A Safe Bayesian Optimization Algorithm for Tuning the Optical Synchronization System at European XFEL

J. O. Lübsen[1] M. Schütte[12] S. Schulz[2] A. Eichler[12]
jannis.luebsen@tuhh.de

[1]Hamburg University of Technology - Institute of Control Systems

[2]Deutsches Elektronen-Synchrotron - Accelerator Beam Controls

2nd collaboration Workshop on RL4AA
February 5th, 2024

TUHH Hamburg University of Technology

DESY.

ICS Institute of Control Systems

MSK

Motivation

Problem Statement

- Model based
  - Optimization with dynamic models involves system identification
  - Optimization accuracy depends on model fidelity
- Online tuning procedures optimize the system directly
- Heuristic tuning is time consuming

Motivation

Problem Statement

- Model based
    - Optimization with dynamic models involves system identification
    - Optimization accuracy depends on model fidelity
- **Online tuning** procedures optimize the system directly
- Heuristic tuning is time consuming

Example: Synchronization System of the EuXFEL

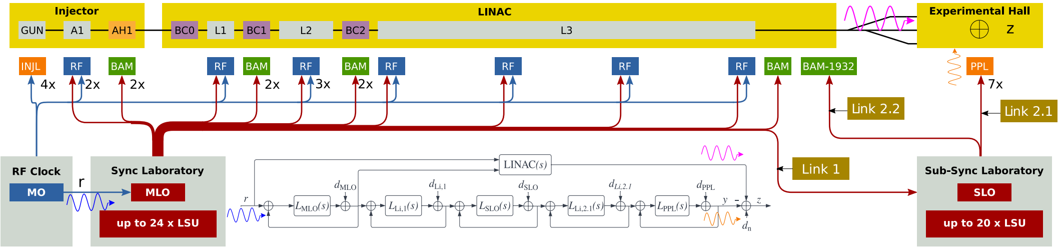

Minimize timing gap between X-Rays and Pump-Probe laser pulses by tuning controllers

- Largest linear particle accelerator in the world
- Measurable timing gap
- Timing gap must be below a constant value T

- Expensive machine time
- Noisy measurements
- Safe optimization

Motivation

Goal

- $\min f(\boldsymbol{x}) \quad \text{s. t.} \quad g(\boldsymbol{x}) \geq 0$

Proposal

- Modified Safe Bayesian optimization
  - Black Box approach
  - Safe during optimization
  - Learns a probabilistic surrogate model
- Increased convergence rate compared to other optimization approaches

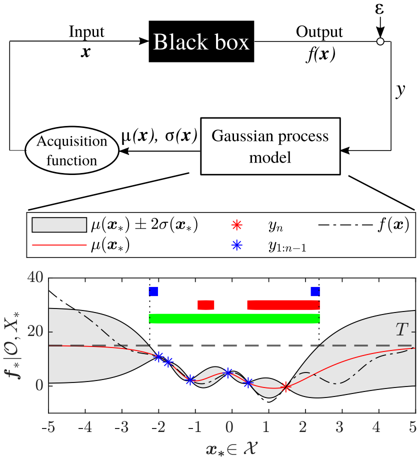

Contents

Black Box

- $y_i = f(\boldsymbol{x}_i) + \varepsilon$, where $\varepsilon \sim \mathcal{N}(0, \sigma_n^2)$
- Set of observations $\mathcal{O} = \{\boldsymbol{x}_i, y_i | i = 1 \ldots n\}$
- Training points $\boldsymbol{x} \in \mathcal{X}$ and test points $\boldsymbol{x}_* \in \mathcal{X}$
- $X = [\boldsymbol{x}_1^T, \ldots, \boldsymbol{x}_n^T]^T$ and $X_* = [\boldsymbol{x}_{*,1}^T, \ldots, \boldsymbol{x}_{*,s}^T]^T$

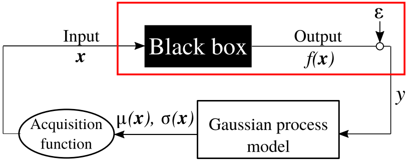

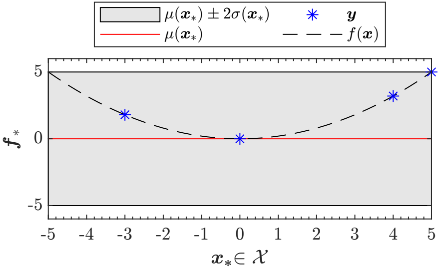

Black Box

- $y_i = f(\boldsymbol{x}_i) + \varepsilon$, where $\varepsilon \sim \mathcal{N}(0, \sigma_n^2)$
- Set of observations $\mathcal{O} = \{\boldsymbol{x}_i, y_i | i = 1 \dots n\}$
- Training points $\boldsymbol{x} \in \mathcal{X}$ and test points $\boldsymbol{x}_* \in \mathcal{X}$
- $X = [\boldsymbol{x}_1^T, \dots, \boldsymbol{x}_n^T]^T$ and $X_* = [\boldsymbol{x}_{*,1}^T, \dots, \boldsymbol{x}_{*,s}^T]^T$

