

# Non-parametric Online AUC Maximization

Balázs Szörényi<sup>1,2(✉)</sup>, Snir Cohen<sup>1</sup>, and Shie Mannor<sup>1</sup>

<sup>1</sup> Technion, Haifa, Israel

szorenyi.balazs@gmail.com, snirc@cs.technion.ac.il, shie@ee.technion.ac.il

<sup>2</sup> Research Group on AI, Hungarian Academy of Sciences, University of Szeged, Szeged, Hungary

**Abstract.** We consider the problems of online and one-pass maximization of the area under the ROC curve (AUC). AUC maximization is hard even in the offline setting and thus solutions often make some compromises. Existing results for the online problem typically optimize for some proxy defined via surrogate losses instead of maximizing the real AUC. This approach is confirmed by results showing that the optimum of these proxies, over the set of all (measurable) functions, maximize the AUC. The problem is that—in order to meet the strong requirements for per round run time complexity—online methods typically work with restricted hypothesis classes and this, as we show, corrupts the above compatibility and causes the methods to converge to suboptimal solutions even in some simple stochastic cases. To remedy this, we propose a different approach and show that it leads to asymptotic optimality. Our theoretical claims and considerations are tested by experiments on real datasets, which provide empirical justification to them.

## 1 Introduction

The *area under the ROC curve* (AUC) [16] measures how well a mapping  $h$  of the instance space to the reals respects the partial order defined by some “ideal” score function  $s$ ; in the special case of bipartite ranking,  $s$  is simply a 0–1 valued function. As such, it has important applications in bioinformatics, information retrieval, anomaly detection, and many other areas.

Maximizing the AUC requires an approach different from maximizing the accuracy, even though there are some connections between the two [3, 5, 11]. Over the last decade, several approaches have been proposed and analyzed, guaranteeing consistency [9] and even optimal learning rates in some restricted cases [19]. Subsequently [22], followed by [13, 21], considered AUC maximization in an online setting, while [14] introduced a one-pass AUC maximization framework.

In this paper we first point out two important shortcomings of the existing methods proposed for online and one-pass AUC optimization:

- (A) None of them guarantees an optimal solution (not even asymptotically).
- (B) They all need to store the whole data. The reason for this is that they require parameter tuning, and thus also multiple passes over the data.

In contrast to (A), the  $k$  nearest neighbor method ( $k$ -NN), as we show, is guaranteed to converge to the optimum. This superiority of  $k$ -NN is also supported by the results of our empirical investigations. What is more, even though it clearly requires storing the whole data, it is not more demanding in terms of space complexity than previous algorithms, according to (B). Finally, one could argue that  $k$ -NN must perform poorly in terms of running time. This is not true, however: efficient solutions exist and, in fact, our experiments suggest that  $k$ -NN is competitive also in this regard. Additionally, dimensionality-originated issues can be taken care of using PCA or related methods.

The rest of the paper is structured as follows. First we introduce the formal framework and the definitions, then we show (A) formally, provide the theoretical justification for the  $k$ -NN method, present our experimental results, sum up the most important results from the literature, and finally we conclude with a short discussion.

## 2 Formal Setup

Given a set of  $n$  samples  $(x_1, y_1), \dots, (x_n, y_n) \in \mathcal{X} \times \mathcal{Y}$ , where  $\mathcal{X} \subseteq \mathbb{R}^d$  for some positive integer  $d$  and  $\mathcal{Y} = \{-1, 1\}$ , and given some mapping  $h : \mathcal{X} \rightarrow \mathbb{R}$ , the *area under the ROC curve* (AUC) [16] is the empirical mean

$$\text{AUC}(h; \mathcal{X}^+, \mathcal{X}^-) = \sum_{x^+ \in \mathcal{X}^+} \sum_{x^- \in \mathcal{X}^-} \left( \frac{\mathbb{I}[h(x^+) > h(x^-)]}{T_+ T_-} + \frac{\mathbb{I}[h(x^+) = h(x^-)]}{2T_+ T_-} \right),$$

where  $\mathcal{X}^+ = \{x_t : y_t = 1, 1 \leq t \leq n\}$ ,  $\mathcal{X}^- = \{x_t : y_t = -1, 1 \leq t \leq n\}$ ,  $T_+ = |\mathcal{X}^+|$ ,  $T_- = |\mathcal{X}^-|$ , and where  $\mathbb{I}[\cdot]$  denotes the indicator function; i.e.,  $\mathbb{I}[E] = 1$  when event  $E$  holds and  $\mathbb{I}[E] = 0$  otherwise. The regret hypothesis  $h$  with respect to some hypothesis set  $\mathcal{H} \subseteq \mathbb{R}^{\mathcal{X}}$  is defined as

$$\text{Regret}^{\mathcal{H}}(h; \mathcal{X}^+, \mathcal{X}^-) = \sup_{h' \in \mathcal{H}} \text{AUC}(h'; \mathcal{X}^+, \mathcal{X}^-) - \text{AUC}(h; \mathcal{X}^+, \mathcal{X}^-) \quad (1)$$

We also denote by  $\text{AUC}^*(\mathcal{X}^+, \mathcal{X}^-)$  the supremum of  $\text{AUC}(h'; \mathcal{X}^+, \mathcal{X}^-)$  over the set of all measurable functions  $h'$ , and introduce the notation

$$\text{Regret}(h; \mathcal{X}^+, \mathcal{X}^-) = \text{AUC}^*(\mathcal{X}^+, \mathcal{X}^-) - \text{AUC}(h; \mathcal{X}^+, \mathcal{X}^-).$$

Maximizing the AUC is also equivalent to minimizing the empirical risk

$$\text{Risk}(h; \mathcal{X}^+, \mathcal{X}^-) = 1 - \text{AUC}(h; \mathcal{X}^+, \mathcal{X}^-) = \sum_{t, t'=1}^n \frac{\ell^{\text{AUC}}(h; (x_t, y_t), (x_{t'}, y_{t'}))}{2T_+ T_-} \quad (2)$$

of the loss function

$$\ell^{\text{AUC}}(h; (x, y), (x', y')) = \mathbb{I}[(h(x) - h(x'))(y - y') < 0] + \mathbb{I}[h(x) = h(x'), y \neq y'].$$

One notorious problem with AUC is that it is non-convex and non-continuous, which makes it hard to work with. Especially in the online and

one-pass settings, where having a low (typically constant or logarithmic, but at least sublinear) per round run time complexity is essential. To resolve this issue, papers that aim for maximizing AUC online [13, 14, 21, 22], choose to replace  $\ell^{\text{AUC}}$  in (2) by some surrogate loss function  $\ell : \mathcal{H} \times (\mathcal{X} \times \mathcal{Y}) \rightarrow \mathbb{R}$ , and instead of maximizing the AUC, they minimize the surrogate risk

$$\text{Risk}^\ell(h; \mathcal{X}^+, \mathcal{X}^-) = \sum_{t, t'=1}^t \frac{\ell(h; (x_t, y_t), (x_{t'}, y_{t'}))}{2T_+ T_-},$$

and derive bounds for  $\text{Regret}^{\ell, \mathcal{H}}(h; \mathcal{X}^+, \mathcal{X}^-)$ , which is obtained by replacing AUC in (1) by  $1 - \text{Risk}^\ell$ .

