

On Adaboost and Optimal Betting Strategies

Pasquale Malacaria
School of Electronic Engineering
and Computer Science
Queen Mary, University of London
Email: pm@dcs.qmul.ac.uk

Fabrizio Smeraldi
School of Electronic Engineering
and Computer Science
Queen Mary, University of London
Email: fabri@dcs.qmul.ac.uk

Abstract

We explore the relation between the Adaboost weight update procedure and Kelly's theory of betting. Specifically, we show that an intuitive optimal betting strategy can easily be interpreted as the solution of the dual of the classical formulation of the Adaboost minimisation problem. This sheds new light over a substantial simplification of Adaboost that had so far only been considered a mere computational development. Besides allowing a natural derivation of the simplification, our approach establishes a connection between Adaboost and Information Theory that we believe may cast a new light on some boosting concepts.

Index Terms

Adaboost, Optimal betting strategies, Kelly's theory of gambling, Information theory, Bregman divergence

1. Introduction

Suppose you want to bet on a series of races of six horses, on which no information is initially available (the same horses participate in all races). Some information becomes available after each race in the form of the names of a subset of the losers for that race. Let us furthermore assume that you are required to bet all of your capital on each race.

Initially, as nothing is known, it seems reasonable to spread one's bets evenly, so that your initial bet will allocate 1/6-th of your capital on each horse.

Suppose now that you learn that horses 3, 4, 5 and 6 lost the first race. As your previous bet was

optimal¹ (w.r.t. your information at the time) you want the new bet to be the "closest" possible to the previous bet, taking into account the outcome of the race. You could for instance for a fraction ρ of your capital ($1/2 \leq \rho \leq 1$) to be bet on the latest "possible" winners (horses 1,2) and $1 - \rho$ on the latest losers (horses 3,4,5,6). This ρ is a measure of how much you believe the result of the last race to be a predictor of the result of the next race.

Notice that, as you have to invest all your capital at each time, $\rho = 1$ (i.e. betting everything on the set of horses containing the previous winner) would be a bad choice: you would then lose all your capital if the following race is won by any other horse. Table 1 presents an algorithm implementing these ideas: the next bet for horse i is obtained from the previous bet by multiplying it by $\frac{\rho}{\epsilon}$ (if the horse was a possible winner) and by $\frac{1-\rho}{(1-\epsilon)}$ (if it was a loser); ϵ is the total amount bet on the latest possible winners in the previous race.

If you need to choose a value of ρ before all races, as no information is available, you may well opt for $\rho = 0.5$. This is a "non-committal" choice: you believe that the subset of losers of the each race revealed to you each time will not to be a good a predictor for the result of the next race. In our example, where 3,4,5 and 6 lost the first race, choosing $\rho = 0.5$ will lead to the bets $(1/4, 1/4, 1/8, 1/8, 1/8, 1/8)$ on the second race. Table 2 shows the bets generated by applying the betting algorithm with $\rho = 0.5$ to a series of 4 races.

In Section 4 we show that the Adaboost weight update cycle is *exactly* equivalent to this betting strategy with $\rho = 0.5$, modulo the trivial identification of

1. Optimal betting is here understood in the information theoretical sense. In probability terms betting all capital on the horse most likely to win will give the highest expected return, but will make you bankrupt with probability 1 in the long term. Maximizing the return has to be balanced with the risk of bankruptcy; see Section 3 for details.

Table 1. The betting algorithm

Betting Algorithm:

Choose the initial bets $b_{1,i} = \frac{1}{m}$, where m is the number of horses;
 For each race:

- 1) Let W be the set of possible winners in the previous race. Compute $\epsilon = \sum_{i \in W} b_{t,i}$;
- 2) Update the bets for race $t + 1$: $b_{t+1,i} = \frac{\rho}{\epsilon} b_{t,i}$ if $i \in W$, $b_{t+1,i} = \frac{1-\rho}{(1-\epsilon)} b_{t,i}$ otherwise.

the horses with training examples and of the possible winners with misclassified training data. As will be shown, this derives straightforwardly from a direct solution of the *dual* of the Adaboost minimisation problem. This result, that is the main contribution of our work, establishes the relation between Kelly’s theory of optimal betting and Adaboost sketched in Figure 1. It also leads directly to the simplified formulation of Adaboost shown in Table 4.

