

# Fast boosting using adversarial bandits

Róbert Busa-Fekete<sup>1,2</sup> Balázs Kégl<sup>1,3</sup>

<sup>1</sup>Linear Accelerator Laboratory (LAL), University of Paris-Sud, CNRS

<sup>2</sup>Research Group on Artificial Intelligence of the Hungarian Academy of Sciences  
and University of Szeged (RGAI)

<sup>3</sup>Computer Science Laboratory (LRI), University of Paris-Sud, CNRS and  
INRIA-Saclay

ANR meeting  
January 26, 2012

- 1 Introduction
  - ADABOOST.MH reminder
  - Base learning, in nutshell
- 2 Accelerating the training of AdaBoost
  - Motivation, Related work
  - The formal setup
  - Adversarial bandits
  - Weak-to-strong-learning result
- 3 Experiments
- 4 Conclusions



ADABOOST( $D_n = \{(\mathbf{x}_i, y_i)\}_{i=1}^n$ , BASE( $\cdot, \cdot$ ),  $T$ )

1  $\mathbf{w}^{(1)} \leftarrow (1/n, \dots, 1/n)$   $\triangleright$  *initial weights*

2 **for**  $t \leftarrow 1$  **to**  $T$

3  $h^{(t)} \leftarrow \text{BASE}(D_n, \mathbf{w}^{(t)})$   $\triangleright$  *calling the base learner*

4  $\gamma^{(t)} \leftarrow \sum_{i=1}^n w_i^{(t)} h^{(t)}(\mathbf{x}_i) y_i$   $\triangleright$  *edge = 1 - 2 × error*

5  $\alpha^{(t)} \leftarrow \frac{1}{2} \ln \left( \frac{1 + \gamma^{(t)}}{1 - \gamma^{(t)}} \right)$   $\triangleright$  *coefficient of  $h^{(t)}$*

6 **for**  $i \leftarrow 1$  **to**  $n$   $\triangleright$  *re-weighting the points*

7 **if**  $h^{(t)}(\mathbf{x}_i) \neq y_i$  **then**

8  $w_i^{(t+1)} \leftarrow w_i^{(t)} \frac{1}{1 - \gamma^{(t)}}$


9 **else**

10  $w_i^{(t+1)} \leftarrow w_i^{(t)} \frac{1}{1 + \gamma^{(t)}}$

11 **return**  $f^{(T)}(\cdot) = \sum_{t=1}^T \alpha^{(t)} h^{(t)}(\cdot)$




# Decision stumps



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




# Decision stumps



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

- 1 Can be learned in  $\theta(ndK)$  time (if features are **pre-sorted**)

# Decision stumps



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

- 1 Can be learned in  $\theta(ndK)$  time (if features are **pre-sorted**)
- 2 Looking at each feature in every boosting iterations


# Accelerating the training of AdaBoost

## Motivation, Related work

- Well boostable base learners
  - Decision stumps  $O(ndK)$
  - Decision trees  $O(ndK \log N)$
  - Decision product  $O(ndKm)$
- Saving on the  $n$  factor
  - stochastic boosting, FILTERBOOST<sup>1</sup>
- Saving on the  $d$  factor
  - LAZYBOOST<sup>2</sup> (random selection)
  - our technique: learn the usefulness of the features in a sequential game using multi-armed bandits (MABs)

---

<sup>1</sup>J.K. Bradley and R.E. Schapire. FilterBoost: Regression and classification on large datasets. In *NIPS 2008*.

<sup>2</sup>G. Escudero, L. Màrquez, and G. Rigau. Boosting applied to word sense disambiguation. In *ECML 2000*. 

# The formal setup

- Partition the base classifier set  $\mathcal{H}$  into  $\{\mathcal{H}_1, \dots, \mathcal{H}_M\}$ 
  - in each iteration  $t$ , use the Multi-armed Bandit (MAB) algorithm to select a subset  $\mathcal{H}_{j^{(t)}}$
  - call the base learner to select  $h^{(t)} \in \mathcal{H}_{j^{(t)}}$
  - compute the edge

$$\gamma^{(t)} = \sum_{i=1}^n w_i^{(t)} h^{(t)}(\mathbf{x}_i) y_i = 1 - 2 \times \text{error}$$

- return the reward

$$r_j^{(t)} = \min \left( 1, -\log \sqrt{1 - \gamma^{(t)2}} \right).$$

to the MAB



ADABOOST( $D_n = \{(\mathbf{x}_i, y_i)\}_{i=1}^n$ , BASE( $\cdot, \cdot$ ),  $T$ )

