Safe Bayesian optimization across noise models: A practical and unified framework

Abdullah Tokmak [1] Thomas B. Schön [2] Dominik Baumann [1,2]

[1]Aalto University, Finland [2]Uppsala University, Sweden

Aalto University School of Electrical Engineering

UPPSALA UNIVERSITET

Contribution

Safe Bayesian optimization (BO) algorithms may be restrictive in practice since they require **homoscedastic sub-Gaussian observation noise**. We propose a **straightforward yet rigorous framework** for safe BO that works **across noise models**, including homoscedastic sub-Gaussian and heteroscedastic heavy-tailed distributions. Hence, we cover a broader noise spectrum [1].

State of the art assumption. The observation noise ϵ_t is homoscedastic R-sub-Gaussian distribution.

Our assumption. The observation noise ϵ_t is defined on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$, from which we can sample.

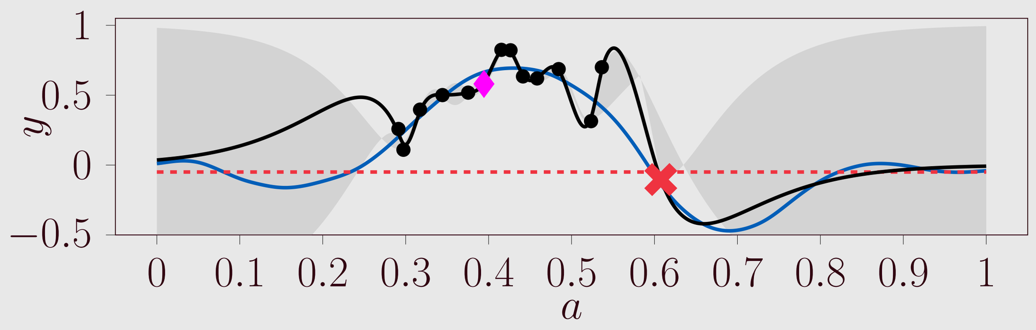

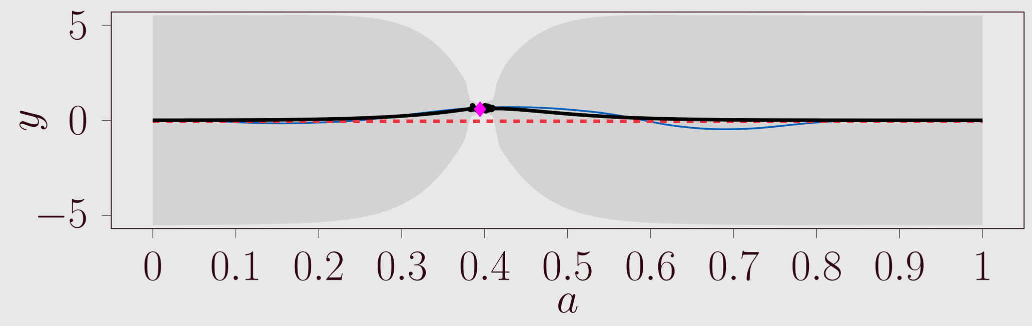

Starting from the initial safe set, we sequentially gather samples to maximize the unknown ground truth while guaranteeing safety under **heteroscedastic heavy-tailed** observation noise. The state of the art (left) can only handle homoscedastic sub-Gaussian noise, which causes **safety violations**. Our safe BO algorithm (right) works across noise models and remains **safe** by exploring conservatively. In practice, the unknown ground truth may be a reward function mapping control parameters to their performance, while safety violations may correspond to experiments that yield hardware damage or harm the environment.

Introduction

- Safe BO algorithms (SafeOpt [2, 3]) only consider **homoscedastic sub-Gaussian noise**
- Modeling network delays or observation noise of radar/LiDAR sensors ⇒ **heavy-tailed**
- Reinforcement learning/parameter tuning: noise depends on input ⇒ **heteroscedastic**

Problem definition

- We use safe BO with Gaussian processes (GPs) to maximize reward f of safety-critical systems
- Kernel k: GP mean μ_t, standard deviation σ_t, covariance matrix K_t, RKHS norm $\|f\|_k$
- Confidence intervals bound difference between reward f and η-regularized GP mean μ_t

$$|f(a) - \mu_t(a)| \leq \left(\|f\|_k + \sqrt{1/\eta}\|\epsilon_{1:t}\|_{\Xi_t}\right) \cdot \sigma_t(a), \quad \Xi_t := K_t(K_t + \eta I_t)^{-1} \qquad (1)$$

- State of the art [4]: **Assume sub-Gaussian** ϵ_t to bound $\|\epsilon_{1:t}\|_{\Xi_t}$ with high probability

How can we bound $\|\epsilon_{1:t}\|_{\Xi_t}$ across noise models without relying on sub-Gaussianity?