The rest of this paper is organised as follows: Section 2 gives a brief introduction to the Adaboost algorithm; in Section 3 we outline the theory of optimal betting underlying the strategy displayed in Table 1. The equivalence between this betting strategy and Adaboost is proved in Section 4. Finally in Section 2 we proceed in the opposite direction, and show how convergence theorems for Adaboost can be used to investigate the asymptotic behaviour of our betting strategy.

2. Background on Adaboost

Boosting concerns itself with the problem of combining several prediction rules, with poor individual performance, into a highly accurate predictor. This is commonly referred to as combining a set of “weak” learners into a “strong” classifier. Adaboost, introduced by Yoav Freund and Robert Schapire [1], differs from previous algorithms in that it does not require previous knowledge of the performance of the weak hypotheses, but it rather adapts accordingly. This adaptation is achieved by maintaining a distribution of weights over the elements of the training set.

Adaboost is structured as an iterative algorithm, that works by maintaining a distribution of weights over the training examples. At each iteration a new weak learner (classifier) is trained, and the distribution of weights is updated according to the examples it misclassifies. The idea is to increase the importance of the training examples that are “difficult” to classify. The weights also determine the contribution of the each weak learner to the final strong classifier.

In the following, we will indicate the m training examples as $\{(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_m, y_m)\}$, where the \mathbf{x}_i are vectors in a feature space X and the $y_i \in \{-1, +1\}$

are the respective class labels. Given a family of weak learners $\mathcal{H} = \{h_t : X \rightarrow \{-1, +1\}\}$, it is convenient to define $\mathbf{u}_{t,i} = y_i h_t(\mathbf{x}_i)$, so that $\mathbf{u}_{t,i}$ is $+1$ if h_t classifies \mathbf{x}_i correctly, and -1 otherwise. If we now let $d_{t,i}$ be the distribution of weights at time t , with $\sum_i d_{t,i} = 1$, the cycle of iterations can be described as in Table 3.

Each iteration of the Adaboost algorithm solves the following minimisation problem [2], [3]:

$$\alpha_t = \arg \min_{\alpha} \log Z_t(\alpha), \quad (1)$$

where

$$Z_t(\alpha) = \sum_{i=1}^m d_{t,i} \exp(-\alpha u_{t,i}) \quad (2)$$

is the partition function that appears in point 3 of Table 3. In other words, the choice of the updated weights $d_{t+1,i}$ is such that α_t will minimise Z_t .

3. The gambler’s problem: optimal betting strategies

Before investigating the formal relation between Adaboost and betting strategies, we review in this Section the basics of the theory of optimal betting, that motivated the heuristic reasoning of Section 1. Our presentation is a summary of the ideas in Chapter 6.1 of [4], based on Kelly’s theory of gambling [5].

Let us consider the problem of determining the optimal betting strategy for a race run by m horses x_1, \dots, x_m . For the i -th horse, let p_i, o_i represents

Table 2. Bets b_i generated by information O_i over four races (w denotes a possible winner).

	x_1	x_2	x_3	x_4	x_5	x_6
b_0	1/6	1/6	1/6	1/6	1/6	1/6
O_0	w	w	ℓ	ℓ	ℓ	ℓ
b_1	1/4	1/4	1/8	1/8	1/8	1/8
O_1	ℓ	ℓ	w	w	ℓ	ℓ
b_2	1/6	1/6	1/4	1/4	1/12	1/12
O_2	w	ℓ	ℓ	ℓ	w	ℓ
b_3	1/3	1/9	1/6	1/6	1/6	1/18
O_3	ℓ	w	ℓ	ℓ	ℓ	ℓ
b_4	27/144	1/2	27/288	27/288	27/288	27/864

