

# Meta-Learning Priors for Safe Bayesian Optimization

Anonymous Author(s)

Affiliation

Address

email

1       **Abstract:** In robotics, optimizing controller parameters under safety constraints  
2 is an important challenge. Safe Bayesian optimization (BO) quantifies uncertainty  
3 in the objective and constraints to safely guide exploration in such settings. Hand-  
4 designing a suitable probabilistic model can be challenging however. In the pres-  
5 ence of unknown safety constraints, it is crucial to choose reliable model hyper-  
6 parameters to avoid safety violations. Here, we propose a data-driven approach to  
7 this problem by *meta-learning priors* for safe BO from offline data. We build on  
8 a meta-learning algorithm, F-PACOH, capable of providing reliable uncertainty  
9 quantification in settings of data scarcity. As core contribution, we develop a novel  
10 framework for choosing *safety-compliant priors* in a data-riven manner via empir-  
11 ical uncertainty metrics and a *frontier search* algorithm. On benchmark functions  
12 and a high-precision motion system, we demonstrate that our meta-learnt priors  
13 accelerate convergence of safe BO approaches while maintaining safety.

14       **Keywords:** Meta-learning, Safety, Controller tuning, Bayesian Optimization

## 15 1 Introduction

16 Optimizing a black-box function with as few queries as possible is a ubiquitous problem in science  
17 and engineering. Bayesian Optimization (BO) is a promising paradigm, which learns a probabilistic  
18 surrogate model (often a Gaussian process, GP) of the unknown function to guide exploration. BO  
19 has been successfully applied for optimizing sensor configurations [1, 2] or tuning the parameters of  
20 robotic controllers [3, 4, 5, 6, 7]. However, such real-world applications are often subject to safety  
21 constraints which must not be violated in the process of optimization, e.g., the robot not getting dam-  
22 aged. Often, the dependence of the safety constraints on the query inputs is a priori unknown and  
23 can only be observed by measurement. To cope with these requirements, safe BO methods [8, 9, 10]  
24 model both the objective *and constraint* functions with GPs. The uncertainty in the constraint is  
25 used to approximate the feasible region from within, to guarantee that no safety violations occur.  
26 Therefore, a critical requirement for safe BO is the reliability of the uncertainty estimates of our  
27 models. Typically, it is assumed that a correct GP prior, upon which the uncertainty estimates hinge,  
28 is exogenously given [e.g., 8, 9, 10]. In practice, however, appropriate choices for the kernel variance  
29 and lengthscale are unknown and typically have to be chosen very conservatively or hand-tuned by  
30 trial and error, a problematic endeavour in a safety-critical setting. A too conservative choice dra-  
31 matically reduces sample efficiency, whereas overestimating smoothness may risk safety violations.

32 Addressing these shortcomings, we develop an approach for *meta-learning informative, but safe,*  
33 *GP priors in a data-driven way* from related offline data. We build on the F-PACOH meta-learning  
34 method [11] which is capable of providing reliable uncertainty estimates, even in the face of  
35 data-scarcity and out-of-distribution data. However, their approach still relies on an appropriate  
36 kernel choice on which it falls back in the absence of sufficient data. We propose a novel framework  
37 for choosing safety-compliant kernel hyper-parameters in a data-driven manner based on calibration  
38 and sharpness metrics of the confidence intervals. To optimize these uncertainty metrics, we devise a  
39 *frontier search* algorithm that efficiently exploits the monotone structure of the problem. The result-  
40 ing *Safe Meta-Bayesian Optimization* (SAMBO) approach can be instantiated with existing safe BO

41 methods and utilize the improved, meta-learned GPs to perform safe optimization more efficiently.  
 42 In our experiments, we evaluate and compare our proposed approach on benchmark functions as  
 43 well as controller tuning for a high-precision motion system. Throughout, SAMBO significantly  
 44 improves the query efficiency of popular safe BO methods, without compromising safety.

## 45 2 Related Work

46 **Safe BO** aims to efficiently optimize a black-box function under safety-critical conditions, where un-  
 47 known safety constraints must not be violated. Constrained variants of standard BO methods [12, 13,  
 48 14] return feasible solutions, but do not reliably exclude unsafe queries. In contrast, SAFEOPT [8, 9]  
 49 and related variants [15] guarantee safety at all times and have been used to safely tune controllers  
 50 in various applications [e.g., 16, 17]. While SAFEOPT explores in an undirected manner, GOOSE  
 51 [10, 18] does so by expanding the safe set in a goal-oriented fashion. All mentioned methods rely on  
 52 GPs to model the target and constraint function and assume that correct kernel hyper-parameters are  
 53 given. Our work is complementary: We show how to use related offline data to obtain informative  
 54 and safe GP priors in a data-driven way that makes the downstream safe BO more query efficient.

55 **Meta-Learning.** Common approaches in meta-learning amortize inference [19, 20, 21], learn a  
 56 shared embedding space [22, 23, 24, 25] or a good neural network initialization [26, 27, 28, 29].  
 57 However, when the available data is limited, these approaches are prone to overfit on the meta-  
 58 level. A body of work studies meta-regularization to prevent overfitting in meta-learning  
 59 [30, 31, 32, 33, 34, 35]. Such meta-regularization methods prevent meta-overfitting for the mean  
 60 predictions, but not for the uncertainty estimates. Recent meta-learning methods aim at providing  
 61 reliable confidence intervals even when data is scarce and non-i.i.d [11, 36, 37]. These methods ex-  
 62 tend meta-learning to interactive and life-long settings. However, they either make unrealistic model  
 63 assumptions or hinge on hyper-parameters whose improper choice is critical in safety constraint set-  
 64 tings. Our work uses F-PACOH [11] to meta-learn reliable GP priors, but removes the need for hand-  
 65 specifying a correct hyper-prior by choosing its parameters in a data-driven, safety-aware manner.

## 66 3 Problem Statement and Background

### 67 3.1 Problem Statement

68 We consider the problem of safe Bayesian Optimization (safe BO), seeking the global minimizer

$$\mathbf{x}^* = \arg \min_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}) \quad \text{s.t.} \quad q(\mathbf{x}) \leq 0 \quad (1)$$

69 of a function  $f : \mathcal{X} \rightarrow \mathbb{R}$  over a bounded domain  $\mathcal{X}$  subject to a safety constraint  $q(\mathbf{x}) \leq 0$  with  
 70 constraint function  $q : \mathcal{X} \rightarrow \mathbb{R}$ . For instance, we may want to optimize the controller parameter of a  
 71 robot without subjecting it to potentially damaging vibrations or collisions. During the optimization,  
 72 we iteratively make queries  $\mathbf{x}_1, \dots, \mathbf{x}_T \in \mathcal{X}$  and observe noisy feedback  $\tilde{f}_1, \dots, \tilde{f}_T$  and  $\tilde{q}_1, \dots, \tilde{q}_T$ ,  
 73 e.g., via  $\tilde{f}_t = f(\mathbf{x}_t) + \epsilon_f$  and  $\tilde{q}_t = q(\mathbf{x}_t) + \epsilon_q$  where  $\epsilon_f, \epsilon_q$  is  $\sigma^2$  sub-Gaussian noise [38, 39].  
 74 In our setting, performing a query is assumed to be costly, e.g., running the robot and observing  
 75 relevant measurements. Hence, we want to find a solution as close to the global minimum in as few  
 76 iterations as possible without violating the safety constraint with any query we make.

77 Additionally, we assume to have access to datasets  $\mathcal{D}_{1,T_1}, \dots, \mathcal{D}_{n,T_n}$  with observations from  $n$  sta-  
 78 tistically similar but distinct data-generating systems, e.g., data of the same robotic platform under  
 79 different conditions. Each dataset  $\mathcal{D}_{i,T_i} = \{(\mathbf{x}_{i,1}, \tilde{f}_{i,1}, \tilde{q}_{i,1}), \dots, (\mathbf{x}_{i,T_i}, \tilde{f}_{i,T_i}, \tilde{q}_{i,T_i})\}$  consists of  $T_i$   
 80 measurement triples, where  $\tilde{f}_{i,t} = f_i(\mathbf{x}_{i,t}) + \epsilon_{f_i}$  and  $\tilde{q}_{i,t} = q_i(\mathbf{x}_{i,t}) + \epsilon_{q_i}$  are the noisy target and  
 81 constraints observations. We assume that the underlying functions  $f_1, \dots, f_n$  and  $q_1, \dots, q_n$  which  
 82 generated the data are representative of our current target and constraint functions  $f$  and  $q$ , e.g., that  
 83 they are all i.i.d. draws from the same stochastic process. However, the data within each dataset  
 84 may be highly dependent (i.e., non-i.i.d.). For instance, each  $\mathcal{D}_{i,T_i}$  may correspond to the queries  
 85 and observations from previous safe BO sessions with the same robot under different conditions.

86 In this paper, we ask the question of how we can harness such related data sources to make safe BO  
 87 on our current problem of interest more query efficient, without compromising safety.

### 88 3.2 Safe Bayesian Optimization Methods

89 Safe BO methods construct a Bayesian *surrogate model* of the functions  $f$  and  $q$  based on previous  
90 observations  $\mathcal{D}_t = \{(\mathbf{x}_{t'}, \tilde{f}_{t'}, \tilde{q}_{t'})\}_{t' < t}$ . Typically, a Gaussian Process GP ( $f(\mathbf{x})|m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}')$ )  
91 with mean  $m(\mathbf{x})$  and kernel function  $k(\mathbf{x}, \mathbf{x}')$  is employed to form posterior beliefs  $p(f(\mathbf{x})|\mathcal{D}_t) =$   
92  $\mathcal{N}(\mu_t^f(\mathbf{x}), (\sigma_t^f(\mathbf{x}))^2)$  and  $p(q(\mathbf{x})|\mathcal{D}_t) = \mathcal{N}(\mu_t^q(\mathbf{x}), (\sigma_t^q(\mathbf{x}))^2)$  over function values [40]. Based on  
93 the predictive posterior, we can form confidence intervals (CIs) to the confidence level  $\alpha \in [0, 1]$

$$CI_\alpha^f(\mathbf{x}|\mathcal{D}_t) := \left[ \mu_t^f(\mathbf{x}) \pm \beta^f(\alpha) \sigma_t^f(\mathbf{x}) \right] \quad (2)$$

94 where  $\beta^f(\alpha)$ , the scaling of the standard deviation, is set such that  $f(x)$  is in the CI with probability  
95  $\alpha$ . For BO, we often employ a shift invariant kernel  $k(\mathbf{x}, \mathbf{x}') = \nu \phi(\|\mathbf{x} - \mathbf{x}'\|/l)$ , where  $\nu$  is its vari-  
96 ance,  $l$  the lengthscale and  $\phi$  a positive function, e.g., squared exponential (SE)  $\phi(z) = \exp(-z^2)$ .

97 BO methods typically choose their query points by maximizing an acquisition function based on  
98  $p(f(\mathbf{x})|\mathcal{D}_t)$ , trading-off exploration and exploitation [41, 42, 43, 44]. When we have safety con-  
99 straints, we need to maintain a *safe set*  $\mathcal{S}_t(\alpha) = \{\mathbf{x} \in \mathcal{X} | \mu_t^q(\mathbf{x}) + \beta^q(\alpha) \sigma_t^q(\mathbf{x}) < 0\}$  which contains  
100 parts of the domain we know to be safe with high-probability  $\alpha$ . To maintain safety, we can only  
101 query points within the current  $\mathcal{S}_t(\alpha)$ . In addition, we need to explore w.r.t.  $q$  so that we can expand  
102  $\mathcal{S}_t$ . E.g., SAFEBOPT [8, 9, 16] computes a safe query candidate that optimizes the acquisition function  
103 for  $f$  as well as a query candidate that promises the best expansion of  $\mathcal{S}_t$ . Then, it selects the one  
104 with the highest uncertainty. While SAFEBOPT expands  $\mathcal{S}_t$  undirectedly, GOOSE [10, 18] does so in  
105 a more directed manner to avoid unnecessary expansion queries, irrelevant for minimizing  $f$ . In ad-  
106 dition to the pessimistic  $\mathcal{S}_t$ , it also maintains an optimistic safe set which is used to calculate a query  
107 candidate  $\mathbf{x}_t^{\text{opt}}$  that maximizes the acquisition function for  $f$ . If  $\mathbf{x}_t^{\text{opt}}$  is outside of  $\mathcal{S}_t$ , it chooses safe  
108 query points aiming at expanding  $\mathcal{S}_t$  in the direction of  $\mathbf{x}_t^{\text{opt}}$ . See Appx. A for more details.

### 109 3.3 Meta-Learning GP Priors

110 In meta-learning [45, 46], we aim to extract prior knowledge (i.e., inductive bias) from a set of  
111 related learning tasks. Typically, the meta-learner is given such learning tasks in the form of  $n$   
112 datasets  $\mathcal{D}_{1, T_1}, \dots, \mathcal{D}_{n, T_n}$  with  $\mathcal{D}_{i, T_i} = \{(\mathbf{x}_{i, t}, y_{i, t})\}_{t=1}^{T_i}$  and outputs a refined prior distribution or  
113 hypothesis space which can then be used to accelerate inference on a new, but related, learning task.

114 Prior work proposes to meta-learning GP priors [47, 48, 34], though, fails to maintain reliable  
115 uncertainty estimates when data is scarce and/or non-i.i.d.. The recently introduced F-PACOH  
116 method [11] overcomes this issue, by using a regularization approach in the function space. As  
117 previous work [e.g., 49, 48], they use a learnable GP prior  $\rho_\theta(h) = \text{GP}(h(\mathbf{x})|m_\theta(\mathbf{x}), k_\theta(\mathbf{x}, \mathbf{x}'))$   
118 where the mean and kernel function are neural networks with parameters  $\theta$  and employ the marginal  
119 log-likelihood to fit the  $\rho_\theta(h)$  to the meta-training data. However, during the meta-learning, they  
120 regularize  $\rho_\theta(h)$  towards a Vanilla GP hyper-prior  $\rho(h) = \text{GP}(h(\mathbf{x})|0, k(\mathbf{x}, \mathbf{x}'))$  with the SE kernel.

121 To do so, they uniformly sample random measurement sets  $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_m] \stackrel{\text{i.i.d.}}{\sim} \mathcal{U}(\mathcal{X})$  from the  
122 domain and compare the GPs finite marginals  $\rho_\theta(\mathbf{h}^{\mathbf{X}}) = p_\theta(h(\mathbf{x}_1), \dots, h(\mathbf{x}_m))$  and  $\rho(\mathbf{h}^{\mathbf{X}})$  through  
123 their KL-divergence. The resulting meta-learning loss with the functional KL regularizer

$$\mathcal{L}(\theta) := \frac{1}{n} \sum_{i=1}^n \left( \underbrace{-\frac{1}{T_i} \ln Z(\mathcal{D}_{i, T_i}, \rho_\theta)}_{\text{marginal log-likelihood}} + \underbrace{\left( \frac{1}{\sqrt{n}} + \frac{1}{nT_i} \right) \mathbb{E}_{\mathbf{X}} [KL[\rho_\theta(\mathbf{h}^{\mathbf{X}}) || \rho(\mathbf{h}^{\mathbf{X}})]]}_{\text{functional KL-divergence}} \right) \quad (3)$$

124 makes sure that, in the absence of sufficient meta-training data, the learned GP behaves like a Vanilla  
125 GP. Overall, this allows us to meta-learn a more informative GP prior which still yields reliable  
126 confidence intervals, even if the meta-training data was collected via safe BO and is thus not i.i.d.

## 127 4 Choosing the Safe Kernel Hyper-Parameters

128 Important for safe BO is the reliability of the uncertainty estimates of our objective and constraint  
129 models. For GPs, the kernel hyper-parameters with the biggest influence on the CIs. If the kernel  
130 variance  $\nu$  is chosen too small and/or the lengthscale  $l$  to large, our models become over-confident

131 and the corresponding BO unsafe. In the reverse case, the CIs become too conservative and safe BO  
 132 requires many queries to progress. Despite the assumption commonly made in earlier work, [e.g., 8,  
 133 9, 10], appropriate choices for  $\nu$  and  $l$  are unknown. In practice, they are typically chosen conserva-  
 134 tively or hand-tuned by trial and error, problematic in safety-critical settings. Aiming to address this  
 135 issue, we develop a framework for choosing the kernel hyper-parameters in a data-driven manner.

