  $(w, b, \alpha)$ , we can separate data points into three disjoint sets:  $\mathcal{M} = \{i, \alpha_i \in [0, C_i], y_i(w^\top x_i + b) = 1\}$ ,  $\mathcal{L} = \{i, \alpha_i = C_i, y_i(w^\top x_i + b) < 1\}$ ,  $\mathcal{R} = \{i, \alpha_i = 0, y_i(w^\top x_i + b) > 1\}$ .

These sets are usually referred to as *active sets* (Boyd and Vandenberghe, 2003). If the sets  $\mathcal{M}$ ,  $\mathcal{L}$  and  $\mathcal{R}$  are known, then it is straightforward to show that, in the domain where the active sets remain constant, the optimal primal-dual variables  $(w, \alpha, b)$  are affine functions of  $(C_+, C_-)$ . This implies that the optimal variables  $(w, \alpha, b)$  are piecewise affine continuous functions of the vector  $(C_+, C_-)$ , with “kinks” where the active sets change.

**Following a path** The active sets remain the same as long as all constraints defining the active sets are satisfied, *i.e.*, (a)  $y_i(w^\top x_i + b) - 1$  is positive for all  $i \in \mathcal{R}$  and negative for all  $i \in \mathcal{L}$ , and (b) for each  $i \in \mathcal{M}$ ,  $\alpha_i$  remains between 0 and  $C_i$ . This defines a set of linear inequalities in  $(C_+, C_-)$ . The facets of the polytope defined by these inequalities can always be found in linear time in  $n$ , when efficient convex hull algorithms are used (Avis et al., 1997). However, when we only follow a straight line in the  $(C_+, C_-)$ -space, the polytope is then a segment and its extremities are trivial to find (also in  $O(n)$ ).

Following Hastie et al. (2005), if a solution is known for one value of  $(C_+, C_-)$ , we can follow the path along a line, by simply monitoring which constraint is violated first and changing the active sets accordingly.

**Path initialization** Hastie et al. (2005) requires that the training datasets are balanced in order to avoid solving a quadratic program to enter the path, *i.e.*, if  $n_+$  (resp.  $n_-$ ) is the number of positive (resp. negative) training examples, then they require that  $n_+ = n_-$ . In our situation, we can explore the two dimensional space  $(C_+, C_-)$ . Thus, the requirements become  $C_+n_+ = C_-n_-$  along the path. We can start the path by following the line  $C_+n_+ = C_-n_-$  and avoid the need to solve a quadratic program.

**Computational complexity** As shown by Hastie et al. (2005), if the appropriate online linear algebra tools are used, the complexity of obtaining one path of classifiers across one line is the same as obtaining the solution for one SVM using classical techniques such as sequential minimal optimization (Platt, 1998).




Figure 3: Lines in the  $(C_+, C_-)$ -space. The line  $C_+n_+ = C_-n_-$  is always followed first; then several lines with constant  $C_+ + C_-$  are followed in parallel, around the optimal line for the validation data (bold curve).

### 3.2 CONSTRUCTING THE ROC CURVE

Given the tools of Section 3.1, we can learn paths of linear classifiers from data. In this section, we present an algorithm to build ROC curves from the paths. We do this by exploring relevant parts of the  $(C_+, C_-)$  space, selecting the best classifiers among the ones that are visited.

We assume that we have two separate datasets, one for training and one for testing. This approach generalizes to cross-validation settings in a straightforward manner.

**Exploration phase** In order to start the path-following algorithm, we need to start at  $C_+ = C_- = 0$  and follow the line  $C_+n_+ = C_-n_-$ . We follow this line up to a large upper bound on  $C_+ + C_-$ . For all classifiers along that line, we compute a misclassification cost on the testing set, with given asymmetry  $(C_+^0, C_-^0)$  (as given by the user, and usually, but not necessarily, close to a point of interest in the ROC space). We then compute the best performing pair  $(C_+^1, C_-^1)$  and we select pairs of the form  $(rC_+^1, rC_-^1)$ , where  $r$  belongs to a set  $R$  of the type  $R = \{1, 10, 1/10, 100, 1/100, \dots\}$ . The set  $R$  provides further explorations for the total amount of regularization  $C_+ + C_-$ .

