
Online F-Measure Optimization

Róbert Busa-Fekete

Department of Computer Science
University of Paderborn, Germany
busarobi@upb.de

Balázs Szörényi

Technion, Haifa, Israel /
MTA-SZTE Research Group on
Artificial Intelligence, Hungary
szorenyibalazs@gmail.com

Krzysztof Dembczyński

Institute of Computing Science
Poznań University of Technology, Poland
kdembczynski@cs.put.poznan.pl

Eyke Hüllermeier

Department of Computer Science
University of Paderborn, Germany
eyke@upb.de

Abstract

The F-measure is an important and commonly used performance metric for binary prediction tasks. By combining precision and recall into a single score, it avoids disadvantages of simple metrics like the error rate, especially in cases of imbalanced class distributions. The problem of optimizing the F-measure, that is, of developing learning algorithms that perform optimally in the sense of this measure, has recently been tackled by several authors. In this paper, we study the problem of F-measure maximization in the setting of online learning. We propose an efficient online algorithm and provide a formal analysis of its convergence properties. Moreover, first experimental results are presented, showing that our method performs well in practice.

1 Introduction

Being rooted in information retrieval [16], the so-called F-measure is nowadays routinely used as a performance metric in various prediction tasks. Given predictions $\hat{\mathbf{y}} = (\hat{y}_1, \dots, \hat{y}_t) \in \{0, 1\}^t$ of t binary labels $\mathbf{y} = (y_1, \dots, y_t)$, the F-measure is defined as

$$F(\mathbf{y}, \hat{\mathbf{y}}) = \frac{2 \sum_{i=1}^t y_i \hat{y}_i}{\sum_{i=1}^t y_i + \sum_{i=1}^t \hat{y}_i} = \frac{2 \cdot \text{precision}(\mathbf{y}, \hat{\mathbf{y}}) \cdot \text{recall}(\mathbf{y}, \hat{\mathbf{y}})}{\text{precision}(\mathbf{y}, \hat{\mathbf{y}}) + \text{recall}(\mathbf{y}, \hat{\mathbf{y}})} \in [0, 1], \quad (1)$$

where $\text{precision}(\mathbf{y}, \hat{\mathbf{y}}) = \sum_{i=1}^t y_i \hat{y}_i / \sum_{i=1}^t \hat{y}_i$, $\text{recall}(\mathbf{y}, \hat{\mathbf{y}}) = \sum_{i=1}^t y_i \hat{y}_i / \sum_{i=1}^t y_i$, and where $0/0 = 1$ by definition. Compared to measures like the error rate in binary classification, maximizing the F-measure enforces a better balance between performance on the minority and majority class; therefore, it is more suitable in the case of imbalanced data. Optimizing for such an imbalanced measure is very important in many real-world applications where positive labels are significantly less frequent than negative ones. It can also be generalized to a weighted harmonic average of precision and recall. Yet, for the sake of simplicity, we stick to the unweighted mean, which is often referred to as the F1-score or the F1-measure.

Given the importance and usefulness of the F-measure, it is natural to look for learning algorithms that perform optimally in the sense of this measure. However, optimizing the F-measure is a quite challenging problem, especially because the measure is not decomposable over the binary predictions. This problem has received increasing attention in recent years and has been tackled by several authors [19, 20, 18, 10, 11]. However, most of this work has been done in the standard setting of batch learning.

In this paper, we study the problem of F-measure optimization in the setting of online learning [4, 2], which is becoming increasingly popular in machine learning. In fact, there are many applications in which training data is arriving progressively over time, and models need to be updated and maintained incrementally. In our setting, this means that in each round t the learner first outputs a prediction \hat{y}_t and then observes the true label y_t . Formally, the protocol in round t is as follows:

1. first an instance $\mathbf{x}_t \in \mathcal{X}$ is observed by the learner,
2. then the predicted label \hat{y}_t for \mathbf{x}_t is computed on the basis of the first t instances $(\mathbf{x}_1, \dots, \mathbf{x}_t)$, the $t-1$ labels (y_1, \dots, y_{t-1}) observed so far, and the corresponding predictions $(\hat{y}_1, \dots, \hat{y}_{t-1})$,
3. finally, the label y_t is revealed to the learner.

The goal of the learner is then to maximize

$$F_{(t)} = F((y_1, \dots, y_t), (\hat{y}_1, \dots, \hat{y}_t)) \quad (2)$$

over time. Optimizing the F-measure in an online fashion is challenging mainly because of the non-decomposability of the measure, and the fact that the \hat{y}_t cannot be changed after round t .

As a potential application of online F-measure optimization consider the recommendation of news from RSS feeds or tweets [1]. Besides, it is worth mentioning that online methods are also relevant in the context of big data and large-scale learning, where the volume of data, despite being finite, prevents from processing each data point more than once [21, 7]. Treating the data as a stream, online algorithms can then be used as single-pass algorithms. Note, however, that single-pass algorithms are evaluated only at the end of the training process, unlike online algorithms that are supposed to learn and predict simultaneously.

We propose an online algorithm for F-measure optimization, which is not only very efficient but also easy to implement. Unlike other methods, our algorithm does not require extra validation data for tuning a threshold (that separates between positive and negative predictions), and therefore allows the entire data to be used for training. We provide a formal analysis of the convergence properties of our algorithm and prove its statistical consistency under different assumptions on the learning process. Moreover, first experimental results are presented, showing that our method performs well in practice.

2 Formal Setting

In this paper, we consider a stochastic setting in which $(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_t, y_t)$ are assumed to be i.i.d. samples from some unknown distribution $\rho(\cdot)$ on $\mathcal{X} \times \mathcal{Y}$, where $\mathcal{Y} = \{0, 1\}$ is the label space and \mathcal{X} is some instance space. We denote the marginal distribution of the feature vector X by $\mu(\cdot)$.¹ Then, the posterior probability of the positive class, i.e., the conditional probability that $Y = 1$ given $X = \mathbf{x}$, is $\eta(\mathbf{x}) = \mathbf{P}(Y = 1 | X = \mathbf{x}) = \frac{\rho(\mathbf{x}, 1)}{\rho(\mathbf{x}, 0) + \rho(\mathbf{x}, 1)}$. The prior distribution of class 1 can be written as $\pi_1 = \mathbf{P}(Y = 1) = \int_{\mathbf{x} \in \mathcal{X}} \eta(\mathbf{x}) \, d\mu(\mathbf{x})$.

Let $\mathcal{B} = \{f : \mathcal{X} \rightarrow \{0, 1\}\}$ be the set of all binary classifiers over the set \mathcal{X} . The F-measure of a binary classifier $f \in \mathcal{B}$ is calculated as

$$F(f) = \frac{2 \int_{\mathcal{X}} \eta(\mathbf{x}) f(\mathbf{x}) \, d\mu(\mathbf{x})}{\int_{\mathcal{X}} \eta(\mathbf{x}) \, d\mu(\mathbf{x}) + \int_{\mathcal{X}} f(\mathbf{x}) \, d\mu(\mathbf{x})} = \frac{2\mathbf{E}[\eta(X)f(X)]}{\mathbf{E}[\eta(X)] + \mathbf{E}[f(X)]}.$$

According to [19], the expected value of (1) converges to $F(f)$ with $t \rightarrow \infty$ when f is used to calculate \hat{y}_t , i.e., $\hat{y}_t = f(\mathbf{x}_t)$. Thus, $\lim_{t \rightarrow \infty} \mathbf{E}[F((y_1, \dots, y_t), (f(\mathbf{x}_1), \dots, f(\mathbf{x}_t)))] = F(f)$.

Now, let $\mathcal{G} = \{g : \mathcal{X} \rightarrow [0, 1]\}$ denote the set of all *probabilistic* binary classifiers over the set \mathcal{X} , and let $\mathcal{T} \subseteq \mathcal{B}$ denote the set of binary classifiers that are obtained by thresholding a classifier $g \in \mathcal{G}$ —that is, classifiers of the form

$$g^\tau(\mathbf{x}) = \llbracket g(\mathbf{x}) \geq \tau \rrbracket \quad (3)$$

for some threshold $\tau \in [0, 1]$, where $\llbracket \cdot \rrbracket$ is the indicator function that evaluates to 1 if its argument is true and 0 otherwise.

¹ \mathcal{X} is assumed to exhibit the required measurability properties.

According to [19], the optimal F-score computed as $\max_{f \in \mathcal{B}} F(f)$ can be achieved by a thresholded classifier. More precisely, let us define the *thresholded F-measure* as

$$F(\tau) = F(\eta^\tau) = \frac{2 \int_{\mathcal{X}} \eta(\mathbf{x}) \mathbb{1}[\eta(\mathbf{x}) \geq \tau] d\mu(\mathbf{x})}{\int_{\mathcal{X}} \eta(\mathbf{x}) d\mu(\mathbf{x}) + \int_{\mathcal{X}} \mathbb{1}[\eta(\mathbf{x}) \geq \tau] d\mu(\mathbf{x})} = \frac{2\mathbf{E}[\eta(X) \mathbb{1}[\eta(X) \geq \tau]]}{\mathbf{E}[\eta(X)] + \mathbf{E}[\mathbb{1}[\eta(X) \geq \tau]]} \quad (4)$$