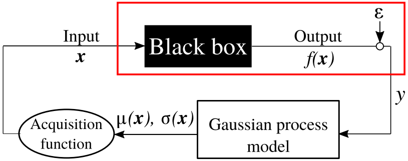

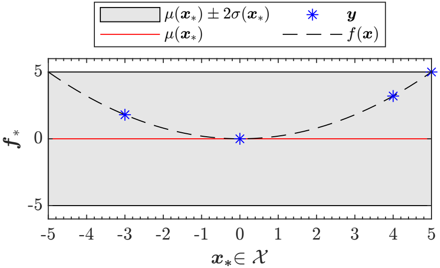

⇒

Kernel

- Prior assumption: $f \sim \mathcal{GP}(\mathbf{0}, k(\boldsymbol{x}, \boldsymbol{x}'))$
- $k(\cdot, \cdot)$ encodes properties of f

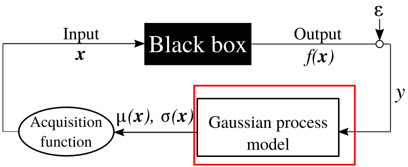

Kernel

- Prior assumption: $f \sim \mathcal{GP}(\mathbf{0}, k(\boldsymbol{x}, \boldsymbol{x}'))$
- $k(\cdot, \cdot)$ encodes properties of f

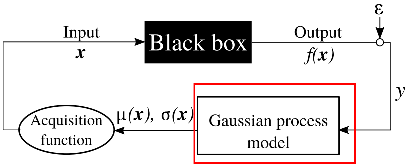

Assumption on the target function

The unknown function f is a member of the Reproducing Kernel Hilbert Space $\mathcal{H}$ defined by the positive definite function $k(\cdot, \cdot)$.

Kernel

- Prior assumption: $f \sim \mathcal{GP}(\mathbf{0}, k(\boldsymbol{x}, \boldsymbol{x}'))$
- $k(\cdot,\cdot)$ encodes properties of f
- $k_{\mathrm{SE}}(\boldsymbol{x}, \boldsymbol{x}') = \sigma_f^2 \exp\left(-0.5\frac{(\boldsymbol{x}-\boldsymbol{x}')^2}{l^2}\right)$
- Adjustable hyperparameters l^2 and σ_f^2

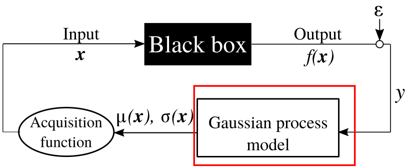

Assumption on the target function

The unknown function f is a member of the Reproducing Kernel Hilbert Space $\mathcal{H}$ defined by the positive definite function $k(\cdot,\cdot)$.

Inference

- Perform inference step as described in[1]

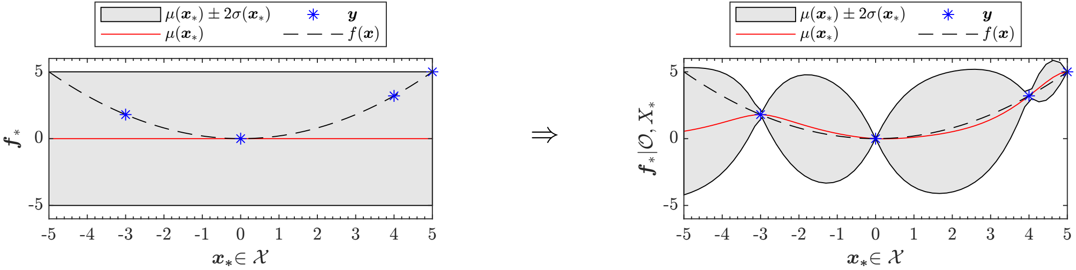

[1]Williams and Rasmussen, *Gaussian processes for machine learning*, 2006.

