

# A SIMPLE DYNAMIC BANDIT ALGORITHM FOR HYPER-PARAMETER TUNING

Xuedong Shang<sup>1,2</sup>, Emilie Kaufmann<sup>2,3</sup> and Michal Valko<sup>4</sup>

<sup>1</sup>Inria Sequel <sup>2</sup>Univ. Lille <sup>3</sup>CNRS <sup>4</sup>DeepMind Paris

## Problem and Objectives

We treat the **hyper-parameter tuning** problem for *supervised learning* tasks.

- global optimisation task:  $\min\{f(\boldsymbol{\lambda}) : \boldsymbol{\lambda} \in \Omega\}$ ;
- $f(\boldsymbol{\lambda}) \triangleq \mathbb{E}[\ell(\mathbf{Y}, g_{\boldsymbol{\lambda}}^{(n)}(\mathbf{X}))]$  measures the generalization power;
- goal**: a simple, robust, (almost) parameter-free bandit algorithm.

## How and Why

### How?

- We see the problem as a *stochastic infinitely many-armed bandit*
- Beta-Bernoulli bandit model
- A Beta reservoir  $\nu_0$  over the means of the arms
- A uniform prior  $\Pi_0$  over the arms  $\rightarrow$  posterior:

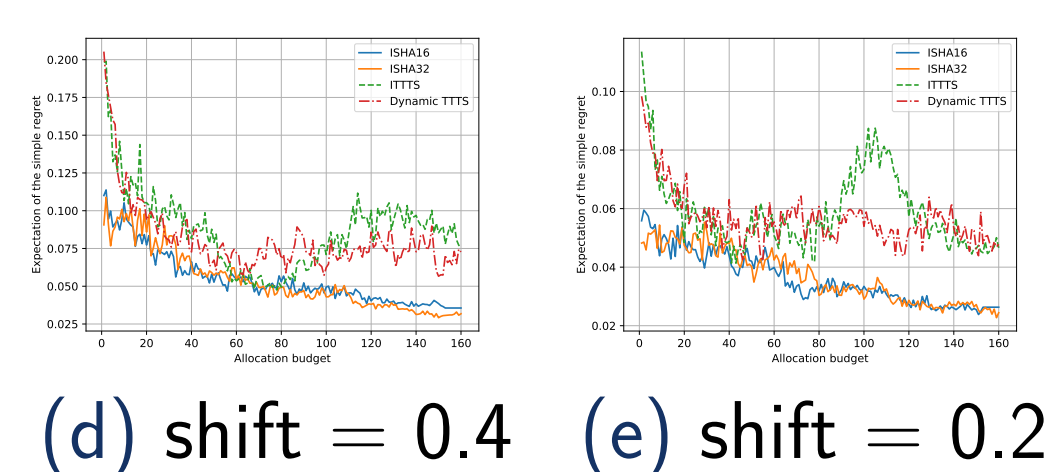
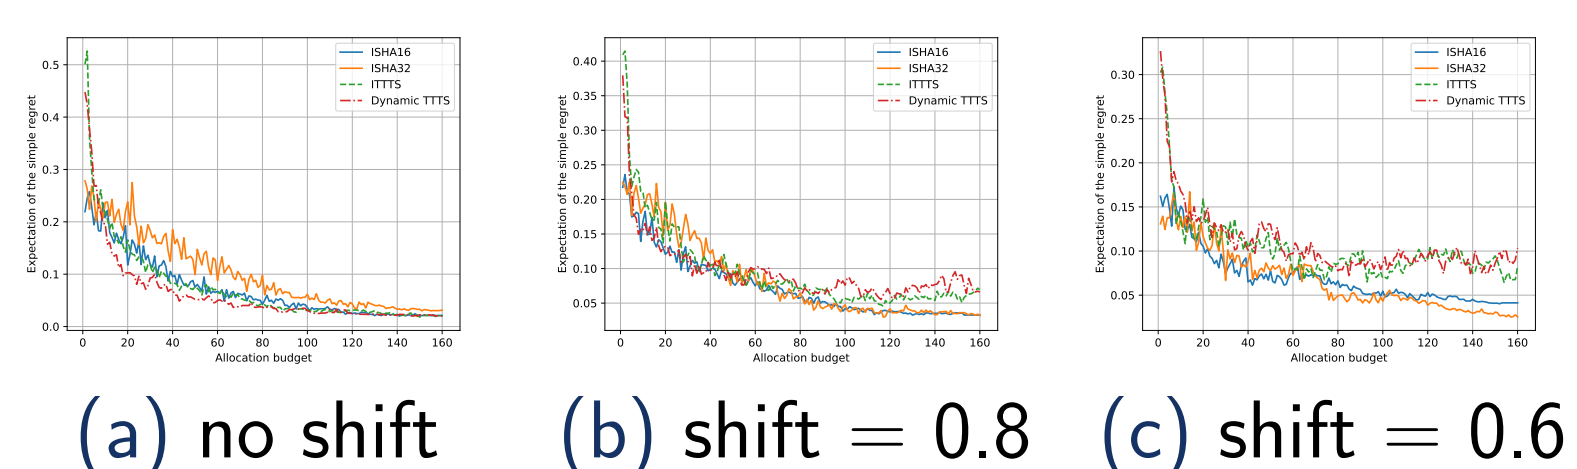
$$\Pi_t = \prod_{i=1}^{k_t} \text{Beta}(1 + S_{t,i}, N_{t,i} - S_{t,i} + 1)$$

