
Optimal Classification with Multivariate Losses

Nagarajan Natarajan

Microsoft Research, INDIA

T-NANATA@MICROSOFT.COM

Oluwasanmi Koyejo

Stanford University, CA & University of Illinois at Urbana-Champaign, IL, USA

SANMI@ILLINOIS.EDU

Pradeep Ravikumar

Inderjit S. Dhillon

The University of Texas at Austin, TX, USA

PRADEEPR@CS.UTEXAS.EDU

INDERJIT@CS.UTEXAS.EDU

Abstract

Multivariate loss functions are extensively employed in several prediction tasks arising in Information Retrieval. Often, the goal in the tasks is to minimize expected loss when retrieving relevant items from a presented set of items, where the expectation is with respect to the joint distribution over item sets. Our key result is that for most multivariate losses, the expected loss is provably optimized by sorting the items by the conditional probability of label being positive and then selecting top k items. Such a result was previously known only for the F -measure. Leveraging on the optimality characterization, we give an algorithm for estimating optimal predictions in practice with runtime quadratic in size of item sets for many losses. We provide empirical results on benchmark datasets, comparing the proposed algorithm to state-of-the-art methods for optimizing multivariate losses.

1. Introduction

A recent flurry of theoretical results and practical algorithms highlights a growing interest in understanding and optimizing general multivariate losses (Joachims, 2005; Petterson and Caetano, 2011; Dembczynski et al., 2011; Ye et al., 2012; Koyejo et al., 2014; Narasimhan et al., 2014). Conventional losses such as the 0-1 loss (or the error) in binary or multiclass classification, and Hamming loss in case of multilabel learning fall short in many applications, such as medical diagnosis, fraud detection, and informa-

tion retrieval, with imbalanced and rare event classification tasks (Lewis and Gale, 1994; Drummond and Holte, 2005; Gu et al., 2009; He and Garcia, 2009). Practitioners employ multivariate loss functions such as the F -measure that capture non-linear trade-off between the entries of the “confusion matrix”, namely, true positives, false positives, true negatives and false negatives. Multivariate loss functions are defined on vector-valued predictions, such as label vectors in multilabel learning, subsets of classes in multiclass classification, etc. The goal in the prediction tasks is to minimize expected loss when retrieving relevant items from a presented set of items, where the expectation is with respect to the joint distribution over item sets (that models uncertainty in the data).

Algorithmic approaches for minimizing expected multivariate losses such as structured support vector machines have been proposed (Joachims, 2005; Petterson and Caetano, 2011). Interestingly, the optimization can be performed efficiently for most multivariate losses that can be written as a function of the entries of the confusion matrix. However, there are no known consistency guarantees for these approaches; in fact, recently, Dembczynski et al. (2013) showed that structured loss minimization is not consistent for the F -measure. On the other hand, an important theoretical question is if and when one can explicitly characterize the optimal solution for the loss minimization problem. A key difficulty in the analysis of general multivariate losses is that they are often *non-decomposable*, i.e., the loss on prediction vectors does not decompose into the sum of losses over individual predictions. Two decades ago, Lewis (1995) showed that the optimal solution to minimizing expected F -measure, under the assumption that the labels are conditionally independent, admits a simple form — in particular, it requires only the knowledge of the marginals $\mathbb{P}(Y_i|\mathbf{x})$. Then, Dembczynski et al. (2011) showed that the optimal solution for F -measure can be de-

scribed using $O(n^2)$ parameters for a general distribution \mathbb{P} over sets of n items. Since then, there has been extensive work focusing on binary and multilabel F -measure (Dembczynski et al., 2012; Ye et al., 2012; Dembczynski et al., 2013; Waegeman et al., 2014; Lipton et al., 2014). Yet, the question of characterizing optimal prediction for many other losses used in practice remains largely unknown.

In this paper, we provide a general theoretical analysis of expected loss minimization for general, non-decomposable, multivariate losses in binary, multiclass and multilabel prediction problems. Under conditional independence assumption on the underlying distribution, we show that the optimal solution for most losses exhibit a simple form, and depends only on the conditional probability of the label being positive. In multiclass classification, the joint distribution is a multinomial, and we show that a similar result holds. In particular, we identify a natural sufficient condition for any loss under which it allows such a simple characterization of optimal — we require the loss to be monotonically decreasing in true positives; this is satisfied by most, if not all, loss functions including the monotonic family studied by Narasimhan et al. (2014), and the linear fractional family studied by Koyejo et al. (2014; 2015). As a special case of our analysis, we naturally recover the F -measure result of Lewis (1995) for binary classification, and the result of Coz Velasco et al. (2009) for multiclass classification.

Minimizing (and even evaluating) expected multivariate losses can involve exponential-time computation, even when given access to the exact label distribution. As we show, in light of our main result characterizing optimal predictions, and with careful implementation, computations can be greatly simplified. We give an algorithm that runs in $O(n^3)$ time for a general loss, where n is the size of the item set (number of instances or labels or classes as the case maybe). For special cases such as F_β and Jaccard, the algorithm can be implemented to run in time $O(n^2)$. We prove that our overall procedure for computing the optimal in practice is consistent. We also support our theoretical results with experimental evaluation on synthetic and real-world datasets.

Related Work: We highlight some of the key results relating to prediction with general multivariate losses. Existing theoretical analysis has focused on two distinct approaches for characterizing the *population* version of multivariate losses: identified by Ye et al. (2012) as decision theoretic analysis (DTA) and empirical utility maximization (EUM). In DTA, the goal is to minimize the expected loss of a classifier on sets of predictions, which is the setting in our work here, while in EUM, the goal is to minimize the loss applied to population confusion matrix with *expected* values as entries. We can interpret DTA as minimizing the

average loss over an infinite set of test sets, each of a fixed size, while EUM as minimizing the loss of a classifier over a single infinitely large test set. More recently, there have been several theoretical and algorithmic advances relating to general performance measures (Parambath et al., 2014; Koyejo et al., 2014; Narasimhan et al., 2014; Kar et al., 2014; Narasimhan et al., 2015; Koyejo et al., 2015) used in binary, multiclass and multilabel classification in the EUM setting. In stark contrast, we know much less about the setting in our paper; several authors have proposed algorithms for empirical optimization of the expected F_β measure including Chai (2005), Jansche (2007) and Dembczynski et al. (2011). Ye et al. (2012) compare the DTA and the EUM analyses for F_β , showing an asymptotic equivalence as the number of test samples goes to infinity. Quevedo et al. (2012) propose a cubic complexity dynamic programming algorithm for computing optimal labeling under conditional label independence assumption, albeit without providing an optimality characterization or consistency.

