

---

# Transductive Active Learning: Theory and Applications

---

**Jonas Hübötter\***

Department of Computer Science  
ETH Zürich, Switzerland

**Bhavya Sukhija**

Department of Computer Science  
ETH Zürich, Switzerland

**Lenart Treven**

Department of Computer Science  
ETH Zürich, Switzerland

**Yarden As**

Department of Computer Science  
ETH Zürich, Switzerland

**Andreas Krause**

Department of Computer Science  
ETH Zürich, Switzerland

## Abstract

We study a generalization of classical active learning to real-world settings with concrete prediction targets where sampling is restricted to an accessible region of the domain, while prediction targets may lie outside this region. We analyze a family of decision rules that sample adaptively to minimize uncertainty about prediction targets. We are the first to show, under general regularity assumptions, that such decision rules converge uniformly to the smallest possible uncertainty obtainable from the accessible data. We demonstrate their strong sample efficiency in two key applications: active fine-tuning of large neural networks and safe Bayesian optimization, where they achieve state-of-the-art performance.

## 1 Introduction

Machine learning, at its core, is about designing systems that can extract knowledge or patterns from data. One part of this challenge is determining not just how to learn given observed data but deciding what data to obtain next, given the information already available. More formally, given an unknown and sufficiently regular function  $f$  over a domain  $\mathcal{X}$ : *How can we learn  $f$  sample-efficiently from (noisy) observations?* This problem is widely studied in *active learning* and *experimental design* (Chaloner & Verdinelli, 1995; Settles, 2009).

Active learning methods commonly aim to learn  $f$  globally, i.e., across the entire domain  $\mathcal{X}$ . However, in many real-world problems, **(i)** the domain is so large that learning  $f$  globally is hopeless or **(ii)** agents have limited information and cannot access the entire domain (e.g., due to restricted access or to act safely). Thus, global learning is often not desirable or even possible. Instead, intelligent systems are typically required to act in a more *directed* manner and *extrapolate* beyond their limited information. This work formalizes the above two aspects of active learning, which have remained largely unaddressed by prior work. We provide a comprehensive overview of related work in Section 6.

**“Directed” transductive active learning** We consider the generalized problem of *transductive active learning*, where given two arbitrary subsets of the domain  $\mathcal{X}$ ; a *target space*  $\mathcal{A} \subseteq \mathcal{X}$ , and a *sample space*  $\mathcal{S} \subseteq \mathcal{X}$ , we study the question:

*How can we learn  $f$  within  $\mathcal{A}$  by actively sampling observations within  $\mathcal{S}$ ?*

---

\*Correspondence to [jonas.huebotter@inf.ethz.ch](mailto:jonas.huebotter@inf.ethz.ch)

This problem is ubiquitous in real-world applications such as safe Bayesian optimization, where  $\mathcal{S}$  is a set of safe parameters and  $\mathcal{A}$  might represent parameters outside  $\mathcal{S}$  whose safety we want to infer. Active fine-tuning of neural networks is another example, where the target space  $\mathcal{A}$  represents the test set over which we want to minimize risk, and the sample space  $\mathcal{S}$  represents the dataset from which we can retrieve data points to fine-tune our model to  $\mathcal{A}$ . Figure 1 visualizes some instances of transductive active learning.

Whereas most prior work has focused on the “global” inductive instance  $\mathcal{X} = \mathcal{A} = \mathcal{S}$ , MacKay (1992) was the first to consider specific target spaces  $\mathcal{A}$  and proposed the principle of selecting points in  $\mathcal{S}$  to minimize the “posterior uncertainty” about points in  $\mathcal{A}$ . Since then, several works have studied this principle empirically (e.g., Seo et al., 2000; Yu et al., 2006; Bogunovic et al., 2016; Wang et al., 2021; Kothawade et al., 2021; Bickford Smith et al., 2023). In this work, we model  $f$  as a Gaussian process or (equivalently) as a function in a reproducing kernel Hilbert space, for which the above principle is analytically and computationally tractable. Our contributions are:

- **Theory (Section 3):** We are the first to give rates for the uniform convergence of uncertainty over the target space  $\mathcal{A}$  to the smallest attainable value, given samples from the sample space  $\mathcal{S}$  (Theorems 3.2 and 3.3). Our results provide a theoretical justification for the principle of minimizing posterior uncertainty in transductive active learning, and indicate that transductive active learning can be more sample efficient than inductive active learning.
- **Applications:** We show that transductive active learning improves upon the state-of-the-art in the batch-wise *active fine-tuning* of neural networks for image classification (Section 4) and in *safe Bayesian optimization* (Section 5).

## 2 Problem Setting

We assume for now that the target space  $\mathcal{A}$  and sample space  $\mathcal{S}$  are finite, and relax these assumptions in the appendices. We model  $f$  as a stochastic process and denote the marginal random variables  $f(\mathbf{x})$  by  $f_{\mathbf{x}}$ , and joint random vectors  $\{f_{\mathbf{x}}\}_{\mathbf{x} \in X}$  for some  $X \subseteq \mathcal{X}$ ,  $|X| < \infty$  by  $\mathbf{f}_X$ . Let  $\mathbf{y}_X$  denote the noisy observations of  $\mathbf{f}_X$ ,  $\{y_{\mathbf{x}} = f_{\mathbf{x}} + \varepsilon_{\mathbf{x}}\}_{\mathbf{x} \in X}$ , where  $\varepsilon_{\mathbf{x}}$  is independent noise.<sup>2</sup> We study the “adaptive” setting, where in round  $n$  the agent selects a point  $\mathbf{x}_n \in \mathcal{S}$  and observes  $y_n = y_{\mathbf{x}_n}$ . The agent’s choice of  $\mathbf{x}_n$  may depend on the outcome of prior observations  $\mathcal{D}_{n-1} \stackrel{\text{def}}{=} \{(\mathbf{x}_i, y_i)\}_{i < n}$ .

**Background on information theory** We briefly recap several important concepts from information theory of which we provide formal definitions in Appendix B. The (differential) entropy  $H[\mathbf{f}]$  is one possible measure of uncertainty about  $\mathbf{f}$  and the conditional entropy  $H[\mathbf{f} | \mathbf{y}]$  is the (expected) posterior uncertainty about  $\mathbf{f}$  after observing  $\mathbf{y}$ . The information gain  $I(\mathbf{f}; \mathbf{y}) = H[\mathbf{f}] - H[\mathbf{f} | \mathbf{y}]$  measures the (expected) reduction in uncertainty about  $\mathbf{f}$  due to  $\mathbf{y}$ . We denote the information gain about  $\mathcal{A}$  from observing  $X$  by  $I(\mathbf{f}_{\mathcal{A}}; \mathbf{y}_X)$ . The maximum information gain about  $\mathcal{A}$  from  $n$  observations within  $\mathcal{S}$  is

$$\gamma_{\mathcal{A}, \mathcal{S}}(n) \stackrel{\text{def}}{=} \max_{\substack{X \subseteq \mathcal{S} \\ |X| \leq n}} I(\mathbf{f}_{\mathcal{A}}; \mathbf{y}_X).$$

This “information capacity” measures the information about  $\mathbf{f}_{\mathcal{A}}$  that is accessible from within  $\mathcal{S}$ , and has been used previously (e.g., by Srinivas et al., 2009; Chowdhury & Gopalan, 2017; Vakili et al., 2021) in the setting where  $\mathcal{X} = \mathcal{A} = \mathcal{S}$ , taking the form of  $\gamma_n \stackrel{\text{def}}{=} \gamma_{\mathcal{X}}(n) \stackrel{\text{def}}{=} \gamma_{\mathcal{X}, \mathcal{X}}(n)$ . We remark that  $\gamma_{\mathcal{A}, \mathcal{S}}(n) \leq \gamma_{\mathcal{S}}(n)$  holds uniformly for all  $\mathcal{A}, \mathcal{S}$ , and  $n$  due to the data processing inequality. Generally,  $\gamma_{\mathcal{A}, \mathcal{S}}(n)$  can be substantially smaller if the target space is a sparse subset of the sample space.

<sup>2</sup> $X$  may be a multiset in which case repeated occurrence of  $\mathbf{x}$  corresponds to independent observations of  $y_{\mathbf{x}}$ .

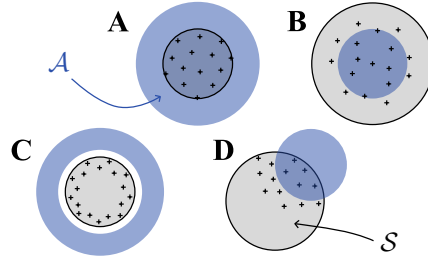


Figure 1: Instances of transductive active learning with target space  $\mathcal{A}$  shown in blue and sample space  $\mathcal{S}$  shown in gray. The points denote plausible observations within  $\mathcal{S}$  to “learn”  $\mathcal{A}$ . In (A), the target space contains “everything” within  $\mathcal{S}$  as well as points outside  $\mathcal{S}$ . In (B, C, D), one makes observations directed towards learning about a particular target. Prior work on inductive active learning has focused on the instance  $\mathcal{A} = \mathcal{S}$ .

### 3 Main Results

We analyze the following principle for transductive active learning:

*Select samples to minimize the posterior “uncertainty” about  $f$  within  $\mathcal{A}$ .* (†)

This principle yields a family of simple and natural decision rules which depend on the chosen measure of “uncertainty”. Two natural measures of uncertainty are (1) the entropy of prediction targets,  $H[\mathbf{f}_{\mathcal{A}}]$ , and (2) their total variance,  $\sum_{\mathbf{x}' \in \mathcal{A}} \text{Var}[f_{\mathbf{x}'}]$ . The corresponding decision rules are

$$(1) \quad \mathbf{x}_n = \arg \min_{\mathbf{x} \in \mathcal{S}} H[\mathbf{f}_{\mathcal{A}} \mid \mathcal{D}_{n-1}, y_{\mathbf{x}}] = \arg \max_{\mathbf{x} \in \mathcal{S}} I(\mathbf{f}_{\mathcal{A}}; y_{\mathbf{x}} \mid \mathcal{D}_{n-1}), \quad (\text{ITL})$$

$$(2) \quad \mathbf{x}_n = \arg \min_{\mathbf{x} \in \mathcal{S}} \text{tr} \text{Var}[\mathbf{f}_{\mathcal{A}} \mid \mathcal{D}_{n-1}, y_{\mathbf{x}}] \quad (\text{VTL})$$

with an implicit expectation over the feedback  $y_{\mathbf{x}}$ . That is, ITL (short for *Information-based Transductive Learning*) and VTL (*Variance-based TL*) select  $\mathbf{x}_n$  so as to minimize the uncertainty about the prediction targets  $\mathbf{f}_{\mathcal{A}}$  (in expectation) after having received the feedback  $y_n$ . Unlike VTL, ITL takes into account the mutual dependence between points in  $\mathcal{A}$ . These decision rules were suggested previously (MacKay, 1992; Seo et al., 2000; Yu et al., 2006) without deriving theoretical guarantees; and they generalize several widely used algorithms which we discuss in more detail in Section 6. Most prominently, in the inductive setting where  $\mathcal{S} \subseteq \mathcal{A}$ , ITL reduces to  $\mathbf{x}_n = \arg \max_{\mathbf{x} \in \mathcal{S}} I(f_{\mathbf{x}}; y_{\mathbf{x}} \mid \mathcal{D}_{n-1})$ , i.e., is “undirected” and reduces to standard uncertainty-based active learning strategies (cf. Appendix C.1). The convergence properties for the special instance of ITL with  $\mathcal{S} = \mathcal{A}$  have been studied extensively. To the best of our knowledge, we are the first to extend these guarantees to the more general setting of transductive active learning.

In our presented results, we make the following assumption.

**Assumption 3.1.** In the case of ITL, the information gain  $\psi_{\mathcal{A}}(X) = I(\mathbf{f}_{\mathcal{A}}; \mathbf{y}_X)$  is submodular. In the case of VTL, the variance reduction  $\psi_{\mathcal{A}}(X) = \text{tr} \text{Var}[\mathbf{f}_{\mathcal{A}}] - \text{tr} \text{Var}[\mathbf{f}_{\mathcal{A}} \mid \mathbf{y}_X]$  is submodular.

Under this assumption,  $\psi_{\mathcal{A}}(\mathbf{x}_{1:n})$  is a constant factor approximation of  $\max_{X \subseteq \mathcal{S}, |X| \leq n} \psi_{\mathcal{A}}(X)$  due to the seminal result on submodular function maximization by Nemhauser et al. (1978). Similar assumptions have been made, e.g., by Bogunovic et al. (2016) and Kothawade et al. (2021). Assumption 3.1 is satisfied exactly for ITL when  $\mathcal{S} \subseteq \mathcal{A}$  and  $f$  is a Gaussian process (cf. Lemma C.9), and we provide an extensive discussion of our results in Appendix C.4 for instances where Assumption 3.1 is satisfied approximately, relying on the notion of weak submodularity (Das & Kempe, 2018).

#### 3.1 Gaussian Process Setting

When  $f \sim \mathcal{GP}(\mu, k)$  is a Gaussian process (GP, Williams & Rasmussen, 2006) with known mean function  $\mu$  and kernel  $k$ , and the noise  $\varepsilon_{\mathbf{x}}$  is mutually independent and zero-mean Gaussian with known variance, the ITL and VTL objectives have a closed form expression (cf. Appendix F) and can be optimized efficiently. Further, the information capacity  $\gamma_n$  is sublinear in  $n$  for a rich class of GPs (Srinivas et al., 2009; Vakili et al., 2021), with rates summarized in Table 3 of the appendix.

**Convergence to irreducible uncertainty** So far, our discussion was centered around the role of the target space  $\mathcal{A}$  in facilitating *directed* learning. An orthogonal contribution of this work is to study *extrapolation* from the sample space  $\mathcal{S}$  to points  $\mathbf{x} \in \mathcal{A} \setminus \mathcal{S}$ . To this end, we derive bounds on the marginal posterior variance  $\sigma_n^2(\mathbf{x}) \stackrel{\text{def}}{=} \text{Var}[f(\mathbf{x}) \mid \mathcal{D}_n]$  for points in  $\mathcal{A}$ . These bounds depend on the instance of transductive active learning (i.e.,  $\mathcal{A}$  and  $\mathcal{S}$ ) and might be of independent interest for active learning. For ITL and VTL, they imply uniform convergence of the variance for a rich class of GPs. To the best of our knowledge, this work is the first to present such bounds.

We define the *irreducible uncertainty* as the variance of  $f(\mathbf{x})$  provided complete knowledge of  $f$  in  $\mathcal{S}$ :

$$\eta_{\mathcal{S}}^2(\mathbf{x}) \stackrel{\text{def}}{=} \text{Var}[f_{\mathbf{x}} \mid \mathbf{f}_{\mathcal{S}}].$$

As the name suggests,  $\eta_{\mathcal{S}}^2(\mathbf{x})$  represents the smallest uncertainty one can hope to achieve from observing only within  $\mathcal{S}$ . For all  $\mathbf{x} \in \mathcal{S}$ , it is easy to see that  $\eta_{\mathcal{S}}^2(\mathbf{x}) = 0$ . However, the irreducible uncertainty of  $\mathbf{x} \notin \mathcal{S}$  may be (and typically is!) strictly positive.

**Theorem 3.2** (Bound on marginal variance for ITL and VTL). *Let Assumption 3.1 hold and the data be selected by either ITL or VTL.. Assume that  $f \sim \mathcal{GP}(\mu, k)$  with known mean function  $\mu$*

and kernel  $k$ , the noise  $\varepsilon_{\mathbf{x}}$  is mutually independent and zero-mean Gaussian with known variance, and  $\gamma_n$  is sublinear in  $n$ . Then there exists a constant  $C$  such that for any  $n \geq 1$  and  $\mathbf{x} \in \mathcal{A}$ ,

$$\sigma_n^2(\mathbf{x}) \leq \underbrace{\eta_{\mathcal{S}}^2(\mathbf{x})}_{\text{irreducible}} + C \underbrace{\frac{\gamma_{\mathcal{A},\mathcal{S}}(n)}{\sqrt{n}}}_{\text{reducible}}. \quad (1)$$

Moreover, if  $\mathbf{x} \in \mathcal{A} \cap \mathcal{S}$ , there exists a constant  $C'$  such that

$$\sigma_n^2(\mathbf{x}) \leq C' \frac{\gamma_{\mathcal{A},\mathcal{S}}(n)}{n}. \quad (2)$$

Intuitively, Equation (1) of Theorem 3.2 can be understood as bounding an epistemic ‘‘generalization gap’’ (Wainwright, 2019) of the learner. The reducible uncertainty converges to zero at all prediction targets  $\mathbf{x} \in \mathcal{A}$ , e.g., for linear, Gaussian, and smooth Matérn kernels. As to be expected, a smaller target space (i.e., more targeted sampling) leads to faster convergence due to a smaller information capacity  $\gamma_{\mathcal{A},\mathcal{S}}(n) \ll \gamma_n$ . Equation (2) matches prior results for the setting  $\mathcal{S} = \mathcal{A}$ . We provide a formal proof of Theorem 3.2 in Appendix C.6.

### 3.2 Agnostic Setting

The result from the GP setting translates also to the agnostic setting, where the ‘‘ground truth’’  $f^*$  may be any sufficiently regular fixed function on  $\mathcal{X}$ .<sup>3</sup> In this case, we use the model  $f$  from Section 3.1 as a (misspecified) model of  $f^*$ , with some kernel  $k$  and zero mean function  $\mu(\cdot) = 0$ . We denote by  $\mu_n(\mathbf{x}) = \mathbb{E}[f(\mathbf{x}) \mid \mathcal{D}_n]$  the posterior mean of  $f$ . W.l.o.g. we assume in the following result that the prior variance is bounded, i.e.,  $\text{Var}[f(\mathbf{x})] \leq 1$ .

**Theorem 3.3** (Bound on approximation error for ITL and VTL, following Abbasi-Yadkori (2013); Chowdhury & Gopalan (2017)). *Let Assumption 3.1 hold and the data be selected by either ITL or VTL. Pick any  $\delta \in (0, 1)$ . Assume that  $f^*$  lies in the reproducing kernel Hilbert space  $\mathcal{H}_k(\mathcal{X})$  of the kernel  $k$  with norm  $\|f^*\|_k < \infty$ , the noise  $\varepsilon_n$  is conditionally  $\rho$ -sub-Gaussian, and  $\gamma_n$  is sublinear in  $n$ . Let  $\beta_n(\delta) = \|f^*\|_k + \rho\sqrt{2(\gamma_n + 1 + \log(1/\delta))}$ . Then for any  $n \geq 1$  and  $\mathbf{x} \in \mathcal{A}$ , jointly with probability at least  $1 - \delta$ ,*

$$|f^*(\mathbf{x}) - \mu_n(\mathbf{x})| \leq \beta_n(\delta) \left[ \underbrace{\eta_{\mathcal{S}}(\mathbf{x})}_{\text{irreducible}} + \underbrace{\nu_{\mathcal{A},\mathcal{S}}(n)}_{\text{reducible}} \right]$$

where  $\nu_{\mathcal{A},\mathcal{S}}^2(n)$  denotes the reducible part of Equation (1).

We provide a formal proof of Theorem 3.3 in Appendix C.7. Theorem 3.3 generalizes approximation error bounds of prior works to the extrapolation setting, where some prediction targets  $\mathbf{x} \in \mathcal{A}$  lie outside the sample space  $\mathcal{S}$ . For prediction targets  $\mathbf{x} \in \mathcal{A} \cap \mathcal{S}$ , the irreducible uncertainty vanishes, and we recover previous results from the setting  $\mathcal{S} = \mathcal{A}$ .

Theorems 3.2 and 3.3 show that ITL and VTL efficiently learn  $f$  at the prediction targets  $\mathcal{A}$  for large classes of ‘‘sufficiently regular’’ functions  $f$ . In the following, we validate these results experimentally by showing that ITL and VTL exhibit strong empirical performance in a broad range of applications.

### 3.3 Experiments in the Gaussian Process Setting

Before demonstrating ITL and VTL on GPs to develop more intuition, we introduce a natural correlation-based baseline, which will later uncover connections to existing approaches:

$$\mathbf{x}_n = \arg \max_{\mathbf{x} \in \mathcal{S}} \sum_{\mathbf{x}' \in \mathcal{A}} \text{Cor}[f_{\mathbf{x}}, f_{\mathbf{x}'} \mid \mathcal{D}_{n-1}]. \quad (\text{CTL})$$

**How does the smoothness of  $f$  affect ITL?** We contrast two ‘‘extreme’’ kernels: the *Gaussian kernel*  $k(\mathbf{x}, \mathbf{x}') = \exp(-\|\mathbf{x} - \mathbf{x}'\|_2^2/2)$  and the *Laplace kernel*  $k(\mathbf{x}, \mathbf{x}') = \exp(-\|\mathbf{x} - \mathbf{x}'\|_1)$ . In the mean-squared sense, the Gaussian kernel yields a smooth process  $f$  whereas the Laplace kernel yields a continuous but non-differentiable  $f$  (Williams & Rasmussen, 2006). Figure 2 shows

<sup>3</sup>Here  $f^*(\mathbf{x})$  denotes the mean observation  $y_{\mathbf{x}} = f^*(\mathbf{x}) + \varepsilon_{\mathbf{x}}$

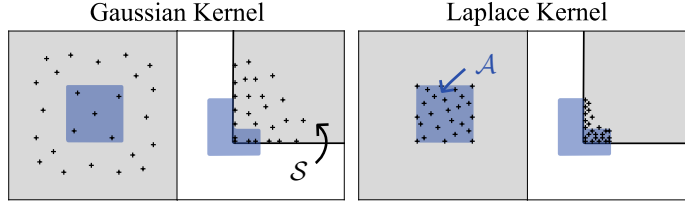


Figure 2: Initial 25 samples of ITL under a Gaussian kernel with lengthscale 1 (left) and a Laplace kernel with lengthscale 10 (right). Shown in gray is the sample space  $\mathcal{S}$  and shown in blue is the target space  $\mathcal{A}$ . In three of the four examples, points outside the target space provide useful information.

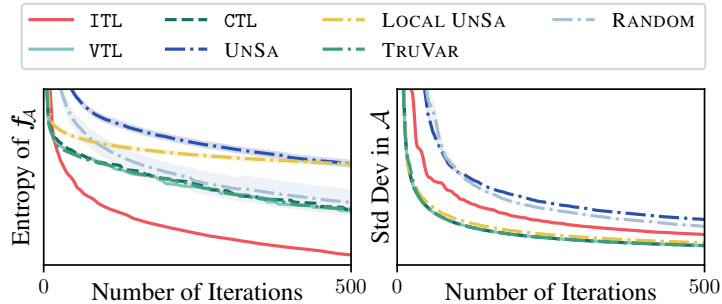


Figure 3: Entropy of  $f_{\mathcal{A}}$  ranging from  $-3850$  to  $-3725$  and the mean marginal standard deviations of  $f_{\mathcal{A}}$  ranging from 0 to 0.15. Experiment is using the Gaussian kernel of the left instance ( $\mathcal{A} \subset \mathcal{S}$ ) from Figure 2. It can be seen that ITL and VTL outperform UNSA and RANDOM. Uncertainty bands correspond to one standard error over 10 random seeds.

how ITL adapts to the smoothness of  $f$ : Under the “smooth” Gaussian kernel, points outside  $\mathcal{A}$  provide higher-order information. In contrast, under the “rough” Laplace kernel and if  $\mathcal{A} \subseteq \mathcal{S}$ , points outside  $\mathcal{A}$  do not provide any additional information, and therefore are not sampled by ITL. If, however,  $\mathcal{A} \not\subseteq \mathcal{S}$ , information “leaks”  $\mathcal{A}$  even under a Laplace kernel prior. That is, even for non-smooth functions, the point with most information need not be in  $\mathcal{A}$ .

**Does ITL outperform uncertainty sampling?** Uncertainty sampling (UNSA, Lewis & Catlett, 1994) is one of the most popular active learning methods. UNSA selects points  $\mathbf{x}$  with high *prior* uncertainty:  $\mathbf{x}_n = \arg \max_{\mathbf{x} \in \mathcal{S}} \sigma_{n-1}^2(\mathbf{x})$ . This is in stark contrast to ITL and VTL which select points  $\mathbf{x}$  that minimize *posterior* (epistemic) uncertainty about  $\mathcal{A}$ . It can be seen that UNSA is the special “undirected” case of ITL when  $\mathcal{S} \subseteq \mathcal{A}$  and observation noise is homoscedastic (cf. Appendix C.1).