Bayesian Optimization

- Acquisition function α searches for promising inputs using the predictive distribution
- $\boldsymbol{x}_{\text{new}} = \arg\max_{\boldsymbol{x}_* \in \mathcal{X}} \alpha(\boldsymbol{x}_*)$
- $\min f(\boldsymbol{x}) \quad \text{s.t.} \quad g(\boldsymbol{x}) \geq 0$

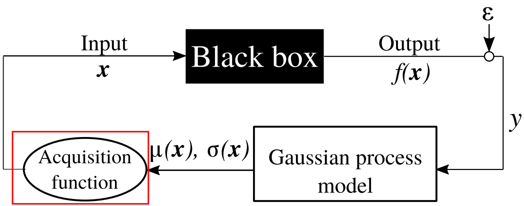

Constraint

- Consider $g(\boldsymbol{x}) = T - f(\boldsymbol{x})$
- T denotes a safety threshold
- Avoid evaluation of unsafe inputs
- One Gaussian process is sufficient as dependency of g and f is known
- Alternatively, two Gaussian processes for g and f respectively

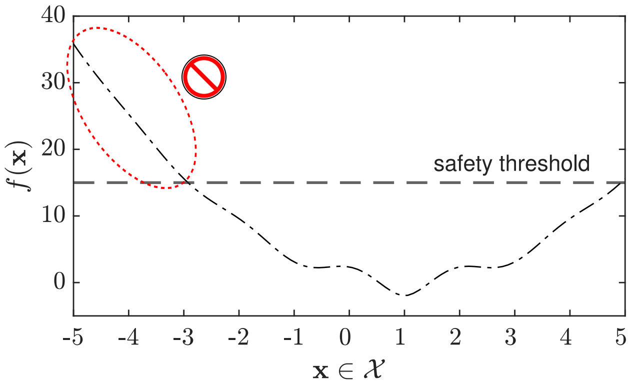

Modified Safe Options (MoSaOpt)

- Safe options[2] evolved to modified safe options
- Safe set: $\mathcal{S} = \{x \in \mathcal{X} | \mathrm{UCB}(x) \leq T\}$
- Minimizer set: $\mathcal{M} = \{x \in \mathcal{S} | \mathrm{LCB}(x) \leq y^*\}$
- Expander set: $\mathcal{G} = \{x \in \mathcal{S} | \delta\mathcal{S}\}$

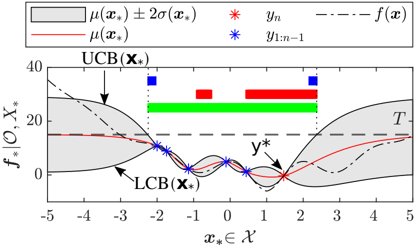

[2] Sui *et al.*, "Safe Exploration for Optimization with Gaussian Processes," 2015

Modified Safe Options (MoSaOpt)

- Safe options[2] evolved to modified safe options
- Safe set: $\mathcal{S} = \{x \in \mathcal{X} | \mathrm{UCB}(x) \leq T\}$
- Minimizer set: $\mathcal{M} = \{x \in \mathcal{S} | \mathrm{LCB}(x) \leq y^*\}$
- Expander set: $\mathcal{G} = \{x \in \mathcal{S} | \delta\mathcal{S}\}$

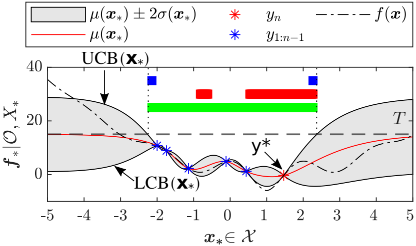

MoSaOpt divided into **exploration** and **exploitation** phase

[2] Sui *et al.*, "Safe Exploration for Optimization with Gaussian Processes," 2015

Exploration

- Observe the reachable set
  $\mathcal{R} = \{x \in \mathcal{X} | f(x) \leq T\}$
- $\boldsymbol{x}_{\text{new}} = \arg\max_{\boldsymbol{x} \in \mathcal{G}} \sigma(\boldsymbol{x})$
- Repeat until entire $\mathcal{R}$ is observed indicated by small uncertainties of expanders