2. Multivariate Loss Minimization

We consider the problem of making multivariate predictions $\mathbf{y} = \{y_1, y_2, \dots, y_n\} \in \{0, 1\}^n$, for a given set of instances (described by their features) $\mathbf{x} \in \mathcal{X}$. Let \mathbf{X} denote the random variable for instances, \mathbf{Y} denote the random variable for label vectors of length n (with Y_i denoting the random variable for i th label). Given a multivariate loss L , the goal is to minimize the expected loss wrt. the underlying joint distribution \mathbb{P} over \mathbf{X} and \mathbf{Y} :

$$\mathbf{h}^* = \arg \min_{\mathbf{h}: \mathbf{x} \mapsto \mathbf{y}} \mathbf{E}_{(\mathbf{X}, \mathbf{Y}) \sim \mathbb{P}} L(\mathbf{h}(\mathbf{X}), \mathbf{Y}).$$

Note that \mathbf{h}^* optimizes the expected loss wrt. to the conditional $\mathbb{P}(\mathbf{Y}|\mathbf{x})$ at each \mathbf{x} ; therefore it is sufficient to analyze the optimal predictions at a given \mathbf{x} .

$$\mathbf{h}^*(\mathbf{x}) = \arg \min_{\mathbf{h}: \mathbf{x} \mapsto \mathbf{y}} \mathbf{E}_{\mathbf{Y} \sim \mathbb{P}(\mathbf{Y}|\mathbf{x})} L(\mathbf{h}(\mathbf{x}), \mathbf{Y}). \quad (1)$$

The choices of \mathbf{x} , \mathbf{y} and L depend on the prediction task:

- In **binary classification**, $\mathbf{x} = \{x_1, x_2, \dots, x_n\}$ is a set of instances, $y_i \in \{0, 1\}$ corresponds to the label of instance x_i and $h_i(\mathbf{x}) \in \{0, 1\}$ denotes the predicted label for x_i . One may be interested in retrieving the positive instances with high precision or high recall (or a function of the two, such as the F -measure). We let $L : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \mathbb{R}_+$.
- In **multiclass classification** with n classes, $\mathbf{x} = \{x\}$ is an instance, $y \in [n]$ corresponds to the class of instance x , whereas the prediction $\mathbf{h}(x) \subseteq [n]$ corresponds to a *subset* of predicted classes (represented by a binary vector in $\{0, 1\}^n$). Here we want to retrieve the true class in a predicted subset of small size

$k \ll n$; so, $L : [n] \times \{0, 1\}^n \rightarrow \mathbb{R}_+$. This problem can be thought of as learning non-deterministic classifiers (Coz Velasco et al., 2009).

- In **multilabel learning** with n labels, $\mathbf{x} = \{x\}$ is an instance, $\mathbf{y}, \mathbf{h}(x) \in \{0, 1\}^n$ correspond to the set of relevant and predicted labels respectively for instance x . Popular choices of loss functions include multilabel F -measure and the Hamming loss. Here, $L : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \mathbb{R}_+$.

2.1. Hardness of the general setting

The optimization problem (1) essentially entails a search over binary vectors of length n . Efficient search for optimal solution or obtaining a consistent estimator can be notoriously hard depending on the loss function and the joint distribution. Of course, if the loss function is *decomposable* over instances (such as the Hamming loss in case of multilabel learning, or 0-1 loss in case of binary classification), a consistent estimator of the optimal solution can be obtained via empirical risk minimization. Dembczynski et al. (2011) showed that for the F -measure, and arbitrary distribution \mathbb{P} , optimal solution to (1) can be obtained provided we can estimate $O(n^2)$ parameters of \mathbb{P} — in particular, $\mathbb{P}(Y_i = 1, \sum_{i=1}^n Y_i = s | \mathbf{x}), i, s \in [n] \times [n]$. In case of the subset 0/1 loss defined as $L(\mathbf{h}(\mathbf{x}), \mathbf{y}) = 1$ if $\mathbf{h}(\mathbf{x}) \neq \mathbf{y}$ and $L(\mathbf{h}(\mathbf{x}), \mathbf{y}) = 0$ otherwise, $\mathbf{h}^*(\mathbf{x})$ is the mode of the distribution $\mathbb{P}(\mathbf{Y} | \mathbf{x})$, which is infeasible to estimate for arbitrary \mathbb{P} . For general multivariate losses, one popular algorithmic approach is to employ structural support vector machines (Joachims, 2005; Peterson and Caetano, 2011) which optimize a convex upper bound of the expected loss on training data. However, there are no consistency results known for the approach. In fact, Dembczynski et al. (2013) show that structural SVMs are inconsistent for arbitrary \mathbb{P} , in case of the F -measure. More recently, Wang et al. (2015) study the multivariate loss minimization problem from an adversarial point of view, and provide a game-theoretic solution. Inevitably, they require solving sub-problems of the form (1), which (as they discuss) can be worked out for a few specific losses (such as the F -measure) but are hard in general.

2.2. Conditional independence

Consider the setting when $\mathbb{P}(\mathbf{Y} | \mathbf{x})$ satisfies conditional independence. In case of binary classification, instances are typically assumed to be i.i.d., and therefore conditional independence $\mathbb{P}(\mathbf{Y} | \mathbf{x}) = \prod_{i=1}^n \mathbb{P}(Y_i | x_i)$ holds. In case of multilabel learning, conditional label independence $\mathbb{P}(\mathbf{Y} | x) = \prod_{i=1}^n \mathbb{P}(Y_i | x)$ may be strong, as labels are likely to be correlated in practice. It has been known for a long time that for the F -measure, under conditional independence (Lewis, 1995), the optimal solution to (1) can be

computed by simply sorting the instances according to $\mathbb{P}(Y | x_i)$ and setting the labels for the top k instances to 1 and the rest to 0 (for some $0 \leq k \leq n$). As a consequence, we only require estimates of the marginals to compute the optimal solution. For convenience, denote $\mathbf{h}(\mathbf{x})$ by $\mathbf{s} \in \{0, 1\}^n$ (and the optimal solution to (1) by \mathbf{s}^*).