We compare UNSA to ITL, VTL, and CTL in Figure 3. We observe that ITL and VTL outperform UNSA which also samples points that are not informative about  $\mathcal{A}$ . Further, ITL and VTL outperform “local” UNSA (i.e., UNSA constrained to  $\mathcal{A} \cap \mathcal{S}$ ) which neglects all information provided by points outside  $\mathcal{A}$ .<sup>4</sup> As one would expect, VTL has an advantage with respect to reducing the total variance of  $f_{\mathcal{A}}$ , whereas ITL reduces the entropy of  $f_{\mathcal{A}}$  faster. We include ablations in Appendix H where we, in particular, observe that the advantage of ITL and VTL over UNSA increases as the volume of prediction targets shrinks in comparison to the size of domain.

## 4 Active Fine-Tuning of Neural Networks

Fine-tuning a large pre-trained model is a cost- and computation-effective approach to improve performance on a given target domain (Lee et al., 2022). While previous work has studied the effectiveness of various training procedures for fine-tuning (e.g., Eustratiadis et al., 2024), we ask: *How can we select the right data for fine-tuning to a specific task?* This *active* fine-tuning problem is an instance of the introduced “directed” transductive learning problem: Concretely, consider a supervised setting, where the function  $f$  maps inputs  $\mathbf{x} \in \mathcal{X}$  to outputs  $y \in \mathcal{Y}$ . We have access

<sup>4</sup>If  $\mathcal{A} \not\subseteq \mathcal{S}$  then “local” UNSA does *not even* converge to the irreducible uncertainty.

to noisy samples from a training set  $\mathcal{S}$  on  $\mathcal{X}$ , and we would like to learn  $f$  such that our estimate minimizes a given risk measure, such as classification error, with respect to a test distribution  $\mathcal{P}_A$  on  $\mathcal{X}$ . The goal is to actively and efficiently sample from  $\mathcal{S}$  to minimize risk with respect to  $\mathcal{P}_A$ .<sup>5</sup> We show in this section that ITL and VTL can learn  $f$  from only *few examples* from  $\mathcal{S}$ .

**How can we leverage the latent structure learned by the pre-trained model?** As common in related work, we approximate the (pre-trained) neural network (NN)  $f(\cdot; \theta)$  as a linear function in a latent embedding space,  $f(\mathbf{x}; \theta) \approx \beta^\top \phi_\theta(\mathbf{x})$ , with weights  $\beta \in \mathbb{R}^p$  and embeddings  $\phi_\theta: \mathcal{X} \rightarrow \mathbb{R}^p$ . Common choices of embeddings include last-layer embeddings (Devlin et al., 2019; Holzmüller et al., 2023), neural tangent embeddings arising from neural tangent kernels (Jacot et al., 2018) which are motivated by their relationship to the training and fine-tuning of ultra-wide NNs (Arora et al., 2019; Lee et al., 2019; Khan et al., 2019; He et al., 2020; Malladi et al., 2023), and loss gradient embeddings (Ash et al., 2020). We provide a comprehensive overview of embeddings in Appendix J.2. Now, supposing the prior  $\beta \sim \mathcal{N}(\mathbf{0}, \Sigma)$ , often with  $\Sigma = \mathbf{I}$  (Khan et al., 2019; He et al., 2020; Antorán et al., 2022; Wei et al., 2022), this approximation of  $f$  is a Gaussian process with kernel  $k(\mathbf{x}, \mathbf{x}') = \phi_\theta(\mathbf{x})^\top \Sigma \phi_\theta(\mathbf{x}')$  which quantifies the similarity between points in terms of their alignment in the learned latent space. Note that the correlation  $k(\mathbf{x}, \mathbf{x}') / \sqrt{k(\mathbf{x}, \mathbf{x})k(\mathbf{x}', \mathbf{x}')}$  between two points  $\mathbf{x}, \mathbf{x}'$  is equal to the cosine similarity of their embeddings.

In this context, Theorem 3.2 bounds the epistemic posterior uncertainty about a prediction using the approximation  $\beta^\top \phi_\theta(\mathbf{x})$ , given that the model is trained using data selected by ITL or VTL. Theorem 3.3 bounds the generalization error when using the posterior mean of  $\beta$  for prediction. This extends recent work which has studied estimators of this generalization error (Wei et al., 2022).

**Batch selection: Diversity via conditional embeddings** Efficient labeling and training necessitates a batch-wise selection of inputs. The selection of a batch of size  $b > 1$  can be seen as an individual *non-adaptive* active learning problem, and significant recent work has shown that batch diversity is crucial in this setting (Ash et al., 2020; Zanette et al., 2021; Holzmüller et al., 2023; Pacchiano et al., 2024). An information-based batch-wise selection strategy is formalized by the following non-adaptive transductive active learning problem (Chen & Krause, 2013) and the greedy approximation of  $B_n$  by ITL which selects elements  $\mathbf{x}_{n,i}$  of the  $n$ -th batch iteratively based on  $\mathbf{x}_{n,1:i-1}$ :

$$B_n = \arg \max_{B \subseteq \mathcal{S}, |B|=b} I(\mathbf{f}_A; \mathbf{y}_B \mid \mathcal{D}_{n-1}); \quad \mathbf{x}_{n,i} = \arg \max_{\mathbf{x} \in \mathcal{S}} I(\mathbf{f}_A; y_{\mathbf{x}} \mid \mathcal{D}_{n-1}, \mathbf{y}_{\mathbf{x}_{n,1:i-1}}). \quad (3)$$

The batch  $B_n$  is diverse and informative by design. We show that under Assumption 3.1,  $B'_n = \mathbf{x}_{n,1:b}$  yields a constant-factor approximation of  $B_n$  (cf. Appendix C.3).

## 4.1 Experiments on Active Fine-Tuning

Our empirical evaluation is motivated by the following practical example: We deploy a pre-trained image classifier to user’s phones who use it within their local environment. We would like to locally fine-tune a user’s model to their environment. Since the users’ images  $\mathcal{A}$  are unlabeled, this requires selecting a small number of relevant and diverse images from the set of labeled images  $\mathcal{S}$ . As such, we will focus here on the setting where the points in our test set do not lie in our training set (i.e.,  $\mathcal{A} \cap \mathcal{S} = \emptyset$ ), and discuss alternative instances such as active domain adaptation in Appendix I.

**Testbeds & architectures** We use the MNIST (LeCun et al., 1998) and CIFAR-100 (Krizhevsky et al., 2009) datasets as testbeds. In both cases, we take  $\mathcal{S}$  to be the training set, and we consider the task of learning the digits 3, 6, and 9 (MNIST) or the first 10 categories of CIFAR-100.<sup>6</sup> For MNIST, we train a simple convolutional neural network with ReLU activations, three convolutional layers with max-pooling, and two fully-connected layers. For CIFAR-100, we fine-tune an EfficientNet-B0 (Tan & Le, 2019) pre-trained on ImageNet (Deng et al., 2009), augmented by a final fully-connected layer. We train the NNs using the cross-entropy loss and the ADAM optimizer (Kingma & Ba, 2014).

**Results** In Figure 4, We compare against (i) active learning methods which largely aim for sample diversity but which are not directed towards the target distribution  $\mathcal{P}_A$  (e.g., BADGE; Ash et al., 2020), and (ii) search methods that aim to retrieve the most relevant samples from  $\mathcal{S}$  with respect to the targets  $\mathcal{P}_A$  (e.g., maximizing cosine similarity to target embeddings as is common in vector databases;

<sup>5</sup>The setting with target distributions  $\mathcal{P}_A$  can be reduced to considering target sets  $\mathcal{A}$  (cf. Appendix E).

<sup>6</sup>That is, we restrict  $\mathcal{P}_A$  to the support of points with labels  $\{3, 6, 9\}$  (MNIST) or labels  $\{0, \dots, 9\}$  (CIFAR-100) and train a neural network using few examples drawn from the training set  $\mathcal{S}$ .

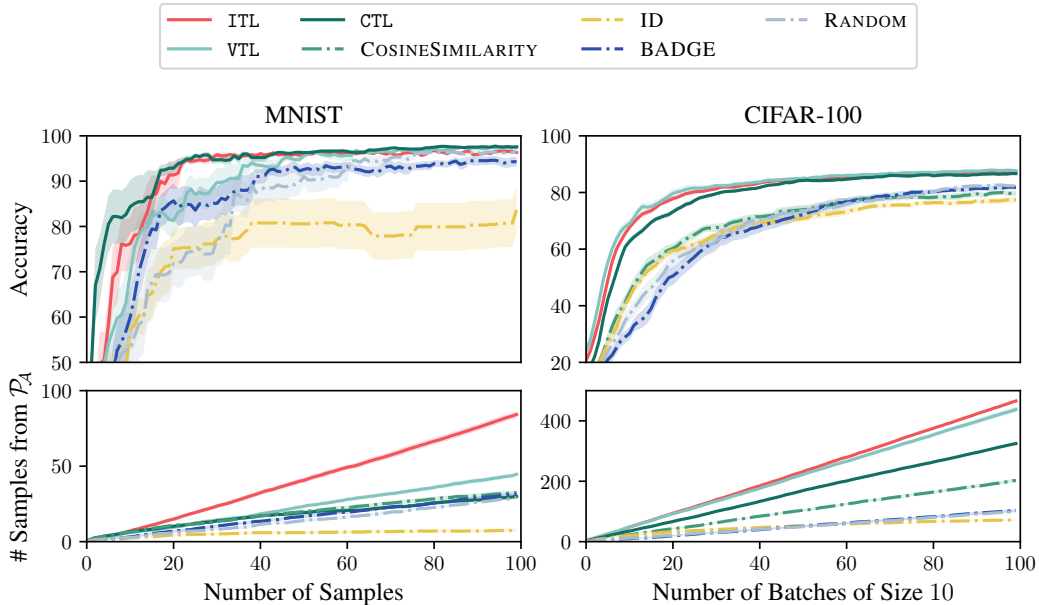


Figure 4: Active fine-tuning on MNIST (left) and CIFAR-100 (right). RANDOM selects each observation uniformly at random from  $\mathcal{S}$ . The batch size is 1 for MNIST and 10 for CIFAR-100. Uncertainty bands correspond to one standard error over 10 random seeds. We see that transductive active learning with ITL and VTL significantly outperforms competing methods, and in particular, retrieves substantially more samples from the support of  $\mathcal{P}_A$ . See Appendix J for details and ablations.

Settles & Craven, 2008; Johnson et al., 2019). INFORMATIONDENSITY (ID, Settles & Craven, 2008) is a heuristic approach aiming to combine (i) diversity and (ii) relevance. In Appendix J.5, we also compare against a wide range of additional baselines (e.g., CORESET (Sener & Savarese, 2017), TYPICLUST (Hacohen et al., 2022), PROBCOVER (Yehuda et al., 2022), etc.) that fall into one of the categories (i) and (ii), and which perform similar to the baselines listed here.

We observe that ITL, VTL, and CTL consistently and significantly outperform random sampling from  $\mathcal{S}$  as well as all baselines. We see that relevance-based methods such as COSINESIMILARITY have an initial advantage over RANDOM but for batch sizes larger than 1 they quickly fall behind due to diminishing informativeness of the selected data. In contrast, diversity-based methods such as BADGE are more competitive with RANDOM but do not explicitly aim to retrieve relevant samples.

Remarkably, transductive active learning outperforms random data selection even in the MNIST experiment where the model is randomly initialized. This suggests that the learned embeddings can be informative for data selection even in the early stages of training, bootstrapping the learning progress.

**Balancing sample relevance and diversity** Our proposed methods unify approaches to coverage (promoting *diverse* samples) and search (aiming for *relevant* samples with respect to a given query  $\mathcal{A}$ ) which leads to the significant improvement upon the state-of-the-art in Figure 4. Notably, for a batch size and query size of 1 and if correlations are non-negative, ITL, VTL, CTL, and the canonical cosine similarity are equivalent. CTL can be seen as a direct generalization of cosine similarity-based retrieval to batch and query sizes larger than one. In contrast to CTL, ITL and VTL may also sample points which exhibit a strong negative correlation (which is also informative).

We observe empirically that ITL obtains samples from  $\mathcal{P}_A$  at more than twice the rate of COSINESIMILARITY, which translates to a significant improvement in accuracy in more difficult learning tasks, while requiring fewer (labeled) samples from  $\mathcal{S}$ . This phenomenon manifests for both MNIST and CIFAR-100, as well as imbalanced datasets  $\mathcal{S}$  or imbalanced reference samples from  $\mathcal{P}_A$  (cf. Appendix J.6). The improvement in accuracy appears to increase in the large-data regime, where the learning tasks become more difficult. Akin to a previously identified scaling trend with size of the pre-training dataset (Tamkin et al., 2022), this suggests a potential scaling trend where the improvement of ITL over random batch selection grows as models are fine-tuned on a larger pool of data.

**Towards task-driven few-shot learning** Being able to efficiently and automatically select data may allow dynamic few-shot fine-tuning to individual tasks (Vinyals et al., 2016; Hardt & Sun, 2024), e.g., fine-tuning the model to each test point / query / prompt. Such task-driven few-shot learning can be seen as a form of “memory recall” akin to associative memory (Hopfield, 1982). Our results are a first indication that task-driven learning can lead to substantial performance gains, and we believe that this is a promising direction for future studies.

## 5 Safe Bayesian Optimization

Another practical problem that can be cast as “directed” learning is safe Bayesian optimization (Safe BO, Sui et al., 2015; Berkenkamp et al., 2021) which has applications in natural science (Cooper & Netoff, 2022) and robotics (Wischnewski et al., 2019; Sukhija et al., 2023; Widmer et al., 2023). Safe BO solves the following optimization problem

$$\max_{\mathbf{x} \in \mathcal{S}^*} f^*(\mathbf{x}) \quad \text{where} \quad \mathcal{S}^* = \{\mathbf{x} \in \mathcal{X} \mid g^*(\mathbf{x}) \geq 0\} \quad (4)$$

which can be generalized to multiple constraints. The functions  $f^*$  and  $g^*$ , and hence also the “safe set”  $\mathcal{S}^*$ , are unknown and have to be actively learned from data. However, it is crucial that the data collection does not violate the constraint, i.e.,  $\mathbf{x}_n \in \mathcal{S}^*, \forall n \geq 1$ .

**Safe Bayesian optimization as Transductive Active Learning** In the agnostic setting from Section 3.2, GPs  $f$  and  $g$  can be used as well-calibrated models of the ground truths  $f^*$  and  $g^*$ , and we denote lower- and upper-confidence bounds by  $l_n^f(\mathbf{x}), l_n^g(\mathbf{x})$  and  $u_n^f(\mathbf{x}), u_n^g(\mathbf{x})$ , respectively. These confidence bounds induce a *pessimistic* safe set  $\mathcal{S}_n \stackrel{\text{def}}{=} \{\mathbf{x} \mid l_n^g(\mathbf{x}) \geq 0\}$  and an *optimistic* safe set  $\hat{\mathcal{S}}_n \stackrel{\text{def}}{=} \{\mathbf{x} \mid u_n^g(\mathbf{x}) \geq 0\}$  which satisfy  $\mathcal{S}_n \subseteq \mathcal{S}^* \subseteq \hat{\mathcal{S}}_n$  with high probability at all times. Similarly, the set of *potential maximizers*

$$\mathcal{A}_n \stackrel{\text{def}}{=} \{\mathbf{x} \in \hat{\mathcal{S}}_n \mid u_n^f(\mathbf{x}) \geq \max_{\mathbf{x}' \in \mathcal{S}_n} l_n^f(\mathbf{x}')\} \quad (5)$$

contains the solution to Equation (4) at all times with high probability.

The (simple) regret  $r_n(\mathcal{S}) \stackrel{\text{def}}{=} \max_{\mathbf{x} \in \mathcal{S}} f^*(\mathbf{x}) - f^*(\hat{\mathbf{x}}_n)$  with  $\hat{\mathbf{x}}_n \stackrel{\text{def}}{=} \arg \max_{\mathbf{x} \in \mathcal{S}_n} l_n^f(\mathbf{x})$  measures the worst-case performance of a decision rule. To achieve small regret, one faces an *exploration-expansion* dilemma wherein one needs to explore points that are known-to-be-safe, i.e., lie in the estimated safe set  $\mathcal{S}_n$ , and might be optimal, while at the same time discovering new safe points by “expanding”  $\mathcal{S}_n$ . Accordingly, a natural choice for the target space of Safe BO is  $\mathcal{A}_n$  since it captures both exploration and expansion *simultaneously*.<sup>7</sup> To prevent constraint violation, the sample space is restricted to the pessimistic safe set  $\mathcal{S}_n$ . In Safe BO, both the target space and sample space change with each round  $n$ , and we generalize our theoretical results from Section 3 in Appendix C to this setting.

**Theorem 5.1** (Convergence to safe optimum). *Pick any  $\epsilon > 0, \delta \in (0, 1)$ . Assume that  $f^*, g^*$  lie in the reproducing kernel Hilbert space  $\mathcal{H}_k(\mathcal{X})$  of the kernel  $k$ , and that the noise  $\varepsilon_n$  is conditionally  $\rho$ -sub-Gaussian. Then, we have with probability at least  $1 - \delta$ ,*

*Safety: for all  $n \geq 1, \mathbf{x}_n \in \mathcal{S}^*$ .*

*Moreover, assume  $\mathcal{S}_0 \neq \emptyset$  and denote with  $\mathcal{R}$  the largest reachable safe set starting from  $\mathcal{S}_0$ . Then, the convergence of reducible uncertainty implies that there exists  $n^* > 0$  such that with probability at least  $1 - \delta$ ,*

*Optimality: for all  $n \geq n^*, r_n(\mathcal{R}) \leq \epsilon$ .*

We provide a formal proof in Appendix C.8. Central to the proof is the application of Theorem 3.3 to show that the safety of parameters *outside* the safe set  $\mathcal{S}_n$  can be inferred efficiently. In Section 3, we outline settings where the reducible uncertainty converges which is the case for a very general class of functions, and for such instances Theorem 5.1 guarantees optimality in the largest reachable safe set  $\mathcal{R}$ .  $\mathcal{R}$  represents the largest set any safe learning algorithm can explore without violating the safety constraints (with high probability) during learning (cf. Definition C.29). Our guarantees are similar to those of other Safe BO algorithms (Berkenkamp et al., 2021) but require fewer assumptions and generalize to continuous domains. We obtain Theorem 5.1 from a more general result (Theorem C.34) which can be specialized to yield “free” novel convergence guarantees for problems other than Bayesian optimization, such as level set estimation, by choosing an appropriate target space.

<sup>7</sup>An alternative possibility is to weigh each point in  $\mathcal{A}_n$  according to how likely it is to be the safe optimum. Which approach performs better is task-dependent, and we include a detailed discussion in Appendix K.1.



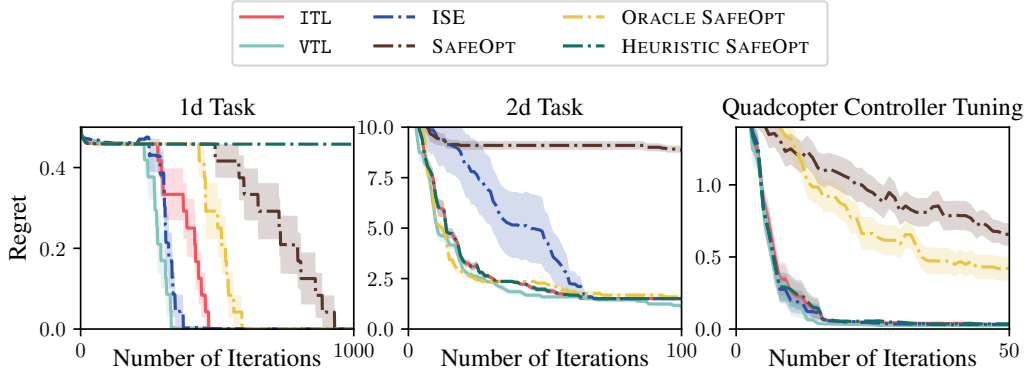


Figure 5: We compare ITL and VTL to ORACLE SAFEOPT, which has oracle knowledge of the Lipschitz constants, SAFEOPT, where the Lipschitz constants are estimated from the GP, as well as HEURISTIC SAFEOPT and ISE, and observe that ITL and VTL systematically perform well. We compare against additional baselines in Appendix K.1. The regret is evaluated with respect to the ground truth objective  $f^*$  and constraint  $g^*$ , and averaged over 10 (in synthetic experiments) and 25 (in the quadcopter experiment) random seeds. Additional details can be found in Appendix K.4.

## 5.1 Experiments on Safe Bayesian Optimization

We evaluate two synthetic experiments for a 1d and 2d parameter space, respectively (cf. Appendix K.4 for details), which demonstrate the various shortcomings of existing Safe BO baselines. Additionally, as third experiment, we safely tune the controller of a quadcopter.

**Safe controller tuning for a quadcopter** We consider a quadcopter with unknown dynamics;  $s_{t+1} = T(s_t, u_t)$  where  $u_t \in \mathbb{R}^{d_u}$  is the control signal and  $s_t \in \mathbb{R}^{d_s}$  is the state at time  $t$ . The inputs  $u_t$  are calculated through a deterministic function of the state  $\pi : \mathcal{S} \rightarrow \mathcal{U}$  which we call the policy. The policy is parameterized via parameters  $x \in \mathcal{X}$ , e.g., PID controller gains, such that  $u_t = \pi_x(s_t)$ . The goal is to find the optimal parameters with respect to an unknown objective  $f^*$  while satisfying some unknown constraint(s)  $g^*(x) \geq 0$ , e.g., the quadcopter does not fall on the ground. This is a typical Safe BO problem which is widely applied for safe controller learning in robotics (Berkenkamp et al., 2021; Baumann et al., 2021; Widmer et al., 2023).

**Results** We compare ITL and VTL to SAFEOPT (Berkenkamp et al., 2021), which is undirected, i.e., expands in all directions including ones that are known-to-be suboptimal, and ISE (Bottero et al., 2022), which is solely expansionist — does not trade-off expansion-exploration. We provide a detailed discussion of baselines in Appendix K.2. In all our experiments, summarized in Figure 5, we observe that ITL and VTL systematically perform well, i.e., better or on par with the state-of-the-art. We attribute this to its directed exploration and less conservative expansion over SAFEOPT (cf. 1d task and quadcopter experiment), and natural trade-off between expansion and exploration as opposed to ISE (see 2d task). Generally, VTL has a slight advantage over ITL, which is because VTL minimizes marginal variances (as opposed to entropy), which are decisive for expanding the safe set. While ITL and VTL do not violate constraints, we observe that other methods that do not explicitly enforce safety such as EIC (Gardner et al., 2014) lead to constraint violation (cf. Appendix K.4.2).

## 6 Related Work

**(Inductive) active learning** The special case of transductive active learning where  $\mathcal{A} = \mathcal{S} = \mathcal{X}$  has been widely studied. We refer to this special instance as *inductive* active learning, since the goal is to extract as much information as possible as opposed to making predictions on a specific target set.

Several works have previously found entropy-based decision rules to be useful for inductive active learning (Krause & Guestrin, 2007; Guo & Greiner, 2007; Krause et al., 2008) and semi-supervised learning (Grandvalet & Bengio, 2004). The variance-based VTL has previously been proposed by Cohn (1993) in the special case of inductive active learning without proving theoretical guarantees. VTL was then recently re-derived by Shoham & Avron (2023) along other experimental design

criteria under the lens of minimizing risk for inductive one-shot learning in overparameterized models. Substantial work on active learning has studied entropy-based criteria in *parameter-space*, most notably BALD (MacKay, 1992; Houlby et al., 2011; Gal et al., 2017; Kirsch et al., 2019), which selects  $\mathbf{x}_n = \arg \max_{\mathbf{x} \in \mathcal{X}} I(\boldsymbol{\theta}; y_{\mathbf{x}} | \mathcal{D}_{n-1})$ , where  $\boldsymbol{\theta}$  is the random parameter vector of a parametric model (e.g., obtained via Bayesian deep learning). Such methods are inherently inductive in the sense that they do not facilitate learning on specific prediction targets.

