


## Near-linear Time Gaussian Process Optimization with Adaptive Batching and Resparsification

D. Calandriello\* <sup>1</sup>, **L. Carratino\*** <sup>2</sup>, A. Lazaric <sup>3</sup>, M. Valko <sup>1</sup>, L. Rosasco <sup>2,4</sup>

\* equal contribution. <sup>1</sup> DeepMind, <sup>2</sup> MaLGA - UniGe, <sup>3</sup> Facebook, <sup>4</sup> MIT - IIT

# Bayesian/Bandit Optimization

Set of candidates  $\mathcal{A}$




# Bayesian/Bandit Optimization

Set of candidates  $\mathcal{A}$

for  $t = 1, \dots, T$ :

- (1) Select candidate
- (2) Receive noisy feedback
- (3) Update model



# Bayesian/Bandit Optimization

Set of candidates  $\mathcal{A}$

for  $t = 1, \dots, T$ :

- (1) Select candidate
- (2) Receive noisy feedback
- (3) Update model




# Bayesian/Bandit Optimization

Set of candidates  $\mathcal{A}$

for  $t = 1, \dots, T$ :

- (1) Select candidate
- (2) Receive noisy feedback
- (3) Update model




# Bayesian/Bandit Optimization

Set of candidates  $\mathcal{A}$

for  $t = 1, \dots, T$ :

- (1) Select candidate
- (2) Receive noisy feedback
- (3) Update model




# Bayesian/Bandit Optimization

Set of candidates  $\mathcal{A}$

for  $t = 1, \dots, T$ :

- (1) Select candidate
- (2) Receive noisy feedback
- (3) Update model




# Bayesian/Bandit Optimization

Set of candidates  $\mathcal{A}$

for  $t = 1, \dots, T$ :

- (1) Select candidate
- (2) Receive noisy feedback
- (3) Update model






# Bayesian/Bandit Optimization

Set of candidates  $\mathcal{A}$

for  $t = 1, \dots, T$ :

- (1) Select candidate
- (2) Receive noisy feedback
- (3) Update model




# Bayesian/Bandit Optimization

Set of candidates  $\mathcal{A}$

for  $t = 1, \dots, T$ :

- (1) Select candidate
- (2) Receive noisy feedback
- (3) Update model




# Bayesian/Bandit Optimization

Set of candidates  $\mathcal{A} = \{x_1, \dots, x_A\} \subset \mathbb{R}^d$ ,

for  $t = 1, \dots, T$ :

- (1) Select candidate
- (2) Receive noisy feedback
- (3) Update model




# Bayesian/Bandit Optimization

Set of candidates  $\mathcal{A} = \{x_1, \dots, x_A\} \subset \mathbb{R}^d$ ,  
**unknown** reward function  $f: \mathcal{A} \rightarrow \mathbb{R}$

for  $t = 1, \dots, T$ :

- (1) Select candidate
- (2) Receive noisy feedback
- (3) Update model




# Bayesian/Bandit Optimization

Set of candidates  $\mathcal{A} = \{x_1, \dots, x_A\} \subset \mathbb{R}^d$ ,  
**unknown** reward function  $f : \mathcal{A} \rightarrow \mathbb{R}$

for  $t = 1, \dots, T$ :

- (1) Select candidate  $x_t$  using model  $u_t$  (ideally  $u_t \approx f$ )
- (2) Receive noisy feedback
- (3) Update model




# Bayesian/Bandit Optimization

Set of candidates  $\mathcal{A} = \{x_1, \dots, x_A\} \subset \mathbb{R}^d$ ,  
**unknown** reward function  $f : \mathcal{A} \rightarrow \mathbb{R}$

for  $t = 1, \dots, T$ :

- (1) Select candidate  $x_t$  using model  $u_t$  (ideally  $u_t \approx f$ )
- (2) Receive noisy feedback  $y_t = f(x_t) + \eta_t$
- (3) Update model




