A General Online Algorithm for Optimizing Complex Performance Metrics

Wojciech Kotłowski [1] **Marek Wydmuch** [1] **Erik Schultheis** [2] **Rohit Babbar** [2 3] **Krzysztof Dembczyński** [1 4]

Abstract

We consider sequential maximization of performance metrics that are general functions of a confusion matrix of a classifier (such as precision, F-measure, or G-mean). Such metrics are, in general, non-decomposable over individual instances, making their optimization very challenging. While they have been extensively studied under different frameworks in the batch setting, their analysis in the online learning regime is very limited, with only a few distinguished exceptions. In this paper, we introduce and analyze a general online algorithm that can be used in a straightforward way with a variety of complex performance metrics in binary, multi-class, and multi-label classification problems. The algorithm's update and prediction rules are appealingly simple and computationally efficient without the need to store any past data. We show the algorithm attains $\mathcal{O}(\frac{\ln n}{n})$ regret for concave and smooth metrics and verify the efficiency of the proposed algorithm in empirical studies.

1. Introduction

Many modern applications of machine learning involve optimization of complex performance metrics that, unlike misclassification error, do not decompose into expectation over instance-wise quantities. Examples of such measures include F-measure (Lewis, 1995), the area under the ROC curve (AUC) (Drummond & Holte, 2005), geometric (Drummond & Holte, 2005; Wang & Yao, 2012; Menon et al., 2013; Cao et al., 2019) and harmonic mean (Kennedy et al., 2009), and Matthews coefficient (Baldi et al., 2000).

Complex performance metrics have been studied in binary (Ye et al., 2012; Koyejo et al., 2014; Busa-Fekete et al., 2015; Dembczynski et al., 2017; Singh & Khim, 2022), multi-class (Narasimhan et al., 2015b; 2022), and multi-label classification (Waegeman et al., 2014; Koyejo et al., 2015; Kotłowski & Dembczyński, 2017) under different theoretical frameworks. The population utility (PU) framework focuses on building a model being close to the optimal one on the population level (Ye et al., 2012). The expected test utility (ETU) framework concerns the optimization of the expected utility computed over a given (known) test set (Lewis, 1995; Jansche, 2007; Ye et al., 2012). These two frameworks are different in the formulation of the objectives (more precisely, they differ in the order in which the expectation and the metric is computed), but they turn out to be asymptotically equivalent (Dembczynski et al., 2017). Still, most of the existing works concern a batch setting.

In this paper, we consider an online setting for optimization of complex performance metrics. In our framework, an algorithm observes a sequence of instances, one at a time, which are drawn *i.i.d.* from some distribution. At each iteration, after receiving the input vector, the algorithm makes a prediction and observes the true label (or vector of labels). After observing the entire sequence of data, the algorithm is evaluated by means of a performance metric computed from its empirical confusion matrix over the data. We evaluate an algorithm in our setting by means of a *regret*, which is the expected (over the data sequence) difference between the algorithm's performance and that of the optimal classifier. The goal is to design no-regret algorithms, that is, algorithms that guarantee regret converges to zero as the sequence length grows.

Contrary to most of the online learning work, we do not assume that predictions of the algorithm are obtained by means of some parametric function of the input vector (e.g., linear or generalized linear models). Instead, we assume that the algorithm has access to a *conditional probability estimator* (CPE), which returns an estimate of the true label conditional distribution at the input vector, which the algorithm will use to issue its prediction. Such a setup would be trivial for decomposable performance metrics, where the optimal decision is fully determined by the label conditional probabilities. For non-decomposable performance metrics, however, the optimal decision on different instances are no longer independent, yet the online setting requires the classi-

[*]Equal contribution [1]Poznan University of Technology, Poznan, Poland [2]Aalto University, Helsinki, Finland [3]University of Bath, Bath, UK [4]Yahoo Research, New York, United States. Correspondence to: Wojciech Kotłowski <wkotlowski@cs.put.poznan.pl>, Marek Wydmuch <mwydmuch@cs.put.poznan.pl>.

Proceedings of the 41st International Conference on Machine Learning, Vienna, Austria. PMLR 235, 2024. Copyright 2024 by the author(s).

fier to commit to a prediction as soon as it gets the instance, without any chance to revise decisions that turn out to be suboptimal in retrospect.

In this paper, we introduce and analyze a general online algorithm called Online Metric Maximization Algorithm (OMMA), which can be used in a straightforward way with a variety of complex performance metrics in binary, multi-class, and multi-label classification problems, with or without budgeted predictions (i.e., requiring exactly k predictions per instance). The algorithm's prediction rule is appealingly simple: at any given trial, the algorithm issues the prediction that maximizes the expected value (with respect to the current, not yet seen label, obtained by means of the CPE) of a *linearized* version of the metric. Thus, OMMA is effectively a cost-sensitive classification rule with costs determined by the gradient of the utility at the current confusion matrix, and computing its prediction boils down to thresholding or sorting linear functions of the conditional probabilities. The algorithm is computationally efficient and does not need to store any past data or predictions beyond the entries of the confusion matrix, update of which can be done online in a straightforward manner. The algorithm attains $\mathcal{O}(\frac{\ln n}{n})$ regret for concave and smooth metrics (up to the estimation error of its CPE), and $\mathcal{O}(\sqrt{\frac{\ln n}{n}})$ regret without smoothness.

We also verify the efficiency of the proposed algorithm in empirical studies, where we evaluate it on a wide range of multi-label benchmarks and performance measures. Our results show that the proposed algorithms achieve attractive performance in comparison to similar algorithms that can only be applied to a limited set of measures or require more computations and memory.

Our contribution can be summarized as follows:

- We formulate an online learning framework in which the classifier itself is not parametric but based on a potentially learnable label probability estimator,
- We provide a simple algorithm with a constant memory footprint that can optimize general non-decomposable performance metrics in an online manner across a wide range of learning tasks, including binary, multi-class, multi-label, and budgeted-at-k predictions,
- We prove regret bounds for this algorithm, showing that for concave and smooth metrics, we converge to the population-optimal at a rate of $\mathcal{O}(\frac{\ln n}{n})$.
- We show that in data sets with large variance, it can be beneficial for the decision process to *ignore* the available label feedback, and instead rely entirely on the biased, but much more stable, CPE instead.

Related work. Despite an extensive and long line of research on complex performance metrics in the batch setting, the analysis of these metrics in the online setting is quite limited, with only a few distinguished exceptions. Below we briefly describe the main results obtained so far.

The online problem has been tackled by optimizing a surrogate loss/reward function convex with respect to model parameters, which are updated in an online fashion. Kar et al. (2014) proposed an online gradient descent method for structural surrogates (Joachims, 2005). Narasimhan et al. (2015a) considered optimizing metrics being functions of true positive and true negative rate, which are then replaced by convex surrogate rewards. Such approaches may have high memory cost and do not converge to an optimal model for a given metric, but to a model optimal for the structural surrogate, and these two do not necessarily coincide (within the considered parametric class of functions).

The setup closest to ours has been considered in the context of optimizing the F-measure, where it is known that the optimal classifier is obtained by thresholding label conditional probabilities. In particular, Busa-Fekete et al. (2015) have introduced an online algorithm to simultaneously train a CPE model and tune the threshold. They have shown that the algorithm is asymptotically consistent, but the convergence rate has not been established. Interestingly, this algorithm can be seen as a special case of our general algorithm. In a follow-up paper, Zhang et al. (2018) have designed an algorithm in which the threshold is updated using stochastic gradient descent of a strongly convex function, which reflects a specific property of the optimal threshold. The algorithm has a convergence rate of $\mathcal{O}(\frac{\ln n}{\sqrt{n}})$.

For a wider class of functions, one can directly exploit the connection to cost-sensitive classification. Yan et al. (2017) has considered simultaneously learning multiple classifiers with different cost vectors together with an online algorithm that learns to select the best one. An alternative approach is to devise a method that finds the optimal costs in the space of confusion matrices, for example, by adapting the Frank-Wolfe algorithm previously used in batch learning for smooth and concave metrics (Narasimhan et al., 2015b; 2022; Schultheis et al., 2024). Such an approach can be then naïvely applied to the online setting by running it in exponentially growing time intervals to amortize the cost of rerunning the algorithm on the entire batch. We use such an approach as a baseline in our experimental studies.

Let us notice that our contribution does not coincide with any of the online/stochastic Frank-Wolfe methods known in the literature (Hazan & Kale, 2012; Reddi et al., 2016), as these concern minimization of an objective that is either a finite sum or an expectation, and thus linearly decomposes between the trials. In our case, the objective is a global, non-linear function of all past predictions and labels.

Online algorithms and regret minimization have been extensively studied in the online convex optimization (OCO) framework (Cesa-Bianchi & Lugosi, 2006; Hazan, 2016; Shalev-Shwartz, 2012). Our setup, however, substantially differs from OCO. On the one hand, we aim at optimizing a more difficult, non-linear function of the entire sequence of prediction, and we need stochastic (i.i.d.) assumption on the data sequence, as otherwise, no algorithm can guarantee vanishing regret (Appendix C). On the other hand, our algorithm is equipped with the CPE, which gives an approximate distribution of labels at the moment of prediction.

