

Online and Stochastic Gradient Methods for Non-decomposable Loss Functions

Purushottam Kar* Harikrishna Narasimhan† Prateek Jain*

*Microsoft Research, INDIA

†Indian Institute of Science, Bangalore, INDIA

{t-purkar,prajain}@microsoft.com, harikrishna@csa.iisc.ernet.in

October 31, 2014

Abstract

Modern applications in sensitive domains such as biometrics and medicine frequently require the use of *non-decomposable* loss functions such as precision@ k , F-measure etc. Compared to point loss functions such as hinge-loss, these offer much more fine grained control over prediction, but at the same time present novel challenges in terms of algorithm design and analysis. In this work we initiate a study of online learning techniques for such non-decomposable loss functions with an aim to enable incremental learning as well as design scalable solvers for batch problems. To this end, we propose an online learning framework for such loss functions. Our model enjoys several nice properties, chief amongst them being the existence of efficient online learning algorithms with sublinear regret and online to batch conversion bounds. Our model is a provable extension of existing online learning models for point loss functions. We instantiate two popular losses, $\text{Prec}_{@k}$ and pAUC, in our model and prove sublinear regret bounds for both of them. Our proofs require a novel structural lemma over ranked lists which may be of independent interest. We then develop scalable stochastic gradient descent solvers for non-decomposable loss functions. We show that for a large family of loss functions satisfying a certain uniform convergence property (that includes $\text{Prec}_{@k}$, pAUC, and F-measure), our methods provably converge to the empirical risk minimizer. Such uniform convergence results were not known for these losses and we establish these using novel proof techniques. We then use extensive experimentation on real life and benchmark datasets to establish that our method can be orders of magnitude faster than a recently proposed cutting plane method.

1 Introduction

Modern learning applications frequently require a level of fine-grained control over prediction performance that is not offered by traditional “per-point” performance measures such as hinge loss. Examples include datasets with mild to severe label imbalance such as spam classification wherein positive instances (spam emails) constitute a tiny fraction of the available data, and learning tasks such as those in medical diagnosis which make it imperative for learning algorithms to be sensitive to class imbalances. Other popular examples include ranking tasks where precision in the top ranked results is valued more than overall precision/recall characteristics. The performance measures of choice in these situations are those that evaluate algorithms over the entire dataset in a holistic manner. Consequently, these measures are frequently *non-decomposable* over data points.

More specifically, for these measures, the loss on a set of points cannot be expressed as the sum of losses on individual data points (unlike hinge loss, for example). Popular examples of such measures include F-measure, Precision@ k , (partial) area under the ROC curve etc.

Despite their success in these domains, non-decomposable loss functions are not nearly as well understood as their decomposable counterparts. The study of point loss functions has led to a deep understanding about their behavior in batch and online settings and tight characterizations of their generalization abilities. The same cannot be said for most non-decomposable losses. For instance, in the popular online learning model, it is difficult to even instantiate a non-decomposable loss function as defining the per-step penalty itself becomes a challenge.

1.1 Our Contributions

Our first main contribution is a framework for online learning with non-decomposable loss functions. The main hurdle in this task is a proper definition of instantaneous penalties for non-decomposable losses. Instead of resorting to canonical definitions, we set up our framework in a principled way that fulfills the objectives of an online model. Our framework has a very desirable characteristic that allows it to recover existing online learning models when instantiated with point loss functions. Our framework also admits online-to-batch conversion bounds.

We then propose an efficient Follow-the-Regularized-Leader [1] algorithm within our framework. We show that for loss functions that satisfy a generic “stability” condition, our algorithm is able to offer vanishing $\mathcal{O}\left(\frac{1}{\sqrt{T}}\right)$ regret. Next, we instantiate within our framework, convex surrogates for two popular performances measures namely, Precision at k (Prec@ k) and partial area under the ROC curve (pAUC) [2] and show, via a stability analysis, that we do indeed achieve sublinear regret bounds for these loss functions. Our stability proofs involve a structural lemma on sorted lists of inner products which proves Lipschitz continuity properties for measures on such lists (see Lemma 2) and might be useful for analyzing non-decomposable loss functions in general.

A key property of online learning methods is their applicability in designing solvers for offline/batch problems. With this goal in mind, we design a stochastic gradient-based solver for non-decomposable loss functions. Our methods apply to a wide family of loss functions (including Prec@ k , pAUC and F-measure) that were introduced in [3] and have been widely adopted [4, 5, 6] in the literature. We design several variants of our method and show that our methods provably converge to the empirical risk minimizer of the learning instance at hand. Our proofs involve uniform convergence-style results which were not known for the loss functions we study and require novel techniques, in particular the structural lemma mentioned above.

Finally, we conduct extensive experiments on real life and benchmark datasets with pAUC and Prec@ k as performance measures. We compare our methods to state-of-the-art methods that are based on cutting plane techniques [7]. The results establish that our methods can be significantly faster, all the while offering comparable or higher accuracy values. For example, on a KDD 2008 challenge dataset, our method was able to achieve a pAUC value of 64.8% within 30ms whereas it took the cutting plane method more than 1.2 seconds to achieve a comparable performance.

1.2 Related Work

Non-decomposable loss functions such as Prec@ k , (partial) AUC, F-measure etc, owing to their demonstrated ability to give better performance in situations with label imbalance etc, have generated significant interest within the learning community. From their role in early works as indicators

of performance on imbalanced datasets [8], their importance has risen to a point where they have become the learning objectives themselves. Due to their complexity, methods that try to indirectly optimize these measures are very common e.g. [9], [10] and [11] who study the F-measure. However, such methods frequently seek to learn a complex probabilistic model, a task arguably harder than the one at hand itself. On the other hand are algorithms that perform optimization directly via structured losses. Starting from the seminal work of [3], this method has received a lot of interest for measures such as the F-measure [3], average precision [4], pAUC [7] and various ranking losses [5, 6]. These formulations typically use cutting plane methods to design dual solvers.

We note that the learning and game theory communities are also interested in non-additive notions of regret and utility. In particular [12] provides a generic framework for online learning with non-additive notions of regret with a focus on showing regret bounds for mixed strategies in a variety of problems. However, even polynomial time implementation of their strategies is difficult in general. Our focus, on the other hand, is on developing efficient online algorithms that can be used to solve large scale batch problems. Moreover, it is not clear how (if at all) can the loss functions considered here (such as $\text{Prec}_{@k}$) be instantiated in their framework.

Recently, online learning for AUC maximization has received some attention [13, 14]. Although AUC is not a point loss function, it still decomposes over pairs of points in a dataset, a fact that [13] and [14] crucially use. The loss functions in this paper do not exhibit any such decomposability.

2 Problem Formulation

Let $\mathbf{x}_{1:t} := \{\mathbf{x}_1, \dots, \mathbf{x}_t\}$, $\mathbf{x}_i \in \mathbb{R}^d$ and $y_{1:t} := \{y_1, \dots, y_t\}$, $y_i \in \{-1, 1\}$ be the observed data points and *true* binary labels. We will use $\hat{y}_{1:t} := \{\hat{y}_1, \dots, \hat{y}_t\}$, $\hat{y}_i \in \mathbb{R}$ to denote the predictions of a learning algorithm. We shall, for sake of simplicity, restrict ourselves to linear predictors $\hat{y}_i = \mathbf{w}^\top \mathbf{x}_i$ for parameter vectors $\mathbf{w} \in \mathbb{R}^d$. A performance measure $\mathcal{P} : \{-1, 1\}^t \times \mathbb{R}^t \rightarrow \mathbb{R}_+$ shall be used to evaluate the the predictions of the learning algorithm against the true labels. Our focus shall be on non-decomposable performance measures such as $\text{Prec}_{@k}$, partial AUC etc.

Since these measures are typically non-convex, convex surrogate *loss functions* are used instead (we will use the terms *loss function* and *performance measure* interchangeably). A popular technique for constructing such loss functions is the *structural SVM* formulation [3] given below. For simplicity, we shall drop mention of the training points and use the notation $\ell_{\mathcal{P}}(\mathbf{w}) := \ell_{\mathcal{P}}(\mathbf{x}_{1:T}, y_{1:T}, \mathbf{w})$.

$$\ell_{\mathcal{P}}(\mathbf{w}) = \max_{\bar{\mathbf{y}} \in \{-1, +1\}^T} \sum_{i=1}^T (\bar{y}_i - y_i) \mathbf{x}_i^\top \mathbf{w} - \mathcal{P}(\bar{\mathbf{y}}, \mathbf{y}). \quad (1)$$

Precision@k. The $\text{Prec}_{@k}$ measure ranks the data points in order of the predicted scores \hat{y}_i and then returns the number of true positives in the top ranked positions. This is valuable in situations where there are very few positives. To formalize this, for any predictor \mathbf{w} and set of points $\mathbf{x}_{1:t}$, define $S(\mathbf{x}, \mathbf{w}) := \{j : \mathbf{w}^\top \mathbf{x}_j > \mathbf{w}^\top \mathbf{x}_i\}$ to be the set of points which \mathbf{w} ranks above \mathbf{x} . Then define

$$\mathbb{T}_{\beta, t}(\mathbf{x}, \mathbf{w}) = \begin{cases} 1, & \text{if } |S(\mathbf{x}, \mathbf{w})| < \lceil \beta t \rceil, \\ 0, & \text{otherwise.} \end{cases} \quad (2)$$

i.e. $\mathbb{T}_{\beta,t}(\mathbf{x}, \mathbf{w})$ is non-zero iff \mathbf{x} is in the top- β fraction of the set. Then we define¹

$$\text{Prec}_{@k}(\mathbf{w}) := \sum_{j:\mathbb{T}_{k,t}(\mathbf{x}_j, \mathbf{w})=1} \mathbb{I}[y_j = 1].$$

The structural surrogate for this measure is then calculated as ²

$$\ell_{\text{Prec}_{@k}}(\mathbf{w}) = \max_{\substack{\bar{y} \in \{-1, +1\}^t \\ \sum_i (\bar{y}_i + 1) = 2kt}} \sum_{i=1}^t (\bar{y}_i - y_i) \mathbf{x}_i^T \mathbf{w} - \sum_{i=1}^t y_i \bar{y}_i. \quad (3)$$

Partial AUC. This measures the area under the ROC curve with the false positive rate restricted to the range $[0, \beta]$. This is in contrast to AUC that allows false positive range in $[0, 1]$. pAUC is useful in medical applications such as cancer detection where a small false positive rate is desirable. Let us extend notation to use $\mathbb{T}_{\beta,t}^-(\mathbf{x}, \mathbf{w})$ to denote the indicator that selects the top β fraction of the *negatively* labeled points i.e. $\mathbb{T}_{\beta,t}^-(\mathbf{x}, \mathbf{w}) = 1$ iff $|\{j : y_j < 0, \mathbf{w}^\top \mathbf{x} > \mathbf{w}^\top \mathbf{x}_j\}| \leq \lceil \beta t_- \rceil$ where t_- is the number of negatives. Then we define

$$\text{pAUC}(\mathbf{w}) = \sum_{i:y_i>0} \sum_{j:y_j<0} \mathbb{T}_{\beta,t}^-(\mathbf{x}_j, \mathbf{w}) \cdot \mathbb{I}[\mathbf{x}_i^\top \mathbf{w} \geq \mathbf{x}_j^\top \mathbf{w}]. \quad (4)$$

The structural surrogate for this performance measure can be equivalently expressed in a simpler form by replacing the indicator functions $\mathbb{I}[\cdot]$ with hinge loss as follows (see [7], Theorem 4)

$$\ell_{\text{pAUC}}(\mathbf{w}) = \sum_{i:y_i>0} \sum_{j:y_j<0} \mathbb{T}_{\beta,t}^-(\mathbf{x}_j, \mathbf{w}) \cdot h(\mathbf{x}_i^\top \mathbf{w} - \mathbf{x}_j^\top \mathbf{w}), \quad (5)$$

where $h(c) = \max(0, 1 - c)$ is the hinge loss function.

In the next section we will develop an online learning framework for non-decomposable performance measures and instantiate our framework with the above mentioned loss functions $\ell_{\text{Prec}_{@k}}$ and ℓ_{pAUC} . Then in Section 4, we will develop stochastic gradient methods for non-decomposable loss functions and prove error bounds for the same. There we will focus on a much larger family of loss functions including $\text{Prec}_{@k}$, pAUC and F-measure.

3 Online Learning with Non-decomposable Loss Functions

We now present our online learning framework for non-decomposable loss functions. Traditional online learning takes place in several rounds, in each of which the player proposes some $\mathbf{w}_t \in \mathcal{W}$ while the adversary responds with a penalty function $\mathcal{L}_t : \mathcal{W} \rightarrow \mathbb{R}$ and a loss $\mathcal{L}_t(\mathbf{w}_t)$ is incurred. The goal is to minimize the *regret* i.e. $\sum_{t=1}^T \mathcal{L}_t(\mathbf{w}_t) - \arg \min_{\mathbf{w} \in \mathcal{W}} \sum_{t=1}^T \mathcal{L}_t(\mathbf{w})$. For point loss functions, the *instantaneous* penalty $\mathcal{L}_t(\cdot)$ is encoded using a data point $(\mathbf{x}_t, y_t) \in \mathbb{R}^d \times \{-1, 1\}$ as $\mathcal{L}_t(\mathbf{w}) = \ell_{\mathcal{P}}(\mathbf{x}_t, y_t, \mathbf{w})$. However, for (surrogates of) non-decomposable loss functions such as ℓ_{pAUC} and $\ell_{\text{Prec}_{@k}}$ the definition of instantaneous penalty itself is not clear and remains a challenge.