# Bayesian/Bandit Optimization

Set of candidates  $\mathcal{A} = \{x_1, \dots, x_A\} \subset \mathbb{R}^d$ ,  
**unknown** reward function  $f: \mathcal{A} \rightarrow \mathbb{R}$

for  $t = 1, \dots, T$ :

- (1) Select candidate  $x_t$  using model  $u_t$  (ideally  $u_t \approx f$ )
- (2) Receive noisy feedback  $y_t = f(x_t) + \eta_t$
- (3) Update model  $u_t$



# Bayesian/Bandit Optimization


Set of candidates  $\mathcal{A} = \{x_1, \dots, x_A\} \subset \mathbb{R}^d$ ,  
**unknown** reward function  $f: \mathcal{A} \rightarrow \mathbb{R}$

for  $t = 1, \dots, T$ :

- (1) Select candidate  $x_t$  using model  $u_t$  (ideally  $u_t \approx f$ )
- (2) Receive noisy feedback  $y_t = f(x_t) + \eta_t$
- (3) Update model  $u_t$

**Performance** measure: cumulative **regret** w.r.t. best  $x_*$

$$R_T = \sum_{t=1}^T f(x_*) - f(x_t).$$





# Bayesian/Bandit Optimization

Set of candidates  $\mathcal{A} = \{x_1, \dots, x_A\} \subset \mathbb{R}^d$ ,  
**unknown** reward function  $f: \mathcal{A} \rightarrow \mathbb{R}$


for  $t = 1, \dots, T$ :

- (1) Select candidate  $x_t$  using model  $u_t$  (ideally  $u_t \approx f$ )
- (2) Receive noisy feedback  $y_t = f(x_t) + \eta_t$
- (3) Update model  $u_t$