2. Problem setup

Let $\boldsymbol{x} \in \mathcal{X}$ be the input instance and $\boldsymbol{y} \in \mathcal{Y} \subseteq \{0,1\}^m$ the output label vector, jointly distributed according to $(\boldsymbol{x}, \boldsymbol{y}) \sim \mathbb{P}$. We let $\mathbb{P}_{\mathcal{X}}$ denote the marginal distribution $\boldsymbol{x}$, and $\boldsymbol{\eta}(\boldsymbol{x})$ the label conditional probability, $\boldsymbol{\eta}(\boldsymbol{x}) = \mathbb{E}_{\boldsymbol{y}|\boldsymbol{x}}[\boldsymbol{y}]$. In *multi-class classification*, the label vector $\boldsymbol{y}$ is a one-hot encoding of one out of m classes, $\mathcal{Y} = \{\boldsymbol{y} \in \{0,1\}^m \colon \sum_j y_j = 1\}$, while $\boldsymbol{\eta}(\boldsymbol{x}) = (\eta_1(\boldsymbol{x}), \ldots, \eta_m(\boldsymbol{x}))$ is the label conditional distribution with $\eta_j(\boldsymbol{x}) = \mathbb{P}(y_j = 1|\boldsymbol{x})$ denoting the conditional probability of class j. Whereas in *multi-label classification* $\boldsymbol{y} \in \{0,1\}^m$ denotes a vector of relevant labels and $\boldsymbol{\eta}(\boldsymbol{x}) = (\eta_1(\boldsymbol{x}), \ldots, \eta_m(\boldsymbol{x}))$ is vector of *label marginal conditional probabilities*, $\eta_j(\boldsymbol{x}) = \mathbb{P}(y_j = 1|\boldsymbol{x})$. Note that in the multi-label case, $\boldsymbol{\eta}(\boldsymbol{x})$ *is not* the conditional distribution $\mathbb{P}(\boldsymbol{y}|\boldsymbol{x})$ (which is a *joint* distribution over $\{0,1\}^m$), but in the context of this paper, it will always suffice to only operate on its marginals; we will thus sometimes use $\boldsymbol{y} \sim \boldsymbol{\eta}(\boldsymbol{x})$ to denote a multi-label vector drawn from $\mathbb{P}(\boldsymbol{y}|\boldsymbol{x})$ with marginals $\boldsymbol{\eta}(\boldsymbol{x})$ if no other property of the distribution matters, and we will call $\boldsymbol{\eta}$ the conditional distribution to simplify the presentation.

We consider an online (stochastic) setting in which the algorithm observes a sequence of instances $(\boldsymbol{x}_1, \boldsymbol{y}_1), (\boldsymbol{x}_2, \boldsymbol{y}_2), \ldots, (\boldsymbol{x}_n, \boldsymbol{y}_n)$, drawn *i.i.d.* from $\mathbb{P}$. At each iteration $t = 1, 2, \ldots, n$, after receiving input $\boldsymbol{x}_t$ drawn from $\mathbb{P}_{\mathcal{X}}$, the algorithm makes a prediction $\widehat{\boldsymbol{y}}_t$, and next observes the true output $\boldsymbol{y}_t$ drawn from $\mathbb{P}(\boldsymbol{y}|\boldsymbol{x}_t)$. We assume that the algorithm has access to a *conditional probability estimator* (CPE) $\widehat{\boldsymbol{\eta}}_t$, which at trial t returns an estimate $\widehat{\boldsymbol{\eta}}_t(\boldsymbol{x}_t)$ of the true conditional distribution $\boldsymbol{\eta}(\boldsymbol{x}_t)$, which the algorithm will use for issuing its prediction $\widehat{\boldsymbol{y}}_t$. We put the time subscript in $\widehat{\boldsymbol{\eta}}_t$ as we allow the CPE to change over time (e.g., it can be produced by an external online learner run on the same sequence of instances); we assume, however, that $\widehat{\boldsymbol{\eta}}_t$ can only depend on the observed data $\boldsymbol{x}_t$, $(\boldsymbol{x}_1, \boldsymbol{y}_1), \ldots, (\boldsymbol{x}_{t-1}, \boldsymbol{y}_{t-1})$ (i.e., does not depend on $\boldsymbol{y}_t$ and all the future instance). The CPE could be, for instance, a neural network or boosted decision trees trained (possibly online) by minimizing some proper scoring loss. In this work, we do not focus on the way the CPE is learned, and all our bounds on the performance of the algorithm will depend on the *estimation error* of $\widehat{\boldsymbol{\eta}}_t$ with respect to true conditional distribution $\boldsymbol{\eta}$. We outline the online protocol in Figure 1.

for $t = 1, \ldots, n$ **do**
  Observe input instance $\boldsymbol{x}_t$ drawn from $\mathbb{P}_{\mathcal{X}}$
  Receive conditional probability estimate $\widehat{\boldsymbol{\eta}}_t(\boldsymbol{x}_t)$
  Predict label $\widehat{\boldsymbol{y}}_t$ based on $\widehat{\boldsymbol{\eta}}_t(\boldsymbol{x}_t)$
  Receive true label $\boldsymbol{y}_t$ drawn from $\mathbb{P}(\cdot|\boldsymbol{x}_t)$
Evaluate based on $\psi(\boldsymbol{C}(\boldsymbol{y}^n, \widehat{\boldsymbol{y}}^n))$

Figure 1. The online protocol.

Given a sequence of labels[1] $\boldsymbol{y}^n = (\boldsymbol{y}_1, \ldots, \boldsymbol{y}_n)$ and a sequence of algorithm's predictions $\widehat{\boldsymbol{y}}^n = (\widehat{\boldsymbol{y}}_1, \ldots, \widehat{\boldsymbol{y}}_n)$, we let $\boldsymbol{C}(\boldsymbol{y}^n, \widehat{\boldsymbol{y}}^n)$ denote the (empirical) *confusion matrix*. In multi-class classification, it is an $m \times m$ matrix $\boldsymbol{C}$ with entries defined as $C_{j\ell} = \frac{1}{n}\sum_{t=1}^n y_{tj}\widehat{y}_{t\ell}$, where y_{tj} is the j-th entry of vector $\boldsymbol{y}_t$. Thus, $C_{j\ell}$ contains the fraction of times an instance from class j was predicted as being in class ℓ. Note that for binary classification ($m = 2$), we get the usual true-positive (C_{11}), false-negative (C_{10}), false-positive (C_{01}), true-negative (C_{00}) entries, where we index the entries from 0 to stick to a more convenient notation. In turn, in multi-label classification, the confusion matrix is an $m \times 2 \times 2$ tensor $\boldsymbol{C}$, where $C_{juv} = \frac{1}{n}\sum_{t=1}^n \mathbf{1}\{y_{tj} = u\}\mathbf{1}\{\widehat{y}_{tj} = v\}$,[2] where $j \in \{1, \ldots, m\}$ and $u, v \in \{0, 1\}$. In other words, $\boldsymbol{C}$ is a sequence of m binary confusion matrices, separately for each label. In a given problem, we let $\mathcal{C}$ denote the set of all achievable confusion matrices, that is the set of confusion matrices that can be formed from allowed label and prediction vectors $\boldsymbol{y}^t$, $\widehat{\boldsymbol{y}}^t$ for any t. Note that, independent of the considered problem, the confusion matrix can always be written as $\boldsymbol{C}(\boldsymbol{y}^t, \widehat{\boldsymbol{y}}^t) = \frac{1}{t}\sum_{i=1}^t \boldsymbol{C}(\boldsymbol{y}_i, \widehat{\boldsymbol{y}}_i)$ (with $\boldsymbol{C}(\boldsymbol{y}_i, \widehat{\boldsymbol{y}}_i)$ being confusion matrices computed out of a single label $\boldsymbol{y}_i$ and prediction $\widehat{\boldsymbol{y}}_i$), which gives a simple online update:

$$\boldsymbol{C}(\boldsymbol{y}^t, \widehat{\boldsymbol{y}}^t) = \frac{t-1}{t}\boldsymbol{C}(\boldsymbol{y}^{t-1}, \widehat{\boldsymbol{y}}^{t-1}) + \frac{1}{t}\boldsymbol{C}(\boldsymbol{y}_t, \widehat{\boldsymbol{y}}_t)\,. \quad (1)$$

In this work, we focus on online maximization of performance metrics that do not decompose into a sum over instances but are general functions of the confusion matrix of the algorithm. Specifically, after observing the entire sequence of n instances, the algorithm is evaluated by means of a *utility metric* $\psi = \psi(\boldsymbol{C}(\boldsymbol{y}^n, \widehat{\boldsymbol{y}}^n))$. For binary classification, examples of such measures include the F-measure, geometric and harmonic mean, area under the ROC curve, recall, precision, etc. We present

[1] Generally, by $\boldsymbol{v}^t$ we denote a sequence $(\boldsymbol{v}_1, \ldots, \boldsymbol{v}_t)$; $f(\boldsymbol{v}^t)$ denotes $(f(\boldsymbol{v}_1), \ldots, f(\boldsymbol{v}_t))$.

[2] $\mathbf{1}(S)$ is the indicator function equal one when S holds.

Table 1. Examples of binary and multi-class confusion matrix measures. Binary measures are expressed in terms of true-positives ($\text{tp} = C_{11}$), false-negatives ($\text{fn} = C_{10}$), false-positives ($\text{fp} = C_{01}$), and true-negatives ($\text{tn} = C_{00}$).