To guide us in this process we turn to some properties of standard online learning frameworks. For point losses, we note that the best solution in hindsight is also the batch optimal solution. This

¹An equivalent definition considers k to be the *number* of top ranked points instead.

²[3] uses a slightly modified, but equivalent, definition that considers labels to be Boolean.

is equivalent to the condition $\arg \min_{\mathbf{w} \in \mathcal{W}} \sum_{t=1}^T \mathcal{L}_t(\mathbf{w}) = \arg \min_{\mathbf{w} \in \mathcal{W}} \ell_{\mathcal{P}}(\mathbf{x}_{1:T}, y_{1:T}, \mathbf{w})$ for non-decomposable losses. Also, since the batch optimal solution is agnostic to the ordering of points, we should expect $\sum_{t=1}^T \mathcal{L}_t(\mathbf{w})$ to be invariant to permutations within the stream. By pruning away several naive definitions of \mathcal{L}_t using these requirements, we arrive at the following definition:

$$\mathcal{L}_t(\mathbf{w}) = \ell_{\mathcal{P}}(\mathbf{x}_{1:t}, y_{1:t}, \mathbf{w}) - \ell_{\mathcal{P}}(\mathbf{x}_{1:(t-1)}, y_{1:(t-1)}, \mathbf{w}). \quad (6)$$

It turns out that the above is a very natural penalty function as it measures the amount of “extra” penalty incurred due to the inclusion of \mathbf{x}_t into the set of points. It can be readily verified that $\sum_{t=1}^T \mathcal{L}_t(\mathbf{w}) = \ell_{\mathcal{P}}(\mathbf{x}_{1:T}, y_{1:T}, \mathbf{w})$ as required. Also, this penalty function seamlessly generalizes online learning frameworks since for point losses, we have $\ell_{\mathcal{P}}(\mathbf{x}_{1:t}, y_{1:t}, \mathbf{w}) = \sum_{i=1}^t \ell_{\mathcal{P}}(\mathbf{x}_i, y_i, \mathbf{w})$ and thus $\mathcal{L}_t(\mathbf{w}) = \ell_{\mathcal{P}}(\mathbf{x}_t, y_t, \mathbf{w})$. We note that our framework also recovers the model for online AUC maximization used in [13] and [14]. The notion of regret corresponding to this penalty is

$$R(T) = \frac{1}{T} \sum_{t=1}^T \mathcal{L}_t(\mathbf{w}_t) - \arg \min_{\mathbf{w} \in \mathcal{W}} \frac{1}{T} \ell_{\mathcal{P}}(\mathbf{x}_{1:T}, y_{1:T}, \mathbf{w}).$$

We note that \mathcal{L}_t , being the difference of two loss functions, is non-convex in general and thus, standard online convex programming regret bounds cannot be applied in our framework. Interestingly, as we show below, by exploiting structural properties of our penalty function, we can still get efficient low-regret learning algorithms, as well as online-to-batch conversion bounds in our framework.

3.1 Low Regret Online Learning

We propose an efficient Follow-the-Regularized-Leader (FTRL) style algorithm in our framework. Let $\mathbf{w}_1 = \arg \min_{\mathbf{w} \in \mathcal{W}} \|\mathbf{w}\|_2^2$ and consider the following update:

$$\mathbf{w}_{t+1} = \arg \min_{\mathbf{w} \in \mathcal{W}} \sum_{t=1}^t \mathcal{L}_t(\mathbf{w}) + \frac{\eta}{2} \|\mathbf{w}\|_2^2 = \arg \min_{\mathbf{w} \in \mathcal{W}} \ell_{\mathcal{P}}(\mathbf{x}_{1:t}, y_{1:t}, \mathbf{w}) + \frac{\eta}{2} \|\mathbf{w}\|_2^2 \quad (\text{FTRL})$$

We would like to stress that despite the non-convexity of \mathcal{L}_t , the FTRL objective is strongly convex if $\ell_{\mathcal{P}}$ is convex and thus the update can be implemented efficiently by solving a regularized batch problem on $\mathbf{x}_{1:t}$. We now present our regret bound analysis for the FTRL update given above.

Theorem 1. *Let $\ell_{\mathcal{P}}(\cdot, \mathbf{w})$ be a convex loss function and $\mathcal{W} \subseteq \mathbb{R}^d$ be a convex set. Assume w.l.o.g. $\|\mathbf{x}_t\|_2 \leq 1, \forall t$. Also, for the penalty function \mathcal{L}_t in (6), let $|\mathcal{L}_t(\mathbf{w}) - \mathcal{L}_t(\mathbf{w}')| \leq G_t \cdot \|\mathbf{w} - \mathbf{w}'\|_2$, for all t and all $\mathbf{w}, \mathbf{w}' \in \mathcal{W}$ for some $G_t > 0$. Suppose we use the update step given in ((FTRL)) to obtain $\mathbf{w}_{t+1}, 0 \leq t \leq T - 1$. Then for all \mathbf{w}^* , we have*

$$\frac{1}{T} \sum_{t=1}^T \mathcal{L}_t(\mathbf{w}_t) \leq \frac{1}{T} \ell_{\mathcal{P}}(\mathbf{x}_{1:T}, y_{1:T}, \mathbf{w}^*) + \|\mathbf{w}^*\|_2 \frac{\sqrt{2 \sum_{t=1}^T G_t^2}}{T}.$$

See Appendix A for a proof. The above result requires the penalty function \mathcal{L}_t to be Lipschitz continuous i.e. be “stable” w.r.t. \mathbf{w} . Establishing this for point losses such as hinge loss is relatively straightforward. However, the same becomes non-trivial for non-decomposable loss functions as \mathcal{L}_t

is now the difference of two loss functions, both of which involve $\Omega(t)$ data points. A naive argument would thus, only be able to show $G_t \leq O(t)$ which would yield vacuous regret bounds.

Instead, we now show that for the surrogate loss functions for $\text{Prec}_{@k}$ and pAUC, this Lipschitz continuity property does indeed hold. Our proofs crucially use a structural lemma given below that shows that sorted lists of inner products are Lipschitz at each fixed position.

Lemma 2 (Structural Lemma). *Let $\mathbf{x}_1, \dots, \mathbf{x}_t$ be t points with $\|\mathbf{x}_i\|_2 \leq 1 \forall t$. Let $\mathbf{w}, \mathbf{w}' \in \mathcal{W}$ be any two vectors. Let $z_i = \langle \mathbf{w}, \mathbf{x}_i \rangle - c_i$ and $z'_i = \langle \mathbf{w}', \mathbf{x}_i \rangle - c_i$, where $c_i \in \mathbb{R}$ are constants independent of \mathbf{w}, \mathbf{w}' . Also, let $\{i_1, \dots, i_t\}$ and $\{j_1, \dots, j_t\}$ be ordering of indices such that $z_{i_1} \geq z_{i_2} \geq \dots \geq z_{i_t}$ and $z'_{j_1} \geq z'_{j_2} \geq \dots \geq z'_{j_t}$. Then for any 1-Lipschitz increasing function $g : \mathbb{R} \rightarrow \mathbb{R}$ (i.e. $|g(u) - g(v)| \leq |u - v|$ and $u \leq v \Leftrightarrow g(u) \leq g(v)$), we have, $\forall k \ |g(z_{i_k}) - g(z'_{j_k})| \leq 3\|\mathbf{w} - \mathbf{w}'\|_2$.*

See Appendix B for a proof. Using this lemma we can show that the Lipschitz constant for $\ell_{\text{Prec}_{@k}}$ is bounded by $G_t \leq 8$ which gives us a $\mathcal{O}\left(\sqrt{\frac{1}{T}}\right)$ regret bound for $\text{Prec}_{@k}$ (see Appendix C for the proof). In Appendix D, we show that the same technique can be used to prove a stability result for the structural SVM surrogate of the Precision-Recall Break Even Point (PRBEP) performance measure [3] as well. The case of pAUC is handled similarly. However, since pAUC discriminates between positives and negatives, our previous analysis cannot be applied directly. Nevertheless, we can obtain the following regret bound for pAUC (a proof will appear in the full version of the paper).

Theorem 3. *Let T_+ and T_- resp. be the number of positive and negative points in the stream and let $\mathbf{w}_{t+1}, 0 \leq t \leq T - 1$ be obtained using the FTRL algorithm ((FTRL)). Then we have*

$$\frac{1}{\beta T_+ T_-} \sum_{t=1}^T \mathcal{L}_t(\mathbf{w}_t) \leq \min_{\mathbf{w} \in \mathcal{W}} \frac{1}{\beta T_+ T_-} \ell_{\text{pAUC}}(\mathbf{x}_{1:T}, y_{1:T}, \mathbf{w}) + \mathcal{O}\left(\sqrt{\frac{1}{T_+} + \frac{1}{T_-}}\right).$$

Notice that the above regret bound depends on both T_+ and T_- and the regret becomes large even if one of them is small. This is actually quite intuitive because if, say $T_+ = 1$ and $T_- = T - 1$, an adversary may wish to provide the lone positive point in the last round. Naturally the algorithm, having only seen negatives till now, would not be able to perform well and would incur a large error.

3.2 Online-to-batch Conversion

To present our bounds we generalize our framework slightly: we now consider the stream of T points to be composed of T/s batches $\mathbf{Z}_1, \dots, \mathbf{Z}_{T/s}$ of size s each. Thus, the instantaneous penalty is now defined as $\mathcal{L}_t(\mathbf{w}) = \ell_{\mathcal{P}}(\mathbf{Z}_1, \dots, \mathbf{Z}_t, \mathbf{w}) - \ell_{\mathcal{P}}(\mathbf{Z}_1, \dots, \mathbf{Z}_{t-1}, \mathbf{w})$ for $t = 1 \dots T/s$ and the regret becomes $R(T, s) = \frac{1}{T} \sum_{t=1}^{T/s} \mathcal{L}_t(\mathbf{w}_t) - \arg \min_{\mathbf{w} \in \mathcal{W}} \frac{1}{T} \ell_{\mathcal{P}}(\mathbf{x}_{1:T}, y_{1:T}, \mathbf{w})$. Let $\mathcal{R}_{\mathcal{P}}$ denote the population risk for the (normalized) performance measure \mathcal{P} . Then we have:

Theorem 4. *Suppose the sequence of points (\mathbf{x}_t, y_t) is generated i.i.d. and let $\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_{T/s}$ be an ensemble of models generated by an online learning algorithm upon receiving these T/s batches. Suppose the online learning algorithm has a guaranteed regret bound $R(T, s)$. Then for $\bar{\mathbf{w}} = \frac{1}{T/s} \sum_{t=1}^{T/s} \mathbf{w}_t$, any $\mathbf{w}^* \in \mathcal{W}$, $\epsilon \in (0, 0.5]$ and $\delta > 0$, with probability at least $1 - \delta$,*

$$\mathcal{R}_{\mathcal{P}}(\bar{\mathbf{w}}) \leq (1 + \epsilon)\mathcal{R}_{\mathcal{P}}(\mathbf{w}^*) + R(T, s) + e^{-\Omega(s\epsilon^2)} + \tilde{\mathcal{O}}\left(\sqrt{\frac{s \ln(1/\delta)}{T}}\right).$$

Algorithm 1 1PMB: Single-Pass with Mini-batches

Input: Step length scale η , Buffer B of size s **Output:** A good predictor $\mathbf{w} \in \mathcal{W}$

- 1: $\mathbf{w}_0 \leftarrow \mathbf{0}$, $B \leftarrow \phi$, $e \leftarrow 0$
 - 2: **while** stream not exhausted **do**
 - 3: Collect s data points $(\mathbf{x}_1^e, y_1^e), \dots, (\mathbf{x}_s^e, y_s^e)$ in buffer B
 - 4: Set step length $\eta_e \leftarrow \frac{\eta}{\sqrt{e}}$
 - 5: $\mathbf{w}_{e+1} \leftarrow \Pi_{\mathcal{W}} [\mathbf{w}_e + \eta_e \nabla_{\mathbf{w}} \ell_{\mathcal{P}}(\mathbf{x}_{1:s}^e, y_{1:s}^e, \mathbf{w}_e)]$
 // $\Pi_{\mathcal{W}}$ projects onto the set \mathcal{W}
 - 6: Flush buffer B
 - 7: $e \leftarrow e + 1$ //start a new epoch
 - 8: **end while**
 - 9: **return** $\bar{\mathbf{w}} = \frac{1}{e} \sum_{i=1}^e \mathbf{w}_i$
-