#### 136 4.1 Assessing kernel hyper-parameters: Calibration and sharpness

137 Our approach is based on the *calibration* and *sharpness* of uncertainty estimates [see e.g. 50, 51, 52,  
 138 53]. Naturally, if we construct CIs to the confidence level  $\alpha$ , we want that at least an  $\alpha$  percentage of  
 139 (unseen) observations to fall within these CIs. If this holds in expectation, we say that the uncertainty  
 140 estimates are *calibrated*. If the empirical percentage is less than  $\alpha$ , it indicates that our model’s un-  
 141 certainty estimates are over-confident and we are likely to underestimate the risk of safety violations.  
 142 To empirically assess how calibrated a probabilistic regression model with hyper-parameters  $\omega$ , con-  
 143 ditioned on a training dataset  $\mathcal{D}^{\text{tr}}$ , is, we compute its *calibration frequency* on a test dataset  $\mathcal{D}^{\text{test}}$ :

$$\text{calib-freq}(\mathcal{D}^{\text{tr}}, \mathcal{D}^{\text{test}}, \omega) := \frac{1}{|A|} \sum_{\alpha \in A} \mathbb{1} \left( \frac{1}{|\mathcal{D}^{\text{test}}|} \sum_{(\mathbf{x}, y) \in \mathcal{D}^{\text{test}}} [\mathbb{1}(y \in CI_{\alpha}^f(\mathbf{x} | \mathcal{D}^{\text{tr}}, \omega))] \geq \alpha \right). \quad (4)$$

144 Here,  $A \subset [0, 1]$  is a set of relevant confidence levels (in our case 20 values equally spaced between  
 145 0.8 and 1.0). Since the CIs of our model need to be calibrated at any iteration  $t$  during the BO  
 146 and for any task we may face, we choose the best empirical estimate we can. We compute the  
 147 average calibration frequency across all meta-training datasets and for any sub-sequence of points  
 148 within a dataset. In particular, for any task  $i = 1, \dots, n$  and  $t = 1, \dots, T_i - 1$  we condition/train  
 149 our model on the data points  $\mathcal{D}_{i, \leq t} = \{(\mathbf{x}_{i, t'}, y_{i, t'})\}_{t'=1}^t$  and use the remaining data points  
 150  $\mathcal{D}_{i, > t} = \{(\mathbf{x}_{i, t'}, y_{i, t'})\}_{t'=t+1}^{T_i}$  to compute the calibration frequency. Overall, this gives us

$$\text{avg-calib}(\{\mathcal{D}_{i, T_i}\}_{i=1}^n, \omega) := \frac{1}{n} \sum_{i=1}^n \frac{1}{T_i - 1} \sum_{t=1}^{T_i - 1} \text{calib-freq}(\mathcal{D}_{i, \leq t}, \mathcal{D}_{i, > t}, \omega). \quad (5)$$

151 While the calibration captures how *reliable* the uncertainty estimates are, it does not reflect how  
 152 *useful* the confidence intervals are for narrowing down the range of possible function values. For  
 153 instance, a predictor that always outputs a mean with sufficiently wide confidence intervals is  
 154 calibrated, but useless for BO. Hence, we also consider the *sharpness* of the uncertainty estimates  
 155 which we empirically quantify through the *average predictive standard deviation*. Similar to (5),  
 156 we average over all tasks and data sub-sequences within each task:

$$\text{avg-std}(\{\mathcal{D}_{i, T_i}\}_{i=1}^n, \omega) := \frac{1}{n} \sum_{i=1}^n \frac{1}{T_i - 1} \sum_{t=1}^{T_i - 1} \frac{1}{|\mathcal{D}_{i, > t}|} \sum_{(\mathbf{x}, y) \in \mathcal{D}_{i, > t}} \sigma(\mathbf{x} | \mathcal{D}_{i, \leq t}, \omega). \quad (6)$$

157 The avg-std measures how concentrated the uncertainty estimates are and, thus, constitutes a natural  
 158 complement to calibration which can be simply achieved by wide/loose confidence intervals.

#### 159 4.2 Choosing good hyper-parameters via Frontier search

160 Based on the two empirical quantities introduced above, we can optimize the hyper-parameters  $\omega$  of  
 161 our model as to maximize sharpness (i.e., minimize the avg-std) subject to calibration [50]:

$$\min_{\omega} \text{avg-std}(\{\mathcal{D}_{i, T_i}\}_{i=1}^n, \omega) \quad \text{s.t.} \quad \text{avg-calib}(\{\mathcal{D}_{i, T_i}\}_{i=1}^n, \omega) \geq 1 \quad (7)$$

162 Since computing  $\text{avg-std}(\{\mathcal{D}_{i, T_i}\}_{i=1}^n, \omega)$  and  $\text{avg-calib}(\{\mathcal{D}_{i, T_i}\}_{i=1}^n, \omega)$  requires solving the the GP  
 163 inference problem many times, each query is computationally demanding. Hence, we need an opti-  
 164 mization algorithm for (7) that requires as few queries as possible to get close to the optimal solution.

165 We develop an efficient *frontier search (FS)* algorithm that exploits the monotonicity properties  
 166 of this optimization problem. Both avg-std and avg-calib are monotonically increasing in the  
 167 kernel variance  $\nu$  and decreasing in the lengthscale  $l^1$ . By setting  $\mathbf{z} = (-l, \nu)$  and writing

<sup>1</sup>Note that the monotonicity of the calibration frequency in  $l$  is only an empirical heuristic that holds in almost all cases if  $\nu$  is at least as big as the variance of the targets  $y$  in a dataset.

---

**Algorithm 1** FRONTIERSEARCH (details in Appendix C)
 

---

**Input:** Domain bounds  $\mathbf{z}^l, \mathbf{z}^u$  s.t.  $\mathbf{z}^l \leq \mathbf{z}^* \leq \mathbf{z}^u$

- 1:  $Q^u \leftarrow \{\mathbf{z}^u\}, Q^l \leftarrow \{\mathbf{z}^l\}$
- 2: **for**  $k = 1, \dots, K$  **do**
- 3:    $(\mathbf{z}_r, \mathbf{z}'_r) \leftarrow \text{LARGESTMAXMINRECT}(Q^l, Q^u)$    // Largest max-min rect betw. frontiers
- 4:    $\mathbf{z}_q \leftarrow \text{BESTWORSTCASEQUERY}(\mathbf{z}_r, \mathbf{z}'_r, Q^l, Q^u)$    // Best query point to split rectangle
- 5:   **if**  $c(\mathbf{z}_q) \geq 1$  **then**
- 6:      $Q_u \leftarrow \text{PRUNE}(Q^u \cup \{\mathbf{z}_q\})$
- 7:   **else**
- 8:      $Q_l \leftarrow \text{PRUNE}(Q^l \cup \{\mathbf{z}_q\})$

**Return:**  $\arg \min_{\mathbf{z} \in Q^u} s(\mathbf{z})$

---

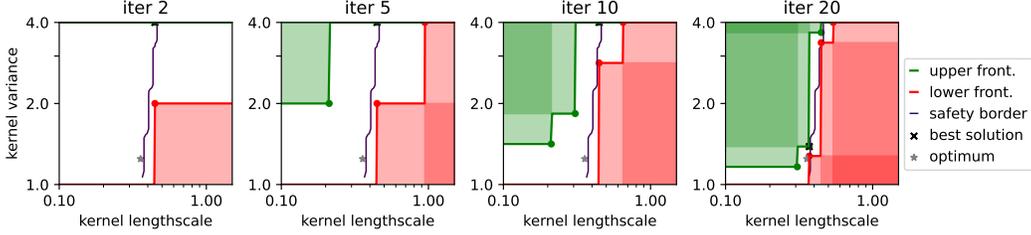


Figure 1: Frontier search (FS) on the kernel lengthscale and variance for the constraint model Argus robot. Red: areas ruled out, because unsafe. Green: Safe areas that are ruled out since dominated by better safe queries. After a few iterations, FS has already shrunk the set of possible optima (white area between fronts) to points close to the safety boarder and picked nearly optimal solution (cross).

168  $s(\mathbf{z}) = \text{avg-std}(\{\mathcal{D}_{i,T_i}\}_{i=1}^n, l, \nu)$  and  $c(\mathbf{z}) = \text{avg-calib}(\{\mathcal{D}_{i,T_i}\}_{i=1}^n, l, \nu)$ , we can turn (7) into

$$\min_{\mathbf{z}} s(\mathbf{z}) \quad \text{s.t.} \quad c(\mathbf{z}) \geq 1 \quad \text{where } s(\mathbf{z}) : \mathbb{R}^2 \mapsto \mathbb{R} \text{ and } c(\mathbf{z}) : \mathbb{R}^2 \mapsto \mathbb{R} \text{ are monotone.} \quad (8)$$

169 We presume an upper and lower bound  $(\mathbf{z}^u, \mathbf{z}^l)$  such that resulting search domain  $\mathcal{Z} = [z_1^l, z_1^u] \times$   
 170  $[z_2^l, z_2^u]$  contains the optimal solution  $\mathbf{z}^* = \arg \min_{\mathbf{z}: c(\mathbf{z}) \geq 1} s(\mathbf{z})$ . Since both  $s(\mathbf{z})$  and  $c(\mathbf{z})$  are mono-  
 171 tone we know that  $\mathbf{z}^*$  must lie on or directly above the constraint boundary  $c(\mathbf{z}) = 1$  (see Lemma 2).

172 In each iteration  $k$  of Algorithm 1 we query a point  $\mathbf{z}_k^q \in \mathcal{Z}$  and observe the corresponding objective  
 173 and constraint values  $s(\mathbf{z}_k^q)$  and  $c(\mathbf{z}_k^q)$ . We separate the queries into two sets  $Q^u$  and  $Q^l$  based on  
 174 whether they lie above or below the constraint boundary. That is, we add  $\mathbf{z}_k^q$  to  $Q^u$  if  $c(\mathbf{z}_k^q) \geq 1$  and  
 175 to  $Q^l$  otherwise. Since the optimal solution lies on the constraint boundary and  $c(\mathbf{z})$  is monotone,  
 176 we can rule out entire corners of the search domain: For each  $\mathbf{z}^q \in Q^u$  we can rule all points  $\mathbf{z}' > \mathbf{z}_k^q$   
 177 as candidates for the optimal solution and, similarly for all  $\mathbf{z}^q \in Q^l$ , we can rule out all  $\mathbf{z}' \leq \mathbf{z}_k^q$ .  
 178 This also allows us to *prune* the sets  $Q^u$  and  $Q^l$  by removing all the points from them that can be  
 179 ruled out by a new query results. To keep track which parts of  $\mathcal{Z}$  have not been ruled out yet, we  
 180 construct an *upper and lower frontiers*, here expressed as functions  $z_1 \mapsto z_2$  and  $z_2 \mapsto z_1$ ,

$$F_2^u(z_1; Q^u) = \min\{z_2 \mid z_1 \geq z_1', \mathbf{z}' \in Q^u\}, \quad F_1^u(z_2; Q^u) = \min\{z_1 \mid z_2 \geq z_2', \mathbf{z}' \in Q^u\} \quad (9)$$

$$F_2^l(z_1; Q^l) = \max\{z_2 \mid z_1 \leq z_1', \mathbf{z}' \in Q^l\}, \quad F_1^l(z_2; Q^l) = \max\{z_1 \mid z_2 \leq z_2', \mathbf{z}' \in Q^l\} \quad (10)$$

181 such that the points  $\Gamma = \{(z_1, z_2) \in \mathcal{Z} \mid F_2^l(z_1; Q^l) \leq z_2 \leq F_2^u(z_1; Q^u)\} \subseteq \mathcal{Z}$  between the  
 182 frontiers are still plausible candidates for the optimal solution. For notational brevity, we define  
 183  $\mathcal{F}^u = \{\mathbf{z} \in \mathcal{Z} \mid F_1^u(z_2; Q^u) = z_1 \vee F_2^u(z_1; Q^u) = z_2\}$  and  $\mathcal{F}^l$  analogously as the set of points that  
 184 lie on the upper and lower frontier respectively.

185 At any point, the best solution to (8) that we can give, is the best query we have made so far that  
 186 fulfills the constraint, i.e.,  $\hat{\mathbf{z}} = \arg \min_{\mathbf{z} \in Q^u} s(\mathbf{z})$ . If we assume Lipschitz continuity for  $s$ , which  
 187 holds in our case holds since  $\mathcal{Z}$  is bounded and the avg-std is differentiable in  $l$  and  $\nu$ , we can bound  
 188 how much our best solution is away from the optimum, i.e.,  $s(\hat{\mathbf{z}})$ :

189 **Lemma 1.** *Let  $s(\mathbf{z})$  be  $L$  Lipschitz and  $d(\Gamma, \mathcal{F}^u) := \max_{\mathbf{z}' \in \Gamma} \min_{\mathbf{z} \in \mathcal{F}^u} \|\mathbf{z} - \mathbf{z}'\|$  the max-min  
 190 distance between the frontiers. Then, the sub-optimality is bounded by  $s(\hat{\mathbf{z}}) - s(\mathbf{z}^*) \leq L d(\Gamma, \mathcal{F}^u)$ .*

---

**Algorithm 2** Safe Meta-BO (SAMBO)

---

**Input:** Safe BO problem with  $f^{tar}$  and  $q^{tar}$ , set of datasets  $\{\mathcal{D}_{i,T_i}^f\}_{i=1}^n, \{\mathcal{D}_{i,T_i}^q\}_{i=1}^n$

1: **for**  $h \in \{f, q\}$  **do**

2:    $(l_h, \nu_h) \leftarrow \text{FRONTIERSEARCH}(\{\mathcal{D}_{i,T_i}^h\}_{i=1}^n)$

3:    $\rho_{\theta_h}(h) \leftarrow \text{F-PACOH}(\{\mathcal{D}_{i,T_i}^h\}_{i=1}^n, \rho_{l_h, \nu_h}(h))$

**Return:**  $\hat{\mathbf{z}}^* \leftarrow \text{SAFEBO}(f^{tar}, q^{tar}, \rho_{\theta_f}(f), \rho_{\theta_q}(q))$

---

191 Here, the key insight is that we can bound the sub-optimality  $s(\hat{\mathbf{z}}) - s(\mathbf{z}^*)$  with the maximum  
 192 distance of any point between the frontiers from its closest point on the upper frontier instead of  $\hat{\mathbf{z}}$ .  
 193 This is the case because  $\hat{\mathbf{z}}$  dominates any point on the upper frontier (i.e.,  $s(\hat{\mathbf{z}}) \leq s(\mathbf{z}') \forall \mathbf{z}' \in \mathcal{F}^u$ ).

194 Hence, we want to choose the next query so that we can shrink the max-min distance  $d(\Gamma, \mathcal{F}^u)$   
 195 between the frontiers the most. For this purpose, we select the largest max-min rectangle between  
 196 the frontiers (see Definition 5). Then, we choose the query point, that reduce the max-min distance  
 197 within the rectangle the best. For this we need to consider two scenarios that will affect the max-min  
 198 distance differently: either the query point satisfies the constraint ( $c(\mathbf{z}^q) \geq 1$ ) or it does not ( $c(\mathbf{z}^q) <$   
 199  $1$ ). We compute the rectangle’s max-min distance for both scenarios and choose the query-point that  
 200 gives us the lowest max-min distance in the less-favorable (worst-case) scenario. For more details we  
 201 refer to Appendix C. Finally, we provide worst-case rates for the proposed frontier search algorithm:

202 **Theorem 1.** *Under the assumptions of Lemma 1, Algorithm 1 needs no more than  $k \leq 3^{\lceil \log_2(1/\epsilon) \rceil} =$   
 203  $\mathcal{O}((1/\epsilon)^{1.59})$  iterations to have a sub-optimality of less than  $s(\hat{\mathbf{z}}) - s(\mathbf{z}^*) \leq L\|\mathbf{z}^u - \mathbf{z}^l\| (1/\epsilon)$ .*

204 The  $\mathcal{O}((1/\epsilon)^{1.59})$  query complexity is more efficient than those of similar algorithms that do not  
 205 make use of the monotonicity properties, e.g., that of grid search:  $\mathcal{O}((1/\epsilon + 1)^2)$ .

## 206 5 Safe BO with meta-learned GP priors

207 In Sec. 4 we discuss how to efficiently choose the kernel variance and lengthscales, so that we obtain  
 208 calibrated and yet sharp uncertainty estimates. Now, we go one step further and use the related  
 209 datasets  $\mathcal{D}_{1,T_1}, \dots, \mathcal{D}_{n,T_n}$  to meta-learn GP priors, aiming to give the downstream safe BO algorithm  
 210 more prior knowledge about the optimization problem at hand so that it can be more query efficient.

211 For this, we use the F-PACOH [11] meta-learning approach, introduced in Sec 3. We choose  
 212 this method since it regularizes the meta-learned model in the function space and, thus, is able  
 213 to maintain reliability of the uncertainty estimates, even for out-of-distribution data. F-PACOH  
 214 requires a Vanilla GP  $\rho(h) = \mathcal{GP}(h(\mathbf{x})|0, k(\mathbf{x}, \mathbf{x}'))$  as a hyper-prior, towards which it regularizes  
 215 the meta-learned GP  $\rho_{\theta}(h) = \mathcal{GP}(h(\mathbf{x})|m_{\theta}(\mathbf{x}), k_{\theta}(\mathbf{x}, \mathbf{x}'))$ . Hence, we face a similar kernel hyper-  
 216 parameter choice problem as addressed in Section 4: If the kernel parameters of the hyper-prior are  
 217 chosen such that the Vanilla GP itself yields over-confident uncertainty estimates, safe BO with the  
 218 meta-learned prior will most likely turn out to be unsafe as well. Hence, we use the calibration-  
 219 sharpness based frontier search to choose the kernel variance and lengthscales of the hyper-prior.

220 Since safe BO maintains separate models for the target  $f(\mathbf{x})$  and constraint  $q(\mathbf{x})$ , we meta-learn  
 221 a prior for each of the functions. We split the sets of data triplets  $\mathcal{D}_{i,T_i} = \{(\mathbf{x}_{i,t}, \tilde{f}_{i,t}, \tilde{q}_{i,t})\}_{t=1}^{T_i}$   
 222 into separate datasets  $\mathcal{D}_{i,T_i}^f = \{(\mathbf{x}_{i,t}, \tilde{f}_{i,t})\}_{t=1}^{T_i}$  and  $\mathcal{D}_{i,T_i}^q = \{(\mathbf{x}_{i,t}, \tilde{q}_{i,t})\}_{t=1}^{T_i}$ . First, we use frontier  
 223 search to find kernel parameters  $(l_f, \nu_f)$  and  $(l_q, \nu_q)$  that give good sharpness subject to calibration  
 224 on the respective set of datasets. Then, we meta-learn GP priors  $\rho_{\theta_f}(f)$  and  $\rho_{\theta_q}(q)$  for  $f$  and  $q$   
 225 with F-PACOH, while using the Vanilla GPs  $\rho_{l_f, \nu_f}(f)$  and  $\rho_{l_q, \nu_q}(q)$  with the chosen kernel  
 226 hyper-parameters as hyper-prior. Finally, we run a safe BO algorithm (either SAFEBO or GOOSE)  
 227 with the meta-learned GP priors and perform safe Bayesian optimization on our target problem of  
 228 interest. This procedure is summarized in Algorithm 2. We refer to our *Safe Meta-BO* (SaMBO)  
 229 algorithm with GOOSE, as SAMBO-G and, when instantiated with SAFEBO as SAMBO-S.