Metric	$\psi(\boldsymbol{C}^{2\times 2})$	$\psi(\boldsymbol{C}^{m\times m})$	Metric	$\psi(\boldsymbol{C}^{2\times 2})$	$\psi(\boldsymbol{C}^{m\times m})$
Accuracy	$\text{tp}+\text{tn}$	$\sum_{i=1}^m C_{ii}$	G-mean	$\sqrt{\frac{\text{tp}\cdot\text{tn}}{(\text{tp}+\text{fn})(\text{tn}+\text{fp})}}$	$\left(\prod_{j=1}^m \frac{C_{jj}}{\sum_{i=1}^m C_{ji}}\right)^{1/m}$
Balanced Acc.	$\frac{\text{tp}}{2(\text{tp}+\text{fn})}+\frac{\text{tn}}{2(\text{tn}+\text{fp})}$	$\sum_{i=1}^m \frac{C_{ii}}{m\sum_{j=1}^m C_{ij}}$	H-mean	$2\left(\frac{\text{tp}+\text{fn}}{\text{tp}}+\frac{\text{tn}+\text{fp}}{\text{tn}}\right)^{-1}$	$m\left(\sum_{j=1}^m \frac{\sum_{i=1}^m C_{ji}}{C_{jj}}\right)^{-1}$
Recall	$\frac{\text{tp}}{\text{tp}+\text{fn}}$	micro- or macro-avg.	Q-mean	$1-\sqrt{\frac{1}{2}\left(\left(\frac{\text{fn}}{\text{tp}+\text{fn}}\right)^2+\left(\frac{\text{fp}}{\text{tn}+\text{fp}}\right)^2\right)}$	$1-\sqrt{\frac{1}{m}\sum_{j=1}^m\left(1-\frac{C_{jj}}{\sum_{i=1}^m C_{ji}}\right)^2}$
Precision	$\frac{\text{tp}}{\text{tp}+\text{fp}}$	micro- or macro-avg.	Jaccard	$\frac{\text{tp}}{\text{tp}+\text{fp}+\text{fn}}$	micro- or macro-avg.
F_β-measure	$\frac{(1+\beta^2)\text{tp}}{(1+\beta^2)\text{tp}+\beta^2\text{fn}+\text{fp}}$	micro- or macro-avg.	Matthews coeff.	$\frac{\text{tp}\cdot\text{tn}-\text{fp}\cdot\text{fn}}{\sqrt{(\text{tp}+\text{fp})(\text{tp}+\text{fn})(\text{tn}+\text{fp})(\text{tn}+\text{fn})}}$	micro- or macro-avg.

their definitions in Table 1. Most of these metrics have multi-class (some also presented in Table 1) as well as multi-label extensions in the micro- and macro-averaged variants. By rewriting the multi-class $\boldsymbol{C}^{m\times m}$ confusion matrix into a multi-label confusion tensor $\boldsymbol{C}^{m\times 2\times 2}$,

$$\boldsymbol{C}^{m\times 2\times 2}(\boldsymbol{C}^{m\times m}) := \left[\begin{pmatrix} C_{jj} & \sum_{i\neq j} C_{ij} \\ \sum_{i\neq j} C_{ji} & \sum_{i\neq j,\ell\neq j} C_{i\ell}\end{pmatrix}\right]_{j=1}^{m},$$

we can define the micro- and macro-averaged metrics for both multi-class and multi-label classification as:

$$\text{Micro-}\psi(\boldsymbol{C}^{m\times 2\times 2}) := \psi\left(\frac{1}{m}\sum_{j=1}^m \boldsymbol{C}_j\right),$$

$$\text{Macro-}\psi(\boldsymbol{C}^{m\times 2\times 2}) := \frac{1}{m}\sum_{j=1}^m \psi\left(\boldsymbol{C}_j\right).$$

All the metrics presented above can be considered under the standard as well as budgeted at k variant (Schultheis et al., 2023; 2024), where the classifier is required to predict exactly k classes per instance ($\widehat{\boldsymbol{y}} \in \widehat{\mathcal{Y}}_k = \{\widehat{\boldsymbol{y}} \in \{0,1\}^m : \sum_j \widehat{y}_j = k\}$). The main difficulty in maximizing ψ comes from the fact that the algorithm is only evaluated at the end of the sequence by a (possibly complex) function of all the labels and predictions, while each prediction $\widehat{\boldsymbol{y}}_t$ must be made immediately upon observing $\boldsymbol{x}_t$, and cannot be changed in the future.

Regret. Since without imposing assumptions on the data distribution $\mathbb{P}$ one cannot meaningfully bound the ψ-accuracy of the algorithm in the absolute terms, the goal is to compare the algorithm's performance *relative* to that of the optimal predictor. Note, however, that defining the optimal predictions $\boldsymbol{y}^{\star n}$ simply as those maximizing the value of the utility, $\boldsymbol{y}^{\star n} = \operatorname{argmax}_{\widehat{\boldsymbol{y}}^n} \psi(\boldsymbol{C}(\boldsymbol{y}^n, \widehat{\boldsymbol{y}}^n))$, leads (for essentially all reasonable utilities) to a trivial solution $\boldsymbol{y}^{\star n} = \boldsymbol{y}^n$, which gives the maximum possible value ψ, and is thus not achievable by any algorithm, no matter how large n is. Thus, we proceed differently, defining a *classifier* $\boldsymbol{h}\colon \mathcal{X} \to \mathcal{Y}$ to be a function from the inputs to the outputs, and the optimal predictions $\boldsymbol{y}^{\star n} = \boldsymbol{h}^\star(\boldsymbol{x}^n)$ to be those generated by a classifier $\boldsymbol{h}^\star$ which maximizes *the expected (with respect to the data sequence) value of the utility*:[3]

$$\boldsymbol{h}^\star := \operatorname*{argmax}_{\boldsymbol{h}} \mathbb{E}_{(\boldsymbol{x},\boldsymbol{y})^n}\left[\psi(\boldsymbol{C}(\boldsymbol{y}^n, \boldsymbol{h}(\boldsymbol{x}^n)))\right], \qquad (2)$$

where the maximization is with respect to all (measurable) classifiers. We define the (expected) *regret* of the algorithm as the difference between its expected performance in terms of ψ, and the performance of the optimal classifier $\boldsymbol{h}^\star$:

$$R_n := \mathbb{E}\left[\psi(\boldsymbol{C}(\boldsymbol{y}^n, \boldsymbol{h}^\star(\boldsymbol{x}^n)))\right] - \mathbb{E}\left[\psi(\boldsymbol{C}(\boldsymbol{y}^n, \widehat{\boldsymbol{y}}^n))\right].$$

Our goal is to design no-regret algorithms, that is, algorithms that guarantee $R_n \to 0$ as $n \to \infty$. In the next section, we propose a computationally- and memory-efficient algorithm, which guarantees (under certain assumptions on the utility) $R_n = \mathcal{O}\left(\frac{\ln n}{n}\right)$ up to the estimation error of CPE.

Note that when $n \to \infty$, $\boldsymbol{h}^\star$ coincides with the ψ-optimal population-level classifier. Approximating $\boldsymbol{h}^\star$ in the batch setting is a task which was tackled in a series of past works spanning binary classification (Ye et al., 2012; Menon et al., 2013; Koyejo et al., 2014; Narasimhan et al., 2014), multi-class classification (Narasimhan et al., 2015b; 2022), and multi-label classification (Waegeman et al., 2014; Koyejo et al., 2015; Kotłowski & Dembczyński, 2017; Schultheis et al., 2024), and boils down to, under certain mild assumptions on the data distribution, finding the optimal threshold on the conditional probability $\boldsymbol{\eta}(\boldsymbol{x})$ (for binary classification), or finding the optimal ensemble of cost-sensitive classifiers with an iterative convex optimization algorithm, such as the Frank-Wolfe method or gradient descent (for multi-class and multi-label classification).[4] In this light, the results

[3] Recall that $\boldsymbol{h}(\boldsymbol{x}^n)$ stands for $(\boldsymbol{h}(\boldsymbol{x}_1), \ldots, \boldsymbol{h}(\boldsymbol{x}_n))$.

[4] In the past papers, the optimal classifier is allowed to be

Algorithm 1 Online Measure Maximization Algorithm

Initialization: confusion matrix $\boldsymbol{C}_0$
for $t = 1, \ldots, n$ **do**
  Receive input $\boldsymbol{x}_t$ and probability estimate $\widehat{\boldsymbol{\eta}}_t(\boldsymbol{x}_t)$
  Predict $\widehat{\boldsymbol{y}}_t = \operatorname{argmax}_{\widehat{\boldsymbol{y}}} \nabla\psi(\boldsymbol{C}_{t-1}) \cdot \left(\mathbb{E}_{\boldsymbol{y}_t \sim \widehat{\boldsymbol{\eta}}_t(\boldsymbol{x}_t)}[\boldsymbol{C}(\boldsymbol{y}_t, \widehat{\boldsymbol{y}})]\right)$
  Receive label $\boldsymbol{y}_t$ and update $\boldsymbol{C}_t = \frac{t-1}{t}\boldsymbol{C}_{t-1} + \frac{1}{t}\boldsymbol{C}(\boldsymbol{y}_t, \widehat{\boldsymbol{y}}_t)$

of our paper can be interpreted as proposing a simple and efficient iterative algorithm, which does only a single pass over the data while achieving performance that is very close to that of the aforementioned batch optimization methods.