Algorithm 2 2PMB: Two-Passes with Mini-batches

Input: Step length scale η , Buffers B_+ , B_- of size s **Output:** A good predictor $\mathbf{w} \in \mathcal{W}$

- Pass 1:** $B_+ \leftarrow \phi$
 - 1: Collect random sample of pos. $\mathbf{x}_1^+, \dots, \mathbf{x}_s^+$ in B_+
 - Pass 2:** $\mathbf{w}_0 \leftarrow \mathbf{0}$, $B_- \leftarrow \phi$, $e \leftarrow 0$
 - 2: **while** stream of negative points not exhausted **do**
 - 3: Collect s negative points $\mathbf{x}_1^{e-}, \dots, \mathbf{x}_s^{e-}$ in B_-
 - 4: Set step length $\eta_e \leftarrow \frac{\eta}{\sqrt{e}}$
 - 5: $\mathbf{w}_{e+1} \leftarrow \Pi_{\mathcal{W}} [\mathbf{w}_e + \eta_e \nabla_{\mathbf{w}} \ell_{\mathcal{P}}(\mathbf{x}_{1:s}^{e-}, \mathbf{x}_{1:s}^+, \mathbf{w}_e)]$
 - 6: Flush buffer B_-
 - 7: $e \leftarrow e + 1$ //start a new epoch
 - 8: **end while**
 - 9: **return** $\bar{\mathbf{w}} = \frac{1}{e} \sum_{i=1}^e \mathbf{w}_i$
-

In particular, setting $s = \tilde{\mathcal{O}}(\sqrt{T})$ and $\epsilon = \sqrt[4]{1/T}$ gives us, with probability at least $1 - \delta$,

$$\mathcal{R}_{\mathcal{P}}(\bar{\mathbf{w}}) \leq \mathcal{R}_{\mathcal{P}}(\mathbf{w}^*) + R(T, \sqrt{T}) + \tilde{\mathcal{O}}\left(\sqrt[4]{\frac{\ln(1/\delta)}{T}}\right).$$

We conclude by noting that for $\text{Prec}_{@k}$ and pAUC, $R(T, \sqrt{T}) \leq \mathcal{O}\left(\sqrt[4]{1/T}\right)$ (see Appendix E).

4 Stochastic Gradient Methods for Non-decomposable Losses

The online learning algorithms discussed in the previous section present attractive guarantees in the sequential prediction model but are required to solve batch problems at each stage. This rapidly becomes infeasible for large scale data. To remedy this, we now present memory efficient stochastic gradient descent methods for batch learning with non-decomposable loss functions. The motivation for our approach comes from mini-batch methods used to make learning methods for *point* loss functions amenable to distributed computing environments [15, 16], we exploit these techniques to offer scalable algorithms for *non-decomposable* loss functions.

4.1 Single-pass Method with Mini-batches

The method assumes access to a limited memory buffer and takes a pass over the data stream. The stream is partitioned into *epochs*. In each epoch, the method accumulates points in the stream, uses them to form gradient estimates and takes descent steps. The buffer is flushed after each epoch. Algorithm 1 describes the **1PMB** method. Gradient computations can be done using Danskin’s theorem (see Appendix H).

4.2 Two-pass Method with Mini-batches

The previous algorithm is unable to exploit relationships between data points across epochs which may help improve performance for loss functions such as pAUC. To remedy this, we observe that several real life learning scenarios exhibit mild to severe label imbalance (see Table 1 in Appendix H) which makes it possible to store all or a large fraction of points of the rare label. Our two pass

method exploits this by utilizing two passes over the data: the first pass collects all (or a random subset of) points of the rare label using some stream sampling technique [13]. The second pass then goes over the stream, restricted to the non-rare label points, and performs gradient updates. See Algorithm 2 for details of the **2PMB** method.

4.3 Error Bounds

Given a set of n labeled data points $(\mathbf{x}_i, y_i), i = 1 \dots n$ and a performance measure \mathcal{P} , our goal is to approximate the empirical risk minimizer $\mathbf{w}^* = \arg \min_{\mathbf{w} \in \mathcal{W}} \ell_{\mathcal{P}}(\mathbf{x}_{1:n}, y_{1:n}, \mathbf{w})$ as closely as possible. In this section we shall show that our methods **1PMB** and **2PMB** provably converge to the empirical risk minimizer. We first introduce the notion of uniform convergence for a performance measure.

Definition 5. *We say that a loss function ℓ demonstrates uniform convergence with respect to a set of predictors \mathcal{W} if for some $\alpha(s, \delta) = \text{poly}(\frac{1}{s}, \log \frac{1}{\delta})$, when given a set of s points $\bar{\mathbf{x}}_1, \dots, \bar{\mathbf{x}}_s$ chosen randomly from an arbitrary set of n points $\{(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_n, y_n)\}$ then w.p. at least $1 - \delta$, we have*

$$\sup_{\mathbf{w} \in \mathcal{W}} |\ell_{\mathcal{P}}(\mathbf{x}_{1:n}, y_{1:n}, \mathbf{w}) - \ell_{\mathcal{P}}(\bar{\mathbf{x}}_{1:s}, \bar{y}_{1:s}, \mathbf{w})| \leq \alpha(s, \delta).$$

Such uniform convergence results are fairly common for decomposable loss functions such as the squared loss, logistic loss etc. However, the same is not true for non-decomposable loss functions barring a few exceptions [17, 10]. To bridge this gap, below we show that a large family of surrogate loss functions for popular non-decomposable performance measures does indeed exhibit uniform convergence. Our proofs require novel techniques and do not follow from traditional proof progressions. However, we first show how we can use these results to arrive at an error bound.

Theorem 6. *Suppose the loss function ℓ is convex and demonstrates $\alpha(s, \delta)$ -uniform convergence. Also suppose we have an arbitrary set of n points which are randomly ordered, then the predictor $\bar{\mathbf{w}}$ returned by **1PMB** with buffer size s satisfies w.p. $1 - \delta$,*

$$\ell_{\mathcal{P}}(\mathbf{x}_{1:n}, y_{1:n}, \bar{\mathbf{w}}) \leq \ell_{\mathcal{P}}(\mathbf{x}_{1:n}, y_{1:n}, \mathbf{w}_*) + 2\alpha\left(s, \frac{s\delta}{n}\right) + \mathcal{O}\left(\sqrt{\frac{s}{n}}\right)$$

We would like to stress that the above result does not assume i.i.d. data and works for arbitrary datasets so long as they are randomly ordered. We can show similar guarantees for the two pass method as well (see Appendix F). Using regularized formulations, we can also exploit logarithmic regret guarantees [18], offered by online gradient descent, to improve this result - however we do not explore those considerations here. Instead, we now look at specific instances of loss functions that possess the desired uniform convergence properties. As mentioned before, due to the combinatorial nature of these performance measures, our proofs do not follow from traditional methods.

Theorem 7 (Partial Area under the ROC Curve). *For any convex, monotone, Lipschitz, classification surrogate $\phi : \mathbb{R} \rightarrow \mathbb{R}_+$, the surrogate loss function for the $(0, \beta)$ -partial AUC performance measure defined as follows exhibits uniform convergence at the rate $\alpha(s, \delta) = \mathcal{O}\left(\sqrt{\log(1/\delta)/s}\right)$:*

$$\frac{1}{\lceil \beta n_- \rceil n_+} \sum_{i: y_i > 0} \sum_{j: y_j < 0} \mathbb{T}_{\beta, t}^-(\mathbf{x}_j, \mathbf{w}) \cdot \phi(\mathbf{x}_i^\top \mathbf{w} - \mathbf{x}_j^\top \mathbf{w})$$

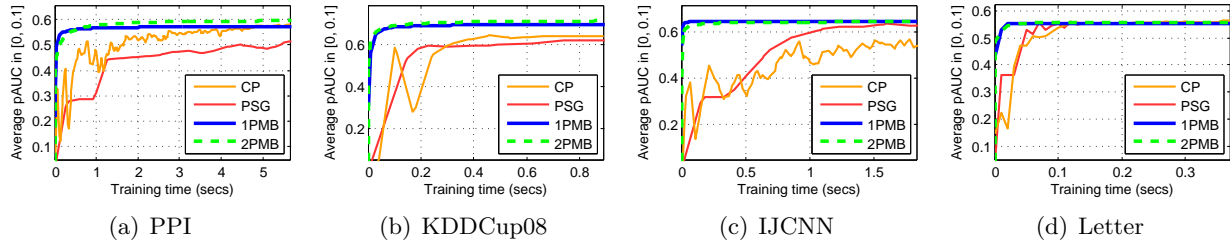


Figure 1: Comparison of stochastic gradient methods with the cutting plane (CP) and projected subgradient (PSG) methods on partial AUC maximization tasks. The epoch lengths/buffer sizes for **1PMB** and **2PMB** were set to 500.

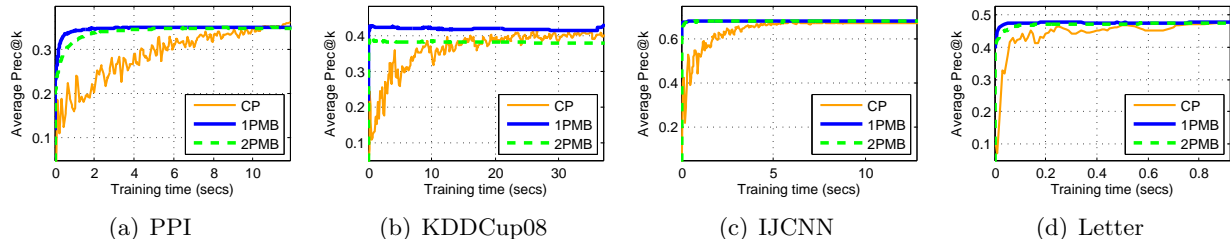


Figure 2: Comparison of stochastic gradient methods with cutting plane (CP) methods on $\text{Prec}@_k$ maximization tasks. The epoch lengths/buffer sizes for **1PMB** and **2PMB** were set to 500.

See Appendix G for a proof sketch. This result covers a large family of surrogate loss functions such as hinge loss (5), logistic loss etc. Note that the insistence on including only top ranked negative points introduces a high degree of non-decomposability into the loss function. A similar result for the special case $\beta = 1$ is due to [17]. We extend the same to the more challenging case of $\beta < 1$.

Theorem 8 (Structural SVM loss for $\text{Prec}@_k$). *The structural SVM surrogate for the $\text{Prec}@_k$ performance measure (see (3)) exhibits uniform convergence at the rate $\alpha(s, \delta) = \mathcal{O}\left(\sqrt{\log(1/\delta)/s}\right)$.*

We defer the proof to the full version of the paper. The above result can be extended to a large family of performance measures introduced in [3] that have been widely adopted [10, 19, 8] such as F-measure, G-mean, and PRBEP. The above indicates that our methods are expected to output models that closely approach the empirical risk minimizer for a wide variety of performance measures. In the next section we verify that this is indeed the case for several real life and benchmark datasets.

5 Experimental Results

We evaluate the proposed stochastic gradient methods on several real-world and benchmark datasets.

Performance measures: We consider three measures, 1) partial AUC in the false positive range $[0, 0.1]$, 2) $\text{Prec}@_k$ with k set to the proportion of positives (PRBEP), and 3) F-measure.

Algorithms: For partial AUC, we compare against the state-of-the-art cutting plane (CP) and projected subgradient methods (PSG) proposed in [7]; unlike the (online) stochastic methods

Dataset	Data Points	Features	Positives
KDDCup08	102,294	117	0.61%
PPI	240,249	85	1.19%
Letter	20,000	16	3.92%
IJCNN	141,691	22	9.57%

Table 1: Statistics of datasets used.

Measure	1PMB	2PMB	CP
pAUC	0.10 (68.2)	0.15 (69.6)	0.39 (62.5)
Prec@ k	0.49 (42.7)	0.55 (38.7)	23.25 (40.8)

Table 2: Comparison of training time (secs) and accuracies (in brackets) of **1PMB**, **2PMB** and cutting plane methods for pAUC (in $[0, 0.1]$) and Prec@ k maximization tasks on the KDD Cup 2008 dataset.

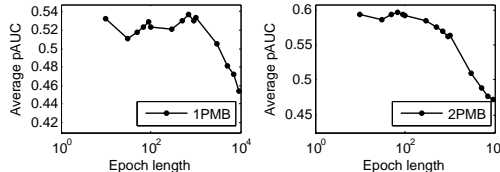


Figure 3: Performance of **1PMB** and **2PMB** on the PPI dataset with varying epoch/buffer sizes for pAUC tasks.

considered in this work, the PSG method is a ‘batch’ algorithm which, at each iteration, computes a subgradient-based update over the entire training set. For Prec@ k and F-measure, we compare our methods against cutting plane methods from [3]. We used structural SVM surrogates for all the measures.

Datasets: We used several data sets for our experiments (see Table 1); of these, KDDCup08 is from the KDD Cup 2008 challenge and involves a breast cancer detection task [20], PPI contains data for a protein-protein interaction prediction task [21], and the remaining datasets are taken from the UCI repository [22].