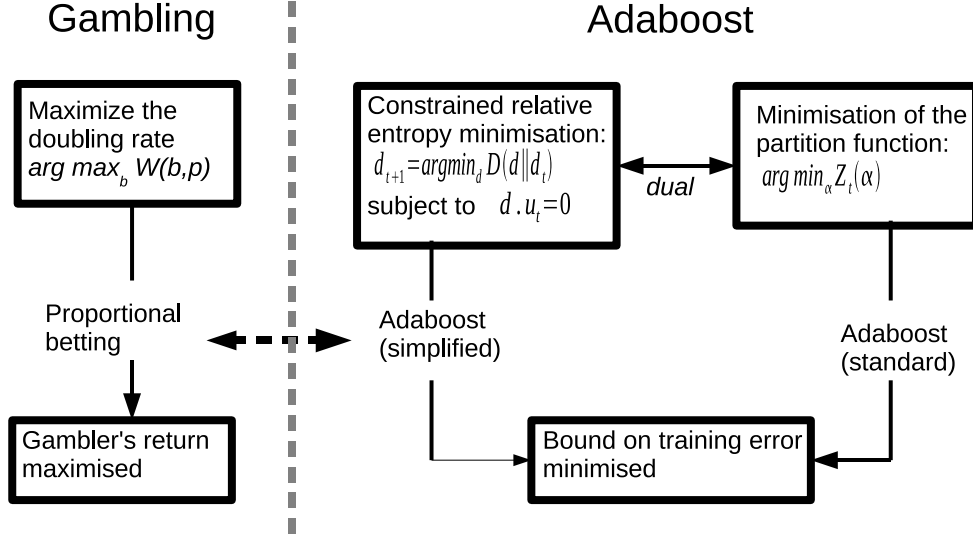


Figure 1. Relation between Kelly's theory of gambling and the standard and dual formulations of Adaboost.

respectively the probability of x_i winning and the *odds*, intended as o_i -for-one (e.g., at 3-for-1 the gambler will get 3 times his bet in case of victory).

Assuming that the gambler always bets all his wealth, let b_i be the fraction that is bet on horse x_i . Then if horse X wins the race, the gambler's capital grows by the *wealth relative* $S(X) = b(X)o(X)$. Assuming all races are independent, the gambler's wealth after n races is therefore $S_n = \prod_{j=1}^n S(X_j)$.

It is at this point convenient to introduce the *doubling rate*

$$W(\mathbf{b}, \mathbf{p}) = E(\log S(X)) = \sum_{i=1}^m p_i \log(b_i o_i). \quad (3)$$

(in this information theoretical context, the logarithm should be understood to be in base 2). From the weak law of large numbers and the definition of S_n , we have that

$$\frac{1}{n} \log S_n = \frac{1}{n} \sum_{j=1}^n \log S(X_j) \rightarrow E(\log S(X)) \quad (4)$$

in probability, which allows us to express the gambler's wealth in terms of the doubling rate:

$$S_n = 2^{n \frac{1}{n} \log S_n} \doteq 2^{n E(\log S(X))} = 2^{n W(\mathbf{p}, \mathbf{b})} \quad (5)$$

where by " \doteq " we indicate equality to the first order in the exponent.

As can be shown the optimal betting scheme, i.e. the choice of the b_i that maximises $W(\mathbf{p}, \mathbf{b})$, is given by $\mathbf{b} = \mathbf{p}$, that is to say *proportional betting*. Note that this is independent of the odds o_i .

