Online Learning for the Black-Box Optimization of Wireless Networks

Public PhD Defense

September 7th, 2023

Presented by **Anthony Bardou**

Jury Members

President: Liva Ralaivola (AMU & Critéo)

Examiners: Patrick Loiseau (Inria & École Polytechnique)
Giovanni Neglia (Inria)
Claire Vernade (UTübingen)

Supervisor: Thomas Begin (UCBL)

Wireless Network

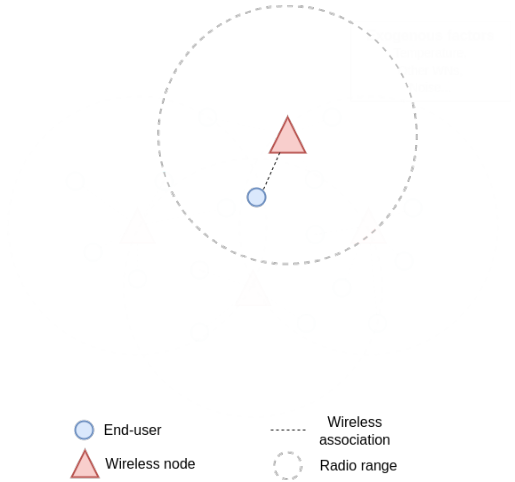

- Performance metrics (e.g. users throughput) easy to describe

Wireless Network

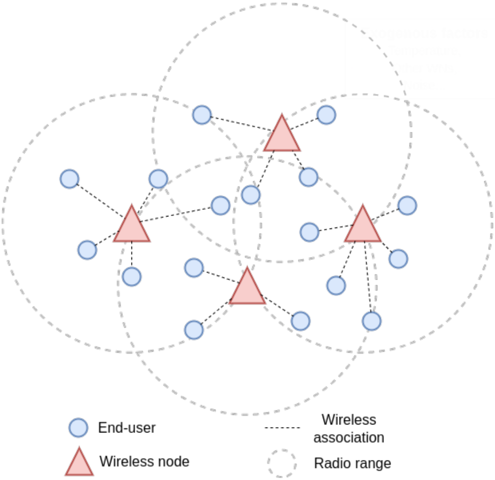

- Performance metrics (e.g. users throughput) hard to describe

Wireless Network

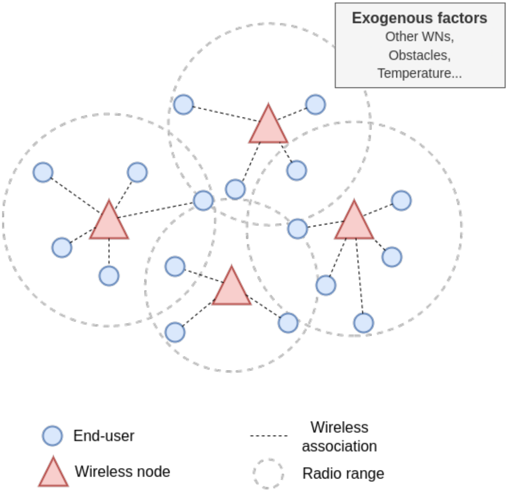

- Performance metrics (e.g. users throughput) *very* hard to describe

Optimization of a Wireless Network

- The network has parameters and $f : \mathcal{C} \to \mathbb{R}$ an objective function
- **Goal**: tune the parameters to maximize the objective f

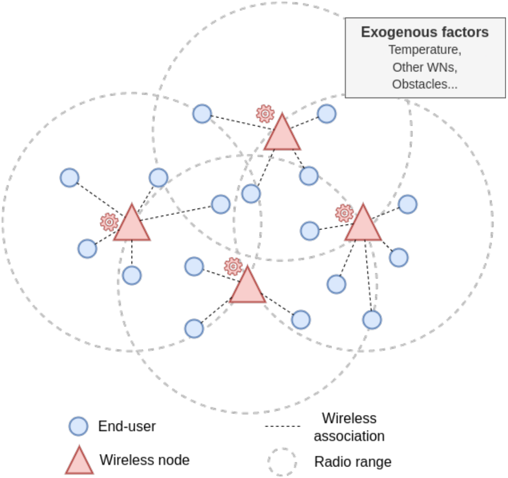

Online, Black-Box Optimization of a Wireless Network?

- **Black-Box**: the closed form of f is unknown (or does not exist)
  - Only noisy-corrupted f-values are observable by query
- **Online**: the learning data is collected during the optimization process

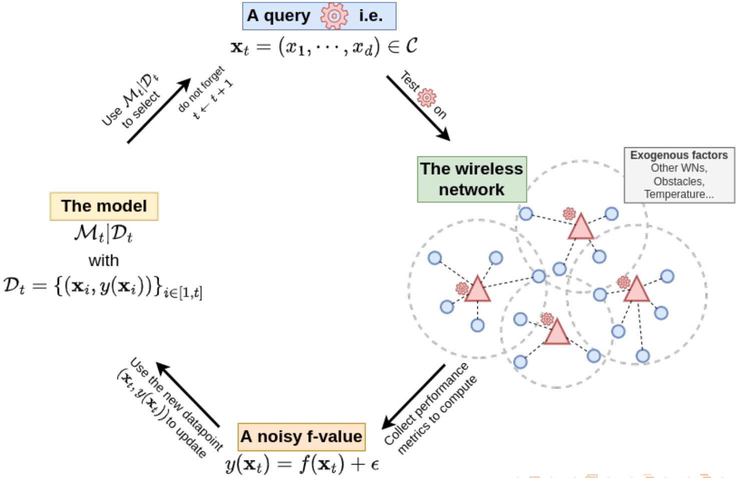

The Notion of Regret

- $\boldsymbol{x}^* = \arg\max_{\boldsymbol{x} \in \mathcal{C}} f(\boldsymbol{x})$
- **Instantaneous regret** at time t:

$$r_t = f(\boldsymbol{x}^*) - f(\boldsymbol{x}_t) \tag{1}$$

- **Cumulative regret** at time t:

$$R_t = \sum_{k=1}^{t} r_k \tag{2}$$

- **Asymptotic optimality**

$$\lim_{t \to +\infty} \frac{R_t}{t} = 0 \tag{3}$$

- Given enough time, $\boldsymbol{x}^*$ will be the most queried configuration (by far!)

Example: Spatial Reuse Optimization in WLANs

- **Maximize an objective function in a WLAN** (e.g. Wi-Fi network)
  - Each 📍 is an *access point* (AP)
  - APs serve *stations* (STAs)
- Each AP i has two parameters denoted $\boldsymbol{x}^{(i)} \in \mathcal{C}^{(i)}$
  - *Transmission power* (`TX_PWR`), *sensibility threshold* (`OBSS_PD`)
  - Dynamical update with IEEE 802.11ax amendment [1] (Wi-Fi 6)
- Spatial reuse optimization
  - is hard
  - must be addressed in next-generation WLANs

[1] "IEEE Standard for Information Technology–Telecommunications and Information Exchange between Systems - Local and Metropolitan Area Networks–Specific Requirements - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications". In: *IEEE Std 802.11-2020 (Revision of IEEE Std 802.11-2016)* (2021).