1  $\mathbf{w}^{(1)} \leftarrow (1/n, \dots, 1/n)$       ▷ *initial weights*

2 **for**  $t \leftarrow 1$  **to**  $T$

3  $h^{(t)} \leftarrow \text{BASE}(D_n, \mathbf{w}^{(t)})$       ▷ *calling the base learner*

4  $\gamma^{(t)} \leftarrow \sum_{i=1}^n w_i^{(t)} h^{(t)}(\mathbf{x}_i) y_i$       ▷ *edge = 1 - 2 × error*

5  $\alpha^{(t)} \leftarrow \frac{1}{2} \ln \left( \frac{1 + \gamma^{(t)}}{1 - \gamma^{(t)}} \right)$       ▷ *coefficient of  $h^{(t)}$*

6 **for**  $i \leftarrow 1$  **to**  $n$       ▷ *re-weighting the points*

7 **if**  $h^{(t)}(\mathbf{x}_i) \neq y_i$  **then**

8  $w_i^{(t+1)} \leftarrow w_i^{(t)} \frac{1}{1 - \gamma^{(t)}}$

9 **else**

10  $w_i^{(t+1)} \leftarrow w_i^{(t)} \frac{1}{1 + \gamma^{(t)}}$

11 **return**  $f^{(T)}(\cdot) = \sum_{t=1}^T \alpha^{(t)} h^{(t)}(\cdot)$

ADABOOST.BA( $D_n = \{(\mathbf{x}_i, y_i)\}_{i=1}^n$ , BASE( $\cdot, \cdot, \cdot$ ),  $T$ ,  $\mathcal{H}$ , BANDITALGO)

1  $\mathbf{w}^{(1)} \leftarrow (1/n, \dots, 1/n)$   $\triangleright$  *initial weights*

2 **for**  $t \leftarrow 1$  **to**  $T$

3  $j \leftarrow \text{BANDITALGO.getArm}()$

4  $h^{(t)} \leftarrow \text{BASE}(D_n, \mathbf{w}^{(t)}, \mathcal{H}_j)$   $\triangleright$  *calling the base learner*

5  $\gamma^{(t)} \leftarrow \sum_{i=1}^n w_i^{(t)} h^{(t)}(\mathbf{x}_i) y_i$   $\triangleright$  *edge = 1 - 2 × error*

6  $r_j^{(t)} = \min \left( 1, -\log \sqrt{1 - \gamma_{\mathcal{H}_j}^{(t)2}} \right)$   $\triangleright$  *calculate reward*

7  $\text{BANDITALGO.receiveReward}(j, r_j^{(t)})$

8  $\alpha^{(t)} \leftarrow \frac{1}{2} \ln \left( \frac{1 + \gamma^{(t)}}{1 - \gamma^{(t)}} \right)$   $\triangleright$  *coefficient of  $h^{(t)}$*

9  $\triangleright$  *re-weighting the points*

10 **return**  $f^{(T)}(\cdot) = \sum_{t=1}^T \alpha^{(t)} h^{(t)}(\cdot)$

# Multi-armed Bandit (MAB) Problem, Adversarial setup

- $(\blacksquare, \blacksquare, \blacksquare, \blacksquare, \blacksquare)$  indexed by  $\{1, \dots, M\}$

# Multi-armed Bandit (MAB) Problem, Adversarial setup

- $\left( \boxed{r_1^{(t)}}, \boxed{r_2^{(t)}}, \dots, \boxed{r_{M-1}^{(t)}}, \boxed{r_M^{(t)}} \right)$  indexed by  $\{1, \dots, M\}$
- An adversary chooses a reward vector  $\mathbf{r}^{(t)} \in \mathbb{R}^M$

# Multi-armed Bandit (MAB) Problem, Adversarial setup

- $\left( \blacksquare, \blacksquare, \blacksquare, r_{j^{(t)}}^{(t)}, \blacksquare \right)$  indexed by  $\{1, \dots, M\}$
- An adversary chooses a reward vector  $\mathbf{r}^{(t)} \in \mathbb{R}^M$
- Decision maker chooses an arm  $j^{(t)}$  and receives reward  $r_{j^{(t)}}^{(t)}$  in each iteration step (**bandit feedback**)