Theorem (Lewis (1995)). *Consider:*

$$L_{F_\beta}(\mathbf{s}, \mathbf{y}) = 1 - \frac{(1 + \beta)^2 \sum_{i=1}^n s_i y_i}{\sum_{i=1}^n s_i + \beta^2 \sum_{i=1}^n y_i}. \quad (2)$$

Let x_i 's be sorted in decreasing order of the marginals $\mathbb{P}(Y | x_i)$. Then, the optimal predictions \mathbf{s}^ (1) for the F_β loss in (2) is given by $s_i^* = 1$, for $i \in [k^*]$, $s_i^* = 0$ otherwise, for some $0 \leq k^* \leq n$ that may depend on \mathbf{x} .*

But such a characterization of optimality is not known for other multivariate losses used in practice.

2.3. Multinomial distributions

Note that in case of multiclass classification with n classes, C_1, C_2, \dots, C_n , one can alternatively think of a joint distribution $\mathbb{P}(\mathbf{Y} | x)$ over label vectors $\mathbf{y} \in \{0, 1\}^n$ such that the distribution is supported only on \mathbf{y} satisfying $\sum_{i=1}^n y_i = 1$. Letting $e_i \in \{0, 1\}^n$ be the indicator vector for class C_i , define:

$$\begin{aligned} \mathbb{P}(\mathbf{Y} = e_i | x) &= \mathbb{P}(Y = C_i | x), \\ \mathbb{P}(\mathbf{Y} = \mathbf{y} | x) &= 0, \text{ otherwise.} \end{aligned} \quad (3)$$

Characterizing optimal solution is simple in this setting as well. Coz Velasco et al. (2009) proved a result very similar to that of (Lewis, 1995) (although the proof turns out to be much simpler in this case), in the context of multiclass classification with F_β -measure.

Theorem (Coz Velasco et al. (2009)). *Consider:*

$$\mathbf{s}^* = \arg \min_{\mathbf{s} \in \{0, 1\}^n} \mathbf{E}_{Y \sim \mathbb{P}(Y | x)} L_{F_\beta}(\mathbf{s}, e_Y),$$

where L_{F_β} is defined as in (2). Let the classes C_i be indexed in the decreasing order of the conditionals $\mathbb{P}(Y = C_i | x)$. Then, \mathbf{s}^ satisfies $s_i^* = 1$, for $i \in [k^*]$, $s_i^* = 0$ otherwise, for some $0 \leq k^* \leq n$ that may depend on \mathbf{x} .*

Next, we show that the above results indeed hold for general multivariate losses.

3. Main Results

We now show that for most multivariate losses L , the corresponding optimal predictions (1) can be explicitly characterized for multinomial distributions (as in the case of multiclass classification) and joint distributions satisfying

conditional independence (as in the cases of binary classification and multilabel learning with conditional label independence assumption). To this end, we consider general losses that are functions of the entries of the so-called confusion matrix, namely true positives, true negatives, false positives and false negatives. The empirical confusion matrix is computed as $\widehat{\mathbf{C}}(\mathbf{s}, \mathbf{y}) = \begin{bmatrix} \widehat{\text{TP}} & \widehat{\text{FN}} \\ \widehat{\text{FP}} & \widehat{\text{TN}} \end{bmatrix}$ with entries:

$$\begin{aligned} \widehat{\text{TP}} &= \frac{1}{n} \sum_{i=1}^n s_i y_i, & \widehat{\text{TN}} &= \frac{1}{n} \sum_{i=1}^n (1 - s_i)(1 - y_i), \\ \widehat{\text{FP}} &= \frac{1}{n} \sum_{i=1}^n s_i (1 - y_i), & \widehat{\text{FN}} &= \frac{1}{n} \sum_{i=1}^n (1 - s_i) y_i. \end{aligned}$$

Most multivariate loss functions used in practice (see Table 1) admit the above form. For example, L_{F_β} in (2) can be written as:

$$L_{F_\beta}(\mathbf{s}, \mathbf{y}) := L_{F_\beta}(\widehat{\mathbf{C}}(\mathbf{s}, \mathbf{y})) = 1 - \frac{(1 + \beta^2)\widehat{\text{TP}}}{(1 + \beta^2)\widehat{\text{TP}} + \beta^2\widehat{\text{FN}} + \widehat{\text{FP}}}.$$

Without loss of generality, we let $L : [0, 1]^4 \mapsto \mathbb{R}_+$ denote a multivariate loss function evaluated on the entries of the confusion matrix. We begin by observing the following equivalent representation for loss L . All the proofs in the manuscript are supplied in the Appendix due to limited space.

Proposition 1. *Let $u(\mathbf{s}, \mathbf{y}) = \widehat{\text{TP}}(\mathbf{s}, \mathbf{y})$, $v = v(\mathbf{s}) := \frac{1}{n} \sum_i s_i$ and $p = p(\mathbf{y}) := \frac{1}{n} \sum_i y_i$, then $\exists \Phi : [0, 1]^3 \rightarrow \mathbb{R}_+$ such that $L(\widehat{\mathbf{C}}(\mathbf{s}, \mathbf{y})) = \Phi(u(\mathbf{s}, \mathbf{y}), v(\mathbf{s}), p(\mathbf{y}))$.*

Next, we define a certain monotonicity property which can be seen to be satisfied by popular loss functions.

Definition 1 (TP Monotonicity). *A loss function L is said to be TP monotonic if for any v, p , and $u_1 > u_2$, it holds that $\Phi(u_1, v, p) < \Phi(u_2, v, p)$.*

In other words, L satisfies TP monotonicity if the corresponding representation Φ (Proposition 1) is monotonically decreasing in its first argument. It is easy to verify, for instance that $\Phi_{F_\beta}(u, v, p) = 1 - \frac{(1 + \beta^2)u}{\beta^2 p + v}$ is monotonically decreasing in u . We are now ready to state our main result regarding minimizing expected losses wrt. distributions satisfying conditional independence.

Theorem 1 (Binary Losses). *Assume the joint distribution \mathbb{P} satisfies conditional independence. Let L be a multivariate loss that satisfies TP monotonicity. Then, the optimal predictions $\mathbf{s}^* := \mathbf{h}^*(\mathbf{x})$ in (1) satisfy:*

$$\min\{\mathbb{P}(Y = 1|x_i)|s_i^* = 1\} \geq \max\{\mathbb{P}(Y = 1|x_i)|s_i^* = 0\}.$$

In other words, $s_i^ = 1$, for all $i \in [k^*]$, for some $0 \leq k^* \leq n$ that may depend on \mathbf{x} , and $s_i^* = 0$ for $i \notin [k^*]$.*

Note that the above result also applies to multilabel classification, albeit under conditional label independence assumption. Next, we give a result for multiclass classification, where the joint distribution is multinomial.