Then, for each  $r$ , we follow the paths of direction  $(1, -1)$  and  $(-1, 1)$  starting from the point  $(rC_+^1, rC_-^1)$ . Those paths have a fixed total amount of regularization but vary in asymmetry. In Figure 3, we show all of lines that are followed in the  $(C_+, C_-)$  space.

**Selection phase** After the exploration phase, we have  $|R| + 1$  different lines in the  $(C_+, C_-)$  space: the line  $C_-n_- = C_+n_+$ , and the  $|R|$  lines  $C_+ + C_- = r(C_+^1 + C_-^1), r \in R$ . For each of these lines, we know the optimal solution  $(w, b)$  for any cost settings on that line. The line  $C_-n_- = C_+n_+$  is used for computational purposes (*i.e.*, to enter the path), so we do not use it for testing purposes.

From  $R$  lines in the  $(C_+, C_-)$ -plane, we build the three ROC curves shown in Figure 4, for a finite sample

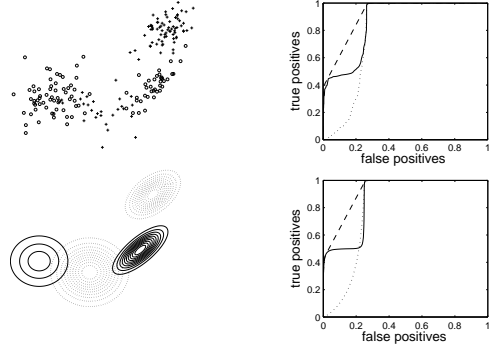


Figure 4: Two examples of ROC curves for bimodal class conditional densities, varying intercept (dotted), varying asymmetry (plain) and varying both (dashed). (Top) obtained from 10 random splits, using the data shown on the left side (one class is plotted as circles, the other one as crosses), (Bottom) obtained from population densities (one class with plain density contours, the other one with dotted contours).

problem and for an infinite sample problem (for the infinite sample, the solution of the SVM was obtained by working directly with densities):

- *Varying intercept*: we extract the slope  $w$  corresponding to the best setting  $(C_+^1 + C_-^1)$ , and vary the intercept  $b$  from  $-\infty$  to  $\infty$ . This is the traditional method for building an ROC curve for an SVM.
- *Varying asymmetry*: we only consider the line  $C_+ + C_- = C_+^1 + C_-^1$  in the  $(C_+, C_-)$ -plane; the classifiers that are used are the optimal solutions of the convex optimization problem. Note that for each value of the asymmetry, we obtain a different value of the slope and the intercept.
- *Varying intercept and asymmetry*: for each of the points on the  $R$  lines in the  $(C_+, C_-)$ -plane, we discard the intercept  $b$  and keep the slope  $w$  obtained from the optimization problem; we then use all possible intercept values; this leads to  $R$  two-dimensional surfaces in the ROC plane. We then compute the convex envelope of these, to obtain a single curve.

Since all classifiers obtained by varying only the intercept (resp. the asymmetry) are included in the set used for varying both the intercept and the asymmetry, the third ROC curve always outperforms the first two curves (*i.e.*, it is always closer to the top left corner). This is illustrated in Figure 4.

Intuitively, the ROC curve obtained by varying the asymmetry should be better than the ROC generated by varying the intercept because, for each point, the slope of the classifier is optimized. Empirically, this is generally true, but is not always the case, as displayed

in Figure 4. This is not a small sample effect, as the infinite sample simulation shows. Another troubling fact is that the ROC curve obtained by varying asymmetry, exhibits strong concavities, *i.e.*, there are many ROC-inconsistent points: for those points, the solution of the SVM with the corresponding asymmetry is far from being the best linear classifier when performance is measured with the same asymmetry but with the exact 0-1 loss. In addition, even for ROC-consistent points, the training asymmetry and the testing asymmetry differ. In the next section, we analyze why they may differ and characterize their relationships in some situations.

