Multi-Fidelity Bayesian Optimization with Unreliable Information Sources

Petrus Mikkola[1], Julien Martinelli[1], Louis Filstroff[2], Samuel Kaski[1,3]
[1]Aalto University [2]ENSAI, CREST [3]University of Manchester
Contact: petrus.mikkola@aalto.fi

1 Problem

- Bayesian Optimization (BO) is a powerful framework for optimizing black-box, expensive-to-evaluate functions.
- Multi-Fidelity Bayesian Optimization (MFBO) integrates cheaper, lower-fidelity auxiliary information sources (ISs) to accelerate optimization over Single-Fidelity BO (SFBO).
- State-of-the-art MFBO algorithms can fail when auxiliary ISs are poor approximations of the primary IS → Leads to higher regret than SFBO, defeating their purpose!

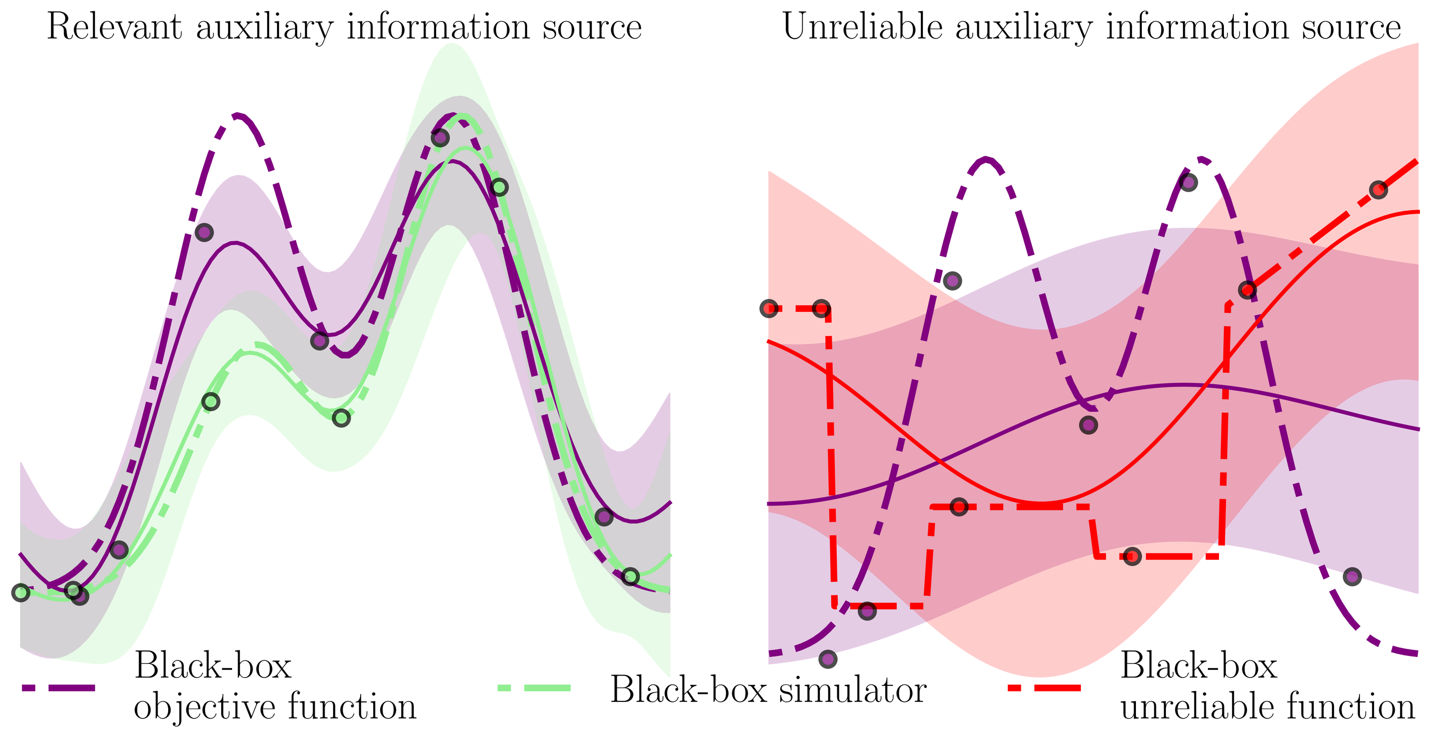

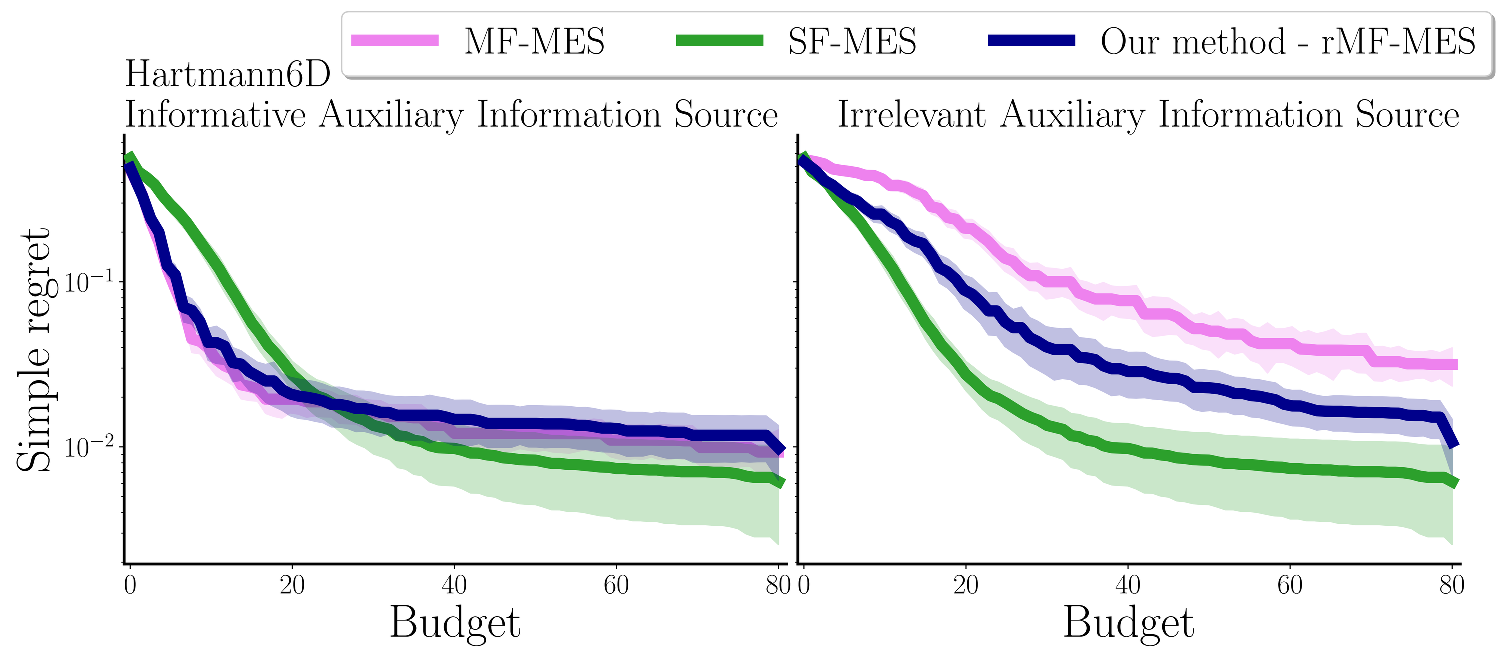

2 Contributions

- We introduce rMFBO, a methodology to make any GP-based MFBO scheme robust to the addition of unreliable ISs.
- rMFBO provides theoretical guarantees that its performance can be tied to its SFBO analog with controllable probability.
- rMFBO outperforms concurrent MFBO methods when unreliable ISs are involved, while speeding up convergence w.r.t. SFBO when including relevant ISs.