Theorem 2 (Multiclass Losses). *Fix a multinomial distribution \mathbb{P} (3). Let L be a multivariate loss that satisfies TP monotonicity. Let the classes C_i be indexed in the decreasing order of the conditionals $\mathbb{P}(Y = C_i|x)$. Then the optimal predictions $\mathbf{s}^* := \mathbf{h}^*(\mathbf{x})$ in (1) is given by: $s_i^* = 1$, for $i \in [k^*]$, for some $0 \leq k^* \leq n$ that may depend on \mathbf{x} , and $s_i^* = 0$ otherwise.*

The proof essentially is a direct consequence of TP monotonicity and thus gives a generalization of the result in (Coz Velasco et al., 2009).

3.1. Recovered and New Results

It is clear that our results generalize those by Lewis (1995) and by Coz Velasco et al. (2009) for expected loss minimization. Now, we draw attention to some of the recent optimality results in classification using general loss functions, where the objective is different from (1). The motivation of this section is to highlight the generality of our analysis and place our contributions in the context of recent results in a closely related setting. For a binary classifier $h : \mathcal{X} \mapsto \{0, 1\}$ and a distribution \mathbb{P} over $\mathcal{X} \times \{0, 1\}$, let $\mathbf{C}(h; \mathbb{P}) = \begin{bmatrix} \text{TP} & \text{FN} \\ \text{FP} & \text{TN} \end{bmatrix}$ represent the *population* confusion matrix with entries:

$$\begin{aligned} \text{TP} &= \mathbb{P}(h(x) = 1, y = 1), & \text{TN} &= \mathbb{P}(h(x) = 0, y = 0), \\ \text{FP} &= \mathbb{P}(h(x) = 1, y = 0), & \text{FN} &= \mathbb{P}(h(x) = 0, y = 1). \end{aligned}$$

Koyejo et al. (2014) and Narasimhan et al. (2014) are interested in optimizing the following notion of expected loss:

$$h^* = \arg \min_{h: \mathcal{X} \rightarrow \{0, 1\}} L(\mathbf{C}(h; \mathbb{P}))$$

i.e., L is applied to the population confusion matrix. Under mild assumptions on the distribution \mathbb{P} , they show for a large family of loss functions, that the optimal solution satisfies a thresholded form, i.e. $h^*(x) = \text{sign}(\mathbb{P}(Y|x) - \delta^*)$, where δ^* is a constant that depends only on the distribution and the loss itself. In contrast, for the expected loss minimization in (1), there need not be a threshold in general that minimizes the expected loss on a given \mathbf{x} (as also discussed in (Lewis, 1995)).

The Fractional Linear Family: Koyejo et al. (2014; 2015) studied a large family L_{FL} of losses that are represented by:

$$\Phi_{\text{FL}}(\widehat{\text{TP}}(\mathbf{s}, \mathbf{y}), v(\mathbf{s}), p(\mathbf{y})) = \frac{c_0 + c_1 \widehat{\text{TP}} + c_2 v + c_3 p}{d_0 + d_1 \widehat{\text{TP}} + d_2 v + d_3 p} \quad (4)$$

for constants $c_i, d_i, i = \{0, 1, 2, 3\}$. We identify a subclass of L_{FL} where our results apply. The following result can be proven by inspection and is stated without proof.

Proposition 2. *If $c_1 < d_1$, then L_{FL} satisfies TP monotonicity.*

Note that this essentially constitutes the most useful losses in L_{FL} where increase in true positives (for a fixed total number of predictions) leads to decrease in loss.

Metrics from Narasimhan et al. (2014): The following alternative three-parameter representation of losses L was studied by Narasimhan et al. (2014). Let $p = p(\mathbf{y}) := \frac{1}{n} \sum_i y_i$, $r_p = \widehat{\text{TPR}}(\mathbf{s}, \mathbf{y}) = \frac{\widehat{\text{TP}}(\mathbf{s}, \mathbf{y})}{p(\mathbf{y})}$ and $r_n = \widehat{\text{TNR}}(\mathbf{s}, \mathbf{y}) = \frac{\widehat{\text{TN}}(\mathbf{s}, \mathbf{y})}{1-p(\mathbf{y})}$, then $\exists \Gamma : [0, 1]^3 \rightarrow \mathbb{R}_+$ such that:

$$L(\widehat{\mathbf{C}}(\mathbf{s}, \mathbf{y})) = \Gamma(\widehat{\text{TPR}}(\mathbf{s}, \mathbf{y}), \widehat{\text{TNR}}(\mathbf{s}, \mathbf{y}), p(\mathbf{y})). \quad (5)$$

As shown in Table 1, many losses used in practice are easily represented in this form. Representation for additional losses is simplified by including the empirical precision, given by $\widehat{\text{Prec}}(\mathbf{s}, \mathbf{y}) = \frac{\widehat{\text{TP}}(\mathbf{s}, \mathbf{y})}{v(\mathbf{s})}$, where $v(\mathbf{s}) := \frac{1}{n} \sum_i s_i = \widehat{\text{TP}} + \widehat{\text{FP}}$. Consider the following monotonicity property relevant to the representation (5).

Definition 2 (TPR/TNR Monotonicity). *A loss L is said to be TPR/TNR monotonic if when $r_{p1} > r_{p2}$ and $r_{n1} > r_{n2}$ and p fixed, then $\Gamma(r_{p1}, r_{n1}, p) < \Gamma(r_{p2}, r_{n2}, p)$.*

In other words, L satisfies TPR/TNR monotonicity if the corresponding Γ in (5) is monotonically decreasing in its first two arguments. It can be shown that all the losses listed in Table 1 satisfy TPR/TNR monotonicity. The following proposition states that any loss function that satisfies TPR/TNR monotonicity also satisfies TP monotonicity.

Proposition 3. *If L satisfies TPR/TNR monotonicity, then L satisfies TP monotonicity.*

We can verify from the third column of Table 1 that each of the TPR/TNR monotonic losses $\Phi(u, v, p)$ is monotonically decreasing in u . So any loss that satisfies TPR/TNR monotonicity admits optimal classifier (1) as stated in Theorems 1 and 2.

