

Optimizing Non-decomposable Performance Measures: A Tale of Two Classes

Harikrishna Narasimhan*
Indian Institute of Science
harikrishna@csa.iisc.ernet.in

Purushottam Kar
Microsoft Research India
t-purkar@microsoft.com

Prateek Jain
Microsoft Research India
prajain@microsoft.com

Abstract

Modern classification problems frequently present mild to severe label imbalance as well as specific requirements on classification characteristics, and require optimizing performance measures that are non-decomposable over the dataset, such as F-measure. Such measures have spurred much interest and pose specific challenges to learning algorithms since their non-additive nature precludes a direct application of well-studied large scale optimization methods such as stochastic gradient descent.

In this paper we reveal that for two large families of performance measures that can be expressed as functions of true positive/negative rates, it is indeed possible to implement point stochastic updates. The families we consider are concave and pseudo-linear functions of TPR, TNR which cover several popularly used performance measures such as F-measure, G-mean and H-mean.

Our core contribution is an adaptive linearization scheme for these families, using which we develop optimization techniques that enable truly point-based stochastic updates. For concave performance measures we propose **SPADE**, a stochastic primal dual solver; for pseudo-linear measures we propose **STAMP**, a stochastic alternate maximization procedure. Both methods have crisp convergence guarantees, demonstrate significant speedups over existing methods - often by an order of magnitude or more, and give similar or more accurate predictions on test data.

1 Introduction

Learning applications with binary classification problems involving severe label imbalance abound, often accompanied with specific requirements in terms of false positive, or negative rates. Examples included spam classification, anomaly detection, and medical applications. Class imbalance is also often introduced as a result of the reduction of a problem to binary classification, such as in multi-class problems Bishop [2006] and multi-label problems due to extreme label sparsity Hsu et al. [2009].

Traditional performance measures such as misclassification rate are ill-suited in such situations as it is usually trivial to optimize them by constantly predicting the majority class. Instead, the performance measures of choice in such cases are those that perform a more holistic evaluation over the entire data. Naturally, these performance measures are *non-decomposable* over the dataset and cannot be expressed as a sum of errors on individual data points. Popular examples include F-measure, G-mean, H-mean etc.

A consistent effort directed at optimizing these performance measures has, over the years, resulted in the development of two broad approaches - 1) surrogate based approaches (e.g. SVMPerf Joachims et al. [2009]) that design convex surrogates for these performance measures, and 2) indirect approaches which include cost-sensitive classification-based approaches Parambath et al. [2014] which solve weighted classification problems, and plug-in approaches Koyejo et al. [2014], Narasimhan et al. [2014] which rely on consistent estimates of class probabilities.

Both these approaches are known to work fairly well on small datasets but do not scale very well to large ones, especially those large enough to not even fit in memory. SVMPerf-style approaches, which employ cutting plane methods do not scale well. On the other hand, plug-in approaches first need to solve a class probability estimation problem optimally and then tune a threshold. This two-stage approach prevents the method from exploiting better classifiers

* Work done while H.N. was an intern at Microsoft Research India, Bangalore.

to automatically obtain better thresholds. Moreover, for multi-class problems with C classes, jointly estimating C parameters can take time exponential in C .

For large datasets, streaming methods such as stochastic gradient descent Shalev-Shwartz et al. [2011] that take only a few passes over the entire data are preferable. However, traditional SGD techniques cannot handle non-decomposable losses. Recently, Kar et al. [2014] proposed optimizing SVMPerf-style surrogates using SGD techniques. Although their method is generic, allowing optimization of performance measures such as F-measure and partial AUC, they require maintaining a *large buffer* to compute online gradient estimates that can be prohibitive.

Motivated by the state of the art, we develop novel methods for optimizing two broad families of non-decomposable performance measures. Our methods incorporate truly point-wise updates, i.e. do not require a buffer, and require only a few passes over data. At an intuitive level, at the core of our work are adaptive linearization strategies for these performance measures, which make these measures amenable to SGD-style point-wise updates. Moreover, our linearizations are able to feed off the improvements made in learning a better classifier, resulting in faster convergence.

We consider two classes of performance measures

Concave Performance Measures (see Table 1): These measures can be written as concave functions of true positive (TPR) and negative (TNR) rates and include G-mean, H-mean etc. We exploit the dual structure of these functions via their Fenchel dual to linearize them in terms of the TPR, TNR variables. Our method then, in parallel, tunes the dual variables in this linearization and maximizes the weighted TPR-TNR combination. These updates are done in an online fashion using stochastic mirror descent steps.

Pseudo-linear Performance Measures (see Table 2): These measures can be written as fractional linear functions of TPR, TNR and include F-measure and the Jaccard coefficient. These functions need not be concave and the techniques outlined above do not apply. Instead, we exploit the pseudo-linear structure to linearize the function and develop a technique to alternately optimize the combination weights and the classifier model via stochastic updates. Although such “alternate-maximization” strategies in general need not converge even to a local optima, we show that our strategy converges to an ϵ -approximate global optimum after $\log(\frac{1}{\epsilon})$ batch updates or $O(1/\epsilon^2)$ stochastic updates.

Finally, we present an empirical validation of our methods. Our experiments reveal that for a range of performance measures in both classes, our methods can be significantly faster than either plug-in or SVMPerf-style methods, as well as give higher or comparable accuracies.

2 Related Works

As noted in Section 1, existing methods for optimizing performance measures that we study can be divided into surrogate-based approaches and indirect approaches based on cost-sensitive classification or plug-in methods. A third approach applicable to certain performance measures is the *decision-theoretic* method that learns a class probability estimate and computes predictions that maximize the expected value of the performance measure on a test set Lewis [1995], Ye et al. [2012]. In addition to these there exist methods dedicated to specific performance measures.

For instance Parambath et al. [2014] focus on optimizing F-measure by exploiting the pseudo-linearity of the function along with a cross validation-based strategy. Our **STAMP** method, on the other hand uses an alternating maximization strategy that does not require cross validation which considerably improves training time (see Figure 3). It is important to note that these performance measures have also been studied in multi-label settings where these no longer remain non-decomposable. For instance, Dembczyński et al. [2013] study plug-in style methods for maximizing F-measure in multi-label settings whereas works such as Koyejo et al. [2014], Narasimhan et al. [2014], Ye et al. [2012] study plug-in approaches for the same problem in the more challenging binary classification setting.

Historically, online learning algorithms have played a key role in designing solvers for large-scale batch problems. However, for non-decomposable loss functions, defining an online learning framework and providing efficient algorithms with small regret itself is challenging. Rakhlin et al. [2011] propose a generic method for such loss functions; however the algorithms proposed therein run in exponential time. Kar et al. [2014] also study such measures with the aim of designing stochastic gradient-style methods. However, their methods require a large buffer to be maintained, which causes them to have poorer convergence guarantees and in practice be slower than our methods.

By exploiting the special structure in our function classes, we are able to do away with such requirements. Our methods make use of standard online convex optimization primitives Zinkevich [2003]. However, their application requires special care in order to avoid divergent behavior.

Table 1: List of concave performance measures $\Psi(P, N)$ along with their monotonicity and Lipschitz properties, sufficient dual regions, and expressions for dual subgradients. $\mathcal{B}(\mathbf{0}, r)$ denotes the ball of radius r around the origin. \mathbb{R}_+^2 denotes the positive quadrant.

Name	Expression (P, N)	Mon.?	Lip.?	$\delta(\epsilon)$	Sufficient dual Region \mathcal{A}_Ψ	$\nabla\Psi^*(\alpha, \beta)$.
Min (Vincent [1994])	$\min\{P, N\}$	✓	✓	ϵ	$\{\alpha + \beta = 1\} \cap \mathbb{R}_+^2$	$\mathbf{0}$
H-mean (Kennedy et al. [2010])	$\frac{2PN}{P+N}$	✓	✓	4ϵ	$\{\sqrt{\alpha} + \sqrt{\beta} \geq \sqrt{2}\} \cap \mathcal{B}(\mathbf{0}, 2)$	$\mathbf{0}$
Q-mean (Liu and Chawla [2011])	$1 - \sqrt{\frac{(1-P)^2 + (1-N)^2}{2}}$	✓	✓	ϵ	$\{\alpha^2 + \beta^2 \leq 1/2\} \cap \mathbb{R}_+^2$	$\mathbf{1}$
G-mean (Daskalaki et al. [2006])	\sqrt{PN}	✓	✗	$3\sqrt{\epsilon}$	$\{\alpha\beta \geq 1/4\} \cap \mathbb{R}_+^2$	$\mathbf{0}$

3 Problem Setting

Let $\mathcal{X} \subset \mathbb{R}^d$ denote the instance space and $\mathcal{Y} = \{-1, +1\}$ the label space, with some distribution \mathcal{D} over $\mathcal{X} \times \mathcal{Y}$. Let $p := \Pr_{(\mathbf{x}, y) \sim \mathcal{D}} [y = +1]$ denote the proportion of positives in the population. Let $\mathcal{T} = \{(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_T, y_T)\}$ denote a sample of training points sampled i.i.d. from \mathcal{D} . For sake of simplicity we shall present our algorithms and analyses for a set of linear models $\mathcal{W} \subseteq \mathbb{R}^d$. Let $R_{\mathcal{X}}$ and $R_{\mathcal{W}}$ denote the radii of the domain \mathcal{X} and hypothesis class \mathcal{W} respectively.

We consider performance measures that can be expressed in terms of the true positive and negative rates of a classifier. To represent these measures, we shall use the notion of a *reward function* r that assigns a *reward* $r(y, \hat{y})$ to a prediction $\hat{y} \in \mathbb{R}$ when the true label is $y \in \mathcal{Y}$. We will use

$$r^+(\mathbf{w}; \mathbf{x}, y) = \frac{1}{p} \cdot r(y, \mathbf{w}^\top \mathbf{x}) \cdot \mathbf{1}(y = 1)$$

$$r^-(\mathbf{w}; \mathbf{x}, y) = \frac{1}{1-p} \cdot r(y, \mathbf{w}^\top \mathbf{x}) \cdot \mathbf{1}(y = -1)$$

to calculate rewards on positive and negative points. Since $\mathbb{E}_{(\mathbf{x}, y)} [r^+(\mathbf{w}; \mathbf{x}, y)] = \mathbb{E}_{(\mathbf{x}, y)} [r(y, \mathbf{w}^\top \mathbf{x}) | y = 1]$, setting $r^{0-1}(y, \hat{y}) = \mathbf{1}(y\hat{y} > 0)$ gives us $\mathbb{E}_{(\mathbf{x}, y)} [r^+(\mathbf{w}; \mathbf{x}, y)] = \text{TPR}(\mathbf{w})$. For sake of convenience, we will use $P(\mathbf{w}) = \mathbb{E}_{(\mathbf{x}, y)} [r^+(\mathbf{w}; \mathbf{x}, y)]$ and $N(\mathbf{w}) = \mathbb{E}_{(\mathbf{x}, y)} [r^-(\mathbf{w}; \mathbf{x}, y)]$ to denote population averages of the reward functions. We shall assume that our reward functions are concave, L_r -Lipschitz, and take values in a bounded range $[-B_r, B_r]$.

4 Concave Performance Measures

The first class of performance measures we analyze are concave performance measures. These measures can be written as concave functions of the TPR and TNR i.e.

$$\mathcal{P}_\Psi(\mathbf{w}) = \Psi(P(\mathbf{w}), N(\mathbf{w}))$$

for some concave link function $\Psi : \mathbb{R}^2 \rightarrow \mathbb{R}$. A large number of popular performance measures fall in this family since these measures are relevant in situations with severe label imbalance or in situations where cost-sensitive classification is required such as detection theory Vincent [1994]. Table 1 gives a list of such performance measures along with some of their relevant properties and references to works that utilize these performance measures.

We shall find it convenient to define the (concave) Fenchel conjugate of the link functions for our performance measures. For any concave function Ψ and $\alpha, \beta \in \mathbb{R}$, define

$$\Psi^*(\alpha, \beta) = \inf_{u, v \in \mathbb{R}} \{\alpha u + \beta v - \Psi(u, v)\}.$$

By the concavity of Ψ , we have, for any $u, v \in \mathbb{R}$,

$$\Psi(u, v) = \inf_{\alpha, \beta \in \mathbb{R}} \{\alpha u + \beta v - \Psi^*(\alpha, \beta)\}.$$

We shall use the notation Ψ to denote, both the link function, as well as the performance measure it induces.