Remark 2.1 (Adversarial sequence of inputs). One might consider a setup in which the input instances $\boldsymbol{x}^n$ are not drawn i.i.d. from a fixed distribution, but rather form an arbitrary (possibly even adversarial) sequence, and the goal is to analyze the worst case with respect to $\boldsymbol{x}^n$ (note that the labels remain stochastic, being drawn from $\boldsymbol{\eta}(\boldsymbol{x})$). Unfortunately, this setting turns out to be too difficult in the worst case: in Appendix C, we show that there exists a conditional probability function $\boldsymbol{\eta}$ (which we even make known to the algorithm), and a sequence of inputs $\boldsymbol{x}^n$, on which no algorithm can have a vanishing regret. Thus, a stochastic data generation mechanism seems crucial for the existence of a no-regret learner.

3. The algorithm

In this section, we introduce our algorithm, called *Online Metric Maximization Algorithm (OMMA)*. To define the algorithm, we need to make the following assumption:

Assumption 3.1. The utility $\psi(\boldsymbol{C})$ is differentiable in $\boldsymbol{C}$.

This property is shared by all performance metrics commonly used in practice.[5] Let $\nabla\psi(\boldsymbol{C})$ be the matrix of derivatives of ψ with respect to $\boldsymbol{C}$, which is of the same shape as $\boldsymbol{C}$ ($m \times m$ for multi-class, and $m \times 2 \times 2$ for multi-label case). To simplify notation, we let $\boldsymbol{C}_t := \boldsymbol{C}(\boldsymbol{y}^t, \widehat{\boldsymbol{y}}^t)$ denote the confusion matrix of the algorithm after t trials. At trial t, given past labels $\boldsymbol{y}^{t-1}$, past predictions $\widehat{\boldsymbol{y}}^{t-1}$, new input $\boldsymbol{x}_t$ and the conditional probability estimate $\widehat{\boldsymbol{\eta}}_t(\boldsymbol{x}_t)$, the OMMA algorithm predicts $\widehat{\boldsymbol{y}}_t$ which maximizes the expected (with respect to $\widehat{\boldsymbol{\eta}}_t$) *linearized* (at $\boldsymbol{C}_{t-1}$) version of the utility:[6]

$$\widehat{\boldsymbol{y}}_t = \operatorname*{argmax}_{\widehat{\boldsymbol{y}}} \nabla\psi(\boldsymbol{C}_{t-1}) \cdot \left(\mathbb{E}_{\boldsymbol{y}_t \sim \widehat{\boldsymbol{\eta}}_t(\boldsymbol{x}_t)}\left[\boldsymbol{C}(\boldsymbol{y}_t, \widehat{\boldsymbol{y}})\right]\right) , \quad (3)$$

where the maximization is over the allowed predictions $\widehat{\boldsymbol{y}}$ (e.g., one-hot vectors in multi-class classification), and the operator "$\cdot$" denotes the matrix/tensor dot product (which is a standard vector dot product between vectorized matrix/tensor arguments). Note that the last element in (3) is a confusion matrix computed from *only a single observation $\boldsymbol{y}_t$ and a single prediction $\widehat{\boldsymbol{y}}$*. The algorithm does not need to store any past data or predictions beyond the entries of the confusion matrix, as the update of matrix can be done online using (1). Due to the linearity of (3), OMMA is effectively a *cost-sensitive classification* rule with costs determined by the gradient of the utility at the current confusion matrix. In what follows, we will show that formula (3) leads to a very simple rule, boiling down to thresholding or sorting linear functions of $\widehat{\boldsymbol{\eta}}_t(\boldsymbol{x}_t)$.

Motivation. OMMA can be derived from considering a greedy method, which, given past labels $\boldsymbol{y}^{t-1}$ and predictions $\widehat{\boldsymbol{y}}^{t-1}$, chooses its prediction $\widehat{\boldsymbol{y}}_t$ in order to maximize the expected utility *including* iteration t (where the expectation is with respect to label $\boldsymbol{y}_t$, which is unknown at the moment of making prediction), that is, to maximize $F(\widehat{\boldsymbol{y}}_t) = \mathbb{E}_{\boldsymbol{y}_t \sim \widehat{\boldsymbol{\eta}}_t(\boldsymbol{x}_t)}\left[\psi(\boldsymbol{C}(\boldsymbol{y}^t, \widehat{\boldsymbol{y}}^t))\right]$. By using a Taylor-expansion of $F(\widehat{\boldsymbol{y}}_t)$ with respect to $\boldsymbol{y}_t$, it turns out that maximizing $F(\widehat{\boldsymbol{y}}_t)$ and maximizing (3) are equivalent up to $\mathcal{O}(1/t^2)$ (see Appendix A for details). At the same time, our algorithmic update (3) is based on a linear objective and, therefore, more straightforward to calculate, as well as easier to analyze. In the experiment section, we show that both updates indeed behave very similarly to each other, being essentially indistinguishable for larger values of n.

Example: binary classification. To simplify the presentation, we switch from one-hot vector notation to scalars and denote by $y_t \in \{0, 1\}$ a label at time t ($y_t = 0$ corresponds to $\boldsymbol{y}_t = (1, 0)$, and $y_t = 1$ to $\boldsymbol{y}_t = (0, 1)$); similarly $\widehat{y}_t \in \{0, 1\}$ denotes the prediction and $\eta(\boldsymbol{x}_t) = P(y_t = 1 | \boldsymbol{x}_t)$. The confusion matrix $\boldsymbol{C}_{t-1} = \boldsymbol{C}(\boldsymbol{y}^{t-1}, \widehat{\boldsymbol{y}}^{t-1})$ consists of four entries

$$C_{t-1,j\ell} = \sum_{i \le t-1} \mathbf{1}\{y_i = j\}\mathbf{1}\{\widehat{y}_i = \ell\}, \quad j, \ell \in \{0, 1\} . \quad (4)$$

Let us abbreviate $\frac{\partial}{\partial C_{j\ell}}\psi(\boldsymbol{C}_{t-1})$ as $\nabla_{j\ell}$, and $\widehat{\eta}_t(\boldsymbol{x}_t)$ as $\widehat{\eta}_t$. Since

$$\mathbb{E}_{y_t \sim \widehat{\eta}_t}[\boldsymbol{C}(y_t, \widehat{y})] = \begin{bmatrix} (1-\widehat{\eta}_t)(1-\widehat{y}) & (1-\widehat{\eta}_t)\widehat{y} \\ \widehat{\eta}_t(1-\widehat{y}) & \widehat{\eta}_t\widehat{y} \end{bmatrix} ,$$

equation (3) boils down to maximizing a cost-sensitive classification accuracy:

$$\nabla_{00}(1-\widehat{\eta}_t)(1-\widehat{y}) + \nabla_{01}(1-\widehat{\eta}_t)\widehat{y} + \nabla_{10}\widehat{\eta}_t(1-\widehat{y}) + \nabla_{11}\widehat{\eta}_t\widehat{y}, \quad (5)$$

which can also be rewritten, up to terms independent of $\widehat{y}$, as $\widehat{y}(\alpha\widehat{\eta}_t - \beta)$, with

$$\alpha = \nabla_{11} + \nabla_{00} - \nabla_{01} - \nabla_{10}, \quad \beta = \nabla_{00} - \nabla_{01},$$

randomized, that is to take values in the convex hull of $\mathcal{Y}$. All of our analysis effortlessly extends to such a setup, but we do not consider it for the clarity of the presentation. Moreover, our algorithm never needs to randomize its predictions.

[5] In case of non-differentiable and concave metrics, one can still use a supergradient in place of gradient.

[6] The algorithm is well-defined, as $\boldsymbol{C}(\boldsymbol{y}_t, \widehat{\boldsymbol{y}})$ is *linear* in labels and thus its expectation only depends on marginals $\widehat{\boldsymbol{\eta}}_t(\boldsymbol{x}_t)$.

a cost-sensitive prediction rule (Elkan, 2001; Natarajan et al., 2018). Since all utilities used in practice are non-decreasing with true positives and true negatives ($\nabla_{11}, \nabla_{00} \geq 0$), and non-increasing with false negatives and false positives ($\nabla_{01}, \nabla_{10} \leq 0$), we get $a \geq 0$, and thus maximizing $\widehat{y}(\alpha\widehat{\eta}_t - \beta)$ boils down to choosing $\widehat{y}_t = 1$ whenever the conditional probability $\widehat{\eta}_t(\boldsymbol{x}_t)$ exceeds a threshold β/α. It is well-known that, under mild assumptions on the input distribution $\mathbb{P}_{\mathcal{X}}$ and the metric ψ, thresholding $\eta(\boldsymbol{x})$ is the optimal classification rule for maximizing the population-level version of ψ (Koyejo et al., 2014; Narasimhan et al., 2014). Thus, OMMA mimics the optimal classification rule, with a threshold computed based on the empirical data.