3 Method

Alongside the MFBO algorithm, we introduce a concurrent pseudo-SFBO algorithm, which keeps track of data from the primary IS only, and so-called *pseudo-observations*. At each round t, we consider both single- and multi-fidelity proposals for an acquisition function α

$$\mathbf{x}_t^{\mathrm{pSF}} = \arg\max_{\mathbf{x}} \alpha(\mathbf{x}, m \mid \mu_{\mathrm{pSF}}, \sigma_{\mathrm{pSF}}),$$
$$(\mathbf{x}_t^{\mathrm{MF}}, \ell_t) = \arg\max_{\mathbf{x},\ell} \alpha(\mathbf{x}, \ell \mid \mu_{\mathrm{MF}}, \sigma_{\mathrm{MF}}).$$

We follow the conservative query from pSFBO, $(\mathbf{x}_t^{\mathrm{pSF}}, m)$, unless both conditions below are satisfied, in which case $(\mathbf{x}_t^{\mathrm{MF}}, \ell_t)$ is queried. When pSFBO is not followed, we add a **pseudo-observation**, $\mu_{\mathrm{MF}}(\mathbf{x}_t^{\mathrm{pSF}}, m)$, to estimate what would have been the value of the SFBO query.

- Condition 1: The accuracy of the pseudo-observation should be high enough: $\sigma_{\mathrm{MF}}(\mathbf{x}_t^{\mathrm{pSF}}, m) \le c_1$.
- Condition 2: The MFBO query proposal should be relevant enough: $s(\mathbf{x}_t^{\mathrm{MF}}, \ell_t) \ge c_2$, where s is a **relevance measure**. In this work, we consider a cost-adjusted information gain [1].

rMFBO acts as an adaptive on/off switch between MFBO and SFBO

4 Theoretical results

Given that the objective function is drawn from a GP with a known smooth kernel, and let $c_1(\varepsilon, q) = \varepsilon/\sqrt{-2\log(1-q)}$:

Theorem ("No harm"). *Assume both algorithms, the robust MFBO and its SFBO variant, return their final proposal. Then,*

$$R(\Lambda + \lambda_m, \mathbf{x}^{rMF}_{choice}) \le R(\Lambda, \mathbf{x}^{SF}_{choice}) + \varepsilon \max\{T\hat{M}_T d^{T+1}, 2\},$$

with probability greater than $q\left(1 - da\exp(-\frac{1}{b^2})\right)$.

$\hat{M}_t$ measures the sensitivity of the next query when moving from pSFBO dataset to SFBO dataset.

- If we tolerate e.g. 0.1 units of regret undershoot with 90% probability, then we can consider $c_1(0.1, 0.9) \approx 0.05$.
- The values $c_1 = c_2 = 0.1$ performed well in the experiments.

5 Experiments

- rMFBO is implemented in the `BoTorch` framework [2].
- We evaluate against several MOGP joint models.