Algorithm 1 SPADE: Stochastic Primal-Dual mEthod

Input: Primal/dual step sizes η_t, η'_t , feasible sets \mathcal{W}, \mathcal{A} **Output:** Classifier $\mathbf{w} \in \mathcal{W}$

```
1:  $\mathbf{w}_0 \leftarrow \mathbf{0}, t \leftarrow 1$ 
2: while data stream has points do
3:   Receive data point  $(\mathbf{x}_t, y_t)$ 
4:   /* Perform primal ascent */
5:   if  $y_t > 0$  then
6:      $\mathbf{w}_{t+1} \leftarrow \Pi_{\mathcal{W}}(\mathbf{w}_t + \eta_t \cdot \alpha_t \nabla_{\mathbf{w}} r^+(\mathbf{w}_t; \mathbf{x}_t, y_t))$ 
7:   else
8:      $\mathbf{w}_{t+1} \leftarrow \Pi_{\mathcal{W}}(\mathbf{w}_t + \eta_t \cdot \beta_t \nabla_{\mathbf{w}} r^-(\mathbf{w}_t; \mathbf{x}_t, y_t))$ 
9:   end if
10:  /* Perform dual descent */
11:   $(a, b) \leftarrow (\alpha_t, \beta_t) - \eta'_t \cdot \nabla_{(\alpha, \beta)} \Psi^*(\alpha_t, \beta_t)$ 
12:  if  $y_t > 0$  then
13:     $a \leftarrow a - \eta'_t \cdot r^+(\mathbf{w}_t; \mathbf{x}_t, y_t)$ 
14:  else
15:     $b \leftarrow b - \eta'_t \cdot r^-(\mathbf{w}_t; \mathbf{x}_t, y_t)$ 
16:  end if
17:   $(\alpha_{t+1}, \beta_{t+1}) \leftarrow \Pi_{\mathcal{A}}((a, b))$ 
18:   $t \leftarrow t + 1$ 
19: end while
20: return  $\bar{\mathbf{w}} = \frac{1}{t} \sum_{\tau=1}^t \mathbf{w}_{\tau}$ 
```

4.1 A Stochastic Primal-dual Method for Optimizing Concave Performance Measures

We now present a novel online stochastic method for optimizing the class of concave performance measures. The use of stochastic gradient techniques for these measures presents specific challenges due to the non-decomposable nature of these measures which makes it difficult to obtain cheap, unbiased estimates of the gradient using a single point. Recent works Kar et al. [2013, 2014] have tried to resolve this issue by looking at mini-batch methods or by using a buffer to maintain a sketch of the stream. However, such techniques bring in a bias into the learning algorithm in the form of buffer size or mini batch length which results in slower convergence. Indeed, the **IPMB** method of Kar et al. [2014] is only able to guarantee a $\bar{\mathcal{O}}(\frac{1}{\sqrt[4]{T}})$ rate of convergence, whereas SGD techniques are usually able to guarantee $\bar{\mathcal{O}}(\frac{1}{\sqrt{T}})$ rates. This is indicative of suboptimal performance and our experiments confirm this (see Figure 3).

Here we show that for the class of concave performance measures, such workarounds are not necessary. To this end we present the **SPADE** algorithm (Algorithm 1) which exploits the dual structure of the performance measures to obtain efficient point updates which do not require the use of mini-batches or online buffers. **SPADE** is able to offer convergence guarantees identical to those that stochastic methods offer for additive performance measures such as least squares, without the presence of any algorithmic bias.

Let $\mathcal{W} \subset \mathcal{X}$ and $\mathcal{A}_{\Psi} \subset \mathbb{R}^2$ be convex regions within the model and dual spaces respectively, and $\Pi_{\mathcal{W}}$ and $\Pi_{\mathcal{A}_{\Psi}}$ denote projection operators for these. Table 1 lists the relevant dual regions for the performance measures listed therein.

4.2 Convergence Analysis for SPADE

This section presents a convergence analysis for the **SPADE** algorithm. The convergence proof is formally stated in Theorem 4. Apart from demonstrating the utility of the algorithm, the proof also sheds light on the choice of algorithm parameters, such as primal/dual feasible regions.

We shall work with performance measures that are monotonically increasing in the true positive and negative rates of the classifier i.e. if $u \geq u', v \geq v'$ then $\Psi(u, v) \geq \Psi(u', v')$. This is a natural assumption and is satisfied by all performance measures considered here (see Table 1). We now introduce two useful concepts.

Definition 1 (Stable Performance Measure). *A performance measure Ψ will be called δ -stable if for some function*

$\delta : \mathbb{R} \rightarrow \mathbb{R}$, we have for all $u, v \in \mathbb{R}$ and $\epsilon \in \mathbb{R}_+$,

$$\Psi(u + \epsilon, v + \epsilon) \leq \Psi(u, v) + \delta(\epsilon).$$

Table 1 lists the stability parameters of all the concave performance measures. Clearly, a performance measure has a linear stability parameter i.e. $\delta(\epsilon) \leq L \cdot \epsilon$ iff its corresponding link function is Lipschitz. We now define the notion of a *sufficient dual region* for a performance measure

Definition 2 (Sufficient Dual Region). *For any link function Ψ , define its sufficient dual region $\mathcal{A}_\Psi \subseteq \mathbb{R}^2$ to be the minimal set such that for all $(u, v) \in \mathbb{R}^2$, we have*

$$\Psi(u, v) = \inf_{(\alpha, \beta) \in \mathcal{A}_\Psi} \{\alpha u + \beta v - \Psi^*(\alpha, \beta)\}.$$

The reason for defining this quantity will become clear in a moment. A closer look at Algorithm 1 indicates that it is performing online gradient descent steps with the dual variables. Clearly, for this procedure to have statistical convergence properties, the magnitude of the updates must be bounded in some sense otherwise the learning procedure may diverge. This motivates the projection step in Step 17. However, in order for the updated dual variables to be informative about the current primal function value, the projection step must be done in a way that does not distort the link function. The notion of a sufficient dual region formally captures the notion of such a projection step.

Having said that, there is no apriori guarantee that the sufficient region for a given performance measure would be bounded, in which case this entire exercise counts for naught. However, the following lemma, by closely linking the stability properties of a performance measure with the size of its sufficient dual region, shows that for well-behaved link functions, this will not be the case .

Lemma 3. *The stability parameter of a performance measure $\Psi(\cdot)$ can be written as $\delta(\epsilon) \leq L_\Psi \cdot \epsilon$ iff its sufficient dual region is bounded in a ball of radius L_Ψ .*

The proof of this result follows from elementary manipulations and can be found in Appendix A. In some sense this result can be seen as a realization of the well known connection between the Fenchel dual of a function and its Lipschitz properties.

To simplify the initial analysis, we shall first concentrate on performance measures whose link functions are Lipschitz. It is easy to see that these are exactly the performance measures whose gradients do not diverge within any compact region of the real plane. Of the performance measures listed in Table 1, all measures except G-mean have associated link functions that are Lipschitz. Subsequently, we shall address the more involved case of non-Lipschitz performance measures such as G-mean as well.

Theorem 4. *Suppose we are given a stream of random samples $(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_T, y_T)$ drawn from a distribution \mathcal{D} over $\mathcal{X} \times \mathcal{Y}$. Let $\Psi(\cdot)$ be a concave, Lipschitz link function. Let Algorithm 1 be executed with a dual feasible set $\mathcal{A} \supseteq \mathcal{A}_\Psi$, $\eta_t = 1/\sqrt{t}$ and $\eta'_t = 1/\sqrt{t}$. Then, the average model $\bar{\mathbf{w}} = \frac{1}{T} \sum_{t=1}^T \mathbf{w}_t$ output by the algorithm satisfies, with probability at least $1 - \delta$,*

$$\mathcal{P}_\Psi(\bar{\mathbf{w}}) \geq \sup_{\mathbf{w}^* \in \mathcal{W}} \mathcal{P}_\Psi(\mathbf{w}^*) - \mathcal{O} \left(\delta_\Psi \left(\sqrt{\frac{1}{T} \log \frac{1}{\delta}} \right) \right).$$

We refer the reader to Appendix B for a proof and explicit constants. The proof closely analyzes the primal ascent and dual descent steps, tying them together using the Fenchel dual of Ψ .

4.3 The Case of non-Lipschitz Link Functions

Non-Lipschitz link functions, such as the one used in the G-mean performance measure, pose a particular challenge to the previous analysis. Owing to their non-Lipschitz nature, their sufficient dual region is unbounded. Indeed as Table 1 indicates, the sufficient region for $\Psi_{\text{G-mean}}$ extends indefinitely along both coordinate axes. More precisely, what happens is that the gradients of the $\Psi_{\text{G-mean}}$ function diverge as either $u \rightarrow 0$, or $v \rightarrow 0$. This poses a stumbling block for the proof of Theorem 4 since the regret and online-to-batch conversion results used therein fail.

Table 2: List of pseudo-linear performance measures $\mathcal{P}_{(a,b)}(P, N)$ along with their popular forms, canonical expressions in terms of (reward functions representative of) true positive (P) and negative (N) rates, monotonicity properties, acceptable range of reward values, and rate of convergence of the Alt-Max procedure for the performance measure when rewards take values in the range $[0, m)$.

Name	Popular Form	Canonical Form (P, N)	Mon.?	Range P, N	Rate $\eta(m)$
F $_{\beta}$ -measure (Manning et al.)	$\frac{(1+\beta^2) \cdot TP}{(1+\beta^2) \cdot TP + \beta^2 \cdot FP + FP}$	$\frac{(1+\beta^2) \cdot P}{\beta^2 + \theta + P - \theta \cdot N}$	✓	$(0, 1 + \frac{\beta^2}{\theta})$	$\frac{m(1+\theta)}{m+\beta^2+\theta}$
Jaccard Coefficient (Koyejo et al.)	$\frac{TP}{TP+FP+FN}$	$\frac{P}{1+\theta-\theta \cdot N}$	✓	$(0, \frac{1+\theta}{\theta})$	$\frac{m\theta}{1+\theta}$
Gower-Legendre $_{\sigma < 1}$ (Sokolova and Lapalme)	$\frac{TP+TN}{TP+\sigma(FP+FN)+TN}$	$\frac{P+\theta \cdot N}{\sigma(1+\theta)+(1-\sigma) \cdot P+\theta(1-\sigma) \cdot N}$	✓	$(0, \infty)$	$\frac{(1-\sigma)m}{(1-\sigma)m+\sigma}$
Gower-Legendre $_{\sigma > 1}$ (Sokolova and Lapalme)	$\frac{TP+TN}{TP+\sigma(FP+FN)+TN}$	$\frac{P+\theta \cdot N}{\sigma(1+\theta)+(1-\sigma) \cdot P+\theta(1-\sigma) \cdot N}$	✓	$(0, \frac{\sigma}{\sigma-1})$	$\frac{(\sigma-1)m}{\sigma}$

A natural way to solve this problem is to ensure that the reward functions r^+, r^- always assign rewards that are bounded away from zero. More specifically, for some $\epsilon > 0$, we have $r^+(\mathbf{w}; \mathbf{x}, y), r^-(\mathbf{w}; \mathbf{x}, y) \geq \epsilon$ for all $\mathbf{w} \in \mathcal{W}$ and $\mathbf{x} \in \mathcal{X}$. For this restricted reward region, one can show, using Lemma 3, that the sufficient dual region can be restricted to a ball of radius $\mathcal{O}(\sqrt{1/\epsilon})$.

The above discussion suggests that we regularize the reward function i.e. at each time step t , we add a small value $\epsilon(t)$ to the original reward function. However, the amount of regularization remains to be decided since over regularization could cause our resulting excess risk bound to be vacuous with respect to the original reward function. It turns out that setting $\epsilon(t) \approx \frac{1}{t^{1/4}}$ strikes a fine balance between regularization and fidelity to the original reward function - this seems intuitive since the regularization becomes milder and milder as learning progresses. The following extension of Theorem 4 formalizes this statement.