Example: multi-label classification (with a budget). Here $\boldsymbol{y}_t$ and $\widehat{\boldsymbol{y}}_t$ are label and prediction vectors of length m, and $\boldsymbol{\eta}(\boldsymbol{x}_t)$ is the vector of marginal label probabilities. $\boldsymbol{C}_{t-1}$ is a sequence of binary confusion matrices $\boldsymbol{C}^j_{t-1}$, $j = 1, \ldots, m$. Abbreviating $\frac{\partial}{\partial C^j_{uv}}\psi(\boldsymbol{C}_{t-1})$ as ∇^j_{uv}, and $\widehat{\boldsymbol{\eta}}_t(\boldsymbol{x}_t)$ as $\widehat{\boldsymbol{\eta}}_t$ allows us to write down the objective (3) as a direct extension of the one from binary classification (5), summed over m labels:

$$\sum_{j=1}^{m} \widehat{y}_j(\alpha_j\widehat{\eta}_{tj} - \beta_j) \tag{6}$$

with $\alpha_j = \nabla^j_{11} + \nabla^j_{00} - \nabla^j_{01} - \nabla^j_{10}$, $\beta_j = \nabla^j_{00} - \nabla^j_{01}$. Maximizing (6) amounts to setting $\widehat{y}_{tj} = 1$ whenever $\alpha_j\widehat{\eta}_{tj} - \beta_j \geq 0$, or $\widehat{\eta}_{tj} \geq \beta_j/\alpha_j$. In the budgeted-at-$k$ variant, where the algorithm must predict *exactly* k labels each time, the prediction amounts to sorting the labels by means of $\alpha_j\widehat{\eta}_{tj} - \beta_j$ in descending order, and setting the top k of them to 1. Interestingly, this is closely related to the optimal population-level prediction rule for the budgeted setting (Schultheis et al., 2024).

4. Theoretical analysis

In order to prove a bound on the regret of OMMA algorithm, we need to make additional assumptions on the utility ψ.

Assumption 4.1. $\psi(\boldsymbol{C})$ is differentiable, concave, L-Lipschitz, and M-smooth in $\boldsymbol{C}$, that is for any $\boldsymbol{C}_1, \boldsymbol{C}_2 \in \mathcal{C}$, $\psi(\boldsymbol{C}_1) \leq \psi(\boldsymbol{C}_2) + \nabla\psi(\boldsymbol{C}_2)^\top(\boldsymbol{C}_1 - \boldsymbol{C}_2)$, $|\psi(\boldsymbol{C}_1) - \psi(\boldsymbol{C}_2)\| \leq L\|\boldsymbol{C}_1 - \boldsymbol{C}_2\|$, and $\|\nabla\psi(\boldsymbol{C}_1) - \nabla\psi(\boldsymbol{C}_2)\| \leq M\|\boldsymbol{C}_1 - \boldsymbol{C}_2\|$, where $\|\cdot\|$ denotes the entrywise L_2-norm.

Remark 4.2. Our algorithm generally requires smoothness of the objective to converge. The assumption can be waived by running it on a *smoothed* version of the metric. This, however, comes at the price of a slower convergence rate $\mathcal{O}(\sqrt{\ln n/n})$. As non-smooth metrics are not commonly employed in machine learning, we relegate the discussion to Appendix B.2

Theorem 4.3. *Let Assumption 4.1 hold. Then the OMMA algorithm has its regret bounded by:*

$$R_n \leq \frac{Ma(1+\ln n)}{n} + \frac{2bL}{n}\sum_{t=1}^{n}\mathbb{E}\left[\|\boldsymbol{\eta}(\boldsymbol{x}_t) - \widehat{\boldsymbol{\eta}}_t(\boldsymbol{x}_t)\|\right]$$

where $a = 1, b = 1$ for multi-class, and $a = m, b = \sqrt{2}$ for multi-label classification.

The proof is given in Appendix B.1. The bound in Theorem 4.3 consists of two parts: (1) the first term of order $\mathcal{O}(\frac{\ln n}{n})$ can be interpreted as the regret of OMMA had it been equipped with the true conditional probability $\boldsymbol{\eta}$ (that is, the estimator is exact); (2) the estimation error of $\widehat{\boldsymbol{\eta}}_t$ averaged over trials $t = 1, \ldots, n$. If the estimation error converges to zero with $n \to \infty$, OMMA becomes a no-regret learning algorithm.

Remark 4.4. In the proof of Theorem 4.3, the Lipschitzness and the smoothness properties are invoked along the parameter path of the algorithm, $\{\boldsymbol{C}_t\}_{t=1}^n$, in order to control the progress in optimizing the utility. These properties might not necessarily hold globally (for every confusion matrix) for utilities given in Table 1, for instance when the observed labels lead to severe class imbalance. However, adding a small constant to the denominator in the definition of a utility stabilizes its values and ensures its global Lipschitzness and smoothness. This is also the approach we take in the experiments to keep our algorithm stable over the initial part of the data sequence.

An alternative approach is to use the fact that most utilities in Table 1 are Lipschitz and smooth when label frequencies in the confusion matrix are bounded away from zero. Taking into account that these properties are applied along the path of the algorithm, and using concentration inequalities on the label frequencies, we show in Appendix E that as long as the probabilities of labels $\mathbb{P}(y_j = 1)$ are bounded away from zero, the regret of OMMA converges at a rate of $\mathcal{O}(\sqrt{\ln n/n})$ with high probability. Motivated by this fact, we also propose to use a regularization technique that, additionally to using a small constant in the denominator of some metrics, also adds the small value λ to the initial entries of the confusion matrix. With the updates of the confusion matrix, this value diminishes, being $\frac{\lambda}{t}$ at iteration t of the algorithm. This simple technique turns out to be helpful for some metrics.

5. Alternative variants of the algorithm

Internal semi-empirical confusion matrix. In the experiment, we also use an alternative version of the OMMA algorithm called OMMA($\widehat{\eta}$), outlined in Algorithm 2. Note that in the algorithm's description we use $\boldsymbol{C}(\widehat{\boldsymbol{\eta}}_t, \widehat{\boldsymbol{y}}_t) = \mathbb{E}_{\boldsymbol{y}_t \sim \widehat{\boldsymbol{\eta}}_t}\boldsymbol{C}(\boldsymbol{y}_t, \widehat{\boldsymbol{y}}_t)$ which follows from the fact that the confusion matrix is linear in labels and that $\mathbb{E}_{\boldsymbol{y}_t \sim \widehat{\boldsymbol{\eta}}_t}[\boldsymbol{y}_t] = \widehat{\boldsymbol{\eta}}_t$.

Algorithm 2 OMMA($\widehat{\eta}$)

Initialization: confusion matrix $\boldsymbol{C}_0$
for $t = 1, \ldots, n$ **do**
  Receive input $\boldsymbol{x}_t$ and probability estimate $\widehat{\boldsymbol{\eta}}_t := \widehat{\boldsymbol{\eta}}_t(\boldsymbol{x}_t)$
  Predict $\widehat{\boldsymbol{y}}_t = \operatorname{argmax}_{\widehat{\boldsymbol{y}}} \nabla\psi(\boldsymbol{C}_{t-1}) \cdot \boldsymbol{C}(\widehat{\boldsymbol{\eta}}_t, \widehat{\boldsymbol{y}})$
  Update $\boldsymbol{C}_t = \frac{t-1}{t}\boldsymbol{C}_{t-1} + \frac{1}{t}\boldsymbol{C}(\widehat{\boldsymbol{\eta}}_t, \widehat{\boldsymbol{y}}_t)$

The only difference between OMMA and its modification is that OMMA($\widehat{\eta}$) updates its running confusion matrix $\boldsymbol{C}_t$ by means of $\widehat{\boldsymbol{\eta}}_t$ instead of the true label $\boldsymbol{y}_t$. In fact, OMMA($\widehat{\eta}$) does not use the true labels at all, fully trusting its CPE. We remark that $\boldsymbol{C}_t$ does *not* correspond to any empirical confusion matrix (as it is not based on the labels), so it should rather be treated as an internal parameter of OMMA($\widehat{\eta}$). In Appendix D we show that if the algorithm's probability estimator is *exact*, that is $\widehat{\boldsymbol{\eta}}_t \equiv \boldsymbol{\eta}$ for all t, OMMA($\widehat{\eta}$) achieves (under the same assumptions as before) a regret bound $R_n \leq \frac{Ma(2+\ln n)}{n}$, which is very similar to that of the original OMMA algorithm in Theorem 4.3. It is, however, unclear whether the modified algorithm converges for a non-exact CPE. Still, it turns out that OMMA($\widehat{\eta}$) performs surprisingly well in the experiments.

Sparse variant for a large number of labels. The introduced OMMA algorithm requires calculating the derivative with respect to every entry of the confusion matrix at each step, resulting in the total complexity of $\mathcal{O}(nm)$ over the entire sequence of data (assuming constant time for the gradient calculation of a single entry). Even this can be fairly expensive in case of a large number of labels (e.g., in recommendation or extreme classification); in these multi-label problems, however, the number of positive labels per sample is much smaller than the total number of label, $\|\boldsymbol{y}\|_1 \ll m$, and often most of the conditional probabilities in $\boldsymbol{\eta}(\boldsymbol{x})$ are 0 or very close to 0. Many recommenders and extreme classifiers are naturally designed to predict only top-k' entries with the highest values of $\eta(\boldsymbol{x})$. We can then leverage the sparsity of top-k' assuming $\eta(\boldsymbol{x}) = 0$ for the rest of the labels and calculate the gradient with respect to entries of the confusion matrix that correspond to non-zero values of $\boldsymbol{\eta}(\boldsymbol{x})$, resulting in a total complexity of $\mathcal{O}(nk')$ on the entire data sequence. With reasonably selected k', according to Theorem 4.3, we should only slightly increase the regret.