Parameters: We used 70% of the data set for training and the remaining for testing, with the results averaged over 5 random train-test splits. Tunable parameters such as step length scale were chosen using a small validation set. All experiments used a buffer of size 500. Epoch lengths were set equal to the buffer size. Since a single iteration of the proposed stochastic methods is very fast in practice, we performed multiple passes over the training data (see Appendix H for details).

Results: The results for pAUC and Prec@ k maximization tasks are shown in the Figures 1 and 2. We found the proposed stochastic gradient methods to be several orders of magnitude faster than the baseline methods, all the while achieving comparable or better accuracies. For example, for the pAUC task on the KDD Cup 2008 dataset, the **1PMB** method achieved an accuracy of 64.81% within 0.03 seconds, while even after 0.39 seconds, the cutting plane method could only achieve an accuracy of 62.52% (see Table 2). As expected, the (online) stochastic gradient methods were faster than the ‘batch’ projected subgradient descent method for pAUC as well. We found similar trends on Prec@ k (see Figure 2) and F-measure maximization tasks as well. For F-measure tasks, on the KDD Cup 2008 dataset, for example, the **1PMB** method achieved an accuracy of 35.92 within 12 seconds whereas, even after 150 seconds, the cutting plane method could only achieve an accuracy of 35.25.

The proposed stochastic methods were also found to be robust to changes in epoch lengths

(buffer sizes) till such a point where excessively long epochs would cause the number of updates as well as accuracy to dip (see Figure 3). The **2PMB** method was found to offer higher accuracies for pAUC maximization on several datasets (see Table 2 and Figure 1), as well as be more robust to changes in buffer size (Figure 3). We defer results on more datasets and performance measures to the full version of the paper.

The cutting plane methods were generally found to exhibit a zig-zag behaviour in performance across iterates. This is because these methods solve the dual optimization problem for a given performance measure; hence the intermediate models do not necessarily yield good accuracies. On the other hand, (stochastic) gradient based methods directly offer progress in terms of the primal optimization problem, and hence provide good intermediate solutions as well. This can be advantageous in scenarios with a time budget in the training phase.

Acknowledgements

The authors thank Shivani Agarwal for helpful comments. They also thank the anonymous reviewers for their suggestions. HN thanks support from a Google India PhD Fellowship.

References

- [1] Alexander Rakhlin. Lecture Notes on Online Learning. http://www-stat.wharton.upenn.edu/~rakhlin/papers/online_learning.pdf, 2009.
- [2] Harikrishna Narasimhan and Shivani Agarwal. A Structural SVM Based Approach for Optimizing Partial AUC. In *30th International Conference on Machine Learning (ICML)*, 2013.
- [3] Thorsten Joachims. A Support Vector Method for Multivariate Performance Measures. In *ICML*, 2005.
- [4] Yisong Yue, Thomas Finley, Filip Radlinski, and Thorsten Joachims. A Support Vector Method for Optimizing Average Precision. In *SIGIR*, 2007.
- [5] Soumen Chakrabarti, Rajiv Khanna, Uma Sawant, and Chiru Bhattacharyya. Structured Learning for Non-Smooth Ranking Losses. In *KDD*, 2008.
- [6] Brian McFee and Gert Lanckriet. Metric Learning to Rank. In *ICML*, 2010.
- [7] Harikrishna Narasimhan and Shivani Agarwal. SVM_{pAUC}^{tight}: A New Support Vector Method for Optimizing Partial AUC Based on a Tight Convex Upper Bound. In *KDD*, 2013.
- [8] Miroslav Kubat and Stan Matwin. Addressing the Curse of Imbalanced. Training Sets: One-Sided Selection. In *24th International Conference on Machine Learning (ICML)*, 1997.
- [9] Krzysztof Dembczyński, Willem Waegeman, Weiwei Cheng, and Eyke Hüllermeier. An Exact Algorithm for F-Measure Maximization. In *NIPS*, 2011.
- [10] 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.

- [11] 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.
- [12] Alexander Rakhlin, Karthik Sridharan, and Ambuj Tewari. Online Learning: Beyond Regret. In *24th Annual Conference on Learning Theory (COLT)*, 2011.
- [13] Purushottam Kar, Bharath K Sriperumbudur, Prateek Jain, and Harish Karnick. On the Generalization Ability of Online Learning Algorithms for Pairwise Loss Functions. In *ICML*, 2013.
- [14] Peilin Zhao, Steven C. H. Hoi, Rong Jin, and Tianbao Yang. Online AUC Maximization. In *ICML*, 2011.
- [15] Ofer Dekel, Ran Gilad-Bachrach, Ohad Shamir, and Lin Xiao. Optimal Distributed Online Prediction Using Mini-Batches. *Journal of Machine Learning Research*, 13:165–202, 2012.
- [16] Yuchen Zhang, John C. Duchi, and Martin J. Wainwright. Communication-Efficient Algorithms for Statistical Optimization. *Journal of Machine Learning Research*, 14:3321–3363, 2013.
- [17] Stéphane Cléménçon, Gábor Lugosi, and Nicolas Vayatis. Ranking and empirical minimization of U-statistics. *Annals of Statistics*, 36:844–874, 2008.
- [18] Elad Hazan, Adam Kalai, Satyen Kale, and Amit Agarwal. Logarithmic Regret Algorithms for Online Convex Optimization. In *COLT*, pages 499–513, 2006.
- [19] Sophia Daskalaki, Ioannis Kostas, and Nikolaos Avouris. Evaluation of Classifiers for an Uneven Class Distribution Problem. *Applied Artificial Intelligence*, 20:381–417, 2006.
- [20] R. Bharath Rao, Oksana Yakhnenko, and Balaji Krishnapuram. KDD Cup 2008 and the Workshop on Mining Medical Data. *SIGKDD Explorations Newsletter*, 10(2):34–38, 2008.
- [21] Yanjun Qi, Ziv Bar-Joseph, and Judith Klein-Seetharaman. Evaluation of Different Biological Data and Computational Classification Methods for Use in Protein Interaction Prediction. *Proteins*, 63:490–500, 2006.
- [22] A. Frank and Arthur Asuncion. The UCI Machine Learning Repository. <http://archive.ics.uci.edu/ml>, 2010. University of California, Irvine, School of Information and Computer Sciences.
- [23] Ankan Saha, Prateek Jain, and Ambuj Tewari. The interplay between stability and regret in online learning. *CoRR*, abs/1211.6158, 2012.
- [24] Martin Zinkevich. Online Convex Programming and Generalized Infinitesimal Gradient Ascent. In *ICML*, pages 928–936, 2003.
- [25] Robert J. Serfling. Probability Inequalities for the Sum in Sampling without Replacement. *Annals of Statistics*, 2(1):39–48, 1974.
- [26] Dimitri P. Bertsekas. *Nonlinear Programming: 2nd Edition*. Belmont, MA: Athena Scientific, 2004.

A Proof of Theorem 1

Broadly, we follow the proof structure of FTRL given in [1, 23]. We first observe that the “forward regret” analysis follows easily despite the non-convexity of \mathcal{L}_t . That is,

$$\sum_{t=1}^T \mathcal{L}_t(\mathbf{w}_{t+1}) \leq \mathbf{x}_{1:T}, y_{1:T}, \mathbf{w}_* + \frac{\eta}{2} \|\mathbf{w}_*\|_2^2, \quad (7)$$

where $\mathbf{w}_* = \arg \min_{\mathbf{w} \in \mathcal{W}} \mathbf{x}_{1:T}, y_{1:T}, \mathbf{w}$. The proof of this statement can be found in [23, Theorem 7] and is reproduced below as Lemma 9 for completeness. Next, using strong convexity of the regularizer $\|\mathbf{w}\|_2^2$ and optimality of \mathbf{w}_t and \mathbf{w}_{t+1} for their respective update steps, we get:

$$\begin{aligned} \ell_{\mathcal{P}}(\mathbf{x}_{1:t}, y_{1:t}, \mathbf{w}_{t+1}) + \frac{\eta}{2} \|\mathbf{w}_{t+1} - \mathbf{w}_t\|_2^2 &\leq \ell_{\mathcal{P}}(\mathbf{x}_{1:t}, y_{1:t}, \mathbf{w}_t) \\ \ell_{\mathcal{P}}(\mathbf{x}_{1:t-1}, y_{1:t-1}, \mathbf{w}_{t+1}) &\geq \ell_{\mathcal{P}}(\mathbf{x}_{1:t-1}, y_{1:t-1}, \mathbf{w}_t) + \frac{\eta}{2} \|\mathbf{w}_{t+1} - \mathbf{w}_t\|_2^2, \end{aligned}$$

which when subtracted, give us

$$\eta \|\mathbf{w}_{t+1} - \mathbf{w}_t\|_2^2 \leq \mathcal{L}_t(\mathbf{w}_t) - \mathcal{L}_t(\mathbf{w}_{t+1}) \leq G_t \|\mathbf{w}_{t+1} - \mathbf{w}_t\|_2, \quad (8)$$

where the last inequality follows using the Lipschitz continuity of \mathcal{L}_t . We now use the fact that

$$\sum_{t=1}^T \mathcal{L}_t(\mathbf{w}_t) = \sum_{t=1}^T \mathcal{L}_t(\mathbf{w}_{t+1}) + \sum_{t=1}^T (\mathcal{L}_t(\mathbf{w}_t) - \mathcal{L}_t(\mathbf{w}_{t+1})),$$

along with (7) and (8) to get

$$\sum_{t=1}^T \mathcal{L}_t(\mathbf{w}_t) \leq \mathbf{x}_{1:T}, y_{1:T}, \mathbf{w}_* + \frac{\eta}{2} \|\mathbf{w}_*\|_2^2 + \frac{\sum_{t=1}^T G_t^2}{\eta}.$$

The result now follows by selecting $\eta = \sqrt{2 \sum_{t=1}^T G_t^2 / \|\mathbf{w}_*\|_2^2}$.

Lemma 9. *For the setting described in Theorem 1, we have*

$$\sum_{t=1}^T \mathcal{L}_t(\mathbf{w}_{t+1}) \leq \mathbf{x}_{1:T}, y_{1:T}, \mathbf{w}_* + \frac{\eta}{2} \|\mathbf{w}_*\|_2^2$$

Proof. Let $\mathcal{L}_0(\mathbf{w}) := \frac{\eta}{2} \|\mathbf{w}\|_2^2$. Thus, we can equivalently write the FTRL update in (FTRL) as

$$\mathbf{w}_{t+1} = \arg \min_{\mathbf{w} \in \mathcal{W}} \sum_{\tau=0}^t \mathcal{L}_\tau(\mathbf{w}).$$

Now, using the optimality of \mathbf{w}_{t+1} at time t , we get

$$\sum_{\tau=0}^t \mathcal{L}_\tau(\mathbf{w}_{t+1}) \leq \sum_{\tau=0}^t \mathcal{L}_\tau(\mathbf{w}_*) \quad (9)$$

Combining this with the optimality of \mathbf{w}_t at time $t - 1$, we get

$$\sum_{\tau=0}^{t-1} \mathcal{L}_\tau(\mathbf{w}_t) + \mathcal{L}_t(\mathbf{w}_{t+1}) \leq \sum_{\tau=0}^t \mathcal{L}_\tau(\mathbf{w}_{t+1}) \leq \sum_{\tau=0}^t \mathcal{L}_\tau(\mathbf{w}_*) \quad (10)$$

Repeating this argument gives us

$$\sum_{\tau=0}^t \mathcal{L}_\tau(\mathbf{w}_{\tau+1}) \leq \sum_{\tau=0}^t \mathcal{L}_\tau(\mathbf{w}_*),$$

which proves the result. \square

B Proof of Lemma 2

We consider the following four exhaustive cases in turn:

Case 1. $z_{i_k} \geq z_{j_k}$ and $z'_{j_k} \geq z'_{i_k}$

We have the following set of inequalities

$$\begin{aligned} g(z_{i_k}) &= g(\langle \mathbf{w}, \mathbf{x}_{i_k} \rangle - c_i) \\ &\leq g(\langle \mathbf{w}', \mathbf{x}_{i_k} \rangle - c_i) + |\langle \mathbf{w} - \mathbf{w}', \mathbf{x}_{i_k} \rangle| \\ &\leq g(\langle \mathbf{w}', \mathbf{x}_{i_k} \rangle - c_i) + \|\mathbf{w} - \mathbf{w}'\|_2 \\ &= g(z'_{i_k}) + \|\mathbf{w} - \mathbf{w}'\|_2 \\ &\leq g(z'_{j_k}) + \|\mathbf{w} - \mathbf{w}'\|_2, \end{aligned}$$

where the first inequality follows by the Lipschitz assumption, the second follows by Cauchy-Schwartz inequality and the last follows by the case assumption $z'_{j_k} \geq z'_{i_k}$ and the fact that g is an increasing function. By renaming $i \leftrightarrow j$ and $\mathbf{w} \leftrightarrow \mathbf{w}'$, we also have $g(z'_{j_k}) \leq g(z_{i_k}) + \|\mathbf{w} - \mathbf{w}'\|_2$. This establishes the result for the specific case.