## 4 TRAINING VS. TESTING ASYMMETRY

We observed in Section 3.2 that the training cost asymmetry can differ from the testing asymmetry. In this section, we analyze their relationships more closely for the population (*i.e.*, infinite sample) case. Although a small sample effect might alter some of the results presented in this section, we argue that most of the discrepancies come from using a convex surrogate to the 0 – 1 loss.

The *Bayes optimal* classifier for a given asymmetry  $(C_+, C_-)$ , is the (usually non-linear) classifier with minimal expected cost. A direct consequence of results in Bartlett et al. (2004) is that, if the Bayes optimal classifier is linear, then using a convex surrogate has no effect, *i.e.*, using the expected  $\phi$ -cost will lead to the minimum expected cost. Thus, if the Bayes optimal classifier is linear, then, in the population case (infinite sample), there should be no difference. However, when the Bayes optimal classifier is not linear, then we might expect to obtain a difference, and we demonstrate that we do have one and quantify it for several situations.

Since we are using population densities, we can get rid of the regularization term and thus only the asymmetry will have an influence on the results, *i.e.*, we can restrict ourselves to  $C_+ + C_- = 1$ . We let  $\gamma = C_+ / (C_+ + C_-) = C_+$  denote the training asymmetry. For a given training asymmetry  $\gamma$  and a loss function  $\phi$ , in Section 2.2, we defined the *optimal testing asymmetry*  $\beta(\gamma)$  for the training asymmetry  $\gamma$ . In this section, we will refer to the  $\beta(\gamma)$  simply as the *testing asymmetry*.

Although a difference might be seen empirically for all possible asymmetries, we analyze the relationship between the testing cost asymmetry and training asymmetry in cases of extreme asymmetry, *i.e.*, in the ROC framework, close to the corner points (0, 0) and (1, 1). We prove that, depending on the class conditional den-

sities, there are two possible different regimes for extreme asymmetries: either the optimal testing asymmetry is more extreme, or it is less extreme. We also provide, under certain conditions, a simple test that can determine the regime given class conditional densities.

In this section, we choose class conditional densities that are either Gaussian or a mixture of Gaussians, because (a) any density can be approximated as well as desired by mixtures of Gaussians (Hastie et al., 2001), and (b) for the square loss and the erf loss, they enable closed-form calculations that lead to Taylor expansions.

### 4.1 OPTIMAL SOLUTIONS FOR EXTREME COST ASYMMETRIES

We assume that the class conditional densities are mixtures of Gaussian, *i.e.*, the density of positive (resp. negative) examples is a mixture of  $k_+$  Gaussians, with means  $\mu_+^i$  and covariance matrix  $\Sigma_+^i$ , and mixing weights  $\pi_+^i$ ,  $i \in \{1, \dots, m_+\}$  (resp.  $k_-$  Gaussians, with means  $\mu_-^i$  and covariance matrix  $\Sigma_-^i$ , and mixing weights  $\pi_-^i$ ,  $i \in \{1, \dots, m_-\}$ ). We denote  $p_+$  and  $p_-$  as the marginal class densities,  $p_+ = P(y = 1)$ ,  $p_- = P(y = -1)$ . We assume that all mixing weights  $\pi_{\pm}^i$  are strictly positives and that all covariance matrices  $\Sigma_{\pm}^i$  have full rank.

In the following sections, we provide Taylor expansions of various quantities around the null training asymmetry  $\gamma = 0$ . They trivially extend around the reverse asymmetry  $\gamma = 1$ . We start with an expansion of the unique global minimum  $(w, b)$  of the  $\phi$ -cost with asymmetry  $\gamma$ . For the square loss,  $(w, b)$  can be obtained in closed form for any class conditional densities so the expansion is easy to obtain, while for the erf loss, an asymptotic analysis of the optimality conditions has to be carried through, and is only valid for mixture of Gaussians (see Bach et al. (2004) for proofs).