6. Empirical study

To demonstrate the practicality and generality of the introduced OMMA algorithm, we test it on a wide range of multi-label and multi-class benchmark datasets that differ substantially in the number of labels, ranging from tens to a few thousands, and in the imbalance of the label distribution [7]. For multi-class experiments, we use News20 (Lang, 1995), Ledgar-LexGlue (Chalkidis et al., 2022) with tf-idf features, Caltech-256 (Griffin et al., 2007) with features obtained using VGG16 (Simonyan & Zisserman, 2014) trained on ImageNet, and for multi-label experiments, we use YouTube, Flickr (Tang & Liu, 2009) with DeepWalk features (Perozzi et al., 2014), Eurlex-LexGlue (Chalkidis et al., 2021), Mediamill (Snoek et al., 2006), RCV1X (Lewis et al., 2004), and AmazonCat (McAuley & Leskovec, 2013; Bhatia et al., 2016) with tf-idf features. We conduct two types of experiments:

1. With fixed conditional probability estimator (CPE) – we train the CPE using L_2-regularized logistic loss on the provided training sets and use the obtained CPE to predict all $\widehat{\boldsymbol{\eta}}$ for the test set, which are then used in the online algorithms on the test set using the protocol from Figure 1. For benchmarks without default train and test sets, we split them randomly in proportion 70/30.

2. With online CPE – instead of training CPE on a separate training set, it incrementally updated on observed instances in the sequence. In this case, we concatenate train and test sets to create one long sequence used in the online protocol. We report the results of this experiment in Appendix H.

Each experiment is repeated five times, each time randomly shuffling the sequence. We report the mean results over all runs. In Appendix F, we include more details regarding the experimental setup.

We compare OMMA with the following algorithms:

- Top-$k(\widehat{\boldsymbol{\eta}}(\boldsymbol{x}))$ – a classifier that selects k labels with the highest values in $\widehat{\boldsymbol{\eta}}(\boldsymbol{x})$. It is used as a baseline in multi-class and multi-label problems with the budget k constrain. For multi-class problems we use $k = 1$ if the budget is not specified.
- $\widehat{\eta}(\boldsymbol{x}) > 0.5$ – a classifier with a constant threshold that predicts a label as positive if its conditional probability is greater than 0.5. It is used as a baseline in multi-label problems without the budget k constrain.
- OFO – a consistent online algorithm for the (micro- and macro-averaged) F-measure in the multi-label setting (Busa-Fekete et al., 2015; Jasinska et al., 2016).
- Greedy – an algorithm that given past labels and predictions chooses its next prediction in order to maximize the expected utility. It might be seen as a close approximation of the OMMA algorithm. It has been motivated

[7] Code to reproduce the experiments: https://github.com/mwydmuch/xCOLUMNs

Table 2. Results of the different online algorithms on *multi-label* problems, averaged over 5 runs, reported as %. In this table we report the final performance obtained on the sequence of n samples. The best result on each metric is in **bold**, the second best is in *italic*. We additionally report basic statistics of the benchmarks: number of labels m and instances in the test sequence n. $\times$ – means that the algorithm does not support the optimization of that metric.

Method	Micro F1	Macro F1	Macro F1@3	Macro Rec.@3	Macro Pr.@3	Macro G-mean	Macro H-mean	Micro F1	Macro F1	Macro F1@3	Macro Rec.@3	Macro Pr.@3	Macro G-mean	Macro H-mean
	YouTube ($m = 46, n = 7926$)							**Eurlex-LexGlue** ($m = 100, n = 5000$)						
Top-k / $\hat{\eta}>0.5$	31.20	22.74	30.99	42.13	26.39	32.82	24.46	70.99	52.43	46.35	36.67	74.70	62.33	55.95
OFO	43.71	36.15	$\times$	$\times$	$\times$	$\times$	$\times$	73.23	58.93	$\times$	$\times$	$\times$	$\times$	$\times$
Greedy	$\times$	36.32	34.72	45.84	*67.18*	**77.98**	**77.93**	$\times$	59.83	54.19	52.67	*88.21*	89.74	89.73
Online-FW	43.67	36.00	34.38	45.83	38.92	*77.96*	*77.91*	**73.68**	59.53	54.18	52.68	58.76	89.69	89.75
Online-FW($\hat{\eta}$)	43.69	**36.47**	**35.43**	**45.89**	50.12	77.92	**77.93**	73.22	59.78	*54.34*	**53.87**	51.37	*89.91*	*89.84*
OMMA	**43.73**	*36.34*	34.81	*45.85*	**67.74**	**77.98**	**77.93**	*73.29*	**59.85**	54.15	52.67	**88.41**	89.74	89.73
OMMA($\widehat{\eta}$)	*43.72*	**36.47**	*35.38*	**45.89**	65.49	**77.98**	**77.93**	73.22	*59.84*	**54.37**	*53.85*	84.99	**89.92**	**89.85**
	Mediamill ($m = 101, n = 12914$)							**Flickr** ($m = 195, n = 24154$)						
Top-k / $\hat{\eta}>0.5$	52.45	4.06	4.43	4.35	7.42	4.62	3.65	29.46	18.27	26.39	38.96	21.38	27.08	20.02
OFO	*56.99*	12.36	$\times$	$\times$	$\times$	$\times$	$\times$	**41.05**	30.46	$\times$	$\times$	$\times$	$\times$	$\times$
Greedy	$\times$	12.43	10.29	*9.41*	*22.54*	*65.68*	**66.48**	$\times$	30.90	*30.42*	**46.41**	*57.55*	*83.39*	**83.37**
Online-FW	56.98	12.22	10.18	9.16	12.70	64.73	*65.68*	**41.05**	30.60	29.58	*46.38*	28.66	83.37	*83.28*
Online-FW($\hat{\eta}$)	*56.99*	*14.33*	**11.97**	**16.43**	17.35	65.37	65.56	*41.02*	**31.17**	29.65	46.28	25.73	83.37	83.20
OMMA	**57.00**	12.39	10.25	9.39	**22.85**	65.67	**66.48**	41.01	30.90	30.39	**46.41**	**58.35**	*83.39*	**83.37**
OMMA($\widehat{\eta}$)	*56.99*	**14.34**	*11.87*	**16.43**	18.31	**65.93**	65.35	*41.02*	*31.15*	**30.55**	46.33	55.56	**83.41**	83.23
	RCV1X ($m = 2456, n = 155962$)							**AmazonCat** ($m = 13330, n = 306784$)						
Top-k / $\hat{\eta}>0.5$	68.57	11.29	5.34	4.59	13.24	16.01	*12.36*	67.77	28.76	14.98	11.18	30.98	*33.93*	*30.03*
OFO	**69.83**	20.26	$\times$	$\times$	$\times$	$\times$	$\times$	*70.38*	39.60	$\times$	$\times$	$\times$	$\times$	$\times$
Greedy	$\times$	**20.80**	*16.01*	21.20	**30.91**	**69.07**	**67.04**	$\times$	44.20	43.64	57.81	**54.00**	**80.86**	**80.04**
Online-FW	**69.83**	19.82	15.33	21.09	19.88	**69.07**	**67.04**	**70.61**	42.42	40.20	57.74	40.10	**80.86**	**80.04**
Online-FW($\hat{\eta}$)	*69.79*	20.40	15.55	*22.21*	22.17	*69.06*	**67.04**	70.34	*47.32*	*44.63*	*58.87*	49.68	**80.86**	**80.04**
OMMA	69.77	20.57	15.87	21.08	*30.39*	**69.07**	**67.04**	69.90	43.04	42.14	57.80	*52.04*	**80.86**	**80.04**
OMMA($\widehat{\eta}$)	69.71	*20.71*	**16.07**	**22.23**	30.38	*69.06*	**67.04**	70.02	**47.67**	**45.06**	**58.89**	51.52	**80.86**	**80.04**

by a similar algorithm recently introduced for batch multi-label classification with predictions budgeted at k (Schultheis et al., 2023). An efficient implementation is possible for metrics that linearly decompose over labels (e.g., macro-averaged metrics).

- Online-FW – an adaptation of the Frank-Wolfe algorithm, used earlier for batch multi-class (Narasimhan et al., 2015b; 2022) and (budgeted at k) multi-label problems (Schultheis et al., 2024), to the online setting. It re-runs the batch algorithm on all instances observed so far (which need to be stored in memory) in exponentially growing intervals of 10×1.1^i instances with $i \in \{0, 1, 2, 3, \dots\}$, which we found to provide frequent enough updates for achieving good predictive performance. Since this algorithm, similar to OMMA, can be run for all metrics considered in the study, we treat it as the main baseline for our algorithm. For a fair comparison, we also introduce an alternative variant that uses $\widehat{\boldsymbol{\eta}}$ instead of true labels $\boldsymbol{y}$ to estimate its internal confusion matrix. We denote this algorithm as Online-FW($\widehat{\eta}$). Note that despite calling this method "online," it requires storing all previously seen instances in memory and thus does not fully adhere to the online paradigm.