Case 2. $z_{i_k} \leq z_{j_k}$ and $z'_{j_k} \leq z'_{i_k}$

This case follows similar to the case above.

Case 3. $z_{i_k} \geq z_{j_k}$ and $z'_{j_k} \leq z'_{i_k}$

Using the above conditions z_{j_k} does not belong to the top k elements of z_1, \dots, z_t , but both z'_{i_k} and z'_{j_k} belong to the top k elements of z'_1, \dots, z'_t . Using the pigeonhole principle, there exists an index s such that $z_s \geq z_{i_k}$ but $z_s \leq z'_{j_k}$. Hence, using arguments similar to Case 1, we get the following two bounds:

$$\begin{aligned} |g(z'_{i_k}) - g(z_s)| &\leq \|\mathbf{w} - \mathbf{w}'\|_2, \\ |g(z_s) - g(z'_{j_k})| &\leq \|\mathbf{w} - \mathbf{w}'\|_2. \end{aligned}$$

We also have $|g(z'_{i_k}) - g(z_{i_k})| \leq |\langle \mathbf{w} - \mathbf{w}', \mathbf{x}_{i_k} \rangle| \leq \|\mathbf{w} - \mathbf{w}'\|_2$. Adding these three inequalities gives us the desired result.

Case 4. $z_{i_k} \leq z_{j_k}$ and $z'_{j_k} \geq z'_{i_k}$

This case follows similar to the case above.

These cases are exhaustive and we thus conclude the proof.

C Stability result for $\text{Prec}_{@k}$

Lemma 10. *Let $\ell_{\text{Prec}_{@k}}$ be the surrogate for $\text{Prec}_{@k}$ as defined in (3), $\|\mathbf{x}_t\|_2 \leq 1, \forall t$ and \mathcal{L}_t be defined as in (6). Then $\forall \mathbf{w}, \mathbf{w}' \in \mathcal{W}$, $|\mathcal{L}_t(\mathbf{w}) - \mathcal{L}_t(\mathbf{w}')| \leq 8\|\mathbf{w} - \mathbf{w}'\|_2$.*

Proof. Recall that, the loss function corresponding to $\text{Prec}_{@k}$ is defined as:

$$\ell_{\text{Prec}_{@k}}(\mathbf{x}_{1:t}, y_{1:t}, \mathbf{w}) = \max_{\substack{q \in \{-1, 1\}^t \\ \sum_i (q_i + 1) = 2\lceil kt \rceil}} \sum_{i=1}^t (q_i - y_i) \mathbf{x}_i^T \mathbf{w} - \sum_{i=1}^t q_i y_i \quad (11)$$

$$= \underbrace{\max_{\substack{q \in \{-1, 1\}^t \\ \sum_i (q_i + 1) = 2\lceil kt \rceil}} \sum_{i=1}^t q_i \mathbf{x}_i^T \mathbf{w} - \sum_{i=1}^t q_i y_i}_{A(\mathbf{x}_{1:t}, y_{1:t}, \mathbf{w})} - \underbrace{\sum_{i=1}^t y_i \mathbf{x}_i^T \mathbf{w}}_{B(\mathbf{x}_{1:t}, y_{1:t}, \mathbf{w})} \quad (12)$$

Since $B(\mathbf{x}_{1:t}, y_{1:t}, \mathbf{w})$ is a decomposable loss function, it can at most add a constant (because of the assumptions made by us, that constant can be shown to be no bigger than 1) to the Lipschitz constant of \mathcal{L}_t . Hence we concentrate on bounding the contribution of $A(\mathbf{x}_{1:t}, y_{1:t}, \mathbf{w})$ to the Lipschitz constant of \mathcal{L}_t . Define $z_i = \langle \mathbf{w}, \mathbf{x}_i \rangle - y_i$ and $z'_i = \langle \mathbf{w}', \mathbf{x}_i \rangle - y_i$. It will be useful to rewrite $A(\mathbf{x}_{1:t}, y_{1:t}, \mathbf{w})$ as follows (and drop mentioning the dependence on $\mathbf{x}_{1:t}$ for notational simplicity):

$$p_t(\mathbf{w}) = 2 \max_{\substack{q \in \{1, 0\}^t \\ \sum_i q_i = \lceil kt \rceil}} \sum_{i=1}^t q_i z_i - \sum_{i=1}^t z_i. \quad (13)$$

Similarly, we can define $p_{t-1}(\mathbf{w})$ as well. Now we have

$$\begin{aligned} \mathcal{L}_t(\mathbf{w}) - \mathcal{L}_t(\mathbf{w}') &= p_t(\mathbf{w}) - p_{t-1}(\mathbf{w}) - p_t(\mathbf{w}') + p_{t-1}(\mathbf{w}') + y_t \mathbf{x}_t^T (\mathbf{w}' - \mathbf{w}) \\ &\leq \underbrace{p_t(\mathbf{w}) - p_{t-1}(\mathbf{w}) - p_t(\mathbf{w}') + p_{t-1}(\mathbf{w}')}_{\Delta_t(\mathbf{w}, \mathbf{w}')} + \|\mathbf{w} - \mathbf{w}'\|_2 \end{aligned}$$

Our main goal in the sequel will be to show that $\Delta_t(\mathbf{w}, \mathbf{w}') \leq \mathcal{O}(\|\mathbf{w} - \mathbf{w}'\|_2)$ which shall establish the desired Lipschitz continuity result. Now for both vectors \mathbf{w}, \mathbf{w}' and time instances $t-1, t$, let us denote the optimal assignments as follows:

$$\begin{aligned} a^t &= \arg \max_{\substack{q \in \{1, 0\}^t \\ \sum_i q_i = \lceil kt \rceil}} \sum_{i=1}^t q_i z_i, & b^t &= \arg \max_{\substack{q \in \{1, 0\}^t \\ \sum_i q_i = \lceil kt \rceil}} \sum_{i=1}^t q_i z'_i, \\ a^{t-1} &= \arg \max_{\substack{q \in \{1, 0\}^{(t-1)} \\ \sum_i q_i = \lceil k(t-1) \rceil}} \sum_{i=1}^{t-1} q_i z_i, & b^{t-1} &= \arg \max_{\substack{q \in \{1, 0\}^{(t-1)} \\ \sum_i q_i = \lceil k(t-1) \rceil}} \sum_{i=1}^{t-1} q_i z'_i. \end{aligned}$$

Also, define indices $1 \leq i_r \leq t-1$ and $1 \leq j_s \leq t-1$ as:

$$\begin{aligned} z_{i_1} &\geq z_{i_2} \cdots \geq z_{i_{t-1}}, \\ z'_{j_1} &\geq z'_{j_2} \cdots \geq z'_{j_{t-1}}. \end{aligned}$$

Now, note that (13) involves maximization of a linear function, hence the optimizing assignment q will always lie on the boundary of the Boolean hypercube with the cardinality constraint. Hence, a^t can be obtained by setting $a_{i_r}^t = 1, \forall 1 \leq r \leq \lceil kt \rceil$ and $a_{i_r}^t = 0, \forall r > \lceil kt \rceil$, similarly for b_t . We consider the following two cases and within each, four subcases which establish the result.

In the rest of the proof, all invocations of Lemma 2 shall use the identity function for $g(\cdot)$ and $c_i = y_i$. Clearly this satisfies the prerequisites of Lemma 2 since the identity function is 1-Lipschitz and increasing.

Case 1 $\lceil kt \rceil = \lceil k(t-1) \rceil = \alpha$

Within this, we have the following four exhaustive subcases:

Case 1.1 $z_t \leq z_{i_\alpha}$ and $z'_t \leq z'_{j_\alpha}$

The above condition implies that both $a_t^t = 0$ and $b_t^t = 0$. Furthermore, $a_{1:(t-1)}^t = a^{t-1}$ and $b_{1:(t-1)}^t = b^{t-1}$. As a result we have

$$\Delta_t(\mathbf{w}, \mathbf{w}') = -z_t + z'_t = -\langle \mathbf{w}, \mathbf{x}_t \rangle + \langle \mathbf{w}', \mathbf{x}_t \rangle \leq \|\mathbf{w} - \mathbf{w}'\|_2.$$

Case 1.2 $z_t > z_{i_\alpha}$ and $z'_t \leq z'_{j_\alpha}$

The above condition implies that $a_t^t = 1$ and $b_t^t = 0$. Hence, $b_{1:(t-1)}^t = b^{t-1}$. Also, as a_t^t is turned on, the cardinality constraint dictates that one previously positive index should be turned off. That is, $a_{i_\alpha}^t = 0$, but $a_{i_\alpha}^{t-1} = 1$. Finally, $a_{i_r}^t = a_{i_r}^{t-1}, r \neq \alpha$ and $r < t$. Using the above observations, we have the following sequence of inequalities:

$$\begin{aligned} \Delta_t(\mathbf{w}, \mathbf{w}') &= (2(z_t - z_{i_\alpha}) - z_t) - (0 - z'_t) \\ &= (z_t - z_{i_\alpha}) + (z'_t - z_{i_\alpha}) \\ &= (z_t - z'_t) + 2(z'_t - z_{i_\alpha}) \\ &\leq (z_t - z'_t) + 2(z'_{j_\alpha} - z_{i_\alpha}) \\ &\leq 7\|\mathbf{w} - \mathbf{w}'\|_2, \end{aligned}$$

where the third inequality follows from the case assumptions and the final inequality follows from an application of Cauchy Schwartz inequality and Lemma 2.

Case 1.3 $z_t \leq z_{i_\alpha}$ and $z'_t > z'_{j_\alpha}$

In this case, we can analyze similarly to get

$$\begin{aligned} \Delta_t(\mathbf{w}, \mathbf{w}') &= (0 - z_t) - (2(z'_t - z'_{j_\alpha}) - z'_t) \\ &= (z'_{j_\alpha} - z_t) + (z'_{j_\alpha} - z'_t) \\ &= (z'_t - z_t) + 2(z'_{j_\alpha} - z'_t) \\ &\leq (z'_t - z_t) \\ &\leq 3\|\mathbf{w} - \mathbf{w}'\|_2. \end{aligned}$$

Case 1.4 $z_t > z_{i_\alpha}$ and $z'_t > z'_{j_\alpha}$

In this case, both $a_t^t = 1$ and $b_t^t = 1$. Hence, both $a_{i_\alpha}^t = 0$ and $b_{j_\alpha}^t = 0$. The remaining terms of a^t and a^{t-1} (similarly for b^t and b^{t-1}) remain the same. That is, we have

$$\Delta_t(\mathbf{w}, \mathbf{w}') = (2(z_t - z_{i_\alpha}) - z_t) - (2(z'_t - z'_{j_\alpha}) - z'_t)$$

$$\begin{aligned}
&= (z_t - z'_t) - 2(z_{i_\alpha} - z'_{j_\alpha}) \\
&\leq 7 \|\mathbf{w} - \mathbf{w}'\|_2.
\end{aligned}$$

Case 2 $\lceil kt \rceil = \lceil k(t-1) \rceil + 1 = \alpha$

Here again, we consider the following four exhaustive subcases:

Case 2.1 $z_t \leq z_{i_\alpha}$ and $z'_t \leq z'_{j_\alpha}$

The above condition implies that $a_t^t = 0$ and $b_t^t = 0$. Also, one new positive is included in both a^t and b^t , i.e., $a_{i_\alpha}^t = 1$ and $b_{j_\alpha}^t = 1$. The remaining entries of a^t and b^t remains the same. Hence,

$$\Delta_t(\mathbf{w}, \mathbf{w}') = (2z_{i_\alpha} - z_t) - (2z'_{j_\alpha} - z'_t) = 2(z_{i_\alpha} - z'_{j_\alpha}) - (z_t - z'_t) \leq 9 \|\mathbf{w} - \mathbf{w}'\|_2.$$

Case 2.2 $z_t > z_{i_\alpha}$ and $z'_t \leq z'_{j_\alpha}$

The above condition implies that $a_t^t = 1$ and $b_t^t = 0$. Also, $b_{j_\alpha}^t = 1$. The remaining entries of a^t and b^t remains the same. Hence we have

$$\begin{aligned}
\Delta_t(\mathbf{w}, \mathbf{w}') &= (2z_t - z_t) - (2z'_{j_\alpha} - z'_t) \\
&= (z_t - z'_{j_\alpha}) + (z'_t - z'_{j_\alpha}) \\
&= (z_t - z'_t) + 2(z'_t - z'_{j_\alpha}) \\
&\leq 3 \|\mathbf{w} - \mathbf{w}'\|_2.
\end{aligned}$$

Case 2.3 $z_t \leq z_{i_\alpha}$ and $z'_t > z'_{j_\alpha}$

In this case we have

$$\begin{aligned}
\Delta_t(\mathbf{w}, \mathbf{w}') &= (2z_{i_\alpha} - z_t) - (2z'_t - z'_t) \\
&= (z_{i_\alpha} - z_t) + (z_{i_\alpha} - z'_t) \\
&= (z'_t - z_t) + 2(z_{i_\alpha} - z'_t) \\
&\leq (z'_t - z_t) + 2(z_{i_\alpha} - z'_{j_\alpha}) \\
&\leq 7 \|\mathbf{w} - \mathbf{w}'\|_2.
\end{aligned}$$

Case 2.4 $z_t > z_{i_\alpha}$ and $z'_t > z'_{j_\alpha}$

The above condition implies that $a_t^t = 1$ and $b_t^t = 1$. The remaining entries of a^t and b^t remains the same. Hence,

$$\Delta_t(\mathbf{w}, \mathbf{w}') = (2z_t - z_t) - (2z'_t - z'_t) = z_t - z'_t \leq 3 \|\mathbf{w} - \mathbf{w}'\|_2.$$

Taking the worst case Lipschitz constants from these 8 subcases and adding the contribution of $B(\mathbf{x}_{1:t}, y_{1:t}, \mathbf{w})$ concludes the proof. \square

D Extension to Precision-Recall Break Even Point (PRBEP)

We note that the above discussion can easily be extended to prove stability results for the structural surrogate loss for the PRBEP performance measure [3]. Recall that the PRBEP measure essentially measures the precision (equivalently recall) of a predictor when thresholded at a point that equates

the precision and recall. Since we have $\text{Prec} = \frac{\text{TP}}{\text{TP} + \text{FP}}$ and $\text{Rec} = \frac{\text{TP}}{\text{TP} + \text{FN}}$, the break even point is reached at a threshold where $\text{TP} + \text{FP} = \text{TP} + \text{FN}$. Notice that the left hand side equals the number of points that are predicted as positive whereas the right hand side equals the number of points that are actual positives.