The area under the ROC curve (AUC): AUC is an important special case, which reduces to Φ_{AUC} (see Table 1) in case of binary predictions (Joachims, 2005). It is clear that following proposition follows directly, and is stated without proof:

Proposition 4. Φ_{AUC} satisfies TP monotonicity.

While prior work on AUC has focused on optimizing prediction of continuous scores, our approach is able to optimize explicit label predictions. Note that optimality results

Table 1. Examples of TP monotonic losses.

LOSS	DEFINITION	$\Phi(u, v, p)$
AM	$1 - \frac{\widehat{\text{TPR}} + \widehat{\text{TNR}}}{2}$	$1 - \frac{u+p(1-v-p)}{p(1-p)}$
F_β	$1 - \frac{1+\beta^2}{\frac{\beta^2}{\widehat{\text{Prec}}} + \frac{1}{\widehat{\text{TPR}}}}$	$1 - \frac{(1+\beta^2)u}{\beta^2 p + v}$
Jaccard	$\frac{\widehat{\text{FP}} + \widehat{\text{FN}}}{\widehat{\text{TP}} + \widehat{\text{FP}} + \widehat{\text{FN}}}$	$\frac{p+v-2u}{p+v-u}$
G-TP/PR	$1 - \sqrt{\widehat{\text{TPR}} \cdot \widehat{\text{Prec}}}$	$1 - \frac{u}{\sqrt{p \cdot v}}$
G-Mean	$1 - \sqrt{\widehat{\text{TPR}} \cdot \widehat{\text{TNR}}}$	$1 - \frac{u(1-v-p+u)}{p(1-p)}$
H-Mean	$1 - 2 / \left(\frac{1}{\widehat{\text{TPR}}} + \frac{1}{\widehat{\text{TNR}}} \right)$	$1 - \frac{2u(1-v-p+u)}{(1-v-p)p+u}$
AUC	$\frac{\widehat{\text{FP}} \cdot \widehat{\text{FN}}}{(\widehat{\text{TP}} + \widehat{\text{FN}})(\widehat{\text{FP}} + \widehat{\text{TN}})}$	$\frac{(v-u)(p-u)}{p(1-p)}$

for continuous scores need not trivially extend to optimality results for discrete label decisions.

4. Algorithms

For multiclass classification, Theorem 2 immediately suggests $O(n^2)$ algorithm for computing the optimal solution when $\mathbb{P}(Y = C_i | x)$ is known — for each $k = 1, 2, \dots, n$, evaluate the expected loss in selecting the top k classes (classes are sorted by $\mathbb{P}(Y = C_i | x)$), which can be done in $O(n)$ time. However, in binary and multilabel learning scenarios, even when $\mathbb{P}(Y_i = 1 | \mathbf{x})$ is known exactly, characterization of optimality in Theorem 1 falls short in practice — it is not obvious how to compute the expectation in (1) without exhaustively enumerating 2^n possible \mathbf{y} vectors. In this section, we present efficient algorithms for computing estimators for optimal predictions given \mathbf{x} and a TP monotonic loss function L . We also prove the consistency of the proposed algorithms.

4.1. Computing Optimal in Practice

We observe that by evaluating the loss L through the function Φ (as defined in Proposition 1), we can generalize the approach suggested by Ye et al. (2012) for F_β -measure to obtain a template algorithm for optimizing any TP monotonic L . Consider the vector $\mathbf{s} \in \{0, 1\}^n$ with the top k values set to 1 and the rest to 0, and let $S_{i:j} := \sum_{l=i}^j y_l$. Note that for any $\mathbf{y} \in \{0, 1\}^n$ that satisfies $S_{1:k} = k_1$ and $S_{k+1:n} = k_2$, $L(\mathbf{s}, \mathbf{y})$ can simply be evaluated as $\Phi(\frac{1}{n}k_1, \frac{1}{n}k, \frac{1}{n}(k_1 + k_2))$. Thus $\sum_{\mathbf{y} \in \{0, 1\}^n} \mathbb{P}(\mathbf{y} | \mathbf{x}) L(\mathbf{s}, \mathbf{y})$ can be evaluated as a sum over possible values of k_1 and k_2 , where the expectation is computed wrt. $P(S_{1:k} = k_1)P(S_{k+1:n} = k_2)$ with $0 \leq k_1 \leq k$ and $0 \leq k_2 \leq n - k$. Now, it remains to compute $P(S_{1:k} = k_1)$ and $P(S_{k+1:n} = k_2)$ efficiently.

Let $\eta_i = \mathbb{P}(Y_i = 1 | \mathbf{x})$. A consistent estimate of this quantity may be obtained by minimizing a strongly proper loss

function such as logistic loss (Reid and Williamson, 2009). Using the conditional independence assumption, we can show that $P(S_{1:k} = k_1)$ and $P(S_{k+1:n} = k_2)$ are the coefficients of z^{k_1} and z^{k_2} in $\prod_{j=1}^k [\eta_j z + (1 - \eta_j)]$ and $\prod_{j=k+1}^n [\eta_j z + (1 - \eta_j)]$ respectively, each of which can be computed in time $O(n^2)$ for fixed k . Note that the loss L itself can be evaluated in constant time. Let L_k denote the expected loss wrt. the estimated conditionals η_i , corresponding to setting the predictions of indices with top k highest conditional probabilities to 1. The resulting algorithm is presented in Algorithm 1. Though the computational complexity of Algorithm 1 is $O(n^3)$, we find that in practice it suffices to run the outer iterations until k^* where k^* is the first k s.t. $L_k \leq L_{k+1}$ (or $k^* = n$ if no such k exists), because for all $k > k^*$, $L_k \geq L_{k^*}$ holds. It is not obvious if this is a property enjoyed by all TP monotonic loss functions under conditional independence, but it would be interesting to theoretically establish this. The improvement in runtime is significant as multivariate losses are typically employed in scenarios where there is heavy imbalance in the distribution (in case of binary classification) or typically a small set of labels are relevant for a given instance (in case of multilabel learning), so $k^* \ll n$, for large n . Finally, we note that for a sub-family of fractional-linear losses studied by Koyejo et al. (2014; 2015), we can get a faster algorithm that runs in time $O(n^2)$ using the trick by Ye et al. (2012). Due to limited space, the resulting algorithm is presented in Appendix B.1.