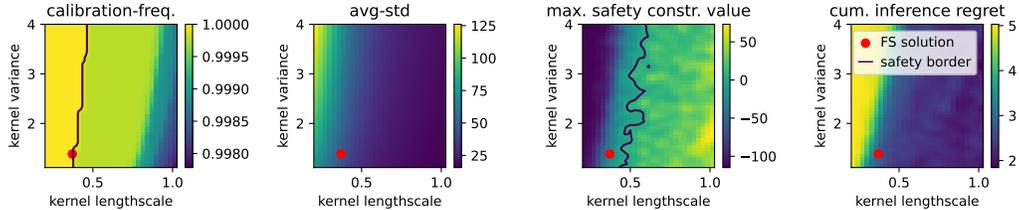


Figure 2: Left: calib and avg-std across a grid of kernel lengthscales  $l_q$  and variances  $\nu_q$  for the constraint model. Right: max safety constraint value and cumulative regret for 200 iterations of safe BO. Calibration and sharpness are good proxies for safety and efficiency of the downstream safe BO.

## 230 6 Experiments

231 First, we investigate if the calibration-/sharpness-based frontier search chooses good and safe kernel  
232 parameters. Second, we evaluate the the full SAMBO algorithm in a range of safe BO environments.

### 233 6.1 Experiment Setup

234 In the following, we describe our experiment setup and methodology. See Appx. F for more details.

235 **Synthetic safe BO benchmark environments.** We consider two synthetic environments, based on  
236 popular benchmark functions with a two-dimensional domain. The first is based on the *Camelback*  
237 function [54] overlaid with products of sinusoids of varying amplitude, phase and shifts. The  
238 constraint function is similar, with an additional random quadratic component to ensure connection  
239 between most of the safe regions of the domain. The second environment is based on the challenging  
240 Eggholder function [55] with many local minima. The function’s shape can be varied by random  
241 sampling of three parameters. The constraint is a quadratic function, overlaid with sinusoids of  
242 varying frequencies. A task corresponds to a pair of randomly drawn target and constraint functions.

243 **Controller Tuning for a high-precision linear robot.** As robotic case study, we tune the  
244 controller of a linear axis in an Argus linear motion system by Schneeberger Linear Technology,  
245 a high-precision/speed robot for wafer inspection. To goal is to tune the three gain parameters  
246 of a cascaded PI controller to achieve minimal position error. We minimize the total variation  
247 (TV) of the position error signal, while constraining its maximum frequency in the FFT which  
248 measures (potentially damaging) instabilities/vibrations in the system. Different tasks correspond  
249 to different step sizes, ranging from  $10\mu m$  to  $10mm$ . At different scales, the robot behaves differently  
250 in response to the controller parameters, resulting in different target and constraint functions.  
251 The experiments are conducted with a simulation of the robot, so we can explore unsafe kernel  
252 hyper-parameters – a key element for finding and visualizing the safety boundary in Figure 2.

253 **Meta-training data.** We generate the benchmark meta-training by running SAFEBOPT with conserva-  
254 tive kernel hyper-parameters choices on each task. The resulting datasets are non-i.i.d. and only  
255 consist of observations where  $q(\mathbf{z}) < 0$ , much more realistic than sampling data uniformly and i.i.d..  
256 With this, we aim to mimic a practical scenario where we have performed various safe BO on related  
257 tasks in the past and now want to harness the corresponding data. For the synthetic environments,  
258 we use  $n = 40$  tasks with  $T_i = 100$  samples in case of the Camelback-Sin and  $T_i = 200$  samples for  
259 the Random Eggholder environment. For the Argus controller tuning, we use  $n = 20$  and  $T_i = 400$ .

260 **Evaluation metrics.** To evaluate the performance of various methods, we run safe BO with them on  
261 at least 4 unseen test tasks with each 5 seeds. To measure a method’s query efficiency, we report the  
262 (safe) inference regret  $r_t = f(\hat{\mathbf{x}}_t^*) - f(\mathbf{x}^*)$  where  $\hat{\mathbf{x}}_t^* = \arg \min_{\mathbf{z} \in \mathcal{S}_t} \mu_t(\mathbf{x})$  is a method’s best current  
263 guess for the safe minimizer of (1). In Fig. 2, we also report the cumulative regret  $R_T := \sum_{t=1}^T r_t$ .

### 264 6.2 Choosing the kernel parameters via calibration & sharpness based frontier search

265 We investigate how well the calibration and sharpness based frontier search (FS) for finding kernel  
266 hyper-parameters works in practice. For that, we consider the Argus controller tuning problem.  
267 Generally, we perform FS in the log-space of  $l$  and  $\nu$  since we found this to work better in practice.

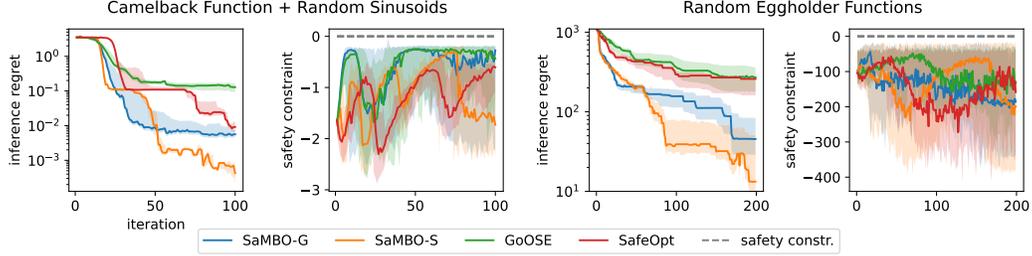


Figure 3: Safe BO regret and safety constraint on the synthetic benchmarks. SAMBO converges significantly faster towards the global optimum than the baselines, without violating the constraints.

268 Fig. 1 displays the frontier search at various iterations. FS quickly shrinks the set of possible safe  
 269 optima (area between upper frontier (green) and lower frontier (red)) to points close to the safety  
 270 border. After only 20 iterations, it already found a solution that is very close to the safe optimum.  
 271 This showcases the efficiency we gain thanks to taking into account the monotonicity of the problem.

272 Furthermore, we investigate how well calibration and sharpness reflect what we care about: 1) No  
 273 safety constraint violations and 2) query efficiency. We compute the maximum constraint value  
 274 across 200 iterations with GOOSE as well as the cumulative regret across a grid of kernel length-  
 275 scales and variances. Fig. 2 holds the results, together with the calib and avg-std. Overall, calibration  
 276 and sharpness of the GPs’ uncertainty estimates are a good proxy for the downstream safety during  
 277 BO and, respectively, the regret. Importantly, all parameters that fulfill the calibration constraint  
 278  $\text{calib}(\{\mathcal{D}_{i,T_i}\}_{i=1}^n, \omega) \geq 1$  lead to safe BO runs without constraint violations. Finally, the kernel pa-  
 279 rameters, chosen by FS, are both safe and lead to a small cumulative regret. Overall, this empirically  
 280 supports the validity of our data-driven approach for choosing good, but safe, kernel parameters.

### 281 6.3 Safe Bayesian Optimization Benchmark & Controller Tuning

282 We compare the two SAMBO instantiations, SAMBO-S and SAMBO-G, with their corre-  
 283 sponding safe BO baseline methods SAFEOPT and GOOSE. For the baselines, we use the GP  
 284 kernel parameters found by FS. Fig. 3 displays the results for the synthetic benchmark func-  
 285 tions and Fig. 4 for the Argus controller tuning. Note that, in case of the regret, the shaded areas  
 286 correspond to confidence intervals while for the safety constraint values they correspond to the  
 287 entire range of values (i.e., max - min). Overall, SAMBO converges to near optimal solutions much  
 288 faster than the baselines without meta-learning. Across all environments, there are no safety  
 289 violations which demonstrates that 1) the kernel parameters by FS are safe and 2) SAMBO  
 290 maintains safety thanks to principled regularization in the function space during meta-learning.  
 291 The improved query efficiency without compromising safety in the controller tuning setting,  
 292 where the constraint (maximum frequency in the signal) is highly non-smooth, demonstrates  
 293 the applicability of SAMBO to challenging real-world robotics problems.

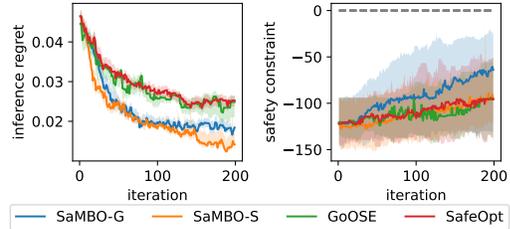


Figure 4: Safe controller tuning for the Argus robot. SAMBO safely finds good controller parameters faster than safe BO baselines.

## 299 7 Summary and Discussion of Limitations

300 We have introduced a data-driven framework for choosing kernel parameters or even meta-learning  
 301 GP priors that are both informative and reliable. When combined with a safe BO algorithm, the  
 302 resulting SAMBO speeds up the optimization of, e.g., controller parameters, without compromising  
 303 safety. Except for the observation noise variance, which is often known or easy to estimate, our  
 304 framework makes safe BO free of hyper-parameters and, thus, more robust and practical. However,  
 305 it relies on the availability of offline data that is both sufficient in quantity and representative of the  
 306 target task. As our approach relies on empirical estimates of the calibration, it may fail to ensure  
 307 safety when given too little data or tasks that are systematically different to the target task.

## References

- [1] R. Garnett, M. A. Osborne, and S. J. Roberts. Bayesian optimization for sensor set selection. In *Proceedings of the 9th ACM/IEEE international conference on information processing in sensor networks*, pages 209–219, 2010.
- [2] E. Cisbani, A. Del Dotto, C. Fanelli, M. Williams, M. Alfred, F. Barbosa, L. Barion, V. Berdnikov, W. Brooks, T. Cao, et al. Ai-optimized detector design for the future electron-ion collider: the dual-radiator rich case. *Journal of Instrumentation*, 15(05):P05009, 2020.
- [3] A. Marco, P. Hennig, J. Bohg, S. Schaal, and S. Trimpe. Automatic lqr tuning based on gaussian process global optimization. In *2016 IEEE international conference on robotics and automation (ICRA)*, pages 270–277. IEEE, 2016.
- [4] M. Neumann-Brosig, A. Marco, D. Schwarzmann, and S. Trimpe. Data-efficient autotuning with bayesian optimization: An industrial control study. *IEEE Transactions on Control Systems Technology*, 28(3):730–740, 2019.
- [5] M. Rowold, A. Wischnewski, and B. Lohmann. Constrained bayesian optimization of a linear feed-forward controller. *IFAC-PapersOnLine*, 52(29):1–6, 2019. ISSN 2405-8963. 13th IFAC Workshop on Adaptive and Learning Control Systems ALCOS 2019.
- [6] M. Khosravi, V. Behrunani, R. S. Smith, A. Rupenyan, and J. Lygeros. Cascade control: Data-driven tuning approach based on bayesian optimization. *IFAC-PapersOnLine*, 53(2):382–387, 2020.
- [7] A. Rupenyan, M. Khosravi, and J. Lygeros. Performance-based trajectory optimization for path following control using bayesian optimization. In *60th IEEE conference on Decision and Control (CDC 2021)*, 2021.
- [8] Y. Sui, A. Gotovos, J. W. Burdick, and A. Krause. Safe exploration for optimization with Gaussian processes. In *ICML*, 2015.
- [9] F. Berkenkamp, A. P. Schoellig, and A. Krause. Safe controller optimization for quadrotors with gaussian processes. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 491–496, 2016.
- [10] M. Turchetta, F. Berkenkamp, and A. Krause. Safe exploration for interactive machine learning. In *Advances in Neural Information Processing Systems*, 2019.
- [11] J. Rothfuss, D. Heyn, J. Chen, and A. Krause. Meta-learning reliable priors in the function space. *Advances in Neural Information Processing Systems*, 34, 2021.
- [12] J. M. Hernández-Lobato, M. A. Gelbart, M. W. Hoffman, R. Adams, and Z. Ghahramani. Predictive entropy search for Bayesian optimization with unknown constraints. In *ICML*, 2015.
- [13] M. Khosravi, C. Koenig, M. Maier, R. S. Smith, J. Lygeros, and A. Rupenyan. Safety-aware cascade controller tuning using constrained bayesian optimization. *IEEE Transactions on Industrial Electronics*, 2022.
- [14] M. Fiducioso, S. Curi, B. Schumacher, M. Gwerder, and A. Krause. Safe contextual bayesian optimization for sustainable room temperature pid control tuning. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*, pages 5850–5856. AAAI Press, December 2019.
- [15] Y. Sui, vincent Zhuang, J. Burdick, and Y. Yue. Stagewise safe Bayesian optimization with Gaussian processes. In *ICML*, volume 80 of *Proceedings of Machine Learning Research*, pages 4781–4789. PMLR, 10–15 Jul 2018.

- 351 [16] F. Berkenkamp, A. Krause, and A. P. Schoellig. Bayesian optimization with safety constraints:  
352 safe and automatic parameter tuning in robotics. *Machine Learning*, pages 1–35, 2021.
- 353 [17] M. Khosravi, A. Eichler, N. A. Schmid, R. S. Smith, and P. Heer. Controller tuning by bayesian  
354 optimization an application to a heat pump. *2019 18th European Control Conference (ECC)*,  
355 pages 1467–1472, 2019.
- 356 [18] C. König, M. Turchetta, J. Lygeros, A. Rupenyan, and A. Krause. Safe and efficient model-  
357 free adaptive control via bayesian optimization. In *2021 IEEE International Conference on*  
358 *Robotics and Automation (ICRA)*, pages 9782–9788. IEEE, 2021.
- 359 [19] S. Hochreiter, A. S. Younger, and P. R. Conwell. Learning To Learn Using Gradient Descent.  
360 In *International Conference on Artificial Neural Networks*, 2001.
- 361 [20] M. Andrychowicz, M. Denil, S. G. Colmenarejo, M. W. Hoffman, D. Pfau, T. Schaul,  
362 B. Shillingford, and N. De Freitas. Learning to learn by gradient descent by gradient descent.  
363 *arXiv*, 2016.
- 364 [21] Y. Chen, M. W. Hoffman, S. G. Colmenarejo, M. Denil, T. P. Lillicrap, M. Botvinick, and  
365 N. De Freitas. Learning to Learn without Gradient Descent by Gradient Descent. In *Interna-*  
366 *tional Conference on Machine Learning*, 2017.
- 367 [22] J. Baxter. A model of inductive bias learning. *Journal of Artificial Intelligence Research*, 2000.
- 368 [23] J. Snell, K. Swersky, and R. Zemel. Prototypical networks for few-shot learning. In *Advances*  
369 *in Neural Information Processing Systems*, 2017.
- 370 [24] O. Vinyals, C. Blundell, T. Lillicrap, D. Wierstra, et al. Matching networks for one shot  
371 learning. In *Advances in Neural Information Processing Systems*, 2016.
- 372 [25] J. Harrison, A. Sharma, and M. Pavone. Meta-learning priors for efficient online bayesian  
373 regression. In *International Workshop on the Algorithmic Foundations of Robotics*, pages  
374 318–337. Springer, 2018.
- 375 [26] C. Finn, P. Abbeel, and S. Levine. Model-agnostic meta-learning for fast adaptation of deep  
376 networks. In *International Conference on Machine Learning*, 2017.
- 377 [27] J. Rothfuss, D. Lee, I. Clavera, T. Asfour, and P. Abbeel. ProMP: Proximal Meta-Policy Search.  
378 In *International Conference on Learning Representations*, 2019.
- 379 [28] A. Nichol, J. Achiam, and J. Schulman. On First-Order Meta-Learning Algorithms. *arXiv*,  
380 2018.
- 381 [29] T. Kim, J. Yoon, O. Dia, S. Kim, Y. Bengio, and S. Ahn. Bayesian model-agnostic meta-  
382 learning. In *Advances in Neural Information Processing Systems*, 2018.
- 383 [30] A. Pentina and C. Lampert. A PAC-Bayesian bound for lifelong learning. In *International*  
384 *Conference on Machine Learning*, 2014.
- 385 [31] R. Amit and R. Meir. Meta-learning by adjusting priors based on extended PAC-Bayes theory.  
386 In *International Conference on Machine Learning*, 2018.
- 387 [32] Y. Balaji, S. Sankaranarayanan, and R. Chellappa. Metareg: Towards domain generalization  
388 using meta-regularization. In *Advances in Neural Information Processing Systems*, volume 31,  
389 2018.
- 390 [33] M. Yin, G. Tucker, M. Zhou, S. Levine, and C. Finn. Meta-learning without memorization. In  
391 *International Conference on Learning Representations*, 2020.
- 392 [34] J. Rothfuss, V. Fortuin, M. Josifoski, and A. Krause. PACOH: Bayes-optimal meta-learning  
393 with PAC-guarantees. In *International Conference for Machine Learning (ICML)*, 2021.