Outline

- **Introduction**
- **Contribution 1**: Multi-Armed Bandit Approaches for 802.11
- **Contribution 2**: Decentralized Bayesian Optimization for 802.11
- **Contribution 3**: Assessing the Benefits of NOMA for Next-Generation Cellular Networks
  - Collaboration with Jean-Marie Gorce, INSA Lyon
- **Contribution 4**: Decentralized, No-Regret Bayesian Optimization of High-Dimensional Functions
  - Collaboration with Patrick Thiran, EPFL
- **Conclusion**

Contribution 1

Multi-Armed Bandit Approaches for 802.11

Context

- WLAN with n APs, m STAs
- Each AP has two discrete parameters, with 20 values each
  - $|\mathcal{C}| = 20^{2n}$ configurations
  - Can be reduced to $|\mathcal{C}| = 200^{n}$ by integrating an IEEE 802.11ax constraint

- **Assumption**. The WLAN is equipped with a controller able to gather the throughputs of the STAs
- **Assumption**. The APs and the STAs do not move in space
  - Mild assumption for stadiums, open-spaces and M2M networks

- $f : \mathcal{C} \to \mathbb{R}$ is an ad-hoc objective function maximized when there is no starvation in the WLAN
  - A STA is said in starvation when its throughput is lower than a given threshold

- The problem is framed as a Multi-Armed Bandit (MAB)
  - Each configuration $\boldsymbol{x}$ is an arm, returning a reward $f(\boldsymbol{x}) + \epsilon$ when pulled

Overview

- $\mathcal{C}$ is too large to explore all the arms in a reasonable amount of time
- To overcome this, we propose two algorithms:
  - **The sampler** must explore the configurations space and gather configurations that appear promising configurations in a reservoir
  - **The optimizer** must identify the best configuration in the reservoir (Thompson sampling [2])

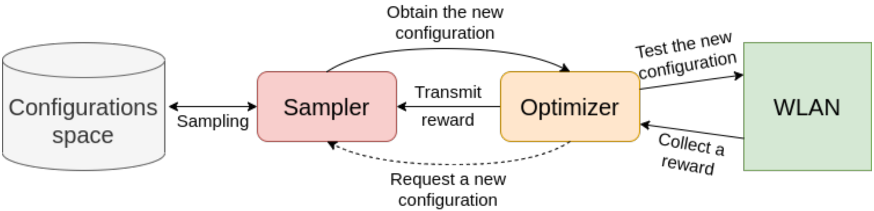

[2] William R Thompson. "On the likelihood that one unknown probability exceeds another in view of the evidence of two samples". In: *Biometrika* 25.3-4 (1933), pp. 285–294.

Building the Reservoir

- **State-of-the-art**: Uniform sampling in $\mathcal{C}$
- **Assumption** (regularity).

$$\exists L > 0, \forall \boldsymbol{x}_i, \boldsymbol{x}_j \in \mathcal{C}, ||\boldsymbol{x}_i - \boldsymbol{x}_j||_1 = 1 \implies |f(\boldsymbol{x}_i) - f(\boldsymbol{x}_j)| < L \quad (4)$$

- Two proposed samplers

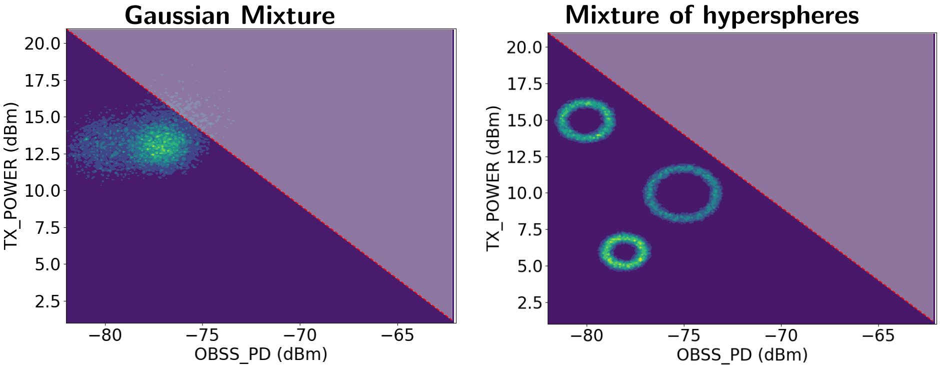

Evaluation

- Evaluation through simulation with ns-3 [3], 22 replications
- Evaluation on a real-world based scenario (selection)

Variable MCS, uplink/downlink traffic

10 APs, 50 STAs

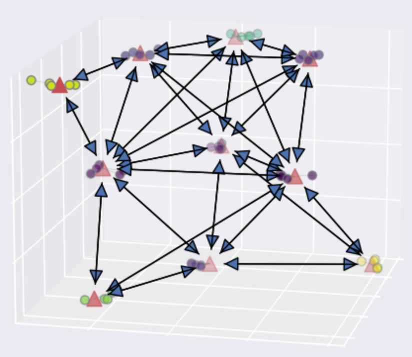

[3] The ns3 Project. *The Network Simulator ns-3*. https://www.nsnam.org/. Accessed: 2021-09-30. 2020.

Evaluation

- Control strategies: DEFAULT
- SOTA strategies: WCNC'15 [4], JNCA'19 [5], GM+NGTS, HM+NGTS
- 75 ms per test, 120 seconds simulated $\implies$ 1,600 iterations

Strategy	**Average Regret** R_t/t
DEFAULT	0.632 $\pm$0.001
WCNC'15	0.438 $\pm$0.001
JNCA'19	0.399 $\pm$0.001
GM+NGTS	0.472 $\pm$0.021
HM+NGTS	**0.237** $\pm$**0.011**

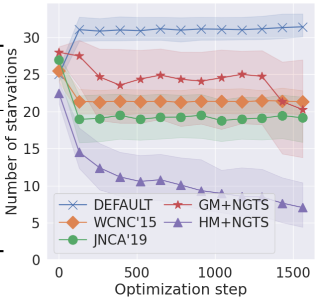

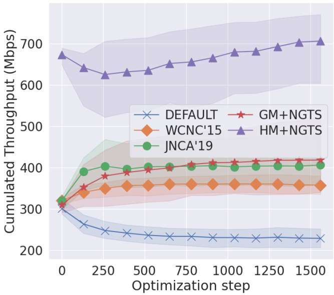

[4] M Shahwaiz Afaqui et al. "Evaluation of dynamic sensitivity control algorithm for IEEE 802.11 ax". In: *2015 IEEE wireless communications and networking conference (WCNC)*. IEEE. 2015, pp. 1060–1065.

[5] Francesc Wilhelmi et al. "Collaborative spatial reuse in wireless networks via selfish multi-armed bandits". In: *Ad Hoc Networks* 88 (2019), pp. 129–141.

Discussion

- **Pros**
  - Two solutions competitive against state-of-the-art strategies
  - Robust evaluation (complex scenarios, credible number of APs)
- **Cons**
  - No theoretical guarantees
  - Many hyperparameters to set
  - Centralized solution

Contribution 2

Decentralized Bayesian Optimization for 802.11

Context

- WLAN with n APs, m STAs
- Each AP i has two continuous parameters, within $\mathcal{C}^{(i)} = [-82, -62] \times [1, 21]$
  - $\mathcal{C} = \mathcal{C}^{(1)} \times \cdots \times \mathcal{C}^{(n)}$
  - $d = \dim \mathcal{C} = 2n$

- **Assumption**. The APs and the STAs do not move in space

- $f : \mathcal{C} \to \mathbb{R}^+$ is built on the proportional fairness of the STAs' throughputs $T(\boldsymbol{x}) = (T_1(\boldsymbol{x}), \cdots, T_m(\boldsymbol{x}))$

$$f(\boldsymbol{x}) = \sum_{i=1}^{m} \log T_i(\boldsymbol{x}) \quad (5)$$

- $\arg\max_{\boldsymbol{x} \in \mathcal{C}} f(\boldsymbol{x})$ is a natural trade-off between
  - a large cumulated throughput $||T(\boldsymbol{x})||_1$
  - a large fairness index $\frac{||T(\boldsymbol{x})||_1^2}{m||T(\boldsymbol{x})||_2^2}$ [6]

[6] Rajendra K Jain, Dah-Ming W Chiu, William R Hawe, et al. "A quantitative measure of fairness and discrimination". In: *Eastern Research Laboratory, Digital Equipment Corporation, Hudson, MA* 21 (1984).