If  $(x, y_1), (x_2, y_2), \dots$  are i.i.d. samples from some probability distribution  $\mathbf{P}$  over  $\mathcal{X} \times \mathcal{Y}$ , then one can define

$$\text{Risk}(h) = \mathbf{E} [\ell^{\text{AUC}}(h; (X, Y), (X', Y')) \mid Y > Y'] \quad (3)$$

and  $\text{AUC}(h) = 1 - \text{Risk}(h)$ . Similarly as above, replacing  $\ell^{\text{AUC}}$  in (3) by some other loss function  $\ell$  one obtains surrogate measures  $\text{Risk}^\ell(h)$  and  $\text{Regret}^{\ell, \mathcal{H}}(h)$ . Along the same analogy, we also use the notation  $\text{AUC}^*$  and  $\text{Regret}(h)$ .

### One-Pass and Online Setting

The one-pass and the online settings both have the same underlying protocol: in each round  $t$ , the learner proposes some hypothesis  $h_t : \mathcal{X} \rightarrow \mathbb{R}$  based on its previous experience, and then it observes the sample  $(x_t, y_t)$ . The two frameworks only differ in their objectives:

- In the *online setting* we are concerned with the evolution of the empirical AUC—that is, with

$$\text{AUC}_t = \text{AUC}(h_t; \mathcal{X}_t^+, \mathcal{X}_t^-)$$

for  $t = 1, \dots, n$ , where  $\mathcal{X}_t^+ = \{x_i : y_i = 1, 1 \leq i \leq t\}$  and  $\mathcal{X}_t^- = \{x_i : y_i = -1, 1 \leq i \leq t\}$ .

- In the *one-pass setting* the generalization ability of the learner is tested after the whole data had been processed. More precisely, the measure of performance is  $\text{AUC}(h_n)$ .

## 3 Surrogate Measures and Restricted Classes

Existing results for the online and one-pass AUC maximization problem optimize for some surrogate risk, instead of working with the AUC directly. In particular, many of them work with the square loss  $\ell_2$  (see Example 1). This approach is also confirmed by results showing *consistency* i.e., that  $\text{Regret}(h_t)$  converges to 0 whenever  $\text{Regret}^\ell(h_t)$  does for some sequence  $h_1, h_2, \dots$  of functions, as  $t$  goes to infinity. (See more about this in the section about the related work.)

These important results require, however, careful interpretation. And this is the starting point of our investigations: we claim that utilization of consistency is

only legitimate when the hypotheses class  $\mathcal{H}$  of our interest contains a global optimizer of the surrogate loss; that is, if  $\sup_h \text{AUC}^\ell(h) = \sup_{h \in \mathcal{H}} \text{AUC}^\ell(h)$ . When working with the set  $\mathcal{H}_{\text{lin}} = \{h^w(x) = w^\top x : w \in \mathbb{R}^d\}$  of linear functions—which is the case for *all* existing results for online and one-pass AUC maximization—this criterion is not fulfilled. Indeed, even though square loss is consistent (see [14]), in Example 1 below it holds that  $\text{AUC}(h') \ll \sup_{h \in \mathcal{H}_{\text{lin}}} \text{AUC}(h)$  for any  $h' \in \text{argmin}_{h \in \mathcal{H}_{\text{lin}}} \text{Risk}^{\ell_2}(h)$ .

The hinge loss  $\ell^\gamma(h; (x, y), (x', y')) = \mathbb{I}[y \neq y'] [\gamma - \frac{1}{2}(y - y')(h(x) - h(x'))]_+$  was also used in algorithmic solutions, but that does not even satisfy consistency [15].

### 3.1 Square Loss with Linear Hypotheses

This section presents the example that existing results can fail completely in maximizing the real AUC even in a simple case. This is demonstrated by the following example.

**Example 1.** Consider the setting when  $\mathcal{X} = \mathbb{R}^2$ , and  $\mathbf{P}[X = (-\epsilon, -1 + \epsilon)|Y = 1] = 1$  and  $\mathbf{P}[X = (0, -1 - \epsilon)|Y = -1] = \mathbf{P}[X = (0, 1)|Y = -1] = \mathbf{P}[X = (1, 0)|Y = -1] = 1/3$  for some small  $\epsilon > 0$ .

[14] first shows that the square loss  $\ell_2(h; (x, y), (x', y')) = (1 - \frac{y-y'}{2})(h(x) - h(x'))^2$  is consistent with AUC, and then use  $\ell_2$  as a surrogate loss to find the best linear score function  $h^w(x) = w^\top x$ . However, in the case above, AUC is maximized at  $h^{w^*}$  where  $w^* = (-1, 0)$ , where it actually takes value 1 (one has a very small freedom though, depending on the size of  $\epsilon$ ), and thus  $\text{Risk}(h^{w^*}) = 0$  for the corresponding linear function  $h^{w^*}$ . On the other hand, the surrogate measure for  $h^{w^*}$  is

$$\begin{aligned} \text{Risk}^{\ell_2}(h^{w^*}) &= \mathbf{E}[\ell_2(h^{w^*}; X, X')|Y > Y'] \\ &= \frac{1}{3}(1 + (w^*)^\top(\epsilon, 2\epsilon))^2 + \frac{1}{3}(1 + (w^*)^\top(\epsilon, 2 - \epsilon))^2 + \frac{1}{3}(1 + (w^*)^\top(1 + \epsilon, 1 - \epsilon))^2 \end{aligned}$$

which evaluates approximately to 2/3. At the same time, the actual optimum  $w'$  of this surrogate measure is around  $(-1/2, -1/2)$ , where it takes the value:

$$\text{Risk}^{\ell_2}(h^{w'}) = \mathbf{E}[\ell_2(h; X, X')|Y > Y'] \approx 1/3,$$

whereas  $\text{AUC}^*(h^{w'}) \approx 2/3$ , which is very far from the true optimum.

That is, when  $\mathcal{H}$  consists of the linear hypotheses (as is the case in [13, 14]), then

$$\text{Regret}^{\ell_2, \mathcal{H}}(h^{w'}) = 0,$$

implying that

$$\text{Regret}^{\mathcal{H}}(h^{w'}) = \text{Regret}(h^{(-1/2, -1/2)}) \approx 1/3.$$

Furthermore, adding a term  $\|w\|^2$  to regularize the surrogate measure does not change on this.