**Performance** measure: cumulative **regret** w.r.t. best  $x_*$


$$R_T = \sum_{t=1}^T f(x_*) - f(x_t).$$

💡 Use Gaussian process/kernelized Bandit to model  $f$




# Gaussian Process Optimization

😊 Well studied: exploration *vs* exploitation → no-regret (low error) 😊



# Gaussian Process Optimization


😊 Well studied: exploration *vs* exploitation → no-regret (low error) 😊



🤔 performance *vs* scalability ? 🤔

# Gaussian Process Optimization

😊 Well studied: exploration **vs** exploitation → no-regret (low error) 😊



😊 performance **vs** scalability ? 😊

😊 Batch BKB: no-regret **and** scalable 😊


# Why Scalable GP Optimization is Hard

Experimental  
scalability

Computational  
scalability

# Why Scalable GP Optimization is Hard

Experimental  
scalability



sequential

vs




batch

Computational  
scalability

# Why Scalable GP Optimization is Hard

Experimental  
scalability



sequential

vs



batch

Computational  
scalability



exact GP


vs



approximate GP

# Why Scalable GP Optimization is Hard

Experimental  
scalability



sequential

vs



batch

Computational  
scalability



exact GP

vs

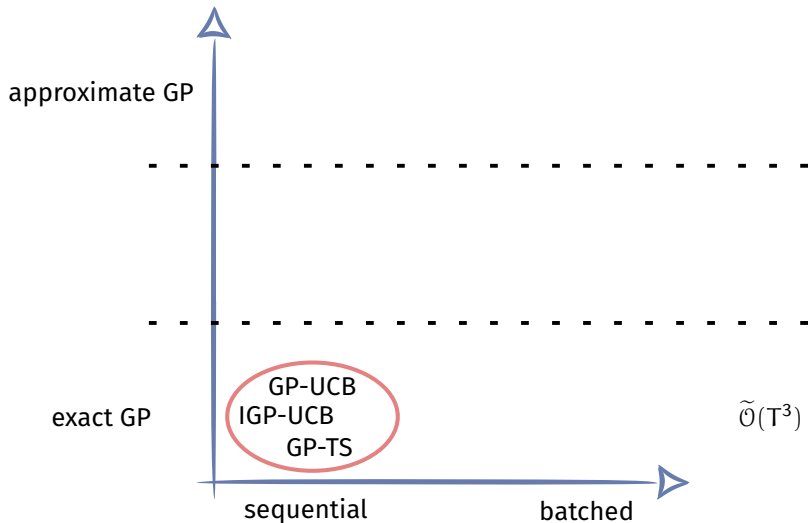


approximate GP

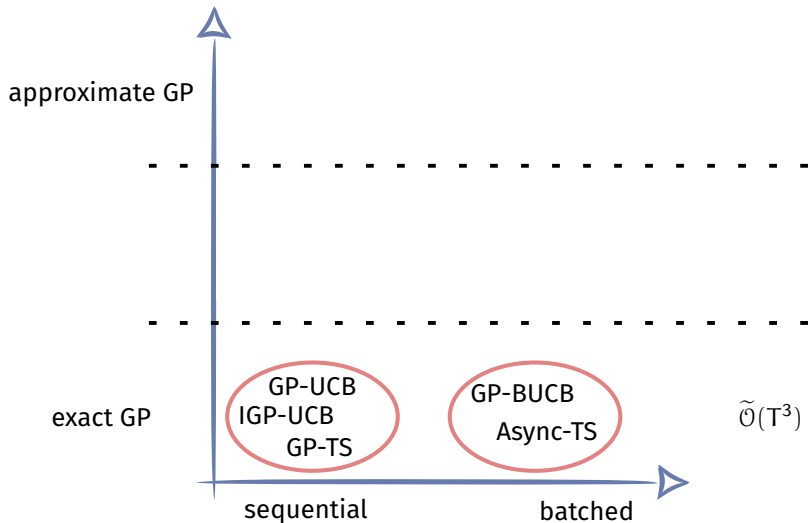
☹️ Batching and approximations increase regret ☹️



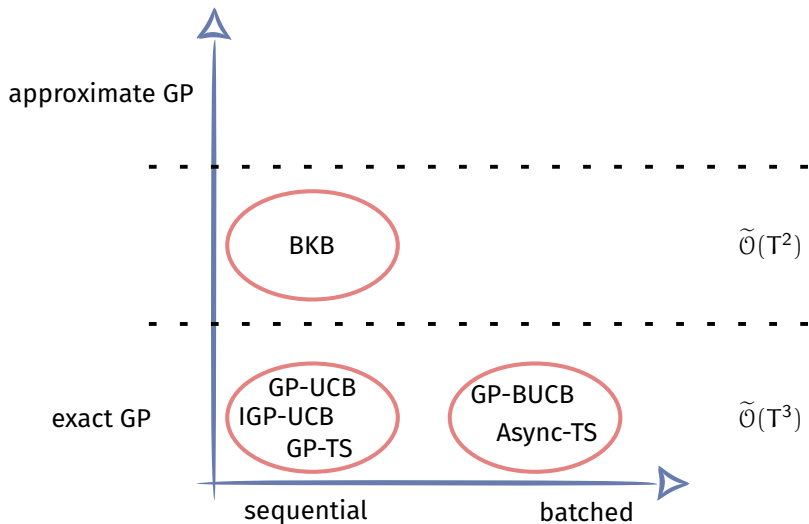
# Landscape of No-Regret GP Optimization