Decentralized Method

- **Assumption**. Each AP can only communicate with APs in its radio range
  - $\mathcal{N}_i$: indices of reachable APs for AP i, including i itself
  - $\mathcal{S}_i$: the STAs associated with AP i

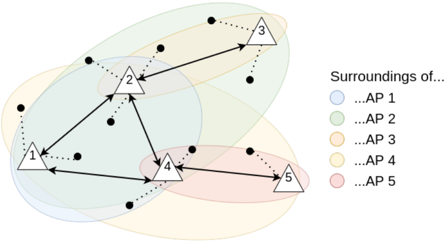

- Additive decomposition of f such as $f(\boldsymbol{x}) = \sum_{i=1}^{n} f^{(i)}(\boldsymbol{x})$
  - $f^{(i)}(\boldsymbol{x}) = \sum_{j \in \mathcal{S}_i} \log T_j(\boldsymbol{x})$?
  - $f^{(i)}(\boldsymbol{x}) = \sum_{j \in \mathcal{N}_i} \frac{1}{|\mathcal{N}_j|} \sum_{k \in \mathcal{S}_j} \log T_k(\boldsymbol{x})$ exploits the whole local information

Bayesian Optimization

- **Assumption**. $\forall i \in [1, n]$, $f^{(i)}$ is a Gaussian Process $\mathcal{GP}\left(0, k^{(i)}\left(\boldsymbol{x}_{\mathcal{N}_i}, \boldsymbol{x}'_{\mathcal{N}_i}\right)\right)$
  - $\boldsymbol{x}_{\mathcal{N}_i} \in \mathcal{C}_{\mathcal{N}_i} = \prod_{j \in \mathcal{N}_i} \mathcal{C}^{(j)}$
  - $\forall \boldsymbol{x} \in \mathcal{C}_{\mathcal{N}_i}, f^{(i)}(\boldsymbol{x}) \sim \mathcal{N}\left(0, \left(\sigma_0^{(i)}(\boldsymbol{x})\right)^2\right)$

- Pioneering work [7] conditions the model on $\mathcal{D}_t$
  - $\forall \boldsymbol{x} \in \mathcal{C}_{\mathcal{N}_i}, f^{(i)}(\boldsymbol{x}) | \mathcal{D}_t \sim \mathcal{N}\left(\mu_t(\boldsymbol{x}), \left(\sigma_t^{(i)}(\boldsymbol{x})\right)^2\right)$

- Acquisition function $\varphi_t^{(i)} : \mathcal{C}_{\mathcal{N}_i} \to \mathbb{R}$
  - Expected Improvement [8]
  - $\varphi_t^{(i)}(\boldsymbol{x}) = \mathbb{E}_{f^{(i)}(\boldsymbol{x}) \sim \mathcal{N}\left(\mu_t^{(i)}(\boldsymbol{x}), \left(\sigma_t^{(i)}(\boldsymbol{x})\right)^2\right)}\left[\left(f^{(i)}(\boldsymbol{x}) - y_t^*\right)^+\right]$
  - We set $\boldsymbol{x}^{(i)} = \arg\max_{\boldsymbol{x} \in \mathcal{C}_{\mathcal{N}_i}} \varphi_t^{(i)}(\boldsymbol{x})$

---

[7] Christopher K. I. Williams and Carl Edward Rasmussen. "Gaussian Processes for Regression". In: *Conference on Neural Information Processing Systems (NeurIPS'95)*. 1995.

[8] Jonas Mockus. "Application of Bayesian approach to numerical methods of global and stochastic optimization". In: *Journal of Global Optimization* 4 (1994), pp. 347–365.

Consensus Function

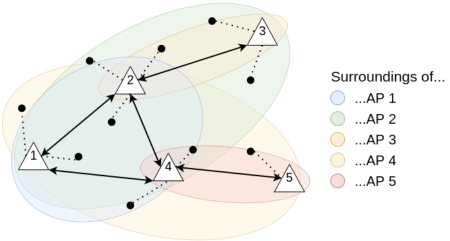

- **Assumption**. $\forall i \in [1, n]$, $f^{(i)}$ is L_i-Lipschitz continuous
- **Theorem**.
  - Let $\mathcal{P}_k = \left\{ x_k^{(i)} \right\}_{i \in \mathcal{N}_j}$ be the prescriptions received by AP $j = \left\lceil \frac{k}{2} \right\rceil$ for its parameter $k \in [1, 2n]$
  - Let $\tilde{x}_k$ be the median of $\mathcal{P}_k$, weighted by the Lipschitz constants $\{L_i\}_{i \in \mathcal{N}_j}$
  - Then, the vector $\tilde{\boldsymbol{x}} = (\tilde{x}_1, \cdots, \tilde{x}_{2n})$ is minimax optimal

Evaluation

- Evaluation through simulation with ns-3, 22 replications
- Two topologies with variable MCS and uplink/downlink traffic

T1: Open spaces, Cisco San Francisco
10 APs, 50 STAs

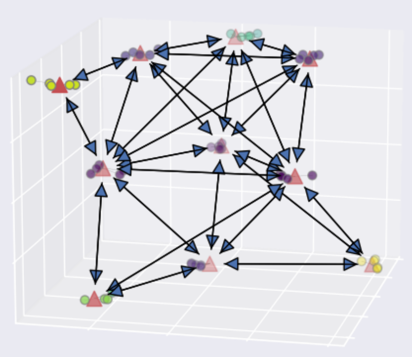

T2: Residential Building
14 APs, 56 STAs

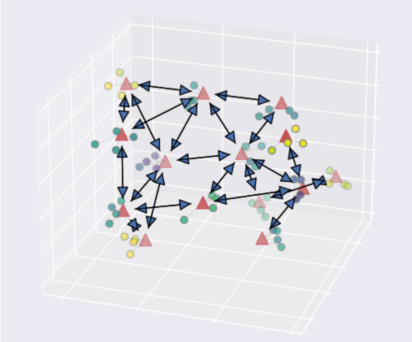

Evaluation

- Control strategy: DEFAULT
- SOTA strategies: WCNC'15 [9], JNCA'19 [10], GM+NGTS, HM+NGTS, INSPIRE, INSPIRE_LIM
- 75 ms per test, 30 seconds of simulated time $\implies$ 400 iterations

Solution	**Average Regret R_t/t**	
	T1	**T2**
DEFAULT	0.652 ± 0.001	0.429 ± 0.004
WCNC'15	0.470 ± 0.001	0.327 ± 0.005
JNCA'19	0.437 ± 0.001	0.398 ± 0.006
GM+NGTS	0.527 ± 0.016	0.375 ± 0.006
HM+NGTS	0.305 ± 0.006	0.379 ± 0.005
INSPIRE	**0.193 ± 0.005**	**0.294 ± 0.006**
INSPIRE_LIM	0.233 ± 0.005	0.329 ± 0.005

[9] Afaqui et al., see n. 4.

[10] Wilhelmi et al., see n. 5.

Network Oriented Metrics

T1

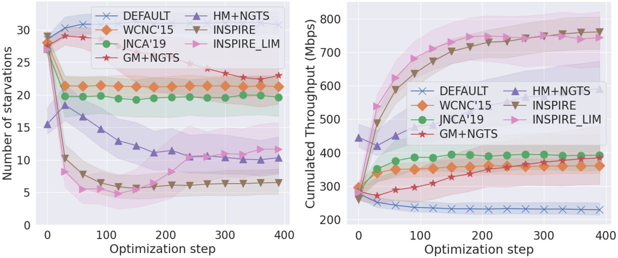

T2

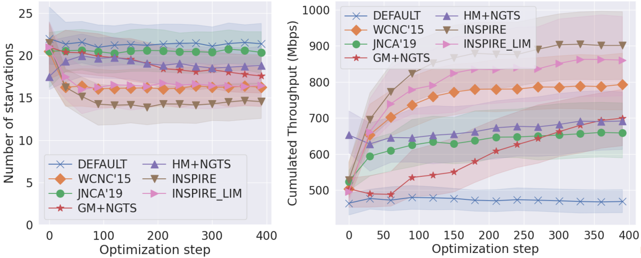

Discussion

- **Pros**
  - Excellent empirical performance
  - Robust evaluation (complex scenarios, credible number of APs)
  - Decentralized
  - Theoretical guarantees (minimax optimal at each iteration)
- **Cons**
  - Asymptotic optimality?
  - Impact of rounding?
  - High-dimensional factors in denser topologies?

Contribution 4

Decentralized, No-Regret Bayesian Optimization of High-Dimensional Functions

BO in High-Dim. Input Spaces

- Recall that BO as defined by [11] involves
  - **Assumption**. $f : \mathcal{C} \subset \mathbb{R}^d \to \mathbb{R}$ is $\mathcal{GP}(\mu, k)$
  - An acquisition function $\varphi_t : \mathcal{C} \to \mathbb{R}$ to discover promising queries
  - $\boldsymbol{x}_{t+1} = \arg\max_{\boldsymbol{x} \in \mathcal{C}} \varphi_t(\boldsymbol{x})$

- Classical BO struggles with high-dimensional input spaces because of the global optimization algorithms used to compute $\boldsymbol{x}_{t+1} = \arg\max_{\boldsymbol{x} \in \mathcal{C}} \varphi_t(\boldsymbol{x})$

- Solution:
  - **Assumption**. An additive decomposition for f, that is

$$f(\boldsymbol{x}) = \sum_{i=1}^{n} f^{(i)}(\boldsymbol{x}_{\mathcal{V}_i}) \tag{6}$$

  - with $f^{(i)} : \mathcal{C}^{(i)} \subseteq \mathbb{R}^{d^{(i)}} \to \mathbb{R}$ being $\mathcal{GP}\left(\mu^{(i)}, k^{(i)}\right)$, $\forall i \in [1, n]$

- **Maximum Factor Size (MFS)**: $\bar{d} = \max_{i \in [1,n]} d^{(i)}$

[11] Williams and Rasmussen, see n. 7.