Theorem 5. *Suppose we have the problem setting in Theorem 4 with the $\Psi_{G\text{-mean}}$ performance measure being optimized for. Consider a modification to Algorithm 1 wherein the reward functions are changed to $r_t^+(\cdot) = r^+(\cdot) + \epsilon(t)$, and $r_t^-(\cdot) = r^-(\cdot) + \epsilon(t)$ for $\epsilon(t) = \frac{1}{t^{1/4}}$. Then, the average model $\bar{\mathbf{w}} = \frac{1}{T} \sum_{t=1}^T \mathbf{w}_t$ output by the algorithm satisfies, with probability at least $1 - \delta$,*

$$\mathcal{P}_{\Psi_{G\text{-mean}}}(\bar{\mathbf{w}}) \geq \sup_{\mathbf{w}^* \in \mathcal{W}} \mathcal{P}_{\Psi_{G\text{-mean}}}(\mathbf{w}^*) - \tilde{\mathcal{O}}\left(\frac{1}{T^{1/4}}\right).$$

The proof of this theorem can be found in Appendix C. We note here that primal dual frameworks have been utilized before in diverse areas such as distributed optimization Jaggi et al. [2014] and multi-objective optimization Mahdavi et al. [2013]. However, these works simply assume the functions involved therein to be Lipschitz and/or smooth and do not address cases where they fail to be so. Theorem 5 on the other hand, is able to recover a non-trivial, albeit weaker, statement even for *locally Lipschitz* functions.

5 Pseudo-linear Performance Measures

The second class of performance measures we analyze are pseudo-linear performance measures. These measures have a fractional linear function as the link function and can be written as follows:

$$\mathcal{P}_{(a,b)}(\mathbf{w}) = \frac{a_0 + a_1 \cdot P(\mathbf{w}) + a_2 \cdot N(\mathbf{w})}{b_0 + b_1 \cdot P(\mathbf{w}) + b_2 \cdot N(\mathbf{w})},$$

for some weighing coefficients \mathbf{a}, \mathbf{b} . Several popularly used performance measures, most notably the F-measure, can be represented as pseudo-linear functions. Table 2 enumerates some popular pseudo-linear performance measures as well as their properties.

We note that these performance measures are usually represented in literature using the entries of the confusion matrix. However, for the sake of our analysis, we shall find it useful to represent them in terms of the true positive and true negative rates. To do so, we shall use p to denote the proportion of positives in the population and $\theta = \frac{1-p}{p}$ to denote the label skew.

Algorithm 2 AMP: Alternate Maximization Procedure

Input: Performance measure $\mathcal{P}_{(\mathbf{a}, \mathbf{b})}$, feasible set \mathcal{W} , tolerance ϵ

Output: An ϵ -optimal classifier $\mathbf{w} \in \mathcal{W}$

- 1: Construct valuation function $V_{(\mathbf{a}, \mathbf{b})}$
 - 2: $\mathbf{w}_0 \leftarrow \mathbf{0}, t \leftarrow 1$
 - 3: **while** $v_t > v_{t-1} + \epsilon$ **do**
 - 4: $\mathbf{w}_{t+1} \leftarrow \arg \max_{\mathbf{w} \in \mathcal{W}} V_{(\mathbf{a}, \mathbf{b})}(\mathbf{w}, v_t)$
 - 5: $v_{t+1} \leftarrow \arg \max_{v > 0} v$ such that $V_{(\mathbf{a}, \mathbf{b})}(\mathbf{w}_{t+1}, v) \geq v$
 - 6: $t \leftarrow t + 1$
 - 7: **end while**
 - 8: **return** \mathbf{w}_t
-

5.1 Alternate-maximization for Optimizing Pseudo-linear Performance Measures

Pseudo-linear functions are named so since their level sets can be defined using linear half-spaces. More specifically, every pseudo-linear function Ψ over \mathbb{R}^d has an associated “level-finder” function $a : \mathbb{R} \rightarrow \mathbb{R}^d$ and $b : \mathbb{R} \rightarrow \mathbb{R}$ such that $\Psi(\mathbf{v}) \geq t$ iff $\langle \mathbf{v}, a(t) \rangle \geq b(t)$. We refer the reader to Parambath et al. [2014] for a more relaxed introduction to these functions and their properties. For our purposes, however, it suffices to notice that this property immediately points toward a cost-sensitive method to optimize these performance measures.

This fact was noticed by Parambath et al. [2014] who exploited this to develop a cost-sensitive classification method for optimizing the F-measure by simply searching for the best weights with which to perform cost-sensitive classification. However, we notice that instead of performing such a brute force search, one can adaptively tune the weights to better and better values and obtain much faster convergence. To develop this intuition, we first define the notion of a *valuation function* below.

Definition 6 (Valuation Function). *The valuation function of a performance measure $\mathcal{P}_{(\mathbf{a}, \mathbf{b})}$, for a classifier $\mathbf{w} \in \mathcal{W}$, and at a level $v \in \mathbb{R}$ is defined as*

$$V_{(\mathbf{a}, \mathbf{b})}(\mathbf{w}, v) := c + (\alpha - v\gamma) \cdot P(\mathbf{w}) + (\beta - v\delta) \cdot N(\mathbf{w}),$$

where $c = \frac{a_0}{b_0}$, $\alpha = \frac{a_1}{b_0}$, $\beta = \frac{a_2}{b_0}$, $\gamma = \frac{b_1}{b_0}$, $\delta = \frac{b_2}{b_0}$.

The following well-known lemma closely links the valuation function to the performance measure.

Lemma 7. *For any performance measure $\mathcal{P}_{(\mathbf{a}, \mathbf{b})}$, $\mathbf{w} \in \mathcal{W}$ and $v \in \mathbb{R}$ we have $\mathcal{P}_{(\mathbf{a}, \mathbf{b})}(\mathbf{w}) \geq v$ iff $V_{(\mathbf{a}, \mathbf{b})}(\mathbf{w}, v) \geq v$. Moreover, in such a situation we say that classifier \mathbf{w} has achieved valuation at level v .*

Lemma 7 indicates that the performance of a classifier is intimately linked to its valuation. This suggests a natural alternate maximization approach wherein we alternate between posing a *challenge level* to the classifier and training a classifier to achieve that level. The resulting algorithm **AMP** is detailed in Algorithm 2. Note that using Lemma 7, step 5 in the algorithm can be executed simply by setting $v_{t+1} = \mathcal{P}_{(\mathbf{a}, \mathbf{b})}(\mathbf{w}_{t+1})$. Thus, in a very natural manner, the current classifier challenges the next classifier to beat its own performance. It turns out that this approach results in rapid convergence as outlined in the following theorem.

Theorem 8. *Let Algorithm 2 be executed with a performance measure $\mathcal{P}_{(\mathbf{a}, \mathbf{b})}$ and reward functions that offer values in the range $[0, m)$. Let $\mathcal{P}^* := \sup_{\mathbf{w} \in \mathcal{W}} \mathcal{P}_{(\mathbf{a}, \mathbf{b})}(\mathbf{w})$. Also let $\Delta_t = \mathcal{P}^* - \mathcal{P}_{(\mathbf{a}, \mathbf{b})}(\mathbf{w}_t)$ be the excess error for the model \mathbf{w}_t generated at time t . Then there exists a value $\eta(m) < 1$ such that $\Delta_t \leq \Delta_0 \cdot \eta(m)^t$.*

The proof of this theorem can be found in Appendix D. Table 2 gives values for the convergence rates of all the pseudo-linear performance measures, as well as the allowed range of values that the reward functions can take for those measures. This is important since performance measures such as the F-measure diverge if the reward function values approach 2. Other performance measures like the Gower-Legendre measure do not impose any such restrictions. Note that the above result shows that Algorithm 2 will always terminate in $\mathcal{O}(\log \frac{1}{\epsilon})$ steps.

At this point it would be apt to make a historical note. Pseudo-linear functions have enjoyed a fair amount of interest in the optimization community Schaible [1976], Dinkelbach [1967], Jagannathan [1966] within the sub-field of

Algorithm 3 STAMP: STochastic Alt-Max Procedure

Input: Feasible set \mathcal{W} , Step sizes η_t , epoch lengths s_e, s'_e
Output: Classifier $\mathbf{w} \in \mathcal{W}$

```

1:  $v \leftarrow 0, t \leftarrow 0, e \leftarrow 0, \mathbf{w}_0 \leftarrow \mathbf{0}$ 
2: repeat
3:   /* Model optimization stage */
4:    $\tilde{\mathbf{w}} \leftarrow \mathbf{w}_e$ 
5:   while  $t < s_e$  do
6:     Receive sample  $(\mathbf{x}, y)$ 
7:      $\tilde{\mathbf{w}} \leftarrow \tilde{\mathbf{w}} + \eta_t \nabla_{\mathbf{w}} \left( (1 - \frac{v_e}{2}) r^+(\tilde{\mathbf{w}}; \mathbf{x}, y) + \frac{v_e}{2} r^-(\tilde{\mathbf{w}}; \mathbf{x}, y) \right)$ 
8:      $t \leftarrow t + 1$ 
9:   end while
10:   $t \leftarrow 0, e \leftarrow e + 1, \mathbf{w}_{e+1} \leftarrow \tilde{\mathbf{w}}$ 
11:  /* Challenge level estimation stage */
12:   $v_+ \leftarrow 0, v_- \leftarrow 0$ 
13:  while  $t < s'_e$  do
14:    Receive sample  $(\mathbf{x}, y)$ 
15:     $v_y \leftarrow v_y + r^y(\mathbf{w}_e; \mathbf{x}, y)$ 
16:     $t \leftarrow t + 1$ 
17:  end while
18:   $t \leftarrow 0, v_e \leftarrow \frac{2v_+}{2+v_+-v_-}$ 
19: until stream is exhausted
20: return  $\mathbf{w}_e$ 

```

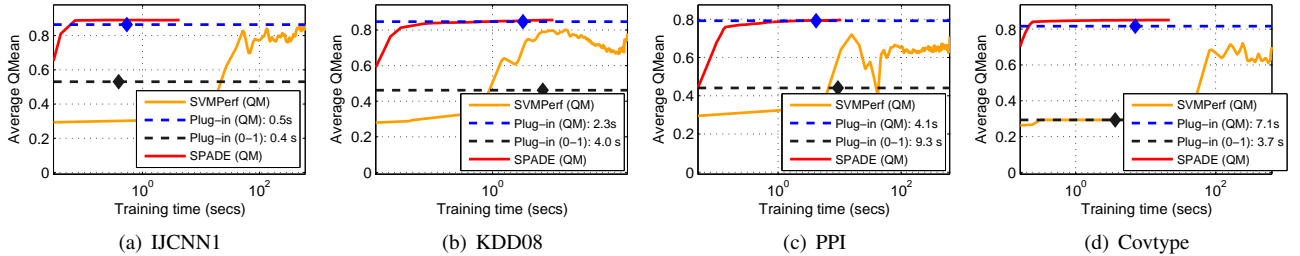


Figure 1: Comparison of stochastic primal-dual method (SPADE) with baseline methods on QMean maximization tasks. SPADE achieves similar/better accuracies while consistently requiring about 3-4x less time than other baseline approaches.

fractional programming. Of the many methods that have been developed to optimize these functions, the Dinkelbach-Jagannathan (DJ) procedure Dinkelbach [1967], Jagannathan [1966] is of specific interest to us. It turns out that the **AMP** method can be seen as performing DJ-style updates over parameterized spaces (the parameter being the model \mathbf{w}). It is known (for instance see Schaible [1976]) that the DJ process is able to offer a linear convergence rates. Our proof of Theorem 8, which was obtained independently, can then be seen as giving a similar result in the parameterized setting.

However, we wish to move one step further and optimize these performance measures in an online stochastic manner. To this end, we observe that the **AMP** algorithm can be executed in an online fashion by using stochastic updates to train the intermediate models. The resulting algorithm **STAMP**, is presented in Algorithm 3. However, this algorithm is much harder to analyze because unlike **AMP** which has the luxury of offering exact updates, **STAMP** offers inexact, even noisy updates. Indeed, even existing works in the optimization community (for example Schaible [1976]) do not seem to have analyzed DJ-style methods with noisy updates.