- 394 [35] A. Farid and A. Majumdar. Generalization bounds for meta-learning via pac-bayes and uniform  
395 stability. In *Advances in Neural Information Processing Systems*, volume 34, pages 2173–  
396 2186, 2021.
- 397 [36] L. Cella, K. Lounici, and M. Pontil. Meta representation learning with contextual linear band-  
398 dits. *arXiv preprint arXiv:2205.15100*, 2022.
- 399 [37] P. Kassraie, J. Rothfuss, and A. Krause. Meta-learning hypothesis spaces for sequential  
400 decision-making. In *International Conference on Machine Learning*, 2022.
- 401 [38] B. Shahriari, K. Swersky, Z. Wang, R. P. Adams, and N. de Freitas. Taking the Human Out  
402 of the Loop: A Review of Bayesian Optimization. *Proceedings of the IEEE*, 104(1):148–175,  
403 2016. doi:10.1109/JPROC.2015.2494218.
- 404 [39] P. I. Frazier. A Tutorial on Bayesian Optimization. *arXiv preprint arXiv:1807.02811*, 2018.
- 405 [40] C. E. Rasmussen and C. K. I. Williams. *Gaussian processes in machine learning*. MIT Press,  
406 2006.
- 407 [41] P. Auer, N. Cesa-Bianchi, and P. Fischer. Finite-time analysis of the multiarmed bandit prob-  
408 lem. *Machine learning*, 47(2):235–256, 2002.
- 409 [42] N. Srinivas, A. Krause, S. Kakade, and M. Seeger. Gaussian Process Optimization in the  
410 Bandit Setting: No Regret and Experimental Design. In *International Conference on Machine  
411 Learning*, pages 1015–1022, 07 2010.
- 412 [43] W. R. Thompson. On the likelihood that one unknown probability exceeds another in view  
413 of the evidence of two samples. *Biometrika*, 25(3/4):285–294, 1933. ISSN 00063444. URL  
414 <http://www.jstor.org/stable/2332286>.
- 415 [44] Z. Wang and S. Jegelka. Max-value entropy search for efficient bayesian optimization. In  
416 *International Conference on Machine Learning*, pages 3627–3635. PMLR, 2017.
- 417 [45] J. Schmidhuber. *Evolutionary principles in self-referential learning. On learning how to learn:  
418 The meta-meta-... hook*. PhD thesis, Technische Universitaet Munchen, 1987.
- 419 [46] S. Thrun and L. Pratt, editors. *Learning to Learn*. Kluwer Academic Publishers, 1998.
- 420 [47] V. Perrone, R. Jenatton, M. W. Seeger, and C. Archambeau. Scalable Hyperparameter Trans-  
421 fer Learning. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and  
422 R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31, 2018.
- 423 [48] V. Fortuin and G. Rätsch. Deep mean functions for meta-learning in gaussian processes. *arXiv  
424 preprint arXiv:1901.08098*, 2019.
- 425 [49] A. G. Wilson, Z. Hu, R. Salakhutdinov, and E. P. Xing. Deep kernel learning. In *International  
426 Conference on Artificial Intelligence and Statistics*, pages 370–378. PMLR, 2016.
- 427 [50] T. Gneiting, F. Balabdaoui, and A. E. Raftery. Probabilistic forecasts, calibration and sharp-  
428 ness. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 69(2):243–  
429 268, 2007.
- 430 [51] T. Gneiting and A. E. Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal  
431 of the American Statistical Association*, 102(477):359–378, 2007.
- 432 [52] V. Kuleshov, N. Fenner, and S. Ermon. Accurate Uncertainties for Deep Learning Using Cali-  
433 brated Regression. In *International Conference on Machine Learning*, 2018.
- 434 [53] M.-O. Pohle. The murphy decomposition and the calibration-resolution principle: A new  
435 perspective on forecast evaluation. *arXiv preprint arXiv:2005.01835*, 2020.

- 436 [54] M. Molga and C. Smutnicki. Test functions for optimization needs. *Test functions for opti-*  
437 *mization needs*, 101:48, 2005.
- 438 [55] M. Jamil and X.-S. Yang. A literature survey of benchmark functions for global optimisation  
439 problems. *International Journal of Mathematical Modelling and Numerical Optimisation*, 4  
440 (2):150–194, 2013.
- 441 [56] P. Moritz, R. Nishihara, S. Wang, A. Tumanov, R. Liaw, E. Liang, M. Elibol, Z. Yang, W. Paul,  
442 M. I. Jordan, and I. Stoica. Ray: A distributed framework for emerging ai applications. In *Pro-*  
443 *ceedings of the 13th USENIX Conference on Operating Systems Design and Implementation*,  
444 OSDI’18, page 561–577, 2018.

## 445 A Safe Bayesian Optimization Algorithms

### 446 A.1 SafeOpt

447 SAFEOPT [9, 16] make use of the GP confidence intervals, by defining a safe set on a discretized  
448 domain  $\mathcal{X}$ :

$$\mathcal{S}_t(\alpha) = \{\mathbf{x} \in \mathcal{X} \mid \mu_t^q(\mathbf{x}) + \beta(\alpha)\sigma_t^q(\mathbf{x}) < 0\}. \quad (11)$$

449 where we write  $\mu_t^q(\mathbf{x})$  and  $\sigma_t^q(\mathbf{x})$  short for  $\mu^q(\mathbf{x}|\mathcal{D}_t)$  and  $\sigma^q(\mathbf{x}|\mathcal{D}_t)$ . Based on the safe set, SAFEOPT  
450 constructs two additional sets. First, the set of potential optimizers

$$\mathcal{M}_t(\alpha) = \{\mathbf{x} \in \mathcal{S}_t(\alpha) \mid \mu_t^f(\mathbf{x}) - \beta(\alpha)\sigma_t^f(\mathbf{x}) < \mu_t^f(\mathbf{x}^\dagger) + \beta(\alpha)\sigma_t^f(\mathbf{x}^\dagger)\}, \quad (12)$$

451 where  $\mathbf{x}^\dagger$  is the best observed input so far. Second, the set of possible expanders

$$\mathcal{G}_t(\alpha) = \{\mathbf{x} \in \mathcal{S}_t(\alpha) \mid g(\mathbf{x}) > 0\}, \quad (13)$$

452 where  $g_t(\mathbf{x})$  is defined as the additional number of inputs that become safe if we would query  $\mathbf{x}$  and  
453 observe an hypothetical optimistic constraint value  $\tilde{q} = \mu_t^q(\mathbf{x}') - \beta(\alpha)\sigma_t^q(\mathbf{x}')$ .

$$g_t(\mathbf{x}) = |\{\mathbf{x} \in \mathcal{X} \setminus \mathcal{S}_t(\alpha) \mid \mu^q(\mathbf{x}'|\mathcal{D}_t \cup (\mathbf{x}, \tilde{q})) + \beta(\alpha)\sigma^q(\mathbf{x}'|\mathcal{D}_t \cup (\mathbf{x}, \tilde{q})) < 0\}|. \quad (14)$$

454 In each iteration of the optimization the sets are updated, w.r.t. the posterior belief of the underlying  
455 GPs a query candidate is selected from each the set of optimizers and the set of expanders. From the  
456 set of optimizers the GP-LCB sample

$$\mathbf{x}_{opt}^* = \arg \min_{\mathbf{z} \in \mathcal{M}_t(\alpha)} \mu_t^f(\mathbf{z}) - \beta(\alpha)\sigma_t^f(\mathbf{z}) \quad (15)$$

457 is selected. From the set of possible expanders

$$\mathbf{x}_{exp}^* = \arg \max_{\mathbf{x} \in \mathcal{G}_t(\alpha)} g_n(x) \quad (16)$$

458 is selected. Finally, SAFEOPT selects

$$\mathbf{x}^* = \arg \max_{\mathbf{x} \in \{\mathbf{x}_{opt}^*, \mathbf{x}_{exp}^*\}} \max\{\sigma_t^f(\mathbf{x}), \sigma_t^q(\mathbf{x})\} \quad (17)$$

459 as the next query.

### 460 A.2 GoOSE

461 In SAFEOPT the expansion of the safe set is traded off against the optimization by the uncertainty  
462 of the prediction. This can lead to part-wise exploration of the safe set without consideration of  
463 the objective function. To cope with this [10] proposes a goal oriented safe exploration (GoOSE)  
464 algorithm. The idea is that the safe set is expanded only if it is also beneficial to the optimizaton  
465 of the objective. Like SAFEOPT it builds a GP model of the objective and constraints from noisy  
466 evaluations based on GP regression. The models of the constraints are used to construct two sets:  
467 First, the pessimistic safe set

$$\mathcal{S}_t^p(\alpha) = \{\mathbf{x} \in \mathcal{X} \mid \mu_t^q(\mathbf{x}) + \beta^q(\alpha)\sigma_t^q(\mathbf{x}) < 0\}, \quad (18)$$

468 which only contains domain points that are with high probability fulfilling the constraint. Sec-  
469 ond, we construct an optimistic safe set which includes all points we can optimistically expand the  
470 pessimistic safe set to. Rather than the version of Turchetta et al. [10] which requires an a-priori  
471 known Lipschitz constant, we use the version of König et al. [18] which only uses quantities that  
472 are provided by our GP model of the constraint function. For instance, the Lipschitz constant is ap-  
473 proximated by the sup-norm  $\|\nabla_{\mathbf{x}}\mu_t(\mathbf{x})\|_\infty$  of the gradient of the GP's posterior mean. This nicely  
474 accommodates non-stationary kernels which might arise during meta-learning. We first make  
475 sure that we don't query the same point over and over by excluding all points within  $\mathcal{S}_t^p(\alpha)$  whose  
476 confidence intervals for  $q$  standard deviation become narrower than a alleatoric noise threshold  $\epsilon$ .  
477 The remaining points are:

$$W_t = \{x \in \mathcal{S}_t^p(\alpha) : 2\beta^q(\alpha)\sigma_t^q(x) > \epsilon\} \quad (19)$$

478 Then, we check for any  $\mathbf{x} \in \mathcal{X}$  whether we still fulfill the safety constraint when we add the distance  
479 times the Lipschitz constant estimate to the optimistic lower bound  $\mu_t^q(\mathbf{x}) - \beta^q(\alpha)\sigma_t^q(\mathbf{x})$  of the  
480 constraint CI, formally:

$$g_t(\mathbf{x}, \mathbf{z}) = \mathbb{I}[\mu_t^q(\mathbf{x}) - \beta^q(\alpha)\sigma_t^q(\mathbf{x}) + \|\nabla_{\mathbf{x}}\mu_t(\mathbf{x})\|_{\infty}\|\mathbf{x} - \mathbf{z}\|_2 < 0] . \quad (20)$$

481 Based on this indicator function, we can construct the optimistic safe set:

$$\mathcal{S}_t^o(\alpha) = \{\mathbf{x} \in \mathcal{X} \mid \exists \mathbf{z} \in W_t : g_t(\mathbf{x}, \mathbf{z}) = 1\} \quad (21)$$

482 Within the optimistic safe set  $\mathcal{S}_t^o(\alpha)$ , we now find the optimizer of a (standard) BO acquisition  
483 function. We use the the UCB acquisition function [42], i.e.  $\text{aq}(\mathbf{x}) := \mu_t^f(\mathbf{x}) - \beta^f(\alpha)\sigma_t^f(\mathbf{x})$ , to find  
484 the next query candidate  $\tilde{\mathbf{x}}^* = \arg \min_{\mathbf{x} \in \mathcal{S}_t^o(\alpha)} \text{aq}(\mathbf{x})$ .

485 If  $\tilde{\mathbf{x}}^*$  is inside  $\mathcal{S}_t^p(\alpha)$ , it is evaluated. If not we query the point in  $W_t$  which is closest to  $\tilde{\mathbf{x}}^*$  and  
486 fulfills the expander criterion in (20). After querying a point and observing the corresponding  $\tilde{f}_t$  and  
487  $\tilde{q}_t$ , the posteriors of the GPs are updated and thus the sets which we defined above. This is repeated  
488 until  $\tilde{\mathbf{x}}^*$  is inside  $\mathcal{S}_t^p(\alpha)$  and can be evaluated or  $\tilde{\mathbf{x}}^*$  is no longer in  $\mathcal{S}_t^p(\alpha)$  and we compute a new  
489 query candidate.

490 The described procedure is summarized in Algorithm 3. Note that in comparison to Turchetta et al.  
491 [10], König et al. [18], Algorithm 3 not exclude points from  $W_t$  that lie at the periphery of the  
492 domain because checking whether a point lies at the periphery is hard when working with uniformly  
493 sampled domain points instead of a grid.

---

### Algorithm 3 GoOSE algorithm

---

**Input:** Initial safe set  $\mathcal{S}_0$

**Input:** GP models  $f \sim \mathcal{GP}(\mu^f, k^f)$ ,  $q \sim \mathcal{GP}(\mu^q, k^q)$

```

1: for  $t = 1, \dots, T$  do
2:    $\mathcal{S}_t^p(\alpha) \leftarrow \{\mathbf{x} \in \mathcal{X} \mid \mu_t^q(\mathbf{x}) + \beta(\alpha)\sigma_t^q(\mathbf{x}) < 0\}$  // pessimistic safe set
3:    $W_t \leftarrow \{x \in \mathcal{S}_t^p(\alpha) \mid 2\beta(\alpha)\sigma_t^q(x) > \epsilon\}$  // expanders
4:    $\mathcal{S}_t^o(\alpha) \leftarrow \{\mathbf{x} \in \mathcal{X} \mid \exists \mathbf{z} \in W_t : g_t(\mathbf{x}, \mathbf{z}) = 1\}$  // optimistic safe set
5:    $\tilde{\mathbf{x}}^* \leftarrow \arg \min_{\mathbf{x} \in \mathcal{S}_t^o(\alpha)} \text{aq}(\mathbf{x})$  // UCB candidate within optimistic safe set
6:   if  $\tilde{\mathbf{x}}^* \in \mathcal{S}_t^p(\alpha)$  then
7:     evaluate  $f(\tilde{\mathbf{x}}^*)$ ,  $q(\tilde{\mathbf{x}}^*)$ ,
8:   else
9:      $\tilde{\mathbf{x}}_w \leftarrow \arg \min_{\mathbf{x} \in W_t} \|\mathbf{x} - \tilde{\mathbf{x}}^*\|_2$  s.t.  $g_t(\mathbf{x}, \tilde{\mathbf{x}}^*) = 1$  // expand  $\mathcal{S}$  in direction of  $\tilde{\mathbf{x}}^*$ 
10:    evaluate  $f(\tilde{\mathbf{x}}_w)$ ,  $q(\tilde{\mathbf{x}}_w)$ ,

```

**Return:**  $\hat{\mathbf{x}}^* \leftarrow \arg \min_{\mathbf{x} \in \mathcal{S}_t^p(\alpha)} \mu_t^f(\mathbf{x})$  // return safe point with best posterior mean

---

## 494 B Meta-Learning reliable priors with F-PACOH

495 The F-PACOH method of Rothfuss et al. [11] uses a set of datasets  $\mathcal{D}_{1,T_1}, \dots, \mathcal{D}_{n,T_n}$  to meta-  
496 learn a GP prior. For that, it requires a parametric family  $\{\rho_{\theta} \mid \theta \in \Theta\}$  of GP priors  $\rho_{\theta}(h) =$   
497  $\text{GP}(h(\mathbf{x}) \mid m_{\theta}(\mathbf{x}), k_{\theta}(\mathbf{x}, \mathbf{x}'))$ . Typically the mean and kernel function of the GP are parameterized  
498 by neural networks. In addition, it presumes a Vanilla GP  $\rho(h) = \text{GP}(h(\mathbf{x}) \mid 0, k(\mathbf{x}, \mathbf{x}'))$  as stochas-  
499 tic process hyper-prior. In our case, the hyper-prior GP has a zero-mean and a SE kernel. During  
500 the meta-training the marginal log-likelihood  $\ln Z(\mathcal{D}_{i,T_i}, \rho_{\theta}) = \ln p(\mathbf{y}_i^{\mathcal{D}} \mid \mathbf{X}_i^{\mathcal{D}}, \theta)$  is used to fit the  
501  $\rho_{\theta}(h)$  to the meta-training data. Here we write  $\mathcal{D}_{i,T_i} = (\mathbf{X}_i^{\mathcal{D}}, \mathbf{y}_i^{\mathcal{D}})$  for the matrix of function inputs  
502 and vector of targets of the respective dataset. Since we use GPs, the marginal log-likelihood can be  
503 computed in closed form as

$$\ln p(\mathbf{y}^{\mathcal{D}} \mid \mathbf{X}^{\mathcal{D}}, \theta) = -\frac{1}{2} (\mathbf{y}^{\mathcal{D}} - \mathbf{m}_{\mathbf{X}^{\mathcal{D}}, \theta})^{\top} \tilde{\mathbf{K}}_{\mathbf{X}^{\mathcal{D}}, \theta}^{-1} (\mathbf{y}^{\mathcal{D}} - \mathbf{m}_{\mathbf{X}^{\mathcal{D}}, \theta}) - \frac{1}{2} \ln |\tilde{\mathbf{K}}_{\mathbf{X}^{\mathcal{D}}, \theta}| - \frac{T}{2} \ln 2\pi \quad (22)$$

504 where  $\tilde{\mathbf{K}}_{\mathbf{X}^{\mathcal{D}}, \theta} = \mathbf{K}_{\mathbf{X}^{\mathcal{D}}, \theta} + \sigma^2 I$ , with kernel matrix  $\mathbf{K}_{\mathbf{X}^{\mathcal{D}}, \theta} = [k_{\theta}(\mathbf{x}_l, \mathbf{x}_k)]_{l,k=1}^{T_i}$ , likelihood variance  
505  $\sigma^2$ , and mean vector  $\mathbf{m}_{\mathbf{X}^{\mathcal{D}}, \theta} = [m_{\theta}(\mathbf{x}_1), \dots, m_{\theta}(\mathbf{x}_{T_i})]^{\top}$ .