# Multi-armed Bandit (MAB) Problem, Adversarial setup

- $\left( \blacksquare, \blacksquare, \blacksquare, \color{green}r_{j^{(t)}}^{(t)}, \blacksquare \right)$  indexed by  $\{1, \dots, M\}$
- An adversary chooses a reward vector  $\mathbf{r}^{(t)} \in \mathbb{R}^M$
- Decision maker chooses an arm  $j^{(t)}$  and receives reward  $r_{j^{(t)}}^{(t)}$  in each iteration step (**bandit feedback**)
- The most common performance measure is the **weak regret**

$$R = \underbrace{\max_j \sum_{t=1}^T r_j^{(t)}}_{\text{Total reward of best arm}} - \underbrace{\sum_{t=1}^T r_{j^{(t)}}^{(t)}}_{G^{(T)} = \text{Total reward}}$$

# Multi-armed Bandit (MAB) Problem, Adversarial setup

- $\left( \blacksquare, \blacksquare, \blacksquare, \color{green}r_{j^{(t)}}^{(t)}, \blacksquare \right)$  indexed by  $\{1, \dots, M\}$
- An adversary chooses a reward vector  $\mathbf{r}^{(t)} \in \mathbb{R}^M$
- Decision maker chooses an arm  $j^{(t)}$  and receives reward  $r_{j^{(t)}}^{(t)}$  in each iteration step (**bandit feedback**)
- The most common performance measure is the **weak regret**

$$R = \underbrace{\max_j \sum_{t=1}^T r_j^{(t)}}_{\text{Total reward of best arm}} - \underbrace{\sum_{t=1}^T r_{j^{(t)}}^{(t)}}_{G^{(T)} = \text{Total reward}}$$

Total reward of best arm     $G^{(T)} =$  Total reward

- A theoretically well-founded algorithm: EXP3.P

# Multi-armed Bandit (MAB) Problem, Adversarial setup

- $\left( \blacksquare, \blacksquare, \blacksquare, \color{green} r_{j^{(t)}}^{(t)}, \blacksquare \right)$  indexed by  $\{1, \dots, M\}$
- An adversary chooses a reward vector  $\mathbf{r}^{(t)} \in \mathbb{R}^M$
- Decision maker chooses an arm  $j^{(t)}$  and receives reward  $r_{j^{(t)}}^{(t)}$  in each iteration step (**bandit feedback**)
- The most common performance measure is the **weak regret**

$$R = \underbrace{\max_j \sum_{t=1}^T r_j^{(t)}}_{\text{Total reward of best arm}} - \underbrace{\sum_{t=1}^T r_{j^{(t)}}^{(t)}}_{G^{(T)} = \text{Total reward}}$$

- A theoretically well-founded algorithm: EXP3.P



# Multi-armed Bandit (MAB) Problem, Adversarial setup

- $\left( \blacksquare, \blacksquare, \blacksquare, \color{green} r_{j^{(t)}}^{(t)}, \blacksquare \right)$  indexed by  $\{1, \dots, M\}$
- An adversary chooses a reward vector  $\mathbf{r}^{(t)} \in \mathbb{R}^M$
- Decision maker chooses an arm  $j^{(t)}$  and receives reward  $r_{j^{(t)}}^{(t)}$  in each iteration step (**bandit feedback**)
- The most common performance measure is the **weak regret**

$$R = \underbrace{\max_j \sum_{t=1}^T r_j^{(t)}}_{\text{Total reward of best arm}} - \underbrace{\sum_{t=1}^T r_{j^{(t)}}^{(t)}}_{G^{(T)} = \text{Total reward}}$$

- A theoretically well-founded algorithm: EXP3.P

Theorem (Auer et al.(1995))

With probability at least  $1 - \delta$

$$R \leq 4\sqrt{MT \log \frac{MT}{\delta}} + 4\sqrt{\frac{5}{3}MT \log M} + 8 \log \frac{MT}{\delta}$$

# Reward definition

Theorem (Schapire, Singer(1998))

For ADABOOST.MH

$$R(\mathbf{f}^{(T)}) = \frac{1}{n} \sum_{i=1}^n \mathbb{I} \{ \ell(\mathbf{x}_i) \neq \hat{\ell}(\mathbf{x}_i) \} \leq \sqrt{K-1} \prod_{t=1}^T \sqrt{1 - \gamma^{(t)2}}$$

# Reward definition

Theorem (Schapire, Singer(1998))

For ADABOOST.MH

$$R(\mathbf{f}^{(T)}) = \frac{1}{n} \sum_{i=1}^n \mathbb{I} \{ \ell(\mathbf{x}_i) \neq \hat{\ell}(\mathbf{x}_i) \} \leq \sqrt{K-1} \prod_{t=1}^T \sqrt{1 - \gamma(t)^2}$$