Decomposing BO Algorithms

Solution	MFS Assumption	Find $\arg\max \varphi_t$
ADD-GPUCB [12]	$\bar{d} = 1$	Yes
QFF [13]	$\bar{d} = 1$	Yes
DEC-HBO [14]	$\bar{d} \leq 3$	Under assumptions
DuMBO (Ours)	None	Under assumptions

- We propose a Decentralized Message-passing Bayesian Optimization algorithm (DuMBO)
  - We demonstrate its asymptotic optimality
  - We demonstrate its competitiveness on synthetic and real-world problems

[12] Kirthevasan Kandasamy, Jeff Schneider, and Barnabás Póczos. "High dimensional Bayesian optimisation and bandits via additive models". In: *International conference on machine learning*. PMLR. 2015, pp. 295–304.

[13] Mojmir Mutny and Andreas Krause. "Efficient high dimensional bayesian optimization with additivity and quadrature fourier features". In: *Advances in Neural Information Processing Systems* 31 (2018).

[14] Trong Nghia Hoang et al. "Decentralized high-dimensional Bayesian optimization with factor graphs". In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 32. 1. 2018.

Decentralized GP-UCB

- The GP-UCB acquisition function is $\varphi_t(\boldsymbol{x}) = \mu_t(\boldsymbol{x}) + \beta_t^{\frac{1}{2}} \sigma_t(\boldsymbol{x})$ [15]
- $\sigma_t(\boldsymbol{x}) = \sqrt{\sum_{i=1}^n \left(\sigma_t^{(i)}(\boldsymbol{x}_{\mathcal{V}_i})\right)^2}$ cannot be computed in a decentralized fashion
- Previous works [16] propose to apply GP-UCB to each factor $f^{(i)}$ individually
- The optimized acquisition function is therefore

$$\begin{aligned}\varphi_t(\boldsymbol{x}) &= \sum_{i=1}^n \varphi_t^{(i)}(\boldsymbol{x}_{\mathcal{V}_i}) \\ &= \sum_{i=1}^n \mu_t^{(i)}(\boldsymbol{x}_{\mathcal{V}_i}) + \beta_t^{\frac{1}{2}} \sigma_t^{(i)}(\boldsymbol{x}_{\mathcal{V}_i}) \\ &= \mu_t(\boldsymbol{x}) + \beta_t^{\frac{1}{2}} \sum_{i=1}^n \sigma_t^{(i)}(\boldsymbol{x}_{\mathcal{V}_i}) \end{aligned} \tag{7}$$

- In (7), $\sigma_t(\boldsymbol{x})$ is replaced by the overestimation $\sum_{i=1}^n \sigma_t^{(i)}(\boldsymbol{x}_{\mathcal{V}_i})$
- The decentralized algorithms explore too much!

[15] Niranjan Srinivas et al. "Information-Theoretic Regret Bounds for Gaussian Process Optimization in the Bandit Setting". In: *IEEE Transactions on Information Theory* 58.5 (2012), pp. 3250–3265. DOI: doi:10.1109/tit.2011.2182033.

[16] Kandasamy, Schneider, and Póczos, see n. 12; Hoang et al., see n. 14.

Reducing the Gap

- The variance term $(\sigma_t^{(i)})^2$ can be decomposed [17] into
  - An **epistemic** term: uncertainty due to the lack of observed data
  - An **aleatoric** term $v_-^{(i)}$: observational noise, natural lower bound of $(\sigma_t^{(i)})^2$
- **Assumption**. $(\sigma_t^{(i)})^2$ is bounded from above by $v_+^{(i)}$
- Then, the optimal linear overestimation of $\sigma_t(\boldsymbol{x})$ on $[v_-, v_+]$ (with $v_- = \sum_{i=1}^n v_-^{(i)}$ and $v_+ = \sum_{i=1}^n v_+^{(i)}$) is

$$\frac{1}{4a} + a \sum_{i=1}^{n} (\sigma_t^{(i)}(\boldsymbol{x}_{\mathcal{V}_i}))^2 \tag{8}$$

- with a the single positive real root of the quartic

$$P(a) = \frac{2\left[u^3\right]_{v_-}^{v_+}}{3} a^4 - \frac{4\left[u^{\frac{5}{2}}\right]_{v_-}^{v_+}}{5} a^3 + \frac{\left[u^{\frac{3}{2}}\right]_{v_-}^{v_+}}{3} a - \frac{\left[u\right]_{v_-}^{v_+}}{8} \tag{9}$$

[17] Eyke Hüllermeier and Willem Waegeman. "Aleatoric and epistemic uncertainty in machine learning: An introduction to concepts and methods". In: *Machine Learning* 110.3 (2021), pp. 457–506.

Proposed Acquisition Function

- **Assumption.** $v_+ \leq \left(\sqrt{v_-} + 2\sum_{i=1}^{n}\sum_{\substack{j=1\\ j\neq i}}^{n}\sqrt{v_-^{(i)} v_-^{(j)}}\right)^2$
- **Theorem.** $\forall \boldsymbol{x} \in \mathcal{C}$, $\forall t \in \mathbb{N}$, we have

$$\sigma_t(\boldsymbol{x}) \leq \frac{1}{4a} + a\sum_{i=1}^{n}(\sigma_t^{(i)}(\boldsymbol{x}_{\mathcal{V}_i}))^2 \leq \sum_{i=1}^{n}\sigma_t^{(i)}(\boldsymbol{x}_{\mathcal{V}_i})$$

- Therefore, an algorithm maximizing

$$\begin{aligned}\varphi_t(\boldsymbol{x}) &= \sum_{i=1}^{n}\varphi_t^{(i)}(\boldsymbol{x}_{\mathcal{V}_i}) \\ &= \sum_{i=1}^{n}\mu_t^{(i)}(\boldsymbol{x}_{\mathcal{V}_i}) + \beta_t^{\frac{1}{2}} a(\sigma_t^{(i)}(\boldsymbol{x}_{\mathcal{V}_i}))^2\end{aligned} \tag{10}$$

- should have a lower regret than current state-of-the-art BO algorithms

Maximizing φ_t

- We use the Alternating Directions Method of Multipliers (ADMM) [18] to maximize φ_t in a decentralized fashion
  - Excellent performance of ADMM on nonconvex problems [19], [20]
  - [21] extends the global maximization guarantee of ADMM to **restricted prox-regular functions**

- **Definition** (restricted prox-regularity). For a lower semi-continuous function f, let $M \in \mathbb{R}^+, f : \mathbb{R}^n \to \mathbb{R} \cup \{+\infty\}$ and ∂f the set of general subgradients of f. Define the exclusion set $S_M = \{\boldsymbol{x} \in \text{dom}(f) : ||\boldsymbol{d}|| > M \text{ for all } \boldsymbol{d} \in \partial f(\boldsymbol{x})\}$. f is called restricted prox-regular if, for any $M > 0$ and bounded set $T \subseteq \text{dom}(f)$, there exists $\gamma > 0$ such that

$$f(\boldsymbol{y}) + \frac{\gamma}{2}||\boldsymbol{x} - \boldsymbol{y}||^2 \geq f(\boldsymbol{x}) + \boldsymbol{d}(\boldsymbol{y} - \boldsymbol{x}), \tag{11}$$

$\forall \boldsymbol{x} \in T \setminus S_M$, $\boldsymbol{y} \in T$, $\boldsymbol{d} \in \partial f(\boldsymbol{x})$, $||\boldsymbol{d}|| \leq M$.

[18] Daniel Gabay and Bertrand Mercier. "A dual algorithm for the solution of nonlinear variational problems via finite element approximation". In: *Computers & mathematics with applications* 2.1 (1976), pp. 17–40.

[19] Athanasios P Liavas and Nicholas D Sidiropoulos. "Parallel algorithms for constrained tensor factorization via alternating direction method of multipliers". In: *IEEE Transactions on Signal Processing* 63.20 (2015), pp. 5450–5463.