**Proposition 1 (square loss)** *Under previous assumptions, we have the following expansions:*

$$\begin{aligned} w &= 2 \frac{p_+}{p_-} \gamma \Sigma_-^{-1} (\mu_+ - \mu_-) + O(\gamma^2) \\ b &= -1 + \frac{p_+}{p_-} \gamma [2 - 2\mu_-^\top (\mu_+ - \mu_-)] + O(\gamma^2) \end{aligned}$$

where  $m = \mu_+ - \mu_-$ , and  $\Sigma_{\pm}$  and  $\mu_{\pm}$  are the class conditional means and variances.

**Proposition 2 (erf loss)** *Under previous assumptions, we have the following expansions:*

$$\begin{aligned} w &= (2 \log(1/\gamma))^{-1/2} \tilde{\Sigma}_- (\tilde{\mu}_+ - \tilde{\mu}_-) + o\left(\log(1/\gamma)^{-1/2}\right) \\ b &= -(2 \log(1/\gamma))^{1/2} + o\left(\log(1/\gamma)^{1/2}\right) \end{aligned}$$

where  $\tilde{m} = \tilde{\mu}_+ - \tilde{\mu}_-$ , and  $\tilde{\Sigma}_\pm$  and  $\tilde{\mu}_\pm$  are convex combinations of the mixture means and covariances, *i.e.*, there exists strictly positive weights  $\tilde{\pi}_\pm^i$ , that sum to one for each sign, such that  $\tilde{\Sigma}_\pm = \sum_i \tilde{\pi}_\pm^i \Sigma_\pm^i$  and  $\tilde{\mu}_\pm = \sum_i \tilde{\pi}_\pm^i \mu_\pm^i$ .

The weights  $\tilde{\pi}_\pm^i$  can be expressed as the solution of a convex optimization problem (see Bach et al. (2004) for more details). When there is only one mixture component (Gaussian densities), then  $\tilde{\pi}_\pm^1 = 1$ .

## 4.2 EXPANSION OF TESTING ASYMMETRIES

Using the expansions of Proposition 1 and 2, we can readily derive an expansion of the ROC coordinates for small  $\gamma$ , as well as the testing asymmetry  $\beta(\gamma)$ . We have (see Bach et al. (2004) for proofs):

**Proposition 3 (square loss)** *Under previous assumptions, we have the following expansion:*

$$\log\left(\frac{p_-}{p_+}(\beta(\gamma)^{-1}-1)\right) = \frac{p_-^2}{8p_+^2\gamma^2} \left( \frac{1}{m^\top \Sigma_-^{-1} \Sigma_-^{i_-} \Sigma_-^{-1} m} - \frac{1}{m^\top \Sigma_-^{-1} \Sigma_+^{i_+} \Sigma_-^{-1} m} \right) + o(1/\gamma^2) \quad (1)$$

where  $i_-$  (resp.  $i_+$ ) is one of the negative (resp. positive) mixture component.

**Proposition 4 (erf loss)** *Under previous assumptions, we have the following expansion:*

$$\log\left(\frac{p_-}{p_+}(\beta(\gamma)^{-1}-1)\right) = 2 \log(1/\gamma) \left( \frac{1}{\tilde{m}^\top \tilde{\Sigma}_-^{-1} \Sigma_-^{i_-} \tilde{\Sigma}_-^{-1} \tilde{m}} - \frac{1}{\tilde{m}^\top \tilde{\Sigma}_-^{-1} \Sigma_+^{i_+} \tilde{\Sigma}_-^{-1} \tilde{m}} \right) + o(\log(1/\gamma)) \quad (2)$$

where  $i_-$  (resp.  $i_+$ ) is one of the negative (resp. positive) mixture component.