- reward should depend on  $\sqrt{1 - \gamma(t)^2}$

# Reward definition

## Theorem (Schapire, Singer(1998))

For ADABOOST.MH

$$R(\mathbf{f}^{(T)}) = \frac{1}{n} \sum_{i=1}^n \mathbb{I} \{ \ell(\mathbf{x}_i) \neq \hat{\ell}(\mathbf{x}_i) \} \leq \sqrt{K-1} \prod_{t=1}^T \sqrt{1 - \gamma(t)^2}$$

- reward should depend on  $\sqrt{1 - \gamma(t)^2}$
- sum of reward is optimized  $\log \sqrt{1 - \gamma(t)^2}$

# Reward definition

## Theorem (Schapire, Singer (1998))

For ADABOOST.MH

$$R(\mathbf{f}^{(T)}) = \frac{1}{n} \sum_{i=1}^n \mathbb{I} \{ \ell(\mathbf{x}_i) \neq \hat{\ell}(\mathbf{x}_i) \} \leq \sqrt{K-1} \prod_{t=1}^T \sqrt{1 - \gamma(t)^2}$$

- reward should depend on  $\sqrt{1 - \gamma(t)^2}$
- sum of reward is optimized  $\log \sqrt{1 - \gamma(t)^2}$
- actually, it is maximized  $-\log \sqrt{1 - \gamma(t)^2}$

# Reward definition

## Theorem (Schapire, Singer(1998))


For ADABOOST.MH

$$R(\mathbf{f}^{(T)}) = \frac{1}{n} \sum_{i=1}^n \mathbb{I} \{ \ell(\mathbf{x}_i) \neq \hat{\ell}(\mathbf{x}_i) \} \leq \sqrt{K-1} \prod_{t=1}^T \sqrt{1 - \gamma(t)^2}$$

- reward should depend on  $\sqrt{1 - \gamma(t)^2}$
- sum of reward is optimized  $\log \sqrt{1 - \gamma(t)^2}$
- actually, it is maximized  $-\log \sqrt{1 - \gamma(t)^2}$
- the rawrds must be bounded  $\min \left( 1, -\log \sqrt{1 - \gamma(t)^2} \right)$

# Reward definition

$$r = \min \left( 1, -\log \sqrt{1 - \gamma^2} \right)$$



# Convergence of ADABOOST.MH.EXP3.P

- Multiclass training error

$$R(\mathbf{f}^{(T)}) = \frac{1}{n} \sum_{i=1}^n \mathbb{I} \left\{ \ell(\mathbf{x}_i) \neq \hat{\ell}(\mathbf{x}_i) \right\}$$



# Convergence of ADABOOST.MH.EXP3.P

- Multiclass training error

$$R(\mathbf{f}^{(T)}) = \frac{1}{n} \sum_{i=1}^n \mathbb{I} \left\{ \ell(\mathbf{x}_i) \neq \widehat{\ell}(\mathbf{x}_i) \right\}$$

- Weak learnability:  $\gamma^{(t)} \geq \rho > 0$

# Convergence of ADABOOST.MH.EXP3.P

- Multiclass training error

$$R(\mathbf{f}^{(T)}) = \frac{1}{n} \sum_{i=1}^n \mathbb{I} \left\{ \ell(\mathbf{x}_i) \neq \widehat{\ell}(\mathbf{x}_i) \right\}$$

- Weak learnability:  $\gamma^{(t)} \geq \rho > 0$

# Convergence of ADABOOST.MH.EXP3.P

- Multiclass training error

$$R(\mathbf{f}^{(T)}) = \frac{1}{n} \sum_{i=1}^n \mathbb{I} \left\{ \ell(\mathbf{x}_i) \neq \widehat{\ell}(\mathbf{x}_i) \right\}$$

- Weak learnability:  $\gamma^{(t)} \geq \rho > 0$

## Theorem

With probability at least  $1 - \delta$ :

$R(\mathbf{f}^{(T)}) = 0$  after


$$T = \max \left( \log^2 \frac{M}{\delta}, \left( \frac{4C}{\rho^2} \right)^4, \frac{4 \log(n\sqrt{K-1})}{\rho^2} \right)$$

iterations, where

$$C = \sqrt{32M} + \sqrt{27M \log M} + 16$$

# Partitioning the base classifier set

- Decision stump: assign a **subset** to each **feature**:  
 $\mathcal{H}_j = \{\varphi_{j,b}(\mathbf{x}) : b \in \mathbb{R}\}$
- For Decision trees and products, the naive solution is to assign a subset of features at size  $N$  to an arm  $\implies$   
Number of arms grows exponentially
- Trees and products are modeled as sequences of decisions over the smaller partitioning used for stumps