[20] Rongjie Lai and Stanley Osher. "A splitting method for orthogonality constrained problems". In: *Journal of Scientific Computing* 58.2 (2014), pp. 431–449.

[21] Yu Wang, Wotao Yin, and Jinshan Zeng. "Global convergence of ADMM in nonconvex nonsmooth optimization". In: *Journal of Scientific Computing* 78 (2019), pp. 29–63.

No-Regret Performance

- **Assumption**. φ_t is restricted prox-regular

- **Theorem**. Let $r_t = f(\boldsymbol{x}^*) - f(\boldsymbol{x}^t)$ denote the instantaneous regret of DuMBO. Let $\delta \in (0,1)$ and $\beta_t = 2\log\left(\frac{|\mathcal{D}|\pi^2 t^2}{6\delta}\right)$. Then $\forall \boldsymbol{x} \in \mathcal{D}, \forall t \in \mathbb{N}$ we have

$$r_t \leq 2\beta_t^{\frac{1}{2}} \left(a \sum_{i=1}^{n} \left(\sigma_t^{(i)}(\boldsymbol{x}^t) \right)^2 + \frac{1}{4a} \right) \tag{12}$$

  with probability at least $1 - \delta$.

- The regret bound (12) is lower than a no-regret algorithm, DEC-HBO [22]
- By piggybacking on the results of DEC-HBO, DuMBO is shown asymptotically optimal, that is

$$\lim_{t \to +\infty} \frac{R_t}{t} = 0 \tag{13}$$

[22] Hoang et al., see n. 14.

Numerical Experiments

- DuMBO: does not have access to the natural additive decomposition of f
  - Must infer it with [23]
- ADD-DuMBO: has access to the natural additive decomposition of f when it exists

- Comparison with two decomposing BO algorithms: ADD-GPUCB [24] and DEC-HBO [25]
  - Recall that they must infer a decomposition when the MFS $\bar{d} > 3$

- Comparison with two solutions that make other assumptions
  - SAASBO [26] and TuRBO [27]
  - Recall that they do not offer no-regret guarantees

[23] Jacob Gardner et al. "Discovering and exploiting additive structure for Bayesian optimization". In: *Artificial Intelligence and Statistics*. PMLR. 2017, pp. 1311–1319.

[24] Kandasamy, Schneider, and Póczos, see n. 12.

[25] Hoang et al., see n. 14.

[26] David Eriksson and Martin Jankowiak. "High-dimensional Bayesian optimization with sparse axis-aligned subspaces". In: *Uncertainty in Artificial Intelligence*. PMLR. 2021, pp. 493–503.

[27] David Eriksson et al. "Scalable global optimization via local bayesian optimization". In: *Advances in neural information processing systems* 32 (2019).

Experiments (Selection)

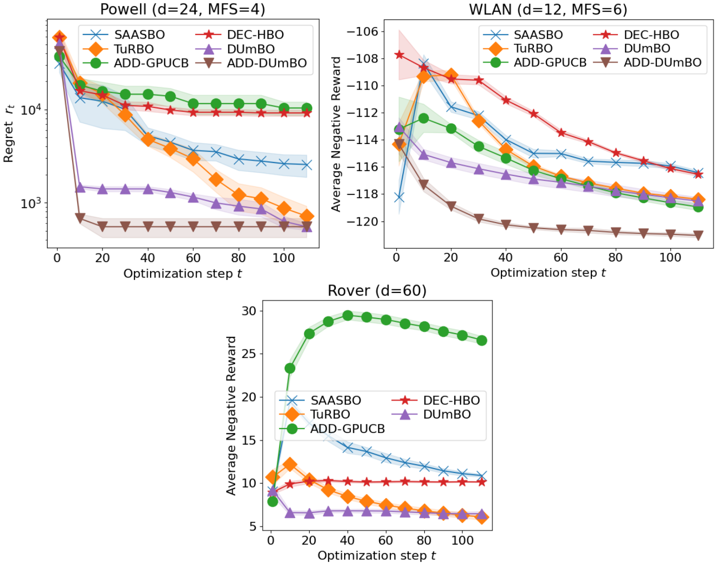

Discussion

- **Pros**
  - Excellent empirical performance
  - Robust evaluation on multiple benchmarks
  - Decentralized
  - Theoretical guarantees (asymptotic optimality)
- **Cons**
  - Restricted prox-regularity of the acquisition function?
  - Wall-clock time larger than SAASBO [28] or TuRBO [29]

[28] Eriksson and Jankowiak, see n. 26.

[29] Eriksson et al., see n. 27.

Conclusion

- We proposed online methods for the optimization of a black-box objective function within a wireless network, that are
  - both centralized and decentralized
  - competitive / better than state-of-the-art solutions
  - able to deal with high-dimensional problems
  - asymptotically optimal (DuMBO)
- Future works
  - More technical applications
  - Dynamic problems
  - Multi-objective problems
  - Student-t processes instead of GPs?

Thank you for your attention

Peer-Reviewed Journals & International Conferences

- **A. Bardou** and T. Begin. "Analysis of a Decentralized Bayesian Optimization Algorithm for Improving Spatial Reuse in Dense WLANs." In: *Computer Communications*. 2023.
- **A. Bardou**, T. Begin and A. Busson. "Mitigating Starvations in Dense WLANs: A Multi-Armed Bandit Solution." In: *Ad Hoc Networks*. 2023.
- **A. Bardou** and T. Begin. "INSPIRE: Distributed Bayesian Optimization for Improving Spatial Reuse in Dense WLANs." In: *MSWiM'22*. 2022. **Best Paper**.
- **A. Bardou**, T. Begin and A. Busson. "Analysis of a Multi-Armed Bandit Solution to Improve the Spatial Reuse in Next-Generation WLANs." In: *Computer Communications*. 2022.
- **A. Bardou**, T. Begin and A. Busson. "Improving the Spatial Reuse in IEEE 802.11ax WLANs: A Multi-Armed Bandit Approach." In: *MSWiM'21*. 2021.

Under Submission

- **A. Bardou**, P. Thiran and T. Begin. "Relaxing the Additivity Constraints in Decentralized No-Regret High-Dimensional Bayesian Optimization." @ *ICLR'24*.
- S. Si-Mohammed, **A. Bardou**, T. Begin, I. Guérin Lassous and P. Vicat-Blanc. "Smart Integration of Network Simulation in Network Digital Twin for Optimizing IoT Networks." @ *Future Generation Computer Systems*.
- **A. Bardou**, J-M. Gorce and T. Begin. "Assessing the Performance of NOMA in a Multi-Cell Context: A General Evaluation Framework." @ *INFOCOM'24*.

National Conferences

- **A. Bardou** and T. Begin. "INSPIRE: Optimisation bayésienne distribuée pour l'amélioration de la réutilisation spatiale des WLANs denses." In: *AlgoTel'22*. 2022. **Best Paper**.
- **A. Bardou**, T. Begin and A. Busson. "Multi-Armed Bandit Algorithm for Spatial Reuse in WLANs: Minimizing Stations in Starvation." In: *ROADEF'22*. 2022.

A Fundamental Problem: The Exploration-Exploitation Dilemma

- **Exploration**: querying policy that maximizes the probability to be surprised
- **Exploitation**: querying policy that maximizes the probability to obtain high f-values according to actual beliefs

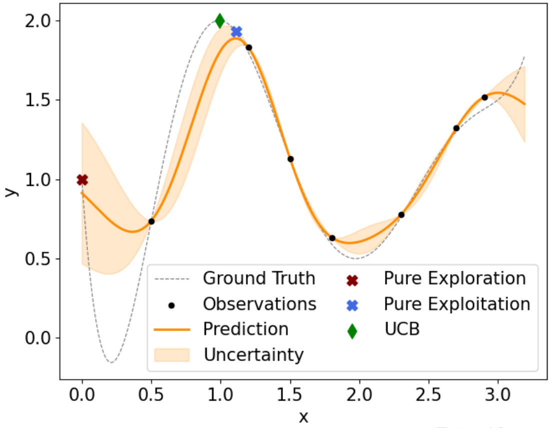

Online Approaches for Spatial Reuse Optimization in Wi-Fi

Proposed solutions	Tuning of OBSS_PD	Tuning of TX_PWR	Dynamic MCS	Traffic Up/Down	Simulator	APs / channels
VTC'04 [30]		✓		Up	Self-made	8/1
Infocom'20 [31]		✓		Up/Down	Self-made	100/11
WCNC'15 [32]	✓			Up	Self-made	100/3
WCNC'21 [33]	✓	✓	✓	Down	ns-3	6/1
Globecom'20 [34]	✓			Up/Down	ns-3	3/1
ADHOC'19 [35]	✓	✓		Down	Self-made	8/1
JNCA'19 [36]	✓	✓		Down	Self-made	8/1

- Solutions evaluated on vanilla scenarios with virtually no theoretical guarantees

[30] Youngsoo Kim, Jeonggyun Yu, and Sunghyun Choi. "SP-TPC: a self-protective energy efficient communication strategy for IEEE 802.11 WLANs". In: *IEEE 60th Vehicular Technology Conference, 2004. VTC2004-Fall. 2004.* Vol. 3. IEEE. 2004, pp. 2078–2082.