Data-driven aleatoric uncertainty quantification via scenario approach [5]

- **Assumption:** Observation noise ϵ_t is defined on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$
- To bound ϵ_t w.h.p., generate m_t i.i.d. **scenarios** $\tilde{\epsilon}_t^{(j)}$ from $(\Omega, \mathcal{F}, \mathbb{P}) \Rightarrow \bar{\epsilon}_t := \max_{j\in[1,m_t]} \tilde{\epsilon}_t^{(j)}$
- If m_t s.t. $(1-\nu)^{m_t} \leq \kappa_t$ with $\nu, \kappa_t \in (0,1)$, then $\mathbb{P}^{m_t}[V(\bar{\epsilon}_t) > \nu] \leq \kappa_t$, $V(\bar{\epsilon}_t) := \mathbb{P}[\bar{\epsilon}_t < |\epsilon_t|]$

Noise bounds that hold simultaneously for all iterations $t \geq 1$

- Confidence level $\kappa \in (0,1) \Rightarrow \kappa_t := \frac{6\kappa}{\pi^2 t^2}$, product probability measure $\tilde{\mathbb{P}} := \otimes_{t=1}^{\infty} \otimes_{j=1}^{m_t} \mathbb{P}$
- $\tilde{\mathbb{P}}[\forall t \geq 1 : V(\bar{\epsilon}_t) > \nu] \leq \sum_{t=1}^{\infty} \tilde{\mathbb{P}}[V(\bar{\epsilon}_t) > \nu] \leq \sum_{t=1}^{\infty} \kappa_t = \sum_{t=1}^{\infty} \frac{6\kappa}{\pi^2 t^2} = \frac{6\kappa}{\pi^2} \cdot \frac{\pi^2}{6} = \kappa$ □

Theorem 1: Scenario-based confidence intervals

Suppose:

- *Observation noise ϵ_t is defined on $(\Omega, \mathcal{F}, \mathbb{P})$, from which we can sample*
- *Reward function f is a member of the RKHS of kernel k with known $\|f\|_k$*

Then, with confidence at least $1-\kappa$ (under $\tilde{\mathbb{P}}$) and probability at least $1-\nu$ (under $\mathbb{P}$):

$$|f(a) - \mu_t(a)| \leq \left(\|f\|_k + \sqrt{\lambda_{\max}(\Xi_t)/\eta}\|\bar{\epsilon}_{1:t}\|_2\right) \cdot \sigma_t(a) \qquad (2)$$

Proof (idea):

- *Take confidence intervals (1) and probabilistically bound $\|\epsilon_{1:t}\|_{\Xi_t}$ using scenario approach*
- *$\|\epsilon_{1:t}\|_{\Xi_t} \leq \sqrt{\lambda_{\max}(\Xi_t)}\|\epsilon_{1:t}\|_2$ (deterministically) $\Rightarrow \|\epsilon_{1:t}\|_{\Xi_t} \leq \sqrt{\lambda_{\max}(\Xi_t)}\|\bar{\epsilon}_{1:t}\|_2$ (w.h.p.)*

Theorem 2: Safe and optimal BO algorithm

Suppose:

- *Hypotheses of Theorem 1 hold*
- *Nonempty initial set of safe policy parameters is given*

Then, our safe BO algorithm with confidence intervals (2) ***safely*** *finds the* ***reachable optimum****.*