Incidentally, it is crucial to this concept of optimality that 4 is true with probability one in the limit of large n . Consider for example the case of 3 horses that win with probability $1/2$, $1/4$ and $1/4$ respectively and that are given at the same odds (say 3). In this case, it would be tempting to bet all our capital on the 1st horse, as this would lead to the highest possible expected gain of $3 \times (1/2) = 3/2$. However, it is easy to see that such a strategy would eventually make us bankrupt with probability one, as the probability that horse 1 wins the first n races is vanishingly small for large n , so that the expected value is not a useful criterion to guide our betting. Equation 4, instead, assures us that the capital will asymptotically grow like 2^{nW} . For the proportional betting strategy, in this case, $W = (1/2) \log(3 \times 1/2) + (1/4) \log(3 \times 1/4) + (1/4) \log(3 \times 1/6) \approx 0.08$, so that the capital will grow to infinity with probability one [4].

In the special case of fair odds with $o_i = \frac{1}{p_i}$ (fair odds are defined as odds such that $\sum 1/o_i = 1$), the doubling rate becomes

$$W(\mathbf{p}, \mathbf{b}) = \sum p_i \log \left(\frac{b_i}{p_i} \right) = -D(\mathbf{p}||\mathbf{b}) / \log_e 2 \quad (6)$$

where $D(\mathbf{p}||\mathbf{b})$ is the Kullback-Leibler divergence (also known as the relative entropy) of the probability of victory \mathbf{p} and the betting strategy \mathbf{b} .

Note that, since the odds are fair, the gambler can at best conserve his wealth, that is $\max_{\mathbf{b}} W(\mathbf{p}, \mathbf{b}) = 0$. In this case, it follows directly from the properties of the Kullback-Leibler divergence that the maximum occurs

Table 3. The standard formulation of the Adaboost algorithm

<p>The Adaboost algorithm:</p> <p>Initialise $d_{1,i} = \frac{1}{m}$. For $t = 1, \dots, T$:</p> <ol style="list-style-type: none"> 1) Select the weak learner h_t that minimises the training error ϵ_t, weighted by current distribution of weights $d_{t,i}$: $\epsilon_t = \sum_{i u_{t,i}=-1} d_{t,i}$. Check that $\epsilon_t < 0.5$. 2) Set $\alpha_t = \frac{1}{2} \log \frac{1-\epsilon_t}{\epsilon_t}$ 3) Update the weights according to $d_{t+1,i} = \frac{1}{Z_t(\alpha_t)} d_{t,i} \exp(-\alpha_t u_{t,i})$, where $Z_t(\alpha_t) = \sum_{i=1}^m d_{t,i} \exp(-\alpha_t u_{t,i})$ is a normalisation factor which ensures that the $d_{t+1,i}$ will be a distribution. <p>Output the final classifier: $H(x) = \text{sign} \left(\sum_{t=1}^T \alpha_t h_t(x) \right)$</p>
--

for $b_i = p_i$, i.e. that the optimal strategy is proportional betting.

In the following, for convenience, we will refer to Equation 6 as if the base of the logarithm were e instead than 2. This modification is inessential, and for the sake of convenience we will continue to refer to the new quantity as the ‘‘doubling rate’’.

4. Equivalence of weight update procedures

We are now in a position to prove the equivalence of the betting algorithm introduced in Section 1 with the Adaboost weight update strategy, points 2) and 3) in Table 3. Modulo the identification of the horses with training examples and of the possible winners of each race with misclassified examples, this leads to the simplification of Adaboost shown in Table 4.

Our proof is based on the duality relations discussed in a number of excellent papers that cast Adaboost in terms of the minimisation of Kullback-Leibler or Bregman divergences [6], [7], [2]. In these, the Adaboost weight update equations are shown to solve the following problem:

$$\mathbf{d}_{t+1} = \arg \min_{\mathbf{d}} D(\mathbf{d}||\mathbf{d}_t) \quad \text{subject to } \mathbf{d} \cdot \mathbf{u}_t = 0. \quad (7)$$

In information theoretical terms, this can be interpreted as the Minimum Discrimination Information principle applied to the estimation of the distribution \mathbf{d}_{t+1} given the initial estimate \mathbf{d}_t and the linear constraint determined by the statistics \mathbf{u}_t [8].