[31] Shuwei Qiu et al. "Joint access point placement and power-channel-resource-unit assignment for 802.11 ax-based dense WiFi with QoS requirements". In: *IEEE INFOCOM 2020-IEEE Conference on Computer Communications.* IEEE. 2020, pp. 2569–2578.

[32] Afaqui et al., see n. 4.

[33] Hyunjoong Lee, Hyung-Sin Kim, and Saewoong Bahk. "LSR: link-aware spatial reuse in IEEE 802.11 ax WLANs". In: *2021 IEEE Wireless Communications and Networking Conference (WCNC).* IEEE. 2021, pp. 1–6.

[34] Elif Ak and Berk Canberk. "FSC: Two-scale AI-driven fair sensitivity control for 802.11 ax networks". In: *GLOBECOM 2020-2020 IEEE Global Communications Conference.* IEEE. 2020, pp. 1–6.

[35] Wilhelmi et al., see n. 5.

[36] Francesc Wilhelmi et al. "Potential and pitfalls of multi-armed bandits for decentralized spatial reuse in WLANs". In: *Journal of Network and Computer Applications* 127 (2019), pp. 26–42.

Finding the Best Arm in the Reservoir

- Our optimizer builds upon [37], which assumes $f(\boldsymbol{x}_i) + \epsilon \sim \mathcal{N}(\mu_i, 1)$
  - A Gaussian conjugate prior is placed on $\theta = \mu_i$ and updated with data
- We assume $f(\boldsymbol{x}_i) + \epsilon \sim \mathcal{N}(\mu_i, \sigma_i^2)$
  - We place a Normal-Gamma conjugate prior on $\boldsymbol{\theta} = \left(\mu_i, \sigma_i^{-2}\right)$ with parameters $\left(\mu_i^0, \lambda_i^0, \alpha_i^0, \beta_i^0\right)$ and update formulas

$$\mu_i^n = \frac{\lambda_i^0 \mu_i^0 + n\bar{y}}{\lambda_i^0 + n}, \tag{14}$$

$$\lambda_i^n = \lambda_i^0 + n, \tag{15}$$

$$\alpha_i^n = \alpha_i^0 + \frac{n}{2}, \tag{16}$$

$$\beta_i^n = \beta_i^0 + \frac{1}{2}\left(ns + \frac{\lambda_i^0 n(\bar{y} - \mu_i^0)^2}{\lambda_i^0 + n}\right). \tag{17}$$

- To identify the best arm, we rely on Thompson sampling, which samples an arm k with probability

$$p_k = \int_\Theta \mathbb{1}_{\mathbb{E}[y|k,\boldsymbol{\theta}] = \max_{k' \in \mathcal{A}} \mathbb{E}[y|k',\boldsymbol{\theta}]} p(\boldsymbol{\theta}|\mathcal{D}) \mathrm{d}\boldsymbol{\theta} \tag{18}$$

[37] Wilhelmi et al., see n. 5.

Evaluation - T1

- Control strategies: DEFAULT, ϵ-GREEDY ($\epsilon \propto \frac{1}{t}$)
- SOTA strategies: UNIF+GTS [38], GM+GTS, GM+NGTS
- 50 ms per test, 120 seconds simulated $\implies$ 2,400 iterations

Strategy	Average Regret R_t/t
DEFAULT	0.376 $\pm$0.008
ϵ-GREEDY	0.458 $\pm$0.011
UNIF+GTS	0.773 $\pm$0.002
GM+GTS	0.354 $\pm$0.007
GM+NGTS	**0.313 $\pm$0.007**

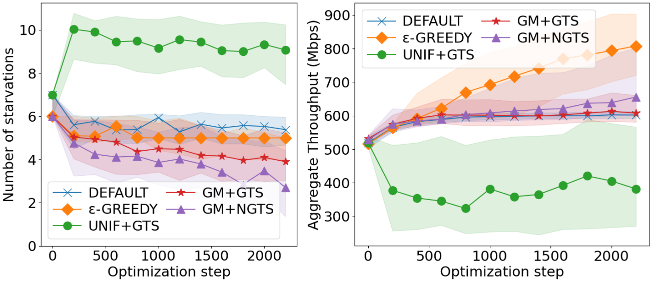

[38] Wilhelmi et al., see n. 5.

INSPIRE - Surrogate Modelling: Gaussian Process

- A Gaussian Process (GP) is a collection of random variables $\{Y(\boldsymbol{x})\}_{\boldsymbol{x}\in\mathcal{C}}$ indexed by a set $\mathcal{C}$
  - Any finite set $\{Y(\boldsymbol{x}_1), \cdots, Y(\boldsymbol{x}_n)\}$ has a joint multivariate Gaussian distribution
  - A GP is fully specified by

$$\mu(\boldsymbol{x}) = \mathbb{E}[Y(\boldsymbol{x})] \tag{19}$$

$$k(\boldsymbol{x}, \boldsymbol{x}') = \mathbb{E}\left[(Y(\boldsymbol{x}) - \mu(\boldsymbol{x}))\left(Y(\boldsymbol{x}') - \mu(\boldsymbol{x}')\right)\right] \tag{20}$$

- **Assumption**. $\forall i \in [1, n]$, $f^{(i)}$ is a Gaussian Process $\mathcal{GP}\left(0, k^{(i)}\left(\boldsymbol{x}_{\mathcal{N}_i}, \boldsymbol{x}'_{\mathcal{N}_i}\right)\right)$
  - $\boldsymbol{x}_{\mathcal{N}_i} \in \mathcal{C}_{\mathcal{N}_i} = \prod_{j\in\mathcal{N}_i} \mathcal{C}^{(j)}$
  - $k^{(i)}$ is a Matérn covariance function [39] with its hyperparameter $\nu = 3/2$

$$k^{(i)}(\boldsymbol{x}, \boldsymbol{x}') = s_i^2 \left(1 + \frac{\sqrt{3}||\boldsymbol{x} - \boldsymbol{x}'||_2}{\rho_i}\right) e^{-\frac{\sqrt{3}||\boldsymbol{x}-\boldsymbol{x}'||_2}{\rho_i}} \tag{21}$$

  - with hyperparameters $\left(s_i^2, \rho_i\right) \in \mathbb{R}_+^* \times \mathbb{R}_+^*$

[39] Marc G Genton. "Classes of kernels for machine learning: a statistics perspective". In: *Journal of machine learning research* 2.Dec (2001), pp. 299–312.