Our next contribution hence, is an analysis of the convergence rate offered by the **AMP** algorithm when neither of the two maximizations is carried out exactly. For the sake of simplicity, we present the **STAMP** algorithm and its

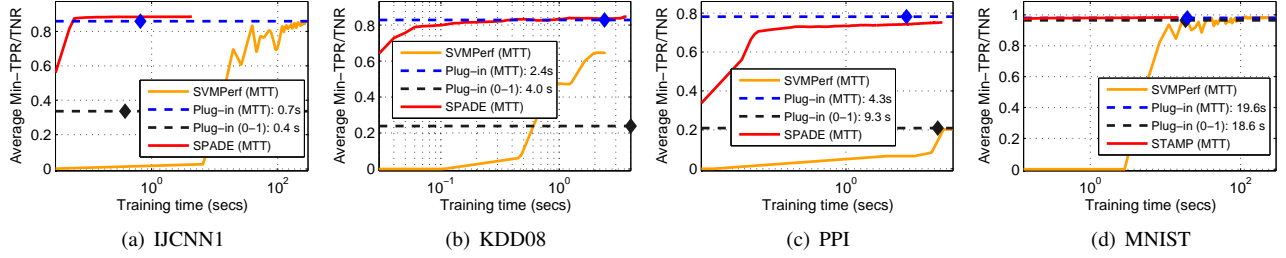


Figure 2: Comparison of stochastic primal-dual method (SPADE) with baseline methods on Min-TPR/TNR maximization tasks

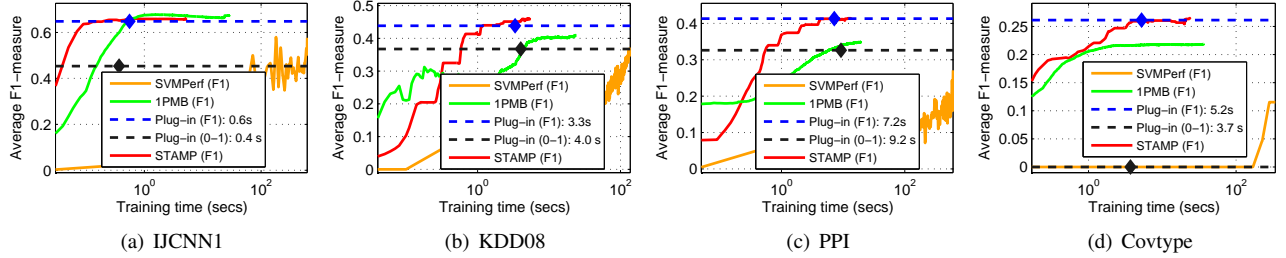


Figure 3: Comparison of stochastic alternating minimization procedure (STAMP) with baseline methods on F1 maximization tasks

analysis for the case of F_1 measure. Suppose at each time step, for some $\epsilon_t \geq 0, \delta_t$, we have

$$\begin{aligned} V(\mathbf{w}_{t+1}, v_t) &= \max_{\mathbf{w} \in \mathcal{W}} V(\mathbf{w}, v_t) - \epsilon_t \\ v_t &= F(\mathbf{w}_t) + \delta_t, \end{aligned}$$

then for some $\eta < 1$, we have

$$\Delta_T \leq \eta^T \Delta_0 + \sum_{i=0}^{T-1} \eta^{T-i} (|\delta_t| + \epsilon_t)$$

As a corollary we present a convergence analysis for the **STAMP** algorithm in Theorem 9.

Theorem 9. *Let Algorithm 3 be executed with a performance measure $\mathcal{P}_{(a,b)}$ and reward functions with range $[0, m)$. Let $\eta = \eta(m)$ be the rate of convergence guaranteed for $\mathcal{P}_{(a,b)}$ by the **AMP** algorithm. Set the epoch lengths to $s_e, s'_e = \tilde{O}\left(\frac{1}{\eta^{2e}}\right)$. Then after $e = \log_{\frac{1}{\eta}}\left(\frac{1}{\epsilon} \log^2 \frac{1}{\epsilon}\right)$ epochs, we can ensure with probability at least $1 - \delta$ that $\mathcal{P}^* - \mathcal{P}_{(a,b)}(\mathbf{w}_e) \leq \epsilon$. Moreover the number of samples consumed till this point is at most $\tilde{O}\left(\frac{1}{\epsilon^2}\right)$.*

The convergence analysis for noisy **AMP** can be found in Appendix E. The proof of this theorem can be found in Appendix F. Both results require a fine grained analysis of how errors accumulate throughout the learning process.

6 Experimental Results

We shall now compare our methods with the state-of-the-art on various performance measures and datasets.

Datasets: We evaluated our methods on 5 publicly available benchmark datasets: a) PPI, b) KDD Cup 2008, c) IJCNN, d) Covtype, e) MNIST. All datasets exhibited moderate to severe label imbalance with the KDD Cup 2008 dataset having just 0.61% positives.

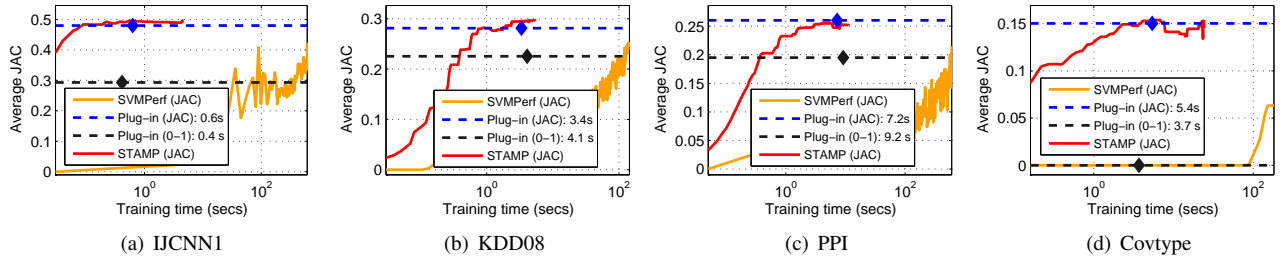


Figure 4: Comparison of stochastic alternating minimization procedure (STAMP) with baseline methods on JAC maximization tasks

Methods: We instantiated the **SPADE** algorithm (Algorithm 1) on the Q-mean and Min-TPR/TNR performance measures. We also instantiated the **STAMP** method (Algorithm 3) on F1-measure and the JAC coefficient. In both cases we compared to the SVMPerf method Joachims et al. [2009] and plug-in method Koyejo et al. [2014] specialized to these measures. For the sake of reference, we also compared to the standard logistic regression method for (unweighted) binary classification. Additionally for F1-measure, we also compared to the **IPMB** stochastic gradient descent method proposed recently by Kar et al. [2014]. All methods were implemented in C.

Parameters: We used 70% of the dataset for training and the rest for testing. Tunable parameters, including thresholds for the plug-in approaches, were cross-validated on a validation set. All results reported here were averaged over 5 random train-test splits. We used hinge-loss based reward functions for our methods. **STAMP** was executed by setting the challenge level to the actual F-measure/JAC at each stage. We used a state of the art LBFGS solver to implement the plug-in methods and used standard implementations of the SVMPerf algorithm. Since our methods are able to take a single pass over the data very rapidly, **SPADE** was allowed to run for 25 passes over the data and **STAMP** was allowed 25 passes with an initial epoch length of 100 which was doubled after every iteration. The SVMPerf algorithm was allowed a runtime of up to $50\times$ of that given to our method after which it was terminated. The LBFGS solver was always allowed to run till convergence.

Figures 1 and 2 compare the **SPADE** method with the baseline methods for the Q-mean and Min-TPR/TNR measures. In general, **SPADE** was found to offer comparable or superior accuracies with greatly accelerated convergence as compared to other methods. On the IJCNN and Covtype datasets, **SPADE** outperformed every other method by about 2-3%. As **SPADE** is a stochastic first order method, it is expected to rapidly find out a fairly accurate solution. Indeed, the method was found to offer greatly accelerated convergence without fail. For instance, on the MNIST dataset, **SPADE** found out the best solution as much as $60\times$ faster than any other method whereas on the KDD Cup and PPI datasets it was $12\times$ and $2\times$ faster respectively. The SVMPerf method, on the other hand, was found to be extremely slow in general and require at least an order of magnitude time more than **SPADE** to find reasonably accurate solutions. It is also notable that in all cases, simple binary classification gave very poor accuracies due to the severe label imbalance in these datasets.

Figures 3 and 4 report the performance of the **STAMP** method applied to pseudo-linear functions. Similar to the concave measures, **STAMP** was found to provide competitive accuracies as compared to the baseline methods but require at least $3-4\times$ less computational time. Interestingly, for the F1-measure, the **IPMB** method, which is another stochastic gradient descent-based method, was found to struggle to obtain accuracies similar to that of **STAMP** or else offer much slower convergence. We suspect two main reasons for the suboptimal behavior of this other stochastic method. Firstly these results confirm the adverse effect of the dependence on an in-memory buffer on these methods. It is notable that this dependence causes even the theoretical convergence rates for these methods to be weaker as was noted earlier in the discussion. Secondly, we note that both SVMPerf and **IPMB** optimize the same “struct-SVM” style surrogate for the F-measure Kar et al. [2014]. This surrogate has been observed to give poor accuracies when compared to plug-in methods in several previous works Koyejo et al. [2014], Narasimhan et al. [2014]. **STAMP** on the other hand, works directly with F-measure in a manner similar to, but faster than, the plug-in methods which might explain its better performance.

Acknowledgements

HN thanks support from a Google India PhD Fellowship.

References

- Christopher M. Bishop. *Pattern Recognition and Machine Learning*. Springer-Verlag New York, Inc., 2006.
- Nicoló Cesa-Bianchi, Alex Conconi, and Claudio Gentile. On the Generalization Ability of On-Line Learning Algorithms. In *15th Annual Conference on Neural Information Processing Systems (NIPS)*, pages 359–366, 2001.
- Sophia Daskalaki, Ioannis Kostas, and Nikolaos Avouris. Evaluation of Classifiers for an Uneven Class Distribution Problem. *Applied Artificial Intelligence*, 20:381–417, 2006.
- Krzysztof Dembczyński, Arkadiusz Jachnik, Wojciech Kotłowski, Willem Waegeman, and Eyke Hüllermeier. Optimizing the F-Measure in Multi-Label Classification: Plug-in Rule Approach versus Structured Loss Minimization. In *30th International Conference on Machine Learning (ICML)*, 2013.
- Werner Dinkelbach. On Nonlinear Fractional Programming. *Management Science*, 13(7, Series A, Sciences):492–498, 1967.
- Daniel Hsu, Sham Kakade, John Langford, and Tong Zhang. Multi-Label Prediction via Compressed Sensing. In *23rd Annual Conference on Neural Information Processing Systems (NIPS)*, pages 772–780, 2009.
- R. Jagannathan. On Some Properties of Programming Problems in Parametric Form Pertaining to Fractional Programming. *Management Science*, 12(7, Series A, Sciences):609–615, 1966.
- Martin Jaggi, Virginia Smith, Martin Takáč, Jonathan Terhorst, Sanjay Krishnan, Thomas Hofmann, and Michael I. Jordan. Communication-Efficient Distributed Dual Coordinate Ascent. In *28th Annual Conference on Neural Information Processing Systems (NIPS)*, pages 694–702, 2014.
- Thorsten Joachims, Thomas Finley, and Chun-Nam John Yu. Cutting-plane training of structural SVMs. *Machine Learning*, 77(1):27–59, 2009.
- Purushottam Kar, Bharath K Sriperumbudur, Prateek Jain, and Harish Karnick. On the Generalization Ability of Online Learning Algorithms for Pairwise Loss Functions. In *30th International Conference on Machine Learning (ICML)*, 2013.
- Purushottam Kar, Harikrishna Narasimhan, and Prateek Jain. Online and Stochastic Gradient Methods for Non-decomposable Loss Functions. In *28th Annual Conference on Neural Information Processing Systems (NIPS)*, pages 694–702, 2014.
- Kenneth Kennedy, Brian Mac Namee, and Sarah Jane Delany. Learning without default: a study of one-class classification and the low-default portfolio problem. In *International Conference on Artificial Intelligence and Cognitive Science (ICAICS)*, volume 6202 of *Lecture Notes in Computer Science*, pages 174–187, 2010.
- Oluwasanmi O. Koyejo, Nagarajan Natarajan, Pradeep K. Ravikumar, and Inderjit S. Dhillon. Consistent Binary Classification with Generalized Performance Metrics. In *28th Annual Conference on Neural Information Processing Systems (NIPS)*, pages 2744–2752, 2014.
- D.D. Lewis. Evaluating and optimizing autonomous text classification systems. In *18th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR)*, 1995.
- W. Liu and S. Chawla. A Quadratic Mean based Supervised Learning Model for Managing Data Skewness. In *11th SIAM International Conference on Data Mining (SDM)*, 2011.