The rest of the analysis is identical for both losses and thus, for simplicity, we focus primarily on the square loss. For the square loss, we have two different regimes, depending on the sign of  $m^\top \Sigma_-^{-1} \Sigma_-^{i_-} \Sigma_-^{-1} m - m^\top \Sigma_-^{-1} \Sigma_+^{i_+} \Sigma_-^{-1} m$ :

- if  $m^\top \Sigma_-^{-1} \Sigma_-^{i_-} \Sigma_-^{-1} m > m^\top \Sigma_-^{-1} \Sigma_+^{i_+} \Sigma_-^{-1} m$ , then from the expansion in Eq. (1) and Eq. (2), we see that the testing asymmetry tends to 1 exponentially fast. Because this is an expansion around the null training asymmetry, the ROC curve must be starting on the bottom right part of the main diagonal and the points close to  $\gamma = 0$  are not ROC-consistent, *i.e.*, the classifiers with training asymmetry too close to zero are useless as they are too extreme.




Figure 5: Training asymmetry vs. testing asymmetry, **square loss**: (Left) Gaussian class conditional densities, (right) testing asymmetry vs. training asymmetry; from top to bottom, the values of  $(m^\top \Sigma_-^{-1} \Sigma_-^{i_-} \Sigma_-^{-1} m)^{-1} - (m^\top \Sigma_-^{-1} \Sigma_+^{i_+} \Sigma_-^{-1} m)^{-1}$  are 0.12, -6, 3, -0.96.

- if  $m^\top \Sigma_-^{-1} \Sigma_-^{i_-} \Sigma_-^{-1} m < m^\top \Sigma_-^{-1} \Sigma_+^{i_+} \Sigma_-^{-1} m$ , then from the expansion in Eq. (1) and Eq. (2), we see that the testing asymmetry tends to 0 exponentially fast, in particular, the derivative  $d\beta/d\gamma$  is null at  $\gamma = 0$ , meaning, that the testing asymmetry is significantly smaller than the training asymmetry, *i.e.*, less extreme.
- if  $m^\top \Sigma_-^{-1} \Sigma_-^{i_-} \Sigma_-^{-1} m = m^\top \Sigma_-^{-1} \Sigma_+^{i_+} \Sigma_-^{-1} m$ , then the asymptotic expansion does not provide any information relating to the behavior of the testing asymmetry. We are currently investigating higher-order expansions in order to study the behavior of this limiting case. Note that when the two class conditional densities are Gaussians with identical covariance (a case where the Bayes optimal classifier with symmetric cost is indeed linear), we are in the present case.

The strength of the effects we have described above depends on the norm of  $m = \mu_+ - \mu_-$ : if  $m$  is large, *i.e.*, the classification problem is simple, then those effects are less strong, while when  $m$  is small, they are stronger. In Figure 5, we provide several examples for the square loss, with the two regimes and different strengths. It is worth noting, that, although the theoretical results obtained in this section are asymptotic expansions around the corners (*i.e.*, extreme asymmetries), the effects also remain valid far from the corners.

We thus must test to identify which regime we are in, namely testing for the sign of  $m^\top \Sigma_-^{-1} \Sigma_-^{i_-} \Sigma_-^{-1} m - m^\top \Sigma_-^{-1} \Sigma_+^{i_+} \Sigma_-^{-1} m$ . This test requires knowledge of the class conditional densities; it can currently always be performed in closed form for the square loss, while for the erf loss, it requires to solve a convex optimization problem, described by Bach et al. (2004).

Dataset	$\gamma$	one asym.	all asym.
PIMA	0.68	$41 \pm 0.4$	$22 \pm 1$
BREAST	0.99	$0.9 \pm 0.03$	$0.09 \pm 0.04$
IONOSPHERE	0.82	$10 \pm 0.5$	$4 \pm 0.8$
LIVER	0.32	$27 \pm 1.8$	$23.8 \pm 0.02$
RINGNORM	0.94	$6.3 \pm 0.06$	$4.3 \pm 0.1$
TWONORM	0.16	$15 \pm 0.2$	$1.2 \pm 0.2$
ADULT	0.70	$12.8 \pm 0.8$	$11.5 \pm 0.3$

Table 1: Training with the testing asymmetry  $\gamma$  vs. training with all cost asymmetries: we report validation costs obtained from 10 half-random splits (pre-multiplied by 100). Only the asymmetry with the largest difference is reported. Given an asymmetry  $\gamma$  we use the cost settings  $C_+ = 2\gamma$ ,  $C_- = 2(1 - \gamma)$  (which leads to the misclassification error if  $\gamma = 1/2$ ).