Thus, the PRBEP is achieved at a threshold that predicts as many points as positive as there are actual positives which gives us the formal definition of this performance measure

$$\text{PRBEP}(\mathbf{w}) := \sum_{j: \mathbb{T}\left(\frac{t_+}{t}, t\right)(\mathbf{x}_j, \mathbf{w})=1} \mathbb{I}[y_j = 1]. \quad (14)$$

Note that this is equivalent to the definition of $\text{Prec}_{@k}$ with $k = \frac{t_+}{t}$. Correspondingly, we can also define the structural SVM surrogate for this performance measure as

$$\ell_{\text{PRBEP}}(\mathbf{w}) = \max_{\substack{\bar{\mathbf{y}} \in \{-1, +1\}^t \\ \sum_i (\bar{y}_i + 1) = 2t_+}} \sum_{i=1}^t (\bar{y}_i - y_i) \mathbf{x}_i^T \mathbf{w} - \sum_{i=1}^t y_i \bar{y}_i. \quad (15)$$

Given this, it is easy to see that the proof of Lemma 10 would apply to this case as well. The only difference in applying the analysis would be that Case 1 and its subcases would apply when $y_t < 0$ which is when the incoming point is negative and hence the number of actual positives in the stream does not go up. Case 2 and its subcases would apply when $y_t > 0$ in which case the number of points to be considered while calculating precision would have to be increased by 1.

E Online-to-batch Conversion

This section presents a proof of the regret bound in the batch model considered in Theorem 4 and a proof sketch of the online-to-batch conversion result. The full proof shall appear in the full version of the paper. We will consider in this section, the pAUC measure in the **2PMB** setting wherein positives are assumed to reside in the buffer and negatives are streaming in. The case of the $\text{Prec}_{@k}$ measure in the usual **1PMB** setting can be handled similarly. Additionally, we will show in Appendix G that for the case of pAUC, the contributions from a large enough buffer of randomly chosen positive points mimics the contributions of the entire population of positive points. Thus, for pAUC, it suffices to show the online-to-batch conversion bounds just with respect to the negatives. We clarify this further in the discussion.

E.1 Regret Bounds in the Modified Framework

We prove the following lemma which will help us in instantiating our online-to-batch conversion proofs.

Lemma 11. *For the surrogate losses of $\text{Prec}_{@k}$ and pAUC, we have $R(T, s) \leq \sqrt{s} \cdot R(T)$*

Proof. The only thing we need to do is analyze one time step for changes in the Lipschitz constant. Fix a time step t and let $\mathbf{Z}_t = \{\mathbf{x}_{t,1}, \mathbf{x}_{t,2}, \dots, \mathbf{x}_{t,s}\}$. Also, let $g_t(\mathbf{w}, i) := \ell_{\mathcal{P}}(\mathbf{Z}_1, \dots, \mathbf{Z}_{t-1}, \mathbf{x}_{t,1:i}, \mathbf{w})$ for any $i = 1 \dots s$ (note that this gives us $g_t(\mathbf{w}, s) = \ell_{\mathcal{P}}(\mathbf{Z}_1, \dots, \mathbf{Z}_t, \mathbf{w})$). Also let us abuse notation to denote $g_t(\mathbf{w}, 0) := \ell_{\mathcal{P}}(\mathbf{Z}_1, \dots, \mathbf{Z}_{t-1}, \mathbf{w}) = g_{t-1}(\mathbf{w}, s)$. Let the Lipschitz constant in the model

with batch size s be denoted as G_t^s . Thus, we have $G_t^1 = G_t$, the Lipschitz constant for the problem in the original model (i.e. for $s = 1$). Then we have, for any $\mathbf{w}, \mathbf{w}' \in \mathcal{W}$,

$$\begin{aligned}
|\mathcal{L}_t(\mathbf{w}) - \mathcal{L}_t(\mathbf{w}')| &= |\ell_{\mathcal{P}}(\mathbf{Z}_{1:t}, \mathbf{w}) - \ell_{\mathcal{P}}(\mathbf{Z}_{1:t-1}, \mathbf{w}) - \ell_{\mathcal{P}}(\mathbf{Z}_{1:t}, \mathbf{w}') + \ell_{\mathcal{P}}(\mathbf{Z}_{1:t-1}, \mathbf{w}')| \\
&= |g_t(\mathbf{w}, s) - g_t(\mathbf{w}, 0) - g_t(\mathbf{w}', s) + g_t(\mathbf{w}', 0)| \\
&= \left| \sum_{i=1}^s g_t(\mathbf{w}, i) - g_t(\mathbf{w}, i-1) - g_t(\mathbf{w}', i) + g_t(\mathbf{w}', i-1) \right| \\
&\leq \sum_{i=1}^s |g_t(\mathbf{w}, i) - g_t(\mathbf{w}, i-1) - g_t(\mathbf{w}', i) + g_t(\mathbf{w}', i-1)| \\
&\leq \sum_{i=1}^s G_t \|\mathbf{w} - \mathbf{w}'\| = G_t \cdot s \|\mathbf{w} - \mathbf{w}'\|,
\end{aligned}$$

where the first inequality follows by triangle inequality and the second inequality follows by a repeated application of the Lipschitz property of these loss functions in the original online model (i.e. with batch size $s = 1$). This establishes the Lipschitz constant in this model as $G_t^s \leq s \cdot G_t$. Now, the usual FTRL analysis gives us the following (note that there are only T/s time steps now)

$$\sum_{t=1}^{T/s} \mathcal{L}_t(\mathbf{w}_t) \leq \ell_{\mathcal{P}}(\mathbf{x}_{1:T}, y_{1:T}, \mathbf{w}_*) + \frac{\eta}{2} \|\mathbf{w}_*\|_2^2 + \frac{\sum_{t=1}^{T/s} (G_t^s)^2}{\eta} \leq \ell_{\mathcal{P}}(\mathbf{x}_{1:T}, y_{1:T}, \mathbf{w}_*) + 2s \|\mathbf{w}_*\|_2 \sqrt{\sum_{t=1}^{T/s} G_t^2},$$

by setting η appropriately. Now, for $\text{Prec}_{@k}$, $G_t \leq 8$. Thus, we have

$$\frac{1}{T} \sum_{t=1}^{T/s} \mathcal{L}_t(\mathbf{w}_t) \leq \frac{1}{T} \ell_{\mathcal{P}}(\mathbf{x}_{1:T}, y_{1:T}, \mathbf{w}_*) + 6 \|\mathbf{w}_*\|_2 \sqrt{\frac{s}{T}},$$

which establishes the result for $\text{Prec}_{@k}$. Similarly, for pAUC, we can show that the regret in the batch model does not worsen by more than a factor of \sqrt{s} . \square

E.2 Online-to-batch Conversion for pAUC

We will consider the **2PMB** setting where negative points come as a stream and positive points reside in an in-memory buffer. At each trial t , the learner receives a batch of s negative points $\mathbf{Z}_t^- = \{\mathbf{x}_{t,1}^-, \dots, \mathbf{x}_{t,s}^-\}$ (we shall assume throughout, for simplicity, that $s\beta$ is an integer). Let us denote the loss w.r.t all the positive points in the buffer by $\phi_+ : \mathcal{W} \times \mathbb{R} \rightarrow [0, B]$. ϕ_+ is defined using a loss function $g(\cdot)$ such as hinge loss or logistic loss as

$$\phi_+(\mathbf{w}, c) = \frac{1}{B} \sum_{i=1}^B g(\mathbf{w}^\top \mathbf{x}_i^+ - c)$$

For sake of brevity, we will abbreviate $\phi_+(\mathbf{w}, c)$ as $\phi_+(c)$, the reference to \mathbf{w} being clear from context. We assume that ϕ_+ is monotonically increasing (as is the case for hinge loss and logistic regression) and bounded i.e. for some fixed $B > 0$, we have, for all $\mathbf{w} \in \mathcal{W}, c \in \mathbb{R}, 0 \leq \phi_+(\mathbf{w}, c) \leq B$.

The empirical (unnormalized) partial AUC loss for a model $\mathbf{w} \in \mathcal{W} \subseteq \mathbb{R}^d$ over the negative points received in t trials is then given by

$$\tilde{\ell}_{\text{pAUC}}(\mathbf{Z}_{1:t}^-, \mathbf{w}) = \sum_{\tau=1}^t \sum_{q=1}^s \mathbb{T}_{\beta,t}^-(\mathbf{x}_{\tau,q}^-, \mathbf{w}) \phi_+(\mathbf{w}^\top \mathbf{x}_{\tau,q}^-),$$

where $\mathbb{T}_{\beta,t}^-(\mathbf{x}^-, \mathbf{w})$ is the (empirical) indicator function that is turned on whenever \mathbf{x}^- appears in the top- β fraction of all the negatives seen till now, ordered by \mathbf{w} , i.e. $\mathbb{T}_{\beta,t}^-(\mathbf{x}^-, \mathbf{w}) = 1$ whenever $|\{\tau \in [t], q \in [s] : \mathbf{w}^\top \mathbf{x}^- > \mathbf{w}^\top \mathbf{x}_{\tau,q}^-\}| \leq t s \beta$. We similarly define a population version of this empirical loss function as

$$\tilde{\mathcal{R}}_{\text{pAUC}}(\mathbf{w}) = \mathbb{E}_{\mathbf{x}^-} \left[\mathbb{T}_{\beta}^-(\mathbf{x}^-, \mathbf{w}) \phi_+(\mathbf{w}^\top \mathbf{x}^-) \right],$$

where $\mathbb{T}_{\beta}^-(\mathbf{x}^-, \mathbf{w})$ is the population indicator function with $\mathbb{T}_{\beta}^-(\mathbf{x}^-, \mathbf{w}) = 1$ whenever $\mathbb{P}_{\tilde{\mathbf{x}}^-}(\mathbf{w}^\top \tilde{\mathbf{x}}^- > \mathbf{w}^\top \mathbf{x}^-) \leq \beta$. Also, we define $\mathcal{L}_t(\mathbf{w}) = \ell_{\text{pAUC}}(\mathbf{Z}_{1:t}^-, \mathbf{w}) - \ell_{\text{pAUC}}(\mathbf{Z}_{1:t-1}^-, \mathbf{w})$, with the regret of a learning algorithm that generates an ensemble of models $\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_{T/s} \in \mathcal{W} \subseteq \mathbb{R}^d$ upon receiving T/s batches of negative points $\mathbf{Z}_{1:T/s}^-$ defined as:

$$R(T, s) = \frac{1}{T} \sum_{t=1}^{T/s} \mathcal{L}_t(\mathbf{w}_t) - \arg \min_{\mathbf{w} \in \mathcal{W}} \frac{1}{T} \tilde{\ell}_{\text{pAUC}}(\mathbf{Z}_{1:T/s}^-, \mathbf{w}).$$

Define $\beta_t = \mathbb{E}_{\mathbf{x}^-} \left[\mathbb{T}_{\beta,t-1}^-(\mathbf{x}^-, \mathbf{w}_t) \right]$ as the fraction of the population that can appear in the top β fraction of the set of points seen till now, i.e. the fraction of the population for which the empirical indicator function is turned on, and

$$\mathcal{Q}_t(\mathbf{w}) = \mathbb{E}_{\mathbf{x}^-} \left[\mathbb{T}_{\beta,t-1}^-(\mathbf{x}^-, \mathbf{w}) \phi_+(\mathbf{w}^\top \mathbf{x}^-) \right]$$

as the population partial AUC computed with respect to the empirical indicator function $\mathbb{T}_{\beta,t-1}^-$ (note that the population risk functional $\tilde{\mathcal{R}}_{\text{pAUC}}(\mathbf{w})$ is computed with respect to $\mathbb{T}_{\beta}^-(\mathbf{x}^-, \mathbf{w})$, the population indicator function instead). We will also find it useful to define the following conditional expectation.

$$\tilde{\mathcal{L}}_t(\mathbf{w}) = \mathbb{E}_{\mathbf{Z}_t^-} \left[\mathcal{L}_t(\mathbf{w}) \mid \mathbf{Z}_{1:t-1}^- \right].$$

We now present a proof sketch of the online-to-batch conversion result in Theorem 4 for pAUC.