Remark on multinomial distributions. For multiclass classification, note that computing the expectation wrt. the multinomial distribution is straight-forward. We simply replace steps 4-6 of Algorithm 1 with the following step:

$$L_k \leftarrow \sum_{j \leq k} \Phi\left(\frac{1}{n}, \frac{k}{n}, \frac{1}{n}\right) \cdot \eta_j + \sum_{k < j \leq n} \Phi\left(0, \frac{k}{n}, \frac{1}{n}\right) \cdot \eta_j$$

4.2. Consistency Analysis

Consider a procedure that minimizes (1) computed with respect to a consistent estimate $\hat{\mathbb{P}}(Y_i | \mathbf{x})$ of the conditional $\mathbb{P}(Y_i | \mathbf{x})$. Here, we show that any such procedure is consistent.

Theorem 3. *Let $\eta_i = \mathbb{P}(Y_i = 1 | \mathbf{x})$, and assume the estimate $\hat{\eta}_i$ satisfies $\hat{\eta}_i \xrightarrow{P} \eta_i$. Given a bounded loss function L and a set of instances \mathbf{x} , let $\mathbf{s}^* = \arg \min_{\mathbf{s} \in \{0,1\}^n} \mathbf{E}_{\mathbf{Y} \sim \mathbb{P}(\cdot | \mathbf{x})} L(\mathbf{s}, \mathbf{Y})$ be the optimal prediction with respect to \mathbb{P} and $\hat{\mathbf{s}} = \arg \min_{\mathbf{s} \in \{0,1\}^n} \mathbf{E}_{\mathbf{Y} \sim \hat{\mathbb{P}}(\cdot | \mathbf{x})} L(\mathbf{s}, \mathbf{Y})$ be the optimal prediction with respect to the consistent estimate $\hat{\mathbb{P}}$, then*

$$\mathbf{E}_{\mathbf{Y} \sim \mathbb{P}(\cdot | \mathbf{x})} [L(\hat{\mathbf{s}}, \mathbf{Y}) - L(\mathbf{s}^*, \mathbf{Y})] \xrightarrow{P} 0.$$

Our consistency result also applies to previous algorithms proposed for F_β in different settings e.g. by Lewis (1995);

Algorithm 1 Computing \mathbf{s}^* for TP Monotonic L

- 1: **Input:** L and estimates of $\eta_i = \mathbb{P}(Y_i = 1 | \mathbf{x})$, $i = 1, 2, \dots, n$ sorted wrt. η_i
 - 2: Init $s_i^* = 0, \forall i \in [n]$.
 - 3: **for** $k = 1, 2, \dots, n$ **do**
 - 4: For $0 \leq i \leq k$, set $C_k[i]$ as the coefficient of z^i in $\prod_{i=1}^k (\eta_i z + (1 - \eta_i))$.
 - 5: For $0 \leq i \leq n - k$, set $D_k[i]$ as the coefficient of z^i in $\prod_{i=k+1}^n (\eta_i z + (1 - \eta_i))$.
 - 6: $L_k \leftarrow \sum_{\substack{0 \leq k_1 \leq k \\ 0 \leq k_2 \leq n - k}} C_k[k_1] D_k[k_2] \Phi\left(\frac{k_1}{n}, \frac{k}{n}, \frac{1}{n}(k_1 + k_2)\right)$.
 - 7: **end for**
 - 8: $k^* \leftarrow \arg \min_k L_k$.
 - 9: return \mathbf{s}^* s.t. $s_i^* \leftarrow 1$ for $i \in [k^*]$.
-

Chai (2005); Jansche (2007); Coz Velasco et al. (2009); Ye et al. (2012), where analysis of consistency with empirical probability estimates was lacking. For TP monotonic multivariate losses, it is immediate from Theorem 3 that Algorithm 1 is consistent.

5. Experiments

We present two sets of experiments. The first is an experimental validation on synthetic data with known ground truth probabilities. The results serve to verify our main result (Theorem 1) for some of the losses in Table 1. The second set is an experimental evaluation of the proposed algorithm for computing optimal prediction on benchmark datasets, with comparisons to baseline and state-of-the-art algorithms for classification with general losses.

5.1. Synthetic data: Verification of the theory

We consider four losses from Table 1 based on the performance metrics AM, Jaccard, F_1 (harmonic mean of Precision and Recall) and G-TP/PR (geometric mean of Precision and Recall). To simulate, we sample a set of ten 2-dimensional vectors $\mathbf{x} = \{x_1, x_2, \dots, x_{10}\}$ from the standard Gaussian. The conditional probability is modeled using a sigmoid function: $\eta_i = \mathbb{P}(Y = 1 | x_i) = \frac{1}{1 + \exp(-w^T x_i)}$, for a random vector w also sampled from the standard Gaussian. The optimal predictions \mathbf{s}^* (1) are then obtained by exhaustive search over the 2^{10} possible label vectors. For each loss, we plot the conditional probabilities (in decreasing order) and \mathbf{s}^* in Figure 1. We observe that it is optimal to assign positive labels to top k^* instances with highest $\mathbb{P}(Y | x)$, as given in Theorem 1, where k^* depends on the instances, the distribution and the loss itself.




Figure 1. Optimal predictions for losses from Table 1 demonstrated on synthetic data. In each case, we verify that s^* conforms to the ordering of $\mathbb{P}(Y_i|\mathbf{x})$ as stated in Theorem 1.