The key observation is now that the Kullback-Leibler divergence of the d_i and the $d_{t,i}$,

$$D(\mathbf{d}||\mathbf{d}_t) = \sum d_i \log \frac{d_i}{d_{t,i}}, \quad (8)$$

bears a strong formal analogy to Equation 6 above.

For the sake of readability, we discard the index t and simply write d_i for $d_{t,i}$ and d_i^* for $d_{t+1,i}$. The Adaboost dual problem Equation 7 then becomes

$$\mathbf{d}^* = \arg \min_{\tilde{\mathbf{d}}} D(\tilde{\mathbf{d}}||\mathbf{d}) \quad \text{subject to } \tilde{\mathbf{d}} \cdot \mathbf{u} = 0. \quad (9)$$

The constraint $\mathbf{d}^* \cdot \mathbf{u} = 0$ can be rewritten as

$$\sum_{i|u_i=-1} d_i^* = \sum_{i|u_i=+1} d_i^* = 1/2. \quad (10)$$

With these notations, we are ready to prove the following

Theorem 1. *The simplified weight update strategy in table 4 leads to the same weights as the original algorithm in Table 3*

Proof: we need to show that the weight update equations in table 4, that are the formal transcription of the betting strategy in Table 2, actually solve the optimisation problem Equation 9.

The constraint Equation 10 allows us to split the objective function in Equation 9 into the sum of two terms:

$$\begin{aligned} D(\mathbf{d}^*||\mathbf{d}) &= \\ &= \sum_{i|u_i=-1} d_i^* \log \frac{d_i^*}{d_i} + \sum_{i|u_i=+1} d_i^* \log \frac{d_i^*}{d_i} \quad (11) \\ &= D_-(\mathbf{d}^*, \mathbf{d}) + D_+(\mathbf{d}^*, \mathbf{d}), \end{aligned}$$

each of which can be maximised separately.

Notice that neither of these terms taken separately represents a Kullback-Leibler divergence, as the partial sums over the d_i^* and d_i are not over distributions. However, this is easily remedied by normalising these sub-sequences separately. To this purpose, we introduce the Adaboost weighted error ϵ_t (see Table 3, point 1), that omitting the t index is equal to $\epsilon = \sum_{i|u_i=-1} d_i$. The first term of the summation can

Table 4. Applying the proportional betting strategy to the solution of the Adaboost dual problem leads to a simpler formulation of the algorithm, equivalent to the original

<p>Adaboost (simplified):</p> <p>Initialise $d_{1,i} = \frac{1}{m}$. For $t = 1, \dots, T$:</p> <ol style="list-style-type: none"> 1) Select the weak learner h_t that minimises the training error ϵ_t, weighted by current distribution of weights $d_{t,i}$: $\epsilon_t = \sum_{i u_{t,i}=-1} d_{t,i}$. Check that $\epsilon_t < 0.5$. 2) Update the weights: $d_{t+1,i} = \frac{1}{2\epsilon_t} d_{t,i}$ if $u_{t,i} = -1$, $d_{t+1,i} = \frac{1}{2(1-\epsilon_t)} d_{t,i}$ otherwise. <p>Output the final classifier: $H(\mathbf{x}) = \text{sign} \left(\sum_{t=1}^T \alpha_t h_t(\mathbf{x}) \right)$, with $\alpha_t = \frac{1}{2} \log \frac{1-\epsilon_t}{\epsilon_t}$</p>
--

then be rewritten as

$$\begin{aligned}
D_-(\mathbf{d}^*, \mathbf{d}) &= \\
&= \frac{1}{2} \sum_{i|u_i=-1} 2d_i^* \log \frac{2d_i^*}{d_i/\epsilon} - \frac{1}{2} \log 2\epsilon \quad (12) \\
&= \frac{1}{2} D_-(2\mathbf{d}^* \parallel \mathbf{d}/\epsilon) - \frac{1}{2} \log 2\epsilon,
\end{aligned}$$

where by D_- we mean the Kullback-Leibler divergence computed over the subset of indexes $\{i|u_i = -1\}$. The minimum point is now immediately evident: $2d_i^* = d_i/\epsilon$ or, in the notation of Table 4,

$$d_{t+1,i} = \frac{1}{2\epsilon_t} d_{t,i} \quad \text{if } u_{t,i} = -1. \quad (13)$$