Note that this does not contradict the consistency of the square loss. The reason is that consistency requires Regret to vanish as  $\text{Regret}^\ell$  approaches 0, but as  $\text{Regret}^\ell$  is huge for all linear hypotheses, this sets no restrictions on how Regret should behave in this example.

### 4 Conditional Probability as Rank Function

In this section we start the investigation of finding alternative algorithmic solutions. With that in mind, we reach back to the fundamentals of AUC, and show that good estimates of the conditional probability function  $\eta(x) = \mathbf{P}[Y = 1|X = x]$  perform well at AUC maximization too.

The particular estimates that we consider here are of the form  $\hat{\eta} : \mathcal{X} \times \mathcal{Z} \rightarrow \mathbb{R}$  for some domain  $\mathcal{Z}$ . Here  $\mathcal{Z}$  is the domain of a variable that is used to encode prior information (e.g., random samples) and internal randomization (used e.g., for tie breaking) of the learner. (In accordance with that, in some cases it will be more convenient to use the notation  $\hat{\eta}_z$  for  $\hat{\eta}(\cdot, z)$ .) For example, given some series  $\{k_n\}_n$  of stepsizes, the  $k_n$ -NN estimate of [12] makes use of some i.i.d. samples  $U_1, \dots, U_n, U$  drawn from the uniform distribution over  $[0, 1]$ . Putting  $Z_n = (X_1, Y_1, \dots, X_n, Y_n, U_1, \dots, U_n, U)$ , their  $k_n$ -NN estimate maps an instance  $x$  to

$$\hat{\eta}_{Z_n}^{\text{DGKL}}(x) = \frac{1}{k_n} \sum_{i=1}^{k_n} Y_{\sigma(Z_n, x, i)}, \tag{4}$$

where  $\sigma(Z_n, x, \cdot)$  is the permutation for which  $(\|X_{\sigma(Z_n, x, 1)} - x\|, \|U_{\sigma(Z_n, x, 1)} - U\|), \dots, (\|X_{\sigma(Z_n, x, n)} - x\|, \|U_{\sigma(Z_n, x, n)} - U\|)$  is in lexicographic order.

Given such an estimate, we show the following result (for the proof see Appendix A).

**Theorem 2.** *Let  $Z$  be some random variable over some domain  $\mathcal{Z}$ , and let  $\hat{\eta} : \mathcal{X} \times \mathcal{Z} \rightarrow \mathbb{R}$  be an estimate of the conditional distribution function  $\eta(x) := \mathbf{E}[Y|X = x]$  as described above. Then  $\mathbf{E}_Z [\text{Regret}(\hat{\eta}_Z)] \leq \frac{3\sqrt{\epsilon}}{\mathbf{P}[Y=1]\mathbf{P}[Y=0]}$ , where  $\epsilon = \mathbf{E}_{X,Z} [|\hat{\eta}(X, Z) - \eta(X)|]$ .*

Similar result has also appeared in [1, 10, 19]. However, this particular estimator requires some small but essential differences in the analysis. Most importantly, kernel estimators are completely determined by the samples, whereas  $k_n$ -NN needs tie breaking. This requires additional randomness and complicates the analysis slightly.

### 5 AUC Maximization Using $k$ -NN

In the previous section we have shown guarantees for the AUC performance of estimators of the conditional probability function  $\eta$ . In this section we review some of the results on estimating  $\eta$  using  $k_n$ -NN, and show what they give combined with Theorem 2.

First of all, Devroye et al. [12] have shown that the  $k_n$ -NN version presented as Algorithm 1 converges under any distribution, assuming some standard restrictions on  $k_n$ .<sup>1</sup>

**Theorem 3 (Theorem 1 in [12]).** *If stepsize  $k_n$  satisfies  $\lim_{n \rightarrow \infty} k_n = \infty$  and  $\lim_{n \rightarrow \infty} k_n/n = 0$ , then  $\mathbf{E} [|\hat{\eta}_{Z_n}^{\text{DGKL}}(X) - \eta(X)|] \rightarrow 0$ , where  $\hat{\eta}_{Z_n}^{\text{DGKL}}$  is defined as in (4).*

Plugging this into Theorem 2 we immediately obtain the following result on  $\text{AUC}(\hat{\eta}_{Z_n}^{\text{DGKL}})$ .

**Corollary 4.** *Let  $\hat{\eta}_{Z_n}^{\text{DGKL}}$  be defined as in (4). Then  $\mathbf{E}_{Z_n} [\text{Regret}(\hat{\eta}_{Z_n}^{\text{DGKL}})] \rightarrow 0$  if  $\lim_{n \rightarrow \infty} k_n = \infty$  and  $\lim_{n \rightarrow \infty} k_n/n = 0$ , where  $\hat{\eta}_{Z_n}^{\text{DGKL}}$  is defined as in (4).*

KNNOAM is thus guaranteed to converge in case of i.i.d. samples.

One can, in fact, derive results also for the rate of convergence based on the work by Chaudhuri and Dasgupta [8]. This would not hold uniformly though, only for some restricted distributions.

---

**Algorithm 1.** KNNOAM( $\{k_t\}_t$ )

---

- 1: Draw a random sample  $U_1$  uniformly at random from  $(0, 1)$
  - 2: **for** round  $t = 1, \dots, n$  **do**
  - 3:   Observe sample  $(x_t, y_t)$
  - 4:   Draw a random sample  $U_{t+1}$  uniformly at random from  $(0, 1)$
  - 5:    $Z_t = (x_1, y_1, x_2, y_2, \dots, x_t, y_t, U_1, \dots, U_t, U_{t+1})$
  - 6:   Construct hypotheses  $h_t : \mathcal{X} \rightarrow \mathbb{R}$ , mapping  $x$  to  $h_t(x) = \hat{\eta}_{Z_t}^{\text{DGKL}}(x)$  {As in (4)}
  - 7: **end for**
- 

**Efficient Implementation**

An important feature of  $k$ -NN-methods is that it can be implemented efficiently. For example, the Cover Tree structure [6] makes it possible to insert a new instance into an existing cover tree or to remove an old one from it in time  $O(\log t)$ , and also to find the  $k_n$  nearest neighbor of some arbitrary point in time  $O(k_n \log t)$ .

**Choosing  $k$**

Choosing the right  $k$  for  $k$ -NN is a hard question.  $k > \log \log n$  is recommended for pointwise convergence (see Remark 1 in 4), but the common practice is to use  $\log n < k < n^{1/2}$ . One can also think about using  $k$  that changes with context; i.e., depends on the particular instance that is queried. See [4] for further details.

For a given dataset, one can also use cross validation or some Bayesian approach to find the best  $k$ . This was used by all the linear methods mentioned in the Introduction, but in a real online setting this is not applicable.

---

<sup>1</sup> They actually show an even stronger equivalence result.