**Transductive active learning** In contrast, ITL operates in *output-space* where it is straightforward to specify prediction targets, and which is computationally easier. Special cases of ITL when  $\mathcal{S} = \mathcal{X}$  and  $|\mathcal{A}| = 1$  have been proposed in the foundational work of MacKay (1992) on “directed” output-space active learning. As generalization to larger target spaces, MacKay (1992) proposed mean-marginal ITL,

$$\mathbf{x}_n = \arg \max_{\mathbf{x} \in \mathcal{S}} \sum_{\mathbf{x}' \in \mathcal{A}} I(f_{\mathbf{x}'}; y_{\mathbf{x}} | \mathcal{D}_{n-1}), \quad (\text{MM-ITL})$$

for which we derive analogous versions of Theorems 3.2 and 3.3 in Appendix D.3. We note that similarly to VTL, MM-ITL disregards the mutual dependence of points in the target space  $\mathcal{A}$  and differs from VTL only in a different weighting of the posterior marginal variances of the prediction targets (cf. Appendix D.3). Recently, Bickford Smith et al. (2023) generalized MM-ITL by treating the prediction target as a random variable, and Kothawade et al. (2021) and Bickford Smith et al. (2024) demonstrated the use of output-space decision rules for image classification tasks in a pre-training context.

**Influence functions** measure the change in a model’s prediction when a single data point is removed from the training data (Cook, 1977; Koh & Liang, 2017; Pruthi et al., 2019). Influence functions have been used for data selection in settings closely related to the transductive active fine-tuning of neural networks proposed in this work (Xia et al., 2024). They select data that reduces a first-order Taylor approximation to the test loss after fine-tuning a neural network, which corresponds to maximizing cosine similarity to the prediction targets in a loss-gradient embedding space. We show in our experiments that transductive active learning can substantially outperform COSINESIMILARITY. We attribute this primarily to influence functions implicitly assuming that the influence of selected data adds linearly (i.e., two equally scored data points are expected to doubly improve the model performance, Xu & Kazantsev, 2019, Section 3.2). This assumption does not hold in practice as seen, e.g., by simply duplicating data. The same limitation applies to the related approach of datamodels (Ilyas et al., 2022).

**Other work on directed active learning** Directed active learning methods have been proposed for the problem of determining the optimum of an unknown function, also known as best-arm identification (Audibert et al., 2010) or pure exploration bandits (Bubeck et al., 2009). Entropy search methods (Hennig & Schuler, 2012; Hernández-Lobato et al., 2014) are widely used and select  $\mathbf{x}_n = \arg \max_{\mathbf{x} \in \mathcal{X}} I(\mathbf{x}^*; y_{\mathbf{x}} | \mathcal{D}_{n-1})$  in *input-space* where  $\mathbf{x}^* = \arg \max_{\mathbf{x}} f_{\mathbf{x}}$ . Similarly to ITL, *output-space* entropy search methods (Hoffman & Ghahramani, 2015; Wang & Jegelka, 2017), which select  $\mathbf{x}_n = \arg \max_{\mathbf{x} \in \mathcal{X}} I(f^*; y_{\mathbf{x}} | \mathcal{D}_{n-1})$  with  $f^* = \max_{\mathbf{x}} f_{\mathbf{x}}$ , are more computationally tractable. In fact, output-space entropy search is a special case of ITL with a stochastic target space (cf. Equation (47) in Appendix K.1). Bogunovic et al. (2016) analyze TRUVAR in the context of Bayesian optimization and level set estimation. TRUVAR is akin to VTL with a similar notion of “target space”, but their algorithm and analysis rely on a threshold scheme which requires that  $\mathcal{A} \subseteq \mathcal{S}$ . Fiez et al. (2019) introduce the *transductive linear bandit* problem, which is a special case of transductive active learning limited to a linear function class and with the objective of determining the maximum within an initial candidate set.<sup>8</sup> We mention additional more loosely related works in Appendix A.

## 7 Conclusion

We investigated the generalization of active learning to settings with concrete prediction targets and/or with limited information due to constrained sample spaces. This provides a flexible framework, applicable also to other domains than were discussed (such as recommender systems, molecular design, robotics, etc.) by varying the choice of target space and sample space. Further, we proved novel generalization bounds which may be of independent interest for active learning. Finally, we demonstrated across broad applications that sampling *relevant and diverse* points (as opposed to only one of the two) leads to a substantial improvement upon the state-of-the-art.

<sup>8</sup>The transductive bandit problem can be solved analogously to Safe BO, by maintaining the set  $\mathcal{A}_n$ .

## Acknowledgements

Many thanks to Armin Lederer, Johannes Kirschner, Jonas Rothfuss, Lars Lorch, Manish Prajapat, Nicolas Emmenegger, Parnian Kassraie, and Scott Sussex for their insightful feedback on different versions of this manuscript, as well as Anton Baumann for helpful discussions. We further thank Freddie Bickford Smith for a constructive discussion regarding the relationship between our work and prior work.

This project was supported in part by the European Research Council (ERC) under the European Union’s Horizon 2020 research and Innovation Program Grant agreement no. 815943, the Swiss National Science Foundation under NCCR Automation, grant agreement 51NF40 180545, and by a grant of the Hasler foundation (grant no. 21039).

## References

- Abbasi-Yadkori, Y. *Online learning for linearly parametrized control problems*. PhD thesis, University of Alberta, 2013.
- Antorán, J., Janz, D., Allingham, J. U., Daxberger, E., Barbano, R. R., Nalisnick, E., and Hernández-Lobato, J. M. Adapting the linearised laplace model evidence for modern deep learning. In *ICML*, 2022.
- Arora, S., Du, S. S., Hu, W., Li, Z., Salakhutdinov, R. R., and Wang, R. On exact computation with an infinitely wide neural net. *NeurIPS*, 32, 2019.
- Arthur, D., Vassilvitskii, S., et al. k-means++: The advantages of careful seeding. In *SODA*, volume 7, 2007.
- Ash, J., Goel, S., Krishnamurthy, A., and Kakade, S. Gone fishing: Neural active learning with fisher embeddings. *NeurIPS*, 34, 2021.
- Ash, J. T., Zhang, C., Krishnamurthy, A., Langford, J., and Agarwal, A. Deep batch active learning by diverse, uncertain gradient lower bounds. *ICLR*, 2020.
- Audibert, J.-Y., Bubeck, S., and Munos, R. Best arm identification in multi-armed bandits. In *COLT*, 2010.
- Balestriero, R., Ibrahim, M., Sobal, V., Morcos, A., Shekhar, S., Goldstein, T., Bordes, F., Bardes, A., Mialon, G., Tian, Y., et al. A cookbook of self-supervised learning. *arXiv preprint arXiv:2304.12210*, 2023.
- Barrett, A. B. Exploration of synergistic and redundant information sharing in static and dynamical gaussian systems. *Physical Review E*, 91(5), 2015.
- Baumann, D., Marco, A., Turchetta, M., and Trimpe, S. Gosafe: Globally optimal safe robot learning. In *ICRA*, 2021.
- Bengio, Y., Louradour, J., Collobert, R., and Weston, J. Curriculum learning. In *ICML*, volume 26, 2009.
- Beraha, M., Metelli, A. M., Papini, M., Tirinzoni, A., and Restelli, M. Feature selection via mutual information: New theoretical insights. In *IJCNN*, 2019.
- Berkenkamp, F., Schoellig, A. P., and Krause, A. Safe controller optimization for quadrotors with gaussian processes. In *ICRA*, 2016.
- Berkenkamp, F., Krause, A., and Schoellig, A. P. Bayesian optimization with safety constraints: safe and automatic parameter tuning in robotics. *Machine Learning*, 2021.
- Berlind, C. and Urner, R. Active nearest neighbors in changing environments. In *ICML*, 2015.
- Bickford Smith, F., Kirsch, A., Farquhar, S., Gal, Y., Foster, A., and Rainforth, T. Prediction-oriented bayesian active learning. In *AISTATS*, 2023.

- Bickford Smith, F., Foster, A., and Rainforth, T. Making better use of unlabelled data in bayesian active learning. In *AISTATS*, 2024.
- Blundell, C., Cornebise, J., Kavukcuoglu, K., and Wierstra, D. Weight uncertainty in neural network. In *ICML*, 2015.
- Bogunovic, I., Scarlett, J., Krause, A., and Cevher, V. Truncated variance reduction: A unified approach to bayesian optimization and level-set estimation. *NeurIPS*, 29, 2016.
- Bottero, A., Luis, C., Vinogradskaja, J., Berkenkamp, F., and Peters, J. R. Information-theoretic safe exploration with gaussian processes. *NeurIPS*, 35, 2022.
- Bottero, A. G., Luis, C. E., Vinogradskaja, J., Berkenkamp, F., and Peters, J. Information-theoretic safe bayesian optimization. *arXiv preprint arXiv:2402.15347*, 2024.
- Bubeck, S., Munos, R., and Stoltz, G. Pure exploration in multi-armed bandits problems. In *ALT*, volume 20, 2009.
- Chaloner, K. and Verdinelli, I. Bayesian experimental design: A review. *Statistical Science*, 1995.
- Chandra, B. Quadrotor simulation, 2023. URL <https://github.com/Bharath2/Quadrotor-Simulation>.
- Chen, Y. and Krause, A. Near-optimal batch mode active learning and adaptive submodular optimization. In *ICML*, 2013.
- Chowdhury, S. R. and Gopalan, A. On kernelized multi-armed bandits. In *ICML*, 2017.
- Cohn, D. Neural network exploration using optimal experiment design. *NeurIPS*, 6, 1993.
- Coleman, C., Chou, E., Katz-Samuels, J., Culatana, S., Bailis, P., Berg, A. C., Nowak, R., Sumbaly, R., Zaharia, M., and Yalniz, I. Z. Similarity search for efficient active learning and search of rare concepts. In *AAAI*, volume 36, 2022.
- Cook, R. D. Detection of influential observation in linear regression. *Technometrics*, 19(1), 1977.
- Cooper, S. E. and Netoff, T. I. Multidimensional bayesian estimation for deep brain stimulation using the safeopt algorithm. *medRxiv*, 2022.
- Cover, T. M. *Elements of information theory*. John Wiley & Sons, 1999.
- Das, A. and Kempe, D. Algorithms for subset selection in linear regression. In *STOC*, volume 40, 2008.
- Das, A. and Kempe, D. Approximate submodularity and its applications: Subset selection, sparse approximation and dictionary selection. *JMLR*, 19(1), 2018.
- Daxberger, E., Kristiadi, A., Immer, A., Eschenhagen, R., Bauer, M., and Hennig, P. Laplace redux-effortless bayesian deep learning. *NeurIPS*, 34, 2021.
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. BERT: Pre-training of deep bidirectional transformers for language understanding. In *NAACL*, 2019.
- Emmenegger, N., Mutn y, M., and Krause, A. Likelihood ratio confidence sets for sequential decision making. *NeurIPS*, 37, 2023.
- Esfandiari, H., Karbasi, A., and Mirrokni, V. Adaptivity in adaptive submodularity. In *COLT*, 2021.
- Eustratiadis, P., Dudziak, Ł., Li, D., and Hospedales, T. Neural fine-tuning search for few-shot learning. *ICLR*, 2024.
- Fiez, T., Jain, L., Jamieson, K. G., and Ratliff, L. Sequential experimental design for transductive linear bandits. *NeurIPS*, 32, 2019.

- Fu, B., Cao, Z., Wang, J., and Long, M. Transferable query selection for active domain adaptation. In *CVPR*, 2021.
- Gal, Y., Islam, R., and Ghahramani, Z. Deep bayesian active learning with image data. In *ICML*, 2017.
- Gao, M., Zhang, Z., Yu, G., Arık, S. Ö., Davis, L. S., and Pfister, T. Consistency-based semi-supervised active learning: Towards minimizing labeling cost. In *ECCV*, 2020.
- Gardner, J. R., Kusner, M. J., Xu, Z. E., Weinberger, K. Q., and Cunningham, J. P. Bayesian optimization with inequality constraints. In *ICML*, volume 2014, 2014.
- Geifman, Y. and El-Yaniv, R. Deep active learning over the long tail. *arXiv preprint arXiv:1711.00941*, 2017.
- Grandvalet, Y. and Bengio, Y. Semi-supervised learning by entropy minimization. *NeurIPS*, 17, 2004.
- Graves, A., Bellemare, M. G., Menick, J., Munos, R., and Kavukcuoglu, K. Automated curriculum learning for neural networks. In *ICML*, 2017.
- Graybill, F. A. *An introduction to linear statistical models*. Literary Licensing, LLC, 1961.
- Guo, Y. and Greiner, R. Optimistic active-learning using mutual information. In *IJCAI*, volume 7, 2007.
- Hacohen, G., Dekel, A., and Weinshall, D. Active learning on a budget: Opposite strategies suit high and low budgets. *ICML*, 2022.
- Hardt, M. and Sun, Y. Test-time training on nearest neighbors for large language models. *ICLR*, 2024.
- He, B., Lakshminarayanan, B., and Teh, Y. W. Bayesian deep ensembles via the neural tangent kernel. *NeurIPS*, 33, 2020.
- Hendrycks, D. and Gimpel, K. A baseline for detecting misclassified and out-of-distribution examples in neural networks. *ICLR*, 2017.
- Hennig, P. and Schuler, C. J. Entropy search for information-efficient global optimization. *JMLR*, 13 (6), 2012.
- Hernández-Lobato, J. M., Hoffman, M. W., and Ghahramani, Z. Predictive entropy search for efficient global optimization of black-box functions. *NeurIPS*, 27, 2014.
- Hoffman, M. W. and Ghahramani, Z. Output-space predictive entropy search for flexible global optimization. In *NeurIPS workshop on Bayesian Optimization*, 2015.
- Holzmüller, D., Zaverkin, V., Kästner, J., and Steinwart, I. A framework and benchmark for deep batch active learning for regression. *JMLR*, 24(164), 2023.
- Hopfield, J. J. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the national academy of sciences*, 79(8), 1982.
- Houlsby, N., Huszár, F., Ghahramani, Z., and Lengyel, M. Bayesian active learning for classification and preference learning. *CoRR*, 2011.
- Hübötter, J., Sukhija, B., Treven, L., As, Y., and Krause, A. Active few-shot fine-tuning. *ICLR workshop on Bridging the Gap Between Practice and Theory in Deep Learning*, 2024.
- Ilyas, A., Park, S. M., Engstrom, L., Leclerc, G., and Madry, A. Datamodels: Predicting predictions from training data. *arXiv preprint arXiv:2202.00622*, 2022.
- Jacot, A., Gabriel, F., and Hongler, C. Neural tangent kernel: Convergence and generalization in neural networks. *NeurIPS*, 31, 2018.

- Johnson, J., Douze, M., and Jégou, H. Billion-scale similarity search with gpus. *IEEE Transactions on Big Data*, 7(3), 2019.
- Kaddour, J., Sæmundsson, S., et al. Probabilistic active meta-learning. *NeurIPS*, 33, 2020.
- Kassraie, P. and Krause, A. Neural contextual bandits without regret. In *AISTATS*, 2022.
- Khan, M. E. E., Immer, A., Abedi, E., and Korzepa, M. Approximate inference turns deep networks into gaussian processes. *NeurIPS*, 32, 2019.
- Khanna, R., Elenberg, E., Dimakis, A., Negahban, S., and Ghosh, J. Scalable greedy feature selection via weak submodularity. In *AISTATS*, 2017.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. In *ICLR*, 2014.
- Kirsch, A. Black-box batch active learning for regression. *arXiv preprint arXiv:2302.08981*, 2023.
- Kirsch, A., Van Amersfoort, J., and Gal, Y. Batchbald: Efficient and diverse batch acquisition for deep bayesian active learning. *NeurIPS*, 32, 2019.
- Kirschner, J., Mutny, M., Hiller, N., Ischebeck, R., and Krause, A. Adaptive and safe bayesian optimization in high dimensions via one-dimensional subspaces. In *ICML*, 2019.
- Koh, P. W. and Liang, P. Understanding black-box predictions via influence functions. In *ICML*, 2017.
- Kothawade, S., Beck, N., Killamsetty, K., and Iyer, R. Similar: Submodular information measures based active learning in realistic scenarios. *NeurIPS*, 34, 2021.
- Krause, A. and Golovin, D. Submodular function maximization. *Tractability*, 3, 2014.
- Krause, A. and Guestrin, C. Nonmyopic active learning of gaussian processes: an exploration-exploitation approach. In *ICML*, volume 24, 2007.
- Krause, A., Singh, A., and Guestrin, C. Near-optimal sensor placements in gaussian processes: Theory, efficient algorithms and empirical studies. *JMLR*, 9(2), 2008.
- Krizhevsky, A., Hinton, G., et al. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009.
- Kumari, L., Wang, S., Das, A., Zhou, T., and Bilmes, J. An end-to-end submodular framework for data-efficient in-context learning. In *NAACL*, 2024.
- Lakshminarayanan, B., Pritzel, A., and Blundell, C. Simple and scalable predictive uncertainty estimation using deep ensembles. *NeurIPS*, 30, 2017.
- LeCun, Y., Cortes, C., and Burges, C. J. The mnist database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>, 1998.
- Lee, J., Bahri, Y., Novak, R., Schoenholz, S. S., Pennington, J., and Sohl-Dickstein, J. Deep neural networks as gaussian processes. *ICLR*, 2018.
- Lee, J., Xiao, L., Schoenholz, S., Bahri, Y., Novak, R., Sohl-Dickstein, J., and Pennington, J. Wide neural networks of any depth evolve as linear models under gradient descent. *NeurIPS*, 32, 2019.
- Lee, Y., Chen, A. S., Tajwar, F., Kumar, A., Yao, H., Liang, P., and Finn, C. Surgical fine-tuning improves adaptation to distribution shifts. *NeurIPS workshop on Distribution Shifts*, 2022.
- Lewis, D. and Gale, W. A sequential algorithm for training text classifiers. In *SIGIR*, 1994.
- Lewis, D. D. and Catlett, J. Heterogeneous uncertainty sampling for supervised learning. In *Machine learning proceedings*. 1994.
- MacKay, D. J. Information-based objective functions for active data selection. *Neural computation*, 4(4), 1992.

- Maddox, W. J., Izmailov, P., Garipov, T., Vetrov, D. P., and Wilson, A. G. A simple baseline for bayesian uncertainty in deep learning. *NeurIPS*, 32, 2019.
- Malladi, S., Wettig, A., Yu, D., Chen, D., and Arora, S. A kernel-based view of language model fine-tuning. In *ICML*, 2023.
- Martens, J. and Grosse, R. Optimizing neural networks with kronecker-factored approximate curvature. In *ICML*, 2015.
- Mehta, R., Shui, C., Nichyporuk, B., and Arbel, T. Information gain sampling for active learning in medical image classification. In *UNSURE*, 2022.
- Murphy, K. P. *Probabilistic machine learning: Advanced topics*. MIT Press, 2023.
- Mutny, M. and Krause, A. Experimental design for linear functionals in reproducing kernel hilbert spaces. *NeurIPS*, 35, 2022.
- Nemhauser, G. L., Wolsey, L. A., and Fisher, M. L. An analysis of approximations for maximizing submodular set functions—i. *Mathematical programming*, 14, 1978.
- Ostrovsky, R., Rabani, Y., Schulman, L. J., and Swamy, C. The effectiveness of lloyd-type methods for the k-means problem. *JACM*, 2013.
- Pacchiano, A., Lee, J. N., and Brunskill, E. Experiment planning with function approximation. *NeurIPS*, 37, 2024.
- Peng, H., Long, F., and Ding, C. Feature selection based on mutual information criteria of max-dependency, max-relevance, and min-redundancy. *IEEE Transactions on pattern analysis and machine intelligence*, 27(8), 2005.
- Prabhu, V., Chandrasekaran, A., Saenko, K., and Hoffman, J. Active domain adaptation via clustering uncertainty-weighted embeddings. In *ICCV*, 2021.
- Pruthi, G., Liu, F., Kale, S., and Sundararajan, M. Estimating training data influence by tracing gradient descent. In *NeurIPS*, 2019.
- Rahimi, A. and Recht, B. Random features for large-scale kernel machines. *NeurIPS*, 20, 2007.
- Rai, P., Saha, A., Daumé III, H., and Venkatasubramanian, S. Domain adaptation meets active learning. In *NAACL HLT workshop on Active Learning for Natural Language Processing*, 2010.
- Rothfuss, J., Koenig, C., Rupenyau, A., and Krause, A. Meta-learning priors for safe bayesian optimization. In *COLT*, 2023.
- Russo, D. J., Van Roy, B., Kazerouni, A., Osband, I., Wen, Z., et al. A tutorial on thompson sampling. *Foundations and Trends® in Machine Learning*, 11(1), 2018.
- Saha, A., Rai, P., Daumé, H., Venkatasubramanian, S., and DuVall, S. L. Active supervised domain adaptation. In *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD*, 2011.
- Scheffer, T., Decomain, C., and Wrobel, S. Active hidden markov models for information extraction. In *IDA*, 2001.
- Schreiter, J., Nguyen-Tuong, D., Eberts, M., Bischoff, B., Markert, H., and Toussaint, M. Safe exploration for active learning with gaussian processes. In *ECML PKDD*, 2015.
- Sener, O. and Savarese, S. Active learning for convolutional neural networks: A core-set approach. *ICLR*, 2017.
- Seo, S., Wallat, M., Graepel, T., and Obermayer, K. Gaussian process regression: Active data selection and test point rejection. In *Mustererkennung 2000*. Springer, 2000.
- Settles, B. Active learning literature survey. Technical report, University of Wisconsin-Madison Department of Computer Sciences, 2009.

- Settles, B. and Craven, M. An analysis of active learning strategies for sequence labeling tasks. In *EMNLP*, 2008.
- Shoham, N. and Avron, H. Experimental design for overparameterized learning with application to single shot deep active learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023.
- Shwartz-Ziv, R. and LeCun, Y. To compress or not to compress—self-supervised learning and information theory: A review. *arXiv preprint arXiv:2304.09355*, 2023.
- Soviany, P., Ionescu, R. T., Rota, P., and Sebe, N. Curriculum learning: A survey. *IJCV*, 2022.
- Srinivas, N., Krause, A., Kakade, S. M., and Seeger, M. Gaussian process optimization in the bandit setting: No regret and experimental design. In *ICML*, volume 27, 2009.
- Strang, G. *Introduction to linear algebra*. SIAM, 5 edition, 2016.
- Su, J.-C., Tsai, Y.-H., Sohn, K., Liu, B., Maji, S., and Chandraker, M. Active adversarial domain adaptation. In *WACV*, 2020.
- Sui, Y., Gotovos, A., Burdick, J., and Krause, A. Safe exploration for optimization with gaussian processes. In *ICML*, 2015.
- Sukhija, B., Turchetta, M., Lindner, D., Krause, A., Trimpe, S., and Baumann, D. Gosafeopt: Scalable safe exploration for global optimization of dynamical systems. *Artificial Intelligence*, 2023.
- Tamkin, A., Nguyen, D., Deshpande, S., Mu, J., and Goodman, N. Active learning helps pretrained models learn the intended task. *NeurIPS*, 35, 2022.
- Tan, M. and Le, Q. Efficientnet: Rethinking model scaling for convolutional neural networks. In *ICML*, 2019.
- Thompson, W. R. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 1933.
- Tu, S., Frostig, R., Singh, S., and Sindhvani, V. JAX: A python library for differentiable optimal control on accelerators, 2023. URL <http://github.com/google/trajax>.
- Turchetta, M., Berkenkamp, F., and Krause, A. Safe exploration for interactive machine learning. *NeurIPS*, 32, 2019.
- Vakili, S., Khezeli, K., and Picheny, V. On information gain and regret bounds in gaussian process bandits. In *AISTATS*, 2021.
- Vapnik, V. *Estimation of dependences based on empirical data*. Springer Science & Business Media, 1982.
- Vergara, J. R. and Estévez, P. A. A review of feature selection methods based on mutual information. *Neural computing and applications*, 24, 2014.
- Vinyals, O., Blundell, C., Lillicrap, T., Wierstra, D., et al. Matching networks for one shot learning. *NeurIPS*, 29, 2016.
- Wainwright, M. J. *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge university press, 2019.
- Wang, C., Sun, S., and Grosse, R. Beyond marginal uncertainty: How accurately can bayesian regression models estimate posterior predictive correlations? In *AISTATS*, 2021.
- Wang, Z. and Jegelka, S. Max-value entropy search for efficient bayesian optimization. In *ICML*, 2017.
- Wei, A., Hu, W., and Steinhardt, J. More than a toy: Random matrix models predict how real-world neural representations generalize. In *ICML*, 2022.