Bayesian Optimization: Inference Formulas

- Pioneering work [40]
- **Assumption**. $\forall i \in [1, n]$, $f^{(i)}$ is a Gaussian Process $\mathcal{GP}\left(0, k^{(i)}\left(\boldsymbol{x}_{\mathcal{N}_i}, \boldsymbol{x}'_{\mathcal{N}_i}\right)\right)$
  - $\boldsymbol{x}_{\mathcal{N}_i} \in \mathcal{C}_{\mathcal{N}_i} = \prod_{j \in \mathcal{N}_i} \mathcal{C}^{(j)}$
  - $\forall \boldsymbol{x} \in \mathcal{C}_{\mathcal{N}_i}, f^{(i)}(\boldsymbol{x}) \sim \mathcal{N}\left(0, \left(\sigma_0^{(i)}(\boldsymbol{x})\right)^2\right)$

- Given $\mathcal{D}_t = (\boldsymbol{X}_t, \boldsymbol{y}_t)$, with the $t \times d$ matrix $\boldsymbol{X}_t = (\boldsymbol{x}_j)^\top_{j \in [1,t]}$ and the t-dimensional vector $\boldsymbol{y}_t = (y_j)^\top_{j \in [1,t]}$, $f^{(i)}(\boldsymbol{x}) | \mathcal{D}_t \sim \mathcal{N}\left(\mu_t^{(i)}(\boldsymbol{x}), \left(\sigma_t^{(i)}(\boldsymbol{x})\right)^2\right)$ with

$$\mu_t^{(i)}(\boldsymbol{x}) = \boldsymbol{k}(\boldsymbol{x}, \boldsymbol{X}_t)^\top \boldsymbol{K}_t^{-1} \boldsymbol{y}_t \tag{22}$$

$$\left(\sigma_t^{(i)}(\boldsymbol{x})\right)^2 = k(\boldsymbol{x}, \boldsymbol{x}) - \boldsymbol{k}(\boldsymbol{x}, \boldsymbol{X}_t) \boldsymbol{K}_t^{-1} \boldsymbol{k}(\boldsymbol{x}, \boldsymbol{X}_t) \tag{23}$$

- with $\boldsymbol{k}(\boldsymbol{x}, \boldsymbol{X}_t) = (k(\boldsymbol{x}, \boldsymbol{x}_j))_{\boldsymbol{x}_j \in X_t}$ and $\boldsymbol{K}_t = (k(\boldsymbol{x}_j, \boldsymbol{x}_k))_{\boldsymbol{x}_j, \boldsymbol{x}_k \in \boldsymbol{X}_t}$

[40] Williams and Rasmussen, see n. 7.

INSPIRE - Acquisition Function

- $\varphi_t^{(i)} : \mathcal{C}_{\mathcal{N}_i} \to \mathbb{R}$
- Many candidates: GP-UCB [41], KG [42], PI [43]

- Expected Improvement [44]

$$\varphi_t^{(i)}(\boldsymbol{x}) = \mathbb{E}_{f^{(i)}(\boldsymbol{x}) \sim \mathcal{N}\left(\mu_t^{(i)}(\boldsymbol{x}), \left(\sigma_t^{(i)}(\boldsymbol{x})\right)^2\right)} \left[\left(f^{(i)}(\boldsymbol{x}) - y_t^*\right)^+\right]$$
$$= \left(\mu_t^{(i)}(\boldsymbol{x}) - y_t^*\right) \Phi(z(\boldsymbol{x})) + \sigma_t^{(i)}(\boldsymbol{x}) \phi(z(\boldsymbol{x})) \quad (24)$$

- with $y_t^* = \max_{j \in [1,t]} y_j$, $(x)^+ = \max(0, x)$, $z(\boldsymbol{x}) = \left(\mu_t^{(i)}(\boldsymbol{x}) - y_t^*\right) / \sigma_t^{(i)}(\boldsymbol{x})$, Φ and ϕ the cdf and pdf of $\mathcal{N}(0,1)$ respectively

- We set $\boldsymbol{x}^{(i)} = \arg\max_{\boldsymbol{x} \in \mathcal{C}_{\mathcal{N}_i}} \varphi_t^{(i)}(\boldsymbol{x})$

[41] Srinivas et al., see n. 15.

[42] Shanti S Gupta and Klaus J Miescke. "Bayesian look ahead one-stage sampling allocations for selection of the best population". In: *Journal of statistical planning and inference* 54.2 (1996), pp. 229–244.

[43] Donald R. Jones, Matthias Schonlau, and William J. Welch. "Efficient global optimization of expensive black-box functions". In: *Journal of Global optimization* 13.4 (1998), 455–492.

[44] Mockus, see n. 8.

INSPIRE - Minimax Optimality

- Objective to minimize: $g(\boldsymbol{x}) = |\sum_{i=1}^{n} f^{(i)}(\boldsymbol{x}^{(i)}) - f(\boldsymbol{x})|$
  - $\boldsymbol{x}^{(i)} \in \mathcal{C}_{\mathcal{N}_i}$ is the prescription of AP i

- Classical formulation of a minimax problem: $\inf_{\boldsymbol{x}} \sup_{\boldsymbol{y}} g(\boldsymbol{x}, \boldsymbol{y})$

- For INSPIRE, we derive $B(\boldsymbol{x}) \geq g(\boldsymbol{x})$
  - Non-uniform upper bound of g
  - Lowest upper bound given the assumed information about g
- We minimize $B(\boldsymbol{x})$ to find a promising consensus

- Given the strong similarity with a minimax optimization task, we call $\tilde{\boldsymbol{x}} = \arg\min_{\boldsymbol{x} \in \mathcal{C}} B(\boldsymbol{x})$ a minimax optimum

INSPIRE - Computational Overhead

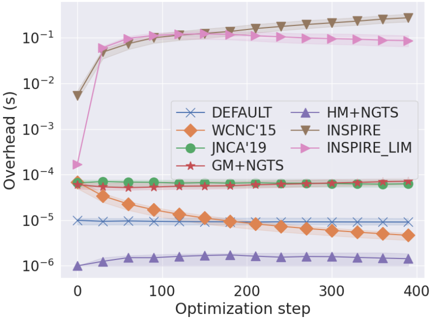

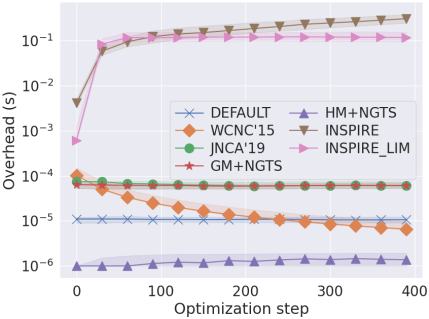

INSPIRE - Different Complexities

- $\boldsymbol{x}^+ \in \mathcal{C}$ recommended by INSPIRE
- Two random vectors $(\boldsymbol{x}_1, \boldsymbol{x}_2) \in \mathcal{C}^2$
  - Plot $\left(a, b, f\left(a\boldsymbol{x}_1 + b\boldsymbol{x}_2 + \boldsymbol{x}^+\right)\right)$

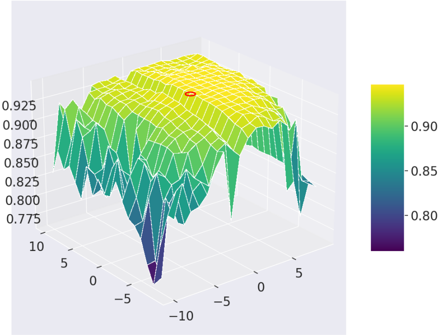

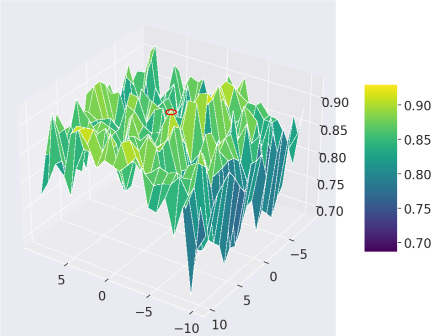

INSPIRE - Alternatives

- `Single GP`: Centralized version of INSPIRE
- `GPs w/o agg.`: Decentralized version without consensus

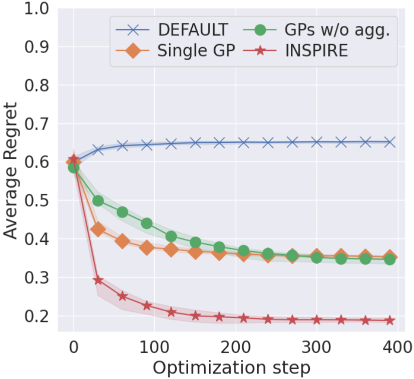

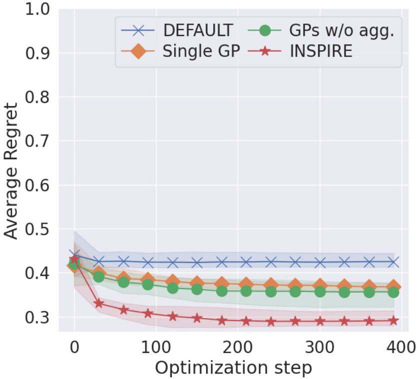

DuMBO - Factor Graph

- Recall that we assume $f = \sum_{i=1}^{n} f^{(i)}$
- This can be represented by a **factor graph**
  - e.g, $f(\boldsymbol{x}) = f_1(x_1, x_3) + f_2(x_2) + f_3(x_2, x_3) + f_4(x_1, x_3)$
  - For this decomposition, the MFS $\bar{d} = 2$
- Each **factor node** i ($i \in [1, n]$) can communicate with the variable nodes in $\mathcal{V}_i$
  - e.g. $\mathcal{V}_1 = \{1, 3\}$
- Each **variable node** j ($j \in [1, d]$) can communicate with the factor nodes in $\mathcal{F}_j$
  - e.g. $\mathcal{F}_3 = \{1, 3, 4\}$
- Therefore, $f(\boldsymbol{x}) = \sum_{i=1}^{n} f^{(i)}(\boldsymbol{x}_{\mathcal{V}_i})$
- We want a decentralized algorithm that can be run on this factor graph