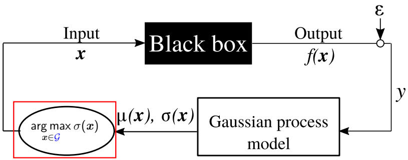

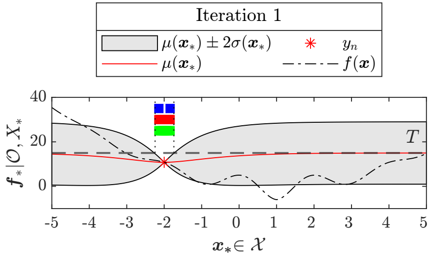

Exploration

- Observe the reachable set $\mathcal{R} = \{x \in \mathcal{X} | f(x) \leq T\}$
- $\boldsymbol{x}_{\text{new}} = \arg\max_{\boldsymbol{x} \in \mathcal{G}} \sigma(\boldsymbol{x})$
- Repeat until entire $\mathcal{R}$ is observed indicated by small uncertainties of expanders

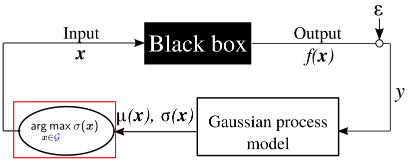

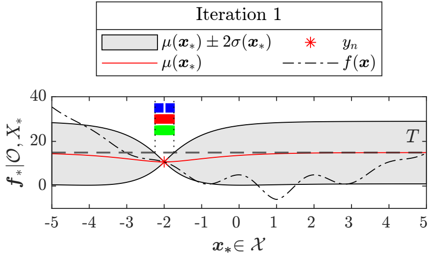

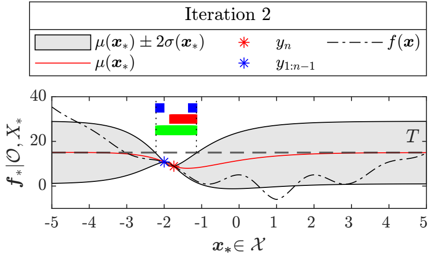

Exploration

- Observe the reachable set $\mathcal{R} = \{x \in \mathcal{X} | f(x) \leq T\}$
- $\boldsymbol{x}_{\text{new}} = \arg\max_{\boldsymbol{x} \in \mathcal{G}} \sigma(\boldsymbol{x})$
- Repeat until entire $\mathcal{R}$ is observed indicated by small uncertainties of expanders

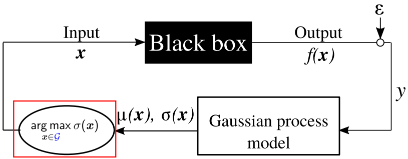

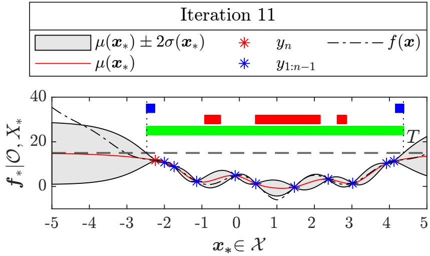

Exploration

- Observe the reachable set $\mathcal{R} = \{x \in \mathcal{X} | f(x) \leq T\}$
- $\boldsymbol{x}_{\text{new}} = \arg\max_{\boldsymbol{x} \in \mathcal{G}} \sigma(\boldsymbol{x})$
- Repeat until entire $\mathcal{R}$ is observed indicated by small uncertainties of expanders

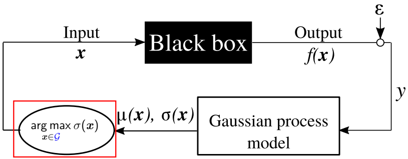

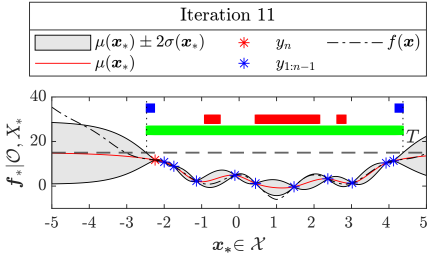

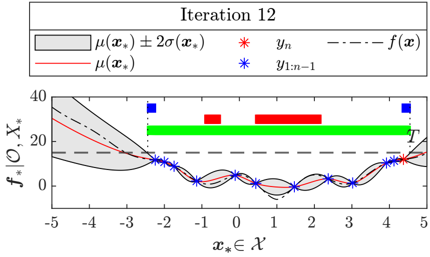

Exploitation

- Optimization step: find the minimum in $\mathcal{R}$
- Freeze the safe set
- Fit the hyperparameters $\theta = \{l, \sigma_f, \sigma_n\}$
- $\min_{\theta} -\log p(\boldsymbol{y}|X, \theta)$
- $\boldsymbol{x}_{\text{new}} = \arg\max_{\boldsymbol{x}\in\mathcal{M}} \text{EI}(\boldsymbol{x})$[3]