- Mehrdad Mahdavi, Tianbao Yang, and Rong Jin. Stochastic Convex Optimization with Multiple Objectives. In *27th Annual Conference on Neural Information Processing Systems (NIPS)*, pages 1115–1123, 2013.
- C. D. Manning, P. Raghavan, and H. Schütze. *Introduction to Information Retrieval*. Cambridge University Press, 2008.
- Harikrishna Narasimhan, Rohit Vaish, and Shivani Agarwal. On the Statistical Consistency of Plug-in Classifiers for Non-decomposable Performance Measures. In *28th Annual Conference on Neural Information Processing Systems (NIPS)*, 2014.
- Shameem Puthiya Parambath, Nicolas Usunier, and Yves Grandvalet. Optimizing F-Measures by Cost-Sensitive Classification. In *28th Annual Conference on Neural Information Processing Systems (NIPS)*, pages 2123–2131, 2014.
- Alexander Rakhlin, Karthik Sridharan, and Ambuj Tewari. Online Learning: Beyond Regret. In *24th Annual Conference on Learning Theory (COLT)*, 2011.
- Siegfried Schaible. Fractional Programming. II, on Dinkelbach’s Algorithm. *Management Science*, 22(8):868–873, 1976.
- Shai Shalev-Shwartz, Yoram Singer, Nathan Srebro, and Andrew Cotter. Pegasos: primal estimated sub-gradient solver for SVM. *Math. Program.*, 127(1):3–30, 2011.
- Marina Sokolova and Guy Lapalme. A systematic analysis of performance measures for classification tasks. *Information Processing & Management*, 45(4):427–437, 2009.
- P.H. Vincent. *An Introduction to Signal Detection and Estimation*. Springer-Verlag New York, Inc., 1994.
- Nan Ye, Kian Ming A. Chai, Wee Sun Lee, and Hai Leong Chieu. Optimizing F-Measures: A Tale of Two Approaches. In *29th International Conference on Machine Learning (ICML)*, 2012.
- Martin Zinkevich. Online Convex Programming and Generalized Infinitesimal Gradient Ascent. In *20th International Conference on Machine Learning (ICML)*, pages 928–936, 2003.

A Proof of Lemma 3

Lemma 3. *The stability parameter of a performance measure $\Psi(\cdot)$ can be written as $\delta(\epsilon) \leq L_\Psi \cdot \epsilon$ iff its sufficient dual region is bounded in a ball of radius $\Theta(L_\Psi)$.*

Proof. Let us denote primal variables using the notation $\mathbf{x} = (u, v)$ and dual variables using the notation $\boldsymbol{\theta} = (\alpha, \beta)$. The proof follows from the fact that any value of $\boldsymbol{\theta}$ for which $\Psi^*(\boldsymbol{\theta}) = -\infty$ can be safely excluded from the sufficient dual region.

For proving the result in one direction suppose Ψ is stable with $\delta(\epsilon) = L\epsilon$ for some $L > 0$. Now consider some $\boldsymbol{\theta} \in \mathbb{R}^2$ such that $\|\boldsymbol{\theta}\|_2 \geq L$. Now set $\mathbf{x}_C = -C \cdot \boldsymbol{\theta}$. Then we have

$$\begin{aligned}
\Psi^*(\boldsymbol{\theta}) &= \inf_{\mathbf{x}} \{\langle \boldsymbol{\theta}, \mathbf{x} \rangle - \Psi(\mathbf{x})\} \\
&\leq \inf_{C>0} \{\langle \boldsymbol{\theta}, \mathbf{x}_C \rangle - \Psi(\mathbf{x}_C)\} \\
&= \inf_{C>0} \{-C \|\boldsymbol{\theta}\|_2^2 - \Psi(\mathbf{x}_C)\} \\
&\leq \inf_{C>0} \{-C \|\boldsymbol{\theta}\|_2^2 - \Psi(\mathbf{0}) + CL \|\boldsymbol{\theta}\|_\infty\} \\
&\leq \inf_{C>0} \{-C \|\boldsymbol{\theta}\|_2^2 - \Psi(\mathbf{0}) + CL \|\boldsymbol{\theta}\|_2\} \\
&= \inf_{C>0} \{-C \|\boldsymbol{\theta}\|_2 (\|\boldsymbol{\theta}\|_2 - L)\} - \Psi(\mathbf{0}) \\
&\leq \inf_{C>0} \{-C \|\boldsymbol{\theta}\|_2 - \Psi(\mathbf{0})\} \\
&= -\infty
\end{aligned}$$

Thus, we can conclude that no dual vector with norm greater than L can be a part of the sufficient dual region. This shows that the sufficient dual region is bounded inside a ball of radius L . For proving the result in the other direction, suppose the dual sufficient region is indeed bounded in a ball of radius R . Consider two points $\mathbf{x}_1, \mathbf{x}_2$ such that

$$\begin{aligned}
\boldsymbol{\theta}_1^* &= \arg \min_{\boldsymbol{\theta} \in \mathcal{A}_\Psi} \{\langle \boldsymbol{\theta}, \mathbf{x}_1 \rangle - \Psi^*(\boldsymbol{\theta})\} \\
\boldsymbol{\theta}_2^* &= \arg \min_{\boldsymbol{\theta} \in \mathcal{A}_\Psi} \{\langle \boldsymbol{\theta}, \mathbf{x}_2 \rangle - \Psi^*(\boldsymbol{\theta})\}
\end{aligned}$$

Now define $f(\boldsymbol{\theta}, \mathbf{x}) := \langle \boldsymbol{\theta}, \mathbf{x} \rangle - \Psi^*(\boldsymbol{\theta})$ so that, by the above definition, $f(\boldsymbol{\theta}_1^*, \mathbf{x}_1) = \Psi(\mathbf{x}_1)$ and $f(\boldsymbol{\theta}_2^*, \mathbf{x}_2) = \Psi(\mathbf{x}_2)$. Now we have

$$\begin{aligned}
\Psi(\mathbf{x}_1) &= f(\boldsymbol{\theta}_1^*, \mathbf{x}_1) \leq f(\boldsymbol{\theta}_2^*, \mathbf{x}_1) \\
&\leq f(\boldsymbol{\theta}_2^*, \mathbf{x}_2) + |\langle \boldsymbol{\theta}_2^*, \mathbf{x}_1 - \mathbf{x}_2 \rangle| \\
&= \Psi(\mathbf{x}_2) + |\langle \boldsymbol{\theta}_2^*, \mathbf{x}_1 - \mathbf{x}_2 \rangle| \\
&\leq \Psi(\mathbf{x}_2) + R \|\mathbf{x}_1 - \mathbf{x}_2\|_2,
\end{aligned}$$

where the fourth step follows from the norm bound on $\boldsymbol{\theta}_2^*$. Similarly we have

$$\Psi(\mathbf{x}_2) \leq \Psi(\mathbf{x}_1) + R \|\mathbf{x}_1 - \mathbf{x}_2\|_2$$

This establishes the result. □

B Proof of Theorem 4

Theorem 4. *Suppose we are given a stream of random samples $(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_T, y_T)$ drawn from a distribution \mathcal{D} over $\mathcal{X} \times \mathcal{Y}$. Let $\Psi(\cdot)$ be a concave, Lipschitz link function. Let Algorithm 1 be executed with a dual feasible set*

$\mathcal{A} \supseteq \mathcal{A}_\Psi$, $\eta_t = 1/\sqrt{t}$ and $\eta'_t = 1/\sqrt{t}$. Then, the average model $\bar{\mathbf{w}} = \frac{1}{T} \sum_{t=1}^T \mathbf{w}_t$ output by the algorithm satisfies, with probability at least $1 - \delta$,

$$\mathcal{P}_\Psi(\bar{\mathbf{w}}) \geq \sup_{\mathbf{w}^* \in \mathcal{W}} \mathcal{P}_\Psi(\mathbf{w}^*) - \delta_\Psi \left(\sqrt{\frac{2B_r^2}{T} \log \frac{1}{\delta}} \right) - (L_\Psi^2 + 4B_r^2) \frac{1}{2\sqrt{T}} - (L_\Psi^2 L_r^2 + R_{\mathcal{W}}^2) \frac{1}{2\sqrt{T}} - \sqrt{\frac{2L_\Psi^2 B_r^2}{T} \log \frac{1}{\delta}}.$$

Proof. For this proof we shall assume that Ψ is L_Ψ -Lipschitz so that its sufficient dual region can be bounded by an application of Lemma 3. Notice that the updates for (α, β) can be written as follows:

$$(\alpha_{t+1}, \beta_{t+1}) \leftarrow \Pi_{\mathcal{A}_\Psi} \left((\alpha_t, \beta_t) - \eta_t \nabla_{(\alpha, \beta)} \ell_t^d(\alpha_t, \beta_t) \right),$$

where

$$\ell_t^d(\alpha, \beta) = \begin{cases} \alpha r^+(\mathbf{w}_t; \mathbf{x}_t, y_t) - \Psi^*(\alpha, \beta) & \text{if } y_t > 0 \\ \beta r^-(\mathbf{w}_t; \mathbf{x}_t, y_t) - \Psi^*(\alpha, \beta) & \text{if } y_t < 0 \end{cases}$$

which can be interpreted as simple gradient descent with ℓ_t . Moreover, since Ψ^* is concave, ℓ_t^d is convex with respect to (α, β) for every t . Note that the terms $r^+(\mathbf{w}_t; \mathbf{x}_t, y_t)$ and $r^-(\mathbf{w}_t; \mathbf{x}_t, y_t)$ do not involve α, β and hence act as arbitrary bounded positive constants for this part of the analysis.

Note that by Lemma 3, we have the radius of \mathcal{A}_Ψ bounded by L_Ψ . Also, since Ψ is a monotone function, by a similar argument, $\Psi^*(\alpha, \beta)$ can be shown to be a $\Psi(B_r, B_r)$ -Lipschitz function. For all the performance measures considered, we have $\Psi(B_r, B_r) \leq B_r$. Thus, $\ell_t^d(\alpha, \beta)$ is a $2B_r$ -Lipschitz function. Hence, using a standard GIGA-style analysis Zinkevich [2003] on the (descent) updates on α_t and β_t in Algorithm 1, we have (for $\eta_t = \frac{1}{\sqrt{t}}$)

$$\begin{aligned} & \frac{1}{T} \sum_{t=1}^T [\alpha_t r^+(\mathbf{w}_t; \mathbf{x}_t, y_t) + \beta_t r^-(\mathbf{w}_t; \mathbf{x}_t, y_t) - \Psi^*(\alpha_t, \beta_t)] \\ & \leq \inf_{(\alpha, \beta) \in \mathcal{A}} \left\{ \frac{1}{T} \sum_{t=1}^T [\alpha r^+(\mathbf{w}_t; \mathbf{x}_t, y_t) + \beta r^-(\mathbf{w}_t; \mathbf{x}_t, y_t) - \Psi^*(\alpha, \beta)] \right\} + (L_\Psi^2 + 4B_r^2) \frac{1}{2\sqrt{T}} \\ & = \inf_{(\alpha, \beta) \in \mathcal{A}} \left\{ \alpha \frac{1}{T} \sum_{t=1}^T r^+(\mathbf{w}_t; \mathbf{x}_t, y_t) + \beta \frac{1}{T} \sum_{t=1}^T r^-(\mathbf{w}_t; \mathbf{x}_t, y_t) - \Psi^*(\alpha, \beta) \right\} + (L_\Psi^2 + 4B_r^2) \frac{1}{2\sqrt{T}} \\ & = \Psi \left(\frac{1}{T} \sum_{t=1}^T r^+(\mathbf{w}_t; \mathbf{x}_t, y_t), \frac{1}{T} \sum_{t=1}^T r^-(\mathbf{w}_t; \mathbf{x}_t, y_t) \right) + (L_\Psi^2 + 4B_r^2) \frac{1}{2\sqrt{T}}, \end{aligned}$$

where the last step follows from Fenchel conjugacy.