Similarly, the minimization of D_+ leads to $2d_i^* = d_i/(1-\epsilon)$, or

$$d_{t+1,i} = \frac{1}{2(1-\epsilon_t)} d_{t,i} \quad \text{if } u_{t,i} = +1. \quad (14)$$

These are precisely the update equations given in Table 4, which proves the theorem.

We have thus shown that the simplified algorithm in Table 4 is completely equivalent to the original Adaboost algorithm in Table 3. Indeed, the two algorithms generate step by step the same sequence of weights, as well as the same decision function.

This simplification of Adaboost, although not so widely used, is already known to researchers in the field (see for instance [9]). However, to the best of our knowledge it has always been derived in a purely computational way by algebraic manipulation of the equations in Table 3. We believe that our derivation as a direct solution of the dual problem Equation 7 is both more straightforward and more illuminating, in that it brings out the direct correspondence with betting strategies represented in Figure 1. Finally, we note that a formally analogous weight update strategy has been employed in online *variants* of the algorithm [10], [11].

5. Adaboost as a betting strategy

The considerations above suggest an interpretation of the betting algorithm in Table 1 as an iterative density estimation procedure in which an initial estimate, represented by the current bets \mathbf{b} , is refined using the information on the previous race. The updated estimate is then used to place new bets (in the notations of Table 1, $b_{t+1,i} = p_i$), in accordance to the proportional betting strategy.

This interpretation is supported by convergence results for the Adaboost weight update cycle, in either its standard form (Table 3) or its simplified form (Table 4). The application to our betting scenario with partial information is straightforward, modulo the the simple formal identifications $b_i = d_{t,i}$ and $p_i = d_{t+1,i}$.

To prove this, we will first state the following convergence theorem, which has been proved in [6], [7], [2]:

Theorem 2. *If the system of the linear constraints has non-trivial solutions, i.e. there is some \mathbf{d} such that $\mathbf{d} \cdot \mathbf{u}_t = 0 \forall t$, then the Adaboost weight distribution converges to the distribution \mathbf{d}^* satisfying all the linear constraints for which $D(\mathbf{d}^* \parallel \mathbf{1}/m)$ is minimal (here by $\mathbf{1}$ we indicate the vector of all ones).*

In other words, the weight update algorithm converges to a probability density that is, among all those that satisfy all the constraints, the closest to the original guess of a uniform distribution. This ‘‘approximation’’ property is best expressed by the following ‘‘Pythagorean theorem’’ for Kullback-Leibler divergences, that was derived in [12] in the context of an application to iterative density estimation (a generalisation to Bregman divergences can be found in [7]). For all solutions \mathbf{d} of the system of linear constraints, we have

$$D(\mathbf{d} \parallel \mathbf{1}/m) = D(\mathbf{d} \parallel \mathbf{d}^*) + D(\mathbf{d}^* \parallel \mathbf{1}/m). \quad (15)$$

In terms of our gambling problem, theorem 2 above

means that the betting algorithm in Table 1 converges to an estimate $\mathbf{p}^* = \mathbf{b}^*$ of the probability of winning of each single horse. Among all the densities that satisfy the constraints imposed by the information on the previous races, \mathbf{p}^* is the one that would maximise the doubling rate of the initial uniform bet, $W(\mathbf{p}^*, \mathbf{1}/m)$ (strictly speaking, theorem 2 proves this result only for the case that $\rho = 1/2$; however, generalisation to the case of a generic ρ is straightforward, see [6]).