## 6 Dimensionality Reduction

In general, any learning method can be applied that approximates  $\eta$  with arbitrary accuracy. Note, however, that all these methods suffer from dimensionality issues; including  $k$ -NN. One way to deal with it is to apply first dimensionality reduction methods. More specifically, the idea is to first feed the obtained sample into some online PCA algorithm (like SGA or CCIPCA—see [7] for a thorough discussion), and then use its output as input for  $k$ -NN, Parzen-Rosenblatt kernels, etc. This way one maintains the good AUC performance guaranteed by the learning algorithms but prevents dimensionality-originated run-time issues thanks to the guarantees of the PCA methods.

The property this task requires from the aforementioned techniques to preserve the good AUC performance guarantees of  $k$ -NN is a kind of stability. More precisely, denoting by  $\Phi_t$  the mapping they induce from the data from the first  $t$  rounds, it should fulfill the property

$$\|\Phi_t(x_t) - \Phi_T(x_t)\| \leq \epsilon(t) \quad \forall T \geq t$$

for some  $\epsilon(t)$  converging to 0 as  $t$  goes to infinity. Maintaining the convergence of  $k$ -NN does not seem possible otherwise.

## 7 Experimental Results

In this section, we evaluate the empirical performance of the proposed KNN Online AUC Maximization (KNNOAM) algorithm on benchmark datasets.

### Compared Algorithms

We compare the proposed KNNOAM algorithm with state-of-the-art online AUC optimization algorithms. Specifically, the compared algorithms in our experiments include:

- **OAM<sub>seq</sub>**: the OAM algorithm with reservoir sampling and sequential updating [22];
- **OAM<sub>gra</sub>**: the OAM algorithm with reservoir sampling and online gradient updating method [22];
- **OPAUC**: the one-pass algorithm AUC optimization algorithm proposed in [14];
- **AdaOAM**: the adaptive gradient AUC optimization algorithm proposed in [13];
- **KNNOAM**: our proposed KNN based algorithm.

### General Experimental Setup

We conduct our experiments on sixteen benchmark datasets that have been used in previous studies on AUC optimization. The details of the datasets are summarized in Table 1. All these datasets can be downloaded from LIBSVM<sup>2</sup>

<sup>2</sup> <https://www.csie.ntu.edu.tw/~cjlin/libsvmtools/>.

and UCI Machine Learning Repository<sup>3</sup>. Note that half of the datasets (segment, satimage, vowel, letter, poker, usps, connect-4, acoustic and vehicle) are originally multi-class. These multi-class datasets have been converted into class-imbalanced binary datasets by choosing one class, setting its label to +1 and the rest to -1. This class has been chosen so that the ratio  $T_-/T_+$  is below 50 and its cardinality has been minimized. (The ratio is kept below 50 to obtain conclusive results.) In case two or more classes have the same size, one has been chosen randomly. Previous studies use similar conversion methods. In addition, the features have been rescaled linearly to  $[-1, 1]$  for all datasets. All experiments are performed with Matlab on a computer workstation with 3.40 GHz CPU and 32 GB memory.

**Table 1.** Details of the benchmark datasets used in the experiments.  $T_+ = |\mathcal{X}^+|$  and  $T_- = |\mathcal{X}^-|$ .

Datasets	# instances	# features	$T_-/T_+$	Datasets	# instances	# features	$T_-/T_+$
fourclass	862	2	1.8078	segment	2310	19	6.0000
svmguide1	3089	4	1.8365	ijcnn1	141691	22	9.4453
magic04	19020	10	1.8439	connect-4	67557	126	9.4756
german	1000	24	2.3333	satimage	4435	36	9.6867
a9a	32561	123	3.1527	vowel	528	10	10.0000
svmguide3	1243	22	3.1993	usps	9298	256	12.1328
vehicle	846	18	3.2513	letter	20000	16	25.7380
acoustic	78823	50	3.3028	poker	25010	10	47.7524

### Experimental Setup of the Online Setting

Our main goal is to compare the performance of the above algorithms in the online setting. In this setting, each algorithm receives a random sample from the dataset, suffers loss and updates its classifier according to this sample. Note that our proposed KNNOAM does not require any parameter tuning, as opposed to the other four existing algorithms. Clearly, parameter tuning requires multiple passes over the data, which is inconsistent with this setting. Although KNNOAM must store at time  $t$  all the samples up to time  $t - 1$ , the rest of the algorithms must store all of the samples in advance for the parameter tuning. Despite the fact that it gives the rest of the algorithms an unfair advantage, we follow the parameter tuning procedures for each algorithm suggested in [13, 14, 22] and use the best obtained parameters for each dataset and algorithm. We apply five-fold cross-validation on the training set to find the best learning rate  $\eta \in 2^{[-10, 10]}$  and the regularization parameter  $\lambda \in 2^{[-10, 2]}$  for both OPAUC and AdaOAM. For OAM<sub>seq</sub> and OAM<sub>gra</sub>, we apply five-fold cross-validation to tune the penalty parameter  $C \in 2^{[-10, 10]}$ , and fix the buffer at 100 as suggested in [22]. For

<sup>3</sup> [www.ics.uci.edu/~mllearn/MLRepository.html](http://www.ics.uci.edu/~mllearn/MLRepository.html).



KNNOAM we only had to choose  $k_n$  that goes to infinity and is of order  $o(n)$ . We choose  $k_n = 2 \log_2 n$ .<sup>4</sup> For every dataset, we average over 20 runs and plot a graph showing the experimental AUC loss as a function of the number of samples. This gives us an opportunity to examine the evolution of the performance of the algorithms. The results of this experiment are presented in Appendix B. The code is available at <https://bitbucket.org/smir/auc2017>.

### Experimental Setup of the One-Pass Setting

We also compare the performance of the algorithm in the one-pass setting. We follow the experimental setup suggested in previous studies [13, 14, 22]. The performance of the compared algorithms is evaluated by four trials of five-fold cross-validation using the parameters received by the parameter tuning procedure explained in the previous section about the online experimental setup. The AUC values are the average of these 20 runs. The results are summarized in Fig. 1.

### 7.1 Evaluation on Benchmark Datasets

In the online setting, KNNOAM outperforms the four other, state-of-the-art online AUC algorithms considered in our experiments in 12 out of 16 datasets. What is more, in 10 out of those 12 datasets, the improvement by KNNOAM is significant. For example, in `ijcnn1` and `acoustic` datasets, KNNOAM converges to a much lower empirical AUC loss than the rest of the algorithms. These are outstanding results for this setting, especially when recalling that the compared algorithms have been tuned before running the experiments. Our graphs demonstrate that KNNOAM has a much smoother convergence. The rest of the algorithms' performance is unstable: after receiving new samples it may improve but might also significantly deteriorate. This is a highly undesirable property.

In the `german`, `svmguide3` and `vehicle` datasets, KNNOAM does not outperform all the compared algorithms. However, these datasets are very small and therefore the results are inconclusive.