Further, noting that $\mathbb{E}_{\mathbf{x}_t, y_t} [r^+(\mathbf{w}_t; \mathbf{x}_t, y_t) \mid \mathbf{x}_{1:t-1}, y_{1:t-1}] = P(\mathbf{w}_t)$, and $\mathbb{E}_{\mathbf{x}_t, y_t} [r^-(\mathbf{w}_t; \mathbf{x}_t, y_t) \mid \mathbf{x}_{1:t-1}, y_{1:t-1}] = N(\mathbf{w}_t)$, we use the standard online-batch conversion bounds Cesa-Bianchi et al. [2001] to the loss functions r^+ and r^- individually to obtain w.h.p.

$$\begin{aligned} \frac{1}{T} \sum_{t=1}^T r^+(\mathbf{w}_t; \mathbf{x}_t, y_t) & \leq \sum_{t=1}^T P(\mathbf{w}_t) + \sqrt{\frac{2B_r^2}{T} \log \frac{1}{\delta}} \\ \frac{1}{T} \sum_{t=1}^T r^-(\mathbf{w}_t; \mathbf{x}_t, y_t) & \leq \sum_{t=1}^T N(\mathbf{w}_t) + \sqrt{\frac{2B_r^2}{T} \log \frac{1}{\delta}} \end{aligned}$$

By monotonicity of Ψ , we get

$$\begin{aligned} & \frac{1}{T} \sum_{t=1}^T [\alpha_t r^+(\mathbf{w}_t; \mathbf{x}_t, y_t) + \beta_t r^-(\mathbf{w}_t; \mathbf{x}_t, y_t) - \Psi^*(\alpha_t, \beta_t)] \\ & \leq \Psi \left(\frac{1}{T} \sum_{t=1}^T P(\mathbf{w}_t) + \sqrt{\frac{2B_r^2}{T} \log \frac{1}{\delta}}, \frac{1}{T} \sum_{t=1}^T N(\mathbf{w}_t) + \sqrt{\frac{2B_r^2}{T} \log \frac{1}{\delta}} \right) + (L_\Psi^2 + 4B_r^2) \frac{1}{2\sqrt{T}} \end{aligned}$$

$$\begin{aligned}
&\leq \Psi\left(\frac{1}{T}\sum_{t=1}^T P(\mathbf{w}_t), \frac{1}{T}\sum_{t=1}^T N(\mathbf{w}_t)\right) + \delta_\Psi\left(\sqrt{\frac{2B_r^2}{T}\log\frac{1}{\delta}}\right) + (L_\Psi^2 + 4B_r^2)\frac{1}{2\sqrt{T}} \\
&\leq \Psi\left(\bar{r}^+\left(\frac{1}{T}\sum_{t=1}^T \mathbf{w}_t\right), \bar{r}^-\left(\frac{1}{T}\sum_{t=1}^T \mathbf{w}_t\right)\right) + \delta_\Psi\left(\sqrt{\frac{2B_r^2}{T}\log\frac{1}{\delta}}\right) + (L_\Psi^2 + 4B_r^2)\frac{1}{2\sqrt{T}} \\
&= \Psi(P(\bar{\mathbf{w}}), N(\bar{\mathbf{w}})) + \delta_\Psi\left(\sqrt{\frac{2B_r^2}{T}\log\frac{1}{\delta}}\right) + (L_\Psi^2 + 4B_r^2)\frac{1}{2\sqrt{T}}, \tag{1}
\end{aligned}$$

where the second inequality follows from stability of Ψ , and the third inequality follows from concavity of \bar{r}^+ and \bar{r}^- , Jensen's inequality, and stability of Ψ .

Similarly, the update to \mathbf{w} can be written as

$$\mathbf{w}_{t+1} \leftarrow \Pi_{\mathcal{W}}(\mathbf{w}_t - \eta'_t \nabla_{\mathbf{w}} \ell_t^p(\mathbf{w}_t)),$$

where $\Pi_{\mathcal{W}}$ is the projection operator for the domain \mathcal{W} and

$$\ell_t^p(\mathbf{w}) = \begin{cases} -\alpha_t r^+(\mathbf{w}; \mathbf{x}_t, y_t) + \Psi^*(\alpha_t, \beta_t) & \text{if } y_t > 0 \\ -\beta_t r^-(\mathbf{w}; \mathbf{x}_t, y_t) + \Psi^*(\alpha_t, \beta_t) & \text{if } y_t < 0 \end{cases}$$

Since r^+, r^- are concave and the term $\Psi^*(\alpha_t, \beta_t)$ does not involve \mathbf{w} , ℓ_t^p is convex in \mathbf{w} for all t . Also, we can show that $\ell_t^p(\mathbf{w})$ is an $(L_\Psi \cdot L_r)$ -Lipschitz function. Hence, applying a standard GIGA analysis Zinkevich [2003] to the (ascent) update on \mathbf{w}_t in Algorithm 1 (with $\eta'_t = \frac{1}{\sqrt{t}}$), we have for any $\mathbf{w}^* \in \mathcal{W}$,

$$\begin{aligned}
&\frac{1}{T}\sum_{t=1}^T [\alpha_t r^+(\mathbf{w}_t; \mathbf{x}_t, y_t) + \beta_t r^-(\mathbf{w}_t; \mathbf{x}_t, y_t) - \Psi^*(\alpha_t, \beta_t)] \\
&\geq \frac{1}{T}\sum_{t=1}^T [\alpha_t r^+(\mathbf{w}^*; \mathbf{x}_t, y_t) + \beta_t r^-(\mathbf{w}^*; \mathbf{x}_t, y_t) - \Psi^*(\alpha_t, \beta_t)] - (L_\Psi^2 L_r^2 + R_{\mathcal{W}}^2)\frac{1}{2\sqrt{T}}.
\end{aligned}$$

Again, observing that by linearity of expectation, we have

$$\mathbb{E}_{\mathbf{x}_t, y_t} [\alpha_t r^+(\mathbf{w}^*; \mathbf{x}_t, y_t) + \beta_t r^-(\mathbf{w}^*; \mathbf{x}_t, y_t) \mid \mathbf{x}_{1:t-1}, y_{1:t-1}] = \alpha_t P(\mathbf{w}^*) + \beta_t N(\mathbf{w}^*),$$

which gives us, through an online-batch conversion argument Cesa-Bianchi et al. [2001] w.h.p,

$$\begin{aligned}
&\frac{1}{T}\sum_{t=1}^T [\alpha_t r^+(\mathbf{w}_t; \mathbf{x}_t, y_t) + \beta_t r^-(\mathbf{w}_t; \mathbf{x}_t, y_t) - \Psi^*(\alpha_t, \beta_t)] \\
&\geq \frac{1}{T}\sum_{t=1}^T [\alpha_t P(\mathbf{w}^*) + \beta_t N(\mathbf{w}^*)] - \frac{1}{T}\sum_{t=1}^T \Psi^*(\alpha_t, \beta_t) - \sqrt{\frac{2L_\Psi^2 B_r^2}{T}\log\frac{1}{\delta}} - (L_\Psi^2 L_r^2 + R_{\mathcal{W}}^2)\frac{1}{2\sqrt{T}} \\
&\geq \frac{1}{T}\sum_{t=1}^T [\alpha_t P(\mathbf{w}^*) + \beta_t N(\mathbf{w}^*)] - \Psi^*\left(\frac{1}{T}\sum_{t=1}^T \alpha_t, \frac{1}{T}\sum_{t=1}^T \beta_t\right) - \sqrt{\frac{2L_\Psi^2 B_r^2}{T}\log\frac{1}{\delta}} - (L_\Psi^2 L_r^2 + R_{\mathcal{W}}^2)\frac{1}{2\sqrt{T}} \\
&= \bar{\alpha}P(\mathbf{w}^*) + \bar{\beta}N(\mathbf{w}^*) - \Psi^*(\bar{\alpha}, \bar{\beta}) - \sqrt{\frac{2L_\Psi^2 B_r^2}{T}\log\frac{1}{\delta}} - (L_\Psi^2 L_r^2 + R_{\mathcal{W}}^2)\frac{1}{2\sqrt{T}} \\
&\geq \inf_{\alpha, \beta} \left\{ \alpha P(\mathbf{w}^*) + \beta N(\mathbf{w}^*) - \Psi^*(\alpha, \beta) \right\} - \sqrt{\frac{2L_\Psi^2 B_r^2}{T}\log\frac{1}{\delta}} - (L_\Psi^2 L_r^2 + R_{\mathcal{W}}^2)\frac{1}{2\sqrt{T}} \\
&= \Psi(P(\mathbf{w}^*), N(\mathbf{w}^*)) - \sqrt{\frac{2L_\Psi^2 B_r^2}{T}\log\frac{1}{\delta}} - (L_\Psi^2 L_r^2 + R_{\mathcal{W}}^2)\frac{1}{2\sqrt{T}}, \tag{2}
\end{aligned}$$

where the second step follows from concavity of Ψ and Jensen's inequality, in the third step $\bar{\alpha} = \frac{1}{T}\sum_{t=1}^T \alpha_t$ and $\bar{\beta} = \frac{1}{T}\sum_{t=1}^T \beta_t$, and the last step follows from Fenchel conjugacy.

Combining Eq. (1) and (2) gives us the desired result. \square

C Proof of Theorem 5

Theorem 5. Suppose we have the problem setting in Theorem 4 with the $\Psi_{G\text{-mean}}$ performance measure being optimized for. Consider a modification to Algorithm 1 wherein the reward functions are changed to $r_t^+(\cdot) = r^+(\cdot) + \epsilon(t)$, and $r_t^-(\cdot) = r^-(\cdot) + \epsilon(t)$ for $\epsilon(t) = \frac{1}{t^{1/4}}$. Then, the average model $\bar{\mathbf{w}} = \frac{1}{T} \sum_{t=1}^T \mathbf{w}_t$ output by the algorithm satisfies, with probability at least $1 - \delta$,

$$\mathcal{P}_{\Psi_{G\text{-mean}}}(\bar{\mathbf{w}}) \geq \sup_{\mathbf{w}^* \in \mathcal{W}} \mathcal{P}_{\Psi_{G\text{-mean}}}(\mathbf{w}^*) - \tilde{\mathcal{O}}\left(\frac{1}{T^{1/4}}\right).$$

Proof. Suppose $\Psi(u + \epsilon, v + \epsilon) \leq \Psi(u, v) + \delta_\Psi(\epsilon)$ as before. Let $r_t^+(\cdot) = r^+(\cdot) + \epsilon(t)$, and $r_t^-(\cdot) = r^-(\cdot) + \epsilon(t)$. Let us make all updates with respect to r_t^+, r_t^- . Let $r(\epsilon)$ be the radius of the sufficient dual domain \mathcal{A} for a given regularization ϵ . Also let $\bar{\epsilon} = \frac{1}{T} \sum_{i=1}^T \epsilon(t)$. We will assume throughout that $\epsilon(t) = O(1)$. Then we have:

$$\begin{aligned} & \frac{1}{T} \sum_{t=1}^T [\alpha_t r_t^+(\mathbf{w}_t; \mathbf{x}_t, y_t) + \beta_t r_t^-(\mathbf{w}_t; \mathbf{x}_t, y_t) - \Psi^*(\alpha_t, \beta_t)] \\ & \leq \inf_{(\alpha, \beta) \in \mathcal{A}} \left\{ \frac{1}{T} \sum_{t=1}^T [\alpha r_t^+(\mathbf{w}_t; \mathbf{x}_t, y_t) + \beta r_t^-(\mathbf{w}_t; \mathbf{x}_t, y_t) - \Psi^*(\alpha, \beta)] \right\} + \mathcal{O}\left(\frac{r(\bar{\epsilon})}{\sqrt{T}}\right) \\ & = \inf_{(\alpha, \beta) \in \mathcal{A}} \left\{ \alpha \frac{1}{T} \sum_{t=1}^T r^+(\mathbf{w}_t; \mathbf{x}_t, y_t) + \bar{\epsilon} + \beta \frac{1}{T} \sum_{t=1}^T r^-(\mathbf{w}_t; \mathbf{x}_t, y_t) + \bar{\epsilon} - \Psi^*(\alpha, \beta) \right\} + \mathcal{O}\left(\frac{r(\bar{\epsilon})}{\sqrt{T}}\right) \\ & = \Psi\left(\frac{1}{T} \sum_{t=1}^T r^+(\mathbf{w}_t; \mathbf{x}_t, y_t) + \bar{\epsilon}, \frac{1}{T} \sum_{t=1}^T r^-(\mathbf{w}_t; \mathbf{x}_t, y_t) + \bar{\epsilon}\right) + \mathcal{O}\left(\frac{r(\bar{\epsilon})}{\sqrt{T}}\right) \quad (3) \\ & = \Psi\left(\frac{1}{T} \sum_{t=1}^T r^+(\mathbf{w}_t; \mathbf{x}_t, y_t), \frac{1}{T} \sum_{t=1}^T r^-(\mathbf{w}_t; \mathbf{x}_t, y_t)\right) + \delta_\Psi(\bar{\epsilon}) + \mathcal{O}\left(\frac{r(\bar{\epsilon})}{\sqrt{T}}\right) \end{aligned}$$

We can now use online to batch conversion bounds Cesa-Bianchi et al. [2001], and monotonicity of Ψ to get

$$\begin{aligned} & \frac{1}{T} \sum_{t=1}^T [\alpha_t r_t^+(\mathbf{w}_t; \mathbf{x}_t, y_t) + \beta_t r_t^-(\mathbf{w}_t; \mathbf{x}_t, y_t) - \Psi^*(\alpha_t, \beta_t)] \\ & \leq \Psi(P(\bar{\mathbf{w}}), N(\bar{\mathbf{w}})) + \delta_\Psi\left(\tilde{\mathcal{O}}\left(\frac{1}{\sqrt{T}}\right)\right) + \delta_\Psi(\bar{\epsilon}) + \mathcal{O}\left(\frac{r(\bar{\epsilon})}{\sqrt{T}}\right), \quad (4) \end{aligned}$$