- Widmer, D., Kang, D., Sukhija, B., Hübotter, J., Krause, A., and Coros, S. Tuning legged locomotion controllers via safe bayesian optimization. *CORL*, 2023.
- Wilks, S. S. Certain generalizations in the analysis of variance. *Biometrika*, 1932.
- Williams, C. K. and Rasmussen, C. E. *Gaussian processes for machine learning*, volume 2. MIT press Cambridge, MA, 2006.
- Wischnewski, A., Betz, J., and Lohmann, B. A model-free algorithm to safely approach the handling limit of an autonomous racecar. In *ICCV*, 2019.
- Xia, M., Malladi, S., Gururangan, S., Arora, S., and Chen, D. Less: Selecting influential data for targeted instruction tuning. In *ICML*, 2024.
- Xu, M. and Kazantsev, G. Understanding goal-oriented active learning via influence functions. In *NeurIPS Workshop on Machine Learning with Guarantees*, 2019.
- Ye, J., Wu, Z., Feng, J., Yu, T., and Kong, L. Compositional exemplars for in-context learning. In *ICML*, 2023.
- Yehuda, O., Dekel, A., Hachohen, G., and Weinshall, D. Active learning through a covering lens. *NeurIPS*, 35, 2022.
- Yu, H. and Kim, S. Passive sampling for regression. In *ICDM*, 2010.
- Yu, K., Bi, J., and Tresp, V. Active learning via transductive experimental design. In *ICML*, volume 23, 2006.
- Zanette, A., Dong, K., Lee, J. N., and Brunskill, E. Design of experiments for stochastic contextual linear bandits. *NeurIPS*, 34, 2021.
- Zheng, H., Liu, R., Lai, F., and Prakash, A. Coverage-centric coreset selection for high pruning rates. *ICLR*, 2023.

# Appendices

A general principle of “transductive learning” was already formulated by the famous computer scientist Vladimir Vapnik in the 20th century. Vapnik proposes the following “imperative for a complex world”:

*When solving a problem of interest, do not solve a more general problem as an intermediate step. Try to get the answer that you really need but not a more general one.*

– Vapnik (1982)

These appendices provide additional background, proofs, experiment details, and ablation studies.

## Contents

<b>A Additional Related Work</b>	<b>20</b>
<b>B Background</b>	<b>20</b>
B.1 Information Theory . . . . .	20
B.2 Gaussian Processes . . . . .	20
<b>C Proofs</b>	<b>21</b>
C.1 Undirected Case of ITL . . . . .	21
C.2 Non-adaptive Data Selection & Submodularity . . . . .	21
C.3 Batch Diversity: Batch Selection as Non-adaptive Data Selection . . . . .	22
C.4 Measures of Synergies & Approximate Submodularity . . . . .	23
C.5 Convergence of Marginal Gain . . . . .	25
C.6 Proof of Theorem 3.2 . . . . .	26
C.7 Proof of Theorem 3.3 . . . . .	31
C.8 Proof of Theorem 5.1 . . . . .	32
C.9 Useful Facts and Inequalities . . . . .	36
<b>D Interpretations &amp; Approximations of Principle (†)</b>	<b>37</b>
D.1 Interpretations of ITL . . . . .	37
D.2 Interpretations of VTL . . . . .	37
D.3 Mean Marginal ITL . . . . .	38
D.4 Correlation-based Transductive Learning . . . . .	40
D.5 Summary . . . . .	40
<b>E Stochastic Target Spaces</b>	<b>41</b>
<b>F Closed-form Decision Rules</b>	<b>41</b>
<b>G Computational Complexity</b>	<b>42</b>
<b>H Additional GP Experiments &amp; Details</b>	<b>42</b>

<b>I</b>	<b>Alternative Settings for Active Fine-Tuning</b>	<b>43</b>
I.1	Prediction Targets are Contained in Sample Space: $\mathcal{A} \subseteq \mathcal{S}$	43
I.2	Active Domain Adaptation	44
<b>J</b>	<b>Additional NN Experiments &amp; Details</b>	<b>44</b>
J.1	Experiment Details	45
J.2	Embeddings and Kernels	46
J.3	Towards Uncertainty Quantification in Latent Space	47
J.4	Batch Selection via Conditional Embeddings	47
J.5	Baselines	48
J.6	Additional experiments	52
J.7	Ablation study of noise standard deviation $\rho$	52
<b>K</b>	<b>Additional Safe BO Experiments &amp; Details</b>	<b>54</b>
K.1	A More Exploitative Stochastic Target Space	54
K.2	Detailed Comparison with Prior Works	55
K.3	Jumping Past Local Barriers	60
K.4	Experiment Details	61

## A Additional Related Work

The general principle of non-active “transductive learning” was introduced by Vapnik (1982). The notion of “target” from transductive active learning is akin to the notion of “task” in curriculum learning (Bengio et al., 2009; Graves et al., 2017; Soviany et al., 2022). The study of settings where the irreducible uncertainty is zero is related to the study of estimability in experimental design (Graybill, 1961; Mutny & Krause, 2022). In feature selection, selecting features that maximize information gain with respect to a to-be-predicted label is a standard approach (Peng et al., 2005; Vergara & Estévez, 2014; Beraha et al., 2019) which is akin to ITL (cf. Appendix D). The themes of relevance and diversity are also important for efficient in-context learning (e.g., Ye et al., 2023; Kumari et al., 2024) and data pruning (Zheng et al., 2023). Transductive active learning is complimentary to other learning methodologies, such as semi-supervised learning (Gao et al., 2020), self-supervised learning (Shwartz-Ziv & LeCun, 2023; Balestrierio et al., 2023), and meta-learning (Kaddour et al., 2020; Rothfuss et al., 2023).

## B Background

### B.1 Information Theory

Throughout this work,  $\log$  denotes the natural logarithm. Given random vectors  $\mathbf{x}$  and  $\mathbf{y}$ , we denote by

$$\begin{aligned} H[\mathbf{x}] &\stackrel{\text{def}}{=} \mathbb{E}_{p(\mathbf{x})}[-\log p(\mathbf{x})], \\ H[\mathbf{x} \mid \mathbf{y}] &\stackrel{\text{def}}{=} \mathbb{E}_{p(\mathbf{x}, \mathbf{y})}[-\log p(\mathbf{x} \mid \mathbf{y})], \quad \text{and} \\ I(\mathbf{x}; \mathbf{y}) &\stackrel{\text{def}}{=} H[\mathbf{x}] - H[\mathbf{x} \mid \mathbf{y}] \end{aligned}$$

the (differential) entropy, conditional entropy, and information gain, respectively (Cover, 1999).<sup>9</sup>

The *multivariate information gain* (Murphy, 2023) between random vectors  $\mathbf{x}, \mathbf{y}, \mathbf{z}$  is given by

$$\begin{aligned} I(\mathbf{x}; \mathbf{y}; \mathbf{z}) &\stackrel{\text{def}}{=} I(\mathbf{x}; \mathbf{y}) - I(\mathbf{x}; \mathbf{y} \mid \mathbf{z}) \\ &= I(\mathbf{x}; \mathbf{y}) + I(\mathbf{x}; \mathbf{z}) - I(\mathbf{x}; \mathbf{y}, \mathbf{z}). \end{aligned} \tag{6} \tag{7}$$

When  $I(\mathbf{x}; \mathbf{y}; \mathbf{z}) \neq 0$  it is said that  $\mathbf{y}$  and  $\mathbf{z}$  *interact* regarding their information about  $\mathbf{x}$ . If the interaction is positive, it is said that the information of  $\mathbf{z}$  about  $\mathbf{x}$  is *redundant* given  $\mathbf{y}$ . Conversely, if the interaction is negative, it is said that the information of  $\mathbf{z}$  about  $\mathbf{x}$  is *synergistic* with  $\mathbf{y}$ . The notion of synergy is akin to the frequentist notion of “suppressor variables” in linear regression (Das & Kempe, 2008).

### B.2 Gaussian Processes

The stochastic process  $f$  is a Gaussian process (GP, Williams & Rasmussen (2006)), denoted  $f \sim \mathcal{GP}(\mu, k)$ , with mean function  $\mu$  and kernel  $k$  if for any finite subset  $X = \{\mathbf{x}_1, \dots, \mathbf{x}_n\} \subseteq \mathcal{X}$ ,  $\mathbf{f}_X \sim \mathcal{N}(\boldsymbol{\mu}_X, \mathbf{K}_{XX})$  is jointly Gaussian with mean vector  $\boldsymbol{\mu}_X(i) = \mu(\mathbf{x}_i)$  and covariance matrix  $\mathbf{K}_{XX}(i, j) = k(\mathbf{x}_i, \mathbf{x}_j)$ .

In the following, we formalize the assumptions from the GP setting (cf. Section 3.1).

**Assumption B.1** (Gaussian prior). We assume that  $f \sim \mathcal{GP}(\mu, k)$  with known mean function  $\mu$  and kernel  $k$ .

**Assumption B.2** (Gaussian noise). We assume that the noise  $\varepsilon_{\mathbf{x}}$  is mutually independent and zero-mean Gaussian with known variance  $\rho^2(\mathbf{x}) > 0$ . We write  $\mathbf{P}_X = \text{diag } \rho^2(\mathbf{x}_1), \dots, \rho^2(\mathbf{x}_n)$ .

Under Assumptions B.1 and B.2, the posterior distribution of  $f$  after observing points  $X$  is  $\mathcal{GP}(\mu_n, k_n)$  with

$$\begin{aligned} \mu_n(\mathbf{x}) &= \mu(\mathbf{x}) + \mathbf{K}_{\mathbf{x}X}(\mathbf{K}_{XX} + \mathbf{P}_X)^{-1}(\mathbf{y}_X - \boldsymbol{\mu}_X), \\ k_n(\mathbf{x}, \mathbf{x}') &= k(\mathbf{x}, \mathbf{x}') - \mathbf{K}_{\mathbf{x}X}(\mathbf{K}_{XX} + \mathbf{P}_X)^{-1}\mathbf{K}_{X\mathbf{x}'}, \\ \sigma_n^2(\mathbf{x}) &= k_n(\mathbf{x}, \mathbf{x}). \end{aligned}$$

<sup>9</sup>One has to be careful to ensure that  $I(\mathbf{x}; \mathbf{y})$  exists, i.e.,  $|I(\mathbf{x}; \mathbf{y})| < \infty$ . We will assume that this is the case throughout this work. When  $\mathbf{x}$  and  $\mathbf{y}$  are jointly Gaussian, this is satisfied when the noise variance  $\rho^2$  is positive.

For Gaussian random vectors  $\mathbf{f}$  and  $\mathbf{y}$ , the entropy is  $H[\mathbf{f}] = \frac{n}{2} \log(2\pi e) + \frac{1}{2} \log |\text{Var}[\mathbf{f}]|$ , the information gain is  $I(\mathbf{f}; \mathbf{y}) = \frac{1}{2}(\log |\text{Var}[\mathbf{y}]| - \log |\text{Var}[\mathbf{y} | \mathbf{f}]|)$ , and

$$\gamma_n = \max_{\substack{X \subseteq \mathcal{X} \\ |X| \leq n}} \frac{1}{2} \log |\mathbf{I} + \mathbf{P}_X^{-1} \mathbf{K}_{XX}|.$$

## C Proofs

We will write

- $\sigma^2 \stackrel{\text{def}}{=} \max_{\mathbf{x} \in \mathcal{X}} \sigma_0^2(\mathbf{x})$ , and
- $\tilde{\sigma}^2 \stackrel{\text{def}}{=} \max_{\mathbf{x} \in \mathcal{X}} \sigma_0^2(\mathbf{x}) + \rho^2(\mathbf{x})$ .

The following is a brief overview of the structure of this section:

1. Appendix C.1 relates ITL in the inductive learning setting ( $\mathcal{S} \subseteq \mathcal{A}$ ) to prior work.
2. Appendix C.2 relates the designs selected by ITL and VTL to the optimal designs for corresponding non-adaptive objectives.
3. Appendix C.3 shows that batch selection via ITL or VTL leads to informative and diverse batches, utilizing the results from Appendix C.2.
4. Appendix C.4 introduces measures of synergies that generalize the submodularity assumption (cf. Assumption 3.1).
5. Appendix C.5 proves key results on the convergence of the ITL and VTL objectives.
6. Appendix C.6 proves Theorem 3.2 (convergence in GP setting).
7. Appendix C.7 proves Theorem 3.3 (convergence in agnostic setting).
8. Appendix C.8 proves Theorem 5.1 (convergence in safe BO application).
9. Appendix C.9 includes useful facts.

### C.1 Undirected Case of ITL

We briefly examine the important special case of ITL where  $\mathcal{S} \subseteq \mathcal{A}$ . In this setting, for all  $\mathbf{x} \in \mathcal{S}$ , the decision rule of ITL simplifies to

$$\begin{aligned} I(\mathbf{f}_{\mathcal{A}}; y_{\mathbf{x}} | \mathcal{D}_n) &\stackrel{(i)}{=} I(\mathbf{f}_{\mathcal{A} \setminus \{\mathbf{x}\}}; y_{\mathbf{x}} | f_{\mathbf{x}}, \mathcal{D}_n) + I(f_{\mathbf{x}}; y_{\mathbf{x}} | \mathcal{D}_n) \\ &\stackrel{(ii)}{=} I(f_{\mathbf{x}}; y_{\mathbf{x}} | \mathcal{D}_n) \\ &= H[y_{\mathbf{x}} | \mathcal{D}_n] - H[\varepsilon_{\mathbf{x}}] \end{aligned}$$

where (i) follows from the chain rule of information gain and  $\mathbf{x} \in \mathcal{S} \subseteq \mathcal{A}$ ; and (ii) follows from the conditional independence  $\mathbf{f}_{\mathcal{A}} \perp y_{\mathbf{x}} | f_{\mathbf{x}}$ .

If additionally  $f$  is a GP then

$$H[y_{\mathbf{x}} | \mathcal{D}_n] - H[\varepsilon_{\mathbf{x}}] = \frac{1}{2} \log \left( 1 + \frac{\text{Var}[f_{\mathbf{x}} | \mathcal{D}_n]}{\text{Var}[\varepsilon_{\mathbf{x}}]} \right).$$

This decision rule has also been termed *total information gain* (MacKay, 1992). When  $\mathcal{S} \subseteq \mathcal{A}$  and observation noise is homoscedastic, this decision rule is equivalent to uncertainty sampling.

### C.2 Non-adaptive Data Selection & Submodularity

Recall the non-myopic information gain  $\psi_{\mathcal{A}}(X) = I(\mathbf{f}_{\mathcal{A}}; \mathbf{y}_X)$  (ITL) and variance reduction  $\psi_{\mathcal{A}}(X) = \text{tr} \text{Var}[\mathbf{f}_{\mathcal{A}}] - \text{tr} \text{Var}[\mathbf{f}_{\mathcal{A}} | \mathbf{y}_X]$  (VTL) objective functions from Assumption 3.1. In this section, we will relate the designs selected by ITL and VTL to the optimal designs for these objectives. To this end, consider the non-adaptive optimization problem

$$X^* = \arg \max_{\substack{X \subseteq \mathcal{S} \\ |X|=k}} \psi_{\mathcal{A}}(X).$$

**Lemma C.1.** For both ITL and VTL,  $\psi_{\mathcal{A}}$  is non-negative and monotone.

*Proof.* For ITL,  $\psi_{\mathcal{A}}(X) \geq 0$  follows from the non-negativity of mutual information. To conclude monotonicity, note that for any  $X' \subseteq X \subseteq \mathcal{S}$ ,

$$I(\mathbf{f}_{\mathcal{A}}; \mathbf{y}_{X'}) = H[\mathbf{f}_{\mathcal{A}}] - H[\mathbf{f}_{\mathcal{A}} | \mathbf{y}_{X'}] \leq H[\mathbf{f}_{\mathcal{A}}] - H[\mathbf{f}_{\mathcal{A}} | \mathbf{y}_X] = I(\mathbf{f}_{\mathcal{A}}; \mathbf{y}_X)$$

due to monotonicity of conditional entropy (which is also called the ‘‘information never hurts’’ principle).

For VTL, recall that  $\text{tr Var}[\mathbf{f}_{\mathcal{A}} | \mathbf{y}_X] \leq \text{tr Var}[\mathbf{f}_{\mathcal{A}} | \mathbf{y}_{X'}]$  for any  $X' \subseteq X \subseteq \mathcal{S}$  (with an implicit expectation over  $\mathbf{y}_X, \mathbf{y}_{X'}$ ). Non-negativity and monotonicity of  $\psi_{\mathcal{A}}$  then follow analogously to ITL.  $\square$

**Lemma C.2.** The marginal gain  $\Delta_{\mathcal{A}}(\mathbf{x} | X) \stackrel{\text{def}}{=} \psi_{\mathcal{A}}(X \cup \{\mathbf{x}\}) - \psi_{\mathcal{A}}(X)$  of  $\mathbf{x} \in \mathcal{S}$  given  $X \subseteq \mathcal{S}$  is the ITL and VTL objective, respectively.

*Proof.* For ITL,

$$\begin{aligned} \Delta_{\mathcal{A}}(\mathbf{x} | X) &= I(\mathbf{f}_{\mathcal{A}}; \mathbf{y}_X, y_{\mathbf{x}}) - I(\mathbf{f}_{\mathcal{A}}; \mathbf{y}_X) \\ &= H[\mathbf{f}_{\mathcal{A}} | \mathbf{y}_X] - H[\mathbf{f}_{\mathcal{A}} | \mathbf{y}_X, y_{\mathbf{x}}] \\ &= I(\mathbf{f}_{\mathcal{A}}; y_{\mathbf{x}} | \mathbf{y}_X) \end{aligned}$$

which is precisely the ITL objective.

For VTL,

$$\begin{aligned} \Delta_{\mathcal{A}}(\mathbf{x} | X) &= \text{tr Var}[\mathbf{f}_{\mathcal{A}} | \mathbf{y}_X] - \text{tr Var}[\mathbf{f}_{\mathcal{A}} | \mathbf{y}_X, y_{\mathbf{x}}] \\ &= -\text{tr Var}[\mathbf{f}_{\mathcal{A}} | \mathbf{y}_X, y_{\mathbf{x}}] + \text{const} \end{aligned}$$

which is precisely the VTL objective.  $\square$

**Definition C.3** (Submodularity).  $\psi_{\mathcal{A}}$  is submodular if and only if for all  $\mathbf{x} \in \mathcal{S}$  and  $X' \subseteq X \subseteq \mathcal{S}$ ,

$$\Delta_{\mathcal{A}}(\mathbf{x} | X') \geq \Delta_{\mathcal{A}}(\mathbf{x} | X).$$

**Theorem C.4** (Nemhauser et al. (1978)). Let Assumption 3.1 hold. For any  $n \geq 1$ , if ITL or VTL selected  $\mathbf{x}_{1:n}$ , respectively, then

$$\psi_{\mathcal{A}}(\mathbf{x}_{1:n}) \geq (1 - 1/e) \max_{\substack{X \subseteq \mathcal{S} \\ |X| \leq n}} \psi_{\mathcal{A}}(X).$$

*Proof.* This is a special case of a canonical result from non-negative monotone submodular function maximization (Nemhauser et al., 1978; Krause & Golovin, 2014).  $\square$

### C.3 Batch Diversity: Batch Selection as Non-adaptive Data Selection

Recall the non-adaptive optimization problem

$$B_{n,k} = \arg \max_{\substack{B \subseteq \mathcal{S} \\ |B|=k}} I(\mathbf{f}_{\mathcal{A}}; \mathbf{y}_B | \mathcal{D}_{n-1})$$

from Equation (3) with batch size  $k > 0$ , and denote by  $B'_{n,k} = \mathbf{x}_{n,1:k}$  the greedy approximation from Equation (3). The selection of an individual batch can be seen as a single non-adaptive optimization problem with marginal gain

$$\begin{aligned} \Delta_n(\mathbf{x} | B) &= I(\mathbf{f}_{\mathcal{A}}; \mathbf{y}_B, y_{\mathbf{x}} | \mathcal{D}_{n-1}) - I(\mathbf{f}_{\mathcal{A}}; \mathbf{y}_B | \mathcal{D}_{n-1}) \\ &= H[\mathbf{f}_{\mathcal{A}} | \mathcal{D}_{n-1}, \mathbf{y}_B] - H[\mathbf{f}_{\mathcal{A}} | \mathcal{D}_{n-1}, \mathbf{y}_B, y_{\mathbf{x}}] \\ &= I(\mathbf{f}_{\mathcal{A}}; y_{\mathbf{x}} | \mathcal{D}_{n-1}, \mathbf{y}_B) \end{aligned}$$

and which is precisely the objective function of ITL from Equation (3). Hence, the approximation guarantees from Theorems C.4 and C.11 apply. The derivation is analogous for VTL.

Prior work has shown that the greedy solution  $B'_n$  is also competitive with a fully sequential ‘‘batchless’’ decision rule (Chen & Krause, 2013; Esfandiari et al., 2021).

## C.4 Measures of Synergies & Approximate Submodularity

We will now show that “downstream synergies”, if present, can be seen as a source of learning complexity, which is orthogonal to the information capacity  $\gamma_n$ .

**Example C.5.** Consider the example where  $f$  is a stochastic process of three random variables  $X, Y, Z$  where  $X$  and  $Y$  are Bernoulli ( $p = \frac{1}{2}$ ), and  $Z$  is the XOR of  $X$  and  $Y$ . Suppose that observations are exact (i.e.,  $\varepsilon_n = 0$ ), that the target space  $\mathcal{A}$  comprises the output variable  $Z$  while the sample space  $\mathcal{S}$  comprises the input variables  $X$  and  $Y$ . Observing any single  $X$  or  $Y$  yields no information about  $Z$ :  $I(Z; X) = I(Z; Y) = 0$ , however, observing both inputs jointly perfectly determines  $Z$ :  $I(Z; X, Y) = 1$ . Thus,  $\gamma_n(\mathcal{A}; \mathcal{S}) = 1$  if  $n \geq 2$  and  $\gamma_n(\mathcal{A}; \mathcal{S}) = 0$  else.

Learning about  $Z$  in examples of this kind is difficult for agents that make decisions greedily, since the next action (observing  $X$  or  $Y$ ) yields no signal about its long-term usefulness. We call a sequence of observations, such as  $\{X, Y\}$ , *synergistic* since its combined information value is larger than the individual values. The prevalence of synergies is not captured by the information capacity  $\gamma_n(\mathcal{A}; \mathcal{S})$  since it measures only the joint information gain of  $n$  samples within  $\mathcal{S}$ . Instead, the prevalence of synergies is captured by the sequence  $\Gamma_n \stackrel{\text{def}}{=} \max_{\mathbf{x} \in \mathcal{S}} \Delta_{\mathcal{A}}(\mathbf{x} | \mathbf{x}_{1:n})$ , which measures the maximum information gain of  $y_{n+1}$ . If  $\Gamma_n > \Gamma_{n-1}$  at any round  $n$ , this indicates a synergy. The following key object measures the additional complexity due to synergies.

**Definition C.6** (Task complexity). For  $n \geq 1$ , assuming  $\Gamma_i > 0$  for all  $1 \leq i \leq n$ , we define the *task complexity* as

$$\alpha_{\mathcal{A}, \mathcal{S}}(n) \stackrel{\text{def}}{=} \max_{i \in \{0, \dots, n-1\}} \frac{\Gamma_{n-1}}{\Gamma_i}.$$

Note that  $\alpha_{\mathcal{A}, \mathcal{S}}(n)$  is large only if the information gain of  $y_n$  is larger than that of a previous observation  $y_i$ . Intuitively, if  $\alpha_{\mathcal{A}, \mathcal{S}}(n)$  is large, the agent had to discover the *implicit* intermediate observations  $y_1, \dots, y_{n-1}$  that lead to downstream synergies. We will subsequently formalize the intimate connections of the task complexity to synergies and submodularity. Note that in the GP setting,  $\alpha_{\mathcal{A}, \mathcal{S}}(n)$  can be computed online by keeping track of the smallest  $\Gamma_i$  during previous rounds  $i$ . Further, note that  $\alpha_{\mathcal{A}, \mathcal{S}}(n) \leq 1$  if  $\psi_{\mathcal{A}}$  is submodular.