In the `a9a` dataset, OPAUC and AdaOAM outperform KNNOAM. Although this dataset cannot be considered small, KNNOAM still behaves as if it were, because the samples are sparse and the dimensionality is relatively big. To strengthen this claim, we can examine the performance of KNNOAM on `connect-4` dataset. This dataset has roughly the same dimensionality as the `a9a` dataset and the samples are as sparse, but has more than twice the samples, and KNNOAM outperforms the rest of the algorithms significantly.

Many of the graphs show that  $OAM_{seq}$ ,  $OAM_{gra}$  are not guaranteed to converge at all.<sup>5</sup> Some of them also show that OPAUC and AdaOAM does not converge smoothly. This behavior is observed when the learning rate  $\eta$  is too high and a single sample could change the classifier drastically and cause a significant

<sup>4</sup> This choice guarantees low time complexity and also turned out to result in a competitive method. (We did try other choices; the results were similar.).

<sup>5</sup> It should be mentioned that the theoretical guarantees of OAM is also doubtful, according to [17].

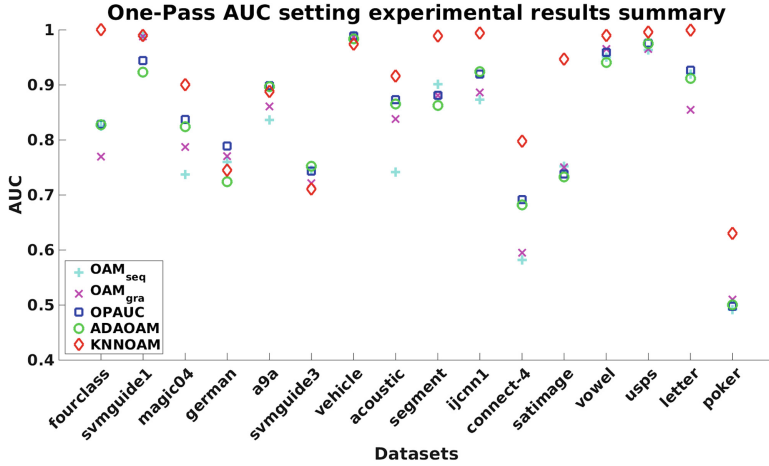


Fig. 1. Experimental results comparing the one-pass AUC performance summary

local deterioration of the performance as can be seen from the performance of OPAUC on svmguide1 dataset for example. On the other hand, if the chosen learning rate  $\eta$  is too small, the performance might be poor, as can be seen from the performance of AdaOAM on the same dataset.

The results for the one-pass setting are similar to the ones for the online setting: KNNOAM is better, or at least competitive in comparison with its competitors. For this setting we also present the running times (see Appendix C) KNNOAM is usually the fastest algorithm and is never the slowest among the compared algorithms.

## 8 Related Work

### Reducing Ranking to Binary Classification

[3, 5] considers the problem of ranking a finite random sample, and reduced this to a related binary classification problem. In particular, they bound the risk of the ranking (which is closely related to AUC of the ranker over this sample) in terms of the classification performance. However, the risk and the classification performance they use is only representative to that particular sample, and they do not consider prediction or generalization bounds. ([5] comment on the case when the rankings are drawn from some distribution, but does not imply any result in our setting.)

### Offline Algorithmic Solutions

[9] aims to minimize (3), and actually obtain asymptotic optimality. Their algorithm constructs a tree in an iterative fashion by solving subsequent challenging optimization problems.

[19] uses kernel estimates of the conditional probability function  $\eta(x) = \mathbf{E}[Y|X = x]$  based on the Parzen-Rosenblatt kernel, and have shown fast and

superfast convergence rates assuming Tsybakov-style noise. They also complement their results by showing lower bounds on the best convergence rate in some situations.

### Consistency of Surrogate Measures

The investigation of consistency with respect to AUC was initiated by [18] showing consistency of a balanced version of the exponential and the logistic loss. Later on [1, 14, 15, 20] investigated the consistency of other loss functions like the exponential, logistic, squared, and  $q$ -normed hinge loss, and variants of them. Finally, [15] shows that hinge loss is not consistent.

### One-Pass and Online Solutions

As mentioned, [22] was the first to analyze online AUC maximization. They have defined the setup, presented algorithmic solutions optimizing for the hinge loss, and provided regret bounds. [21] also uses the hinge loss with a perceptron-like algorithm which, in round  $t$ , achieves regret  $O(1/\sqrt{t})$ . [17] works with the same setting as [21], and achieves several improvements in terms of different parameters.

[14] uses square loss in the one-pass setting, and obtains a convergence rate of order  $O(1/t)$  for the linearly separable case and  $O(\sqrt{1/t})$  for the general one. [13] obtains similar results, but using the Adaptive Gradient method.

All these papers work with the set of linear hypothesis.

### Uniform Convergence Bounds

Uniform convergence bounds like the ones in [2] show how fast the empirical AUC-risk converges to the actual AUC over a given class of hypothesis. Consequently, they do not provide any practical guidance on how one could acquire some hypothesis with small risk.

## 9 Concluding Remarks

We have shown that existing methodology for maximizing AUC in an online or one-pass setting can fail already in very simple situations. To remedy this, we have proposed to reach back to the fundamentals of AUC, and suggested an algorithmic solution based on the celebrated  $k$ -NN-estimate of the conditional probability function. This has guarantees in the stochastic setting, has efficient implementations, and outperforms previous methods on several real datasets. The latter is even more surprising in view of the fact that, unlike its competitors, it requires no parameter tuning.

Nevertheless, we feel that this should not be considered as an ultimate solution, but rather as an encouragement for future research to explore further alternative solutions. To mention a few:

- Combining KNNOAM with metric learning arises naturally, and could extend its applicability to more exotic domains.
- Maximizing the objective function  $AUC_n$  in the adversarial setting is another important question, which existing results do not tell anything about.

**Acknowledgements.** This research was supported in part by the European Communities Seventh Framework Programme (FP7/2007-2013) under grant agreement 306638 (SUPREL).