### 4.3 BUILDING THE ENTIRE ROC CURVE FOR A SINGLE POINT

As shown empirically in Section 3.2, and demonstrated theoretically in this section, training and testing asymmetries differ; and this difference suggests that even when the user is interested in only one cost asymmetry, the training procedure should explore more cost asymmetries, i.e. build the ROC curve as presented in Section 3.2 and chose the best classifier as follows: a given asymmetry in cost for the test case leads to a unique slope in the ROC space, and the optimal point for this asymmetry is the point on the ROC curve whose tangent has the corresponding slope and which is closest to the upper-left corner.

We compare in Figure 1, for various datasets and linear classifiers, the performance with cost asymmetry  $\gamma$  of training a classifier with cost asymmetry  $\gamma$  to the performance of training with all cost asymmetries. Using all asymmetries always perform better than assuming a single asymmetry—we simply have more classifiers to choose from. In Figure 1, we report only the performance for the cost asymmetries which show the greatest differences, showing that in some cases, it is very significant, and that a simple change in the training procedure may lead to substantial gains.

## 5 CONCLUSION

We have presented an efficient algorithm to build ROC curves by varying the training cost asymmetries for SVMs. The algorithm is based on the piecewise linearity of the path of solutions when the cost of false positives and false negatives vary. We have also provided a theoretical analysis of the relationship between the potentially different cost asymmetries assumed in training and testing classifiers, showing that they differ under certain circumstances. We have characterized

key relationships, and have worked to highlight the potential value of building the entire ROC curve even when a single point may be needed. Such an approach can lead to a significant improvement of performance with little added computational cost. Finally, we note that, although we have focused in this paper on the single kernel learning problem, our approach can be readily extended to the multiple kernel learning setting (Bach et al., 2005) with appropriate numerical path following techniques.

### Acknowledgments

We thank John Platt for sharing his experiences and insights with considerations of cost functions in training and testing support vector machines.

### References

- D. Avis, D. Bremner, and R. Seidel. How good are convex hull algorithms? In *Computational Geometry: Theory and Applications*, volume 7, 1997.
- F. R. Bach, D. Heckerman, and E. Horvitz. On the path to an ideal ROC curve: Considering cost asymmetry in learning classifiers. Technical Report MSR-TR-2004-124, Microsoft Research, 2004.
- F. R. Bach, R. Thibaux, and M. I. Jordan. Computing regularization paths for learning multiple kernels. In *NIPS 17*, 2005.
- P. L. Bartlett, M. I. Jordan, and J. D. McAuliffe. Large margin classifiers: convex loss, low noise, and convergence rates. In *NIPS 16*, 2004.
- S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge Univ. Press, 2003.
- P. A. Flach. The geometry of ROC space: understanding machine learning metrics through ROC isometrics. In *ICML*, 2003.
- T. Hastie, S. Rosset, R. Tibshirani, and J. Zhu. The entire regularization path for the support vector machine. *JMLR*, 5:1391–1415, 2005.
- T. Hastie, R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning*. Springer-Verlag, 2001.
- M. S. Pepe. Receiver operating characteristic methodology. *J. Am. Stat. Assoc.*, 95(449):308–311, 2000.
- J. Platt. Fast training of support vector machines using sequential minimal optimization. In *Advances in Kernel Methods: Support Vector Learning*, Cambridge, MA, 1998. MIT Press.
- F. Provost and T. Fawcett. Robust classification for imprecise environments. *Machine Learning*, 42(3): 203–231, 2001.
- B. Schölkopf and A. J. Smola. *Learning with Kernels*. MIT Press, 2002.