### C.4.1 The Information Ratio

Another object will prove useful in our analysis of synergies.

Consider an alternative multiplicative interpretation of the multivariate information gain (cf. Equation (7)), which we call the *information ratio* of  $X \subseteq \mathcal{S}$  given  $D \subseteq \mathcal{S}$ ,  $|X|, |D| < \infty$ :

$$\bar{\kappa}(X | D) \stackrel{\text{def}}{=} \frac{\sum_{\mathbf{x} \in X} \Delta_{\mathcal{A}}(\mathbf{x} | D)}{\Delta_{\mathcal{A}}(X | D)} \in [0, \infty). \quad (8)$$

Observe that  $\bar{\kappa}(X | D)$  measures the synergy properties of  $\mathbf{y}_X$  with respect to  $\mathbf{f}_{\mathcal{A}}$  given  $\mathbf{y}_D$  in a multiplicative sense. That is, if  $\bar{\kappa}(X | D) > 1$  then information in  $\mathbf{y}_X$  is redundant, whereas if  $\bar{\kappa}(X | D) < 1$  then information in  $\mathbf{y}_X$  is synergistic, and if  $\bar{\kappa}(X | D) = 1$  then  $\mathbf{y}_X$  do not mutually interact with respect to  $\mathbf{f}_{\mathcal{A}}$  (all given  $\mathbf{y}_D$ ). In the degenerate case where  $\Delta_{\mathcal{A}}(X | D) = 0$  (which implies  $\sum_{\mathbf{x} \in X} \Delta_{\mathcal{A}}(\mathbf{x} | D) = 0$ ) we therefore let  $\bar{\kappa}(X | D) = 1$ .

**The information ratio of ITL is strictly positive in the Gaussian case** We prove the following straightforward lower bound to the information ratio of ITL.

**Lemma C.7.** *Let  $X, D \subseteq \mathcal{S}, |X|, |D| < \infty$ . If  $\mathbf{f}_{\mathcal{A}}$  and  $\mathbf{y}_{X \cup D}$  are jointly Gaussian then  $\bar{\kappa}(X | D) > 0$ .*

*Proof.* W.l.o.g. assume  $D = \emptyset$ . We let  $X = \{\mathbf{x}_1, \dots, \mathbf{x}_k\}$  and prove lower and upper bound separately. We assume w.l.o.g. that  $I(\mathbf{f}_{\mathcal{A}}; \mathbf{y}_X) > 0$  which implies  $|\text{Var}[\mathbf{f}_{\mathcal{A}} | \mathbf{y}_X]| < |\text{Var}[\mathbf{f}_{\mathcal{A}}]|$ . Thus, there exists some  $i$  such that  $\mathbf{f}_{\mathcal{A}}$  and  $y_{\mathbf{x}_i}$  are dependent, so  $|\text{Var}[\mathbf{f}_{\mathcal{A}} | y_{\mathbf{x}_i}]| < |\text{Var}[\mathbf{f}_{\mathcal{A}}]|$  which implies  $I(\mathbf{f}_{\mathcal{A}}; y_{\mathbf{x}_i}) > 0$ . We therefore conclude that  $\bar{\kappa}(X) > 0$ .  $\square$

The following example shows that this lower bound is tight.

**Example C.8** (Synergies of Gaussian random variables, inspired by Section 3 of Barrett (2015)). Consider the three random variables  $X$ ,  $Y$ , and  $Z$  (think  $\mathcal{A} = \{X\}$  and  $\mathcal{S} = \{Y, Z\}$ ) which are jointly Gaussian with mean vector  $\mathbf{0}$  and covariance matrix

$$\Sigma = \begin{bmatrix} 1 & a & a \\ a & 1 & 0 \\ a & 0 & 1 \end{bmatrix}, \quad \text{for } 2a^2 < 1$$

where the constraint on  $a$  is to ensure that  $\Sigma$  is positive definite. Computing the mutual information, we have

$$I(X; Y) = I(X; Z) = -\frac{1}{2} \log(1 - a^2)$$

and  $I(X; Y, Z) = -\frac{1}{2} \log(1 - 2a^2)$ . Therefore,

$$\frac{I(X; Y) + I(X; Z)}{I(X; Y, Z)} = \frac{\log(1 - 2a^2 + a^4)}{\log(1 - 2a^2)} < 1.$$

Note that

$$\lim_{a \rightarrow \frac{1}{\sqrt{2}}} \frac{\log(1 - 2a^2 + a^4)}{\log(1 - 2a^2)} = 0,$$

and hence — perhaps unintuitively — even if  $Y$  and  $Z$  are uncorrelated, their information about  $X$  may be arbitrarily synergistic.

#### C.4.2 The Submodularity of the Special “Undirected” Case of ITL

In the inductive active learning problem considered in most prior works, where  $\mathcal{S} \subseteq \mathcal{A}$  and  $f$  is a Gaussian process, it holds for ITL that  $\alpha_{\mathcal{A}, \mathcal{S}}(n) = 1$  since all learning targets appear *explicitly* in  $\mathcal{S}$ :

**Lemma C.9.** *Let  $\mathcal{S} \subseteq \mathcal{A}$ . Then  $\psi_{\mathcal{A}}$  of ITL is submodular.*

*Proof.* Fix any  $\mathbf{x} \in \mathcal{S}$  and  $X' \subseteq X \subseteq \mathcal{S}$ . Let  $\bar{X} \stackrel{\text{def}}{=} X \setminus X'$ . By the definition of conditional information gain, we have

$$\Delta_{\mathcal{A}}(\mathbf{x} | X) = I(y_{\mathbf{x}}; \mathbf{f}_{\mathcal{A}} | \mathbf{y}_X) = I(y_{\mathbf{x}}; \mathbf{f}_{\mathcal{A}}, \mathbf{y}_{X'} | \mathbf{y}_{\bar{X}}) - I(y_{\mathbf{x}}; \mathbf{y}_{X'} | \mathbf{y}_{\bar{X}}).$$

Since for any  $\mathbf{x} \in \mathcal{S}$  and  $X \subseteq \mathcal{S}$ ,  $\mathbf{y}_{\mathbf{x}} \perp \mathbf{y}_X | \mathbf{f}_{\mathcal{A}}$ , this simplifies to

$$I(y_{\mathbf{x}}; \mathbf{f}_{\mathcal{A}} | \mathbf{y}_X) = I(y_{\mathbf{x}}; \mathbf{f}_{\mathcal{A}} | \mathbf{y}_{\bar{X}}) - I(y_{\mathbf{x}}; \mathbf{y}_{X'} | \mathbf{y}_{\bar{X}}).$$

It then follows from  $I(y_{\mathbf{x}}; \mathbf{y}_{X'} | \mathbf{y}_{\bar{X}}) \geq 0$  that

$$\Delta_{\mathcal{A}}(\mathbf{x} | X) = I(y_{\mathbf{x}}; \mathbf{f}_{\mathcal{A}} | \mathbf{y}_X) \leq I(y_{\mathbf{x}}; \mathbf{f}_{\mathcal{A}} | \mathbf{y}_{\bar{X}}) = \Delta_{\mathcal{A}}(\mathbf{x} | X').$$

□

This implies that  $\alpha_{\mathcal{A}, \mathcal{S}}(n) \leq 1$  for any  $n$  and  $\bar{\kappa}(X | D) \geq 1$  for any  $X, D \subseteq \mathcal{S}$  when  $\mathcal{S} \subseteq \mathcal{A}$ .

#### C.4.3 The Submodularity Ratio

Building upon the theory of maximizing non-negative monotone submodular functions (Nemhauser et al., 1978; Krause & Golovin, 2014), Das & Kempe (2018) define the following notion of “approximate” submodularity:

**Definition C.10** (Submodularity ratio). The *submodularity ratio* of  $\psi_{\mathcal{A}}$  up to cardinality  $n \geq 1$  is

$$\kappa_{\mathcal{A}}(n) \stackrel{\text{def}}{=} \min_{\substack{D \subseteq \mathbf{x}_{1:n} \\ X \subseteq \mathcal{S}: |X| \leq n \\ D \cap X = \emptyset}} \bar{\kappa}(X | D), \quad (9)$$

where they define  $\frac{0}{0} \equiv 1$ .  $\psi_{\mathcal{A}}$  is said to be  $\kappa$ -weakly submodular for some  $\kappa > 0$  if  $\inf_{n \in \mathbb{N}} \kappa_{\mathcal{A}}(n) \geq \kappa$ .

As a special case of Theorem 6 from Das & Kempe (2018), applying that  $\psi_{\mathcal{A}}$  is non-negative and monotone, we obtain the following result.

**Theorem C.11** (Das & Kempe (2018)). *For any  $n \geq 1$ , if ITL or VTL selected  $\mathbf{x}_{1:n}$ , respectively, then*

$$\psi_{\mathcal{A}}(\mathbf{x}_{1:n}) \geq (1 - e^{-\kappa_{\mathcal{A}}(n)}) \max_{\substack{X \subseteq \mathcal{S} \\ |X| \leq n}} \psi_{\mathcal{A}}(X).$$

If  $\psi_{\mathcal{A}}$  is submodular, it is implied that  $\kappa_{\mathcal{A}}(n) \geq 1$  for all  $n \geq 1$  in which case Theorem C.11 recovers Theorem C.4.



## C.5 Convergence of Marginal Gain

Our following analysis allows for changing target spaces  $\mathcal{A}_n$  and sample spaces  $\mathcal{S}_n$  (cf. Section 5), and to this end, we redefine  $\Gamma_n \stackrel{\text{def}}{=} \max_{\mathbf{x} \in \mathcal{S}_n} \Delta_{\mathcal{A}_n}(\mathbf{x} \mid \mathbf{x}_{1:n})$ . The following theorems show that the marginal gains of ITL and VTL converge to zero, and will serve as the main tool for establishing Theorems 3.2 and 3.3. We will abbreviate  $\alpha_{\mathcal{A},\mathcal{S}}(n)$  by  $\alpha_n$ .

**Theorem C.12** (Convergence of Marginal Gain for ITL). *Assume that Assumptions B.1 and B.2 are satisfied. Fix any integers  $n_1 > n_0 \geq 0$ ,  $\Delta = n_1 - n_0 + 1$  such that for all  $i \in \{n_0, \dots, n_1 - 1\}$ ,  $\mathcal{A}_{i+1} \subseteq \mathcal{A}_i$  and  $\mathcal{S} \stackrel{\text{def}}{=} \mathcal{S}_{i+1} = \mathcal{S}_i$ . Further, assume  $|\mathcal{A}_{n_0}| < \infty$ . Then, if the sequence  $\{\mathbf{x}_{i+1}\}_{i=n_0}^{n_1}$  was generated by ITL,*

$$\Gamma_{n_1} \leq \alpha_{n_1} \frac{\gamma_{\Delta}}{\Delta}. \quad (10)$$

Moreover, if  $n_0 = 0$ ,

$$\Gamma_{n_1} \leq \alpha_{n_1} \frac{\gamma_{\mathcal{A}_0, \mathcal{S}}(\Delta)}{\Delta}. \quad (11)$$

*Proof.* We have

$$\begin{aligned} \Gamma_{n_1} &= \frac{1}{\Delta} \sum_{i=n_0}^{n_1} \Gamma_{n_1} \\ &\stackrel{(i)}{\leq} \frac{\alpha_{n_1}}{\Delta} \sum_{i=n_0}^{n_1} \Gamma_i \\ &= \frac{\alpha_{n_1}}{\Delta} \sum_{i=n_0}^{n_1} \max_{\mathbf{x} \in \mathcal{S}} \mathbb{I}(\mathbf{f}_{\mathcal{A}_i}; y_{\mathbf{x}} \mid \mathbf{y}_{1:i}) \\ &\stackrel{(ii)}{=} \frac{\alpha_{n_1}}{\Delta} \sum_{i=n_0}^{n_1} \mathbb{I}(\mathbf{f}_{\mathcal{A}_i}; y_{\mathbf{x}_{i+1}} \mid \mathcal{D}_i) \\ &\stackrel{(iii)}{\leq} \frac{\alpha_{n_1}}{\Delta} \sum_{i=n_0}^{n_1} \mathbb{I}(\mathbf{f}_{\mathcal{A}_{n_0}}; y_{\mathbf{x}_{i+1}} \mid \mathcal{D}_i) \\ &\stackrel{(iv)}{=} \frac{\alpha_{n_1}}{\Delta} \sum_{i=n_0}^{n_1} \mathbb{I}(\mathbf{f}_{\mathcal{A}_{n_0}}; y_{\mathbf{x}_{i+1}} \mid \mathbf{y}_{\mathbf{x}_{n_0+1:i}}, \mathcal{D}_{n_0}) \\ &\stackrel{(v)}{=} \frac{\alpha_{n_1}}{\Delta} \mathbb{I}(\mathbf{f}_{\mathcal{A}_{n_0}}; \mathbf{y}_{\mathbf{x}_{n_0+1:n_1+1}} \mid \mathcal{D}_{n_0}) \\ &\leq \frac{\alpha_{n_1}}{\Delta} \max_{\substack{X \subseteq \mathcal{S} \\ |X|=\Delta}} \mathbb{I}(\mathbf{f}_{\mathcal{A}_{n_0}}; \mathbf{y}_X \mid \mathcal{D}_{n_0}) \\ &\stackrel{(vi)}{\leq} \frac{\alpha_{n_1}}{\Delta} \max_{\substack{X \subseteq \mathcal{S} \\ |X|=\Delta}} \mathbb{I}(\mathbf{f}_X; \mathbf{y}_X \mid \mathcal{D}_{n_0}) \\ &\stackrel{(vii)}{\leq} \frac{\alpha_{n_1}}{\Delta} \max_{\substack{X \subseteq \mathcal{S} \\ |X|=\Delta}} \mathbb{I}(\mathbf{f}_X; \mathbf{y}_X) \\ &= \alpha_{n_1} \frac{\gamma_{\Delta}}{\Delta} \end{aligned}$$

where (i) follows from the definition of the task complexity  $\alpha_{n_1}$  (cf. Definition C.6); (ii) uses the objective of ITL and that the posterior variance of Gaussians is independent of the realization and only depends on the *location* of observations; (iii) uses  $\mathcal{A}_{i+1} \subseteq \mathcal{A}_i$  and monotonicity of information gain; (iv) uses that the posterior variance of Gaussians is independent of the realization and only depends on the *location* of observations; (v) uses the chain rule of information gain; (vi) uses  $\mathbf{y}_X \perp \mathbf{f}_{\mathcal{A}_{n_0}} \mid \mathbf{f}_X$  and the data processing inequality. The conditional independence follows from the assumption that the observation noise is independent. Similarly,  $\mathbf{y}_X \perp \mathcal{D}_{n_0} \mid \mathbf{f}_X$  which implies (vii).

If  $n_0 = 0$ , then the bound before line (vi) simplifies to  $\alpha_{n_1} \gamma_{\mathcal{A}_0, \mathcal{S}}(\Delta)/\Delta$ .  $\square$

The result for VTL is stated, for simplicity, only for the case where the target space and sample space are fixed.

**Theorem C.13** (Convergence of Marginal Gain for VTL). *Assume that Assumptions B.1 and B.2 are satisfied. Then for any  $n \geq 1$ , if the sequence  $\{\mathbf{x}_i\}_{i=1}^n$  is generated by VTL,*

$$\Gamma_{n-1} \leq \frac{2\sigma^2\alpha_n}{n} \sum_{\mathbf{x}' \in \mathcal{A}} \gamma_{\{\mathbf{x}'\}, \mathcal{S}}(n). \quad (12)$$

We remark that  $\sum_{\mathbf{x}' \in \mathcal{A}} \gamma_{\{\mathbf{x}'\}, \mathcal{S}}(n) \leq |\mathcal{A}| \gamma_{\mathcal{A}, \mathcal{S}}(n)$ .

*Proof.* We have

$$\begin{aligned} \Gamma_{n-1} &= \frac{1}{n} \sum_{i=0}^{n-1} \Gamma_{n-1} \\ &\stackrel{(i)}{\leq} \frac{\alpha_n}{n} \sum_{i=0}^{n-1} \Gamma_i \\ &= \frac{\alpha_n}{n} \sum_{i=0}^{n-1} \left[ \text{tr} \text{Var}[\mathbf{f}_{\mathcal{A}} \mid \mathbf{y}_{1:i}] - \min_{\mathbf{x} \in \mathcal{S}} \text{tr} \text{Var}[\mathbf{f}_{\mathcal{A}} \mid \mathbf{y}_{1:i}, \mathbf{y}_{\mathbf{x}}] \right] \\ &\stackrel{(ii)}{=} \frac{\alpha_n}{n} \sum_{i=0}^{n-1} \left[ \text{tr} \text{Var}[\mathbf{f}_{\mathcal{A}} \mid \mathcal{D}_i] - \text{tr} \text{Var}[\mathbf{f}_{\mathcal{A}} \mid \mathcal{D}_{i+1}] \right] \\ &\stackrel{(iii)}{\leq} \frac{\sigma^2\alpha_n}{n} \sum_{\mathbf{x}' \in \mathcal{A}} \sum_{i=0}^{n-1} \log \left( \frac{\text{Var}[f_{\mathbf{x}'} \mid \mathcal{D}_n]}{\text{Var}[f_{\mathbf{x}'} \mid \mathcal{D}_{n+1}]} \right) \\ &= \frac{2\sigma^2\alpha_n}{n} \sum_{\mathbf{x}' \in \mathcal{A}} \sum_{i=0}^{n-1} \text{I}(f_{\mathbf{x}'}; \mathbf{y}_{\mathbf{x}_{n+1}} \mid \mathcal{D}_n) \\ &\stackrel{(iv)}{=} \frac{2\sigma^2\alpha_n}{n} \sum_{\mathbf{x}' \in \mathcal{A}} \sum_{i=0}^{n-1} \text{I}(f_{\mathbf{x}'}; \mathbf{y}_{\mathbf{x}_{n+1}} \mid \mathbf{y}_{\mathbf{x}_{1:n}}) \\ &\stackrel{(v)}{=} \frac{2\sigma^2\alpha_n}{n} \sum_{\mathbf{x}' \in \mathcal{A}} \text{I}(f_{\mathbf{x}'}; \mathbf{y}_{\mathbf{x}_{1:n}}) \\ &\leq \frac{2\sigma^2\alpha_n}{n} \sum_{\mathbf{x}' \in \mathcal{A}} \max_{\substack{X \subseteq \mathcal{S} \\ |X|=n}} \text{I}(f_{\mathbf{x}'}; \mathbf{y}_X) \\ &= \frac{2\sigma^2\alpha_n}{n} \sum_{\mathbf{x}' \in \mathcal{A}} \gamma_{\{\mathbf{x}'\}, \mathcal{S}}(n) \end{aligned}$$

where (i) follows from the definition of the task complexity  $\alpha_{n_1}$  (cf. Definition C.6); (ii) follows from the VTL decision rule and that the posterior variance of Gaussians is independent of the realization and only depends on the *location* of observations; (iii) follows from Lemma C.38 and monotonicity of variance; (iv) uses that the posterior variance of Gaussians is independent of the realization and only depends on the *location* of observations; and (v) uses the chain rule of mutual information. The remainder of the proof is analogous to the proof of Theorem C.12 (cf. Appendix C.5).  $\square$

**Keeping track of the task complexity online** In general, the task complexity  $\alpha_n$  may be larger than one in the “directed” setting (i.e., when  $\mathcal{S} \not\subseteq \mathcal{A}$ ). However, note that  $\alpha_n$  can easily be evaluated online by keeping track of the smallest  $\Gamma_i$  during previous rounds  $i$ .

## C.6 Proof of Theorem 3.2

We will now prove Theorem 3.2. We first prove the convergence of marginal variance within  $\mathcal{S}$  for ITL, before proving the convergence outside  $\mathcal{S}$  in Appendix C.6.1.

**Lemma C.14** (Uniform convergence of marginal variance within  $\mathcal{S}$  for ITL). *Assume that Assumptions B.1 and B.2 are satisfied. For any  $n \geq 0$  and  $\mathbf{x} \in \mathcal{A} \cap \mathcal{S}$ ,*

$$\sigma_n^2(\mathbf{x}) \leq 2\tilde{\sigma}^2 \cdot \Gamma_n. \quad (13)$$

*Proof.* We have

$$\begin{aligned} \sigma_n^2(\mathbf{x}) &= \text{Var}[f_{\mathbf{x}} \mid \mathcal{D}_n] - \underbrace{\text{Var}[f_{\mathbf{x}} \mid f_{\mathbf{x}}, \mathcal{D}_n]}_0 \\ &\stackrel{(i)}{=} \text{Var}[y_{\mathbf{x}} \mid \mathcal{D}_n] - \rho^2(\mathbf{x}) - (\text{Var}[y_{\mathbf{x}} \mid f_{\mathbf{x}}, \mathcal{D}_n] - \rho^2(\mathbf{x})) \\ &= \text{Var}[y_{\mathbf{x}} \mid \mathcal{D}_n] - \text{Var}[y_{\mathbf{x}} \mid f_{\mathbf{x}}, \mathcal{D}_n] \\ &\stackrel{(ii)}{\leq} \tilde{\sigma}^2 \log\left(\frac{\text{Var}[y_{\mathbf{x}} \mid \mathcal{D}_n]}{\text{Var}[y_{\mathbf{x}} \mid f_{\mathbf{x}}, \mathcal{D}_n]}\right) \\ &= 2\tilde{\sigma}^2 \cdot \text{I}(f_{\mathbf{x}}; y_{\mathbf{x}} \mid \mathcal{D}_n) \\ &\stackrel{(iii)}{\leq} 2\tilde{\sigma}^2 \cdot \text{I}(\mathbf{f}_{\mathcal{A}}; y_{\mathbf{x}} \mid \mathcal{D}_n) \\ &\stackrel{(iv)}{\leq} 2\tilde{\sigma}^2 \cdot \max_{\mathbf{x}' \in \mathcal{S}} \text{I}(\mathbf{f}_{\mathcal{A}}; y_{\mathbf{x}'} \mid \mathcal{D}_n) \\ &\stackrel{(v)}{=} 2\tilde{\sigma}^2 \cdot \max_{\mathbf{x}' \in \mathcal{S}} \text{I}(\mathbf{f}_{\mathcal{A}}; y_{\mathbf{x}'} \mid \mathbf{y}_{1:n}) \\ &= 2\tilde{\sigma}^2 \cdot \Gamma_n \end{aligned}$$

where (i) follows from the noise assumption (cf. Assumption B.2); (ii) follows from Lemma C.38 and using monotonicity of variance; (iii) follows from  $\mathbf{x} \in \mathcal{A}$  and monotonicity of information gain; (iv) follows from  $\mathbf{x} \in \mathcal{S}$ ; and (v) uses that the posterior variance of Gaussians is independent of the realization and only depends on the *location* of observations.  $\square$

### C.6.1 Convergence outside $\mathcal{S}$ for ITL

We will now show convergence of marginal variance to the irreducible uncertainty for points outside the sample space.

Our proof roughly proceeds as follows: We construct an ‘‘approximate Markov boundary’’ of  $\mathbf{x}$  in  $\mathcal{S}$ , and show (1) that the size of this Markov boundary is independent of  $n$ , and (2) that a small uncertainty reduction within the Markov boundary implies that the marginal variances at the Markov boundary and(!)  $\mathbf{x}$  are small.

**Definition C.15** (Approximate Markov boundary). For any  $\epsilon > 0$ ,  $n \geq 0$ , and  $\mathbf{x} \in \mathcal{X}$ , we denote by  $B_{n,\epsilon}(\mathbf{x})$  the smallest (multi-)subset of  $\mathcal{S}$  such that

$$\text{Var}[f_{\mathbf{x}} \mid \mathcal{D}_n, \mathbf{y}_{B_{n,\epsilon}(\mathbf{x})}] \leq \eta_{\mathcal{S}}^2(\mathbf{x}) + \epsilon. \quad (14)$$

We call  $B_{n,\epsilon}(\mathbf{x})$  an  $\epsilon$ -approximate Markov boundary of  $\mathbf{x}$  in  $\mathcal{S}$ .