However, for convergence to occur the information \mathbf{u}_t on the various races must be provided in the order determined by Point 1) in either Table 3 or Table 4, corresponding to the Adaboost requirement that the weak learner should minimise the error with respect to the weights. In the betting case, of course, this appears problematic as the gambler has no influence on the order in which the results occur. The corresponding assumption corresponding to the minimal error requirement would be, in a sense, one of “minimal luck”: that is, once we decide on a betting strategy, the next race is the one for which the set of the possible winners has, according to our estimate, the lowest total probability.

When the conditions for convergence are met, the “Pythagorean” equality Equation 15 quantifies the advantage of the estimated probability \mathbf{p}^* over any other density \mathbf{p} compatible with the constraint, if the initial uniform distribution (or indeed any other density obtained in an intermediate iteration) are used for betting:

$$W(\mathbf{p}, \mathbf{1}/m) = W(\mathbf{p}^*, \mathbf{1}/m) - D(\mathbf{p} \parallel \mathbf{p}^*). \quad (16)$$

6. Conclusions

We discussed the relation between Adaboost and Kelly’s theory of gambling. Specifically, we showed how the dual formulation of Adaboost formally corresponds to maximising the doubling rate for a gambler’s capital, in a partial information situation. The applicable modification of the optimal proportional betting strategy therefore maps to a direct, intuitive solution of the dual of the Adaboost problem. This leads naturally to a substantial formal simplification of the Adaboost weight update cycle. Previous derivations of this simplification are, to the best of our knowledge, purely computational, and therefore, besides being more cumbersome, lack any meaningful interpretation.

The implications of the correspondence between Adaboost and gambling theory works in both directions; as we have shown, it is for instance possible to use convergence results obtained for Adaboost for investigating the asymptotic behaviour of the betting strategy we introduced in Table 1.

We believe that a further investigation of the relation between Adaboost and Kelly’s theory, or indeed information theory in general, may lead to a deeper insight into some boosting concepts, and possibly a cross-fertilization between these two domains.

References

- [1] Y. Freund and R. Schapire, “A decision-theoretic generalization of on-line learning and an application to boosting,” *Journal of Computer and System Science*, vol. 55, no. 1, 1997.
- [2] M. W. J. Kivinen, “Boosting as entropy projection,” in *Proceedings of COLT’99*. ACM, 1999.
- [3] R. Schapire and Y. Singer, “Improved boosting algorithms using confidence-rated predictions,” *Machine Learning*, vol. 37, no. 3, 1999.
- [4] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. Wiley Interscience, 1991.
- [5] J. J. L. Kelly, “A new interpretation of information rate,” *Bell System Technical Journal*, vol. 35, pp. 917–926, July 1956.
- [6] M. Collins, R. E. Schapire, and Y. Singer, “Logistic regression, adaboost and bregman distances,” in *Proceedings of the Thirteenth Annual Conference on Computational Learning Theory*, 2000.
- [7] S. D. Pietra, V. D. Pietra, and H. Lafferty, “Inducing features of random field,” *IEEE Transactions Pattern Analysis and Machine Intelligence*, vol. 19, no. 4, April 1997.
- [8] S. Kullback, *Information Theory and Statistics*. Wiley, 1959.
- [9] C. Rudin, R. E. Schapire, and I. Daubechies, “On the dynamics of boosting,” in *Advances in Neural Information Processing Systems*, vol. 16, 2004.
- [10] N. C. Oza and S. Russell, “Online bagging and boosting,” in *Proceedings of the IEEE International Conference on Systems, Man and Cybernetics*, vol. 3, 2005, pp. 2340–2345.
- [11] H. Grabner and H. Bischof, “On-line boosting and vision,” in *Proceedings of CVPR*, vol. 1, 2006, pp. 260–267.
- [12] S. Kullback, “Probability densities with given marginals,” *The Annals of Mathematical Statistics*, vol. 39, no. 4, 1968.