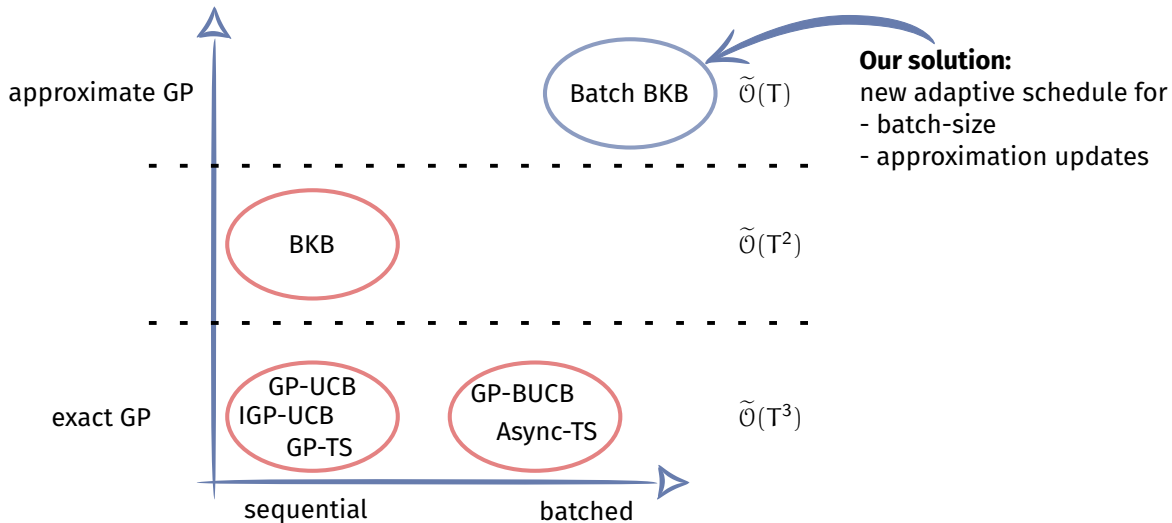
# Landscape of No-Regret GP Optimization




# Landscape of No-Regret GP Optimization



# Landscape of No-Regret GP Optimization



## Choosing good candidates with GP-UCB




## Choosing good candidates with GP-UCB

$$X_t = \{x_1, \dots, x_t\}, Y_t = \{y_1, \dots, y_t\}$$

**Exact GP-UCB:**

$$u_t(\cdot) = \mu(\cdot | X_t, Y_t)$$




## Choosing good candidates with GP-UCB

$$X_t = \{x_1, \dots, x_t\}, Y_t = \{y_1, \dots, y_t\}$$

**Exact GP-UCB:**

$$u_t(\cdot) = \mu(\cdot | X_t, Y_t) + \beta_t \sigma(\cdot | X_t)$$




## Choosing good candidates with GP-UCB

$$X_t = \{x_1, \dots, x_t\}, Y_t = \{y_1, \dots, y_t\}$$

**Exact GP-UCB:**

$$u_t(\cdot) = \mu(\cdot | X_t, Y_t) + \beta_t \sigma(\cdot | X_t)$$

[Sri+10]:  $u_t$  valid UCB.





## Choosing good candidates with GP-UCB

$$X_t = \{x_1, \dots, x_t\}, Y_t = \{y_1, \dots, y_t\}$$

### Exact GP-UCB:


$$u_t(\cdot) = \mu(\cdot | X_t, Y_t) + \beta_t \sigma(\cdot | X_t)$$

[Sri+10]:  $u_t$  valid UCB.

### Sparse GP-UCB:

$$\tilde{u}_t(\cdot) = \tilde{\mu}(\cdot | X_t, Y_t, \mathcal{D}_t) + \tilde{\beta}_t \tilde{\sigma}(\cdot | X_t, \mathcal{D}_t)$$

with  $\mathcal{D}_t \subset X_t$  inducing points



## Choosing good candidates with GP-UCB

$$X_t = \{x_1, \dots, x_t\}, Y_t = \{y_1, \dots, y_t\}$$

### Exact GP-UCB:

$$u_t(\cdot) = \mu(\cdot | X_t, Y_t) + \beta_t \sigma(\cdot | X_t)$$


[Sri+10]:  $u_t$  valid UCB.

