ABSTRACT

Relative Loss Bounds for On-line Learning Using Divergence Measures
Manfred Warmuth

An on-line learning algorithm receives a sequence of alternating instances and labels: $x_1, y_1, x_2, y_2, \ldots, x_T, y_T$. After observing the sequence up to x_t, the algorithm outputs $\hat{y}_t$ which is its prediction for the label y_t. The algorithm then receives y_t and suffers a *loss* $L_t = L(y_t, \hat{y}_t)$. The loss maps a label and a prediction to a non-negative real value that represents the discrepancy between the two.

Relative loss bounds are upper bounds on the difference

$$\sum_{t=1}^{T} L_t \quad - \quad \min_{f \in F} \sum_{t=1}^{T} L(y_t, f(x_t))$$

where F is a parameterized comparison class of predictors. These bounds hold uniformly for all sequences of instances and labels.

This talk will overview some of the work on on-line algorithms that have good relative loss bounds starting from Littlestone's Winnow algorithm. The emphasis will be on a unified framework that is based on a divergence measure between the members of F. The divergence is used to derive the parameter update of the on-line learner and it becomes the potential function in the proof of the relative loss bound for the same update.

We use Bregman distances as our divergences. The axiomatic properties of these divergences have been investigated by Csiszar. We will introduce families of update algorithms that are characterized by different divergences and discuss the merits of the different families.