- Offline-FW – the batch variant of the Frank-Wolfe algorithm used earlier in (Narasimhan et al., 2015b; Schultheis et al., 2024), trained using the same training set that was used to obtain CPE. It then uses the same probability estimates as online methods for inference on the test sequence without further updates to the classifier. The comparison with this algorithm is reported in Appendix G.

As discussed in Remark 4.4, we add both small constant $\epsilon =$ 1e-9 to the denominators in all metrics, as well as the regularization value $\lambda \in \{0, \text{1e-6}, \text{1e-3}, 0.1, 1\}$ to the entries of the confusion matrix used in all online algorithms (Greedy, OFO, Online-FW, and OMMA).

We use the following popular metrics for evaluating the algorithms: Micro and Macro-averaged F1, which are widely used to evaluate classifiers in many domains such as information retrieval; budgeted at $k = 3$ Macro F1, Macro Recall,

Table 3. Results of the online algorithms on *multi-class* problems, averaged over 5 runs, reported as %. In this table we report the final performance obtained on the sequence of n samples.

Method	Macro				Multi-class means		
	F1	F1@3	Rec.@3	Pr.@3	G-	H-	Q-
	NEWS20 ($m = 20, n = 7532$)						
Top-k	**83.37**	49.23	94.78	33.95	82.51	81.66	80.12
Greedy	82.44	72.69	*94.89*	**83.89**	×	×	×
On.-FW	82.88	54.10	94.88	19.90	82.68	*82.35*	80.95
On.-FW($\widehat{\eta}$)	82.55	56.06	**95.01**	15.58	**82.78**	**82.57**	*80.97*
OMMA	82.11	*72.84*	*94.89*	*83.88*	82.72	82.08	80.95
OMMA($\widehat{\eta}$)	*83.07*	**73.41**	**95.01**	83.30	*82.77*	82.18	**81.00**
	LEDGAR-LEXGLUE ($m = 100, n = 10000$)						
Top-k	79.06	51.80	92.04	38.88	0.00	0.00	69.21
Greedy	*79.30*	78.08	93.26	*91.17*	×	×	×
On.-FW	79.22	70.94	93.30	54.52	62.31	75.81	**74.59**
On.-FW($\widehat{\eta}$)	79.22	73.85	*93.38*	49.63	*78.02*	**77.58**	74.50
OMMA	79.28	*78.10*	93.26	**91.41**	77.48	74.62	**74.59**
OMMA($\widehat{\eta}$)	**79.34**	**78.22**	**93.39**	90.10	**78.03**	*76.08*	*74.53*
	CALTECH-256 ($m = 256, n = 14890$)						
Top-k	79.45	46.82	89.85	32.58	77.32	75.69	74.53
Greedy	*79.59*	*79.02*	90.20	96.61	×	×	×
On.-FW	79.15	70.29	90.20	63.22	78.31	*77.99*	76.00
On.-FW($\widehat{\eta}$)	79.29	72.97	*90.34*	65.42	**78.41**	**78.08**	*76.07*
OMMA	79.54	78.96	90.20	*96.77*	78.33	77.12	**76.15**
OMMA($\widehat{\eta}$)	**79.66**	**79.11**	**90.35**	**96.96**	*78.36*	77.01	75.99

and Macro Precision being well-suited to recommendation systems and extreme classification; and Macro G-mean and H-mean for multi-label problems and Multi-class G-mean, H-mean and Q-mean for multi-class problems, which are frequently used in class imbalance problems. Note that OMMA and Online-FW can target these metrics directly, while the scope of other algorithms is limited.

We present the results of the experiment with fixed CPE in Tables 2 and 3. We report the mean performance on the entire sequence of n instances for the best value of λ. Additionally, we present incremental performance on Figure 2 and the effect of using different λ values on Figure 3 for the Flickr dataset. The introduced OMMA algorithm matches the performance of the online Frank-Wolfe algorithm for most of the measures, performing much better on Macro-Precision@3, which is not Lipschitz. Concurrently, OMMA uses less time and memory as it does not require storing all previously seen instances. We also observe that OMMA, as an approximation of Greedy, matches its performance on metrics supported by this method. This is additionally confirmed by the plots where we can observe that the performance of these algorithms is very close to each other at each iteration t. Surprisingly, OMMA($\widehat{\eta}$) often performs better on macro-averaged metrics, especially on benchmarks with a large number of labels, where many of them have only a small number of positive samples. In these cases, the

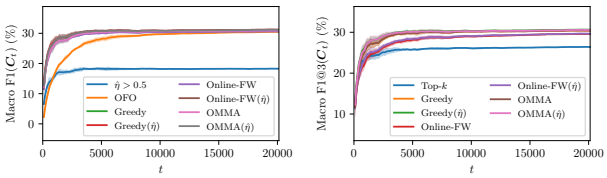

Figure 2. Comparison of the incremental performance of the online algorithms on the Flickr dataset. Averaged over 5 runs, the opaque fill indicates the standard deviation at given iteration t.

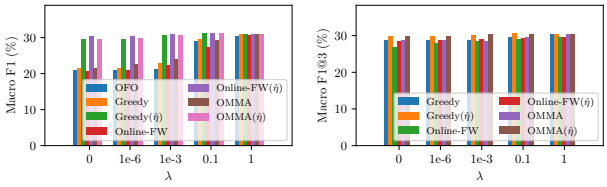

Figure 3. Impact of λ on the results of the online algorithms on the Flickr dataset. Averaged over 5 runs.

variance reduction effect of using $\widehat{\boldsymbol{\eta}}$ seems to compensate the estimation error of $\boldsymbol{\eta}$. It is also much less sensitive to the selection of λ, as small probability values that are added to many entries of the confusion matrix give the same effect as the regularization value λ. Extended results with additional algorithms, standard deviations, running times, and plots for the rest of the datasets can be found in Appendix G.

7. Discussion

In our framework, the algorithm aims to maximize the empirical value of the utility of the observed data sequence, $\psi(\boldsymbol{C}(\boldsymbol{y}^n, \widehat{\boldsymbol{y}}^n))$. An interesting direction of research is an alternative online framework, in which the goal is to sequentially learn a *classifier* $\widehat{\boldsymbol{h}}_n \colon \mathcal{X} \to \mathcal{Y}$, which is then evaluated on the entire population by means of $\psi(\boldsymbol{C}(\widehat{\boldsymbol{h}}_n))$, where $\boldsymbol{C}(\widehat{\boldsymbol{h}}_n) = \mathbb{E}_{(\boldsymbol{x},\boldsymbol{y})}[\boldsymbol{C}(\boldsymbol{y}, \widehat{\boldsymbol{h}}_n(\boldsymbol{x}))]$ is the population confusion matrix of $\widehat{\boldsymbol{h}}_n$.

In the analysis of the algorithm, we assumed the concavity of the metric. On one hand, *some* assumption of this type seems essential for no-regret learning; for instance, in multi-class macro-averaged F-measure learning, the utility belongs to function class, optimizing which is, in general, NP-hard (Narasimhan et al., 2016). On the other hand, our concavity assumption excludes an important class of linear-fractional functions (such as binary or micro-averaged F-measure), for which a specialized version of our algorithm is known to converge, although without providing a convergence rate (Busa-Fekete et al., 2015). The analysis of our method for linear-fractional utilities is an interesting direction for future research, especially given that our algorithm works very well on these metrics in the experiments.

Impact Statement

This paper presents work whose goal is to advance the field of Machine Learning. There are many potential societal consequences of our work, none which we feel must be specifically highlighted here.

Acknowledgments

This research has been supported by Poznan University of Technology via grant no. 0311/SBAD/0758 and by Academy of Finland via grants no. 347707 and 348215.

References

Babbar, R. and Schölkopf, B. Dismec: Distributed sparse machines for extreme multi-label classification. In *Proceedings of the tenth ACM international conference on web search and data mining*, pp. 721–729, 2017.

Baldi, P., Brunak, S., Chauvin, Y., Andersen, C. A., and Nielsen, H. Assessing the accuracy of prediction algorithms for classification: an overview. *Bioinformatics*, 16 (5):412–424, 2000.

Beck, A. *First-order methods in optimization*. SIAM, 2017.

Beck, A. and Teboulle, M. Smoothing and first order methods: A unified framework, 2012.

Bhatia, K., Dahiya, K., Jain, H., Mittal, A., Prabhu, Y., and Varma, M. The extreme classification repository: Multi-label datasets and code, 2016. URL http://manikvarma.org/downloads/XC/XMLRepository.html.

Bubeck, S. Convex optimization: Algorithms and complexity. 2015.

Busa-Fekete, R., Szörényi, B., Dembczynski, K., and Hüllermeier, E. Online F-measure optimization. In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pp. 595–603, 2015.

Cao, K., Wei, C., Gaidon, A., Arechiga, N., and Ma, T. Learning imbalanced datasets with label-distribution-aware margin loss. *Advances in neural information processing systems*, 32, 2019.

Cesa-Bianchi, N. and Lugosi, G. *Prediction, Learning, and Games*. Cambridge University Press, New York, NY, USA, 2006.

Chalkidis, I., Jana, A., Hartung, D., Bommarito, M., Androutsopoulos, I., Katz, D. M., and Aletras, N. Lexglue: A benchmark dataset for legal language understanding in english. *arXiv preprint arXiv:2110.00976*, 2021.