DATASET	T	Proposed F_1	Baseline F_1	Tuning δ F_1	S-SVM F_1	Proposed Jaccard	Baseline Jaccard	Tuning δ Jaccard	S-SVM Jaccard
REUTERS (65)	1	0.4125	0.4849	0.5020	0.3254	0.5239	0.5692	0.5691	0.4977
	10	0.1753	0.2376	0.2401	0.1604	0.3199	0.3591	0.3090	0.2983
	50	0.1003	0.1572	0.1490	0.1533	0.2485	0.2552	0.2422	0.2528
	100	0.0144	0.0325	0.0331	0.0360	0.0602	0.0625	0.0643	0.0689
LETTERS (26)	1	0.2890	0.5173	0.4255	0.4395	0.5728	0.6368	0.5682	0.5029
SCENE (6)	1	0.3084	0.8023	0.5603	0.3937	0.6460	0.9794	0.7920	0.5541
WEB PAGE	1	0.1606	0.3145	0.3191	0.4700	0.5363	0.4785	0.4806	0.5758
SPAMBASE	1	0.1552	0.1202	0.1161	0.2263	0.2686	0.2133	0.1997	0.3461
IMAGE	1	0.1458	0.1429	0.1419	0.1227	0.2545	0.2500	0.2377	0.2235
BREAST CANCER	1	0.0207	0.0411	0.0234	0.0270	0.0658	0.0789	0.0519	0.0526

Table 2. Comparison of methods: Fractional-linear losses, F_1 and Jaccard defined in Table 1. Reported values correspond to losses on test data (lower values are better). Baseline refers to thresholding $\hat{\eta}(x)$ at 0.5; ‘Tuning δ ’ refers to the plugin-estimator combined with threshold selection method in Koyejo et al. (2014); S-SVM refers to the structured SVM method proposed by Joachims (2005). First three are multiclass datasets (number of classes indicated in parenthesis).

5.2. Benchmark data: Proposed algorithms

We perform classification using the proposed approach: (i) obtain a model for the conditional distribution $\eta(x) = \mathbb{P}(Y = 1|x)$ using training data and (ii) compute s^* for the test data using estimated conditionals in the proposed Algorithm 1 (and in the faster Algorithm 2 in the Appendix, when applicable). We use logistic loss on the training samples (with L_2 regularization) to obtain an estimate $\hat{\eta}(x)$ of $\mathbb{P}(Y = 1|x)$. In our experiments, we consider losses based on four performance metrics AM, F_1 , Jaccard and G-TP/PR. For AM and G-TP/PR we use Algorithm 1, while for the fractional-linear losses Jaccard and F_1 we use the more efficient Algorithm 2. We report the loss on the test data.

We compare our approach with that of structured hinge loss minimization for multivariate losses (Joachims, 2005). We use the MATLAB wrapper provided by Vedaldi (2011), which internally uses the fast `svm-struct` implementation of (Joachims, 2005). The solver is based on cutting-plane method and for each of the aforementioned losses, we can implement the constraint generation step efficiently. We also compare with that of the plugin-estimator combined with threshold selection (which we refer to as ‘Tun-

ing δ^*) proposed by Koyejo et al. (2014) and Narasimhan et al. (2014), that minimizes the loss on the *population* confusion matrix (as discussed in Section 3.1). In this case, the optimal classifier is given by $\text{sign}(\hat{\eta}(x) - \delta^*)$. The training data is split into two sets, one set is used for estimating $\hat{\eta}(x)$ and the other for selecting the optimal δ for the loss function. The predictions are then made by thresholding $\hat{\eta}(x)$ of the test data points at δ . We also compare to the standard baseline method for binary classification — thresholding $\hat{\eta}(x)$ at $1/2$.

We report results on seven benchmark datasets (used in (Koyejo et al., 2014; Ye et al., 2012)). (1) REUTERS, consisting of 8293 news articles categorized into 65 topics. We present results for averaging over topics with at least T positives in the training (5946 articles) as well as the test (2347 articles) data; (2) LETTERS dataset consisting of 20000 handwritten characters (16000 training and 4000 test instances) categorized into 26 letters; (3) SCENE (a UCI benchmark dataset) consisting of 2230 images (1137 training and 1093 test instances) categorized into 6 scene types; (4) WEBPAGE binary dataset, consisting of 34780 web pages (6956 train and 27824 test); highly imbalanced, with only about 182 positive instances in the

DATASET	T	Proposed AM	Baseline AM	Tuning δ AM	S-SVM AM	Proposed G-TP/PR	Baseline G-TP/PR	Tuning δ G-TP/PR	S-SVM G-TP/PR
REUTERS (65)	1	0.1166	0.2777	0.2267	0.0598	0.2711	0.4553	0.4782	0.2700
	10	0.0480	0.1640	0.0889	0.0368	0.1934	0.2200	0.1924	0.1579
	50	0.0341	0.0983	0.0418	0.0374	0.1505	0.1559	0.1309	0.1527
	100	0.0217	0.0239	0.0219	0.0235	0.0313	0.0325	0.0328	0.0369
LETTERS (26)	1	0.1285	0.2980	0.1280	0.1800	0.4213	0.4936	0.4098	0.3465
SCENE (6)	1	0.4160	0.4935	0.4190	0.2250	0.4931	0.9395	0.6152	0.3935
WEB PAGE	1	0.1311	0.1795	0.1250	0.1528	0.3383	0.3133	0.3114	0.4384
SPAMBASE	1	0.1220	0.0990	0.0910	0.1650	0.1506	0.1169	0.1087	0.2204
IMAGE	1	0.1959	0.1808	0.1931	0.1599	0.1324	0.1423	0.1298	0.1198
BREAST CANCER	1	0.0204	0.0399	0.0170	0.0205	0.0340	0.0410	0.0266	0.0270

Table 3. Comparison of methods: AM and G-TP/PR losses defined in Table 1. Reported values correspond to losses on test data (lower values are better). Baseline refers to thresholding $\hat{\eta}(x)$ at 0.5; ‘Tuning δ ’ refers to the plugin-estimator combined with threshold selection method in Narasimhan et al. (2014). First three are multiclass datasets (number of classes indicated in parenthesis).

train; (5) IMAGE, with 1300 train and 1010 test images; (6) BREAST CANCER, with 463 train and 220 test instances, and (7) SPAMBASE with 3071 train and 1530 test instances. See (Koyejo et al., 2014; Ye et al., 2012) for more details on the datasets. Note that in case of the multiclass datasets, we report results (using *one-versus-all* classifiers) averaged over classes.

The results for F_1 and Jaccard losses are presented in Table 2. Note that the reported values correspond to losses — smaller values are better. We find that our proposed algorithm almost always achieves the least F_1 loss compared to other methods. However, in case of the Jaccard loss, tuning a threshold on training data seems to perform better though it is theoretically not optimal. The structured SVM method, though not known to have strong consistency properties, is competitive here. The results for AM and G-TP/PR losses are presented in Table 3. We find that the proposed approach is competitive across many datasets for the AM loss. Again, we find that the structured SVM method is competitive in case of the G-TP/PR loss. Perhaps, structured loss minimization may hold promise for certain data distributions and it is an interesting future research direction.