Theorem 12 (Online-to-batch Conversion for pAUC). *Suppose the sequence of negative points $\mathbf{x}_1^-, \dots, \mathbf{x}_T^-$ is generated i.i.d.. Let us partition this sequence into T/s batches of size s and let $\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_{T/s}$ be an ensemble of models generated by an online learning algorithm upon receiving these T/s batches. Suppose the online learning algorithm has a guaranteed regret bound $R(T, s)$. Then for $\bar{\mathbf{w}} = \frac{1}{T/s} \sum_{t=1}^{T/s} \mathbf{w}_t$, any $\mathbf{w}^* \in \mathcal{W} \subseteq \mathbb{R}^d$, $\epsilon \in (0, 1]$ and $\delta > 0$, with probability at least $1 - \delta$,*

$$\tilde{\mathcal{R}}_{\text{pAUC}}(\bar{\mathbf{w}}) \leq (1 + \epsilon) \tilde{\mathcal{R}}_{\text{pAUC}}(\mathbf{w}^*) + \frac{1}{\beta} R(T, s) + e^{-\Omega(s\epsilon^2)} + \tilde{\mathcal{O}} \left(\sqrt{\frac{s \ln(1/\delta)}{T}} \right).$$

In particular, setting $s = \tilde{\mathcal{O}}(\sqrt{T})$ and $\epsilon = \sqrt[4]{1/T}$ gives us, with probability at least $1 - \delta$,

$$\tilde{\mathcal{R}}_{\text{pAUC}}(\bar{\mathbf{w}}) \leq \tilde{\mathcal{R}}_{\text{pAUC}}(\mathbf{w}^*) + \frac{1}{\beta} R(T, \sqrt{T}) + \tilde{\mathcal{O}} \left(\sqrt[4]{\frac{\ln(1/\delta)}{T}} \right).$$

Proof (Sketch). Fix $\epsilon \in (0, 0.5]$. We wish to bound the difference

$$(1 - \epsilon)s\beta \sum_{t=1}^{T/s} \tilde{\mathcal{R}}_{\text{pAUC}}(\mathbf{w}_t) - T\beta\tilde{\mathcal{R}}_{\text{pAUC}}(\mathbf{w}_*) \quad (16)$$

and do so by decomposing (16) into four terms as shown below.

$$(16) \leq \sum_{t=1}^{T/s} RE_t(\mathbf{w}_t) + MC(\mathbf{w}_{1:T/s}) + R(\mathbf{w}_{1:T/s}) + UC(\mathbf{w}_*),$$

where we have

$$\begin{aligned} UC(\mathbf{w}_*) &= \tilde{\ell}_{\text{pAUC}}(\mathbf{Z}_{1:T/s}^-, \mathbf{w}_*) - T\beta\tilde{\mathcal{R}}_{\text{pAUC}}(\mathbf{w}_*) && \text{(Uniform Convergence Term)} \\ R(\mathbf{w}_{1:T/s}) &= \sum_{t=1}^{T/s} \mathcal{L}_t(\mathbf{w}_t) - \sum_{t=1}^{T/s} \mathcal{L}_t(\mathbf{w}_*) && \text{(Regret Term)} \\ MC(\mathbf{w}_{1:T/s}) &= \sum_{t=1}^{T/s} \tilde{\mathcal{L}}_t(\mathbf{w}_t) - \sum_{t=1}^{T/s} \mathcal{L}_t(\mathbf{w}_t) && \text{(Martingale Convergence Terms)} \\ RE_t(\mathbf{w}_t) &= (1 - \epsilon)s\beta\tilde{\mathcal{R}}_{\text{pAUC}}(\mathbf{w}_t) - \tilde{\mathcal{L}}_t(\mathbf{w}_t) && \text{(Residual Error Terms)} \end{aligned}$$

Note that the above has used the fact that $\tilde{\ell}_{\text{pAUC}}(\mathbf{Z}_{1:T/s}^-, \mathbf{w}_*) = \sum_{t=1}^{T/s} \mathcal{L}_t(\mathbf{w}_*)$. We will bound these terms in order below. First we look at the term $UC(\mathbf{w}_*)$. Bounding this simply requires a batch generalization bound of the form we prove in Theorem 7. Thus, we can show, that with probability $1 - \delta/3$, we have

$$UC(\mathbf{w}_*) \leq \mathcal{O}\left(\sqrt{T \log(1/\delta)}\right).$$

We now move on the term $R(\mathbf{w}_{1:T/s})$. This is simply bounded by the regret of the ensemble $\mathbf{w}_{1:T/s}$. This gives us

$$R(\mathbf{w}_{1:T/s}) \leq T \cdot R(T, s).$$

The next term we bound is $MC(\mathbf{w}_{1:T/s})$. Note that by definition of $\tilde{\mathcal{L}}_t(\mathbf{w})$, if we define

$$v_t = \tilde{\mathcal{L}}_t(\mathbf{w}_t) - \mathcal{L}_t(\mathbf{w}_t),$$

then the terms $\{v_t\}$ form a martingale difference sequence. Since $|\tilde{\mathcal{L}}_t(\mathbf{w}_t) - \mathcal{L}_t(\mathbf{w}_t)| \leq \mathcal{O}(s)$, we get, by an application of the Azuma-Hoeffding inequality, with probability at least $1 - \delta/3$,

$$MC(\mathbf{w}_{1:T/s}) \leq \mathcal{O}\left(s\sqrt{\frac{T}{s} \ln \frac{1}{\delta}}\right) = \mathcal{O}\left(\sqrt{sT \ln(1/\delta)}\right).$$

The last step requires us to bound the residual term $RE_t(\mathbf{w}_t)$ which will again require uniform convergence techniques. We shall show, that with probability, at least $1 - (\delta \cdot s/3T)$, we have

$$\beta_t \geq \beta - \tilde{\mathcal{O}}\left(\sqrt{\frac{\log \frac{1}{\delta}}{s(t-1)}}\right).$$

This shall allow us to show that with the same probability, we have

$$\mathcal{Q}_t(\mathbf{w}_t) - \tilde{\mathcal{R}}_{\text{pAUC}}(\mathbf{w}_t) \leq \tilde{\mathcal{O}} \left(\sqrt{\frac{\log \frac{1}{\delta}}{s(t-1)}} \right).$$

The last ingredient in the proof shall involve showing that the following holds for any $\epsilon > 0$

$$\tilde{\mathcal{L}}_t(\mathbf{w}_t) \geq (1 - \epsilon)s\beta_t\mathcal{Q}_t(\mathbf{w}_t) - \Omega(s \exp(-s\beta_t^2\epsilon^2))$$

Combining the above with a union bound will show us that, with probability at least $1 - \delta/3$,

$$\sum_{i=1}^{T/s} RE_t(\mathbf{w}_t) \leq \mathcal{O}(T \exp(-s\epsilon^2)) + \tilde{\mathcal{O}} \left(\sqrt{sT \log(1/\delta)} \right)$$

A final union bound and some manipulations would then establish the claimed result. \square

F Proof of Theorem 6

The proof proceeds in two parts: the first part uses the fact that the **1PMB** method essentially simulates the GIGA method of [24] with the non-decomposable loss function and the second part uses the uniform convergence properties of the loss function to establish the error bound. To proceed, let us set up some notation. Consider the e^{th} epoch of the **1PMB** algorithm. Let us denote the set of points considered in this epoch by $X_e = \{x_1^e, \dots, x_s^e\}$. With this notation it is clear that the **1PMB** algorithm can be said to be performing online gradient descent with respect to the instantaneous loss functions $\mathcal{L}_e(\mathbf{w}) = \mathcal{L}(X_e, \mathbf{w}) := \ell_{\mathcal{P}}(\mathbf{x}_{1:s}^e, y_{1:s}^e, \mathbf{w})$.

Since the loss function $\mathcal{L}_e(\mathbf{w})$ is convex, the standard analysis for online convex optimization would apply under mild boundedness assumptions on the domain and the (sub)gradients of the loss function. Since there are n/s epochs (assuming for simplicity that n is a multiple of s), this allows us to use the standard regret bounds [24] to state the following:

$$\frac{s}{n} \sum_{e=1}^{n/s} \mathcal{L}_e(\mathbf{w}_e) \leq \frac{s}{n} \sum_{e=1}^{n/s} \mathcal{L}_e(\mathbf{w}_*) + \mathcal{O} \left(\sqrt{\frac{s}{n}} \right).$$

Now we will invoke uniform convergence properties of the loss function. However, doing so requires clarifying certain aspects of the problem setting. The statement of Theorem 6 assumes only a random ordering of training data points whereas uniform convergence properties typically require i.i.d. samples. We reconcile this by noticing that all our uniform convergence proofs use the Hoeffding's lemma to establish statistical convergence and that the Hoeffding's lemma holds when random variables are sampled without replacement as well (e.g. see [25]). Since a random ordering of the data provides, for each epoch, a uniformly random sample without replacement, we are able to invoke the uniform convergence proofs.

Thus, if we denote $\mathcal{L}(\mathbf{w}) := \ell_{\mathcal{P}}(\mathbf{x}_{1:n}, y_{1:n}, \mathbf{w})$, then by using the uniform convergence properties of the loss function, for every e , with probability at least $1 - \frac{s\delta}{n}$, we have $\mathcal{L}_e(\mathbf{w}_e) \geq \mathcal{L}(\mathbf{w}_e) - \alpha \left(s, \frac{s\delta}{n} \right)$ as well as $\mathcal{L}_e(\mathbf{w}_*) \leq \mathcal{L}(\mathbf{w}_*) + \alpha \left(s, \frac{s\delta}{n} \right)$. Applying the union bound and Jensen's inequality gives us, with probability at least $1 - \delta$, the desired result:

$$\mathcal{L}(\bar{\mathbf{w}}) \leq \frac{s}{n} \sum_{e=1}^{n/s} \mathcal{L}(\mathbf{w}_e) \leq \mathcal{L}(\mathbf{w}_*) + 2\alpha \left(s, \frac{s\delta}{n} \right) + \mathcal{O} \left(\sqrt{\frac{s}{n}} \right).$$

We note that we can use similar arguments as above to give error bounds for the **2PMB** procedure as well. Suppose $\bar{\mathbf{x}}_{1:s_+}^+$ and $\bar{\mathbf{x}}_{1:s_-}^-$ are the positive and negative points sampled in the process (note that here the number of positive and negatives points (i.e. s_+ and s_- respectively) are random quantities as well). Also suppose $\mathbf{x}_{1:n_+}^+$ and $\mathbf{x}_{1:n_-}^-$ are the positive and negative points in the population. Then recall that Definition 5 requires, for a uniform (but possibly without replacement) sample,

$$\sup_{\mathbf{w} \in \mathcal{W}} \left| \ell_{\mathcal{P}}(\mathbf{x}_{1:n_+}^+, \mathbf{x}_{1:n_-}^-, \mathbf{w}) - \ell_{\mathcal{P}}(\bar{\mathbf{x}}_{1:s_+}^+, \bar{\mathbf{x}}_{1:s_-}^-, \mathbf{w}) \right| \leq \tilde{\mathcal{O}}(\alpha(s, \delta)).$$

To prove bounds for **2PMB**, we require that for arbitrary choice of $s_+, s_- \geq \Omega(s)$, when $\bar{\mathbf{x}}_{1:s_+}^+$ and $\bar{\mathbf{x}}_{1:s_-}^-$ are chosen separately and uniformly (but yet again possibly without replacement) from $\mathbf{x}_{1:n_+}^+$ and $\mathbf{x}_{1:n_-}^-$ respectively, we still obtain a similar result as above. Since the first pass and each epoch of the second pass provide such a sample, we can use this result to prove error bounds for the **2PMB** procedure. We defer the detailed arguments for such results to the full version of the paper.

We however note that the proof of Theorem 7 below does indeed prove such a result for the pAUC loss function by effectively proving (see Section G.1) the following two results

$$\begin{aligned} \sup_{\mathbf{w} \in \mathcal{W}} \left| \ell_{\mathcal{P}}(\mathbf{x}_{1:n_+}^+, \bar{\mathbf{x}}_{1:s_-}^-, \mathbf{w}) - \ell_{\mathcal{P}}(\bar{\mathbf{x}}_{1:s_+}^+, \bar{\mathbf{x}}_{1:s_-}^-, \mathbf{w}) \right| &\leq \tilde{\mathcal{O}}(\alpha(s, \delta)) \\ \sup_{\mathbf{w} \in \mathcal{W}} \left| \ell_{\mathcal{P}}(\mathbf{x}_{1:n_+}^+, \mathbf{x}_{1:n_-}^-, \mathbf{w}) - \ell_{\mathcal{P}}(\mathbf{x}_{1:n_+}^+, \bar{\mathbf{x}}_{1:s_-}^-, \mathbf{w}) \right| &\leq \tilde{\mathcal{O}}(\alpha(s, \delta)). \end{aligned}$$

G Uniform Convergence Bounds for Partial Area under the ROC Curve

In this section we present a proof sketch of Theorem 7 which we restate below for convenience.