For the primal updates, we get, for any $\mathbf{w}^* \in \mathcal{W}$,

$$\begin{aligned} & \frac{1}{T} \sum_{t=1}^T [\alpha_t r_t^+(\mathbf{w}_t; \mathbf{x}_t, y_t) + \beta_t r_t^-(\mathbf{w}_t; \mathbf{x}_t, y_t) - \Psi^*(\alpha_t, \beta_t)] \\ & \geq \frac{1}{T} \sum_{t=1}^T [\alpha_t r_t^+(\mathbf{w}^*; \mathbf{x}_t, y_t) + \beta_t r_t^-(\mathbf{w}^*; \mathbf{x}_t, y_t) - \Psi^*(\alpha_t, \beta_t)] - \tilde{\mathcal{O}}\left(\frac{r(\bar{\epsilon})}{\sqrt{T}}\right) \\ & = \frac{1}{T} \sum_{t=1}^T [\alpha_t r^+(\mathbf{w}^*; \mathbf{x}_t, y_t) + \beta_t r^-(\mathbf{w}^*; \mathbf{x}_t, y_t) - \Psi^*(\alpha_t, \beta_t)] + \frac{1}{T} \sum_{t=1}^T \epsilon(t)(\alpha_t + \beta_t) - \tilde{\mathcal{O}}\left(\frac{r(\bar{\epsilon})}{\sqrt{T}}\right) \\ & \geq \frac{1}{T} \sum_{t=1}^T [\alpha_t r^+(\mathbf{w}^*; \mathbf{x}_t, y_t) + \beta_t r^-(\mathbf{w}^*; \mathbf{x}_t, y_t) - \Psi^*(\alpha_t, \beta_t)] - \tilde{\mathcal{O}}\left(\frac{r(\bar{\epsilon})}{\sqrt{T}}\right), \end{aligned}$$

since $\epsilon(t), \alpha_t, \beta_t \geq 0$. Again using an online-batch conversion argument Cesa-Bianchi et al. [2001] we get w.h.p,

$$\frac{1}{T} \sum_{t=1}^T [\alpha_t r_t^+(\mathbf{w}_t; \mathbf{x}_t, y_t) + \beta_t r_t^-(\mathbf{w}_t; \mathbf{x}_t, y_t) - \Psi^*(\alpha_t, \beta_t)] \geq \Psi(P(\mathbf{w}^*), N(\mathbf{w}^*)) - \tilde{\mathcal{O}}\left(\frac{r(\bar{\epsilon})}{\sqrt{T}}\right). \quad (5)$$

Combining Eq. (4) and (5) gives us

$$\Psi(P(\bar{\mathbf{w}}), N(\bar{\mathbf{w}})) \geq \Psi(P(\mathbf{w}^*), N(\mathbf{w}^*)) - \tilde{\mathcal{O}}\left(\frac{r(\bar{\epsilon})}{\sqrt{T}}\right) - \delta_{\Psi}\left(\tilde{\mathcal{O}}\left(\frac{1}{\sqrt{T}}\right)\right) - \delta_{\Psi}(\bar{\epsilon})$$

For G-mean, $\delta_{\Psi}(x) = \sqrt{x}$, and by an application of Lemma 3, we have $r(\epsilon) = O(1/\sqrt{\epsilon})$. Thus we have

$$\Psi(P(\bar{\mathbf{w}}), N(\bar{\mathbf{w}})) \geq \Psi(P(\mathbf{w}^*), N(\mathbf{w}^*)) - \tilde{\mathcal{O}}\left(\frac{1}{\sqrt{T\bar{\epsilon}}}\right) - \tilde{\mathcal{O}}\left(\frac{1}{\sqrt[4]{T}}\right) - \sqrt{\bar{\epsilon}}$$

For $\bar{\epsilon} = \mathcal{O}\left(\frac{1}{\sqrt[4]{T}}\right)$, we get

$$\Psi(P(\bar{\mathbf{w}}), N(\bar{\mathbf{w}})) \geq \Psi(P(\mathbf{w}^*), N(\mathbf{w}^*)) - \tilde{\mathcal{O}}\left(\frac{1}{\sqrt[4]{T}}\right)$$

This can be achieved with $\epsilon(t) = \frac{1}{\sqrt[4]{t}}$. □

D Proof of Theorem 8

Theorem 8. *Let Algorithm 2 be executed with a performance measure $\mathcal{P}_{(\mathbf{a}, \mathbf{b})}$ and reward functions that offer values in the range $[0, m]$. Let $\mathcal{P}^* := \sup_{\mathbf{w} \in \mathcal{W}} \mathcal{P}_{(\mathbf{a}, \mathbf{b})}(\mathbf{w})$. Also let $\Delta_t = \mathcal{P}^* - \mathcal{P}_{(\mathbf{a}, \mathbf{b})}(\mathbf{w}_t)$ be the excess error for the model \mathbf{w}_t generated at time t . Then there exists a value $\eta(m) < 1$ such that for $\Delta_t \leq \Delta_0 \cdot \eta(m)^t$.*

Proof. In order to be generic in its treatment, the proof will require the following regularity conditions on the performance measure

1. $b_0 \neq 0$
2. $\alpha - \mathcal{P}(\mathbf{w}) \cdot \gamma \geq 0$ for all $\mathbf{w} \in \mathcal{W}$
3. $\beta - \mathcal{P}(\mathbf{w}) \cdot \delta \geq 0$ for all $\mathbf{w} \in \mathcal{W}$
4. $-1 < f \leq \gamma \cdot P(\mathbf{w}) + \delta \cdot N(\mathbf{w}) \leq g$ for all $\mathbf{w} \in \mathcal{W}$

Define $e_t := V(\mathbf{w}_{t+1}, v_t) - v_t$. Then we can state the following lemmata which together yield the convergence bound proof.

Lemma 10. $\frac{e_t}{1+f} \geq \mathcal{P}^* - v_t$

Proof. Assume that for some \mathbf{w}^* , $\mathcal{P}(\mathbf{w}^*) = v_t + e_t + e'$ where $e' > 0$. Then we have

$$\begin{aligned} V(\mathbf{w}^*, v_t) &= \left(\frac{e_t}{1+f} + e'\right)(1 + \gamma \cdot P(\mathbf{w}^*) + \delta \cdot N(\mathbf{w}^*)) - e_t \\ &\geq \left(\frac{e_t}{1+f} + e'\right)(1+f) - e_t \\ &= e'(1+f) > 0, \end{aligned}$$

which contradicts the fact that no classifier can achieve a valuation greater than $v_t + e_t$ at level v_t , thus proving the desired result. □

Lemma 11. *For any \mathbf{w} that achieves $V(\mathbf{w}, v) = v + e$ such that $e \geq 0$, we have*

$$\mathcal{P}(\mathbf{w}) \geq v + \frac{e}{g+1}$$

Proof. Let $v' = v + \frac{e}{g+1}$. We will show that $V(\mathbf{w}, v') \geq v'$ which will establish the result by pseudo-linearity. We have

$$\begin{aligned}
V(\mathbf{w}, v') - v' &= c + (\alpha - v'\gamma) \cdot P(\mathbf{w}) + (\beta - v'\delta) \cdot N(\mathbf{w}) - v' \\
&= c + (\alpha - v\gamma) \cdot P(\mathbf{w}) + (\beta - v\delta) \cdot N(\mathbf{w}) - v' - \frac{e}{g+1}(\gamma \cdot P(\mathbf{w}) + \delta \cdot N(\mathbf{w})) \\
&= v + e - v' - \frac{e}{g+1}(\gamma \cdot P(\mathbf{w}) + \delta \cdot N(\mathbf{w})) \\
&\geq v + e - v' - \frac{ge}{g+1} = 0,
\end{aligned}$$

where we have used the bounds on $\gamma \cdot P(\mathbf{w}) + \delta \cdot N(\mathbf{w})$ and the fact that $1 + g > 0$. \square

Given the above results we can establish the convergence bound. More specifically, we can show the following: let $\Delta_t = \mathcal{P}^* - \mathcal{P}(\mathbf{w}_t)$. Then we have

$$\Delta_{t+1} \leq \frac{g-f}{g+1} \cdot \Delta_t$$

To see this, consider the following

$$\begin{aligned}
\Delta_{t+1} &= \mathcal{P}^* - \mathcal{P}(\mathbf{w}_{t+1}) \leq \mathcal{P}^* - \left(v_t + \frac{e_t}{g+1} \right) \leq \mathcal{P}^* - \left(v_t + \frac{(1+f)(\mathcal{P}^* - v_t)}{g+1} \right) \\
&= \mathcal{P}^* - \left(\mathcal{P}(\mathbf{w}_t) + \frac{(1+f)(\mathcal{P}^* - \mathcal{P}(\mathbf{w}_t))}{g+1} \right) = \Delta_t - \frac{1+f}{g+1} \cdot \Delta_t = \frac{g-f}{g+1} \cdot \Delta_t,
\end{aligned}$$

which proves the result. Notice that Table 2 gives the rates of convergence for the different performance measures by calculating bounds on the value of $\frac{g-f}{g+1}$ for those performance measures. \square

E An analysis of the AMP Algorithm under Inexact Maximizations

For this and the next section, we will, for the sake of simplicity, we will focus only on the F-measure for $\beta = 1$ and $p = 1/2$ so that $\theta = 1$. For this setting, the F-measure looks like the following: $F(P, N) = \frac{2P}{2+P-N}$, and the valuation function looks like $V(\mathbf{w}, v) = (1 - v/2) \cdot P(\mathbf{w}) + v/2 \cdot N(\mathbf{w})$. We shall denote the performance measure as $F(\mathbf{w})$, and its optimal value as F^* . We will assume that the reward functions give bounded rewards in the range $[0, m)$.

So far we assumed that Step 4 in the Algorithm **AMP** gave us \mathbf{w}_{t+1} such that

$$V(\mathbf{w}_{t+1}, v_t) = \max_{\mathbf{w} \in \mathcal{W}} V(\mathbf{w}, v_t)$$

Now we will only assume that \mathbf{w}_{t+1} satisfies

$$V(\mathbf{w}_{t+1}, v_t) = \max_{\mathbf{w} \in \mathcal{W}} V(\mathbf{w}, v_t) - \epsilon_t$$

We also assume that the level v_t is only approximated in Step 5 of **AMP**, i.e. using Lemma 7 we have

$$v_t = F(\mathbf{w}_t) + \delta_t$$

where δ_t is a signed real number.

Given these approximations, we can prove the following results

Lemma 12. *The following hold for the setting described above*

1. If $\delta_t \leq 0$ then $e_t \geq 0$
2. If $\delta_t > 0$ then $e_t \geq -\delta_t \left(1 + \frac{m}{2}\right)$

3. If $F^* < v_t$ (which can happen only if $\delta_t > 0$), then $e_t < 0$

4. If $e_t < 0$ then $F^* < v_t$

5. We have

(a) If $e_t \geq 0$, then $e_t \geq \left(\frac{2-m}{2}\right) (F^* - v_t)$.

(b) If $e_t < 0$, then $e_t \geq \left(\frac{2+m}{2}\right) (F^* - v_t)$.

6. If $V(\mathbf{w}, v) = v + e$, then

(a) If $e \geq 0$ then $F(\mathbf{w}) \geq v + \frac{2e}{2+m}$

(b) If $e < 0$ then $F(\mathbf{w}) \geq v + \frac{2e}{2-m}$

Proof. We give the proof in parts

1. If $\delta_t \leq 0$ then this means that there exists a \mathbf{w} such that $F(\mathbf{w}) \geq v_t$. The result then follows from pseudo linearity.

2. $v_t = F(\mathbf{w}_t) + \delta_t$ gives us, by pseudo linearity of F-measure,

$$(1 - v_t/2) \cdot P(\mathbf{w}_t) + v_t/2 \cdot N(\mathbf{w}_t) = v_t - \delta_t \left(1 + \frac{P(\mathbf{w}_t) - N(\mathbf{w}_t)}{2}\right) \geq v_t - \delta_t \left(1 + \frac{m}{2}\right).$$

The bound on e_t now follows from its definition.