6. Conclusions

We study optimizing expected multivariate losses used in several prediction tasks. Our analysis shows that for losses that satisfy an intuitive monotonicity property, optimal predictions can be computed given the knowledge of the conditional probability of the positive class. We propose efficient and consistent estimators for computing optimal predictions in practice. Our results are complementary to some of the recent advances in the understanding of optimal classification (Koyejo et al., 2014; Narasimhan et al., 2014) and a leap forward in the direction initiated two decades ago by Lewis (1995).

Acknowledgments: We acknowledge the support of NSF via CCF-1320746 and IIS-1546459, and NIH via R01 GM117594-01 as part of the Joint DMS/NIGMS Initiative to Support Research at the Interface of the Biological and Mathematical Sciences.

References

- Kian Ming Adam Chai. Expectation of F-measures: tractable exact computation and some empirical observations of its properties. In *Proceedings of the 28th Annual Intl. ACM SIGIR Conf. on Research and Development in Information Retrieval*, pages 593–594. ACM, 2005.
- Juan José del Coz Velasco, Jorge Díez Peláez, and Antonio Bahamonde Rionda. Learning nondeterministic classifiers. *Journal of Machine Learning Research*, 10, 2009.
- Krzysztof Dembczyński, Willem Waegeman, Weiwei Cheng, and Eyke Hüllermeier. On label dependence and loss minimization in multi-label classification. *Machine Learning*, 88(1-2):5–45, 2012.
- Krzysztof Dembczynski, Arkadiusz Jachnik, Wojciech Kotlowski, Willem Waegeman, and Eyke Hüllermeier. Optimizing the F-measure in multi-label classification: Plug-in rule approach versus structured loss minimization. In *Proceedings of The 30th Intl. Conf. on Machine Learning*, pages 1130–1138, 2013.
- Krzysztof J Dembczynski, Willem Waegeman, Weiwei Cheng, and Eyke Hüllermeier. An exact algorithm for F-measure maximization. In *Advances in Neural Information Processing Systems*, pages 1404–1412, 2011.
- Chris Drummond and Robert C Holte. Severe class imbalance: Why better algorithms aren’t the answer? In *Machine Learning: ECML 2005*, pages 539–546. Springer, 2005.

- Qiong Gu, Li Zhu, and Zhihua Cai. Evaluation measures of the classification performance of imbalanced data sets. In *Computational Intelligence and Intelligent Systems*, pages 461–471. Springer, 2009.
- Haibo He and Eduardo A Garcia. Learning from imbalanced data. *Knowledge and Data Engineering, IEEE Transactions on*, 21(9):1263–1284, 2009.
- Martin Jansche. A maximum expected utility framework for binary sequence labeling. In *Annual Meeting of the Association of Computational Linguistics*, volume 45, page 736, 2007.
- Thorsten Joachims. A support vector method for multivariate performance measures. In *Proceedings of the 22nd Intl. Conf. on Machine Learning*, pages 377–384. ACM, 2005.
- Purushottam Kar, Harikrishna Narasimhan, and Prateek Jain. Online and stochastic gradient methods for non-decomposable loss functions. In *Advances in Neural Information Processing Systems*, pages 694–702, 2014.
- Oluwasanmi O Koyejo, Nagarajan Natarajan, Pradeep K Ravikumar, and Inderjit S Dhillon. Consistent binary classification with generalized performance metrics. In *Advances in Neural Information Processing Systems*, pages 2744–2752, 2014.
- Oluwasanmi O Koyejo, Nagarajan Natarajan, Pradeep K Ravikumar, and Inderjit S Dhillon. Consistent multilabel classification. In *Advances in Neural Information Processing Systems*, pages 3303–3311, 2015.
- David D Lewis. Evaluating and optimizing autonomous text classification systems. In *Proceedings of the 18th Intl. ACM SIGIR Conf. on Research and Development in Information Retrieval*, pages 246–254. ACM, 1995.
- David D Lewis and William A Gale. A sequential algorithm for training text classifiers. In *Proceedings of the 17th Annual Intl. ACM SIGIR Conf.*, pages 3–12. Springer-Verlag New York, Inc., 1994.
- Zachary Chase Lipton, Charles Elkan, and Balakrishnan Narayanaswamy. Thresholding classifiers to maximize F1 score. *arXiv*, abs/1402.1892, 2014.
- Harikrishna Narasimhan, Rohit Vaish, and Shivani Agarwal. On the statistical consistency of plug-in classifiers for non-decomposable performance measures. In *Advances in Neural Information Processing Systems*, pages 1493–1501, 2014.
- Harikrishna Narasimhan, Harish Ramaswamy, Aadirupa Saha, and Shivani Agarwal. Consistent multiclass algorithms for complex performance measures. In *Proceedings of the 32nd Intl. Conf. on Machine Learning*, pages 2398–2407, 2015.
- Shameem Puthiya Parambath, Nicolas Usunier, and Yves Grandvalet. Optimizing F-measures by cost-sensitive classification. In *Advances in Neural Information Processing Systems*, pages 2123–2131, 2014.
- James Petterson and Tibério S Caetano. Submodular multilabel learning. In *Advances in Neural Information Processing Systems*, pages 1512–1520, 2011.
- José Ramón Quevedo, Oscar Luaces, and Antonio Bahamonde. Multilabel classifiers with a probabilistic thresholding strategy. *Pattern Recognition*, 45(2):876–883, 2012.
- Mark D Reid and Robert C Williamson. Surrogate regret bounds for proper losses. In *Proceedings of the 26th Intl. Conf. on Machine Learning*, pages 897–904. ACM, 2009.
- A. Vedaldi. A MATLAB wrapper of SVM^{struct}. <http://www.vlfeat.org/~vedaldi/code/svm-struct-matlab>, 2011.
- Willem Waegeman, Krzysztof Dembczynski, Arkadiusz Jachnik, Weiwei Cheng, and Eyke Hüllermeier. On the bayes-optimality of F-measure maximizers. *Journal of Machine Learning Research*, 15:3333–3388, 2014.
- Hong Wang, Wei Xing, Kaiser Asif, and Brian Ziebart. Adversarial prediction games for multivariate losses. In *Advances in Neural Information Processing Systems*, pages 2710–2718, 2015.
- Nan Ye, Kian Ming A Chai, Wee Sun Lee, and Hai Leong Chieu. Optimizing F-measures: a tale of two approaches. In *Proceedings of the Intl. Conf. on Machine Learning*, 2012.