---

**Algorithm 4** F-PACOH [11]

---

**Input:** Datasets  $\mathcal{D}_{1,T_1}, \dots, \mathcal{D}_{n,T_n}$ , parametric family  $\{\rho_{\theta} | \theta \in \Theta\}$  of priors, learning rate  $\alpha$

**Input:** Stochastic process hyper-prior with marginals  $\rho(\cdot)$

- 1: Initialize the parameters  $\theta$  of prior  $\rho_{\theta}$
  - 2: **while** not converged **do**
  - 3:   **for**  $i = 1, \dots, n$  **do**
  - 4:      $\mathbf{X}_i = [\mathbf{X}_{i,s}^{\mathcal{D}}, \mathbf{X}_i^{\mathcal{M}}]$ , where  $\mathbf{X}_{i,s}^{\mathcal{D}} \subseteq \mathbf{X}_i^{\mathcal{D}}, \mathbf{X}_i^{\mathcal{M}} \stackrel{iid}{\sim} \mathcal{U}(\mathcal{X})$  // Sample measurement set
  - 5:     Estimate or compute  $\nabla_{\theta} \ln Z(\mathbf{X}_i^{\mathcal{D}}, \rho_{\theta})$  and  $\nabla_{\theta} KL[\rho_{\theta}(\mathbf{h}^{\mathbf{X}_i}) || \rho(\mathbf{h}^{\mathbf{X}_i})]$
  - 6:      $\nabla_{\theta} J_{F,i} = -\frac{1}{T_i} \nabla_{\theta} \ln Z(\mathcal{D}_{i,T_i}, \rho_{\theta}) + \left(\frac{1}{\sqrt{n}} + \frac{1}{nT_i}\right) \nabla_{\theta} KL[\rho_{\theta}(\mathbf{h}^{\mathbf{X}_i}) || \rho(\mathbf{h}^{\mathbf{X}_i})]$
  - 7:      $\theta \leftarrow \theta - \alpha \frac{1}{n} \sum_{i=1}^n \nabla_{\theta} J_{F,i}$  // Update prior parameter
- 

506 At the same time, during meta-training,  $\rho_{\theta}(h)$  is regularized towards the hyper-prior  $\rho(h)$ . We  
507 can only tractably assess the stochastic processes  $\rho_{\theta}(h)$  and  $\rho(h)$  in a finite set of (measurement)  
508 points  $\mathbf{X} := [\mathbf{x}_1, \dots, \mathbf{x}_k] \in \mathcal{X}^k, k \in \mathbb{N}$  through their finite marginal distributions of function values  
509  $\rho(\mathbf{h}^{\mathbf{X}}) := \rho(h(\mathbf{x}_1), \dots, h(\mathbf{x}_k))$  and  $\rho_{\theta}(\mathbf{h}^{\mathbf{X}}) := \rho_{\theta}(h(\mathbf{x}_1), \dots, h(\mathbf{x}_k))$  respectively. In particular, for  
510 each task, we construct measurement sets  $\mathbf{X}_i = [\mathbf{X}_{i,s}^{\mathcal{D}}, \mathbf{X}_i^{\mathcal{M}}]$  by selecting a random subset  $\mathbf{X}_{i,s}^{\mathcal{D}}$   
511 of the meta-training inputs  $\mathbf{X}_i^{\mathcal{D}}$  as well as random points  $\mathbf{X}_i^{\mathcal{M}} \stackrel{iid}{\sim} \mathcal{U}(\mathcal{X})$  sampled independently  
512 and uniformly from the bounded domain  $\mathcal{X}$ .<sup>2</sup> In each iteration, we compute the KL-divergence be-  
513 tween the marginal distributions of the stochastic processes in the sampled measurement sets. Since  
514 both stochastic processes are GPs, their finite marginals are Gaussians  $\rho(\mathbf{h}_i^{\mathbf{X}}) = \mathcal{N}(\mathbf{0}, \mathbf{K}_{\mathbf{X}_i})$  and  
515  $\rho_{\theta}(\mathbf{h}_i^{\mathbf{X}}) = \mathcal{N}(\mathbf{m}_{\mathbf{X}_i, \theta}, \mathbf{K}_{\mathbf{X}_i, \theta})$  and thus the KL-divergence available in close form. In expectation,  
516 over many iteration, we effectively minimize  $\mathbb{E}_{\mathbf{X}_i} [KL [q(\mathbf{h}^{\mathbf{X}_i}) || \rho(\mathbf{h}^{\mathbf{X}_i})]]$ .

517 Overall, the loss with the functional KL regularizer which we minimize in F-PACOH reads as:

$$\mathcal{L}(\theta) = \frac{1}{n} \sum_{i=1}^n \left( \underbrace{-\frac{1}{T_i} \ln Z(\mathcal{D}_{i,T_i}, \rho_{\theta})}_{\text{marginal log-likelihood}} + \underbrace{\left(\frac{1}{\sqrt{n}} + \frac{1}{nT_i}\right) \mathbb{E}_{\mathbf{X}_i} [KL[\rho_{\theta}(\mathbf{h}^{\mathbf{X}}) || \rho(\mathbf{h}^{\mathbf{X}})]]}_{\text{functional KL-divergence}} \right). \quad (23)$$

518 The stochastic minimization procedure for (23) which we have described is summarized in Algo-  
519 rithm 4.

## 520 C Frontier Search

521 We aim to solve the constraint optimization problem

$$\min_{\omega} s(\mathbf{z}) \quad \text{s.t.} \quad c(\mathbf{z}) \geq 1 \quad (24)$$

522 where  $s : \mathbb{R}^2 \mapsto \mathbb{R}$  and  $c : \mathbb{R}^2 \mapsto \mathbb{R}$  are monotonically increasing functions. Formally, we define the  
523 monotonicity w.r.t. to the partial order on  $\mathbb{R}^2$ :

524 **Definition 1** (Monotone Function). *A function  $h(\mathbf{z}) : \mathbb{R}^2 \mapsto \mathbb{R}$  is said to be monotone if for all*  
525  *$(z_1, z_2), (z'_1, z'_2) \in \mathbb{R}^2$*

$$z_1 \leq z'_1 \wedge z_2 \leq z'_2 \Rightarrow h(z_1, z_2) \leq h(z'_1, z'_2). \quad (25)$$

526 For brevity we write  $\mathbf{z} \leq \mathbf{z}'$  short for  $z_1 \leq z'_1 \wedge z_2 \leq z'_2$  and  $\mathbf{z} \geq \mathbf{z}' \Leftrightarrow z_1 \geq z'_1 \wedge z_2 \geq z'_2$ .

527 For our algorithm, we assume knowledge of an upper bound  $\mathbf{z}^u$  and lower bound  $\mathbf{z}^l$  of the optimal so-  
528 lution  $\mathbf{z}^*$  such that the rectangle  $\mathcal{Z} = [z_1^l, z_1^u] \times [z_2^l, z_2^u]$  that is spanned by the bounds and contains  $\mathbf{z}^*$ .

529 **Assumption 1** (Valid Search Domain). *The search domain  $\mathcal{Z} = [z_1^l, z_1^u] \times [z_2^l, z_2^u]$  is valid if it*  
530 *contains the optimum  $\mathbf{z}^* = \arg \min_{\mathbf{z}: c(\mathbf{z}) \geq 1} s(\mathbf{z})$  of the constraint optimization problem.*

---

<sup>2</sup>In Section 3.3, we have portrayed the measurement sets as only the uniform domain points  $\mathbf{X}_i^{\mathcal{M}}$  for brevity of the exposition. However, our implementation uses  $\mathbf{X}_i = [\mathbf{X}_{i,s}^{\mathcal{D}}, \mathbf{X}_i^{\mathcal{M}}]$ , as described here in the appendix.

531 Due to the monotonicity of  $s(\mathbf{z})$  and  $c(\mathbf{z})$  we know that the optimal solution must lie on or imme-  
 532 diately above the constraint boundary. The latter case is rather a technical detail which is due to the  
 533 fact that we do not assume continuity of  $c$ . Formally we have:

534 **Lemma 2** ( $\mathbf{z}^*$  lies on or directly above the constraint boundary). *Let  $\mathbf{z}^* = \arg \min_{\mathbf{z}: c(\mathbf{z}) \geq 1} s(\mathbf{z})$  be*  
 535 *the unique minimizer, then there exists no  $z_2 \in [z_2^l, z_2^u]$  with  $1 \leq c(z_1^*, z_2) < c(\mathbf{z}^*)$ .*

536 *Proof.* The proof of Lemma 2 follows by reduction: Assume  $\exists z_2 \in [z_2^l, z_2^u]$  with  $1 \leq c(z_1^*, z_2) <$   
 537  $c(\mathbf{z}^*)$ . Since  $c$  is monotonically increasing  $c(z_1^*, z_2) < c(\mathbf{z}^*)$  implies that  $z_2 < z_2^*$ . Thus, by  
 538 the monotonicity of  $s$ , we have that  $s(z_1^*, z_2) \leq s(\mathbf{z}^*)$  and  $\mathbf{z}^*$  cannot be the unique constrained  
 539 minimizer.  $\square$

540 In each iteration  $k$  of Algorithm 1 we query a point  $\mathbf{z}_k^q \in \mathcal{Z}$  and observe the corresponding objective  
 541 and constraint values  $s(\mathbf{z}_k^q)$  and  $c(\mathbf{z}_k^q)$ . Due to Lemma 2 and the monotonicity of  $q$  we can rule out  
 542 an entire corner of the search domain. In particular, if  $c(\mathbf{z}_k^q) \geq 0$ , we can rule all points  $\mathbf{z}' > \mathbf{z}_k^q$  as  
 543 candidates for the optimal solution and, similarly, if  $c(\mathbf{z}_k^q) < 0$ , we can rule out all  $\mathbf{z}' \leq \mathbf{z}_k^q$ .

544 To keep track of all the areas of the search domain we can rule out, we construct an upper and a  
 545 lower frontier such that all ruled-out points lie either above the upper and below the lower frontier.  
 546 To construct these frontiers, we separate all queries based on whether they fulfill the constraint or  
 547 not. Initially, we set  $\mathcal{Q}^u = \{\mathbf{z}^u\}$  and  $\mathcal{Q}^l = \{\mathbf{z}^l\}$ , since Assumption 1 together with Lemma 2 imply  
 548 that  $c(\mathbf{z}^u) \geq 1$  and  $c(\mathbf{z}^l) \leq 1$ . Then, for every query, we add  $\mathbf{z}_k^q$  to  $\mathcal{Q}^u$  if  $c(\mathbf{z}^q) \geq 1$  and to  $\mathcal{Q}^l$   
 549 otherwise.

550 We define the upper and lower frontiers as maps  $z_1 \mapsto z_2$  and  $z_2 \mapsto z_1$ :

$$F_2^u(z_1; \mathcal{Q}^u) = \min\{z_2' \mid z_1 \geq z_1', \mathbf{z}' \in \mathcal{Q}^u\}, \quad F_1^u(z_2; \mathcal{Q}^u) = \min\{z_1' \mid z_2 \geq z_2', \mathbf{z}' \in \mathcal{Q}^u\} \quad (26)$$

$$F_2^l(z_1; \mathcal{Q}^l) = \max\{z_2' \mid z_1 \leq z_1', \mathbf{z}' \in \mathcal{Q}^l\}, \quad F_1^l(z_2; \mathcal{Q}^l) = \max\{z_1' \mid z_2 \leq z_2', \mathbf{z}' \in \mathcal{Q}^l\} \quad (27)$$

551 For convenience, we define the sets of points that lie on the frontiers as

$$\mathcal{F}^u(\mathcal{Q}^u) = \{\mathbf{z} \in \mathcal{Z} \mid F_1^u(z_2; \mathcal{Q}^u) = z_1 \vee F_2^u(z_1; \mathcal{Q}^u) = z_2\}, \quad (28)$$

$$\mathcal{F}^l(\mathcal{Q}^l) = \{\mathbf{z} \in \mathcal{Z} \mid F_1^l(z_2; \mathcal{Q}^l) = z_1 \vee F_2^l(z_1; \mathcal{Q}^l) = z_2\}. \quad (29)$$

552 If we assume Lipschitz continuity for  $s$ , we can bound how much our best solution is away from the  
 553 optimum, i.e.,  $s(\mathbf{z}^*)$ :

554 **Lemma 3** (Long version of Lemma 1). *Let  $\mathbf{z}^* = \arg \min_{\mathbf{z}: c(\mathbf{z}) \geq 1} s(\mathbf{z})$  be the solution of the con-*  
 555 *straint optimization problem where  $s : \mathbb{R}^2 \mapsto \mathbb{R}$  is monotone and  $L$  Lipschitz, and  $c : \mathbb{R}^2 \mapsto \mathbb{R}$  is*  
 556 *monotone constraint. Let*

$$\Gamma(\mathcal{Q}^l, \mathcal{Q}^u) = \{(z_1, z_2) \in \mathcal{Z} \mid F^l(z_1; \mathcal{Q}^l) \leq z_2 \leq F^u(z_1; \mathcal{Q}^u)\} \quad (30)$$

557 *be the set of points that lie between the frontiers and  $\hat{\mathbf{z}} = \arg \min_{\mathbf{z}^q \in \mathcal{Q}^l} s(\mathbf{z}^q)$  the current best*  
 558 *solution. Then, we always have that*

$$s(\hat{\mathbf{z}}) - s(\mathbf{z}^*) \leq L \underbrace{\max_{\mathbf{z}' \in \Gamma} \min_{\mathbf{z} \in \mathcal{F}^u} \|\mathbf{z} - \mathbf{z}'\|}_{d(\Gamma, \mathcal{F}^u)}. \quad (31)$$

559 *where  $d(\Gamma, \mathcal{F}^u) := \max_{\mathbf{z}' \in \Gamma} \min_{\mathbf{z} \in \mathcal{F}^u} \|\mathbf{z} - \mathbf{z}'\|$  is the max-min distance between the frontiers.*

560 *Proof.* By Assumption 1, we know that  $\mathbf{z}^* \in \mathcal{Z}$ . Due to Lemma 2, and the construction of the upper  
 561 and lower Frontiers (i.e.  $F^u(z_1) < z_2 \Rightarrow c(z_1, z_2) \geq 1$  and  $F^u(z_1) > z_2 \Rightarrow c(z_1, z_2) \geq 1$ ) we  
 562 always have that  $F^l(z_1^*) \leq z_2^* \leq F^u(z_1^*)$ , that is,  $\mathbf{z}^* \in \Gamma$ . Due to the monotonicity of  $s$ , we have  
 563  $\forall \mathbf{z} \in \mathcal{F}^u$  that

$$s(\hat{\mathbf{z}}) - s(\mathbf{z}^*) = \underbrace{(s(\hat{\mathbf{z}}) - s(\mathbf{z}))}_{\leq 0} - (s(\mathbf{z}^*) - s(\mathbf{z})) \quad (32)$$

$$\leq s(\mathbf{z}) - s(\mathbf{z}^*) \quad (33)$$

$$\leq L \|\mathbf{z} - \mathbf{z}^*\| \quad (34)$$

564 where the last step follows from the Lipschitz property of  $s$ . Finally, as  $\mathbf{z}^* \in \Gamma$  we have take the  
 565 maximum over  $\mathbf{z}' \in \Gamma$  so that the bound holds in the worst case. However, at the same time, we can  
 566 take the minimum over  $\mathbf{z} \in \mathcal{F}^u$  since (34) holds for all points  $\mathbf{z}$  on the upper frontier. Both steps  
 567 yield the final result

$$s(\hat{\mathbf{z}}) - s(\mathbf{z}^*) \leq L \max_{\mathbf{z}' \in \Gamma} \min_{\mathbf{z} \in \mathcal{F}^u} \|\mathbf{z} - \mathbf{z}'\|, \quad (35)$$

568 which concludes the proof.  $\square$

569 The key insight of Lemma 3 is that we can bound the sub-optimality  $s(\hat{\mathbf{z}}) - s(\mathbf{z}^*)$  with the maximum  
 570 distance of any point between the frontiers from its closest point on the upper frontier instead of  $\hat{\mathbf{z}}$ .  
 571 This is the case because  $\hat{\mathbf{z}}$  dominates any point on the upper frontier (i.e.,  $s(\hat{\mathbf{z}}) \leq s(\mathbf{z}') \forall \mathbf{z}' \in \mathcal{F}^u$ ). In  
 572 fact, we can further reduce the set of points which we have to consider for computing the max-min  
 573 distance  $d(\Gamma, \mathcal{F}^u)$  to the outer corner points of  $\Gamma$ , as defined in the following:

574 **Definition 2** (Outer Corner Points of  $\Gamma$ ). Let  $\mathbf{b}^l = (F_1^l(z_2^u), z_2^u)$  be the left upper and  $\mathbf{b}^r =$   
 575  $(z_1^u, F_2^l(z_1^u))$  the right outer corner points of  $\Gamma$ . Let  $\mathcal{Q}_{\text{ub}}^l = \mathcal{Q}^l \cup \{\mathbf{b}^l, \mathbf{b}^r\}$  and  $(z_1, \dots, z_{|\mathcal{Q}^l|+2})$   
 576 its ordering such that  $z_{1,i}, \leq z_{1,i+1}$ . We define

$$\Gamma_{\text{out}}^l(\mathcal{Q}^l) := \{(z_{k-1,1}, z_{k,2}) \mid i = 2, \dots, |\mathcal{Q}^l| + 2\} \cup \{\mathbf{b}^l, \mathbf{b}^r\} \quad (36)$$

577 as the the outer corner points of  $\Gamma$ .