Then the *optimal threshold* τ^* can be obtained as

$$\tau^* = \operatorname{argmax}_{0 \leq \tau \leq 1} F(\tau) . \quad (5)$$

Clearly, for the classifier in the form of (3) with $g(\mathbf{x}) = \eta(\mathbf{x})$ and $\tau = \tau^*$, we have $F(g^\tau) = F(\tau^*)$. Then, as shown by [19] (see their Theorem 4), the performance of any binary classifier $f \in \mathcal{B}$ cannot exceed $F(\tau^*)$, i.e., $F(f) \leq F(\tau^*)$ for all $f \in \mathcal{B}$. Therefore, estimating posteriors first and adjusting a threshold afterward appears to be a reasonable strategy. In practice, this seems to be the most popular way of maximizing the F-measure in a batch mode; we call it the 2-stage F-measure maximization approach, or 2S for short. More specifically, the 2S approach consists of two steps: first, a classifier is trained for estimating the posteriors, and second, a threshold is tuned on the posterior estimates. For the time being, we are not interested in the training of this classifier but focus on the second step, that is, the labelling of instances via thresholding posterior probabilities. For doing this, suppose a finite set $\mathcal{D}_N = \{(\mathbf{x}_i, y_i)\}_{i=1}^N$ of labeled instances are given as training information. Moreover, suppose estimates $\hat{p}_i = g(\mathbf{x}_i)$ of the posterior probabilities $p_i = \eta(\mathbf{x}_i)$ are provided by a classifier $g \in \mathcal{G}$. Next, one might define the F-score obtained by applying the threshold classifier g^τ on the data \mathcal{D}_N as follows:

$$F(\tau; g, \mathcal{D}_N) = \frac{\sum_{i=1}^N y_i \mathbb{1}[\tau \leq g(\mathbf{x}_i)]}{\sum_{i=1}^N y_i + \sum_{i=1}^N \mathbb{1}[\tau \leq g(\mathbf{x}_i)]} \quad (6)$$

In order to find an optimal threshold $\tau_N \in \operatorname{argmax}_{0 \leq \tau \leq 1} F(\tau; g, \mathcal{D}_N)$, it suffices to search the finite set $\{\hat{p}_1, \dots, \hat{p}_N\}$, which requires time $\mathcal{O}(N \log N)$. In [19], it is shown that $F(\tau; g, \mathcal{D}_N) \xrightarrow{P} F(g^\tau)$ as $N \rightarrow \infty$ for any $\tau \in (0, 1)$, and [11] provides an even stronger result: If a classifier $g_{\mathcal{D}_N}$ is induced from \mathcal{D}_N by an L_1 -consistent learner,² and a threshold τ_N is obtained by maximizing (6) on an independent set \mathcal{D}'_N , then $F(g_{\mathcal{D}_N}^{\tau_N}) \xrightarrow{P} F(\tau^*)$ as $N \rightarrow \infty$ (under mild assumptions on the data distribution).

3 Maximizing the F-Measure on a Population Level

In this section we assume that the data distribution is known. According to the analysis in the previous section, optimizing the F-measure boils down to finding the optimal threshold τ^* . At this point, an observation is in order.

Remark 1. *In general, the function $F(\tau)$ is neither convex nor concave. For example, when \mathcal{X} is finite, then the denominator and numerator of (4) are step functions, whence so is $F(\tau)$. Therefore, gradient methods cannot be applied for finding τ^* .*

Nevertheless, τ^* can be found based on a recent result of [20], who show that finding the root of

$$h(\tau) = \int_{\mathbf{x} \in \mathcal{X}} \max(0, \eta(\mathbf{x}) - \tau) d\mu(\mathbf{x}) - \tau \pi_1 \quad (7)$$

is a necessary and sufficient condition for optimality. Note that $h(\tau)$ is continuous and strictly decreasing, with $h(0) = \pi_1$ and $h(1) = -\pi_1$. Therefore, $h(\tau) = 0$ has a unique solution which is τ^* . Moreover, [20] also prove an interesting relationship between the optimal threshold and the F-measure induced by that threshold: $F(\tau^*) = 2\tau^*$.

The marginal distribution of the feature vectors, $\mu(\cdot)$, induces a distribution $\zeta(\cdot)$ on the posteriors: $\zeta(p) = \int_{\mathbf{x} \in \mathcal{X}} \mathbb{1}[\eta(\mathbf{x}) = p] d\mu(\mathbf{x})$ for all $p \in [0, 1]$. By definition, $\mathbb{1}[\eta(\mathbf{x}) = p]$ is the Radon-Nikodym derivative of $\frac{d\mu}{d\zeta}$, and $\zeta(p)$ the density of observing an instance \mathbf{x} for which the probability of the

²A learning algorithm, viewed as a map from samples \mathcal{D}_N to classifiers $g_{\mathcal{D}_N}$, is called L_1 -consistent w.r.t. the data distribution ρ if $\lim_{N \rightarrow \infty} \mathbf{P}_{\mathcal{D}_N \sim \rho} \left(\int_{\mathbf{x} \in \mathcal{X}} |g_{\mathcal{D}_N}(\mathbf{x}) - \eta(\mathbf{x})| d\mu(\mathbf{x}) > \epsilon \right) = 0$ for all $\epsilon > 0$.

positive label is p . We shall write concisely $d\nu(p) = \zeta(p) dp$. Since $\nu(\cdot)$ is an *induced probability measure*, the measurable transformation allows us to rewrite the notions introduced above in terms of $\nu(\cdot)$ instead of $\mu(\cdot)$ —see, for example, Section 1.4 in [17]. For example, the prior probability $\int_{\mathcal{X}} \eta(\mathbf{x}) d\mu$ can be written equivalently as $\int_0^1 p d\nu(p)$. Likewise, (7) can be rewritten as follows:

$$\begin{aligned} h(\tau) &= \int_0^1 \max(0, p - \tau) d\nu(p) - \tau \int_0^1 p d\nu(p) = \int_\tau^1 p - \tau d\nu(p) - \tau \int_0^1 p d\nu(p) \\ &= \int_\tau^1 p d\nu(p) - \tau \left[\int_\tau^1 1 d\nu(p) - \int_0^1 p d\nu(p) \right] \end{aligned} \quad (8)$$

Equation (8) will play a central role in our analysis. Note that precise knowledge of $\nu(\cdot)$ suffices to find the maxima of $F(\tau)$. This is illustrated by two examples presented in Appendix E, in which we assume specific distributions for $\nu(\cdot)$, namely uniform and Beta distributions.

4 Algorithmic Solution

In this section, we provide an algorithmic solution to the online F-measure maximization problem. For this, we shall need in each round t some classifier $g_t \in \mathcal{G}$ that provides us with some estimate $\hat{p}_t = g_t(\mathbf{x}_t)$ of the probability $\eta(\mathbf{x}_t)$. We would like to stress again that the focus of our analysis is on optimal thresholding instead of classifier learning. Thus, we assume the sequence of classifiers g_1, g_2, \dots to be produced by an external online learner, for example, logistic regression trained by stochastic gradient descent.

As an aside, we note that F-measure maximization is not directly comparable with the task that is most often considered and analyzed in online learning, namely regret minimization [4]. This is mainly because the F-measure is a non-decomposable performance metric. In fact, the cumulative regret is a summation of a per-round regret r_t , which only depends on the prediction \hat{y}_t and the true outcome y_t [11]. In the case of the F-measure, the score $F_{(t)}$, and therefore the optimal prediction \hat{y}_t , depends on the entire history, that is, all observations and decisions made by the learner till time t . This is discussed in more detail in Section 6.

The most naive way of forecasting labels is to implement online learning as repeated batch learning, that is, to apply a batch learner (such as 2S) to $\mathcal{D}_t = \{(\mathbf{x}_i, y_i)\}_{i=1}^t$ in each time step t . Obviously, however, this strategy is prohibitively expensive, as it requires storage of all data points seen so far (at least in mini-batches), as well as optimization of the threshold τ_t and re-computation of the classifier g_t on an ever growing number of examples.

In the following, we propose a more principled technique to maximize the online F-measure. Our approach is based on the observation that $h(\tau^*) = 0$ and $h(\tau)(\tau - \tau^*) < 0$ for any $\tau \in [0, 1]$ such that $\tau \neq \tau^*$ [20]. Moreover, it is a monotone decreasing continuous function. Therefore, finding the optimal threshold τ^* can be viewed as a root finding problem. In practice, however, $h(\tau)$ is not known and can only be estimated. Let us define $h(\tau, y, \hat{y}) = y\hat{y} - \tau(y + \hat{y})$. For now, assume $\eta(\mathbf{x})$ to be known and write concisely $\hat{h}(\tau) = h(\tau, y, \mathbb{I}[\eta(\mathbf{x}) \geq \tau])$. We can compute the expectation of $\hat{h}(\tau)$ with respect to the data distribution for a fixed threshold τ as follows:

$$\begin{aligned} \mathbf{E}[\hat{h}(\tau)] &= \mathbf{E}[h(\tau, y, \mathbb{I}[\eta(\mathbf{x}) \geq \tau])] = \mathbf{E}[y \mathbb{I}[\eta(\mathbf{x}) \geq \tau] - \tau(y + \mathbb{I}[\eta(\mathbf{x}) \geq \tau])] \\ &= \int_0^1 p \mathbb{I}[p \geq \tau] d\nu(p) - \tau \left(\int_0^1 p + \mathbb{I}[p \geq \tau] d\nu(p) \right) \\ &= \int_\tau^1 p d\nu(p) - \tau \left[\int_0^1 p d\nu(p) + \int_\tau^1 1 d\nu(p) \right] = h(\tau) \end{aligned} \quad (9)$$

Thus, an unbiased estimate of $h(\tau)$ can be obtained by evaluating $\hat{h}(\tau)$ for an instance \mathbf{x} . This suggests designing a stochastic approximation algorithm that is able to find the root of $h(\cdot)$ similarly to the Robbins-Monro algorithm [12]. Exploiting the relationship between the optimal F-measure and the optimal threshold, $F(\tau^*) = 2\tau^*$, we define the threshold in time step t as

$$\tau_t = \frac{1}{2} F_{(t)} = \frac{a_t}{b_t} \quad \text{where} \quad a_t = \sum_{i=1}^t y_i \hat{y}_i, \quad b_t = \sum_{i=1}^t y_i + \sum_{i=1}^t \hat{y}_i. \quad (10)$$

With this threshold, the first differences between thresholds, i.e. $\tau_{t+1} - \tau_t$, can be written as follows.

Proposition 2. *If thresholds τ_t are defined according to (10) and \hat{y}_{t+1} as $\llbracket \eta(\mathbf{x}_{t+1}) > \tau_t \rrbracket$, then*

$$(\tau_{t+1} - \tau_t)b_{t+1} = h(\tau_t, y_{t+1}, \hat{y}_{t+1}) . \quad (11)$$

The proof of Prop. 2 is deferred to Appendix A. According to (11), the method we obtain ‘‘almost’’ coincides with the update rule of the Robbins-Monro algorithm. There are, however, some notable differences. In particular, the sequence of coefficients, namely the values $1/b_{t+1}$, does not consist of predefined real values converging to zero (as fast as $1/t$). Instead, it consists of random quantities that depend on the history, namely the observed labels y_1, \dots, y_t and the predicted labels $\hat{y}_1, \dots, \hat{y}_t$. Moreover, these ‘‘coefficients’’ are not independent of $h(\tau_t, y_{t+1}, \hat{y}_{t+1})$ either. In spite of these additional difficulties, we shall present a convergence analysis of our algorithm in the next section.

The pseudo-code of our online F-measure optimization algorithm, called Online F-measure Optimizer (OFO), is shown in Algorithm 1. The forecast rule can be written in the form of $\hat{y}_t = \llbracket p_t \geq \tau_{t-1} \rrbracket$ for \mathbf{x}_t where the threshold is defined in (10) and $p_t = \eta(\mathbf{x}_t)$. In practice, we use $\hat{p}_t = g_{t-1}(\mathbf{x}_t)$ as an estimate of the true posterior p_t . In line 8 of the code, an online learner $\mathcal{A} : \mathcal{G} \times \mathcal{X} \times \mathcal{Y} \rightarrow \mathcal{G}$ is assumed, which produces classifiers g_t by incrementally updating the current classifier with the newly observed example, i.e., $g_t = \mathcal{A}(g_{t-1}, \mathbf{x}_t, y_t)$.

In our experimental study, we shall test and compare various state-of-the-art online learners as possible choices for \mathcal{A} .

Algorithm 1 OFO

- 1: Select g_0 from \mathcal{B} , and set $\tau_0 = 0$
 - 2: **for** $t = 1 \rightarrow \infty$ **do**
 - 3: Observe the instance \mathbf{x}_t
 - 4: $\hat{p}_t \leftarrow g_{t-1}(\mathbf{x}_t)$ \triangleright estimate posterior
 - 5: $\hat{y}_t \leftarrow \llbracket \hat{p}_t \geq \tau_{t-1} \rrbracket$ \triangleright current prediction
 - 6: Observe label y_t
 - 7: Calculate $F_{(t)} = \frac{2a_t}{b_t}$ and $\tau_t = \frac{a_t}{b_t}$
 - 8: $g_t \leftarrow \mathcal{A}(g_{t-1}, \mathbf{x}_t, y_t)$ \triangleright update the classifier
 - 9: **return** τ_T
-

5 Consistency

In this section, we provide an analysis of the online F-measure optimizer proposed in the previous section. More specifically, we show the statistical consistency of the OFO algorithm: The sequence of online thresholds and F-scores produced by this algorithm converge, respectively, to the optimal threshold τ^* and the optimal thresholded F-score $F(\tau^*)$ in probability. As a first step, we prove this result under the assumption of knowledge about the true posterior probabilities; then, in a second step, we consider the case of estimated posteriors.

Theorem 3. *Assume the posterior probabilities $p_t = \eta(\mathbf{x}_t)$ of the positive class to be known in each step of the online learning process. Then, the sequences of thresholds τ_t and online F-scores $F_{(t)}$ produced by OFO both converge in probability to their optimal values τ^* and $F(\tau^*)$, respectively: For any $\epsilon > 0$, we have $\lim_{t \rightarrow \infty} \mathbf{P}(|\tau_t - \tau^*| > \epsilon) = 0$ and $\lim_{t \rightarrow \infty} \mathbf{P}(|F_{(t)} - F(\tau^*)| > \epsilon) = 0$.*

Here is a sketch of the proof of this theorem, the details of which can be found in the supplementary material (Appendix B):

- We focus on $\{\tau_t\}_{t=1}^{\infty}$, which is a stochastic process the filtration of which is defined as $\mathcal{F}_t = \{y_1, \dots, y_t, \hat{y}_1, \dots, \hat{y}_t\}$. For this filtration, one can show that $\hat{h}(\tau_t)$ is \mathcal{F}_t -measurable and $\mathbf{E}[\hat{h}(\tau_t) | \mathcal{F}_t] = h(\tau_t)$ based on (9).
- As a first step, we can decompose the update rule given in (11) as follows: $\mathbf{E} \left[\frac{1}{b_{t+1}} \hat{h}(\tau_t) \middle| \mathcal{F}_t \right] = \frac{1}{b_{t+2}} h(\tau_t) + \mathcal{O} \left(\frac{1}{b_t^2} \right)$ conditioned on the filtration \mathcal{F}_t (see Lemma 7).
- Next, we show that the sequence $1/b_t$ behaves similarly to $1/t$, in the sense that $\sum_{t=1}^{\infty} \mathbf{E} [1/b_t^2] < \infty$ (see Lemma 8). Moreover, one can show that $\sum_{t=1}^{\infty} \mathbf{E} [1/b_t] \geq \sum_{t=1}^{\infty} \frac{1}{2t} = \infty$.
- Although $h(\tau)$ is not differentiable on $[0, 1]$ in general (it can be piecewise linear, for example), one can show that its finite difference is between $-1 - \pi_1$ and $-\pi_1$ (see Proposition 9 in the appendix). As a consequence of this result, our process defined in (11) does not get stuck even close to τ^* .
- The main part of the proof is devoted to analyzing the properties of the sequence of $\beta_t = \mathbf{E} [(\tau_t - \tau^*)^2]$ for which we show that $\lim_{t \rightarrow \infty} \beta_t = 0$, which is sufficient for the statement

of the theorem. Our proof follows the convergence analysis of [12]. Nevertheless, our analysis essentially differs from theirs, since in our case, the coefficients cannot be chosen freely. Instead, as explained before, they depend on the labels observed and predicted so far. In addition, the noisy estimation of $h(\cdot)$ depends on the labels, too, but the decomposition step allows us to handle this undesired effect.

Remark 4. *In principle, the Robbins-Monro algorithm can be applied for finding the root of $h(\cdot)$ as well. This yields an update rule similar to (11), with $1/b_{t+1}$ replaced by C/t for a constant $C > 0$. In this case, however, the convergence of the online F-measure is difficult to analyze (if at all), because the empirical process cannot be written in a nice form. Moreover, as it has been found in the analysis, the coefficient C should be set $\approx 1/\pi_1$ (see Proposition 9 and the choice of $\{k_t\}$ at the end of the proof of Theorem 3). Yet, since π_1 is not known beforehand, it needs to be estimated from the samples, which implies that the coefficients are not independent of the noisy evaluations of $h(\cdot)$ —just like in the case of the OFO algorithm. Interestingly, OFO seems to properly adjust the values $1/b_{t+1}$ in an adaptive manner (b_t is a sum of two terms, the first of which is $t\pi_1$ in expectation), which is a very nice property of the algorithm. Empirically, based on synthetic data, we found the performance of the original Robbins-Monro algorithm to be on par with OFO.*

As already announced, we are now going to relax the assumption of known posterior probabilities $p_t = \eta(\mathbf{x}_t)$. Instead, estimates $\hat{p}_t = g_t(\mathbf{x}_t) \approx p_t$ of these probabilities are obtained by classifiers g_t that are provided by the external online learner in Algorithm 1. More concretely, assume an online learner $\mathcal{A} : \mathcal{G} \times \mathcal{X} \times \mathcal{Y} \rightarrow \mathcal{G}$, where \mathcal{G} is the set of probabilistic classifiers. Given a current model g_t and a new example (\mathbf{x}_t, y_t) , this learner produces an updated classifier $g_{t+1} = \mathcal{A}(g_t, \mathbf{x}_t, y_t)$. Showing a consistency result for this scenario requires some assumptions on the online learner. With this formal definition of online learner, a statistical consistency result similar to Theorem 3 can be shown. The proof of the following theorem is again deferred to supplementary material (Appendix C).

Theorem 5. *Assume that the classifiers $(g_t)_{t=1}^\infty$ in the OFO framework are provided by an online learner for which the following holds: There is a $\lambda > 0$ such that $\mathbf{E} \left[\int_{\mathbf{x} \in \mathcal{X}} |\eta(\mathbf{x}) - g_t(\mathbf{x})| d\mu(\mathbf{x}) \right] = \mathcal{O}(t^{-\lambda})$. Then, $F_{(t)} \xrightarrow{P} F(\tau^*)$ and $\tau_t \xrightarrow{P} \tau^*$.*

This theorem’s requirement on the online learner is stronger than what is assumed by [11] and recalled in Footnote 2. First, the learner is trained online and not in a batch mode. Second, we also require that the L_1 error of the learner goes to 0 with a convergence rate of order $t^{-\lambda}$.

It might be interesting to note that a universal rate of convergence cannot be established without assuming regularity properties of the data distribution, such as smoothness via absolute continuity. Results of that kind are beyond the scope of this study. Instead, we refer the reader to [5, 6] for details on L_1 consistency and its connection to the rate of convergence.

6 Discussion

Regret optimization and stochastic approximation: Stochastic approximation algorithms can be applied for finding the optimum of (4) or, equivalently, to find the unique root of (8) based on noisy evaluations—the latter formulation is better suited for the classic version of the Robbins-Monro root finding algorithm [12]. These algorithms are iterative methods whose analysis focuses on the difference of $F(\tau_t)$ from $F(\tau^*)$, where τ_t denotes the estimate of τ^* in iteration t , whereas our online setting is concerned with the distance of $F((y_1, \dots, y_t), (\hat{y}_1, \dots, \hat{y}_t))$ from $F(\tau^*)$, where \hat{y}_i is the prediction for y_i in round i . This difference is crucial because $F(\tau_t)$ only depends on τ_t and in addition, if τ_t is close to τ^* then $F(\tau_t)$ is also close to $F(\tau^*)$ (see [19] for concentration properties), whereas in the online F-measure optimization setup, $F((y_1, \dots, y_t), (\hat{y}_1, \dots, \hat{y}_t))$ can be very different from $F(\tau^*)$ even if the current estimate τ_t is close to τ^* in case the number of previous incorrect predictions is large.

In *online learning* and *online optimization* it is common to work with the notion of (cumulative) regret. In our case, this notion could be interpreted either as $\sum_{i=1}^t |F((y_1, \dots, y_i), (\hat{y}_1, \dots, \hat{y}_i)) - F(\tau^*)|$ or as $\sum_{i=1}^t |y_i - \hat{y}_i|$. After division by t , the former becomes the average accuracy of the F-measure over time and the latter the accuracy of our predictions. The former is hard to interpret because $|F((y_1, \dots, y_i), (\hat{y}_1, \dots, \hat{y}_i)) - F(\tau^*)|$ itself is an aggregate measure of our performance

Table 1: Main statistics of the benchmark datasets and one pass F-scores obtained by OFO and 2S methods on various datasets. The bold numbers indicate when the difference is significant between the performance of OFO and 2S methods. The significance level is set to one sigma that is estimated based on the repetitions.

Dataset	#instances	#pos	#neg	#features	Learner:		Pegasos		Perceptron	
					OFO	2S	OFO	2S	OFO	2S
gisette	7000	3500	3500	5000	0.954	0.955	0.950	0.935	0.935	0.920
news20.bin	19996	9997	9999	1355191	0.879	0.876	0.879	0.883	0.908	0.930
Replab	45671	10797	34874	353754	0.924	0.923	0.926	0.928	0.914	0.914
WebspamUni	350000	212189	137811	254	0.912	0.918	0.914	0.910	0.927	0.912
epsilon	500000	249778	250222	2000	0.878	0.872	0.884	0.886	0.862	0.872
covtype	581012	297711	283301	54	0.761	0.762	0.754	0.760	0.732	0.719
url	2396130	792145	1603985	3231961	0.962	0.963	0.951	0.950	0.971	0.972
SUSY	5000000	2287827	2712173	18	0.762	0.762	0.754	0.745	0.710	0.720
kdda	8918054	7614730	1303324	20216830	0.927	0.926	0.921	0.926	0.913	0.927
kddb	20012498	17244034	2768464	29890095	0.934	0.934	0.930	0.929	0.923	0.928

over the first t rounds, which thus makes no sense to aggregate again. The latter, on the other hand, differs qualitatively from our ultimate goal; in fact, $|F((y_1, \dots, y_t), (\hat{y}_1, \dots, \hat{y}_t)) - F(\tau^*)|$ is the alternate measure that we are aiming to optimize for instead of the accuracy.

Online optimization of non-decomposable measures: Online optimization of the F-measure can be seen as a special case of optimizing non-decomposable loss functions as recently considered by [9]. Their framework essentially differs from ours in several points. First, regarding the data generation process, the adversarial setup with oblivious adversary is assumed, unlike our current study where a stochastic setup is assumed. From this point of view, their assumption is more general since the oblivious adversary captures the stochastic setup. Second, the set of classifiers is restricted to differentiable parametric functions, which may not include the F-measure maximizer. Therefore, their proof of vanishing regret does in general not imply convergence to the optimal F-score. Seen from this point of view, their result is weaker than our proof of consistency (i.e., convergence to the optimal F-measure in probability if the posterior estimates originate from a consistent learner). There are some other non-decomposable performance measures which are intensively used in many practical applications. Their optimization had already been investigated in the online or one-pass setup. The most notable such measure might be the area under the ROC curve (AUC) which had been investigated in an online learning framework by [21, 7].

7 Experiments

In this section, the performance of the OFO algorithm is evaluated in a one-pass learning scenario on benchmark datasets, and compared with the performance of the 2-stage F-measure maximization approach (2S) described in Section 2. We also assess the rate of convergence of the OFO algorithm in a pure online learning setup.³

The online learner \mathcal{A} in OFO was implemented in different ways, using Logistic Regression (LOGREG), the classical Perceptron algorithm (PERCEPTRON) [13] and an online linear SVM called PEGASOS [14]. In the case of LOGREG, we applied the algorithm introduced in [15] which handles L1 and L2 regularization. The hyperparameters of the methods and the validation procedures are described below and in more detail in Appendix D. If necessary, the raw outputs of the learners were turned into probability estimates, i.e., they were rescaled to $[0, 1]$ using logistic transform.

We used in the experiments nine datasets taken from the LibSVM repository of binary classification tasks.⁴ Many of these datasets are commonly used as benchmarks in information retrieval where the F-score is routinely applied for model selection. In addition, we also used the textual data released in the Replab challenge of identifying relevant tweets [1]. We generated the features used by the winner team [8]. The main statistics of the datasets are summarized in Table 1.

³Additional results of experiments conducted on synthetic data are presented in Appendix F.

⁴<http://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/binary.html>




Figure 1: Online F-scores obtained by OFO algorithm on various dataset. The dashed lines represent the one-pass performance of the OFO algorithm from Table 1 which we considered as baseline.

One-pass learning. In one-pass learning, the learner is allowed to read the training data only once, whence online learners are commonly used in this setting. We run OFO along with the three classifiers trained on 80% of the data. The learner obtained by OFO is of the form g_t^T , where t is the number of training samples. The rest 20% of the data was used to evaluate g_t^T in terms of the F-measure. We run every method on 10 randomly shuffled versions of the data and averaged results. The means of the F-scores computed on the test data are shown in Table 1. As a baseline, we applied the 2S approach. More concretely, we trained the same set of learners on 60% of the data and validated the threshold on 20% by optimizing (6). Since both approaches are consistent, the performance of OFO should be on par with the performance of 2S. This is confirmed by the results, in which significant differences are observed in only 7 of 30 cases. These differences in performance might be explained by the finiteness of the data. The advantage of our approach over 2S is that there is no need of validation and the data needs to be read only once, therefore it can be applied in a pure one-pass learning scenario. The hyperparameters of the learning methods are chosen based on the performance of 2S. We tuned the hyperparameters in a wide range of values which we report in Appendix D.

Online learning. The OFO algorithm has also been evaluated in the online learning scenario in terms of the online F-measure (2). The goal of this experiment is to assess the convergence rate of OFO. Since the optimal F-measure is not known for the datasets, we considered the test F-scores reported in Table 1. The results are plotted in Figure 1 for four benchmark datasets (the plots for the remaining datasets can be found in Appendix G). As can be seen, the online F-score converges to the test F-score obtained in one-pass evaluation in almost every case. There are some exceptions in the case of PEGASOS and PERCEPTRON. This might be explained by the fact that SVM-based methods as well as the PERCEPTRON tend to produce poor probability estimates in general (which is a main motivation for calibration methods turning output scores into valid probabilities [3]).

8 Conclusion and Future Work

This paper studied the problem of online F-measure optimization. Compared to many conventional online learning tasks, this is a specifically challenging problem, mainly because of the non-decomposable nature of the F-measure. We presented a simple algorithm that converges to the optimal F-score when the posterior estimates are provided by a sequence of classifiers whose L_1 error converges to zero as fast as $t^{-\lambda}$ for some $\lambda > 0$. As a key feature of our algorithm, we note that it is a purely online approach; moreover, unlike approaches such as 2S, there is no need for a hold-out validation set in batch mode. Our promising results from extensive experiments validate the empirical efficacy of our algorithm.

For future work, we plan to extend our online optimization algorithm to a broader family of complex performance measures which can be expressed as ratios of linear combinations of true positive, false positive, false negative and true negative rates [10]; the F-measure also belongs to this family. Moreover, going beyond consistency, we plan to analyze the rate of convergence of our OFO algorithm. This might be doable thanks to several nice properties of the function $h(\tau)$. Finally, an intriguing question is what can be said about the case when some bias is introduced because the classifier g_t does not converge to η .

Acknowledgments. Krzysztof Dembczyński is supported by the Polish National Science Centre under grant no. 2013/09/D/ST6/03917. The research leading to these results has received funding

from the European Research Council under the European Union’s Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement n. 306638.

References

- [1] E. Amigó, J. C. de Albornoz, I. Chugur, A. Corujo, J. Gonzalo, T. Martín-Wanton, E. Meij, M. de Rijke, and D. Spina. Overview of RepLab 2013: Evaluating online reputation monitoring systems. In *CLEF*, volume 8138, pages 333–352, 2013.
- [2] S. Bubeck and N. Cesa-Bianchi. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends in Machine Learning*, 5(1):1–122, 2012.
- [3] R. Busa-Fekete, B. Kégl, T. Éltető, and Gy. Szarvas. Tune and mix: Learning to rank using ensembles of calibrated multi-class classifiers. *Machine Learning*, 93(2–3):261–292, 2013.
- [4] N. Cesa-Bianchi and G. Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, 2006.
- [5] L. Devroye and L. Györfi. *Nonparametric Density Estimation: The L_1 View*. Wiley, NY, 1985.
- [6] L. Devroye, L. Györfi, and G. Lugosi. *A Probabilistic Theory of Pattern Recognition*. Springer, NY, 1996.
- [7] W. Gao, R. Jin, S. Zhu, and Z.-H. Zhou. One-pass AUC optimization. In *ICML*, volume 30:3, pages 906–914, 2013.
- [8] V. Hangya and R. Farkas. Filtering and polarity detection for reputation management on tweets. In *Working Notes of CLEF 2013 Evaluation Labs and Workshop*, 2013.
- [9] P. Kar, H. Narasimhan, and P. Jain. Online and stochastic gradient methods for non-decomposable loss functions. In *NIPS*, 2014.
- [10] N. Nagarajan, S. Koyejo, R. Ravikumar, and I. Dhillon. Consistent binary classification with generalized performance metrics. In *NIPS*, pages 2744–2752, 2014.
- [11] H. Narasimhan, R. Vaish, and Agarwal S. On the statistical consistency of plug-in classifiers for non-decomposable performance measures. In *NIPS*, 2014.
- [12] H. Robbins and S. Monro. A stochastic approximation method. *Ann. Math. Statist.*, 22(3):400–407, 1951.
- [13] F. Rosenblatt. The perceptron: A probabilistic model for information storage and organization in the brain. *Psychological Review*, 65(6):386–408, 1958.
- [14] S. Shalev-Shwartz, Y. Singer, and N. Srebro. Pegasos: Primal estimated sub-gradient solver for SVM. In *ICML*, pages 807–814, 2007.
- [15] Y. Tsuruoka, J. Tsujii, and S. Ananiadou. Stochastic gradient descent training for L1-regularized log-linear models with cumulative penalty. In *ACL*, pages 477–485, 2009.
- [16] C.J. van Rijsbergen. Foundation and evaluation. *Journal of Documentation*, 30(4):365–373, 1974.
- [17] S. R. S. Varadhan. *Probability Theory*. New York University, 2000.
- [18] W. Waegeman, K. Dembczyński, A. Jachnik, W. Cheng, and E. Hüllermeier. On the Bayes-optimality of F-measure maximizers. *Journal of Machine Learning Research*, 15(1):3333–3388, 2014.
- [19] N. Ye, K. M. A. Chai, W. S. Lee, and H. L. Chieu. Optimizing F-measure: A tale of two approaches. In *ICML*, 2012.
- [20] M. Zhao, N. Edakunni, A. Pocock, and G. Brown. Beyond Fano’s inequality: Bounds on the optimal F-score, BER, and cost-sensitive risk and their implications. *JMLR*, pages 1033–1090, 2013.
- [21] P. Zhao, S. C. H. Hoi, R. Jin, and T. Yang. Online AUC maximization. In *ICML*, pages 233–240, 2011.

Supplementary material for “Online F-Measure Optimization”

A Proposition 2

For reading convenience we restate Proposition 2.

Proposition 6. *If thresholds τ_t are defined according to (10) and \hat{y}_{t+1} as $\llbracket \eta(\mathbf{x}_{t+1}) > \tau_t \rrbracket$, then*

$$(\tau_{t+1} - \tau_t)b_{t+1} = h(\tau_t, y_{t+1}, \hat{y}_{t+1}) .$$

Proof. The proof is a simple analysis of the three cases $y_{t+1} = \hat{y}_{t+1} = 1$, $y_{t+1} \neq \hat{y}_{t+1}$ and $y_{t+1} = \hat{y}_{t+1} = 0$:

- If $y_{t+1} = \hat{y}_{t+1} = 1$, then we have $a_{t+1} = a_t + 1$ and $b_{t+1} = b_t + 2$, and so

$$\tau_{t+1} - \tau_t = \frac{a_t + 1}{b_t + 2} - \frac{a_t}{b_t} = \frac{b_t - 2a_t}{(b_t + 2)b_t} = \frac{b_t - 2a_t}{b_{t+1}b_t}$$

and

$$h(\tau_t, y_{t+1}, \hat{y}_{t+1}) = y_t \hat{y}_t - \tau_t(y_t + \hat{y}_t) = 1 - 2\tau_t = \frac{b_t - 2a_t}{b_t} = (\tau_{t+1} - \tau_t)b_{t+1}.$$

- If $y_{t+1} \neq \hat{y}_{t+1} = 1$, then we have $a_{t+1} = a_t$ and $b_{t+1} = b_t + 1$, and so

$$\tau_{t+1} - \tau_t = \frac{a_t}{b_t + 1} - \frac{a_t}{b_t} = -\frac{a_t}{(b_t + 1)b_t} = -\frac{a_t}{b_{t+1}b_t}$$

and

$$h(\tau_t, y_{t+1}, \hat{y}_{t+1}) = -\tau_t = -\frac{a_t}{b_t} = (\tau_{t+1} - \tau_t)b_{t+1}.$$

- If $y_{t+1} = \hat{y}_{t+1} = 0$, then we have $a_{t+1} = a_t$ and $b_{t+1} = b_t$, and so

$$\tau_{t+1} - \tau_t = \frac{a_t}{b_t} - \frac{a_t}{b_t} = 0,$$

and

$$h(\tau_t, y_{t+1}, \hat{y}_{t+1}) = 0 = (\tau_{t+1} - \tau_t)b_{t+1},$$

□

B Consistency with knowledge about the true posterior probabilities

For the proof of Theorem 3 we will need the following two lemmas. As a first step, one might decompose the update rule of the threshold τ_t given in (11). Based on Lemma 7, the difference sequence of the thresholds conditioned on the filtration can be rewritten such that the terms after the decomposition only depend on the data distribution, but not on the filtration.

Lemma 7.

$$\mathbf{E} \left[\frac{1}{b_{t+1}} \hat{h}(\tau_t) \middle| \mathcal{F}_t \right] = \frac{1}{b_t + 2} h(\tau_t) + \mathcal{O} \left(\frac{1}{b_t^2} \right)$$

where the $\mathcal{O}(\cdot)$ notation hides only universal constants.

Proof. Simple calculation yields that

$$\begin{aligned} \mathbf{E} \left[\frac{1}{b_{t+1}} \hat{h}(\tau_t) \middle| \mathcal{F}_t \right] &= \mathbf{E} \left[\frac{1}{b_t + 2} \hat{h}(\tau_t) \middle| \mathcal{F}_t \right] - \mathbf{E} \left[\left(\frac{1}{b_t + 2} - \frac{1}{b_{t+1}} \right) \hat{h}(\tau_t) \middle| \mathcal{F}_t \right] \\ &= \frac{1}{b_t + 2} h(\tau_t) - \mathbf{E} \left[\left(\frac{1}{b_t + 2} - \frac{1}{b_{t+1}} \right) \hat{h}(\tau_t) \middle| \mathcal{F}_t \right] \\ &= \frac{1}{b_t + 2} h(\tau_t) + \mathcal{O} \left(\frac{1}{b_t^2} \right) \end{aligned} \tag{12}$$

where (12) follows from the fact that

$$\frac{1}{b_t + 2} - \frac{1}{b_{t+1}} = \begin{cases} -\frac{2}{b_t^2 + 2b_t}, & \text{if } y_{t+1} = \hat{y}_{t+1} = 0 \\ 0, & \text{if } y_{t+1} = \hat{y}_{t+1} = 1 \\ -\frac{1}{b_t^2 + b_t - 2}, & \text{if } y_{t+1} \neq \hat{y}_{t+1} \end{cases}$$

and $\hat{h}(\tau) \geq -1$. This completes the proof. \square

In the update rule of the thresholds given in (11), the $1/b_{t+1}$ values controls the step size, and in fact, these values play a similar role like the coefficient in the stochastic approximation algorithm of [12] (the coefficients are denoted by a_t in the original paper). Therefore, the sequence of $1/b_{t+1}$ should behave similar to $1/t$, that is $\sum_{t=1}^{\infty} 1/b_t = \infty$ and $\sum_{t=1}^{\infty} 1/b_t^2 < \infty$. Since $b_t \leq 2t$, we have that $\sum_{t=1}^{\infty} 1/b_t \geq \sum_{t=1}^{\infty} 1/2t = \infty$. The infinite sum of the square of $1/b_{t+1}$ is analysed in the next lemma.

Lemma 8.

$$\mathbf{E} \left[\frac{1}{b_t^2} \right] = \mathcal{O} \left(\frac{1}{t^2} \right)$$

where the $\mathcal{O}(\cdot)$ notation hides only universal constants. Consequently, it holds that

$$\sum_{t=1}^{\infty} \mathbf{E} \left[\frac{1}{b_t^2} \right] < \infty$$

Proof. By applying the Chernoff bound for $c_t = \sum_{t'=1}^t y_{t'}$ with $t\pi_1/3$ as error term, we have

$$\mathbf{P} \left(c_t \leq \frac{2}{3} t\pi_1 \right) \leq \underbrace{\exp \left(\frac{-2t\pi_1^2}{9} \right)}_{\gamma_t}. \quad (13)$$

Therefore, based on (13), we can upper bound $\mathbf{E}[1/(c_t^2 + 1)]$ as

$$\mathbf{E} \left[\frac{1}{(c_t^2 + 1)} \right] \leq \gamma_t + (1 - \gamma_t) \frac{9}{4t^2\pi_1^2} \leq \gamma_t + \frac{9}{4t^2\pi_1^2} \quad (14)$$

Next, note that $1 \leq 1/b_t^2$ since $\tau_0 = 0$ which results in that $\hat{y}_1 = 1$. Moreover, $1/b_t^2 \leq 1/(c_t^2 + 1)$ since $b_t \geq (c_t + 1)$ which proves the first claim.

In order to proof the second claim, we might rewrite

$$\begin{aligned} \sum_{t=1}^{\infty} \mathbf{E} \left[\frac{1}{b_t^2} \right] &\leq \sum_{t=1}^{\infty} \mathbf{E} \left[\frac{1}{c_t^2 + 1} \right] \\ &\leq \sum_{t=1}^{\infty} \gamma_t + \frac{9}{\pi_1^2 4} \sum_{t=1}^{\infty} \frac{1}{t^2} < \infty \end{aligned} \quad (15)$$

where we obtained (15) by applying (14). Based on the definition γ_t , it readily follows that (15) is finite which completes the proof. \square

Before we turn to the proof of the Theorem 3, we show that function $h(\cdot)$ is not “too flat” anywhere. In general, function $h(\cdot)$ might be not differentiable everywhere (for example, it can be piecewise linear, or it might be discontinuous for some $\tau \in [0, 1]$). Nevertheless, its finite difference is strictly negative based on the next proposition.

Proposition 9. For any $\tau, \tau' \in [0, 1]$ such that $\tau \neq \tau'$, it holds

$$-1 - \pi_1 \leq \frac{h(\tau) - h(\tau')}{\tau - \tau'} \leq -\pi_1.$$

Proof. Based on (8), we can write $h(\tau)$ in the form of

$$h(\tau) = \int_0^1 \max(0, p - \tau) d\nu(p) - \tau\pi_1 .$$

Next, define $g(\tau, \tau')$ as

$$\begin{aligned} g(\tau, \tau') &= \frac{1}{\tau - \tau'} \int_0^1 \max(0, p - \tau) - \max(0, p - \tau') d\nu(p) \\ &= \int_0^1 \frac{\max(0, p - \tau) - \max(0, p - \tau')}{\tau - \tau'} d\nu(p) . \end{aligned}$$

Simple calculation yields that

$$\frac{h(\tau) - h(\tau')}{\tau - \tau'} = g(\tau, \tau') - \pi_1$$

Now there are two cases:

1. When $\tau < \tau'$, we have

$$\max(0, p - \tau) - \max(0, p - \tau') = \begin{cases} \tau' - \tau & \text{if } p > \tau' > \tau \\ p - \tau & \text{if } \tau' > p > \tau \\ 0 & \text{otherwise} \end{cases}$$

thus

$$\frac{\max(0, p - \tau) - \max(0, p - \tau')}{\tau - \tau'} = \begin{cases} \frac{\tau' - \tau}{\tau - \tau'} = -1 & \text{if } p > \tau' > \tau \\ -1 \leq \frac{p - \tau}{\tau - \tau'} \leq 0 & \text{if } \tau' > p > \tau \\ 0 & \text{otherwise} \end{cases}$$

And so $-1 \leq g(\tau, \tau') \leq 0$, which yields that $-1 - \pi_1 \leq \frac{h(\tau) - h(\tau')}{\tau - \tau'} \leq -\pi_1$ in this case.

2. When $\tau > \tau'$, we have

$$\max(0, p - \tau) - \max(0, p - \tau') = \begin{cases} \tau' - \tau & \text{if } p > \tau > \tau' \\ \tau' - p & \text{if } \tau > p > \tau' \\ 0 & \text{otherwise} \end{cases}$$

thus

$$\frac{\max(0, p - \tau) - \max(0, p - \tau')}{\tau - \tau'} = \begin{cases} \frac{\tau' - \tau}{\tau - \tau'} = -1 & \text{if } p > \tau > \tau' \\ -1 \leq \frac{\tau' - p}{\tau - \tau'} \leq 0 & \text{if } \tau > p > \tau' \\ 0 & \text{otherwise} \end{cases}$$

And similarly to the previous case, we have $-1 \leq g(\tau, \tau') \leq 0$, which yields that $-1 - \pi_1 \leq \frac{h(\tau) - h(\tau')}{\tau - \tau'} \leq -\pi_1$ in this case as well.

This completes the proof. □

The proof of Theorem 3 follows the proof of [12] given for the stochastic approximation method for root finding. However, our proof differs essentially from theirs, because of the lack of independence between $1/b_{t+1}$ and $\widehat{h}(\tau_t)$.

Proof. of Theorem 3

Our goal is to show that $\lim_{t \rightarrow \infty} \beta_t = 0$ where

$$\beta_t = \mathbf{E} [(\tau_t - \tau^*)^2]$$

which yields the convergence of the thresholds in probability, formally, $\tau_t \xrightarrow{P} \tau^*$, and consequently $F_{(t)} \xrightarrow{P} F(\tau^*)$, since $F(\tau^*) = 2\tau^*$ and, in addition, $F_{(t)} = 2\tau_t$ based on the definition of $F_{(t)}$ and τ_t given in (2) and (10), respectively.

We can decompose β_t as follows:

$$\begin{aligned} \beta_{t+1} &= \mathbf{E} [(\tau_{t+1} - \tau^*)^2] \\ &= \mathbf{E} [\mathbf{E} [(\tau_{t+1} - \tau^*)^2 | \mathcal{F}_t]] \\ &= \mathbf{E} \left[\mathbf{E} \left[\left(\tau_t - \tau^* + \frac{1}{b_{t+1}} \widehat{h}(\tau_t) \right)^2 \middle| \mathcal{F}_t \right] \right] \end{aligned} \quad (16)$$

$$\begin{aligned} &= \mathbf{E} [\mathbf{E} [(\tau_t - \tau^*)^2 | \mathcal{F}_t]] + 2\mathbf{E} \left[\mathbf{E} \left[(\tau_t - \tau^*) \frac{1}{b_{t+1}} \widehat{h}(\tau_t) \middle| \mathcal{F}_t \right] \right] + \mathbf{E} \left[\mathbf{E} \left[\left(\frac{\widehat{h}(\tau_t)}{b_{t+1}} \right)^2 \middle| \mathcal{F}_t \right] \right] \\ &= \beta_t + 2\mathbf{E} \left[\frac{\tau_t - \tau^*}{b_t + 2} h(\tau_t) + \mathcal{O} \left(\frac{1}{b_t^2} \right) \right] + \mathbf{E} \left[\mathbf{E} \left[\left(\frac{\widehat{h}(\tau_t)}{b_{t+1}} \right)^2 \middle| \mathcal{F}_t \right] \right] \end{aligned} \quad (17)$$

$$\begin{aligned} &= \beta_t - 2\mathbf{E} \left[\frac{\tau^* - \tau_t}{b_t + 2} h(\tau_t) \right] + \mathcal{O} \left(\mathbf{E} \left[\frac{1}{b_t^2} \right] \right) + \mathbf{E} \left[\mathbf{E} \left[\left(\frac{\widehat{h}(\tau_t)}{b_{t+1}} \right)^2 \middle| \mathcal{F}_t \right] \right] \\ &= \beta_t - 2\mathbf{E} \left[\frac{\tau^* - \tau_t}{b_t + 2} h(\tau_t) \right] + \mathcal{O} \left(\mathbf{E} \left[\frac{1}{b_t^2} \right] \right) \end{aligned} \quad (18)$$

where (16) follows from the update rule defined in (11), and (17) is obtained by applying Lemma 7. Finally, (18) follows from the fact that $\widehat{h}(\tau_t)$ is bounded and $1/(b_t + 2) \leq 1/b_{t+1} \leq 1/b_t$. Note that \mathcal{O} hides only universal constants, therefore $\mathbf{E}(\mathcal{O}(\cdot)) = \mathcal{O}(\mathbf{E}(\cdot))$. So the error term β_{t+1} can be written as

$$\beta_{t+1} = \beta_t - 2\delta_t + \mathcal{O} \left(\mathbf{E} \left[\frac{1}{b_t^2} \right] \right) \quad (19)$$

where

$$\delta_t = \mathbf{E} \left[\frac{\tau^* - \tau_t}{b_t + 2} h(\tau_t) \right] .$$

Summing up (19), we have

$$\beta_{t+1} = \beta_1 - 2 \sum_{t'=1}^t \delta_{t'} + \sum_{t'=1}^t \mathcal{O} \left(\mathbf{E} \left[\frac{1}{b_{t'}^2} \right] \right)$$

Hence $h(\tau)(\tau^* - \tau) > 0$, it holds that $\delta_t > 0$, thus we have

$$\begin{aligned} \sum_{t'=1}^t \delta_{t'} &\leq \frac{1}{2} \left[\beta_1 + \sum_{t'=1}^t \mathcal{O} \left(\mathbf{E} \left[\frac{1}{b_{t'}^2} \right] \right) \right] \\ &\leq \frac{1}{2} \left[\beta_1 + \sum_{t'=1}^{\infty} \mathcal{O} \left(\mathbf{E} \left[\frac{1}{b_{t'}^2} \right] \right) \right] \end{aligned} \quad (20)$$

Based on Lemma 8, $\sum_{t=1}^{\infty} \mathcal{O} \left(\mathbf{E} \left[\frac{1}{b_t^2} \right] \right)$ is finite. Therefore, $\sum_{t'=1}^{\infty} \delta_{t'}$ is also finite, and so there exists a $\beta \geq 0$, such that

$$\lim_{t \rightarrow \infty} \beta_t = \beta_1 - 2 \sum_{t'=1}^{\infty} \delta_{t'} + \sum_{t'=1}^{\infty} \mathcal{O} \left(\mathbf{E} \left[\frac{1}{b_{t'}^2} \right] \right) = \beta .$$

since all sums are finite above, and $b \geq 0$ based on (20).

As a next step, we show that $\beta = 0$. For doing this, first note that $\frac{1}{2(t+1)} \leq \frac{1}{b_{t+2}}$ for any $t > 0$, thus it holds

$$\sum_{t=1}^{\infty} \frac{1}{2(t+1)} \mathbf{E} [(\tau^* - \tau_t) h(\tau_t)] \leq \sum_{t=1}^{\infty} \delta_t < \infty$$

Now suppose that there exists a sequence $\{k_t\}$ of non-negative real numbers such that

$$k_t \beta_t \leq \mathbf{E}[(\tau^* - \tau_t)h(\tau_t)] \quad (21)$$

and

$$\sum_{t=1}^{\infty} \frac{k_t}{2(t+1)} = \infty. \quad (22)$$

From (21), we have

$$\begin{aligned} \sum_{t=1}^{\infty} \frac{k_t}{2(t+1)} \beta_t &\leq \sum_{t=1}^{\infty} \frac{1}{2(t+1)} \mathbf{E}[(\tau^* - \tau_t)h(\tau_t)] \\ &\leq \sum_{t=1}^{\infty} \delta_t < \infty. \end{aligned}$$

From (22), it follows that for any $\epsilon > 0$ there must exist infinitely many values such that $\beta_t < \epsilon$. Since we know that $\lim_{t \rightarrow \infty} \beta_t = b$ exists, thus $\beta = 0$ necessarily (if $\{k_n\}$ does exist).

The only thing remained to be shown is the existence of $\{k_n\}$ which satisfies (21) and (22). Based on Proposition 9, we have

$$-\pi_1 \leq \sup_{\tau \neq \tau^*} \left[\frac{h(\tau) - h(\tau^*)}{\tau - \tau^*} \right] = \sup_{\tau \neq \tau^*} \left[\frac{h(\tau)}{\tau - \tau^*} \right].$$

So we can lower bound $\mathbf{E}[(\tau^* - \tau_t)h(\tau_t)]$ as

$$\begin{aligned} \mathbf{E}[(\tau^* - \tau_t)h(\tau_t)] &= \mathbf{E} \left[(\tau^* - \tau_t)^2 \frac{h(\tau_t)}{\tau^* - \tau_t} \right] \\ &\geq \pi_1 \mathbf{E}[(\tau^* - \tau_t)^2] \\ &\geq \pi_1 \beta_t \end{aligned}$$

Therefore the constant sequence $k_t = \pi_1$ satisfies (21) and (22). This completes the proof. \square

C Consistency with estimated posterior probabilities

Assuming that the posterior estimate is provided as $p_{t+1} = g_t(\mathbf{x}_{t+1})$ in time step t , the update rule of the thresholds can be written in the form of

$$\tau_{t+1} - \tau_t = \frac{1}{b_{t+1}} h(\tau_t, y_{t+1}, \llbracket g_t(\mathbf{x}_{t+1}) \geq \tau_t \rrbracket). \quad (23)$$

Let us write concisely $\tilde{h}(\tau, g) = h(\tau, y, \llbracket g(\mathbf{x}) \geq \tau \rrbracket)$. As a first step, we can upper bound the difference between $\tilde{h}(\tau, g)$ and $\hat{h}(\tau_t)$ in terms of the L_1 error of g with respect to the data distribution as follows.

Lemma 10.

$$\left| \mathbf{E} \left[\tilde{h}(\tau_t, g_t) - \hat{h}(\tau_t) | \mathcal{F}_t \right] \right| \leq \int_{\mathbf{x} \in \mathcal{X}} |\eta(\mathbf{x}) - g_t(\mathbf{x})| d\mu(\mathbf{x}) = \|\eta - g_t\|_1$$

where $\tilde{h}(\tau, g) = h(\tau, y, \llbracket g(\mathbf{x}) \geq \tau \rrbracket) = h(\tau, y, g^\tau(\mathbf{x}))$.

Proof. First, compute

$$\begin{aligned} \tilde{h}(\tau, g) - \hat{h}(\tau) &= y \llbracket g(\mathbf{x}) \geq \tau \rrbracket - \tau (y + \llbracket g(\mathbf{x}) \geq \tau \rrbracket) - y \llbracket \eta(\mathbf{x}) \geq \tau \rrbracket + \tau (y + \llbracket \eta(\mathbf{x}) \geq \tau \rrbracket) \\ &= y \llbracket g(\mathbf{x}) \geq \tau \rrbracket - \tau \llbracket g(\mathbf{x}) \geq \tau \rrbracket - y \llbracket \eta(\mathbf{x}) \geq \tau \rrbracket + \tau \llbracket \eta(\mathbf{x}) \geq \tau \rrbracket \\ &= y (\llbracket g(\mathbf{x}) \geq \tau \rrbracket - \llbracket \eta(\mathbf{x}) \geq \tau \rrbracket) - \tau (\llbracket g(\mathbf{x}) \geq \tau \rrbracket - \llbracket \eta(\mathbf{x}) \geq \tau \rrbracket) \\ &= (y - \tau) (\llbracket g(\mathbf{x}) \geq \tau \rrbracket - \llbracket \eta(\mathbf{x}) \geq \tau \rrbracket) \end{aligned}$$

Now, note that

$$\llbracket g(\mathbf{x}) \geq \tau \rrbracket - \llbracket \eta(\mathbf{x}) \geq \tau \rrbracket = \begin{cases} 1 & \text{if } g(\mathbf{x}) > \tau > \eta(\mathbf{x}) \\ -1 & \text{if } \eta(\mathbf{x}) > \tau > g(\mathbf{x}) \end{cases}$$

therefore, we have

$$\begin{aligned} \left| \tilde{h}(\tau, g) - \hat{h}(\tau) \right| &= |(y - \tau) (\llbracket g(\mathbf{x}) \geq \tau \rrbracket - \llbracket \eta(\mathbf{x}) \geq \tau \rrbracket)| \\ &\leq |\eta(\mathbf{x}) - g(\mathbf{x})| . \end{aligned}$$

Consequently, we can upper bound the absolute value of

$$\mathbf{E} \left[\tilde{h}(\tau_t, g_t) - \hat{h}(\tau_t) | \mathcal{F}_t, \mathbf{x}_{t+1} \right] = (\eta(\mathbf{x}_{t+1}) - \tau_t) (\llbracket g(\mathbf{x}) \geq \tau \rrbracket - \llbracket \eta(\mathbf{x}) \geq \tau \rrbracket)$$

as follows:

$$\left| \mathbf{E} \left[\tilde{h}(\tau_t, g_t) - \hat{h}(\tau_t) | \mathcal{F}_t, \mathbf{x}_{t+1} \right] \right| \leq |\eta(\mathbf{x}_{t+1}) - g_t(\mathbf{x}_{t+1})| .$$

Therefore we can upper bound the following $\mathbf{E} \left[\tilde{h}(\tau_t, g_t) - \hat{h}(\tau_t) | \mathcal{F}_t \right]$ as

$$\begin{aligned} \left| \mathbf{E} \left[\tilde{h}(\tau_t, g_t) - \hat{h}(\tau_t) | \mathcal{F}_t \right] \right| &= \mathbf{E} \left[\left| \mathbf{E} \left[\tilde{h}(\tau_t, g_t) - \hat{h}(\tau_t) | \mathcal{F}_t, \mathbf{x}_{t+1} \right] \right| | \mathcal{F}_t \right] \\ &\leq \mathbf{E} \left[|\eta(\mathbf{x}) - g_t(\mathbf{x})| | \mathcal{F}_t \right] \\ &= \int_{\mathbf{x} \in \mathcal{X}} |\eta(\mathbf{x}) - g_t(\mathbf{x})| d\mu(\mathbf{x}) \\ &= \|\eta - g_t\|_1 \end{aligned}$$

□

We need to following results for decomposing the update rule given 10 similarly to Lemma 7.

Lemma 11.

$$\mathbf{E} \left[\frac{1}{b_{t+1}} (\tilde{h}(\tau_t, g_t) - \hat{h}(\tau_t)) | \mathcal{F}_t \right] = \frac{1}{b_t + 2} \|\eta - g_t\|_1 + \mathcal{O} \left(\frac{1}{b_t^2} \right)$$

where the $\mathcal{O}(\cdot)$ notation hides only universal constants.

Proof. Simple calculation yields that

$$\begin{aligned} \mathbf{E} \left[\frac{1}{b_{t+1}} (\tilde{h}(\tau_t, g_t) - \hat{h}(\tau_t)) | \mathcal{F}_t \right] &= \mathbf{E} \left[\frac{1}{b_t + 2} (\tilde{h}(\tau_t, g_t) - \hat{h}(\tau_t)) | \mathcal{F}_t \right] \\ &\quad - \mathbf{E} \left[\left(\frac{1}{b_t + 2} - \frac{1}{b_{t+1}} \right) (\tilde{h}(\tau_t, g_t) - \hat{h}(\tau_t)) | \mathcal{F}_t \right] \\ &\leq \frac{1}{b_t + 2} \|\eta - g_t\|_1 \\ &\quad - \mathbf{E} \left[\left(\frac{1}{b_t + 2} - \frac{1}{b_{t+1}} \right) (\tilde{h}(\tau_t, g_t) - \hat{h}(\tau_t)) | \mathcal{F}_t \right] \quad (24) \\ &= \frac{1}{b_t + 2} \|\eta - g_t\|_1 + \mathcal{O} \left(\frac{1}{b_t^2} \right) \quad (25) \end{aligned}$$

where (24) follows from Lemma 10 and (25) can be computed as (12) in the proof of Lemma 7. This completes the proof. □

The proof of Theorem 5 goes in a similar way like the proof of Theorem 3. The basic difference is that there is a new term in the decomposition of β_{t+1} which stems from the error of the classifiers which provides the posterior estimates. But this term can be upper bounded based on Lemma 10, and based on the fact that the L_1 error of the learner is vanishing as fast as $t^{-\lambda}$ where $1 > \lambda > 0$.