- At each round, **D-TTTS** either **samples a new arm** or **re-samples a previous one**, and runs **TTTS** on the increasing set of arms

### Why?

- $\rightarrow$  **TTTS** is *anytime* for finitely-armed bandits
- $\rightarrow$  The number of arms added by **D-TTTS** depends on the difficulty of the task (the reservoir)
- $\rightarrow$  e.g. if the reservoir is difficult (like **Beta**(5,1)), the pseudo-arm  $\mu_0$  will be sampled more often
- $\rightarrow$  **D-TTTS** **does not** need to fix the number of arms sampled in advance, and naturally **adapts** to the difficulty of the task

**When it fails?**  $\leftarrow$  If  $\mu^* \neq 1$ ?



## Notation and Glossary

- $\Omega$  is the hyper-parameter space
- $\boldsymbol{\lambda}$  is a hyper-parameter configuration
- $g_{\boldsymbol{\lambda}}$  is a classifier
- $\ell$  is the cross-validation error in this work
- $\mu_1, \mu_2, \dots$  denote the true means
- $\Theta$  is the parameter space
- $\Theta_i \triangleq \{\theta \in \Theta \mid \theta_i > \max_{j \neq i} \theta_j\}$  is the subset of arm  $i$  being optimal

## Recommendation Rule

We recommend the arm with the largest *posterior probability of being optimal*:

$$\bar{I}_n \triangleq \arg \max_{i \in \mathcal{A}} \Pi_n(\Theta_i).$$

## Sampling Rule

```

1: Input:  $\beta$ 
2: Initialization:  $\mu_1 \sim \nu_0$ ;  $\mathcal{A} = \{\mu_1\}$ ;  $m = 1$ ;  $S_1, N_1 = 0$ 
3: while budget still available do
4:    $\mu_{m+1} \sim \nu_0$ ;  $\mathcal{A} \leftarrow \mathcal{A} \cup \{\mu_{m+1}\}$ 
5:    $S_{m+1}, N_{m+1} \leftarrow 0$ ;  $m \leftarrow m + 1$ 
6:    $\forall i \in \mathcal{A}, \theta_i \sim \text{Beta}(S_i + 1, N_i - S_i + 1)$ 
7:    $I^{(1)} = \arg \max_{i=0, \dots, m} \theta_i$ 
8:   if  $U(\sim \mathcal{U}([0, 1])) > \beta$  then
9:     while  $I^{(2)} \neq I^{(1)}$  do
10:       $\forall i \in \mathcal{A}, \theta'_i \sim \text{Beta}(S_i + 1, N_i - S_i + 1)$ 
11:       $I^{(2)} \leftarrow \arg \max_{i=0, \dots, m} \theta'_i$ 
12:     end while
13:     $I^{(1)} \leftarrow I^{(2)}$ 
14:   end if
15:    $Y \leftarrow \text{evaluate arm } I^{(1)}$ ;  $X \sim \text{Ber}(Y)$ 
16:    $S_{I^{(1)}} \leftarrow S_{I^{(1)}} + X$ ;  $N_{I^{(1)}} \leftarrow N_{I^{(1)}} + 1$ 
17: end while

```

## Some Tricks

- Binarization trick:** When a reward  $Y_{t,i} \in [0, 1]$  is observed, the algorithm is updated with a fake reward  $Y'_{t,i} \sim \text{Ber}(Y_{t,i} \in \{0, 1\})$ .
- Order statistic trick:** At time  $t$ , let  $\mathcal{L}_{t-1}$  be the list of arms that have been efficiently sampled, we run **TTTS** on the set  $\mathcal{L}_{t-1} \cup \{\mu_0\}$  where  $\mu_0$  is a pseudo-arm with posterior distribution **Beta**( $t - |\mathcal{L}_{t-1}|, 1$ ).