Equation (14) is akin to the notion of the smallest Markov blanket in  $\mathcal{S}$  of some  $\mathbf{x} \in \mathcal{X}$  (called a *Markov boundary*) which is the smallest set  $\mathcal{B} \subseteq \mathcal{S}$  such that  $f_{\mathbf{x}} \perp \mathbf{f}_{\mathcal{S}} \mid \mathbf{f}_{\mathcal{B}}$ .

**Lemma C.16** (Existence of an approximate Markov boundary). *For any  $\epsilon > 0$ , let  $k$  be the smallest integer satisfying*

$$\frac{\gamma_k}{k} \leq \frac{\epsilon \lambda_{\min}(\mathbf{K}_{\mathcal{S}\mathcal{S}})}{2|\mathcal{S}|\sigma^2\tilde{\sigma}^2}. \quad (15)$$

*Then, for any  $n \geq 0$  and  $\mathbf{x} \in \mathcal{X}$ , there exists an  $\epsilon$ -approximate Markov boundary  $B_{n,\epsilon}(\mathbf{x})$  of  $\mathbf{x}$  in  $\mathcal{S}$  with size at most  $k$ .*

Lemma C.16 shows that for any  $\epsilon > 0$  there exists a universal constant  $b_{\epsilon}$  (with respect to  $n$  and  $\mathbf{x}$ ) such that

$$|B_{n,\epsilon}(\mathbf{x})| \leq b_{\epsilon} \quad \forall n \geq 0, \mathbf{x} \in \mathcal{X}. \quad (16)$$

We defer the proof of Lemma C.16 to Appendix C.6.3 where we also provide an algorithm to compute  $B_{n,\epsilon}(\mathbf{x})$ .

**Lemma C.17.** For any  $\epsilon > 0$ ,  $n \geq 0$ , and  $\mathbf{x} \in \mathcal{X}$ ,

$$\sigma_n^2(\mathbf{x}) \leq 2\sigma^2 \cdot \mathbb{I}(f_{\mathbf{x}}; \mathbf{y}_{B_{n,\epsilon}(\mathbf{x})} \mid \mathcal{D}_n) + \eta_S^2(\mathbf{x}) + \epsilon \quad (17)$$

where  $B_{n,\epsilon}(\mathbf{x})$  is an  $\epsilon$ -approximate Markov boundary of  $\mathbf{x}$  in  $\mathcal{S}$ .

*Proof.* We have

$$\begin{aligned} \sigma_n^2(\mathbf{x}) &= \text{Var}[f_{\mathbf{x}} \mid \mathcal{D}_n] - \eta_S^2(\mathbf{x}) + \eta_S^2(\mathbf{x}) \\ &\stackrel{(i)}{\leq} \text{Var}[f_{\mathbf{x}} \mid \mathcal{D}_n] - \text{Var}[f_{\mathbf{x}} \mid \mathbf{y}_{B_{n,\epsilon}(\mathbf{x})}, \mathcal{D}_n] + \eta_S^2(\mathbf{x}) + \epsilon \\ &\stackrel{(ii)}{\leq} \sigma^2 \log \left( \frac{\text{Var}[f_{\mathbf{x}} \mid \mathcal{D}_n]}{\text{Var}[f_{\mathbf{x}} \mid \mathbf{y}_{B_{n,\epsilon}(\mathbf{x})}, \mathcal{D}_n]} \right) + \eta_S^2(\mathbf{x}) + \epsilon \\ &= 2\sigma^2 \cdot \mathbb{I}(f_{\mathbf{x}}; \mathbf{y}_{B_{n,\epsilon}(\mathbf{x})} \mid \mathcal{D}_n) + \eta_S^2(\mathbf{x}) + \epsilon \end{aligned}$$

where (i) follows from the defining property of an  $\epsilon$ -approximate Markov boundary (cf. Equation (14)); and (ii) follows from Lemma C.38 and using monotonicity of variance.  $\square$

**Lemma C.18.** For any  $\epsilon > 0$ ,  $n \geq 0$ , and  $\mathbf{x} \in \mathcal{A}$ ,

$$\mathbb{I}(f_{\mathbf{x}}; \mathbf{y}_{B_{n,\epsilon}(\mathbf{x})} \mid \mathcal{D}_n) \leq \frac{b_\epsilon}{\bar{\kappa}_n(B_{n,\epsilon}(\mathbf{x}))} \Gamma_n \quad (18)$$

where  $B_{n,\epsilon}(\mathbf{x})$  is an  $\epsilon$ -approximate Markov boundary of  $\mathbf{x}$  in  $\mathcal{S}$ ,  $|B_{n,\epsilon}(\mathbf{x})| \leq b_\epsilon$ , and where  $\bar{\kappa}_n(\cdot) \stackrel{\text{def}}{=} \bar{\kappa}(\cdot \mid \mathbf{x}_{1:n})$  denotes the information ratio from Equation (8).

We remark that  $\bar{\kappa}_n(\cdot) > 0$  as is shown in Lemma C.7, and hence, the right-hand side of the inequality is well-defined.

*Proof.* We use the abbreviated notation  $B = B_{n,\epsilon}(\mathbf{x})$ . We have

$$\begin{aligned} \mathbb{I}(f_{\mathbf{x}}; \mathbf{y}_B \mid \mathcal{D}_n) &\stackrel{(i)}{\leq} \mathbb{I}(\mathbf{f}_{\mathcal{A}}; \mathbf{y}_B \mid \mathcal{D}_n) \\ &\stackrel{(ii)}{\leq} \frac{1}{\bar{\kappa}_{n,b_\epsilon}} \sum_{\bar{\mathbf{x}} \in B} \mathbb{I}(\mathbf{f}_{\mathcal{A}}; y_{\bar{\mathbf{x}}} \mid \mathcal{D}_n) \\ &\stackrel{(iii)}{\leq} \frac{b_\epsilon}{\bar{\kappa}_{n,b_\epsilon}} \max_{\bar{\mathbf{x}} \in B} \mathbb{I}(\mathbf{f}_{\mathcal{A}}; y_{\bar{\mathbf{x}}} \mid \mathcal{D}_n) \\ &\stackrel{(iv)}{\leq} \frac{b_\epsilon}{\bar{\kappa}_{n,b_\epsilon}} \max_{\bar{\mathbf{x}} \in \mathcal{S}} \mathbb{I}(\mathbf{f}_{\mathcal{A}}; y_{\bar{\mathbf{x}}} \mid \mathcal{D}_n) \\ &\stackrel{(v)}{=} \frac{b_\epsilon}{\bar{\kappa}_{n,b_\epsilon}} \max_{\bar{\mathbf{x}} \in \mathcal{S}} \mathbb{I}(\mathbf{f}_{\mathcal{A}}; y_{\bar{\mathbf{x}}} \mid \mathbf{y}_{1:n}) \\ &= \frac{b_\epsilon}{\bar{\kappa}_{n,b_\epsilon}} \Gamma_n \end{aligned}$$

where (i) follows from monotonicity of mutual information; (ii) follows from the definition of the information ratio  $\bar{\kappa}_{n,b_\epsilon}$  (cf. Equation (8)); (iii) follows from  $b \leq b_\epsilon$ ; (iv) follows from  $B \subseteq \mathcal{S}$ ; and (v) uses that the posterior variance of Gaussians is independent of the realization and only depends on the *location* of observations.  $\square$

*Proof of Theorem 3.2 for ITL.* The case where  $\mathbf{x} \in \mathcal{A} \cap \mathcal{S}$  is shown by Lemma C.14 with  $C = 2\bar{\sigma}^2$ .

To prove the more general result, fix any  $\mathbf{x} \in \mathcal{A}$  and  $\epsilon > 0$ . By Lemma C.16, there exists an  $\epsilon$ -approximate Markov boundary  $B_{n,\epsilon}(\mathbf{x})$  of  $\mathbf{x}$  in  $\mathcal{S}$  such that  $|B_{n,\epsilon}(\mathbf{x})| \leq b_\epsilon$ . We have

$$\begin{aligned} \sigma_n^2(\mathbf{x}) &\stackrel{(i)}{\leq} 2\sigma^2 \cdot \mathbb{I}(f_{\mathbf{x}}; \mathbf{y}_{B_{n,\epsilon}(\mathbf{x})} \mid \mathcal{D}_n) + \eta_S^2(\mathbf{x}) + \epsilon \\ &\stackrel{(ii)}{\leq} \frac{2\sigma^2 b_\epsilon}{\bar{\kappa}_n(B_{n,\epsilon}(\mathbf{x}))} \Gamma_n + \eta_S^2(\mathbf{x}) + \epsilon \end{aligned}$$

where (i) follows from Lemma C.17; and (ii) follows from Lemma C.18.

Let  $\epsilon = c \frac{\gamma \sqrt{n}}{\sqrt{n}}$  with  $c = 2|\mathcal{S}|\sigma^2\tilde{\sigma}^2/\lambda_{\min}(\mathbf{K}_{SS})$ . Then, by Equation (15),  $b_\epsilon$  can be bounded for instance by  $\sqrt{n}$ . Together with Theorem C.12 this implies for ITL that

$$\begin{aligned}\sigma_n^2(\mathbf{x}) &\leq \eta_S^2(\mathbf{x}) + 2\sigma^2\sqrt{n}\Gamma_n + c\gamma\sqrt{n}/\sqrt{n} \\ &\leq \eta_S^2(\mathbf{x}) + c'\gamma_n/\sqrt{n}\end{aligned}$$

for a constant  $c'$ , e.g.,  $c' = 2\sigma^2 + c$ .  $\square$

### C.6.2 Convergence outside $S$ for VTL

*Proof of Theorem 3.2 for VTL.* Analogously to Lemma C.17, we have

$$\begin{aligned}\sigma_n^2(\mathbf{x}) &= \text{Var}[f_{\mathbf{x}} | \mathcal{D}_n] - \eta_S^2(\mathbf{x}) + \eta_S^2(\mathbf{x}) \\ &\stackrel{(i)}{\leq} \text{Var}[f_{\mathbf{x}} | \mathcal{D}_n] - \text{Var}[f_{\mathbf{x}} | \mathbf{y}_{B_{n,\epsilon}(\mathbf{x})}, \mathcal{D}_n] + \eta_S^2(\mathbf{x}) + \epsilon\end{aligned}$$

where (i) follows from the defining property of an  $\epsilon$ -approximate Markov boundary (cf. Equation (14)). Further, we have

$$\begin{aligned}&\text{Var}[f_{\mathbf{x}} | \mathcal{D}_n] - \text{Var}[f_{\mathbf{x}} | \mathbf{y}_{B_{n,\epsilon}(\mathbf{x})}, \mathcal{D}_n] \\ &\stackrel{(i)}{\leq} \sum_{\tilde{\mathbf{x}} \in B_{n,\epsilon}(\mathbf{x})} (\text{Var}[f_{\mathbf{x}} | \mathcal{D}_n] - \text{Var}[f_{\mathbf{x}} | y_{\tilde{\mathbf{x}}}, \mathcal{D}_n]) \\ &\stackrel{(ii)}{\leq} \sum_{\tilde{\mathbf{x}} \in B_{n,\epsilon}(\mathbf{x})} (\text{tr Var}[f_{\mathcal{A}} | \mathbf{y}_{1:n}] - \text{tr Var}[f_{\mathcal{A}} | y_{\tilde{\mathbf{x}}}, \mathbf{y}_{1:n}]) \\ &\stackrel{(iii)}{\leq} b_\epsilon \Gamma_n\end{aligned}$$

where (i) follows from the submodularity of  $\psi_{\mathcal{A}}$ ; (ii) uses that the posterior variance of Gaussians is independent of the realization and only depends on the *location* of observations; and (iii) follows from the definition of  $\Gamma_n$  and Lemma C.16.

The remainder of the proof is analogous to the result for ITL, using Theorem C.13 to bound  $\Gamma_n$ .  $\square$

### C.6.3 Existence of an Approximate Markov Boundary

We now derive Lemma C.16 which shows the existence of an approximate Markov boundary of  $\mathbf{x}$  in  $\mathcal{S}$ .

**Lemma C.19.** *For any  $S \subseteq \mathcal{S}$  and  $k \geq 0$ , there exists  $B \subseteq S$  with  $|B| = k$  such that for all  $\mathbf{x}' \in S$ ,*

$$\text{Var}[f_{\mathbf{x}'} | \mathbf{y}_B] \leq 2\tilde{\sigma}^2 \frac{\gamma^k}{k}. \quad (19)$$

*Proof.* We choose  $B \subseteq S$  greedily using the acquisition function

$$\tilde{\mathbf{x}}_k \stackrel{\text{def}}{=} \arg \max_{\tilde{\mathbf{x}} \in S} \text{I}(\mathbf{f}_S; y_{\tilde{\mathbf{x}}} | \mathbf{y}_{B_{k-1}})$$

where  $B_k = \tilde{\mathbf{x}}_{1:k}$ . Note that this is the ‘‘undirected’’ special case of ITL, and hence, we have

$$\begin{aligned}\text{Var}[f_{\mathbf{x}'} | \mathbf{y}_{B_k}] &\stackrel{(i)}{\leq} 2\tilde{\sigma}^2 \Gamma_k \\ &\stackrel{(ii)}{\leq} 2\tilde{\sigma}^2 \frac{\gamma^k}{k}\end{aligned}$$

where (i) is due to Lemma C.14; and (ii) is due to Theorem C.12 and  $\alpha_{S,S}(k) \leq 1$ .  $\square$

**Lemma C.20.** *Given any  $\epsilon > 0$  and  $B \subseteq S \subseteq \mathcal{S}$  with  $|S| < \infty$ , such that for any  $\mathbf{x}' \in S$ ,*

$$\text{Var}[f_{\mathbf{x}'} | \mathbf{y}_B] \leq \frac{\epsilon \lambda_{\min}(\mathbf{K}_{SS})}{|S|\sigma^2}. \quad (20)$$

*Then for any  $\mathbf{x} \in \mathcal{X}$ ,*

$$\text{Var}[f_{\mathbf{x}} | \mathbf{y}_B] \leq \text{Var}[f_{\mathbf{x}} | \mathbf{f}_S] + \epsilon. \quad (21)$$

*Proof.* We will denote the right-hand side of Equation (20) by  $\epsilon'$ . We have

$$\begin{aligned}
& \text{Var}[f_{\mathbf{x}} \mid \mathbf{y}_B] \\
& \stackrel{(i)}{=} \mathbb{E}_{\mathbf{f}_S} [\text{Var}_{f_{\mathbf{x}}} [f_{\mathbf{x}} \mid \mathbf{f}_S, \mathbf{y}_B] \mid \mathbf{y}_B] \\
& \quad + \text{Var}_{\mathbf{f}_S} [\mathbb{E}_{f_{\mathbf{x}}} [f_{\mathbf{x}} \mid \mathbf{f}_S, \mathbf{y}_B] \mid \mathbf{y}_B] \\
& \stackrel{(ii)}{=} \text{Var}_{f_{\mathbf{x}}} [f_{\mathbf{x}} \mid \mathbf{f}_S, \mathbf{y}_B] + \text{Var}_{\mathbf{f}_S} [\mathbb{E}_{f_{\mathbf{x}}} [f_{\mathbf{x}} \mid \mathbf{f}_S, \mathbf{y}_B] \mid \mathbf{y}_B] \\
& \stackrel{(iii)}{=} \underbrace{\text{Var}_{f_{\mathbf{x}}} [f_{\mathbf{x}} \mid \mathbf{f}_S]}_{\text{irreducible uncertainty}} + \underbrace{\text{Var}_{\mathbf{f}_S} [\mathbb{E}_{f_{\mathbf{x}}} [f_{\mathbf{x}} \mid \mathbf{f}_S] \mid \mathbf{y}_B]}_{\text{reducible (epistemic) uncertainty}}
\end{aligned}$$

where (i) follows from the law of total variance; (ii) uses that the conditional variance of a Gaussian depends only on the location of observations and not on their value; and (iii) follows from  $f_{\mathbf{x}} \perp \mathbf{y}_B \mid \mathbf{f}_S$  since  $B \subseteq S$ . It remains to bound the reducible uncertainty.

Let  $h_{\mathbf{x}} : \mathbb{R}^d \rightarrow \mathbb{R}$ ,  $\mathbf{f}_S \mapsto \mathbb{E}[f_{\mathbf{x}} \mid \mathbf{f}_S]$  where we write  $d \stackrel{\text{def}}{=} |S|$ . Using the formula for the GP posterior mean, we have

$$h_{\mathbf{x}}(\mathbf{f}_S) = \mathbb{E}[f_{\mathbf{x}}] + \mathbf{z}^\top (\mathbf{f}_S - \mathbb{E}[\mathbf{f}_S])$$

where  $\mathbf{z} \stackrel{\text{def}}{=} \mathbf{K}_{SS}^{-1} \mathbf{K}_{S\mathbf{x}}$ . Because  $h$  is a linear function in  $\mathbf{f}_S$  we have for the reducible uncertainty that

$$\begin{aligned}
\text{Var}_{\mathbf{f}_S} [h_{\mathbf{x}}(\mathbf{f}_S) \mid \mathbf{y}_B] &= \mathbf{z}^\top \text{Var}[\mathbf{f}_S \mid \mathbf{y}_B] \mathbf{z} \\
& \stackrel{(i)}{\leq} d \cdot \mathbf{z}^\top \text{diag} \text{Var}[\mathbf{f}_S \mid \mathbf{y}_B] \mathbf{z} \\
& \stackrel{(ii)}{\leq} \epsilon' d \mathbf{z}^\top \mathbf{z} \\
& = \epsilon' d \mathbf{K}_{\mathbf{x}S} \mathbf{K}_{SS}^{-1} \mathbf{K}_{SS}^{-1} \mathbf{K}_{S\mathbf{x}} \\
& \leq \frac{\epsilon' d}{\lambda_{\min}(\mathbf{K}_{SS})} \mathbf{K}_{\mathbf{x}S} \mathbf{K}_{SS}^{-1} \mathbf{K}_{S\mathbf{x}} \\
& \stackrel{(iii)}{\leq} \frac{\epsilon' d \sigma^2}{\lambda_{\min}(\mathbf{K}_{SS})}
\end{aligned}$$

where (i) follows from Lemma C.37; (ii) follows from the assumption that  $\text{Var}[f_{\mathbf{x}'} \mid \mathbf{y}_B] \leq \epsilon'$  for all  $\mathbf{x}' \in S$ ; and (iii) follows from

$$\mathbf{K}_{\mathbf{x}S} \mathbf{K}_{SS}^{-1} \mathbf{K}_{S\mathbf{x}} \leq \mathbf{K}_{\mathbf{x}\mathbf{x}} = \sigma^2$$

since  $\mathbf{K}_{\mathbf{x}\mathbf{x}} - \mathbf{K}_{\mathbf{x}S} \mathbf{K}_{SS}^{-1} \mathbf{K}_{S\mathbf{x}} \geq 0$ . □

*Proof of Lemma C.16.* Let  $B \subseteq S$  be the set of size  $k$  generated by Lemma C.19 to satisfy  $\text{Var}[f_{\mathbf{x}'} \mid \mathbf{y}_B] \leq 2\tilde{\sigma}^2 \gamma_k / k$  for all  $\mathbf{x}' \in S$ . We have for any  $\mathbf{x} \in \mathcal{X}$ ,

$$\begin{aligned}
\text{Var}[f_{\mathbf{x}} \mid \mathcal{D}_n, \mathbf{y}_B] & \stackrel{(i)}{\leq} \text{Var}[f_{\mathbf{x}} \mid \mathbf{y}_B] \\
& \stackrel{(ii)}{\leq} \text{Var}[f_{\mathbf{x}} \mid \mathbf{f}_S] + \epsilon
\end{aligned}$$

where (i) follows from monotonicity of variance; and (ii) follows from Lemma C.20; using  $|S| < \infty$  and the condition on  $k$ . □

We remark that Lemma C.19 provides an algorithm (just “undirected” ITL!) to compute an approximate Markov boundary, and the set  $B$  returned by this algorithm is a valid approximate Markov boundary for all  $\mathbf{x} \in \mathcal{X}$ . One can simply swap-in ITL with target space  $\{\mathbf{x}\}$  for “undirected” ITL to obtain tighter (but instance-dependent) bounds on the size of the approximate Markov boundary.

#### C.6.4 Generalization to Continuous $S$ for Finite Dimensional RKHSs

In this subsection we generalize Theorem 3.2 to continuous sample spaces  $S$ . We will make the following assumption:

**Assumption C.21.** The RKHS of the kernel  $k$  is finite dimensional. In other words, the kernel  $k$  can be expressed as  $k(\mathbf{x}, \mathbf{x}') = \phi(\mathbf{x})^\top \phi(\mathbf{x}')$  for some feature map  $\phi : \mathcal{X} \rightarrow \mathbb{R}^d$  with  $d < \infty$ .

In the following, we will denote the design matrix of the sample space  $\mathcal{S}$  by  $\Phi \stackrel{\text{def}}{=} [\phi(\mathbf{x}) : \mathbf{x} \in \mathcal{S}]^\top \in \mathbb{R}^{|\mathcal{S}| \times d}$ , and we denote by  $\Pi_\Phi$  its orthogonal projection onto the orthogonal complement of the span of  $\Phi$ . In particular, it holds that

1.  $\Pi_\Phi \mathbf{v} = \mathbf{0}$  for all  $\mathbf{v} \in \text{span } \Phi$ , and
2.  $\Pi_\Phi \mathbf{v} = \mathbf{v}$  for all  $\mathbf{v} \in (\text{span } \Phi)^\perp$ .

Especially,  $\mathbf{v} \in \ker \Pi_\Phi$  if and only if  $\mathbf{v} \in \text{span } \Phi$ . This projection can be computed as follows:

**Lemma C.22.** *It holds that*

$$\Pi_\Phi = I - \Phi^\top (\Phi \Phi^\top)^{-1} \Phi. \quad (22)$$

*Proof.*  $\Phi^\top (\Phi \Phi^\top)^{-1} \Phi$  is the orthogonal projection onto the span of  $\Phi$  (see, e.g., [Strang, 2016](#), page 211).  $\square$

**Lemma C.23.** *Under Assumption C.21, the irreducible uncertainty  $\eta_{\mathcal{S}}^2(\mathbf{x})$  of  $\mathbf{x} \in \mathcal{X}$  is*

$$\eta_{\mathcal{S}}^2(\mathbf{x}) = \|\phi(\mathbf{x})\|_{\Pi_\Phi}^2 \quad (23)$$

where  $\|\mathbf{v}\|_A = \sqrt{\mathbf{v}^\top A \mathbf{v}}$  denotes the Mahalanobis distance.

*Proof.* This is an immediate consequence of the formula for the conditional variance of multivariate Gaussians (cf. [Appendix B.2](#)), applied to the linear kernel.  $\square$

Lemmas [C.22](#) and [C.23](#) imply that  $\eta_{\mathcal{S}}^2(\mathbf{x}^\parallel) = 0$  for all  $\mathbf{x}^\parallel \in \mathcal{X}$  with  $\phi(\mathbf{x}^\parallel) \in \text{span } \Phi$ . That is, the irreducible uncertainty is zero for points in the span of the sample space. In contrast, for points  $\mathbf{x}^\perp$  with  $\phi(\mathbf{x}^\perp) \in (\text{span } \Phi)^\perp$ , the irreducible uncertainty equals the initial uncertainty:  $\eta_{\mathcal{S}}^2(\mathbf{x}^\perp) = \sigma_0^2(\mathbf{x}^\perp)$ . The irreducible uncertainty of any other point  $\mathbf{x}$  can be computed by simple decomposition of  $\phi(\mathbf{x})$  into parallel and orthogonal components.