### A Proof of Theorem 2

Let us first introduce the notation  $\mathcal{X}_z = \{x : |\widehat{\eta}(x, z) - \eta(x)| < \sqrt{\epsilon}\}$  for  $z \in \mathcal{X}$  and define for  $h : \mathcal{X} \rightarrow \mathbb{R}$  measurable and  $x, x' \in \mathcal{X}$

$$a(h, x, x') = \frac{1}{2}\mathbb{I}[h(x) = h(x')] [\eta(x)(1 - \eta(x'))] + \mathbb{I}[h(x) > h(x')] [\eta(x)(1 - \eta(x'))]$$

and its symmetrization

$$b(h, x, x') = \frac{1}{2}a(h, x, x') + \frac{1}{2}a(h, x', x).$$

It then holds that

$$\begin{aligned} & \mathbf{E}_{(X,Y)} \mathbf{E}_{(X',Y')} \left[ \frac{1}{2} \mathbb{I}[h(X) > h(X')] \mathbb{I}[Y = 1, Y' = 0] \right. \\ & \quad \left. + \mathbb{I}[h(X) > h(X')] \mathbb{I}[Y = 1, Y' = 0] \right] \\ &= \mathbf{E}_{X,X'} \left[ \frac{1}{2} \mathbb{I}[h(X) = h(X')] \mathbf{E}_{Y,Y'} [\mathbb{I}[Y = 1, Y' = 0]] \right. \\ & \quad \left. + \mathbb{I}[h(X) > h(X')] \mathbf{E}_{Y,Y'} [\mathbb{I}[Y = 1, Y' = 0]] \right] \\ &= \mathbf{E}_{X,X'} [a(h, X, X')] \\ &= \mathbf{E}_{X,X'} [b(h, X, X')], \end{aligned}$$

where the last equation follows because  $X$  and  $X'$  are i.i.d. This then gives

$$\text{AUC}(h) = \frac{\mathbf{E}_{X,X'} [b(h, X, X')]}{\mathbf{P}[Y = 1] \mathbf{P}[Y = 0]}. \tag{5}$$

Now, note that  $x, x' \in \mathcal{X}_z$  implies  $|\eta(x)(1 - \eta(x')) - \eta(x')(1 - \eta(x))| \leq 2\sqrt{\epsilon}$  because of the  $\alpha\beta - \alpha'\beta' = (\alpha - \alpha')\beta + \alpha'(\beta - \beta')$  equality. Combining this with the fact that  $\mathbb{I}[h(x) = h(x')] + \mathbb{I}[h(x) < h(x')] + \mathbb{I}[h(x) > h(x')] = 1$  for any  $h : \mathcal{X} \rightarrow \mathbb{R}$  and any  $x, x' \in \mathcal{X}$ , it follows that

$$b(\eta, x, x') - b(\widehat{\eta}_z, x, x') \leq \sqrt{\epsilon}, \quad z \in \mathcal{Z}, \quad x, x' \in \mathcal{X}_z.$$

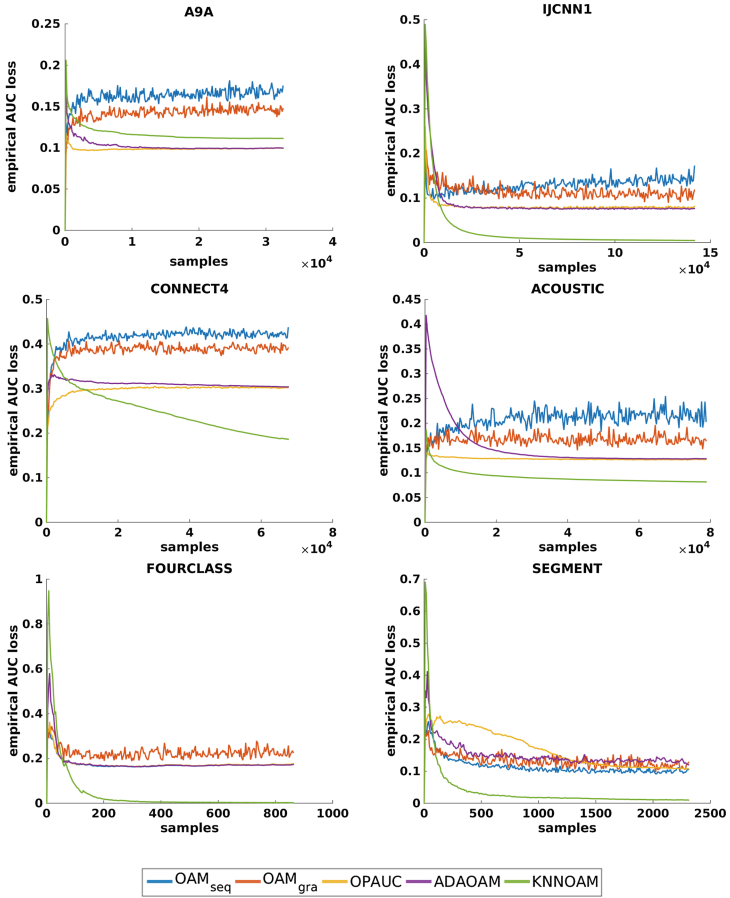
Accordingly,

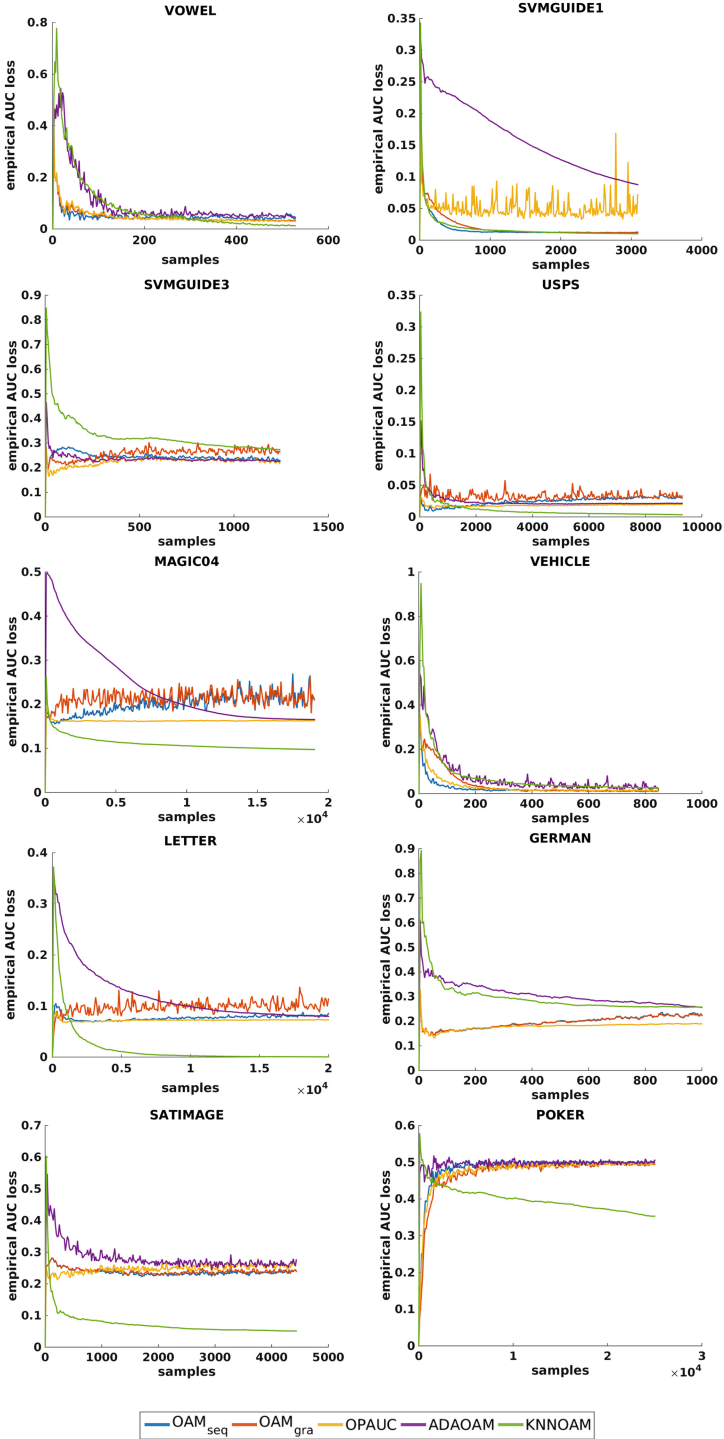
$$\begin{aligned} & \mathbf{E}_{X,X'} [b(\eta, X, X')] - \mathbf{E}_{X,X',Z} [b(\widehat{\eta}_Z, X, X')] \\ & \quad \leq \sqrt{\epsilon} + \mathbf{P}_{X,X',Z} [X \notin \mathcal{X}_Z \text{ or } X' \notin \mathcal{X}_Z] \\ & \quad \leq 3\sqrt{\epsilon}, \end{aligned} \tag{6}$$