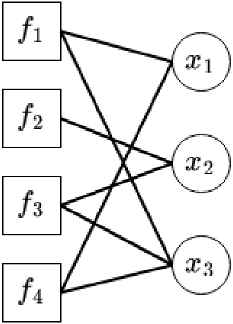

DuMBO - Inferring the Decomposition

- The additive decomposition can be inferred by MCMC from data [45]

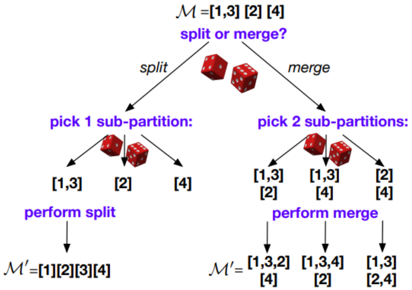

- Acceptance probability: $\mathcal{P}\left(\mathcal{M}'|\mathcal{M}\right) = \min\left(1, \frac{p(\boldsymbol{y}|\boldsymbol{X},\mathcal{M}')g(\mathcal{M}'|\mathcal{M})}{p(\boldsymbol{y}|\boldsymbol{X},\mathcal{M})g(\mathcal{M}|\mathcal{M}')}\right)$
- Given $\mathcal{M}_1, \cdots, \mathcal{M}_k$ additive decompositions, we optimize

$$\varphi_t(\boldsymbol{x}) = \frac{1}{k}\sum_{i=1}^{k} \varphi_t^{\mathcal{M}_i}(\boldsymbol{x}) \tag{25}$$

[45] Gardner et al., see n. 23.

DuMBO - Maximizing φ_t

$$\begin{aligned} &\text{maximize} \sum_{i=1}^{n} \varphi_t^{(i)}(\boldsymbol{x}^{(i)}) \\ &\text{such that } \boldsymbol{x}^{(i)}_{\mathcal{V}_i \cap \mathcal{V}_j} = \boldsymbol{x}^{(j)}_{\mathcal{V}_i \cap \mathcal{V}_j}, \forall i, j \in [1, n] \end{aligned} \tag{26}$$

- with $\boldsymbol{x}^{(i)} \in \mathcal{C}^{(i)}$
- Introducing a global consensus variable $\bar{\boldsymbol{x}} \in \mathcal{C}$,

$$\begin{aligned} &\text{maximize} \sum_{i=1}^{n} \varphi_t^{(i)}(\boldsymbol{x}^{(i)}) \\ &\text{such that } \boldsymbol{x}^{(i)}_{\mathcal{V}_i} = \bar{\boldsymbol{x}}_{\mathcal{V}_i}, \forall i \in [1, n] \end{aligned} \tag{27}$$

- Augmented Lagrangian relaxation $\mathcal{L}_\eta(\boldsymbol{x}^{(1)}, \cdots, \boldsymbol{x}^{(n)}, \bar{\boldsymbol{x}}, \boldsymbol{\lambda}) = \sum_{i=1}^{n} \mathcal{L}_\eta^{(i)}$ with

$$\mathcal{L}_\eta^{(i)}(\boldsymbol{x}^{(i)}, \boldsymbol{\lambda}^{(i)}, \bar{\boldsymbol{x}}_{\mathcal{V}_i}) = \varphi_t^{(i)}(\boldsymbol{x}^{(i)}) - \boldsymbol{\lambda}^{(i)\top}(\boldsymbol{x}^{(i)} - \bar{\boldsymbol{x}}_{\mathcal{V}_i}) - \frac{\eta}{2}||\boldsymbol{x}^{(i)} - \bar{\boldsymbol{x}}_{\mathcal{V}_i}||_2^2 \tag{28}$$

- with $\boldsymbol{\lambda}^{(i)}$ the Lagrange multipliers for $\mathcal{L}_\eta^{(i)}$

DuMBO - ADMM

- We use ADMM [46] to maximize $\mathcal{L}_\eta$
- Iterative method that successively finds, at iteration k,

$$\boldsymbol{x}_{k+1}^{(1)} = \arg\max_{\boldsymbol{x}^{(1)}} \mathcal{L}_\eta^{(1)}(\boldsymbol{x}^{(1)}, \bar{\boldsymbol{x}}_k, \boldsymbol{\lambda}_k)$$

$$\vdots$$

$$\boldsymbol{x}_{k+1}^{(n)} = \arg\max_{\boldsymbol{x}^{(n)}} \mathcal{L}_\eta^{(n)}(\boldsymbol{x}^{(n)}, \bar{\boldsymbol{x}}_k, \boldsymbol{\lambda}_k)$$

$$\bar{\boldsymbol{x}}_{k+1} = \arg\max_{\bar{\boldsymbol{x}}} \mathcal{L}_\eta(\boldsymbol{x}_{k+1}^{(1)}, \cdots, \boldsymbol{x}_{k+1}^{(n)}, \bar{\boldsymbol{x}}, \boldsymbol{\lambda}_k) \tag{29}$$

$$\boldsymbol{\lambda}_{k+1} = \arg\max_{\boldsymbol{\lambda}} \mathcal{L}_\eta(\boldsymbol{x}_{k+1}^{(1)}, \cdots, \boldsymbol{x}_{k+1}^{(n)}, \bar{\boldsymbol{x}}_{k+1}, \boldsymbol{\lambda}) \tag{30}$$

- Note that $\boldsymbol{x}_{k+1}^{(1)}, \cdots, \boldsymbol{x}_{k+1}^{(n)}$ can be found concurrently
- Note that (29) and (30) have closed forms

[46] Gabay and Mercier, see n. 18.

DuMBO - Asymptotic Optimality

- Piggybacking on the results of [47]
- **Theorem 1**: the instantaneous regret of DEC-HBO (discrete domain) is

$$r_t \leq 2\beta_t^{\frac{1}{2}} \sum_{i=1}^{n} \sigma_t^{(i)}(\boldsymbol{x}^t) \tag{31}$$

- Note that this regret bound is larger than ours
- **Theorem 2**: DEC-HBO is asymptotically optimal in a discrete domain
- **Theorem 3**: DEC-HBO is asymptotically optimal in a continuous domain
  - **Assumption**. f is Lipschitz-continuous

[47] Hoang et al., see n. 14.

DuMBO - Other Experiments

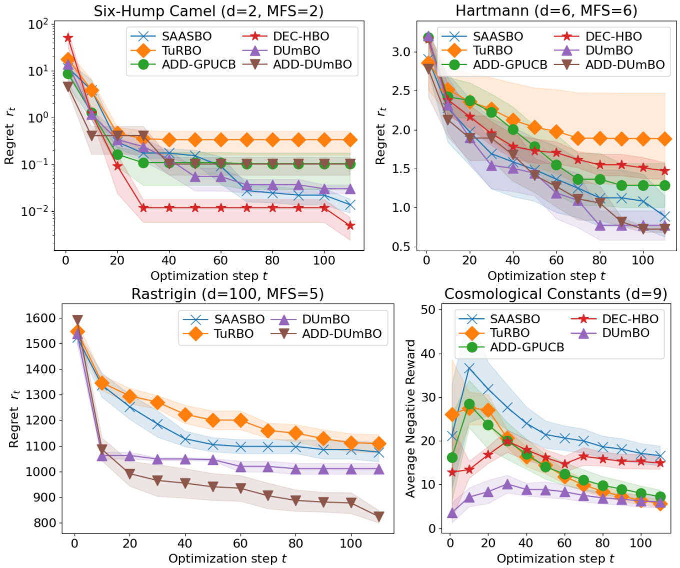

DuMBO - Wall-Clock Time