Synthetic experiments

Classic safe BO

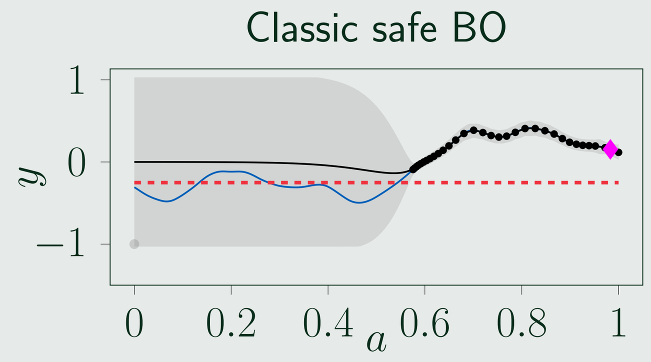

Our algorithm

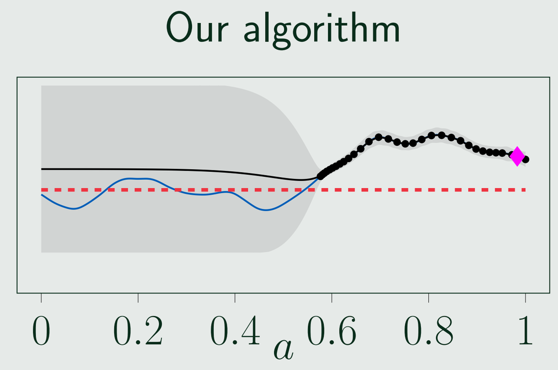

- Our algorithm (right) exhibits **similar performance** to classic safe BO algorithms (left) under sub-Gaussian measurement noise, working with **more general noise assumptions**
- Our algorithm remain **safe** while classic safe BO algorithms fail under heteroscedastic heavy-tailed noise (see Contribution)