# Experiments

- **Stochastic bandit algorithms**: assuming that the rewards are drawn from a stationary probability distribution, **UCB**<sup>3</sup>, **UCBV**<sup>4</sup>
- **Stochastic bandits** does not fit to our setup since the *edge depends on the weights*

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

- Random feature selection  $\approx$  LAZYBOOST
- Synthetic data, UCI datasets and MNIST handwritten digits recognition





---

<sup>3</sup>Auer, P., Cesa-Bianchi, N., and Fischer, P.: Finite-time analysis of the multiarmed bandit problem. *Machine Learning*, 47:235–256, 2002

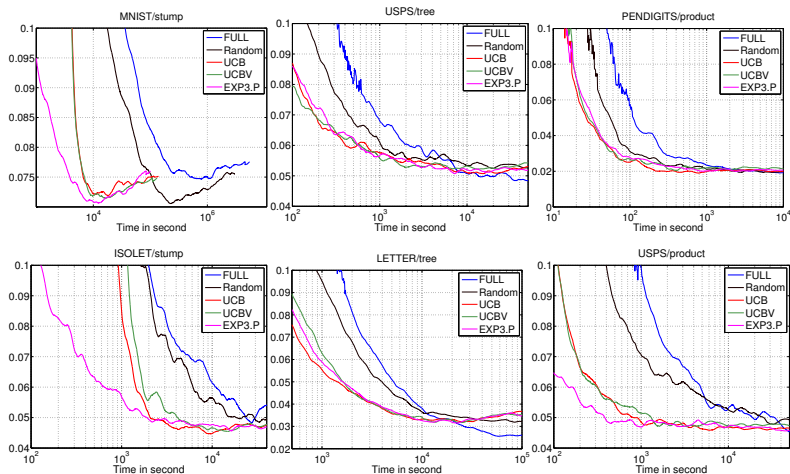
<sup>4</sup>Audibert, J.-Y., Munos, R., and Szepesvári, Cs.: Exploration-exploitation tradeoff using variance estimates in multi-armed bandits. *Theor. Comput. Sci.*, 410(19):1876–1902, 2009

# Synthetic datasets

- $d = 10$ , number of **useful** feature = 3, **stumps**,  $T = 1000$



# Test error vs. CPU time



## Concluding remarks

- our multiboost implementation [multiboost.org](http://multiboost.org) includes the bandit based setup
- ICML'10 [Yahoo Learning to Rank Challenge](#)
- [top ten](#) performance using [regression-calibrated bandit boosting](#)
- [High-dimensional, structured](#) feature spaces (linear, Haar): [continuous, "metric"](#) bandits?
- MABs are [stateless](#), boosting is not → [MDPs?](#)



EXP3.P( $\eta, \lambda, T$ )

1     **for**  $j \leftarrow 1$  **to**  $M$      ▷ *initialization*

2              $\omega_j^{(1)} \leftarrow \exp\left(\frac{\eta\lambda}{3}\sqrt{\frac{T}{M}}\right)$

3     **for**  $t \leftarrow 1$  **to**  $T$

4             **for**  $j \leftarrow 1$  **to**  $M$

5                      $p_j^{(t)} \leftarrow (1 - \lambda) \frac{\omega_j^{(t)}}{\sum_{j'=1}^M \omega_{j'}^{(t)}} + \frac{\lambda}{M}$

6              $j^{(t)} \leftarrow \text{RANDOM}(p_1^{(t)}, \dots, p_M^{(t)})$

7             Receive reward  $r_{j^{(t)}}^{(t)}$

8             **for**  $j \leftarrow 1$  **to**  $M$

9                      $\hat{r}_j^{(t)} \leftarrow \begin{cases} r_j^{(t)} / p_j^{(t)} & \text{if } j = j^{(t)} \\ 0 & \text{otherwise} \end{cases}$

10              $\omega_j^{(t+1)} \leftarrow \omega_j^{(t)} \exp\left(\frac{\lambda}{3M} \left(\hat{r}_j^{(t)} + \frac{\eta}{p_j^{(t)}\sqrt{MT}}\right)\right)$