578 This allows to compute the max-min distance  $d(\Gamma, \mathcal{F}^u)$  by maximizing over a finite set of outer  
 579 corner points, instead of the whole inter-frontier area  $\Gamma$ :

580 **Lemma 4** (Version of the max-min distance that is easier to compute). Let  $\Gamma_{\text{out}}^l$  be the outer corner  
 581 points of  $\Gamma$ , as defined in (40). Then, we have that

$$d(\Gamma, \mathcal{F}^u) = \max_{\mathbf{z} \in \Gamma_{\text{out}}^l} \min_{\mathbf{z}' \in \mathcal{F}^u} \|\mathbf{z} - \mathbf{z}'\|. \quad (37)$$

582 *Proof.* Due the construction of the frontiers, for every tuple  $(\mathbf{z}, \mathbf{z}') \in \Gamma \times \mathcal{F}^u$  that is the optimal  
 583 solution of the max-min problem in (31) there must be a  $(\tilde{\mathbf{z}}, \mathbf{z}') \in \Gamma_{\text{out}}^l \times \mathcal{F}^u$  that is the optimal  
 584 solution of the max-min problem on the RHS of (37). Again, we can show this by reduction:  
 585 Assume this is not the case, and  $\mathbf{z} \in \Gamma \setminus \Gamma_{\text{out}}^l$ , then, by the geometrical properties of  $\Gamma$  there exists a  
 586  $\tilde{\mathbf{z}} \in \Gamma_{\text{out}}^l$  which we can obtain by reducing  $z_1$  or  $z_2$  of  $\mathbf{z} = (z_1, z_2)$ , such that  $\|\tilde{\mathbf{z}} - \mathbf{z}'\| > \|\mathbf{z} - \mathbf{z}'\|$ .  
 587 Thus  $(\mathbf{z}, \mathbf{z}')$  cannot be the optimal solution of the max-min problem.  $\square$

588 As we only have to consider a finite set of outer corners, instead of the whole  $\Gamma$ , this makes the  
 589 max-min distance easier to implement in practice.

590 Since the max-min distance bounds how sub-optimal our current best solution can be, we want to  
 591 choose the next query so that we can shrink the max-min distance  $d(\Gamma, \mathcal{F}^u)$  between the frontiers.  
 592 For this purpose, we find a rectangle

$$\mathcal{R}_{\mathbf{z}, \mathbf{z}'} = \{\tilde{\mathbf{z}} \in \mathcal{Z} \mid \min\{z_1, z_1'\} \leq \tilde{z}_1 \leq \max\{z_1, z_1'\} \wedge \min\{z_2, z_2'\} \leq \tilde{z}_2 \leq \max\{z_2, z_2'\}\} \quad (38)$$

593 within  $\Gamma$  that has the largest max-min distance

$$d_{\mathcal{R}_{\mathbf{z}, \mathbf{z}'}}(\Gamma, \mathcal{F}^u) = d(\Gamma \cap \mathcal{R}_{\mathbf{z}, \mathbf{z}'}, \mathcal{F}^u \cap \mathcal{R}_{\mathbf{z}, \mathbf{z}'}). \quad (39)$$

594 To make finding this rectangle tractable, we can narrow down the points on the upper frontier which  
 595 we have to consider to its corner points, defined in the following:

**Definition 3** (Boundary Points of Upper Frontier). Let  $\mathcal{Z} = [z_1^l, z_1^u] \times [z_2^l, z_2^u]$  be the search domain.  
 Then

$$\mathbf{b}_v^u = (z_1^u, \min\{z_2 \mid (z_1, z_2) \in \mathcal{F}^u\}), \quad \mathbf{b}_h^u = (\min\{z_1 \mid (z_1, z_2) \in \mathcal{F}^u\}, z_2^u)$$

596 are the points of the upper frontier that intersect the vertical and horizontal domain boundary.

597 **Definition 4** (Corner Points of  $\mathcal{F}^u$ ). Let  $(\mathbf{z}_1, \dots, \mathbf{z}_{|\mathcal{Q}^u|+2})$  the ordering of  $\mathcal{Q}^u \cup \{\mathbf{b}_v^u, \mathbf{b}_h^u\}$  such that  
 598  $z_{1,i} \leq z_{1,i+1}$  and  $z_{2,i} \geq z_{2,i+1}$ . Then

$$\mathcal{F}_{cor}^u(\mathcal{Q}^u) := \{(z_{k-1,1}, z_{k,2}) \mid i = 2, \dots, |\mathcal{Q}^u| + 2\} \cup \mathcal{Q}^u \cup \{\mathbf{b}^l, \mathbf{b}^r\} \quad (40)$$

599 are the corner points of the upper frontier  $\mathcal{F}^u$ .

600 Now, we consider all outer corners of the inter frontier area  $\Gamma_{out}^l(\mathcal{Q}^l)$  and corners of the upper frontier  
 601  $\mathcal{F}_{cor}^u(\mathcal{Q}^u)$  to find the largest max-min rectangle:

602 **Definition 5** (Largest Max-Min Rectangle). The largest max-min rectangle, defined by its corner  
 603 points  $(\mathbf{z}, \mathbf{z}') \in \Gamma_{out}^l \times \mathcal{F}_{cor}^u$  is the rectangle in  $\Gamma$  with the largest max-min distance such that it fulfills  
 604 condition (3), formally:

$$\text{LARGESTMAXMINRECT}(\mathcal{Q}^l, \mathcal{Q}^u) = \arg \max_{(\mathbf{z}, \mathbf{z}') \in \Gamma_{out}^l(\mathcal{Q}^l) \times \mathcal{F}_{cor}^u(\mathcal{Q}^u)} d_{\mathcal{R}_{\mathbf{z}, \mathbf{z}'}}(\Gamma, \mathcal{F}^u) \text{ s.t. } (1) \wedge (2) \wedge (3)$$

605 (1) rectangle lies in  $\Gamma$ , i.e.  $(z_1, z'_1) \in \Gamma \wedge (z'_1, z_2) \in \Gamma$

606 (2) rectangle has a positive area, i.e.  $|z_1 - z'_1| \cdot |z_2 - z'_2| > 0$

607 (3) we cannot obtain a smaller rectangle  $(\mathbf{z}, \tilde{\mathbf{z}})$  with a upper frontier evaluation  $\tilde{\mathbf{z}} \in \mathcal{Q}^u$  that  
 608 matches  $(\mathbf{z}, \mathbf{z}')$  in one side of the rectangle but has a smaller second side, i.e.

$$\neg \exists \tilde{\mathbf{z}} \in \mathcal{Q}^u : (z_1 < \tilde{z}_1 < z'_1 \wedge z'_2 = \tilde{z}_2) \vee (z_2 < \tilde{z}_2 < z'_2 \wedge z'_1 = \tilde{z}_1) \quad (41)$$

609 Given the largest max-min rectangle  $\mathcal{R}_{\mathbf{z}, \mathbf{z}'}$ , we want to choose the next query point so that we  
 610 can reduce the max-min distance  $d_{\mathcal{R}_{\mathbf{z}, \mathbf{z}'}}(\Gamma, \mathcal{F}^u)$  within this rectangle as efficiently as possible. We  
 611 consider the set of query candidates

$$\mathcal{Q}_{\mathbf{z}, \mathbf{z}'} = \left\{ \underbrace{(z_1/2 + z'_1/2, z_2/2 + z'_2/2)}_{\text{center of rect.}}, \underbrace{(z'_1, z_2/2 + z'_2/2)}_{\text{middle of right side}}, \underbrace{(z_1/2 + z'_1/2, z'_2)}_{\text{middle of upper side}} \right\}, \quad (42)$$

612 consisting of the center point, together with the middle points of its right/upper sides. From this  
 613 query set, we choose the candidate that minimizes the worst-case max-min distance. In particular, if  
 614 we query a point  $\mathbf{z}^q$  there are two possible scenarios that will affect the max-min distance differently:  
 615 either the point satisfies the constraint ( $c(\mathbf{z}^q) \leq 1$ ) or it does not ( $c(\mathbf{z}^q) < 1$ ). We compute the  
 616 rectangle's max-min distance for both scenarios and choose the query-point that gives us the lowest  
 617 max-min distance in the less-favorable (worst-case) scenario:

$$\begin{aligned} \text{BESTWORSTCASEQUERY}(\mathbf{z}, \mathbf{z}', \mathcal{Q}^l, \mathcal{Q}^u) = \\ \arg \min_{\mathbf{z}^q \in \mathcal{Q}_{\mathbf{z}, \mathbf{z}'}} \max \left\{ d_{\mathcal{R}_{\mathbf{z}, \mathbf{z}'}}(\Gamma(\mathcal{Q}^l \cup \mathbf{z}^q, \mathcal{Q}^u), \mathcal{F}^u(\mathcal{Q}^u)), \right. \\ \left. d_{\mathcal{R}_{\mathbf{z}, \mathbf{z}'}}(\Gamma(\mathcal{Q}^l, \mathcal{Q}^u \cup \mathbf{z}^q), \mathcal{F}^u(\mathcal{Q}^u \cup \mathbf{z}^q)) \right\}. \end{aligned} \quad (43)$$

618 If one of the query candidates already lies on one of the frontiers, it cannot expand the frontiers and  
 619 thus also not improve the worst-case distance. Hence, we directly exclude such query candidates by  
 620 removing them from the argmin in (43).

621 **Theorem 2** (Appendix version of Theorem 1). Under the assumptions of Lemma 3 the Algorithm,  
 622 1 needs no more than  $k \leq 3^{\lceil \log_2(1/\epsilon) \rceil} = \mathcal{O}(\lceil (1/\epsilon) \rceil^{1.59})$  iterations to get

$$s(\hat{\mathbf{z}}_k) - s(\mathbf{z}^*) \leq L \|\mathbf{z}^u - \mathbf{z}^l\| (1/\epsilon). \quad (44)$$

623 close to the optimal solution.

624 *Proof.* In the worst case, it requires Algorithm 1 no more than three queries to half the max-min  
 625 distance within a max-min rectangle. The process of doing so maximally segments the inter frontier  
 626 area  $\Gamma$  within the rectangle into three new max-min rectangles. This can be checked by going though  
 627 the possible cases of how a max-min rectangle is split up by Algorithm 1. When a max-min rectangle

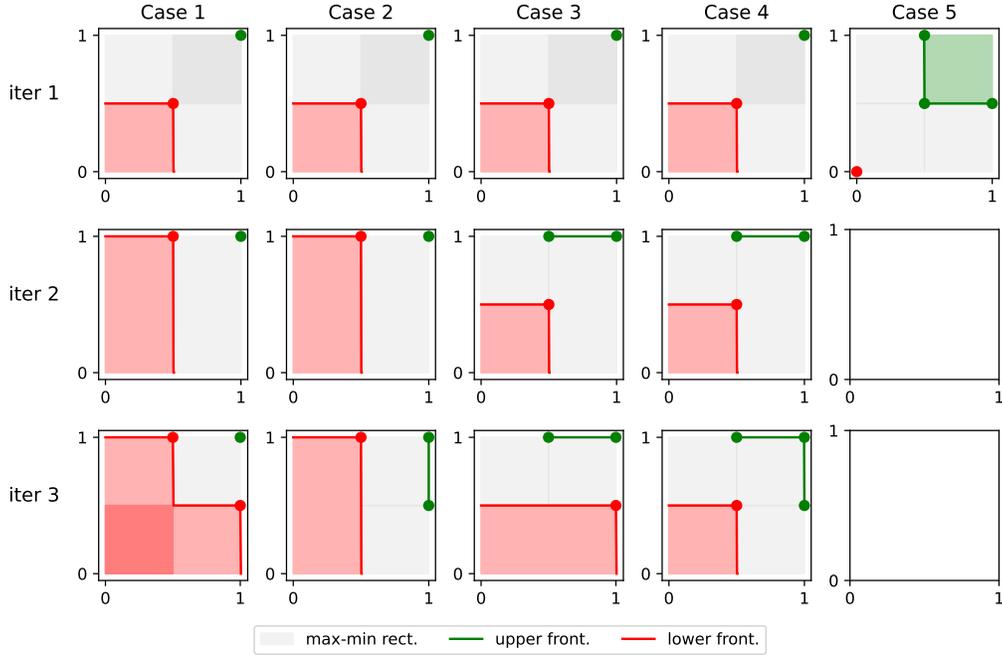


Figure 5: All possible cases how Algorithm 1 halves the max-min distance within a max-min rectangle when only the upper right corner of the rectangle belongs to the upper frontier.

628 is split up in three iterations, there are, in principle,  $8 = 2^3$  cases since each query can either lie  
629 above or below the constraint boundary. However, when the first query fulfills the constraint, it  
630 immediately halves the max-min distance. Thus no further queries are necessary and we effectively  
631 only have to consider 5 cases. Fig. 5 illustrates these five split-up cases for a max-min rectangle  
632 where only the upper right corner belongs to  $\mathcal{F}^u$ . Similarly, Fig. 6 and 7 illustrate the split-up cases  
633 when either half or the entire upper side of the rectangle belongs to  $\mathcal{F}^u$ . The cases when half or  
634 the entire right side of the rectangle belong to  $\mathcal{F}^u$  are analogous. In all of these scenarios, only  
635 two further queries are necessary to half the max-min distance. Finally, in the case where both the  
636 upper and the right side of the rectangle belong half or full to  $\mathcal{F}^u$  is trivial, as it only requires one  
637 more query. Thus, to shrink the max-min distance of all max-min rectangles by a factor of  $n$ , we  
638 need  $k \leq 3^{\lceil \log_2(n) \rceil}$  iterations at most. If we set  $n = 1/\epsilon$  it follows by Lemma 3 that we need  
639  $k \leq 3^{\lceil \log_2(1/\epsilon) \rceil}$  to obtain  $d(\Gamma, \mathcal{F}^u) \leq \|\mathbf{z}^u - \mathbf{z}^l\|/\epsilon$ . Finally, the result in the theorem follows from  
640 combining this with Lemma 3.  $\square$

## 641 D Implementation Details of SaMBO

642 **Calibration & Sharpness based Frontier Search.** To compute the calibration (5) and sharpness  
643 (6) across multiple datasets efficiently, we distribute the computation across individual processes  
644 for each dataset, using the ray distributed computing framework [56]. A model that always gives  
645 valid confidence intervals should always be calibrated, no matter what is the ordering of the datasets  
646  $\mathcal{D}_i$  which we split into train sets  $\mathcal{D}_{i,\leq t}$  and test sets  $\mathcal{D}_{i,>t}$  to compute (5). Based on this principle,  
647 we compute avg-std and calib based on the given ordering as well as the reversed ordering of each  
648 dataset, and average both results. In principle, more permutations can be done, but come with  
649 additional computational cost.

650 We perform frontier search to search for the kernel lengthscale  $l$  and variance  $\nu$  parameter for  
651 both the target and constraint function model. For the constraint function model, we use the  
652 safety constraint  $\text{calib}(\{\mathcal{D}_{i,T_i}^q\}_{i=1}^n, l_q, \nu_q) \geq 1$  (i.e.  $\text{calib} = 1$  since  $\text{calib} \in [0, 1]$ ). For the

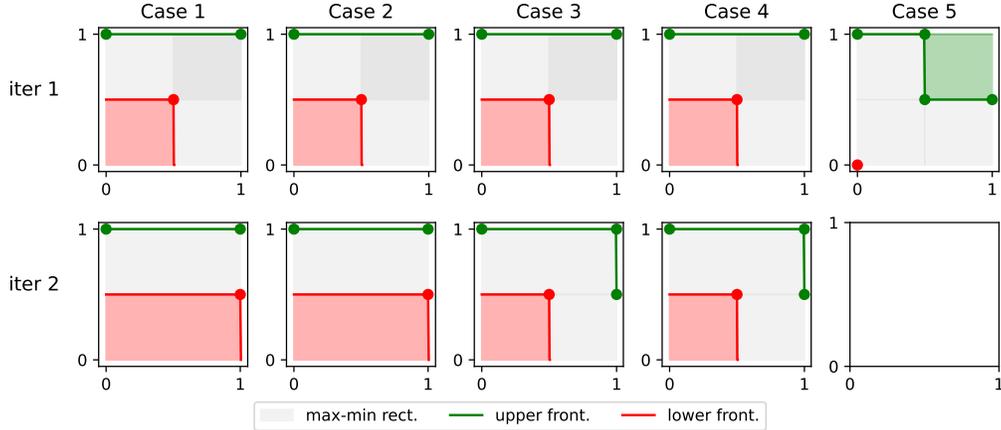


Figure 6: All possible cases how Algorithm 1 halves the max-min distance within a max-min rectangle when the upper side of the rectangle belongs to the upper frontier.