where the last inequality is true because  $X$  and  $X'$  are i.i.d. and because  $\mathbf{P}_{X,Z} [X \notin \mathcal{X}_Z] = \mathbf{P}_{X,Z} [|\widehat{\eta}(X, Z) - \eta(X)| \geq \sqrt{\epsilon}] \leq \sqrt{\epsilon}$  by the definition of  $\epsilon$ .

Finally, according to [10],  $\text{AUC}(h)$  is maximized when  $h = \eta$ . The theorem thus follows by combining (5) and (6).

## B Figures for Benchmark Datasets





## C Experimental Results for the One-Pass AUC Setup

See Table 2.

**Table 2.** Experimental results comparing the one-pass AUC performance (i.e.,  $AUC(h_n)$ ) of  $OAM_{seq}$ ,  $OAM_{gra}$ , OPAUC, ADAOM and KNNOAM

Algorithm	fourclass		Algorithm	segment	
	AUC	Time (s)		AUC	Time (s)
$OAM_{seq}$	$0.8253 \pm 0.0290$	0.5125	$OAM_{seq}$	$0.9010 \pm 0.0205$	1.4031
$OAM_{gra}$	$0.7694 \pm 0.0518$	0.1797	$OAM_{gra}$	$0.8811 \pm 0.0262$	0.4719
OPAUC	$0.8281 \pm 0.0308$	0.0422	OPAUC	$0.8808 \pm 0.0237$	0.1289
AdaOAM	$0.8278 \pm 0.0269$	0.0641	AdaOAM	$0.8629 \pm 0.0337$	0.1922
KNNOAM	<b><math>1.0000 \pm 0.0000</math></b>	<b>0.0070</b>	KNNOAM	<b><math>0.9891 \pm 0.0087</math></b>	<b>0.0227</b>
Algorithm	svmguidel		Algorithm	ijcnn1	
	AUC	Time (s)		AUC	Time (s)
$OAM_{seq}$	$0.9885 \pm 0.0026$	2.0953	$OAM_{seq}$	$0.8733 \pm 0.0253$	101.9695
$OAM_{gra}$	$0.9884 \pm 0.0039$	0.6922	$OAM_{gra}$	$0.8862 \pm 0.0345$	34.1859
OPAUC	$0.9440 \pm 0.0175$	0.1516	OPAUC	$0.9201 \pm 0.0080$	<b>10.3078</b>
AdaOAM	$0.9229 \pm 0.0338$	0.2336	AdaOAM	$0.9237 \pm 0.0036$	14.2531
KNNOAM	<b><math>0.9900 \pm 0.0041</math></b>	<b>0.0156</b>	KNNOAM	<b><math>0.9945 \pm 0.0005</math></b>	66.9656
Algorithm	magic04		Algorithm	connect-4	
	AUC	Time (s)		AUC	Time (s)
$OAM_{seq}$	$0.7370 \pm 0.0910$	13.4914	$OAM_{seq}$	$0.5817 \pm 0.0205$	50.7984
$OAM_{gra}$	$0.7875 \pm 0.0530$	4.6008	$OAM_{gra}$	$0.5948 \pm 0.0232$	<b>17.4992</b>
OPAUC	$0.8372 \pm 0.0063$	1.1969	OPAUC	$0.6918 \pm 0.0077$	117.2625
AdaOAM	$0.8240 \pm 0.0085$	1.7000	AdaOAM	$0.6824 \pm 0.0166$	144.8133
KNNOAM	<b><math>0.9003 \pm 0.0061</math></b>	<b>0.3875</b>	KNNOAM	<b><math>0.7977 \pm 0.0053</math></b>	77.9141
Algorithm	german		Algorithm	satimage	
	AUC	Time (s)		AUC	Time (s)
$OAM_{seq}$	$0.7603 \pm 0.0395$	0.6086	$OAM_{seq}$	$0.7520 \pm 0.0250$	2.9531
$OAM_{gra}$	$0.7704 \pm 0.0368$	0.2172	$OAM_{gra}$	$0.7504 \pm 0.0262$	1.1398
OPAUC	<b><math>0.7890 \pm 0.0274</math></b>	0.0578	OPAUC	$0.7380 \pm 0.0166$	0.4508
AdaOAM	$0.7238 \pm 0.0691$	0.0859	AdaOAM	$0.7328 \pm 0.0267$	0.5719
KNNOAM	$0.7454 \pm 0.0268$	<b>0.0086</b>	KNNOAM	<b><math>0.9466 \pm 0.0108</math></b>	<b>0.1172</b>
Algorithm	a9a		Algorithm	vowel	
	AUC	Time (s)		AUC	Time (s)
$OAM_{seq}$	$0.8368 \pm 0.0199$	25.3867	$OAM_{seq}$	$0.9503 \pm 0.0329$	0.1195
$OAM_{gra}$	$0.8611 \pm 0.0108$	<b>8.3141</b>	$OAM_{gra}$	$0.9654 \pm 0.0157$	0.0547
OPAUC	<b><math>0.8989 \pm 0.0036</math></b>	54.7289	OPAUC	$0.9592 \pm 0.0245$	0.0273
AdaOAM	$0.8967 \pm 0.0036$	67.2898	AdaOAM	$0.9410 \pm 0.0290$	0.0422
KNNOAM	$0.8879 \pm 0.0040$	17.4187	KNNOAM	<b><math>0.9895 \pm 0.0092</math></b>	<b>0.0047</b>
Algorithm	svmguidel3		Algorithm	usps	
	AUC	Time (s)		AUC	Time (s)
$OAM_{seq}$	$0.7453 \pm 0.0410$	0.7477	$OAM_{seq}$	$0.9634 \pm 0.0121$	7.3531
$OAM_{gra}$	$0.7217 \pm 0.0428$	0.2703	$OAM_{gra}$	$0.9650 \pm 0.0097$	<b>2.3188</b>
OPAUC	$0.7433 \pm 0.0384$	0.0703	OPAUC	$0.9759 \pm 0.0066$	35.9055
AdaOAM	<b><math>0.7519 \pm 0.0442</math></b>	0.1047	AdaOAM	$0.9752 \pm 0.0081$	43.8656
NNOAM	$0.7112 \pm 0.0339$	<b>0.0102</b>	KNNOAM	<b><math>0.9957 \pm 0.0029</math></b>	3.0055
Algorithm	vehicle		Algorithm	letter	
	AUC	Time (s)		AUC	Time (s)
$OAM_{seq}$	<b><math>0.9898 \pm 0.0067</math></b>	0.4680	$OAM_{seq}$	$0.9200 \pm 0.0088$	13.1813
$OAM_{gra}$	$0.9872 \pm 0.0084$	0.1617	$OAM_{gra}$	$0.8548 \pm 0.0554$	4.3805
OPAUC	$0.9886 \pm 0.0060$	0.0469	OPAUC	$0.9267 \pm 0.0062$	1.2367
AdaOAM	$0.9838 \pm 0.0101$	0.0711	AdaOAM	$0.9120 \pm 0.0315$	1.7828
KNNOAM	$0.9745 \pm 0.0074$	<b>0.0063</b>	KNNOAM	<b><math>0.9993 \pm 0.0003</math></b>	<b>1.0063</b>
Algorithm	acoustic		Algorithm	poker	
	AUC	Time (s)		AUC	Time (s)
$OAM_{seq}$	$0.7421 \pm 0.0587$	58.3867	$OAM_{seq}$	$0.4916 \pm 0.0266$	15.6133
$OAM_{gra}$	$0.8379 \pm 0.0277$	19.5711	$OAM_{gra}$	$0.5099 \pm 0.0318$	5.6508
OPAUC	$0.8729 \pm 0.0036$	<b>8.0313</b>	OPAUC	$0.4982 \pm 0.0368$	1.4664
AdaOAM	$0.8653 \pm 0.0120$	10.4055	AdaOAM	$0.5001 \pm 0.0321$	2.1391
KNNOAM	<b><math>0.9166 \pm 0.0023</math></b>	42.7477	KNNOAM	<b><math>0.6299 \pm 0.0309</math></b>	<b>1.1422</b>

## References

1. Agarwal, S.: Surrogate regret bounds for bipartite ranking via strongly proper losses. *J. Mach. Learn. Res.* **15**(1), 1653–1674 (2014)
2. Agarwal, S., Graepel, T., Herbrich, R., Har-Peled, S., Roth, D.: Generalization bounds for the area under the ROC curve. *JMLR* **6**, 393–425 (2005)
3. Ailon, N., Mohri, M.: An efficient reduction of ranking to classification. In: *COLT 2008, Helsinki, Finland, 9–12 July 2008*, pp. 87–98 (2008)
4. Anava, O., Levy, K.:  $k^*$ -nearest neighbors: from global to local. In: Lee, D.D., Sugiyama, M., Luxburg, U.V., Guyon, I., Garnett, R. (eds.) *Advances in Neural Information Processing Systems*, vol. 29, pp. 4916–4924. Curran Associates Inc., Red Hook (2016)
5. Balcan, M.F., Bansal, N., Beygelzimer, A., Coppersmith, D., Langford, J., Sorkin, G.B.: Robust reductions from ranking to classification. *Mach. Learn.* **72**(1), 139–153 (2008)
6. Beygelzimer, A., Kakade, S., Langford, J.: Cover trees for nearest neighbor. In: *ICML*, pp. 97–104. ACM, New York (2006)
7. Cardot, H., Degras, D.: Online principal component analysis in high dimension: which algorithm to choose? *CoRR* abs/1511.03688 (2015). <http://arxiv.org/abs/1511.03688>
8. Chaudhuri, K., Dasgupta, S.: Rates of convergence for nearest neighbor classification. In: *NIPS 2014*, pp. 3437–3445 (2014)
9. Cléménçon, S., Vayatis, N.: Tree-based ranking methods. *IEEE Trans. Inf. Theory* **55**(9), 4316–4336 (2009)
10. Cléménçon, S., Lugosi, G., Vayatis, N.: Ranking and empirical minimization of U-statistics. *Ann. Stat.* **36**(2), 844–874 (2008)
11. Cortes, C., Mohri, M.: AUC optimization vs. error rate minimization. In: Thrun, S., Saul, L., Schölkopf, B. (eds.) *NIPS*, pp. 313–320. MIT Press, Cambridge (2004)
12. Devroye, L., Györfi, L., Krzyżak, A., Lugosi, G.: On the strong universal consistency of nearest neighbor regression function estimates. *Ann. Stat.* **22**(3), 1371–1385 (1994)
13. Ding, Y., Zhao, P., Hoi, S.C.H., Ong, Y.: An adaptive gradient method for online AUC maximization. In: *AAAI*, pp. 2568–2574 (2015)
14. Gao, W., Jin, R., Zhu, S., Zhou, Z.: One-pass AUC optimization. In: *ICML 2013*, pp. 906–914 (2013)
15. Gao, W., Zhou, Z.: On the consistency of AUC pairwise optimization. In: *IJCAI 2015*, pp. 939–945 (2015)
16. Hanley, J.A., Mcneil, B.J.: The meaning and use of the area under a receiver operating characteristic (ROC) curve. *Radiology* **143**, 29–36 (1982)
17. Kar, P., Sriperumbudur, B.K., Jain, P., Karnick, H.: On the generalization ability of online learning algorithms for pairwise loss functions. In: *30th ICML 2013, 16–21 June 2013, Atlanta, GA, USA*, pp. 441–449 (2013)
18. Kotlowski, W., Dembczynski, K., Hüllermeier, E.: Bipartite ranking through minimization of univariate loss. In: *ICML*, pp. 1113–1120. Omnipress (2011)
19. Robbiano, S., Cléménçon, S.: Minimax learning rates for bipartite ranking and plug-in rules. *ICML 2011*, pp. 441–448 (2011)
20. Uematsu, K., Lee, Y.: On theoretically optimal ranking functions in bipartite ranking. Technical report 863, Department of Statistics, The Ohio State University, December 2011
21. Wang, Y., Kharon, R., Pechyony, D., Jones, R.: Generalization bounds for online learning algorithms with pairwise loss functions. In: *COLT*, pp. 13.1–13.22 (2012)
22. Zhao, P., Hoi, S.C.H., Jin, R., Yang, T.: Online AUC maximization. In: *ICML*, pp. 233–240 (2011)