Theorem 13. *Consider any convex, monotonic and Lipschitz classification surrogate $\phi : \mathbb{R} \rightarrow \mathbb{R}_+$. Then the loss function for the $(0, \beta)$ -partial AUC performance measure defined as follows exhibits uniform convergence at the rate $\alpha(s) = \tilde{\mathcal{O}}(1/\sqrt{s})$:*

$$\ell_{\mathcal{P}}(\mathbf{x}_{1:n}, y_{1:n}, \mathbf{w}) = \frac{1}{\beta n_+ n_-} \sum_{i=1}^n \mathbb{I}[y_i > 0] \sum_{j=1}^n \mathbb{I}[y_j < 0] \mathbb{T}_{\beta, n}^-(\mathbf{x}_j, \mathbf{w}) \phi\left(\mathbf{w}^\top (\mathbf{x}_i - \mathbf{x}_j)\right),$$

where $n_+ = |\{i : y_i > 0\}|$ and $n_- = |\{i : y_i < 0\}|$.

Proof (Sketch). We shall use the notation $\hat{\mathbb{T}}_{\beta, s}^-$ to denote the indicator function for the top β fraction of the negative elements in the smaller sample of size s . Thus, over the smaller sample $(\bar{\mathbf{x}}_1, \bar{y}_1) \dots (\bar{\mathbf{x}}_s, \bar{y}_s)$, the pAUC is calculated as

$$\ell_{\mathcal{P}}(\bar{\mathbf{x}}_{1:s}, \bar{y}_{1:s}, \mathbf{w}) = \frac{1}{\beta s_+ s_-} \sum_{i=1}^s \mathbb{I}[\bar{y}_i > 0] \sum_{j=1}^s \mathbb{I}[\bar{y}_j < 0] \hat{\mathbb{T}}_{\beta, s}^-(\bar{\mathbf{x}}_j, \mathbf{w}) \phi\left(\mathbf{w}^\top (\bar{\mathbf{x}}_i - \bar{\mathbf{x}}_j)\right).$$

Our goal would be to show that with probability at least $1 - \delta$, for all $\mathbf{w} \in \mathcal{W}$

$$\left| \ell_{\mathcal{P}}(\mathbf{x}_{1:n}, y_{1:n}, \mathbf{w}) - \ell_{\mathcal{P}}(\bar{\mathbf{x}}_{1:s}, \bar{y}_{1:s}, \mathbf{w}) \right| \leq \tilde{\mathcal{O}}\left(\frac{1}{\sqrt{s}}\right)$$

We shall demonstrate this by establishing the following three statements:

1. For any fixed $\mathbf{w} \in \mathcal{W}$, w.h.p., we have $|\ell_{\mathcal{P}}(\mathbf{x}_{1:n}, y_{1:n}, \mathbf{w}) - \ell_{\mathcal{P}}(\bar{\mathbf{x}}_{1:s}, \bar{y}_{1:s}, \mathbf{w})| \leq \tilde{\mathcal{O}}\left(\frac{1}{\sqrt{s}}\right)$
2. For any two $\mathbf{w}, \mathbf{w}' \in \mathcal{W}$, we have $|\ell_{\mathcal{P}}(\mathbf{x}_{1:n}, y_{1:n}, \mathbf{w}) - \ell_{\mathcal{P}}(\mathbf{x}_{1:n}, y_{1:n}, \mathbf{w}')| \leq \mathcal{O}(\|\mathbf{w} - \mathbf{w}'\|_2)$
3. For any two $\mathbf{w}, \mathbf{w}' \in \mathcal{W}$, we have $|\ell_{\mathcal{P}}(\bar{\mathbf{x}}_{1:s}, \bar{y}_{1:s}, \mathbf{w}) - \ell_{\mathcal{P}}(\bar{\mathbf{x}}_{1:s}, \bar{y}_{1:s}, \mathbf{w}')| \leq \mathcal{O}(\|\mathbf{w} - \mathbf{w}'\|_2)$

With these three results established, we would be able to conclude the proof by an application of a standard covering number argument. We now prove these three statements in parts.

G.1 Part 1: Pointwise Convergence for pAUC

Fix a predictor $\mathbf{w} \in \mathcal{W}$ and S_+ and S_- denote the set of positive and negative samples. We shall assume that $s_+, s_- \geq \Omega(s)$ which holds with high probability. Denote, for any \mathbf{x}_i such that $y_i > 0$,

$$\ell^+(\mathbf{x}_i, \mathbf{w}) = \frac{1}{\beta n_-} \sum_{j=1}^n \mathbb{I}[y_j < 0] \mathbb{T}_{\beta, n}^-(\mathbf{x}_j, \mathbf{w}) \phi\left(\mathbf{w}^\top(\mathbf{x}_i - \mathbf{x}_j)\right),$$

and for any $\bar{\mathbf{x}}_i \in S_+$,

$$\ell_{S_-}^+(\bar{\mathbf{x}}_i, \mathbf{w}) = \frac{1}{\beta s_-} \sum_{j=1}^s \mathbb{I}[\bar{y}_j < 0] \hat{\mathbb{T}}_{\beta, s}^-(\bar{\mathbf{x}}_j, \mathbf{w}) \phi\left(\mathbf{w}^\top(\bar{\mathbf{x}}_i - \bar{\mathbf{x}}_j)\right).$$

Notice that $\ell_{\mathcal{P}}(\bar{\mathbf{x}}_{1:s}, \bar{y}_{1:s}, \mathbf{w}) = \frac{1}{n_+} \sum_{i=1}^n \mathbb{I}[y_i > 0] \ell^+(\mathbf{x}_i, \mathbf{w})$ and $\ell_{\mathcal{P}}(\bar{\mathbf{x}}_{1:s}, \bar{y}_{1:s}, \mathbf{w}) = \frac{1}{s_+} \sum_{i=1}^s \mathbb{I}[\bar{y}_i > 0] \ell_{S_-}^+(\bar{\mathbf{x}}_i, \mathbf{w})$. We shall now show the following holds w.h.p. over S_- :

1. For any \mathbf{x}_i such that $y_i > 0$, $\left| \ell^+(\mathbf{x}_i, \mathbf{w}) - \ell_{S_-}^+(\mathbf{x}_i, \mathbf{w}) \right| \leq \tilde{\mathcal{O}}\left(\frac{1}{\sqrt{s}}\right)$.
2. $\frac{1}{n_+} \left| \sum_{i=1}^n \mathbb{I}[y_i > 0] \ell^+(\mathbf{x}_i, \mathbf{w}) - \mathbb{I}[y_i > 0] \ell_{S_-}^+(\mathbf{x}_i, \mathbf{w}) \right| \leq \tilde{\mathcal{O}}\left(\frac{1}{\sqrt{s}}\right)$.
3. $\left| \frac{1}{n_+} \sum_{i=1}^n \mathbb{I}[y_i > 0] \ell_{S_-}^+(\mathbf{x}_i, \mathbf{w}) - \frac{1}{s_+} \sum_{i=1}^s \mathbb{I}[\bar{y}_i > 0] \ell_{S_-}^+(\bar{\mathbf{x}}_i, \mathbf{w}) \right| \leq \tilde{\mathcal{O}}\left(\frac{1}{\sqrt{s}}\right)$.

The second part follows from the first part by an application of the triangle inequality. The third part also can be shown to hold by an application of Hoeffding's inequality and other arguments. This leaves the first part for which we provide a proof in the full version of the paper.

G.2 Parts 2 and 3: Establishing an ϵ -net for pAUC

For simplicity, we assume that the domain is finite. This does not affect the proof in any way since it still allows the domain to be approximated arbitrarily closely by an ϵ -net of (arbitrarily) large size. However, we note that we can establish the same result for infinite domains as well, but choose not to for sake of simplicity. We prove the second part, the proof of the first part being similar. We have

$$\begin{aligned} \left| \ell_{\mathcal{P}}(\mathbf{x}_{1:n}, y_{1:n}, \mathbf{w}) - \ell_{\mathcal{P}}(\mathbf{x}_{1:n}, y_{1:n}, \mathbf{w}') \right| &= \frac{1}{s_+} \left| \sum_{i=1}^s \mathbb{I}[\bar{y}_i > 0] \ell_{S_-}^+(\mathbf{x}_i, \mathbf{w}) - \mathbb{I}[\bar{y}_i > 0] \ell_{S_-}^+(\mathbf{x}_i, \mathbf{w}') \right| \\ &\leq \frac{1}{s_+} \sum_{i=1}^s \left| \mathbb{I}[\bar{y}_i > 0] \left(\ell_{S_-}^+(\mathbf{x}_i, \mathbf{w}) - \ell_{S_-}^+(\mathbf{x}_i, \mathbf{w}') \right) \right| \\ &\leq \mathcal{O}(\|\mathbf{w} - \mathbf{w}'\|_2), \end{aligned}$$

using Lemma 2 with $g(a) = \phi(\mathbf{w}^\top \mathbf{x}_i - a)$ and $c_i = 0$. This concludes the proof. \square

H Methodology for implementing 1PMB and 2PMB for pAUC tasks

In this section we clarify the mechanisms used to implement the **1PMB** and **2PMB** routines. Going as per the dataset statistics (see Table 1), we will consider the variant of the **2PMB** routine with the positive class as the rare class. Recall the definition of the surrogate loss function for pAUC (5)

$$\ell_{\text{pAUC}}(\mathbf{w}) = \sum_{i:y_i>0} \sum_{j:y_j<0} \mathbb{T}_{\beta,t}^-(\mathbf{x}_j, \mathbf{w}) \cdot h(\mathbf{x}_i^\top \mathbf{w} - \mathbf{x}_j^\top \mathbf{w}).$$

We now rewrite this in a slightly different manner. Define, for any $i : y_i > 0$

$$\ell_{S_-}^+(x_i, \mathbf{w}) = \sum_{j:y_j<0} \mathbb{T}_{\beta,t}^-(\mathbf{x}_j, \mathbf{w}) \cdot h(\mathbf{x}_i^\top \mathbf{w} - \mathbf{x}_j^\top \mathbf{w}),$$

so that we can write $\ell_{\text{pAUC}}(\mathbf{w}) = \sum_{i:y_i>0} \ell_{S_-}^+(x_i, \mathbf{w})$. This shows that a subgradient to $\ell_{\text{pAUC}}(\mathbf{w})$ can be found by simply finding and summing up, subgradients for $\ell_{S_-}^+(x_i, \mathbf{w})$. For now, fix an i such that $y_i > 0$ and define $g(a) = h(\mathbf{x}_i^\top \mathbf{w} - a)$. Using the properties of the hinge loss function, it is clear that $g(a)$ is an increasing function of a . Since $\ell_{S_-}^+(x_i, \mathbf{w})$ is defined on the top ranked $\lceil \beta t_- \rceil$ negatives, we can, using the monotonicity of $g(\cdot)$, equivalently write it as follows. Let $\mathcal{Z}_\beta = \binom{S_-}{\lceil \beta t_- \rceil}$ be the set of all sets of negative points of negative training points of size $\lceil \beta t_- \rceil$. Then we can write

$$\ell_{S_-}^+(x_i, \mathbf{w}) = \max_{S \in \mathcal{Z}_\beta} \sum_{\mathbf{x}^- \in S} g(\mathbf{x}^{-\top} \mathbf{w})$$

Since the maximum in the above formulation is achieved at $S = \left\{ j : y_j < 0, \mathbb{T}_{\beta,t}^-(\mathbf{x}_j, \mathbf{w}) = 1 \right\}$, by Danskin's theorem (see, for example [26]), we get the following result: let $\mathbf{v}_{ij} \in \delta h(\mathbf{x}_i^\top \mathbf{w} - \mathbf{x}_j^\top \mathbf{w})$ be a subgradient to the hinge loss function, then for the following vector

$$\mathbf{v}_i := \sum_{j:y_j<0} \mathbb{T}_{\beta,t}^-(\mathbf{x}_j, \mathbf{w}) \cdot \mathbf{v}_{ij},$$

we have $\mathbf{v}_i \in \delta \ell_{S_-}^+(x_i, \mathbf{w})$ and consequently, for $\mathbf{v} := \sum_{i:y_i>0} \mathbf{v}_i$, we have $\mathbf{v} \in \delta \ell_{\text{pAUC}}(\mathbf{w})$. This gives us a straightforward way to implement **1PMB**: for each epoch, we take all the negatives in that epoch, filter out the top β fraction of them according to the scores assigned to them by the current iterate \mathbf{w}_e and then calculate the (sub)gradients between all the positives in that epoch and these filtered negatives. This takes at most $\mathcal{O}(s \log s)$ time per epoch.