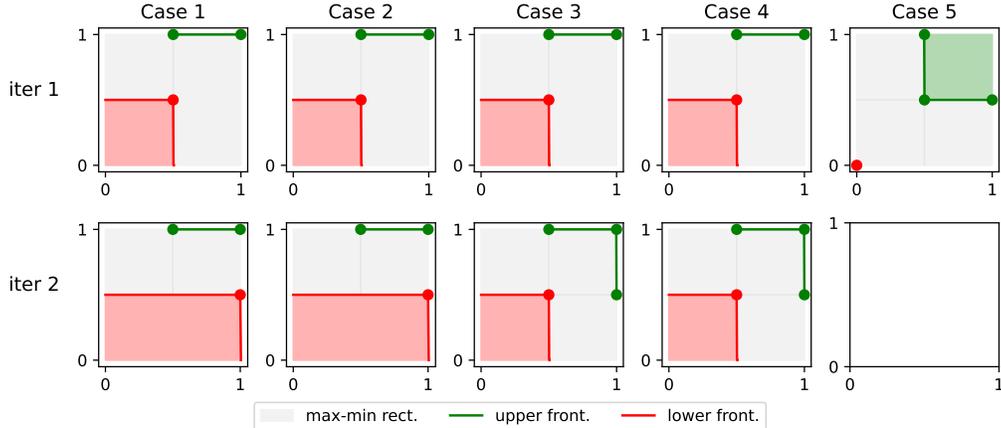


Figure 7: All possible cases how Algorithm 1 halves the max-min distance within a max-min rectangle when half of the upper side of the rectangle belongs to the upper frontier.

653 target function model, the calibration requirements are not safety related and thus not as strict  
654 as in case of the constraint function. Hence, we only use a calibration frequency constraint  
655  $\text{calib}(\{\mathcal{D}_{i,T_i}^f\}_{i=1}^n, l_f, \nu_f) \geq 0.95$  of 95 % when performing frontier search on  $l_f$  and  $\nu_f$ . For both  
656 models, we run frontier search for 20 iterations as this gives us already a good solution without  
657 incurring to much computational cost. For the variances, we use an lower and upper boundary of 1.0  
658 and 6.0 whereas for the lengthscale we consider the range [0.01, 5.0]. The frontier search is per-  
659 formed in the log-space to the base 10 of both variance and lengthscale, i.e., we transform the both  
660 variance and lengthscale with  $t(z) = \log_{10}(z)$  during the frontier search and in the end transform  
661 back the result by  $t^{-1}(z) = 10^z$ .

662 **Parallelizing the sharpness and calibration computations.** Computing the calibration and sharp-  
663 ness metrics in (5) and (6) is computationally expensive since it requires computing the calib-freq  
664 and avg-std over all  $n$  tasks and  $T_i - 1$  subsequences within each task. Fortunately, each of these  
665 computations can be done independently. Thus, we use compute each of the summands in (5) and  
666 (6) in parallel processes and aggregate the results in the end. This significantly reduces the compu-  
667 tation time for the calibration and sharpness which allows us to perform 20 steps of frontier search  
668 relatively fast.

669 **The NN-based GP prior.** Following [49, 11], we parameterize the GP prior  $\rho_{\theta}(h) =$   
670  $\mathcal{GP}(h|m_{\theta}(\mathbf{x}), k_{\theta}(\mathbf{x}, \mathbf{x}'))$ , particularly the mean  $m_{\theta}$  and kernel function  $k_{\theta}$ , as neural networks  
671 (NN). Here, the parameter vector  $\theta$  corresponds to the weights and biases of the NN. To ensure the  
672 positive-definiteness of the kernel, we use the neural network as feature map  $\Phi_{\theta}(\mathbf{x}) : \mathcal{X} \mapsto \mathbb{R}^d$   
673 that maps to a  $d$ -dimensional real-values feature space in which we apply a squared exponential  
674 kernel. We choose the dimensionality of the feature space the same as the domain. Accordingly, the  
675 parametric kernel reads as

$$k_{\theta}(x, x') = \nu_P \exp(-\|\Phi_{\theta}(\mathbf{x}) - \Phi_{\theta}(\mathbf{x}')\|^2 / (2l_P)) . \quad (45)$$

676 Both  $m_{\theta}(\mathbf{x})$  and  $\Phi_{\theta}(\mathbf{x})$  are fully-connected neural networks with 3 layers with each 32 neurons  
677 and tanh non-linearities. The kernel variance  $\nu_P$  and lengthscale  $l_P$  are also learnable parameters  
678 which are appended to the NN parameters  $\theta$ . Since  $l_P$  and  $\sigma_P^2$  need to be positive, we represent  
679 and optimize them in log-space. Unlike Rothfuss et al. [11], we set the variance  $\sigma^2$  of the Gaussian  
680 likelihood  $p(y|h(\mathbf{x})) = \mathcal{N}(y; h(\mathbf{x}), \sigma^2)$  as a fixed parameter rather than meta-learning it. This is,  
681 because, in order to ensure safety, the likelihood variance needs to be the same as the one that was  
682 used during frontier search, when computing then calibration and sharpness.

683 **The hyper-prior.** We use a Vanilla GP  $\mathcal{GP}(0, k(x, x'))$  as hyper-prior stochastic process. In that,

$$k(x, x') = \nu \exp(-\|x - x'\|^2 / (2l)) \quad (46)$$

684 is a SE kernel with variance  $\nu$  and lengthscale  $l$  chosen by the frontier search procedure, discussed  
685 in Sec. 4.2.

686 **Minimization of the F-PACOH objective.** To estimate the F-PACOH objective in (23), we use  
687 measurement sets of size  $k = 20$ , i.e., considering a subset of 10 training points and 10 uniformly  
688 sampled domain points per iteration. We minimize the loss by performing 5000 iterations with the  
689 Adam optimizer with a learning rate of 0.001.

690 **Code and Data.** We provide implementations of the SaMBO components as well as the experi-  
691 ment scripts and data under <https://tinyurl.com/safe-meta-bo>.

## 692 E Further discussions about SaMBO

693 **Discussion on the applicability to higher-dimensional problems.** Similar to other BO / safe BO  
694 methods that are based on GPs, the sample complexity grows exponentially with the number of  
695 dimensions. The meta-learning in our case alleviates this issue to some degree by making the GP  
696 prior more informed about our environment of tasks in areas where meta-training data is available.  
697 However, to maintain safety, we regularize the meta-learned GP prior towards a Vanilla GP with SE  
698 kernel (cf. Eq. 3) in areas without or little meta-training data. The higher-dimensional our safe BO  
699 problem, the sparser is the meta-training data in the domain. Thus, the areas where our meta-learned  
700 GP resembles a Vanilla GP become larger in proportion and we face again the general issue of poor  
701 sample complexity in high dimensions. Generally, without additional assumption that, e.g., there  
702 exists some lower-dimensional sub-space in which the functions we are optimizing lie, we do not  
703 believe that the ‘curse of dimensionality’ problem can be solved while, at the same time, assuring  
704 safety. Generally, we not aware of any safe BO method that has been successfully applied to a  
705 high-dimensional optimization problem with safety constraint.

706 **Discussion on the monotonicity of the calibration-sharpness constraint optimization problem.**  
707 As discussed in Section 4.2, the introduced frontier search algorithm aims to exploit the mono-  
708 tonicity properties of the calibration-sharpness constraint optimization problem in 7. For large  
709 ranges of  $l$  and  $\nu$ , the avg-std and avg-calib are monotonically increasing in the kernel variance  
710  $\nu$  and decreasing in the lengthscale  $l$ . While the monotonicity for the avg-std provably holds across  
711 the spectrum, the monotonicity of the calibration frequency in  $l$  is only an empirical heuristic that  
712 holds in almost all cases if  $\nu$  is at least as big as the variance of the targets  $y$  in a dataset. If the  
713 kernel variance is chosen smaller than the variance of the data, significant parts of the function

714 that underlie the data are outside of the high-probability  
715 regions of the GP prior and the GPs predictive distribution  
716 systematically underestimates the corresponding variance  
717 in the absence of close-by data points. In such a cases,  
718 a larger lengthscale can actually improve the calibration  
719 by increasing the sphere of local influence of data points  
720 on the predictive distribution and, thus, slowing down the  
721 reversion of the posterior towards the prior with too small  
722 variance.

723 Figure 8 displays the calibration frequency in response  
724 to  $l$  and  $\nu$  in the same setting as Figure 2, but with  
725 an extended range where the kernel variance  $\nu$  becomes  
726 smaller than 1 (data is standardized) and thus smaller than  
727 the variance of the data. The lower left corner where  $\nu <$   
728 1 illustrates the described breakdown of monotonicity.

729 In practice, this is not an issue since we can easily we  
730 standardize our data and then choose the lower bound of  
731 kernel variance in the frontier search to be 1. In the re-  
732 sulting search space, our constraint optimization problem  
733 is monotone and the frontier search converges as expected.

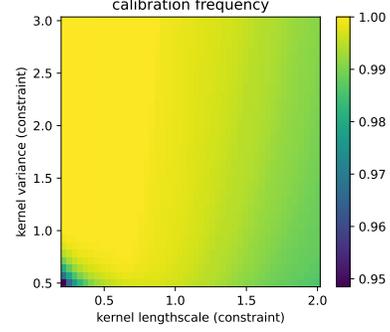


Figure 8: Calibration frequency for varying kernel lengthscales  $l$  and variances  $\nu$ . The data underlying the calibration computations has been standardized. While avg-calib is monotone in  $l$  for  $\nu > 1$ , the monotonicity breaks down if the kernel variance is smaller than the variance of the data.

## 734 F Experiment Details

### 735 F.1 SafeBO environments

#### 736 F.1.1 The Camelback + Random Sinusoids environment

737 The environment objective/target function corresponds to a Camelback function [54]

$$g(x_1, x_2) = \max \left( -(4 - 2.1 \cdot x_1^2 + x_1^4/3) * x_1^2 - x_1 x_2 - (4 \cdot x_2^2 - 4) * x_2^2, -2.5 \right). \quad (47)$$

738 plus random sinusoid functions, defined over the 2-dimensional cube  $\mathcal{X} = [-2, 2] \times [-1, 2]$  as  
739 domain. Specifically, the target function is defined as

$$f(x_1, x_2) = g(x_1, x_2) + a \sin(\omega_f * (x_1 - \rho)) \sin(\omega_f * (x_2 - \rho)) \quad (48)$$

740 wherein the parameters are sampled independently as

$$a \sim \mathcal{U}(0.3, 0.5), \quad \omega_f \sim \mathcal{U}(0.2, 2.0), \quad \rho \sim \mathcal{N}(0, 1.0). \quad (49)$$

741 The constraint function  $q(\mathbf{x})$  is a linear combination of the camelback function  $g(\mathbf{x})$ , a product of  
742 sinusoids along the two dimensions and a quadratic component that ensures that the safe regions are  
743 sufficiently connected so that they can be reached:

$$q(x_1, x_2) = 3 \cdot \sin(0.4 \cdot \pi \cdot \omega_q - 2) \cdot \sin(2\pi \cdot \omega_q) - b \cdot (x_1^2 + x_2^2) + 1.2 \cdot g(x_1, x_2) - 0.7. \quad (50)$$

744 Here the parameters  $\omega_q$  and  $b$  are sampled as follows:

$$\omega_q \sim \mathcal{U}(0.45, 0.5), \quad b \sim \mathcal{U}(0.3, 0.5) \quad (51)$$

745 The optimization domain is  $\mathcal{X} = [-2, 2] \times [-1, 1]$ . As initial safe point we use  $\mathcal{S}_0 =$   
746  $\{(-1.5, -0.5)\}$ . Figure 9 displays an example task of the Camelback + Random Sinusoids en-  
747 vironment.

#### 748 F.1.2 The Random Eggholder environment

749 The random Eggholder environemnt is based on the Eggholder function [55], a popular benchmark  
750 function used in global optimization:

$$f(x_1, x_2) = -(x_2 + c) \cdot \sin \left( \sqrt{|ax_2 + x_1/2 + 47|} \right) - b \cdot x_1 \cdot \sin \left( \sqrt{|x_1 - x_2 - 47|} \right) \quad (52)$$

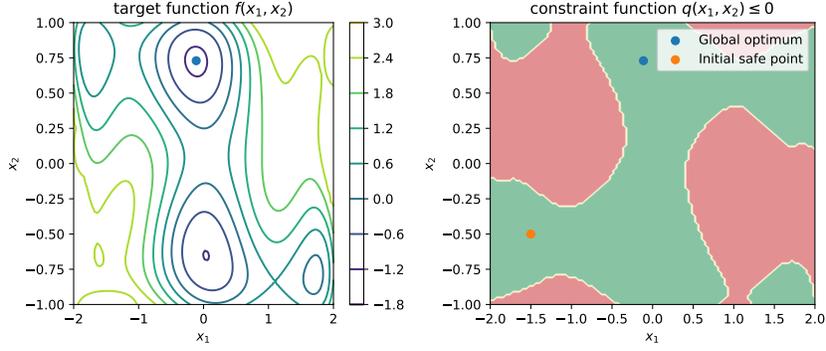


Figure 9: Example task of the Camelback + Random Sinusoids environment. Left: target function. Right: Safe regions in green and unsafe regions in red.

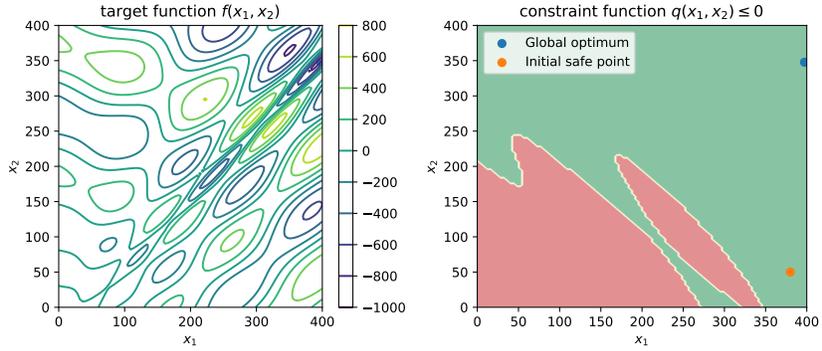


Figure 10: Example task of the Random Eggholder environment. Left: target function. Right: Safe regions in green and unsafe regions in red.

751 We randomly sample its parameters  $a, b, c$  as follows:

$$a \sim \mathcal{U}(0.6, 1.4), \quad b \sim \mathcal{U}(0.6, 1.4), \quad c \sim \mathcal{N}(47, 5^2) \quad (53)$$

752 to obtain different tasks. The corresponding constraint function is defined as

$$q(x_1, x_2) = 300 - \sqrt{x_1^2 + 2x_2^2} + 50 \sin((\omega_1 x_1 + \omega_2 x_2)/20), \quad (54)$$

753 where the frequencies are sampled independently as  $\omega_1, \omega_2 \sim \mathcal{U}(0.8, 1.2)$ . The optimization domain  
 754 is  $\mathcal{X} = [0, 400] \times [0, 400]$ . As initial safe point we use  $\mathcal{S}_0 = \{(380, 50)\}$ . Figure 10 displays an  
 755 example task of the Random Eggholder environment. The environment is particularly challenging  
 756 since the Eggholder function has many local minima.

### 757 E.1.3 Tuning controller parameters for the Argus linear robotic platform

758 The high-precision motion system Argus from Schneeberger Linear Technology is shown on Fig.  
 759 11). It is a 3 axes positioning system with 2 orthogonal linear axes and a rotational axes on top. The  
 760 system has an accuracy of  $\pm 10 \mu\text{m}$ , bidirectional repeatability of  $0.7 \mu\text{m}$  and  $3\sigma$  position stability of  
 761  $< 1 \text{nm}$ . In our experiments we focus on the upper linear axis of the system.

762 The exact optimization problem is

$$\mathbf{x}^* = \arg \min_{\mathbf{x}=[\text{PKP}, \text{VKP}, \text{VKI}]} T_s \cdot \frac{d}{dt} \text{pe}[t_{\text{move}} : t_{\text{end}}] \quad (55)$$

$$\text{s.t.} \quad \text{FFT}_{\max}(\mathbf{x}^*) = \max_{\omega \in [0.03, 0.07]} |\text{fft}(\text{ve})|(\omega) + b \cdot \max_{\omega \in [0.08, 0.1]} |\text{fft}(\text{ve})|(\omega) - \kappa < 0 \quad (56)$$



Figure 11: Argus motion system

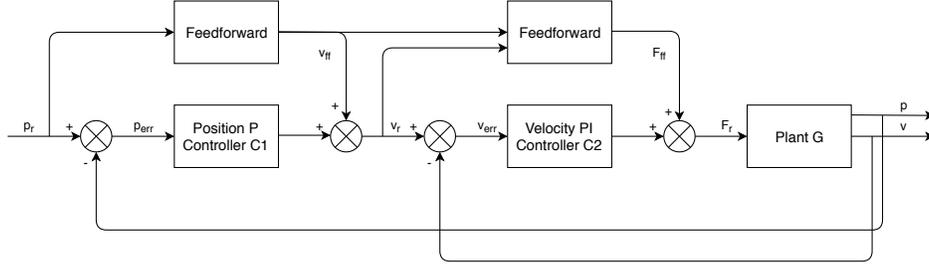


Figure 12: Argus controller structure

763 where PKP is the proportional gain of the position control loop, VKP is the proportional gain of  
 764 the velocity control loop, VKI is the integral gain of the velocity control loop,  $T_s$  is the settling  
 765 time,  $p_e$  is the position error measurement,  $v_e$  is the velocity error measurement,  $t_{\text{move}}$  is the move-  
 766 time defined as the timepoint where the reference movement is finished,  $t_{\text{end}}$  is the time endpoint  
 767 of the movement measurement, set to 1.2s,  $\omega$  is the frequency,  $\text{fft}$  is the fast Fourier transform in a  
 768 given frequency window,  $b = 5$  is a scaling factor, and  $\kappa$  is the constraint limit, dependent on the  
 769 reference stepsize. The position reference trajectory is a step-wise constant jerk based s-curve with  
 770 jerk =  $200 \frac{\text{m}}{\text{s}^3}$ , maximum acceleration  $a = 20 \frac{\text{m}}{\text{s}^2}$  and maximum velocity  $v = 1 \frac{\text{m}}{\text{s}}$ .