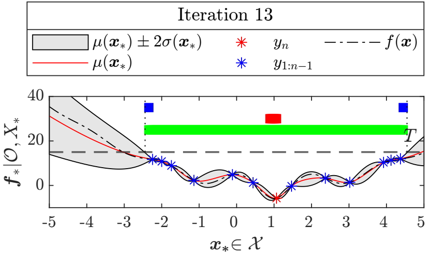

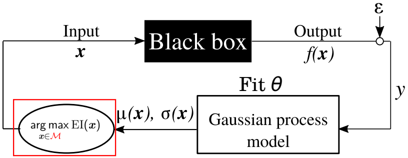

[3] Jones *et al.*, "Efficient Global Optimization of Expensive Black-Box Functions," 1998

Exploitation

- Optimization step: find the minimum in $\mathcal{R}$
- Freeze the safe set
- Fit the hyperparameters $\theta = \{l, \sigma_f, \sigma_n\}$
- $\min_{\theta} - \log p(\boldsymbol{y}|X, \theta)$
- $\boldsymbol{x}_{\text{new}} = \arg\max_{\boldsymbol{x} \in \mathcal{M}} \text{EI}(\boldsymbol{x})$[3]

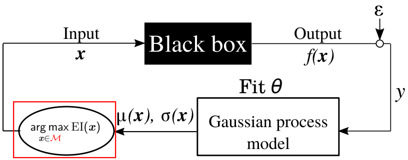

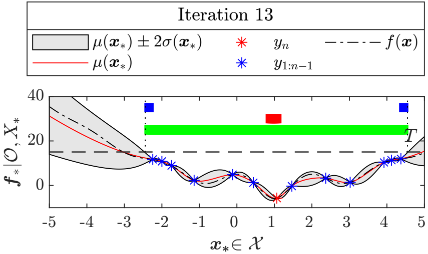

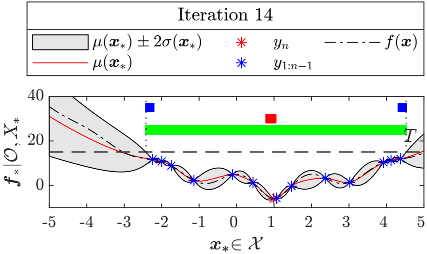

[3] Jones *et al.*, "Efficient Global Optimization of Expensive Black-Box Functions," 1998

Modified Safe Options (MoSaOpt)

- Safe options[2] evolved to modified safe options
- Safe set: $\mathcal{S} = \{x \in \mathcal{X} | \mathrm{UCB}(x) \leq T\}$
- Minimizer set: $\mathcal{M} = \{x \in \mathcal{S} | \mathrm{LCB}(x) \leq y^*\}$
- Expander set: $\mathcal{G} = \{x \in \mathcal{S} | \delta\mathcal{S}\}$

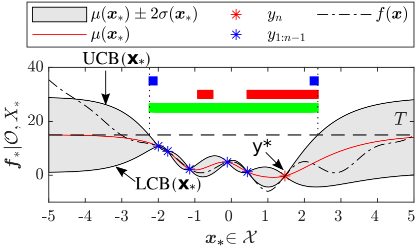

`MoSaOpt` divided into **exploration** and **exploitation** phase

- Efficient exploration by evaluating points at the boundaries
- Efficient exploitation by fitting the kernel on the data

$\Rightarrow$ Increased convergence speed

[2]Sui *et al.*, "Safe Exploration for Optimization with Gaussian Processes," 2015

Feasibility

Challenge

- Consider $\mathcal{X} \subseteq \mathbb{R}^D$
- Calculation of sets $\mathcal{S}$, $\mathcal{M}$, $\mathcal{G}$ are not numerically tractable for high D

Solution

- Apply Bayesian optimization on an iteratively changing subspace[4] $\mathcal{L} \subset \mathcal{X}$
- $\dim(\mathcal{L}) = 1 \rightarrow$ `LineBO`
- $\dim(\mathcal{L}) = 2 \rightarrow$ `PlaneBO`
- $x_{\text{opt}} = \underset{\boldsymbol{x}_i, y_i \in \mathcal{O}}{\arg\min}(1-\kappa)y_i + \kappa\mu(\boldsymbol{x}_i),\ 1 \geq \kappa \geq 0$

[4] Kirschner *et al.*, "Adaptive and Safe Bayesian Optimization in High Dimensions via One-Dimensional Subspaces," 2019.