3. Suppose $e_t \geq 0$ then by pseudo linearity of F-measure, we have, for some \mathbf{w} , $V(\mathbf{w}, v_t) \geq v_t$ which means $F(\mathbf{w}) \geq v_t$ which contradicts the assumption.

4. Suppose there exists \mathbf{w}^* with $F(\mathbf{w}^*) = v_t + e'$ with $e' \geq 0$ then we have

$$(1 - v_t/2) \cdot P(\mathbf{w}^*) + v_t/2 \cdot N(\mathbf{w}^*) = v_t + e' \left(1 + \frac{P(\mathbf{w}^*) - N(\mathbf{w}^*)}{2}\right) \geq 0,$$

which contradicts the fact that $e_t < 0$.

5. Part (a) is simply Lemma 10. For part (b), we will prove that $F^* \leq v_t + \frac{2e_t}{2+m}$. Since $\frac{2}{2+m} > 0$, the result will follow. Assume the contrapositive that some \mathbf{w}^* achieves $F(\mathbf{w}^*) = v_t + \frac{2e_t}{2+m} + e'$ for some $e' > 0$. Using the pseudo linearity of F-measure (and using the shorthand $v' = v_t + \frac{2e_t}{2+m} + e'$), this can be expressed as

$$(1 - v'/2) \cdot P(\mathbf{w}^*) + v'/2 \cdot N(\mathbf{w}^*) = v'$$

where for some $e' > 0$. Then we have

$$\begin{aligned} (1 - v_t/2) \cdot P(\mathbf{w}^*) + v_t/2 \cdot N(\mathbf{w}^*) - v_t - e_t &= v' - v_t - e_t + \frac{1}{2} \left(\frac{2e_t}{2+m} + e' \right) (P(\mathbf{w}^*) - N(\mathbf{w}^*)) \\ &= \frac{2e_t}{2+m} + e' - e_t + \frac{1}{2} \left(\frac{2e_t}{2+m} + e' \right) (P(\mathbf{w}^*) - N(\mathbf{w}^*)) \\ &\geq \frac{2e_t}{2+m} + e' - e_t + \frac{m}{2} \left(\frac{2e_t}{2+m} + e' \right) \\ &= e' \left(1 + \frac{m}{2}\right) + e_t \left(\frac{2}{2+m} - 1 + \frac{m}{2+m} \right) \\ &= e' \left(1 + \frac{m}{2}\right) > 0, \end{aligned}$$

where we have assumed that e' is chosen small enough so that $\frac{2e_t}{2+m} + e' < 0$ still and used the fact that $P(\mathbf{w}^*) - N(\mathbf{w}^*) \leq m$.

6. Part (a) is simply Lemma 11. To prove part (b), we let $v' = v + \frac{2e}{2-m}$, then we have

$$\begin{aligned}
\left(1 - \frac{v'}{2}\right) \cdot P(\mathbf{w}) + \frac{v'}{2} \cdot N(\mathbf{w}) - v' &= \left(1 - \frac{v}{2}\right) \cdot P(\mathbf{w}) + \frac{v}{2} \cdot N(\mathbf{w}) - v' + \frac{e}{2-m} (N(\mathbf{w}) - P(\mathbf{w})) \\
&\geq \left(1 - \frac{v}{2}\right) \cdot P(\mathbf{w}) + \frac{v}{2} \cdot N(\mathbf{w}) - v' + \frac{me}{2-m} \\
&= v + e - \left(v + \frac{2e}{2-m}\right) + \frac{me}{2-m} \\
&= e \left(1 - \frac{2}{2-m} + \frac{m}{2-m}\right) \\
&= 0,
\end{aligned}$$

where the second inequality follows since $N(\mathbf{w}) - P(\mathbf{w}) \leq m$ and $e < 0$ by using the bounds on the reward functions. This proves the result. □

E.1 Convergence analysis

We have the following cases with us

1. Case 1 ($\delta_t \leq 0$): In this case we are setting v_t to a value less than the F-measure of the current classifier. This should hurt performance - we know that $v_t = F(\mathbf{w}_t) + \delta_t$ which gives us, on applying part (a) of the previous lemma using $F^* - v_t = \Delta_t - \delta_t$, the following

$$e_t \geq \frac{2-m}{2} (\Delta_t - \delta_t).$$

Note that we are guaranteed that $e_t \geq 0$ in this case. Now since the maximization in step 4 is also carried out approximately, we have $V(\mathbf{w}_{t+1}, v_t) = v_t + e_t - \epsilon_t$. Now we have two sub cases

- (a) Case 1.1 ($\epsilon_t \leq e_t$): In this case we can apply part 6(a) of the previous lemma to get the following result

$$\Delta_{t+1} \leq \frac{2m}{2+m} \Delta_t - \frac{2m}{2+m} \delta_t + \frac{2\epsilon_t}{2m}$$

- (b) Case 1.2 ($\epsilon_t > e_t$): In this case we are actually making negative progress in the maximization step (since we have $V(\mathbf{w}_{t+1}, v_t) \leq v_t$) and we can only invoke Lemma 5.6(b) to get

$$\Delta_{t+1} \leq \frac{2\epsilon_t}{2-m}$$

Note that the above result should not be interpreted as a one shot step to a very good classifier. The above result holds along with the condition that $\epsilon_t > e_t$. Thus the performance of the classifier is lower bounded by e_t which depends on how far the current classifier is from the best.

2. Case 2 ($\delta_t > 0$): In this case we are setting v_t to the value higher than the F-measure of the current classifier. This can mislead the classifier and results in the following two sub-cases

- (a) Case 2.1 ($F^* \geq v_t$): In this case we are still setting v_t to a legitimate value, i.e. one that is a valid F-measure for some classifier in the hypothesis class. This can only benefit the next optimization stage (in fact if we set $v_t = F^*$, then we would obtain the best classifier in this very iteration!). In this case $e_t \geq 0$ and we can use the analyses of Cases 1.1 and 1.2.

- (b) Case 2.2 ($F^* < v_t$): In this case we are setting v_t to an illegal value, one that is an unachievable value of F-measure. Consequently, using part 3 of the previous lemma, $e_t < 0$ and using part(b) of the previous lemma we get

$$e_t \geq \frac{2+m}{2}(\Delta_t - \delta_t),$$

which, upon applying part 6(b) of the previous lemma (since $e_t - \epsilon_t \leq e_t < 0$) will give us

$$\begin{aligned} \Delta_{t+1} &\leq \frac{2m}{2-m}(\delta_t - \Delta_t) + \frac{2\epsilon_t}{2-m} \\ &\leq \frac{2m}{2-m}\delta_t + \frac{2\epsilon_t}{2-m} \end{aligned}$$

We can combine the cases together as follows

$$\begin{aligned} \Delta_{t+1} &\leq \max \left\{ \mathbf{1}\{\delta \leq 0\} \cdot \left\{ \frac{2m}{2+m}\Delta_t - \frac{2m}{2+m}\delta_t + \frac{2\epsilon_t}{2+m} \right\}, \mathbf{1}\{\epsilon_t > e_t\} \cdot \frac{2\epsilon_t}{2-m}, \mathbf{1}\{\delta > 0\} \cdot \left\{ \frac{2m}{2-m}\delta_t + \frac{2\epsilon_t}{2-m} \right\} \right\} \\ &\leq \max \left\{ \frac{2m}{2+m}\Delta_t + \frac{2m}{2+m}|\delta_t| + \frac{2\epsilon_t}{2+m}, \mathbf{1}\{\epsilon_t > e_t\} \cdot \frac{2\epsilon_t}{2-m}, \frac{2m}{2-m}|\delta_t| + \frac{2\epsilon_t}{2-m} \right\} \\ &\leq \frac{2m}{2+m}\Delta_t + \frac{2m}{2-m}|\delta_t| + \frac{2\epsilon_t}{2-m} \end{aligned}$$

If we let $\eta = \frac{2m}{2+m}$, $\eta' = \frac{2m}{2-m}$, and $\xi_t = |\delta_t| + \epsilon_t/m$, then this gives us

$$\Delta_{t+1} \leq \eta\Delta_t + \eta'\xi_t,$$

which gives us

$$\Delta_T \leq \eta^T \Delta_0 + \frac{\eta'}{\eta} \cdot \sum_{i=0}^{T-1} \eta^{T-i} \xi_i$$

This concludes our analysis.

F Proof of Theorem 9

Theorem 9. Let Algorithm 3 be executed with a performance measure $\mathcal{P}_{(a,b)}$ and reward functions with range $[0, m)$. Let $\eta = \eta(m)$ be the rate of convergence guaranteed for $\mathcal{P}_{(a,b)}$ by the AMP algorithm. Set the epoch lengths to $s_e, s'_e = \tilde{O}\left(\frac{1}{\eta^{2e}}\right)$. Then after $e = \log_{\frac{1}{\eta}}\left(\frac{1}{\epsilon} \log^2 \frac{1}{\epsilon}\right)$ epochs, we can ensure with probability at least $1 - \delta$ that $\mathcal{P}^* - \mathcal{P}_{(a,b)}(\mathbf{w}_e) \leq \epsilon$. Moreover the number of samples consumed till this point is at most $\tilde{O}\left(\frac{1}{\epsilon^2}\right)$.

Proof. Using Hoeffding's inequality, standard regret and online-to-batch guarantees Cesa-Bianchi et al. [2001], Zinkevich [2003], we can ensure that, if the stream lengths for the Model optimization stage and Challenge level estimation stage procedures are s_e and s'_e respectively, then for some fixed $c > 0$ that is independent of the stream length, we have

$$|\delta_t| \leq c \cdot \sqrt{\frac{\log \frac{1}{\delta}}{s'_e}}, |\epsilon_t| \leq c \sqrt{\frac{\log \frac{1}{\delta}}{s_e}}$$

Let $T = \log_{\frac{1}{\eta}}\left(\frac{1}{\epsilon} \log^2 \frac{1}{\epsilon}\right)$ and $s_e = \left(\frac{2c}{m}\right)^2 \left(\frac{1}{\eta}\right)^{2e} \log \frac{T}{\delta}$ and $s'_e = 4c^2 \left(\frac{1}{\eta}\right)^{2e} \log \frac{T}{\delta}$ - this gives us, for each e , with probability at least $1 - \delta/T$,

$$\xi_e \leq \eta^e$$

Thus, using a union bound, with probability at least $1 - \delta$, we have, by the discussion in the previous section,

$$\Delta_T \leq \eta^T \Delta_0 + \frac{\eta'}{\eta} \sum_{i=0}^{T-1} \eta^{T-i} \xi_i \leq \eta^T \Delta_0 + \frac{\eta'}{\eta} T \eta^T$$

$$\begin{aligned}
&\leq \epsilon \Delta_0 \log^{-2} \frac{1}{\epsilon} + \frac{\eta'}{\eta} \log \frac{1}{\eta} \left(\frac{1}{\epsilon} \log^2 \frac{1}{\epsilon} \right) \epsilon \log^{-2} \frac{1}{\epsilon} \\
&\leq \epsilon \left(\Delta_0 + \frac{\eta'}{\eta \log \frac{1}{\eta}} \right),
\end{aligned}$$

where the last step follows from the fact that for any $\epsilon < 1/e^2$, we have

$$\log \left(\frac{1}{\epsilon} \log^2 \frac{1}{\epsilon} \right) \leq \log^2 \frac{1}{\epsilon}$$

Let $d = \left(\Delta_0 + \frac{\eta'}{\eta \log \frac{1}{\eta}} \right)$ so that we can later set $\epsilon' = \epsilon/d$, and $s = 4c^2 \left(1 + \frac{1}{m^2} \right)$ so that $s_e + s'_e = s \left(\frac{1}{\eta} \right)^{2e} \log \frac{T}{\delta}$. The total number of samples required can then be calculated as

$$\sum_{e=1}^T s_e + s'_e = s \log \frac{T}{\delta} \sum_{e=1}^T \left(\frac{1}{\eta} \right)^{2e} = s \log \frac{T}{\delta} \frac{1}{1 - \eta^2} \left(\frac{1}{\eta^2} \right)^T \leq s \log \frac{T}{\delta} \frac{1}{1 - \eta^2} \frac{1}{\epsilon^2} \log^4 \frac{1}{\epsilon}$$

This gives the number of samples required as

$$\mathcal{O} \left(\frac{1}{\epsilon^2} \log^4 \frac{1}{\epsilon} \left(\log \log \frac{1}{\epsilon} + \log \frac{1}{\delta} \right) \right),$$

to get an ϵ -accurate solution with confidence $1 - \delta$. □