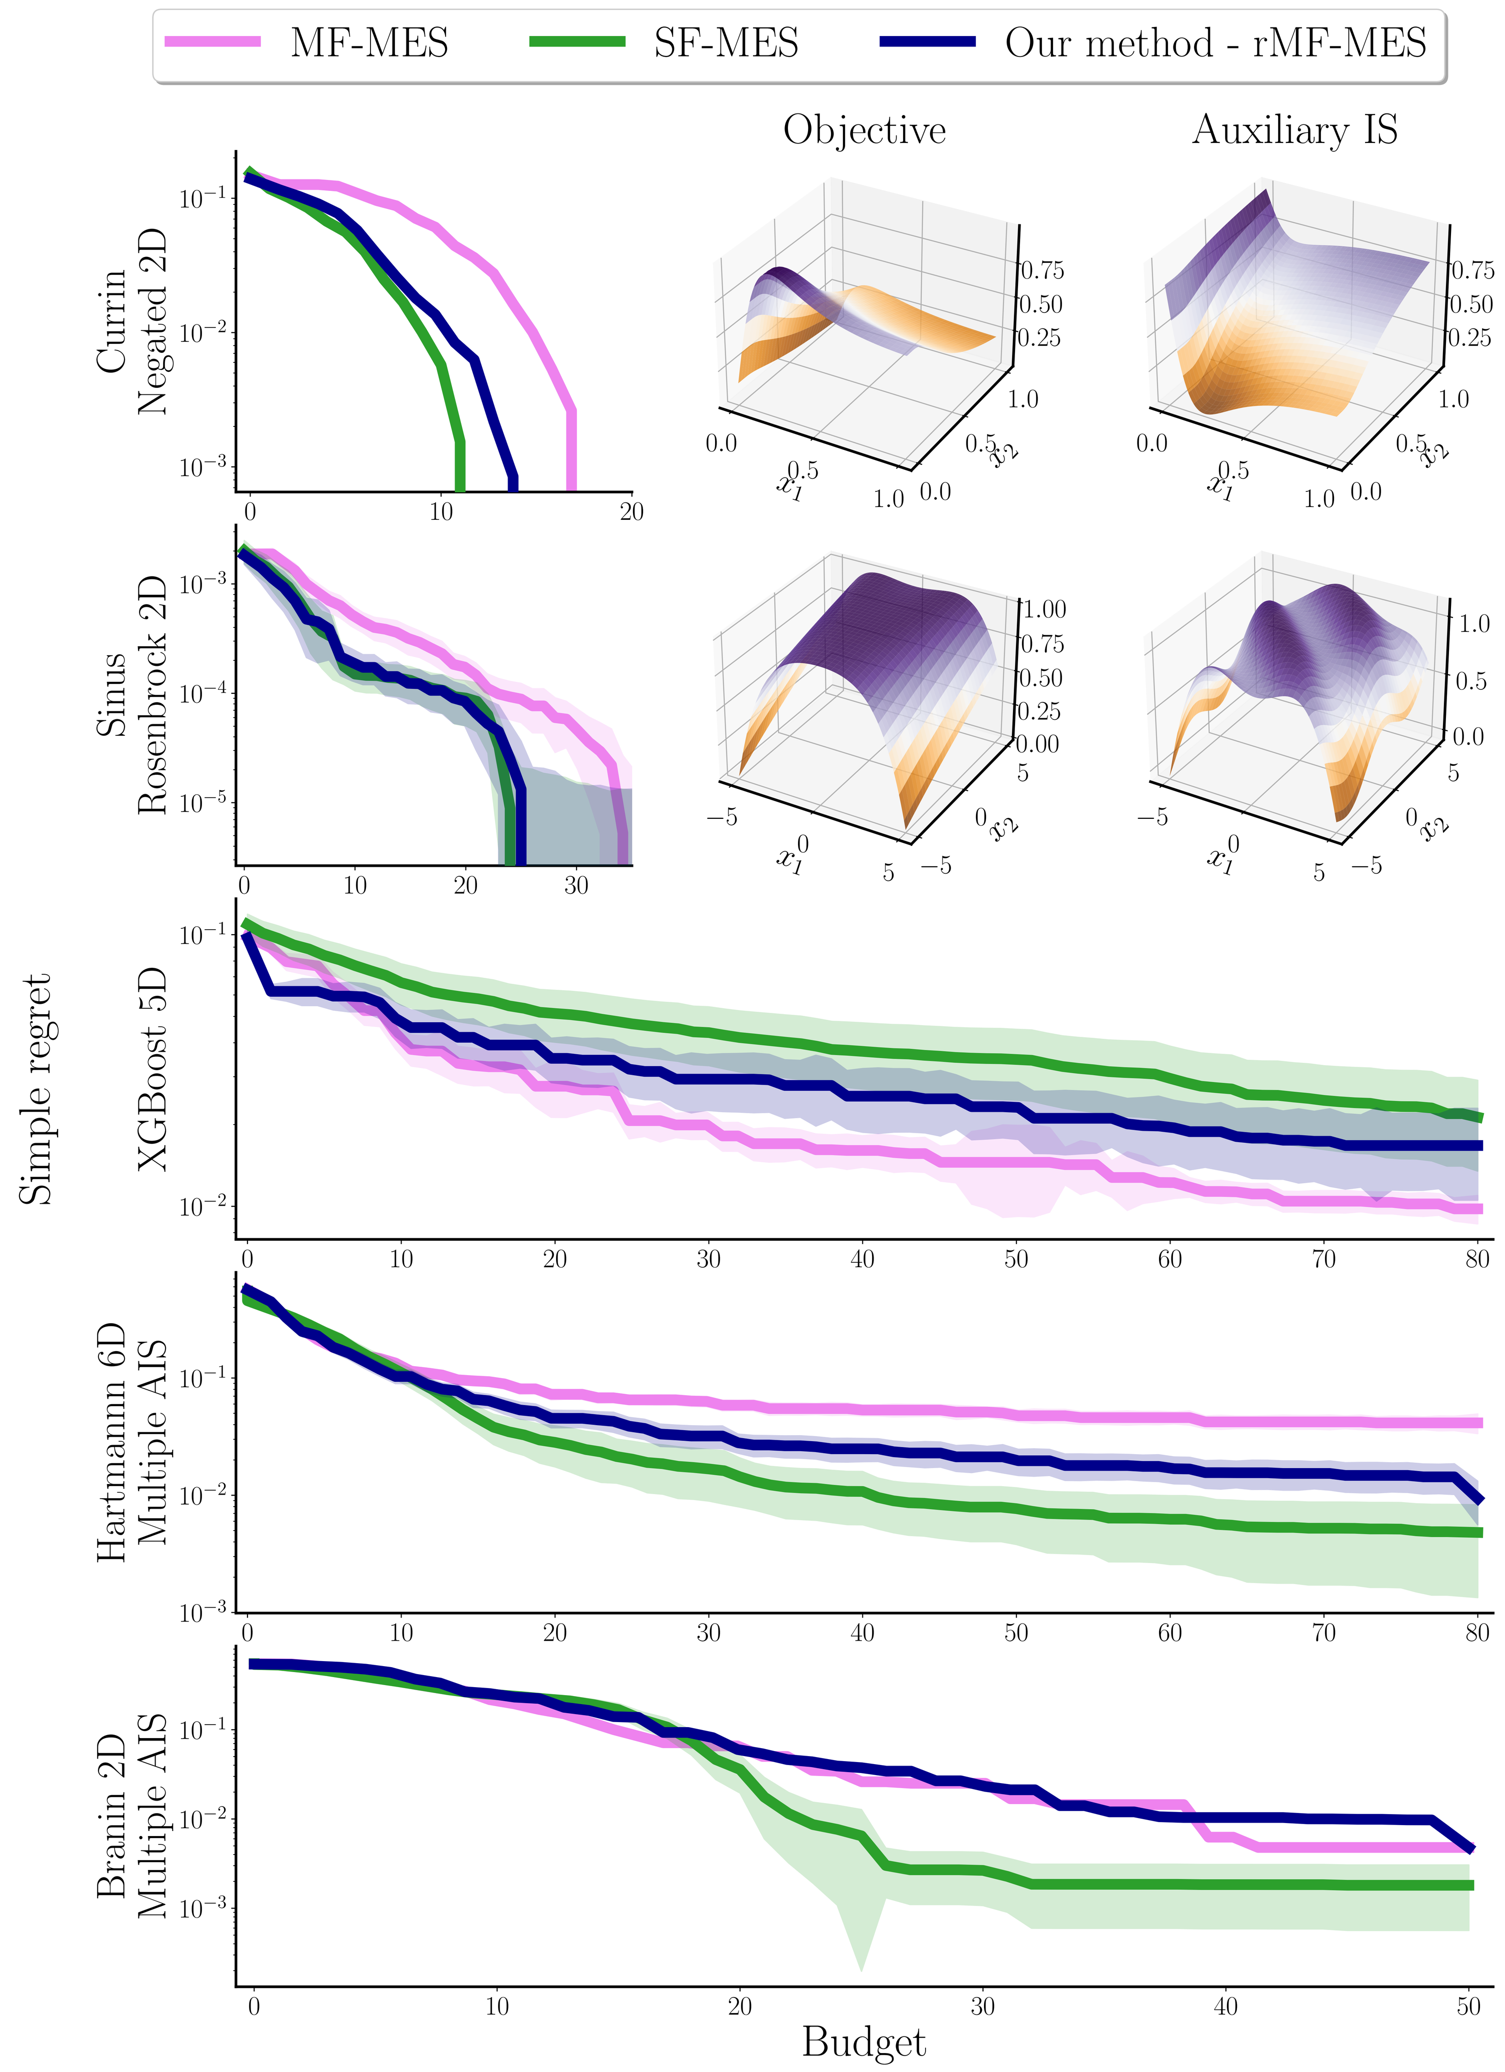

References

[1] Shion Takeno, Hitoshi Fukuoka, Yuhki Tsukada, Toshiyuki Koyama, Motoki Shiga, Ichiro Takeuchi, and Masayuki Karasuyama. *Multi-fidelity Bayesian optimization with max-value entropy search and its parallelization*. ICML, 2020.
[2] Maximilian Balandat, Brian Karrer, Daniel Jiang, Samuel Daulton, Ben Letham, Andrew G Wilson, and Eytan Bakshy. *BoTorch: A Framework for Efficient Monte-Carlo Bayesian Optimization*. NeurIPS, 2020.