Assuming that Assumption [C.21](#) holds and given any (non-finite)  $\mathcal{S} \subseteq \mathcal{X}$ , there exists a basis  $\Omega_{\mathcal{S}} \subseteq \mathcal{X}$  in the space of embeddings  $\phi(\cdot)$  such that  $\text{span } \mathcal{S} = \text{span } \Omega_{\mathcal{S}}$  and  $|\Omega_{\mathcal{S}}| \leq d$ . The generalized existence of an approximate Markov boundary for continuous domains can then be shown analogously to [Lemma C.16](#):

**Lemma C.24** (Existence of an approximate Markov boundary for a continuous domain). *Let  $\mathcal{S}$  be any (continuous) subset of  $\mathcal{X}$  and let Assumption C.21 hold with  $d < \infty$ . Further, for any  $\epsilon > 0$ , let  $k$  be the smallest integer satisfying*

$$\frac{\gamma_k}{k} \leq \frac{\epsilon \lambda_{\min}(\mathbf{K}_{\Omega_{\mathcal{S}} \Omega_{\mathcal{S}}})}{2d\sigma^2\bar{\sigma}^2}. \quad (24)$$

*Then, for any  $n \geq 0$  and  $\mathbf{x} \in \mathcal{X}$ , there exists an  $\epsilon$ -approximate Markov boundary  $B_{n,\epsilon}(\mathbf{x})$  of  $\mathbf{x}$  in  $\mathcal{S}$  with size at most  $k$ .*

*Proof sketch.* The proof follows analogously to [Lemma C.16](#) by conditioning on the finite set  $\Omega_{\mathcal{S}}$  as opposed to  $\mathcal{S}$ .  $\square$

### C.7 Proof of Theorem 3.3

We first formalize the assumptions of [Theorem 3.3](#):

**Assumption C.25** (Regularity of  $f^*$ ). We assume that  $f^*$  is in a reproducing kernel Hilbert space  $\mathcal{H}_k(\mathcal{X})$  associated with a kernel  $k$  and has bounded norm, that is,  $\|f^*\|_k \leq B$  for some finite  $B \in \mathbb{R}$ .

**Assumption C.26** (Sub-Gaussian noise). We further assume that each  $\varepsilon_n$  from the noise sequence  $\{\varepsilon_n\}_{n=1}^\infty$  is conditionally zero-mean  $\rho(\mathbf{x}_n)$ -sub-Gaussian with known constants  $\rho(\mathbf{x}) > 0$  for all  $\mathbf{x} \in \mathcal{X}$ . Concretely,

$$\forall n \geq 1, \lambda \in \mathbb{R} : \quad \mathbb{E}[e^{\lambda \varepsilon_n} \mid \mathcal{D}_{n-1}] \leq \exp\left(\frac{\lambda^2 \rho^2(\mathbf{x}_n)}{2}\right)$$

where  $\mathcal{D}_{n-1}$  corresponds to the  $\sigma$ -algebra generated by the random variables  $\{\mathbf{x}_i, \varepsilon_i\}_{i=1}^{n-1}$  and  $\mathbf{x}_n$ .

We make use of the following foundational result, showing that under the above two assumptions the (misspecified) Gaussian process model from Section 3.1 is an all-time well-calibrated model of  $f^*$ :

**Lemma C.27** (Well-calibrated confidence intervals; Abbasi-Yadkori (2013); Chowdhury & Gopalan (2017)). *Pick  $\delta \in (0, 1)$  and let Assumptions C.25 and C.26 hold. Let*

$$\beta_n(\delta) = \|f^*\|_k + \rho\sqrt{2(\gamma_n + 1 + \log(1/\delta))}$$

where  $\rho = \max_{\mathbf{x} \in \mathcal{X}} \rho(\mathbf{x})$ .<sup>10</sup> Then, for all  $\mathbf{x} \in \mathcal{X}$  and  $n \geq 0$  jointly with probability at least  $1 - \delta$ ,

$$|f^*(\mathbf{x}) - \mu_n(\mathbf{x})| \leq \beta_n(\delta) \cdot \sigma_n(\mathbf{x})$$

where  $\mu_n(\mathbf{x})$  and  $\sigma_n^2(\mathbf{x})$  are mean and variance (as defined in Appendix B.2) of the GP posterior of  $f(\mathbf{x})$  conditional on the observations  $\mathcal{D}_n$ , pretending that  $\varepsilon_i$  is Gaussian with variance  $\rho^2(\mathbf{x}_i)$ .

The proof of Theorem 3.3 is a straightforward application of Lemma C.27 and Theorem 3.2:

*Proof of Theorem 3.3.* By Theorem 3.2, we have that for all  $\mathbf{x} \in \mathcal{A}$ ,

$$\sigma_n(\mathbf{x}) \leq \sqrt{\eta_S^2(\mathbf{x}) + \nu_{\mathcal{A},S}^2(n)} \leq \eta_S(\mathbf{x}) + \nu_{\mathcal{A},S}(n).$$

The result then follows by application of Lemma C.27.  $\square$

## C.8 Proof of Theorem 5.1

In this section, we derive our main result on Safe BO. In Appendix C.8.1, we give the definition of the reachable safe set  $\mathcal{R}$  and derive the conditions under which convergence to the reachable safe set is guaranteed. Then, in Appendix C.8.2, we prove Theorem 5.1.

**Notation** In the agnostic setting from Section 3.2 (i.e., under Assumptions C.25 and C.26), Lemma C.27 provides us with the following  $(1 - \delta)$ -confidence intervals (CIs)

$$\mathcal{C}_n(\mathbf{x}) \stackrel{\text{def}}{=} \mathcal{C}_{n-1}(\mathbf{x}) \cap [\mu_n(\mathbf{x}) \pm \beta_n(\delta) \cdot \sigma_n(\mathbf{x})] \quad (25)$$

where  $\mathcal{C}_{-1}(\mathbf{x}) = \mathbb{R}$ . We write  $u_n(\mathbf{x}) \stackrel{\text{def}}{=} \max \mathcal{C}_n(\mathbf{x})$ ,  $l_n(\mathbf{x}) \stackrel{\text{def}}{=} \min \mathcal{C}_n(\mathbf{x})$ , and  $w_n(\mathbf{x}) \stackrel{\text{def}}{=} u_n(\mathbf{x}) - l_n(\mathbf{x})$  for its upper bound, lower bound, and width, respectively.

We learn separate statistical models  $f$  and  $\{g_1, \dots, g_q\}$  for the ground truth objective  $f^*$  and ground truth constraints  $\{g_1^*, \dots, g_q^*\}$ . We write  $\mathcal{I} \stackrel{\text{def}}{=} \{f, 1, \dots, q\}$  and collect the constraints in  $\mathcal{I}_s \stackrel{\text{def}}{=} \{1, \dots, q\}$ . Without loss of generality, we assume that the confidence intervals include the ground truths with probability at least  $1 - \delta$  jointly for all  $i \in \mathcal{I}$ .<sup>11</sup> For  $i \in \mathcal{I}$ , denote by  $u_{n,i}, l_{n,i}, w_{n,i}, \eta_i, \beta_{n,i}$  the respective quantities. In the following, we do not explicitly denote the dependence of  $\beta_n$  on  $\delta$ .

To improve clarity, we will refer to the set of potential maximizers defined in Equation (5) as  $\mathcal{M}_n$  and denote by  $\mathcal{A}_n$  an arbitrary target space.

We point out the following corollary:

**Corollary C.28** (Safety). *With high probability, jointly for any  $n \geq 0$  and any  $i \in \mathcal{I}_s$ ,*

$$\forall \mathbf{x} \in \mathcal{S}_n : g_i^*(\mathbf{x}) \geq 0. \quad (26)$$

### C.8.1 Convergence to Reachable Safe Set

**Definition C.29** (Reachable safe set). Given any pessimistic safe set  $\mathcal{S} \subseteq \mathcal{X}$  and any  $\epsilon \geq 0$  and  $\beta \geq 0$ , we define the *reachable safe set* up to  $(\epsilon, \beta)$ -slack and its closure as

$$\begin{aligned} \mathcal{R}_{\epsilon,\beta}(\mathcal{S}) &\stackrel{\text{def}}{=} \mathcal{S} \cup \{\mathbf{x} \in \mathcal{X} \setminus \mathcal{S} \mid \\ &\quad g_i^*(\mathbf{x}) - \beta(\eta_i(\mathbf{x}; \mathcal{S}) + \epsilon) \geq 0 \text{ for all } i \in \mathcal{I}_s\} \\ \bar{\mathcal{R}}_{\epsilon,\beta}(\mathcal{S}) &\stackrel{\text{def}}{=} \lim_{n \rightarrow \infty} (\mathcal{R}_{\epsilon,\beta})^n(\mathcal{S}) \end{aligned}$$

where  $(\mathcal{R}_{\epsilon,\beta})^n$  denotes the  $n$ -th composition of  $\mathcal{R}_{\epsilon,\beta}$  with itself.

<sup>10</sup> $\beta_n(\delta)$  can be tightened adaptively (Emmenegger et al., 2023).

<sup>11</sup>This can be achieved by taking a union bound and rescaling  $\delta$ .



*Remark C.30.* Convergence of the safe set to the closure of the reachability operator can only be guaranteed for finite safe sets ( $|\mathcal{S}^*| < \infty$ ). The following proofs readily generalize to continuous domains by considering convergence within the  $k$ -th composition of the reachability operator with itself for some  $k < \infty$ . In this case the sample complexity grows with  $k$  rather than  $|\mathcal{S}^*|$ . The only required modification is to lift the assumption of Theorem C.12 that information is gained only while safe sets remain constant (i.e.,  $\mathcal{S}_{i+1} = \mathcal{S}_i$  for all  $i$ ). This assumption is straightforward to lift since for any  $n \geq 0$  and  $T \geq 1$ ,

$$\max_{\mathbf{x} \in \mathcal{S}_n} \Delta_{\mathcal{A}}(\mathbf{x} \mid \mathbf{x}_{1:n+T}) \leq \frac{1}{T} \sum_{t=1}^T \max_{\mathbf{x} \in \mathcal{S}_n} \Delta_{\mathcal{A}}(\mathbf{x} \mid \mathbf{x}_{1:n+t}) \leq \frac{1}{T} \sum_{t=1}^T \max_{\mathbf{x} \in \mathcal{S}_{n+t}} \Delta_{\mathcal{A}}(\mathbf{x} \mid \mathbf{x}_{1:n+t}) \leq \frac{\gamma T}{T},$$

using submodularity for the first inequality and the monotonicity of the safe set for the second inequality. In particular, this shows that one continues learning about points in the original safe set — even as the safe set grows.

We denote by  $\mathcal{S}_0$  the initial pessimistic safe set induced by the (prior) statistical model  $g$  (cf. Section 5) and write  $\bar{\mathcal{R}}_{\epsilon, \beta} \stackrel{\text{def}}{=} \bar{\mathcal{R}}_{\epsilon, \beta}(\mathcal{S}_0)$ .

**Lemma C.31** (Properties of the reachable safe set). *For all  $\mathcal{S}, \mathcal{S}' \subseteq \mathcal{X}$ ,  $\epsilon \geq 0$ , and  $\beta \geq 0$ :*

- (i)  $\mathcal{S}' \subseteq \mathcal{S} \implies \mathcal{R}_{\epsilon, \beta}(\mathcal{S}') \subseteq \mathcal{R}_{\epsilon, \beta}(\mathcal{S})$ ,
- (ii)  $\mathcal{R}_{\epsilon, \beta}(\mathcal{S}) \subseteq \mathcal{S} \implies \bar{\mathcal{R}}_{\epsilon, \beta}(\mathcal{S}) \subseteq \mathcal{S}$ , and
- (iii)  $\mathcal{R}_{0, 0}(\emptyset) = \bar{\mathcal{R}}_{0, 0} = \mathcal{S}^*$ .

*Proof (adapted from lemma 7.1 of Berkenkamp et al. (2021)).*

1. Let  $\mathbf{x} \in \mathcal{R}_{\epsilon, \beta}(\mathcal{S}')$ . If  $\mathbf{x} \in \mathcal{S}$  then  $\mathbf{x} \in \mathcal{R}_{\epsilon, \beta}(\mathcal{S})$ , so let  $\mathbf{x} \notin \mathcal{S}$ . Then, by definition, for all  $i \in \mathcal{I}_s$ ,  $f_i^*(\mathbf{x}) - \beta \eta_i(\mathbf{x}; \mathcal{S}') - \epsilon \geq 0$ . By the monotonicity of variance,  $\eta_i(\mathbf{x}; \mathcal{S}') \geq \eta_i(\mathbf{x}; \mathcal{S})$  for all  $i \in \mathcal{I}$ , and hence  $f_i^*(\mathbf{x}) - \beta \eta_i(\mathbf{x}; \mathcal{S}) - \epsilon \geq 0$  for all  $i \in \mathcal{I}_s$ . It follows that  $\mathbf{x} \in \mathcal{R}_{\epsilon, \beta}(\mathcal{S})$ .
2. By the monotonicity of variance,  $\eta_i(\mathbf{x}; \mathcal{R}_{\epsilon, \beta}(\mathcal{S})) \geq \eta_i(\mathbf{x}; \mathcal{S})$  for all  $\mathbf{x} \in \mathcal{X}$  and  $i \in \mathcal{I}$ . Thus, by definition of the safe region, we have that  $\mathcal{R}_{\epsilon, \beta}(\mathcal{R}_{\epsilon, \beta}(\mathcal{S})) \subseteq \mathcal{S}$ . The result follows by taking the limit.
3. The result follows directly from the definition of the true safe set  $\mathcal{S}^*$  (cf. Equation (4)).  $\square$

Clearly, we cannot expand the safe set beyond  $\bar{\mathcal{R}}_{0, 0}$ . The following is our main intermediate result, showing that either we expand the safe set at some point or the uncertainty converges to the irreducible uncertainty.

**Lemma C.32.** *Given any  $n_0 \geq 0$ ,  $\epsilon > 0$ , let  $n'$  be the smallest integer such that  $\nu_{n', \bar{\epsilon}^2} \leq \bar{\epsilon}$  where  $\bar{\epsilon} = \epsilon/2$ . Let  $\beta_{n_0+n'} = \max_{i \in \mathcal{I}_s} \beta_{n_0+n', i}$ . Assume that the sequence of target spaces is monotonically decreasing, i.e.,  $\mathcal{A}_{n+1} \subseteq \mathcal{A}_n$ . Then, we have with high probability (at least) one of*

$$\left( \forall \mathbf{x} \in \mathcal{A}_{n_0+n'}, \forall i \in \mathcal{I} : \right. \\ \left. w_{n_0+n', i}(\mathbf{x}) \leq \beta_{n_0+n'} [\eta_i(\mathbf{x}; \mathcal{S}_{n_0+n'}) + \epsilon] \right. \\ \left. \text{and } \mathcal{A}_{n_0+n'} \cap \mathcal{R}_{\epsilon, \beta_{n_0+n'}}(\mathcal{S}_{n_0+n'}) \subseteq \mathcal{S}_{n_0+n'} \right)$$

or  $|\mathcal{S}_{n_0+n'+1}| > |\mathcal{S}_{n_0}|$ .

*Proof.* Suppose that  $|\mathcal{S}_{n_0+n'+1}| = |\mathcal{S}_{n_0}|$ . Then, by Theorem 3.3 (using that the sequence of target spaces is monotonically decreasing), for any  $\mathbf{x} \in \mathcal{A}_{n_0+n'}$  and  $i \in \mathcal{I}$ ,

$$w_{n_0+n', i}(\mathbf{x}) \leq \beta_{n_0+n'} [\eta_i(\mathbf{x}; \mathcal{S}_{n_0+n'}) + \epsilon].$$

As  $\mathcal{S}_{n_0+n'+1} = \mathcal{S}_{n_0+n'}$  we have for all  $\mathbf{x} \in \mathcal{A}_{n_0+n'} \setminus \mathcal{S}_{n_0+n'}$  and  $i \in \mathcal{I}_s$ , with high probability that

$$\begin{aligned} 0 > l_{n_0+n', i}(\mathbf{x}) &\geq g_i^*(\mathbf{x}) - w_{n_0+n', i}(\mathbf{x}) \\ &\geq g_i^*(\mathbf{x}) - \beta_{n_0+n'} [\eta_i(\mathbf{x}; \mathcal{S}_{n_0+n'}) + \epsilon]. \end{aligned}$$

It follows that  $\mathcal{A}_{n_0+n'} \cap \mathcal{R}_{\epsilon, \beta_{n_0+n'}}(\mathcal{S}_{n_0+n'}) \subseteq \mathcal{S}_{n_0+n'}$ .  $\square$

To gather more intuition about the above lemma, consider the target space

$$\mathcal{E}_n \stackrel{\text{def}}{=} \widehat{\mathcal{S}}_n \setminus \mathcal{S}_n. \quad (27)$$

We call  $\mathcal{E}_n$  the *potential expanders* since it contains all points which might be safe, but are not yet known to be safe. Under this target space, the above lemma simplifies slightly:

**Lemma C.33.** *For any  $n \geq 0$  and  $\epsilon, \beta \geq 0$ , if  $\mathcal{E}_n \subseteq \mathcal{A}_n$  then with high probability,*

$$\mathcal{S}_n \cup (\mathcal{A}_n \cap \mathcal{R}_{\epsilon, \beta}(\mathcal{S}_n)) = \mathcal{R}_{\epsilon, \beta}(\mathcal{S}_n).$$

*Proof.* With high probability,  $\mathcal{R}_{\epsilon, \beta}(\mathcal{S}_n) \subseteq \widehat{\mathcal{S}}_n = \mathcal{S}_n \cup \mathcal{E}_n$ . The lemma is a direct consequence.  $\square$

The above lemmas can be combined to yield our main result of this subsection, establishing the convergence of ITL to the reachable safe set.

**Theorem C.34** (Convergence to reachable safe set). *For any  $\epsilon > 0$ , let  $n'$  be the smallest integer satisfying the condition of Lemma C.32, and define  $n^* \stackrel{\text{def}}{=} (|\mathcal{S}^*| + 1)n'$ . Let  $\bar{\beta}_{n^*} \geq \beta_{n, i}$  for all  $n \leq n^*, i \in \mathcal{I}_s$ . Assume that the sequence of target spaces is monotonically decreasing, i.e.,  $\mathcal{A}_{n+1} \subseteq \mathcal{A}_n$ . Then, the following inequalities hold jointly with probability at least  $1 - \delta$ :*

$$(i) \quad \forall n \geq 0, \forall i \in \mathcal{I}_s : g_i^*(\mathbf{x}_n) \geq 0,$$

*safety*

$$(ii) \quad \mathcal{A}_{n^*} \cap \bar{\mathcal{R}}_{\epsilon, \bar{\beta}_{n^*}} \subseteq \mathcal{S}_{n^*} \subseteq \bar{\mathcal{R}}_{0, 0} = \mathcal{S}^*,$$

*convergence to safe region*

$$(iii) \quad \forall \mathbf{x} \in \mathcal{A}_{n^*}, \forall i \in \mathcal{I} : w_{n^*, i}(\mathbf{x}) \leq \bar{\beta}_{n^*} \eta_i(\mathbf{x}; \bar{\mathcal{R}}_{\epsilon, \bar{\beta}_{n^*}}) + \epsilon,$$

*convergence of width*

$$(iv) \quad \forall \mathbf{x} \in \bar{\mathcal{R}}_{\epsilon, \bar{\beta}_{n^*}}, \forall i \in \mathcal{I} : \eta_i(\mathbf{x}; \bar{\mathcal{R}}_{\epsilon, \bar{\beta}_{n^*}}) = 0.$$

*convergence of width within safe region*

*Proof.* (i) is a direct consequence of Corollary C.28.  $\mathcal{S}_{n^*} \subseteq \mathcal{S}^*$  follows directly from the pessimistic safe set  $\mathcal{S}_{n^*}$  from (ii) being a subset of the true safe set  $\mathcal{S}^*$ . (iv) follows directly from the definition of irreducible uncertainty. Thus, it remains to establish  $\mathcal{A}_{n^*} \cap \bar{\mathcal{R}}_{\epsilon, \bar{\beta}_{n^*}} \subseteq \mathcal{S}_{n^*}$  and (iii).

Recall that with high probability  $|\mathcal{S}_n| \in [0, |\mathcal{S}^*|]$  for all  $n \geq 0$ . Thus, the size of the pessimistic safe set can increase at most  $|\mathcal{S}^*|$  many times. By Lemma C.32, using the assumption on  $n'$ , the size of the pessimistic safe set increases at least once every  $n'$  iterations, or else:

$$\begin{aligned} \forall \mathbf{x} \in \mathcal{A}_{n_0+n'}, \forall i \in \mathcal{I} : w_{n_0+n', i}(\mathbf{x}) &\leq \beta_{n_0+n'} [\eta_i(\mathbf{x}; \mathcal{S}_{n_0+n'}) + \epsilon] \\ \text{and } \mathcal{A}_{n_0+n'} \cap \mathcal{R}_{\epsilon, \beta_{n_0+n'}}(\mathcal{S}_{n_0+n'}) &\subseteq \mathcal{S}_{n_0+n'}. \end{aligned} \quad (28)$$

Because the safe set can expand at most  $|\mathcal{S}^*|$  many times, Equation (28) occurs eventually for some  $n_0 \leq |\mathcal{S}^*|n'$ . In this case, since  $\bar{\beta}_{n^*} \geq \beta_{n_0+n'}$  and  $\mathcal{A}_{n^*} \subseteq \mathcal{A}_{n_0+n'}$  (as  $n_0 + n' \leq n^*$ ) we have that

$$\begin{aligned} \mathcal{A}_{n^*} \cap \bar{\mathcal{R}}_{\epsilon, \bar{\beta}_{n^*}}(\mathcal{S}_{n_0+n'}) &\subseteq \mathcal{A}_{n_0+n'} \cap \mathcal{R}_{\epsilon, \beta_{n_0+n'}}(\mathcal{S}_{n_0+n'}) \\ &\subseteq \mathcal{S}_{n_0+n'}. \end{aligned}$$

By Lemma C.31 (ii), this implies

$$\mathcal{A}_{n^*} \cap \bar{\mathcal{R}}_{\epsilon, \bar{\beta}_{n^*}} \subseteq \mathcal{S}_{n_0+n'} \subseteq \mathcal{S}_{n^*}.$$

$\square$

We emphasize that Theorem C.34 holds for arbitrary target spaces  $\mathcal{A}_n$ . If additionally,  $\mathcal{E}_n \subseteq \mathcal{A}_n$  for all  $n \geq 0$  then by Lemma C.33, Theorem C.34 (ii) strengthens to  $\bar{\mathcal{R}}_{\epsilon, \bar{\beta}_{n^*}} \subseteq \mathcal{S}_{n^*}$ . Intuitively,  $\mathcal{E}_n \subseteq \mathcal{A}_n$  ensures that one aims to expand the safe set in *all* directions. Conversely, if  $\mathcal{E}_n \not\subseteq \mathcal{A}_n$  then one aims only to expand the safe set in the direction of  $\mathcal{A}_n$  (or not at all if  $\mathcal{A}_n \subseteq \mathcal{S}_n$ ).

**“Free” convergence guarantees in many applications** Theorem C.34 can be specialized to yield convergence guarantees in various settings by choosing an appropriate target space  $\mathcal{A}_n$ . Straightforward application of Theorem C.34 (informally) requires that the sequence of target spaces is monotonically decreasing (i.e.,  $\mathcal{A}_{n+1} \subseteq \mathcal{A}_n$ ), and that each target space  $\mathcal{A}_n$  is an “over-approximation” of the actual set of targeted points (such as the set of optimas in the Bayesian optimization setting). We discuss two such applications in the following.

1. *Pure expansion:* For example, for the target space  $\mathcal{E}_n$ , Theorem C.34 bounds the convergence of the safe set to the reachable safe set. In this case, the transductive active learning problem corresponds to the “pure expansion” setting, also addressed by the ISE baseline discussed in Section 5. The ISE baseline, however, does not establish convergence guarantees of the kind of Theorem C.34. Note that  $\mathcal{E}_n$  satisfies the (informal) requirements laid out previously, since it is monotonically decreasing by definition, and with high probability, any point  $\mathbf{x} \in \mathcal{S}^*$  that is not in  $\mathcal{S}_n$  is contained within  $\mathcal{E}_n$ .
2. *Level set estimation:* Given any  $\tau \in \mathbb{R}$ , we denote the (safe)  $\tau$ -level set of  $f^*$  by  $\mathcal{L}^\tau \stackrel{\text{def}}{=} \{\mathbf{x} \in \mathcal{S}^* \mid f^*(\mathbf{x}) = \tau\}$ . We define the *potential level set* as

$$\mathcal{L}_n^\tau \stackrel{\text{def}}{=} \{\mathbf{x} \in \widehat{\mathcal{S}}_n \mid l_n^f(\mathbf{x}) \leq \tau \leq u_n^f(\mathbf{x})\}. \quad (29)$$

That is,  $\mathcal{L}_n^\tau$  is the subset of the optimistic safe set  $\widehat{\mathcal{S}}_n$  where the  $\tau$ -level set of  $f^*$  may be located. Analogously to the potential expanders, it is straightforward to show that  $\mathcal{L}_n^\tau$  over-approximates the true  $\tau$ -level set and is monotonically decreasing.

