Safe Risk-averse Bayesian Optimization for Controller Tuning

Christopher König Miks Ozols Anastasia Makarova Efe C. Balta Andreas Krause Alisa Rupenyan

ICRA2024 YOKOHAMA | JAPAN

ETH zürich

Black-box optimization arises in high-stakes applications

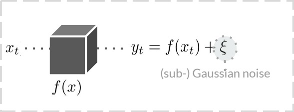

- Tuning complex systems

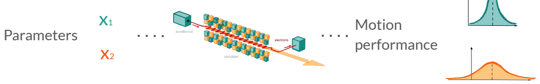

Problem 1: Two inputs x_1 and x_2 have **similar expected output values**, but one produces **much noisier realizations**

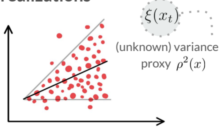

Problem 2: Input x_1 is **safe** and input x_2 is **unsafe**

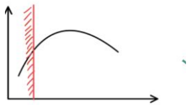

Takeaway

Practical optimization method that avoids choosing unsafe actions and takes into account unknown heteroscedastic noise.

The RaGoose Algorithm

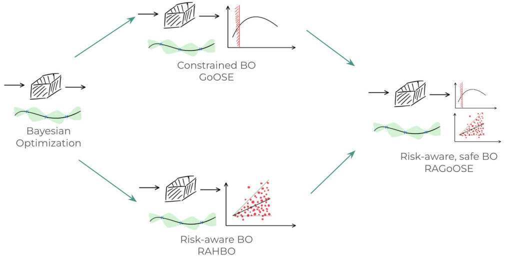

1. Sketch RAHBO

solutions

Challenge 1: confidence bounds for $f(x)$ rely on **known** $\rho^2(x)$
Challenge 2: estimate noise variance proxy $\rho^2(x)$.

- use conservative estimate to build $\mathrm{ucb}^f_t(x)$
- repeat experiment k times to get sample variance

- for $t = 1, 2, \dots$
  - choose $x_t \in \arg\max_{x \in \mathcal{X}} \mathrm{ucb}^f_{t-1}(x) - \alpha\, \mathrm{lcb}^{var}_{t-1}(x)$
  - collect k evaluations $\{y_i(x_t)\}_{i=1}^k$
  - use sample mean and variance of $\{y_i(x_t)\}$ for GP posteriors
  - update conf. bounds $\mathrm{lcb}^{var}_t(x)$ and $\mathrm{ucb}^{var}_t(x)$ for $\rho^2(x)$
  - update $\mathrm{ucb}^f_t(x) = \mu_t(x|\hat{\Sigma}_t) + \beta_t \sigma_t(x|\hat{\Sigma}_t)$

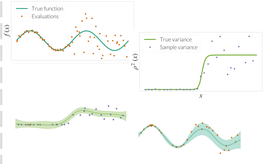

2. Sketch GoOSE

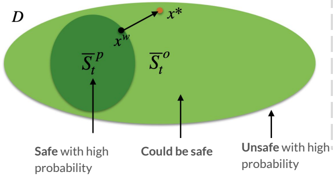

- for $t = 1, 2, \dots$
  - choose action $x^* \in \overline{S}_t^o$
  - if $x^* \in \overline{S}_t^p$: evaluate it and update $\overline{S}_t^p$ and $\overline{S}_t^o$
  - else: choose another action x^*
  - else: choose $\mathrm{argmin}_{x \in \overline{S}_t^p}(d(x, x^*))$

Experimental results

1. **Synthetic experiment:** sin function with three optima and sigmoid noise

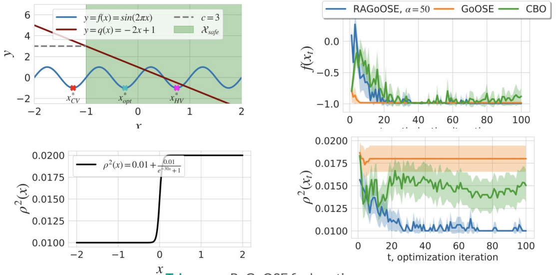

Takeaway: RaGoOSE finds optimum with lower observation noise .

2. **Real system:** linear axis of a high precision positioning system

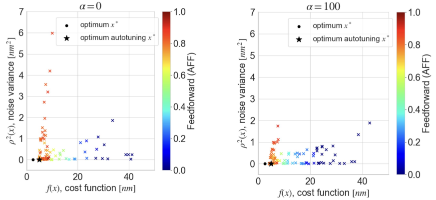

References

[1] Risk-averse heteroscedastic bayesian optimization; *Makarova et al; Neurips 2021*

[2] Safe and efficient model-free adaptive control via bayesian optimization; *Konig et al; ICRA 2021*

[3] Bayesian optimization with inequality constraints (CBO); *Gardner et al; ICML 2014.*

Paper:

Safe Risk-averse Bayesian Optimization for Controller Tuning

ICRA 2024 YOKOHAMA | JAPAN

ETH zürich

Christopher König Miks Ozols Anastasia Makarova Efe C. Balta Andreas Krause Alisa Rupenyan

Black-box optimization arises in high-stakes applications

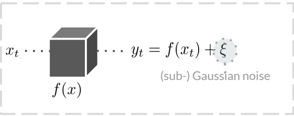

- Tuning complex systems

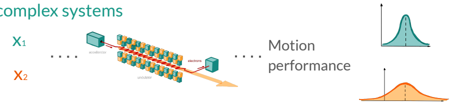

Problem 1: Two inputs x_1 and x_2 have **similar expected output values**, but one produces **much noisier realizations**

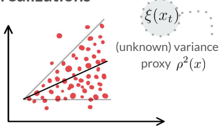

Problem 2: Input x_1 is **safe** and input x_2 is **unsafe**

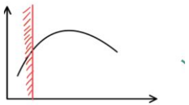

2. Sketch GoOSE

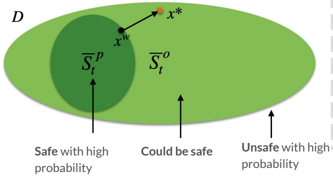

- for $t = 1, 2, \ldots$
  - choose action $x^* \in \overline{S}_t^o$
  - if $x^* \in \overline{S}_t^p$: evaluate it and update $\overline{S}_t^p$ and $\overline{S}_t^o$
  - else: choose another action x^*
  - else: choose $\operatorname{argmin}_{x \in \overline{S}_t^p} (d(x, x^*))$

Takeaway

Practical optimization method that avoids choosing unsafe actions and takes into account unknown heteroscedastic noise.

The RaGoose Algorithm

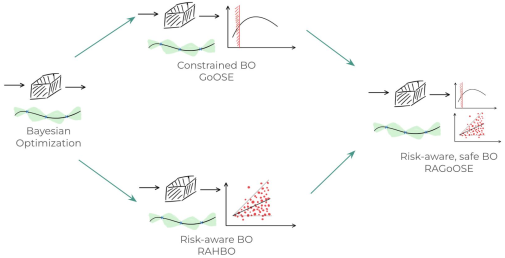

1. Sketch RAHBO

solutions

Challenge 1: confidence bounds for $f(x)$ rely on **known** $\rho^2(x)$

Challenge 2: estimate noise variance proxy $\rho^2(x)$.

- use conservative estimate to build $\mathrm{ucb}_t^f(x)$
- repeat experiment k times to get sample variance

- for $t = 1, 2, \ldots$
  - choose $x_t \in \arg\max_{x \in \mathcal{X}} \mathrm{ucb}_{t-1}^f(x) - \alpha\, \mathrm{lcb}_{t-1}^{var}(x)$
  - collect k evaluations $\{y_i(x_t)\}_{i=1}^k$
  - use sample mean and variance of $\{y_i(x_t)\}$ for GP posteriors
  - update conf. bounds $\mathrm{lcb}_t^{var}(x)$ and $\mathrm{ucb}_t^{var}(x)$ for $\rho^2(x)$
  - update $\mathrm{ucb}_t^f(x) = \mu_t(x|\hat{\Sigma}_t) + \beta_t \sigma_t(x|\hat{\Sigma}_t)$

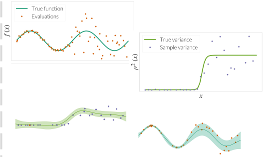

Experimental results

1. **Synthetic experiment:** sin function with three optima and sigmoid noise

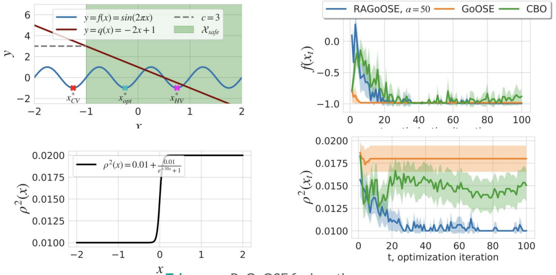

Takeaway: RaGoOSE finds optimum with lower observation noise .

2. **Real system:** linear axis of a high precision positioning system

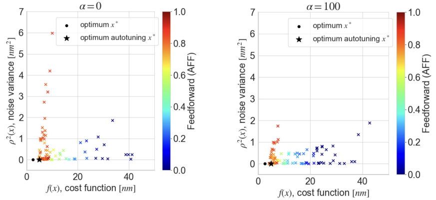

References

[1] Risk-averse heteroscedastic bayesian optimization; *Makarova et al; Neurips 2021*

[2] Safe and efficient model-free adaptive control via bayesian optimization; *Konig et al; ICRA 2021*

[3] Bayesian optimization with inequality constraints (CBO); *Gardner et al; ICML 2014.*

Paper: