

# Accelerating AdaBoost using multi-armed bandits

Róbert Busa-Fekete and Balázs Kégl

{busarobi, balazs.kegl}@gmail.com

## Motivation

ADABOOST.MH with trees or products is **state-of-the-art but slow** if the number of features is large

Idea: learn the “utility“ of features using **multi-armed bandits (MAB)** → the base learner can be aided in finding good base classifiers fast

## ADABOOST.MH

```

ADABOOST.MH( $\mathbf{X}, \mathbf{Y}, \mathbf{W}^{(1)}, \text{BASE}(\cdot, \cdot, \cdot), T$ )
1   for  $t \leftarrow 1$  to  $T$ 
2      $\mathbf{h}^{(t)}(\cdot) \leftarrow \alpha^{(t)} \mathbf{v}^{(t)} \varphi^{(t)}(\cdot) \leftarrow \text{BASE}(\mathbf{X}, \mathbf{Y}, \mathbf{W}^{(t)})$ 
3     for  $i \leftarrow 1$  to  $n$  for  $\ell \leftarrow 1$  to  $K$ 
4        $w_{i,\ell}^{(t+1)} \leftarrow w_{i,\ell}^{(t)} \frac{\exp(-h_{\ell}^{(t)}(\mathbf{x}_i) y_{i,\ell})}{\sum_{i'=1}^n \sum_{\ell'=1}^K w_{i',\ell'}^{(t)} \exp(-h_{\ell'}^{(t)}(\mathbf{x}_{i'}) y_{i',\ell'})}$ 
5   return  $\mathbf{f}^{(T)}(\cdot) = \sum_{t=1}^T \mathbf{h}^{(t)}(\cdot)$ 

```

## Base learners

The goal of the **base learner** is to maximize the **edge**


$$\gamma = \sum_{i=1}^n w_i h(\mathbf{x}_i) y_i.$$

DECISION STUMP :

$$\varphi_{j,b}(\mathbf{x}) = \begin{cases} 1 & \text{if } x^{(j)} \geq b, \\ -1 & \text{otherwise.} \end{cases}$$


PRODUCT :  $h(\mathbf{x}) = \prod_{j=1}^m \varphi_{j,b_j}(\mathbf{x})$

DECISION TREE :



## Multi-armed bandits

Each **feature** corresponds to an **arm**.



Each time a feature is selected, its **reward**


$$r_j^{(t)} = 1 - \sqrt{1 - \gamma^2}$$

is recorded.

We select  $k$  features that maximize the sum of the reward and a **confidence interval term**

$$c_j^{(t)} = \sqrt{\frac{2 \ln t}{T_j^{(t)}}}.$$

## Validation curves



## Running time in minutes

Base learner	STUMP	PRODUCT	TREE
Running time	274	456	384

## Results

Hyperparameters:

- Number of terms in PRODUCT:  $m = 3$
- Number of leaves in DECISION TREE:  $N = 8$
- Number of features chosen by UCB( $k$ ):  $k = 50$
- Number of iterations (validated):  $T = 8000$

learner \ data set	Churn (AUC)		Appetency (B. Acc.)		Upselling (AUC)	
	Validation	Evaluation	Validation	Evaluation	Validation	Evaluation
STUMP	0.7424	0.6633	0.7051	0.6617	0.8938	0.8840
PRODC	0.6702	0.6314	0.6999	0.6807	0.8888	0.8666
TREE	0.7088	0.6819	0.7424	0.7343	0.8956	0.8883
WINNER	–	0.7611	–	0.883	–	0.9038