Simulation - Synchronization System of the EuXFEL

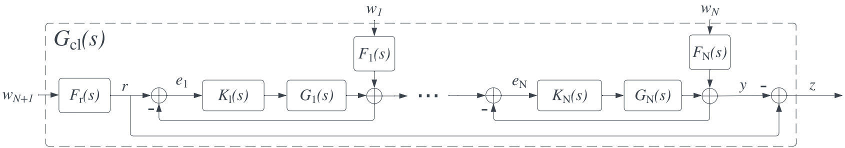

- $N = 5$ subsystems, each equipped with a PI controller
- Inputs $w_{1:N+1}$ are white Gaussian noise
- $\min \|z\|_{\mathrm{RMS}} = \min \|G_{\mathrm{cl}}(s)\|_{\mathcal{H}_2}$
- Compared to SafeOpt[5]

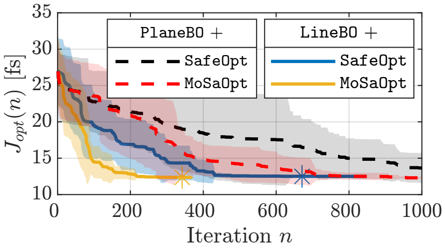

[5] Berkenkamp *et al.*, "Safe controller optimization for quadrotors with Gaussian processes," 2016

Experimental - Small Scale Synchronization System

- $N = 2$ subsystems
- $\|\cdot\|_{\mathrm{RMS}}$ of multiple measurements is averaged
- Compared to Nelder-Mead[6]
- Nelder-Mead shows higher noise sensitivity
- `MoSaOpt` finds the optimum approx. 4x faster

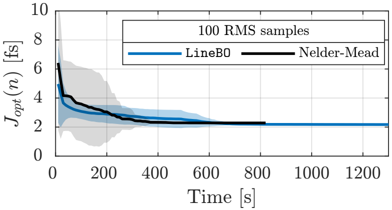

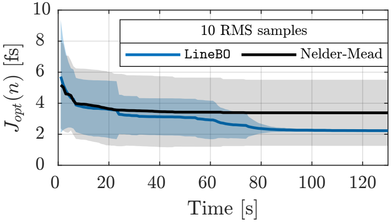

[6] Lagarias *et al.*, "Convergence Properties of the Nelder–Mead Simplex Method in Low Dimensions," 1998

Conclusion

Summary

- Sample efficient and noise robust Bayesian optimization procedure
- Increased convergence rate compared to other methods
- Applicable to high-dimensional optimization problems
- **All safeness guarantees are only valid if the true hyperparameters are known**

Outlook

- Extension to multitask Bayesian optimization
- Taking simulation into account
- Samples from simulator are cheap
- Find dependency between both tasks to increase convergence
- **How can theoretical guarantees be preserved?**

The End

Thank you very much for your attention!

[1] C. K. Williams and C. E. Rasmussen, *Gaussian processes for machine learning*. MIT press Cambridge, MA, 2006, vol. 2.

[2] Y. Sui, A. Gotovos, J. Burdick, and A. Krause, "Safe exploration for optimization with Gaussian processes," in *32nd Int. Conf. Mach. Learn. (ICML)*, ser. Proceedings of Machine Learning Research, vol. 37, Lille, France: PMLR, Jul. 2015, pp. 997–1005.

[3] D. R. Jones, M. Schonlau, and W. J. Welch, "Efficient global optimization of expensive black-box functions," *Journal of Global Optimization*, vol. 13, no. 4, pp. 455–492, Dec. 1998, ISSN: 1573-2916.

[4] J. Kirschner, M. Mutny, N. Hiller, R. Ischebeck, and A. Krause, "Adaptive and safe bayesian optimization in high dimensions via one-dimensional subspaces," in *36th Int. Conf. Mach. Learn. (ICML)*, 2019, pp. 3429–3438.

[5] F. Berkenkamp, A. P. Schoellig, and A. Krause, "Safe controller optimization for quadrotors with Gaussian processes," in *IEEE Int. Conf. Robot. Autom (ICRA)*, 2016, pp. 491–496.

[6] J. C. Lagarias, J. A. Reeds, M. H. Wright, and P. E. Wright, "Convergence properties of the nelder–mead simplex method in low dimensions," *SIAM Journal on Optimization*, vol. 9, no. 1, pp. 112–147, 1998.