Chalkidis, I., Jana, A., Hartung, D., Bommarito, M., Androutsopoulos, I., Katz, D., and Aletras, N. LexGLUE: A benchmark dataset for legal language understanding in English. In Muresan, S., Nakov, P., and Villavicencio, A. (eds.), *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 4310–4330, Dublin, Ireland, May 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.acl-long.297.

Dembczynski, K., Kotlowski, W., Koyejo, O., and Natarajan, N. Consistency analysis for binary classification revisited. In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, volume 70 of *Proceedings of Machine Learning Research*, pp. 961–969. PMLR, 2017.

Dembczyński, K., Kotłowski, W., Koyejo, O., and Natarajan, N. Consistency analysis for binary classification revisited. In *Proceedings of the 34th International Conference on Machine Learning (ICML 2017)*, volume 70 of *Proceedings of Machine Learning Research*, pp. 961–969. PMLR, 2017.

Drummond, C. and Holte, R. C. Severe class imbalance: Why better algorithms aren't the answer. In *European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML PKDD)*, pp. 539–546. Springer, 2005.

Elkan, C. The foundations of cost-sensitive learning. JCAI'01, pp. 973–978, San Francisco, CA, USA, 2001. Morgan Kaufmann Publishers Inc.

Fan, R., Chang, K., Hsieh, C., Wang, X., and Lin, C. LIBLINEAR: A library for large linear classification. *Journal of Machine Learning Research*, 9:1871–1874, 2008.

Griffin, G., Holub, A., and Perona, P. Caltech-256 object category dataset. Technical Report 7694, California Institute of Technology, 2007.

Hazan, E. Introduction to online convex optimization. *Foundations and Trends® in Optimization*, 2(3-4):157–325, 2016.

Hazan, E. and Kale, S. Projection-free online learning. In *Proceedings of the 29th International Coference on International Conference on Machine Learning*, ICML'12, pp. 1843–1850, Madison, WI, USA, 2012. Omnipress. ISBN 9781450312851.

Hoeffding, W. Probability inequalities for sums of bounded random variables. *Journal of the American statistical association*, 58(301):13–30, 1963.

Jansche, M. A maximum expected utility framework for binary sequence labeling. In Zaenen, A. and van den Bosch, A. (eds.), *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics*, pp. 736–743, Prague, Czech Republic, June 2007. Association for Computational Linguistics.

Jasinska, K., Dembczynski, K., Busa-Fekete, R., Pfannschmidt, K., Klerx, T., and Hüllermeier, E. Extreme F-measure maximization using sparse probability estimates. In *Proceedings of the 33nd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016*, volume 48 of *JMLR Workshop and Conference Proceedings*, pp. 1435–1444. JMLR.org, 2016.

Joachims, T. A support vector method for multivariate performance measures. In *Proceedings of the 22nd International Conference on Machine Learning (ICML 2005), August 7-11, 2005, Bonn, Germany*, 2005.

Kar, P., Narasimhan, H., and Jain, P. Online and stochastic gradient methods for non-decomposable loss functions. *Advances in Neural Information Processing Systems*, 27, 2014.

Kennedy, K., Mac Namee, B., and Delany, S. J. Learning without default: A study of one-class classification and the low-default portfolio problem. In *Irish conference on artificial intelligence and cognitive science*, pp. 174–187. Springer, 2009.

Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. In Bengio, Y. and LeCun, Y. (eds.), *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.

Kotłowski, W. and Dembczyński, K. Surrogate regret bounds for generalized classification performance metrics. *Machine Learning*, 10:549–572, 2017.

Koyejo, O., Natarajan, N., Ravikumar, P., and Dhillon, I. S. Consistent binary classification with generalized performance metrics. In *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, pp. 2744–2752, 2014.

Koyejo, O., Natarajan, N., Ravikumar, P., and Dhillon, I. S. Consistent multilabel classification. In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pp. 3321–3329, 2015.

Lang, K. Newsweeder: Learning to filter netnews. In *Machine learning proceedings 1995*, pp. 331–339. Elsevier, 1995.

Lewis, D. D. Evaluating and optimizing autonomous text classification systems. In *Proceedings of the 18th Intl. ACM SIGIR Conf. on Research and Development in Information Retrieval*, pp. 246–254. ACM, 1995.

Lewis, D. D., Yang, Y., Russell-Rose, T., and Li, F. RCV1: A new benchmark collection for text categorization research. *Journal of Machine Learning Research*, 5(Apr): 361–397, 2004.

McAuley, J. and Leskovec, J. Hidden factors and hidden topics: understanding rating dimensions with review text. In *Proceedings of the 7th ACM conference on Recommender systems*, pp. 165–172, 2013.

Menon, A. K., Narasimhan, H., Agarwal, S., and Chawla, S. On the statistical consistency of algorithms for binary classification under class imbalance. In *International Conference on Machine Learning (ICML)*, 2013.

Narasimhan, H., Vaish, R., and Agarwal, S. On the statistical consistency of plug-in classifiers for non-decomposable performance measures. In *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, pp. 1493–1501, 2014.

Narasimhan, H., Kar, P., and Jain, P. Optimizing non-decomposable performance measures: A tale of two classes. In *ICML*, volume 37, pp. 199–208, 2015a.

Narasimhan, H., Ramaswamy, H., Saha, A., and Agarwal, S. Consistent multiclass algorithms for complex performance measures. In Bach, F. and Blei, D. (eds.), *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pp. 2398–2407, Lille, France, 07 2015b. PMLR.

Narasimhan, H., Pan, W., Kar, P., Protopapas, P., and Ramaswamy, H. G. Optimizing the multiclass f-measure via biconcave programming. In *2016 IEEE 16th International Conference on Data Mining (ICDM)*, pp. 1101–1106, 2016.

Narasimhan, H., Ramaswamy, H. G., Tavker, S. K., Khurana, D., Netrapalli, P., and Agarwal, S. Consistent multiclass algorithms for complex metrics and constraints. 2022. doi: 10.48550/ARXIV.2210.09695.

Natarajan, N., Dhillon, I. S., Ravikumar, P., and Tewari, A. Cost-sensitive learning with noisy labels. *Journal of Machine Learning Research*, 18(155):1–33, 2018.

Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S. Pytorch: An imperative style, high-performance deep learning library. In Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 32*, pp. 8024–8035. Curran Associates, Inc., 2019.

Perozzi, B., Al-Rfou, R., and Skiena, S. Deepwalk: Online learning of social representations. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '14, pp. 701–710. ACM, 2014.

Reddi, S. J., Sra, S., Póczos, B., and Smola, A. Stochastic frank-wolfe methods for nonconvex optimization. *Allerton Conference on Communication, Control, and Computing (Allerton)*, pp. 1244–1251, 2016.

Schultheis, E., Wydmuch, M., Kotlowski, W., Babbar, R., and Dembczynski, K. Generalized test utilities for long-tail performance in extreme multi-label classification. In Oh, A., Neumann, T., Globerson, A., Saenko, K., Hardt, M., and Levine, S. (eds.), *Advances in Neural Information Processing Systems*, volume 36, pp. 22269–22303. Curran Associates, Inc., 2023.

Schultheis, E., Kotłowski, W., Wydmuch, M., Babbar, R., Borman, S., and Dembczyński, K. Consistent algorithms for multi-label classification with macro-at-k metrics. In *The Twelfth International Conference on Learning Representations*, 2024.

Shalev-Shwartz, S. Online learning and online convex optimization. *Foundations and Trends in Machine Learning*, 4(2):107–194, 2012.

Simonyan, K. and Zisserman, A. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.

Singh, S. and Khim, J. T. Optimal binary classification beyond accuracy. *Advances in Neural Information Processing Systems*, 35:18226–18240, 2022.

Snoek, C. G. M., Worring, M., van Gemert, J. C., Geusebroek, J.-M., and Smeulders, A. W. M. The challenge problem for automated detection of 101 semantic concepts in multimedia. In *ACM International Conference on Multimedia*, pp. 421–430. Association for Computing Machinery, 2006.

Tang, L. and Liu, H. Relational learning via latent social dimensions. In *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '09, pp. 817–826. Association for Computing Machinery, 2009.

Waegeman, W., Dembczynski, K., Jachnik, A., Cheng, W., and Hüllermeier, E. On the bayes-optimality of F-measure maximizers. *Journal of Machine Learning Research*, 15(1):3333–3388, 2014.

Wang, S. and Yao, X. Multiclass imbalance problems: Analysis and potential solutions. *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics*, 42(4): 1119–1130, 2012.

Yan, Y., Yang, T., Yang, Y., and Chen, J. A framework of online learning with imbalanced streaming data. *Proceedings of the AAAI Conference on Artificial Intelligence*, 31 (1), 2017.

Ye, N., Chai, K. M. A., Lee, W. S., and Chieu, H. L. Optimizing F-measure: A tale of two approaches. In *Proceedings of the 29th International Conference on Machine Learning, ICML 2012, Edinburgh, Scotland, UK, June 26 - July 1, 2012*. icml.cc / Omnipress, 2012.

Zhang, X., Liu, M., Zhou, X., and Yang, T. Faster online learning of optimal threshold for consistent f-measure optimization. 31, 2018.