Proof. of Theorem 5

Similarly to the proof of Theorem 3, let us decompose β_{t+1} which now depend on $\tilde{h}(\cdot)$. We have

$$\begin{aligned}
\beta_{t+1} &= \mathbf{E} [(\tau_{t+1} - \tau^*)^2] \\
&= \mathbf{E} [\mathbf{E} [(\tau_{t+1} - \tau^*)^2 | \mathcal{F}_t]] \\
&= \mathbf{E} \left[\mathbf{E} \left[\left(\tau_t - \tau^* + \frac{1}{b_{t+1}} \tilde{h}(\tau_t) \right)^2 \middle| \mathcal{F}_t \right] \right] \\
&= \mathbf{E} [\mathbf{E} [(\tau_t - \tau^*)^2 | \mathcal{F}_t]]
\end{aligned} \tag{26}$$

$$\begin{aligned}
&\quad + 2\mathbf{E} \left[\mathbf{E} \left[(\tau_t - \tau^*) \frac{1}{b_{t+1}} \tilde{h}(\tau_t) \middle| \mathcal{F}_t \right] \right] + \mathbf{E} \left[\mathbf{E} \left[\left(\frac{\tilde{h}(\tau_t)}{b_{t+1}} \right)^2 \middle| \mathcal{F}_t \right] \right] \\
&= \beta_t + 2\mathbf{E} \left[\mathbf{E} \left[\frac{\tau_t - \tau^*}{b_{t+1}} \hat{h}(\tau_t) \middle| \mathcal{F}_t \right] \right] \\
&\quad + 2\mathbf{E} \left[\mathbf{E} \left[\frac{\tau_t - \tau^*}{b_{t+1}} (\tilde{h}(\tau_t) - \hat{h}(\tau_t)) \middle| \mathcal{F}_t \right] \right] + \mathbf{E} \left[\mathbf{E} \left[\left(\frac{\tilde{h}(\tau_t)}{b_{t+1}} \right)^2 \middle| \mathcal{F}_t \right] \right] \\
&= \beta_t + 2\mathbf{E} \left[\frac{\tau_t - \tau^*}{b_t + 2} h(\tau_t) \right] + 2\mathbf{E} \left[\mathbf{E} \left[\frac{\tau_t - \tau^*}{b_{t+1}} (\tilde{h}(\tau_t) - \hat{h}(\tau_t)) \middle| \mathcal{F}_t \right] \right] + \mathcal{O} \left(\mathbf{E} \left[\frac{1}{b_t^2} \right] \right)
\end{aligned} \tag{27}$$

$$\begin{aligned}
&= \beta_t - 2\mathbf{E} \left[\frac{\tau^* - \tau_t}{b_t + 2} h(\tau_t) \right] + \mathbf{E} \left[\frac{\tau_t - \tau^*}{b_t + 2} \|\eta - g_t\|_1 + \mathcal{O} \left(\frac{1}{b_t^2} \right) \right] + \mathcal{O} \left(\mathbf{E} \left[\frac{1}{b_t^2} \right] \right) \\
&= \beta_t - 2\mathbf{E} \left[\frac{\tau^* - \tau_t}{b_t + 2} h(\tau_t) \right] + \mathbf{E} \left[\frac{\tau_t - \tau^*}{b_t + 2} \|\eta - g_t\|_1 \right] + \mathcal{O} \left(\mathbf{E} \left[\frac{1}{b_t^2} \right] \right)
\end{aligned} \tag{28}$$

where (26) follows from the update rule defined in (23), and (27) follows from the fact that $\tilde{h}(\tau_t)$ is bounded and $1/(b_t + 2) \leq 1/b_{t+1} \leq 1/b_t$. Finally, Lemma 11 is applied in (28).

As a next step, we show that

$$\begin{aligned}
\sum_{t=1}^{\infty} \mathbf{E} \left[\frac{\tau_t - \tau^*}{b_t + 2} \|\eta - g_t\|_1 \right] &\leq \sum_{t=1}^{\infty} \mathbf{E} \left[\frac{\|\eta - g_t\|_1}{b_t + 2} \right] \\
&\leq \sum_{t=1}^{\infty} \sqrt{\mathbf{E} \left[\frac{1}{(b_t + 2)^2} \right]} \sqrt{\mathbf{E} [\|\eta - g_t\|_1^2]}
\end{aligned} \tag{29}$$

$$\leq \sum_{t=1}^{\infty} \mathcal{O} \left(\frac{1}{t} \right) \sqrt{\mathbf{E} [\|\eta - g_t\|_1^2]} \tag{30}$$

$$\leq \sum_{t=1}^{\infty} \mathcal{O} \left(\frac{1}{t} \right) \sqrt{\mathbf{E} [2\|\eta - g_t\|_1]} \tag{31}$$

$$= \sum_{t=1}^{\infty} \mathcal{O} \left(\frac{1}{t^{1+\lambda/2}} \right) < \infty . \tag{32}$$

where the Cauchy-Schwarz inequality is applied in (29), and (30) follows from Lemma 8 and, (31) follows from the fact that $\|\eta - g_t\|_1 \leq 2$. Finally, (32) follows from the assumption regarding the online learner, namely, there exists an $\lambda > 0$, such that

$$\mathbf{E} \left[\int_{\mathbf{x} \in \mathcal{X}} |\eta(\mathbf{x}) - g_t(\mathbf{x})| d\mu(\mathbf{x}) \right] = \mathcal{O}(t^{-\lambda})$$

Therefore, in this case, β_{t+1} can be decomposed in a similar way like in (19). Plugging this result into (20), the rest of the proof is analogous to the proof of Theorem 3. \square

D Hyperparameters of the online learners

The regularized version of the Logistic Regression was trained as it is introduced in [15]. This model has three hyperparameters the initial learning rate η , and the regularization coefficients for L1 and L2 penalties that are denoted by λ_1 and λ_2 respectively. The PEGASOS learning method has a single hyper-parameter λ , which controls the step size of the parameter update. Finally, the only hyperparameter of PERCEPTRON algorithm is the learning rate which we denote by γ . The hyperparameters and their values on which they were optimized are summarized in Table 2.

Table 2: The hyperparameters of the learning methods and their ranges used in hyperparmater optimization.

Learner	hyperparameter	Validation range
LR	η	{0.1, 0.01, 0.001, 0.0001}
	λ_1	{0.1, 0.01, 0.001, 0.0001}
	λ_2	{0.1, 0.01, 0.001, 0.0001}
PEGASOS	λ	{0.5, 0.4, 0.3, 0.2, 0.1, , 0.05, 0.03, 0.01, 0.005, 0.001, 0.0001}
PERCEPTRON	γ	{10.0, 5.0, 3.0, 1.0, 0.5, 0.3, 0.1, 0.05, 0.01, 0.005, 0.001, 0.0001}

E Optimal threshold τ^* with Beta and uniform distributions as $\nu(\cdot)$

Example 12. Assume that $\nu(\cdot)$ is uniform on $[0, 1]$. Then, according to (8), $h(\tau)$ is given as follows:

$$\begin{aligned} h(\tau) &= 1/2 - \tau^2/2 - \tau(1 - \tau) - \tau/2 \\ &= \tau^2 - 3\tau + 1 \end{aligned}$$

Setting $h(\tau) = 0$ and noting that the solution has to be in $[0, 1]$, the optimal threshold is $\tau^* = (3 - \sqrt{5})/2$, whence $F(\tau^*) = 3 - \sqrt{5} \approx 0.76$.

Example 13. Assume that $\nu(\cdot)$ is the Beta density function with parameters α and β : $\nu(p) = p^{\alpha-1}(1-p)^{\beta-1}/B(\alpha, \beta)$, where $B(\alpha, \beta) = \Gamma(\alpha + \beta)/(\Gamma(\alpha) + \Gamma(\beta))$ and $\Gamma(\cdot)$ is the Gamma function. The cumulative distribution function of the Beta distribution is denoted by $F(x; \alpha, \beta)$ and defined as $F(x; \alpha, \beta) = B(x; \alpha, \beta)/B(\alpha, \beta)$, where $B(x; \alpha, \beta) = \int_0^x t^{\alpha-1}(1-t)^{\beta-1} dt$. Thus, $h(\tau)$ can be written as follows:

$$\begin{aligned} h(\tau) &= \frac{1}{B(\alpha, \beta)} \int_{\tau}^1 p^{\alpha}(1-p)^{\beta-1} dp - \tau(1 - F(\tau; \alpha, \beta)) - \tau \frac{\alpha}{\alpha + \beta} \\ &= \frac{B(\alpha + 1, \beta)}{B(\alpha, \beta)} (1 - F(\tau; \alpha + 1, \beta)) - \tau(1 - F(\tau; \alpha, \beta)) - \tau \frac{\alpha}{\alpha + \beta} \\ &= \frac{\tau(1 - \tau)B(\alpha + 1, \beta)}{B(\alpha, \beta)} - \tau \frac{\alpha}{\alpha + \beta} \end{aligned}$$

Noting that $\tau^* = 0$ can be excluded, a simple calculation yields $\tau^* = 1 - \frac{\alpha}{\alpha + \beta} \frac{B(\alpha, \beta)}{B(\alpha + 1, \beta)}$.

F Synthetic Data

We sampled instance vectors from a two-dimensional Gaussian with mean $(1, 1)$ and the identity as covariance matrix. The class conditional probability for an instance $\mathbf{x} = (x_1, x_2)$ is computed by a logistic function as follows: $\eta(\mathbf{x}) = \mathbf{P}(Y = 1 | X = \mathbf{x}) = \frac{1}{1 + \exp(-x_1 + x_2 + 3)}$. Figure 2 (a) and (b) show a sample dataset consisting of 30000 instances and the contour plot of their posteriors probabilities, respectively.

We run the OFO algorithm with logistic regression without regularization as online learner on 30000 random examples \mathcal{D}_N in order to satisfy the required consistency condition. The initial learning rate was set to 0.01 and no regularization was applied here. Since the precise posteriors are known, the




Figure 2: Experiments with synthetic data. (a) The data points sampled from a two-dimensional normal distribution. The blue line indicates the Bayes optimal decision surface for which $\mathbf{P}(Y = 1 | X = \mathbf{x}) = 1/2$. The red points represent the positive instances. (b) The contour plot of the posterior probabilities $\mathbf{P}(Y = 1 | X = \mathbf{x})$. (c) The solid and dashed black lines show the F-measure and the validated threshold, respectively, in batch mode (note that $F(\tau^*) = 2\tau^*$). The red lines show the online F-score and the threshold found by OFO with logistic regression. Similarly, the solid lines represent the F-score and the dashed one the threshold.

validated threshold and F-measure were computed as a baseline. These are obtained by optimizing $F(\tau; \eta(\cdot), \mathcal{D}_N)$ as described in Section 2. The dashed black lines in Figure 2 (c) show the validated threshold (below) and the F-measure (above) found (note that $F(\tau^*) = 2\tau^*$). The online F-measure and thresholds found by OFO are plotted with solid and dashed red lines, respectively. As the results show, the online algorithm converges to the validation-based approach in terms of both F-measure and threshold, which in accordance with our theoretical results. All results are obtained by averaging over 100 repetitions.

G Online F-scores obtained by OFO algorithm on benchmark datasets




Figure 3: Online F-scores obtained by OFO algorithm with various online learners on benchmark datasets