## Experimental Setting

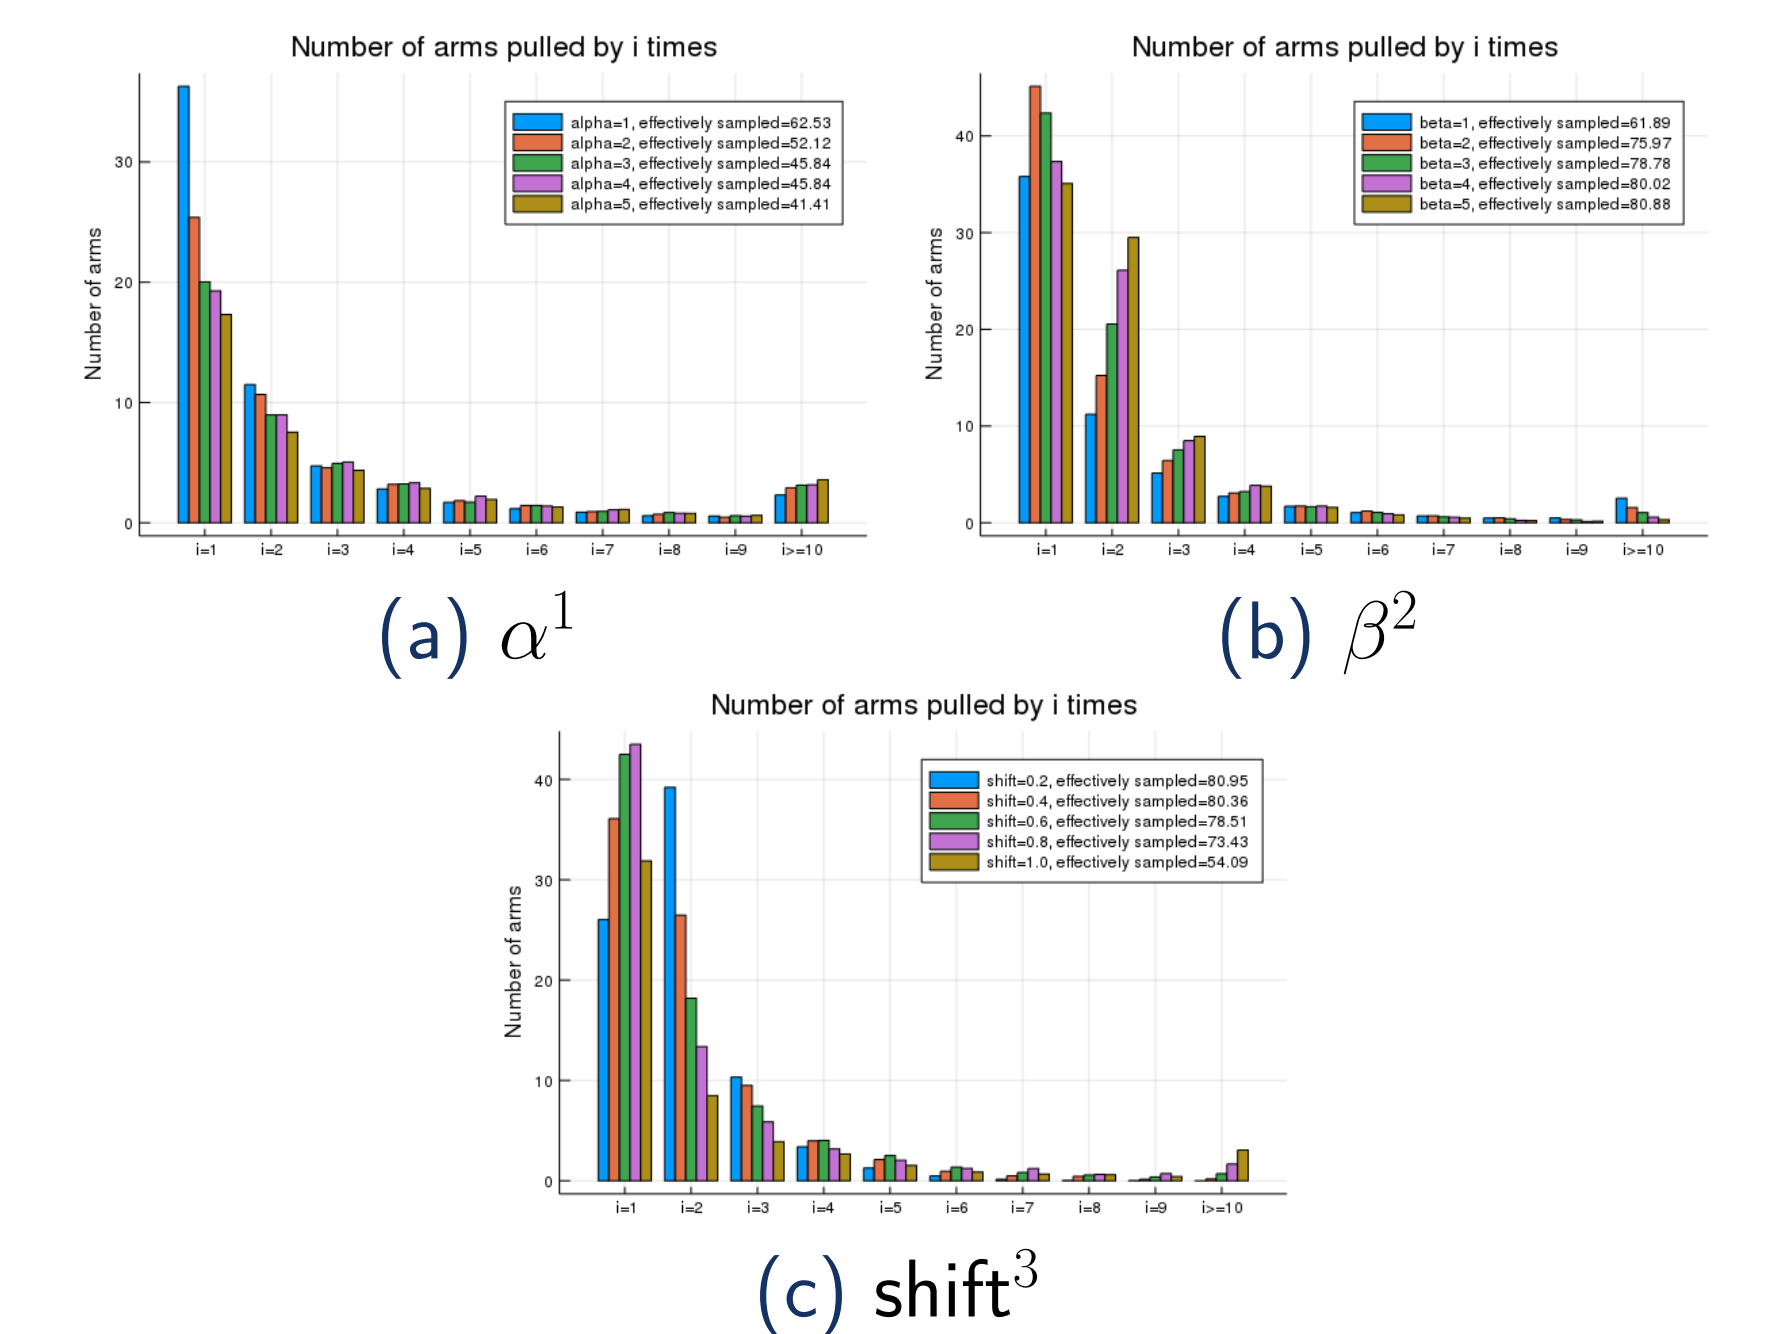
Classifier	Hyper-parameter	Type	Bounds
SVM	$C$	$\mathbb{R}^+$	$[10^{-5}, 10^5]$
	$\gamma$	$\mathbb{R}^+$	$[10^{-5}, 10^5]$