### Sparse GP-UCB:

$$\tilde{u}_t(\cdot) = \tilde{\mu}(\cdot | X_t, Y_t, \mathcal{D}_t) + \tilde{\beta}_t \tilde{\sigma}(\cdot | X_t, \mathcal{D}_t)$$

with  $\mathcal{D}_t \subset X_t$  **inducing points**

[Cal+19]:  $\tilde{u}_t$  valid UCB if  $\mathcal{D}_t$  updated at every  $t$ .



## Performance vs Scalability

Better performance: collect more feedback,      update inducing points (resparsify)

$$\tilde{u}_t(\cdot) = \tilde{\mu}(\cdot | X_t, Y_t, \mathcal{D}_t) + \tilde{\beta}_t \tilde{\sigma}(\cdot | X_t, \mathcal{D}_t)$$

## Performance vs Scalability

Better performance: collect more feedback, update inducing points (resparsify)

$$\tilde{u}_t(\cdot) = \tilde{\mu}(\cdot | X_t, Y_t, \mathcal{D}_t) + \tilde{\beta}_t \tilde{\sigma}(\cdot | X_t, \mathcal{D}_t)$$

Worse scalability: experimental cost, resparsification cost

# Performance vs Scalability

Better performance: collect more feedback, update inducing points (resparsify)

$$\tilde{u}_t(\cdot) = \tilde{\mu}(\cdot | X_t, Y_t, \mathcal{D}_t) + \tilde{\beta}_t \tilde{\sigma}(\cdot | X_t, \mathcal{D}_t)$$

Worse scalability: experimental cost, resparsification cost

Improve scalability: batching feedback (GP-BUCB), batching resparsification ?

# Delayed Resparsification

## New adaptive batching rule

no-resparsify until  $\sum_{i \in \text{Batch}} \tilde{\sigma}^2(x_i) \geq 1$

# Delayed Resparsification

## New adaptive batching rule

no-resparsify until  $\sum_{i \in \text{Batch}} \tilde{\sigma}^2(x_i) \geq 1$

**“Not too big” Lemma:** valid UCB


# Delayed Resparsification

## New adaptive batching rule

no-resparsify until  $\sum_{i \in \text{Batch}} \tilde{\sigma}^2(x_i) \geq 1$

“Not too big” Lemma: valid UCB

“Not too small” Lemma: batch-size =  $\Omega(t)$





## Batch-BKB

### Theorem

With high probability **Batch-BKB** achieves *no-regret* with time complexity  $O(Td_{\text{eff}}^2)$ , where  $d_{\text{eff}} \ll T$  is the effective dimension / degrees of freedom of the GP.

# Batch-BKB

## Theorem

With high probability **Batch-BKB** achieves *no-regret* with time complexity  $O(Td_{\text{eff}}^2)$ , where  $d_{\text{eff}} \ll T$  is the effective dimension / degrees of freedom of the GP.


## Comparisons:

- 😊 Same regret of GP-UCB/IGP-UCB and better scalability (from  $O(T^3)$  to  $O(Td_{\text{eff}}^2)$ )
- 😊 Larger batches than GP-BUCB
- 😊 Better regret and better scalability than async-TS

## In practice: Scalability

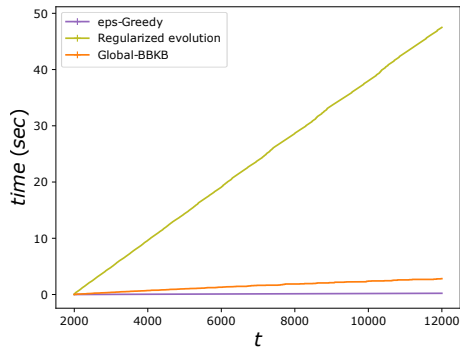
### Cadata

$A = 20640, d = 8, T = 2000$




### NAS-bench-101

$A = 12416, d = 19, T = 12000$




# In practice: Performance

NAS-bench-101  
Regret / Regret uniform



NAS-bench-101  
Simple Regret



# Thank you