771 The simulation model of the Argus used in section 6.3 contains a cascade controller and a model of  
 772 the Argus system (plant). The cascaded controller (see figure 12) was rebuilt from the real controller  
 773 design and the plant (see figure 13) was modelled by 1) a fitted linear transfer function  $G(s)$ , con-  
 774 taining a double integrator, the five most dominant resonances of form  $T(s) = \frac{s^2/\omega_{ni} + 2\lambda_{ni}/\omega_{ni} + 1}{s^2/\omega_{di} + 2\lambda_{di}/\omega_{di} + 1}$   
 775 and a dead time, using frequency domain data and 2) a nonlinear position-dependent cogging  
 776 model  $F_c(p)$ , based on linear interpolation of a lookup table for 16000 points at positions between  
 777  $\pm 200\text{mm}$  of the axis using a discretization of 0.025mm, where  $p$  is position,  $v$  is velocity, and  $F_r$   
 778 is the force reference (proportional to current) applied to the motor of the axis. The linear transfer  
 779 function has order 13. The resonant and anti-resonant frequencies are displayed in table 1, with  
 780  $f_{n/d} = \omega_{n/d} \cdot \sqrt{1 - 2\lambda_{n/d}^2}$  and dead time  $t_{\text{dead}} = 2\text{ms}$

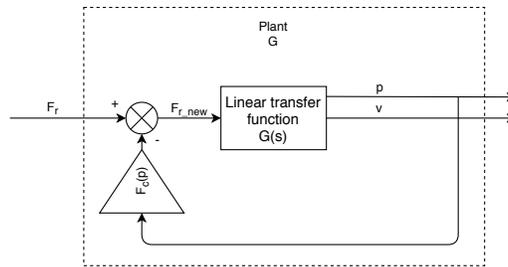


Figure 13: Argus controller structure

	$f_n$ in Hz	$\lambda_n$	$f_d$ in Hz	$\lambda_d$
1	390	0.1	400	0.1
2	475	0.03	500	0.05
3	690	0.03	800	0.06
4	870	0.03	900	0.04
5	1050	0.03	1100	0.06

Table 1: Argus simulation linear transfer function parameters

Environment	# meta-train tasks	# points per task	lengthscale constr model
Camelback + Random Sin	40	100	0.5
Random Eggolder	40	200	0.4
Argus Controller Tuning	20	400	0.4

Table 2: Specifications of the meta-train data

781 The optimization domain is 3-dimensional and corresponds to the three controller gains PKP, VKP,  
782 VKI, restricted to the following ranges:  $\mathcal{X} = [100, 400] \times [300, 1200] \times [500, 4000]$ . As initial safe  
783 set, we take  $\mathcal{S}_0 = \{(200, 800, 1000)\}$ .

## 784 F.2 Environment Normalization and Data Collection

785 **Environment Normalization.** The three environments, specified in Appx F.1 have vastly different  
786 scales, both in  $\mathbf{x}$ ,  $\tilde{f}$  and  $\tilde{q}$ . To alleviate the problems arising from different value ranges we *stan-*  
787 *dardize* the environment data such that the value ranges are roughly those of a standard normally  
788 distributed random variable, before we pass on the data to the GP model. Hence, the kernel variances  
789 and lengthscales displayed in Fig. 6.2 correspond to the standardized value ranges.

790 To determine the standardization statistics (i.e., mean and standard deviation (std)) for  $\mathbf{x}$  we use  
791 the domain ranges  $([x_i^l, x_i^u])$  for dimension  $i = 1, \dots, d$  of the environments. Assuming a uniform  
792 distribution over the domain, we use  $\mu_{x_i} = \frac{x_i^u + x_i^l}{2}$  as mean and  $\sigma_{x_i} = \sqrt{\frac{(x_i^u - x_i^l)^2}{12}}$ . For  $\tilde{f}$  and  $\tilde{q}$   
793 we are more conservative, because we do not want to underestimate their std which could lead to  
794 getting stuck in local optima or safety violations. In particular, we consider the respective minimum  
795 and maximum values of  $\tilde{f}$  and  $\tilde{q}$  observed in the union of the meta-train datasets, i.e.  $\tilde{f}^{\min}$ ,  $\tilde{f}^{\max}$ ,  
796  $\tilde{q}^{\min}$ ,  $\tilde{q}^{\max}$ . For  $\tilde{f}$ , we use  $\mu_{\tilde{f}} = \frac{\tilde{f}^{\max} + \tilde{f}^{\min}}{2}$  and  $\sigma_{\tilde{f}} = \frac{\tilde{f}^{\max} - \tilde{f}^{\min}}{3}$ . For the constraint values, we set  
797  $\mu_{\tilde{q}} = 0$  because we do not want to distort the safety threshold and use  $\sigma_{\tilde{q}} = \frac{\max\{|\tilde{q}^{\max}|, |\tilde{q}^{\min}|\}}{2}$  as std  
798 to re-scale the constraint values.

799 **Collection of meta-train data.** We collect the meta-training data by running SAFEOPT on each  
800 task. Table 2 holds the specifications for each of the three environments, i.e. the number of tasks  
801  $n$ , the number of points/iterations  $T_i$  per task and the constraint model lengthscale that is used for  
802 SAFEOPT. In all cases, we use a conservative lengthscale of 0.2 for the target model and a likelihood  
803 std of 0.1. Conservative parameter choices like these ensure that we continue exploring (within the  
804 confines of SAFEOPT) throughout the data collection and don not start to over-exploit, e.g. by  
805 querying the same data point many times.

## 806 F.3 Parameter Choices for the Experiments

807 **Domain discretization.** Since both SAFEOPT and GOOSE require finite domains, we discretize  
808 the continuous domains of the safe BO tasks with 40000 uniform points each.

809 **Likelihood Std.** The standard deviation of the Gaussian likelihood of the GP is the only hyper-  
810 parameter we have to choose when employing our proposed frontier search in combination with  
811 Vanilla GPs. However, the observation noise is often known or easy to estimate by querying the

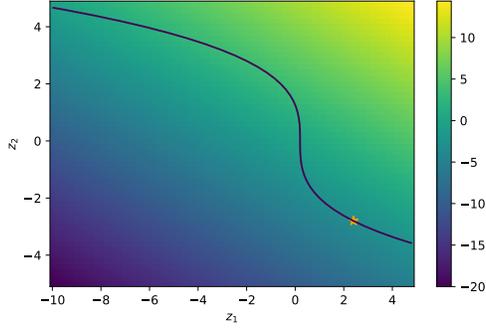


Figure 14: Monotone optimization problem in Eq. 57

812 same point multiple times and observing the variation of responses. We use the latter approach to  
 813 determine the likelihood std in our experiments, obtaining the following settings:  $\sigma = 0.02$  for  
 814 the Camelback + Random Sin environment,  $\sigma = 0.05$  for the Random Eggholder environment and  
 815  $\sigma = 0.1$  for the Argus Controller parameter tuning.

816 **GoOSE expander sets.** For computing the GoOSE expander sets  $W_t$ , we always use  $\epsilon = 0.2$ . Note,  
 817 that this applies after the environment standardization.

## 818 G Further Experiment Results

### 819 G.1 Frontier Search

820 In this section, we aim to investigate how fast the proposed Frontier Search algorithm (c.f. Algorithm  
 821 1) approaches the optimal solution of a monotone optimization problem like the one in (24). To this  
 822 end, we use the following monotone, constraint optimization problem to test the algorithm:

$$\begin{aligned} \min_{\omega} s(\mathbf{z}) \quad \text{s.t.} \quad c(\mathbf{z}) \geq 1 \\ \text{with } s(\mathbf{z}) := z_1 + 2z_2 \\ q(\mathbf{z}) := 5 * z_1 + 0.5 * z_2^3 - 3 \end{aligned} \tag{57}$$

823 Figure 14 visualizes the  $s(\mathbf{z})$  and the constraint boundary  $q(\mathbf{z}) = 0$ . In Figure 15 we visualize the  
 824 first 33 iterations of Frontier Search on the optimization problem in (57). As we can see, Algorithm  
 825 1 quickly shrinks the area between the frontiers and returns solutions close to the optimum after only  
 826 a handful of queries.

827 Finally, Figure 16 visualizes the convergence of Frontier Search, both in terms of the max-min  
 828 distance and the sub-optimality  $s(\hat{\mathbf{z}}^*) - s(\mathbf{z}^*)$  of the solution, together with the bounds provided in  
 829 Theorem 1.

### 830 G.2 Compute Time Analysis

831 Here, we provide a computational analysis of various components of the proposed SaMBO algo-  
 832 rithm. Table 3 reports the time it takes to finish the computations associated with the frontier search  
 833 as well as F-PACOH for different sizes of the meta-training data. Each experiment has been run on  
 834 16 processor cores of an Intel Xeon 8360Y 2.4GHz CPU. As in all the experiments, we use 20 itera-  
 835 tions of frontier search and 5000 iterations of stochastic gradient descent (with the Adam optimized)  
 836 on the F-PACOH objective.

837 As already discussed in Appendix D, we can use parallelization for the computation of the avg-calib  
 838 and avg-std metric which makes frontier search relatively fast. For instance, for  $n = 20$  tasks and  
 839  $T = 200$  samples, the frontier search part of Algorithm 2 finishes in ca. 35 seconds. The compu-  
 840 tational complexity of the calibration and sharpness computations grows with  $T$  since the number

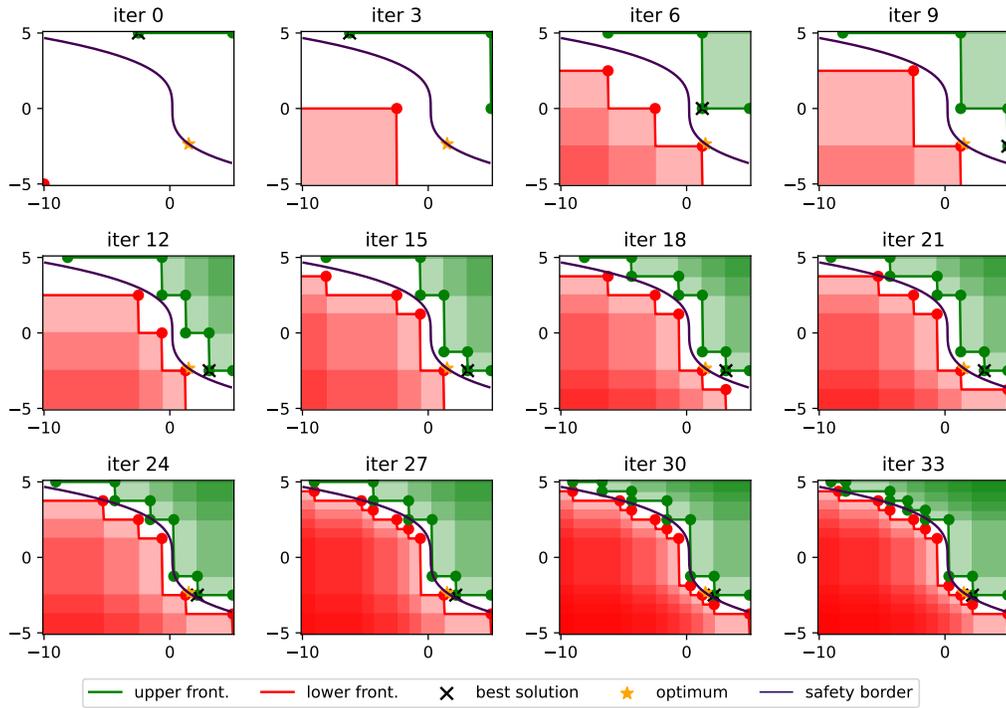


Figure 15: Frontier Search on the monotone optimization problem in in Eq. 57. Algorithm 1 quickly shrinks the area between the frontiers and returns solutions close to the optimum after only a handful of queries.

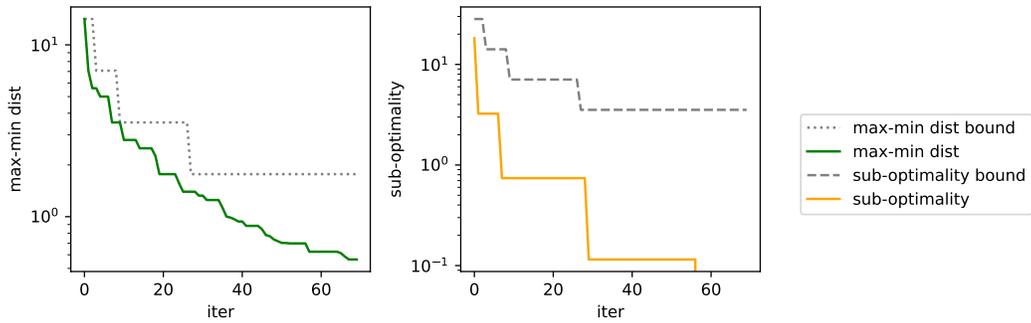


Figure 16: Convergence of Frontier Search for the optimization problem in Eq. 57.

$n$	$T$	Duration Frontier Search	Duration F-PACOH
10	50	$14.19 \pm 0.52$	$169.55 \pm 1.99$
10	100	$18.83 \pm 0.62$	$183.55 \pm 2.21$
10	200	$33.82 \pm 1.26$	$230.90 \pm 4.57$
10	400	$75.23 \pm 10.42$	$247.89 \pm 11.05$
20	50	$13.96 \pm 0.87$	$178.47 \pm 4.93$
20	100	$22.14 \pm 0.68$	$187.57 \pm 1.84$
20	200	$37.44 \pm 2.54$	$224.54 \pm 8.75$
20	400	$95.15 \pm 11.65$	$243.12 \pm 8.94$

Table 3: Compute time in seconds for Frontier Search and F-PACOH for different number of tasks  $n$  and samples per task  $T$ . Reported is the mean and standard deviation over 5 sees / repetitions.

841 of sub-sequences of each dataset, as well as the maximum size of a sub-sequence used for GP in-  
842 ference grows. While the former can be parallelized, GP inference cannot be easily parallelized.  
843 The meta-training of the GP priors, typically takes ca. 3-4 minutes. Finally, the compute time for  
844 the safe BO part strongly varies depending on the employed algorithm (e.g. GoOSE or SafeOpt)  
845 and how time intensive/costly it is to query the  $f$  and  $q$ . In the best case, when we use simulated  
846 toy functions that are very cheap to evaluate and run GoOSE for 200 steps, the safe BO takes ca.  
847 500 seconds (8-9 min). In case of SafeOpt, the compute time is highly variable depending on the  
848 expander computation and ranges from 15 minutes to 1.5 hours. Generally, when employed with a  
849 real system or expensive simulation, the time for the safe BO outweighs the compute time for the  
850 frontier search and F-PACOH steps of SaMBO by a order of magnitude.

### 851 G.3 Study on amount of meta-train data

852 In this section, we aim to study how varying amounts of meta-training data affect the overall per-  
853 formance of SaMBO. For this purpose, we employ SaMBO-G on the the Camelback Function +  
854 Random Sinusoid environment with a varying number of meta-training tasks and data points per  
855 task. Figure 17 displays the terminal inference regret at ( $T = 100$ ) for a varying number of tasks  
856  $n$  with each  $m = 100$  data points (left) and a varying number of tasks  $n$  with each  $m = 100$  data  
857 points (right).

858 In all cases, the performance of SaMBO-G is substantially better than the inference regret of  $> 0.1$   
859 we obtain when running GoOSE with a Vanilla GP, as done in Figure 4. This demonstrates that we  
860 can also successfully perform positive transfer when the amount of meta-training data is smaller.  
861 Finally, we observe that the performance improvement we gain by doubling the number of tasks is  
862 much bigger than when we double the number of points per task. This is consistent with learning  
863 theory in e.g. Pentina and Lampert [30] and Rothfuss et al. [34].

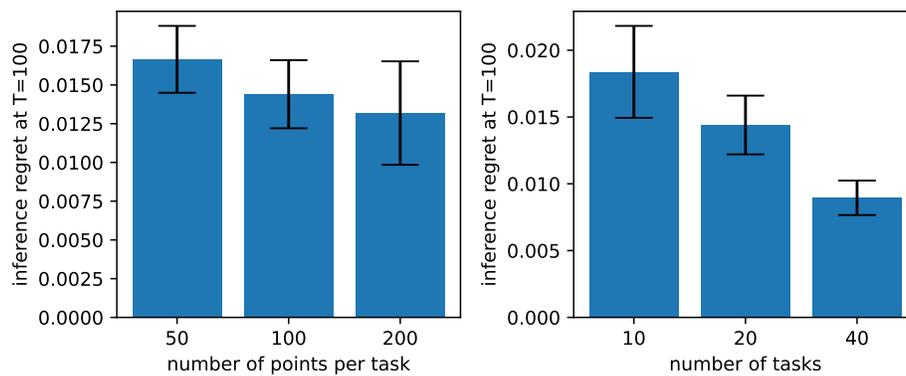


Figure 17: Overall performance for different amounts of meta-training. Displayed is the inference regret at last step ( $T = 100$ ) of SaMBO-G in the Camelback Function + Random Sinusoid environment. Left: Varying number of data points per task  $m$  for  $n = 20$  tasks. Right: varying number of tasks  $n$  with each  $m = 100$  data points. Doubling the number of tasks leads to more performance improvements than doubling the number of points per tasks.