Control parameter tuning on the Franka Emika robot

Experiment setup

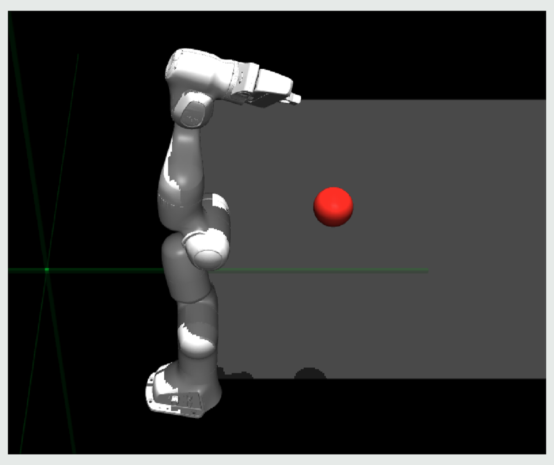

Starting position

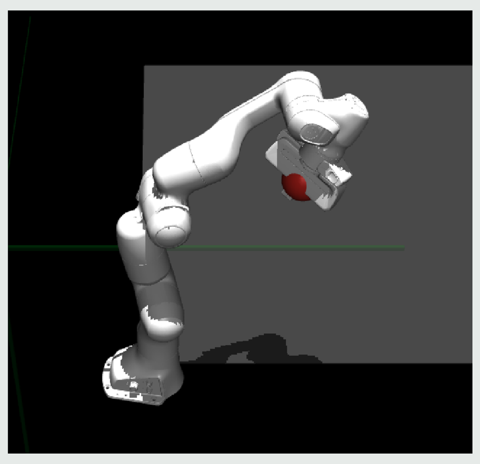

End position

Safe exploration of the domain to find the reachable optimum

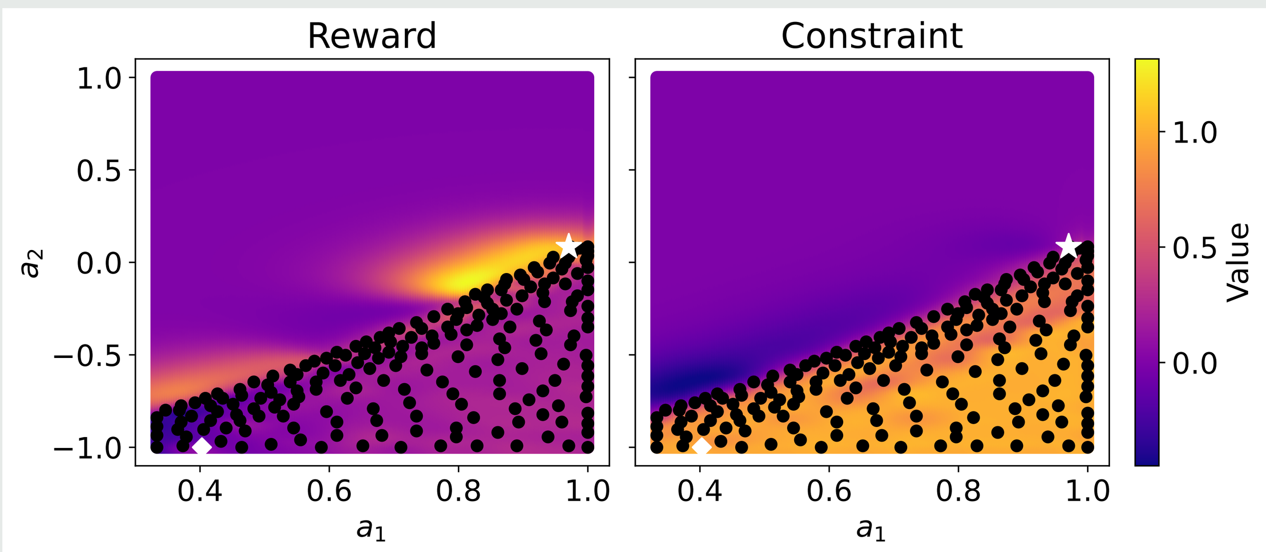

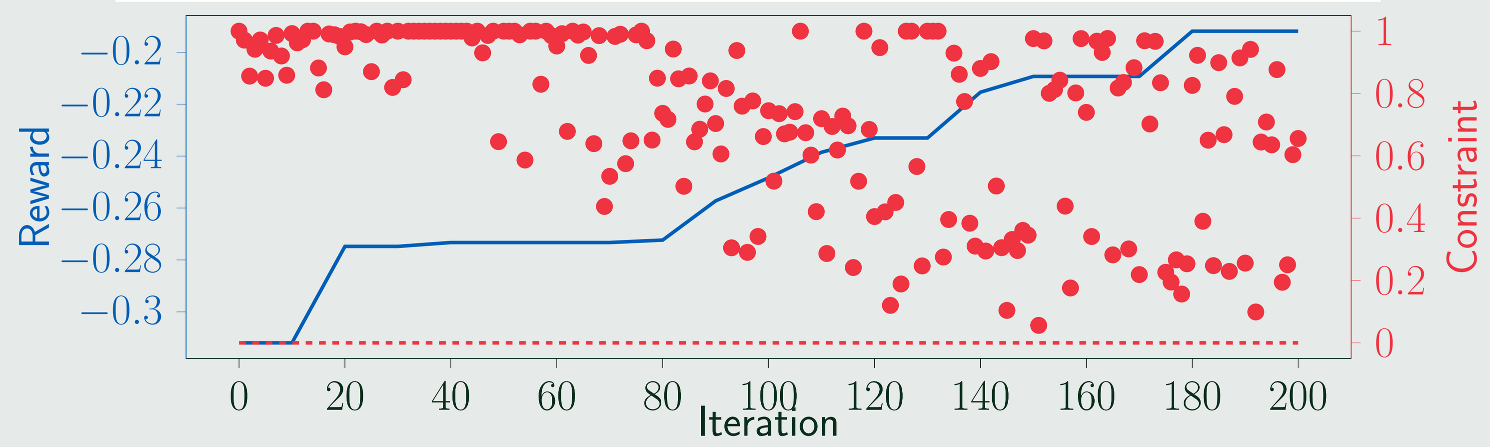

- LQR parameter tuning to improve set-point tracking performance of the Franka Emika robot
- Starting from **low-performing policy** (white diamond), identify **optimal policy** (white star)
- Continuous reward improvement while only conducting safe experiments

Conclusion

- We assume that the **observation noise lives on a probability space** $(\Omega, \mathcal{F}, \mathbb{P})$ from which we can sample.
- We bound observation noise from homoscedastic sub-Gaussian and heteroscedastic heavy-tailed distributions via **scenario approach**.
- We develop high probability **confidence intervals** (Theorem 1) and prove that our safe BO algorithm **remains safe** and finds the safely **reachable optimum** (Theorem 2).

References

[1] A. Tokmak, T. B. Schön, and D. Baumann. "Safe Bayesian optimization across noise models: A practical and unified framework". under review for IEEE L-CSS and ACC. 2025.

[2] Y. Sui, A. Gotovos, J. Burdick, and A. Krause. "Safe exploration for optimization with Gaussian processes". In: *International Conference on Machine Learning*. 2015.

[3] F. Berkenkamp, A. Krause, and A. P. Schoellig. "Bayesian optimization with safety constraints: Safe and automatic parameter tuning in robotics". In: *Machine Learning* (2023).

[4] S. R. Chowdhury and A. Gopalan. "On kernelized multi-armed bandits". In: *International Conference on Machine Learning*. 2017.

[5] M. C. Campi and S. Garatti. *Introduction to the scenario approach*. SIAM, 2018.