Table: hyper-parameters to be tuned for UCI experiments.

Classifier	Hyper-parameter	Type	Bounds
MLP	hidden_layer_size	Integer	$[5, 50]$
	alpha	$\mathbb{R}^+$	$[0, 0.9]$
	learning_rate_init	$\mathbb{R}^+$	$[10^{-5}, 10^{-1}]$

Table: hyper-parameters to be tuned for MNIST experiments.

## Illustrations of Efficiently Sampled Arms



- efficiently sampled arms for **Beta**( $\alpha, 1$ ) reservoirs
- efficiently sampled arms for **Beta**(1,  $\beta$ ) reservoirs
- efficiently sampled arms for shifted **Beta** reservoirs

## References

- Daniel Russo. Simple Bayesian algorithms for best arm identification. In *Proceedings of the 29th Conference on Learning Theory (CoLT)*, 2016.
- Lisha Li, Kevin Jamieson, Giulia DeSalvo, Amey Talwalkar, and Afshin Rostamizadeh. Hyperband: Bandit-based configuration evaluation for hyperparameter optimization. In *Proceedings of the 5th International Conference on Learning Representations (ICLR)*, 2017.

## Acknowledgements

The research presented was supported by European CHIST-ERA project DELTA, French Ministry of Higher Education and Research, Nord-Pas-de-Calais Regional Council, Inria and Otto-von-Guericke-Universität Magdeburg associated-team north-European project Allocate, and French National Research Agency projects BADASS (n.ANR-16-CE40-0002), and BoB (n.ANR-16-CE23-0003).

## Results for HPO