We remark that our guarantees from this section also apply to the standard (“unsafe”) setting where  $\mathcal{S}^* = \mathcal{S}_0 = \mathcal{X}$ .

## C.8.2 Convergence to Safe Optimum

In this section, we specialize Theorem C.34 for the case that the target space contains the potential maximizers  $\mathcal{M}_n$  (cf. Equation (5)). It is straightforward to see that the sequence  $\mathcal{M}_n$  is monotonically decreasing (i.e.,  $\mathcal{M}_{n+1} \subseteq \mathcal{M}_n$ ). The following lemma shows that the potential maximizers over-approximate the set of safe maxima  $\mathcal{X}^* \stackrel{\text{def}}{=} \arg \max_{\mathbf{x} \in \mathcal{S}^*} f^*(\mathbf{x})$ .

**Lemma C.35** (Potential maximizers over-approximate safe maxima). *For all  $n \geq 0$  and with probability at least  $1 - \delta$ ,*

- (i)  $\mathbf{x} \in \mathcal{X}^*$  implies  $\mathbf{x} \in \mathcal{M}_n$  and
- (ii)  $\mathbf{x} \notin \mathcal{M}_n$  implies  $\mathbf{x} \notin \mathcal{X}^*$ .

*Proof.* If  $\mathbf{x} \notin \mathcal{M}_n$  then

$$u_{n,f}(\mathbf{x}) < \max_{\mathbf{x}' \in \mathcal{S}_n} l_{n,f}(\mathbf{x}') \leq \max_{\mathbf{x}' \in \mathcal{S}^*} l_{n,f}(\mathbf{x}')$$

where we used  $\mathcal{S}_n \subseteq \mathcal{S}^*$  with high probability, which directly implies with high probability that  $\mathbf{x} \notin \mathcal{X}^*$ .

For the other direction, if  $\mathbf{x} \in \mathcal{X}^*$  then

$$u_{n,f}(\mathbf{x}) \geq \max_{\mathbf{x}' \in \mathcal{S}^*} l_{n,f}(\mathbf{x}') \geq \max_{\mathbf{x}' \in \mathcal{S}_n} l_{n,f}(\mathbf{x}')$$

with high probability. □

We denote the set of optimal actions which are safe up to  $(\epsilon, \beta)$ -slack by

$$\mathcal{X}_{\epsilon,\beta}^* \stackrel{\text{def}}{=} \arg \max_{\mathbf{x} \in \mathcal{R}_{\epsilon,\beta}} f^*(\mathbf{x}),$$

and by  $f_{\epsilon,\beta}^*$  the maximum value attained by  $f^*$  at any of the points in  $\mathcal{X}_{\epsilon,\beta}^*$ . The regret can be expressed as

$$r_n(\bar{\mathcal{R}}_{\epsilon,\beta}) = f_{\epsilon,\beta}^* - f^*(\widehat{\mathbf{x}}_n)$$

The following theorem formalizes Theorem 5.1 and establishes convergence to the safe optimum.

**Theorem C.36** (Convergence to safe optimum). *For any  $\epsilon > 0$ , let  $n'$  be the smallest integer satisfying the condition of Lemma C.32, and define  $n^* \stackrel{\text{def}}{=} (|\mathcal{S}^*| + 1)n'$ . Let  $\bar{\beta}_{n^*} \geq \beta_{n,i}$  for all  $n \leq n^*$ ,  $i \in \mathcal{I}_s$ . Then, the following inequalities hold jointly with probability at least  $1 - \delta$ :*

$$(i) \quad \forall n \geq 0, \forall i \in \mathcal{I}_s : g_i^*(\mathbf{x}_n) \geq 0,$$

safety

$$(ii) \quad \forall n \geq n^* : r_n(\bar{\mathcal{R}}_{\epsilon, \bar{\beta}_{n^*}}) \leq \epsilon.$$

convergence to safe optimum

*Proof.* Fix any  $\mathbf{x}^* \in \mathcal{X}_{\epsilon, \bar{\beta}_{n^*}}^* \subseteq \bar{\mathcal{R}}_{\epsilon, \bar{\beta}_{n^*}}$ . Assume w.l.o.g. that  $\mathbf{x}^* \in \mathcal{M}_{n^*}$ .<sup>12</sup> Then, with high probability,

$$\begin{aligned} f_{\epsilon, \bar{\beta}_{n^*}}^* &= f^*(\mathbf{x}^*) \leq u_{n^*, f}(\mathbf{x}^*) \\ &= l_{n^*, f}(\mathbf{x}^*) + w_{n^*, f}(\mathbf{x}^*) \\ &\stackrel{(i)}{\leq} l_{n^*, f}(\hat{\mathbf{x}}_{n^*}) + w_{n^*, f}(\mathbf{x}^*) \\ &\leq f^*(\hat{\mathbf{x}}_{n^*}) + w_{n^*, f}(\mathbf{x}^*) \\ &\stackrel{(ii)}{\leq} f^*(\hat{\mathbf{x}}_{n^*}) + \epsilon \end{aligned}$$

where (i) follows from the definition of  $\hat{\mathbf{x}}_{n^*}$ ; and (ii) follows from Theorem C.34 and noting that  $\mathbf{x}^* \in \mathcal{M}_{n^*} \cap \bar{\mathcal{R}}_{\epsilon, \bar{\beta}_{n^*}}$ .

We have shown that  $f^*(\hat{\mathbf{x}}_{n^*}) \geq f_{\epsilon, \bar{\beta}_{n^*}}^* - \epsilon$ , which implies  $r_{n^*}(\bar{\mathcal{R}}_{\epsilon, \bar{\beta}_{n^*}}) \leq \epsilon$ . Since the upper- and lower-confidence bounds are monotonically decreasing / increasing, respectively, we have that for all  $n \geq n^*$ ,  $r_n(\bar{\mathcal{R}}_{\epsilon, \bar{\beta}_{n^*}}) \leq \epsilon$ .  $\square$

## C.9 Useful Facts and Inequalities

We denote by  $\preceq$  the Loewner partial ordering of symmetric matrices.

**Lemma C.37.** *Let  $\mathbf{A} \in \mathbb{R}^{n \times n}$  be a positive definite matrix with diagonal  $\mathbf{D}$ . Then,  $\mathbf{A} \preceq n\mathbf{D}$ .*

*Proof.* Equivalently, one can show  $n\mathbf{D} - \mathbf{A} \succeq \mathbf{0}$ . We write  $\mathbf{A} \stackrel{\text{def}}{=} \mathbf{D}^{1/2} \mathbf{Q} \mathbf{D}^{1/2}$ , and thus,  $\mathbf{Q} = \mathbf{D}^{-1/2} \mathbf{A} \mathbf{D}^{-1/2}$  is a positive definite symmetric matrix with all diagonal elements equal to 1. It remains to show that

$$n\mathbf{D} - \mathbf{A} = \mathbf{D}^{1/2} (n\mathbf{I} - \mathbf{Q}) \mathbf{D}^{1/2} \succeq \mathbf{0}.$$

Note that  $\sum_{i=1}^n \lambda_i(\mathbf{Q}) = \text{tr } \mathbf{Q} = n$ , and hence, all eigenvalues of  $\mathbf{Q}$  belong to  $(0, n)$ .  $\square$

**Lemma C.38.** *If  $a, b \in (0, M]$  for some  $M > 0$  and  $b \geq a$  then*

$$b - a \leq M \cdot \log\left(\frac{b}{a}\right). \quad (30)$$

*If additionally,  $a \geq M'$  for some  $M' > 0$  then*

$$b - a \geq M' \cdot \log\left(\frac{b}{a}\right). \quad (31)$$

*Proof.* Let  $f(x) \stackrel{\text{def}}{=} \log x$ . By the mean value theorem, there exists  $c \in (a, b)$  such that

$$\frac{1}{c} = f'(c) = \frac{f(b) - f(a)}{b - a} = \frac{\log b - \log a}{b - a} = \frac{\log\left(\frac{b}{a}\right)}{b - a}.$$

Thus,

$$b - a = c \cdot \log\left(\frac{b}{a}\right) < M \cdot \log\left(\frac{b}{a}\right).$$

<sup>12</sup>Otherwise, with high probability,  $f^*(\hat{\mathbf{x}}_{n^*}) > f_{\epsilon, \bar{\beta}_{n^*}}^*$ .

Under the additional condition that  $a \geq M'$ , we obtain

$$b - a = c \cdot \log\left(\frac{b}{a}\right) > M' \cdot \log\left(\frac{b}{a}\right).$$

□

## D Interpretations & Approximations of Principle (†)

We give a brief overview of interpretations and approximations of ITL, as well as alternative decision rules adhering to the fundamental principle (†).

The discussed interpretations of (†) differ mainly in how they quantify the “uncertainty” about  $\mathcal{A}$ . In the GP setting, this “uncertainty” is captured by the covariance matrix  $\Sigma$  of  $f_{\mathcal{A}}$ , and we consider two main ways of “scalarizing”  $\Sigma$ :

1. the total (marginal) variance  $\text{tr } \Sigma$ , and
2. the “generalized variance”  $|\Sigma|$ .

The generalized variance — which was originally suggested by Wilks (1932) as a generalization of variance to multiple dimensions — takes into account correlations. In contrast, the total variance discards all correlations between points in  $\mathcal{A}$ .

All discussed decision rules following principle (†) (i.e., ITL, VTL, MM-ITL) differ only in their weighting of the points in  $\mathcal{A}$ , and they coincide when  $|\mathcal{A}| = 1$ .

### D.1 Interpretations of ITL

We briefly discuss three interpretations of ITL.

**Minimizing generalized variance** In the GP setting, ITL can be equivalently characterized as minimizing generalized posterior variance:

$$\begin{aligned} \mathbf{x}_n &= \arg \max_{\mathbf{x} \in \mathcal{S}} I(\mathbf{f}_{\mathcal{A}}; y_{\mathbf{x}} \mid \mathcal{D}_n) \\ &= \arg \max_{\mathbf{x} \in \mathcal{S}} \frac{1}{2} \log \left( \frac{|\text{Var}[\mathbf{f}_{\mathcal{A}} \mid \mathcal{D}_{n-1}]|}{|\text{Var}[\mathbf{f}_{\mathcal{A}} \mid \mathcal{D}_{n-1}, y_{\mathbf{x}}]|} \right) \\ &= \arg \min_{\mathbf{x} \in \mathcal{S}} |\text{Var}[\mathbf{f}_{\mathcal{A}} \mid \mathcal{D}_{n-1}, y_{\mathbf{x}}]|. \end{aligned} \quad (32)$$

**Maximizing relevance and minimizing redundancy** An alternative interpretation of ITL is

$$I(\mathbf{f}_{\mathcal{A}}; y_{\mathbf{x}} \mid \mathcal{D}_n) = \underbrace{I(\mathbf{f}_{\mathcal{A}}; y_{\mathbf{x}})}_{\text{relevance}} - \underbrace{I(\mathbf{f}_{\mathcal{A}}; y_{\mathbf{x}}; \mathcal{D}_n)}_{\text{redundancy}} \quad (33)$$

where  $I(\mathbf{f}_{\mathcal{A}}; y_{\mathbf{x}}; \mathcal{D}_n) = I(\mathbf{f}_{\mathcal{A}}; y_{\mathbf{x}}) - I(\mathbf{f}_{\mathcal{A}}; y_{\mathbf{x}} \mid \mathcal{D}_n)$  denotes the *multivariate information gain* (cf. Appendix B). In this way, ITL can be seen as maximizing observation relevance while minimizing observation redundancy. This interpretation is common in the literature on feature selection (Peng et al., 2005; Vergara & Estévez, 2014; Beraha et al., 2019).

**Steepest descent in measure spaces** ITL can be seen as performing steepest descent in the space of probability measures over  $\mathbf{f}_{\mathcal{A}}$ , with the KL divergence as metric:

$$I(\mathbf{f}_{\mathcal{A}}; y_{\mathbf{x}} \mid \mathcal{D}_n) = \mathbb{E}_{y_{\mathbf{x}}} [\text{KL}(p(\mathbf{f}_{\mathcal{A}} \mid \mathcal{D}_n, y_{\mathbf{x}}) \| p(\mathbf{f}_{\mathcal{A}} \mid \mathcal{D}_n))].$$

That is, ITL finds the observation yielding the “largest update” to the current density.

### D.2 Interpretations of VTL

Quantifying the uncertainty about  $\mathbf{f}_{\mathcal{A}}$  by the marginal variance of points in  $\mathcal{A}$  rather than entropy (or generalized variance), the principle (†) leads to VTL. Note that if  $|\mathcal{A}| = 1$ , then VTL is equivalent to ITL. Unlike the similar, but more technical, TRUVAR algorithm proposed by Bogunovic et al. (2016), VTL does not require truncated variances, and hence, VTL can be applied to constrained settings (where  $\mathcal{A} \not\subseteq \mathcal{S}$ ) as well.

**Relationship to ITL** Note that the ITL criterion in the GP setting can be expressed as

$$\mathbf{x}_n = \arg \min_{\mathbf{x} \in \mathcal{S}} \text{tr} \log \text{Var}[\mathbf{f}_{\mathcal{A}} \mid \mathcal{D}_{n-1}, y_{\mathbf{x}}] \quad (34)$$

where for a positive semi-definite matrix  $\mathbf{A}$  with spectral decomposition  $\mathbf{A} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^\top$  we write  $\log \mathbf{A} = \mathbf{V} \log \mathbf{\Lambda} \mathbf{V}^\top$  for the logarithmic matrix function. To derive Equation (34) we use that  $\log |\mathbf{A}| = \sum_i \log \lambda_i(\mathbf{A}) = \text{tr} \log \mathbf{A}$ . Hence, ITL and VTL are identical up to a different weighting of the eigenvalues of the posterior covariance matrix.

**Minimizing a bound to the approximation error** Chowdhury & Gopalan (2017) (page 19) bound the approximation error  $|f^*(\mathbf{x}) - \mu_n(\mathbf{x})|$  by

$$\underbrace{|\mathbf{k}_t(\mathbf{x})^\top (\mathbf{K}_t + \mathbf{P}_t)^{-1} \boldsymbol{\varepsilon}_{1:t}|}_{\text{variance}} + \underbrace{|f^*(\mathbf{x}) - \mathbf{k}_t(\mathbf{x})^\top (\mathbf{K}_t + \mathbf{P}_t)^{-1} \mathbf{f}_{1:t}|}_{\text{bias}}$$

where  $\mathbf{k}_t(\mathbf{x}) \stackrel{\text{def}}{=} \mathbf{K}_{\mathbf{x}\mathbf{x}_{1:t}}$ ,  $\mathbf{K}_t \stackrel{\text{def}}{=} \mathbf{K}_{\mathbf{x}_{1:t}\mathbf{x}_{1:t}}$ , and  $\mathbf{P}_t \stackrel{\text{def}}{=} \mathbf{P}_{\mathbf{x}_{1:t}}$ . Similar to a standard bias-variance decomposition, the first term measures variance and the second term measures bias. Following Lemma C.27, VTL can be seen as greedily minimizing this bound to the approximation error (i.e., both bias and variance).

**Maximizing correlation to prediction targets weighted by their variance** It can be shown (see the proof below) that the VTL decision rule is equivalent to

$$\mathbf{x}_n = \arg \max_{\mathbf{x} \in \mathcal{S}} \sum_{\mathbf{x}' \in \mathcal{A}} \text{Var}[f_{\mathbf{x}'} \mid \mathcal{D}_{n-1}] \cdot \text{Cor}[f_{\mathbf{x}'}, y_{\mathbf{x}} \mid \mathcal{D}_{n-1}]^2. \quad (35)$$

That is, VTL maximizes the squared correlation between the next observation and the prediction targets, weighted by their variance. Intuitively, prediction targets are weighted by their variance since more can be learned about a prediction target with higher variance. This is precisely what leads to the “diverse” sample selection, and is akin to “uncertainty sampling” among the prediction targets and then selecting the observation which is most correlated with the selected prediction target.

*Proof.* Starting with the VTL objective, we have

$$\begin{aligned} \arg \min_{\mathbf{x} \in \mathcal{S}} \sum_{\mathbf{x}' \in \mathcal{A}} \text{Var}[f_{\mathbf{x}'} \mid \mathcal{D}_n, y_{\mathbf{x}}] &= \arg \min_{\mathbf{x} \in \mathcal{S}} \sum_{\mathbf{x}' \in \mathcal{A}} \left( \text{Var}[f_{\mathbf{x}'} \mid \mathcal{D}_n] - \frac{\text{Cov}[f_{\mathbf{x}'}, y_{\mathbf{x}} \mid \mathcal{D}_n]^2}{\text{Var}[y_{\mathbf{x}} \mid \mathcal{D}_n]} \right) \\ &= \arg \max_{\mathbf{x} \in \mathcal{S}} \sum_{\mathbf{x}' \in \mathcal{A}} \frac{\text{Var}[f_{\mathbf{x}'} \mid \mathcal{D}_n] \cdot \text{Cov}[f_{\mathbf{x}'}, y_{\mathbf{x}} \mid \mathcal{D}_n]^2}{\text{Var}[f_{\mathbf{x}'} \mid \mathcal{D}_n] \cdot \text{Var}[y_{\mathbf{x}} \mid \mathcal{D}_n]} + \text{const} \\ &= \arg \max_{\mathbf{x} \in \mathcal{S}} \sum_{\mathbf{x}' \in \mathcal{A}} \text{Var}[f_{\mathbf{x}'} \mid \mathcal{D}_n] \cdot \text{Cor}[f_{\mathbf{x}'}, y_{\mathbf{x}} \mid \mathcal{D}_n]^2 + \text{const}. \end{aligned}$$

□

### D.3 Mean Marginal ITL

MacKay (1992) previously proposed “mean-marginal” ITL (MM-ITL) in the setting where  $\mathcal{S} = \mathcal{X}$ , which selects

$$\mathbf{x}_n = \arg \max_{\mathbf{x} \in \mathcal{S}} \sum_{\mathbf{x}' \in \mathcal{A}} I(f_{\mathbf{x}'}; y_{\mathbf{x}} \mid \mathcal{D}_{n-1}) \quad (36)$$

and which simplifies in the GP setting to

$$\begin{aligned} \mathbf{x}_n &= \arg \max_{\mathbf{x} \in \mathcal{S}} \frac{1}{2} \sum_{\mathbf{x}' \in \mathcal{A}} \log \left( \frac{\text{Var}[f_{\mathbf{x}'} \mid \mathcal{D}_{n-1}]}{\text{Var}[f_{\mathbf{x}'} \mid \mathcal{D}_{n-1}, y_{\mathbf{x}}]} \right) \\ &= \arg \min_{\mathbf{x} \in \mathcal{S}} \sum_{\mathbf{x}' \in \mathcal{A}} \log \text{Var}[f_{\mathbf{x}'} \mid \mathcal{D}_{n-1}, y_{\mathbf{x}}] \\ &= \arg \min_{\mathbf{x} \in \mathcal{S}} \text{tr} \log \text{diag} \text{Var}[\mathbf{f}_{\mathcal{A}} \mid \mathcal{D}_{n-1}, y_{\mathbf{x}}]. \end{aligned} \quad (37)$$

Analogously to the derivation of Equation (34), this can also be expressed as

$$\mathbf{x}_n = \arg \min_{\mathbf{x} \in \mathcal{S}} |\text{diag Var}[\mathbf{f}_{\mathcal{A}} | \mathcal{D}_{n-1}, y_{\mathbf{x}}]|. \quad (38)$$

Effectively, MM-ITL ignores the mutual interaction between points in  $\mathcal{A}$ . As can be seen from Equation (37) and as is also mentioned by MacKay (1992), MM-ITL is equivalent to VTL up to a different weighting of the points in  $\mathcal{A}$ : instead of minimizing the average posterior variance (as in VTL), MM-ITL minimizes the average posterior log-variance. Under the lens of principle (†), this can be seen as minimizing the average marginal entropy of predictions within the target space:

$$\mathbf{x}_n = \arg \min_{\mathbf{x} \in \mathcal{S}} \sum_{\mathbf{x}' \in \mathcal{A}} \text{H}[f_{\mathbf{x}'} | \mathcal{D}_{n-1}, y_{\mathbf{x}}].$$

We remark that MM-ITL is a special case of EPIG (Bickford Smith et al., 2023, Appendix E.2).

**Not a generalization of uncertainty sampling** Unlike ITL, MM-ITL is not a generalization of uncertainty sampling. The reason is precisely that MM-ITL ignores the mutual interaction between points in  $\mathcal{A}$ . Consider the example where  $\mathcal{X} = \mathcal{S} = \mathcal{A} = \{1, \dots, 10\}$  where  $f_{1:9}$  are highly correlated while  $f_{10}$  is mostly independent of the other points. Visually, imagine a smooth function (i.e., under a Gaussian kernel) with points 1 through 9 close to each other and point 10 far away. Further, suppose that point 10 has a slightly larger marginal variance than the others. Then, MM-ITL would select one of the points 1 : 9 since this leads to the largest reduction in the marginal (log-)variances (i.e., to a small posterior “uncertainty”).<sup>13</sup> In contrast, ITL selects the point with the largest prior marginal variance (cf. Appendix C.1), point 10, since this leads to the largest reduction in entropy.<sup>14</sup>

**Similarity to VTL** Observe that the following decision rule is equivalent to VTL:

$$\mathbf{x}_n = \arg \max_{\mathbf{x} \in \mathcal{S}} \text{tr Var}[\mathbf{f}_{\mathcal{A}} | \mathcal{D}_{n-1}] - \text{tr Var}[\mathbf{f}_{\mathcal{A}} | \mathcal{D}_{n-1}, y_{\mathbf{x}}].$$

By Lemma C.38, for any  $\mathbf{x} \in \mathcal{S}$ , this objective value can be tightly lower- and upper-bounded (up to constant-factors) by

$$\begin{aligned} & \sum_{\mathbf{x}' \in \mathcal{A}} \log \left( \frac{\text{Var}[f_{\mathbf{x}'} | \mathcal{D}_{n-1}]}{\text{Var}[f_{\mathbf{x}'} | \mathcal{D}_{n-1}, y_{\mathbf{x}}]} \right) \\ &= 2 \sum_{\mathbf{x}' \in \mathcal{A}} \text{I}(f_{\mathbf{x}'}; y_{\mathbf{x}} | \mathcal{D}_{n-1}) \quad (\text{see MM-ITL}) \\ &\stackrel{(i)}{=} - \sum_{\mathbf{x}' \in \mathcal{A}} \log \left( 1 - \text{Cor}[f_{\mathbf{x}'}, y_{\mathbf{x}} | \mathcal{D}_{n-1}]^2 \right) \end{aligned} \quad (39)$$

where (i) is detailed in example 8.5.1 of Cover (1999). Thus, VTL and MM-ITL are closely related.

**Experiments** In our experiments with Gaussian processes from Figures 2 and 6, we observe that MM-ITL performs similarly to VTL and CTL.

**Convergence of uncertainty** We derive a convergence guarantee for MM-ITL which is analogous to the guarantees for ITL from Theorem C.12 and for VTL from Theorem C.13. We will assume for simplicity that  $\Gamma_n$  is monotonically decreasing in  $n$  (i.e.,  $\alpha_n \leq 1$ ).

**Theorem D.1** (Convergence of uncertainty reduction of MM-ITL). *Assume that Assumptions B.1 and B.2 are satisfied. Then for any  $n \geq 1$ , if  $\Gamma_0 \geq \dots \geq \Gamma_{n-1}$  and the sequence  $\{\mathbf{x}_i\}_{i=1}^n$  is generated by MM-ITL, then*

$$\Gamma_{n-1} \leq \frac{1}{n} \sum_{\mathbf{x}' \in \mathcal{A}} \gamma_n(\{\mathbf{x}'\}; \mathcal{S}). \quad (40)$$

*Proof.* We have

$$\Gamma_{n-1} = \frac{1}{n} \sum_{i=0}^{n-1} \Gamma_{n-1}$$

<sup>13</sup>This is because the observation reduces uncertainty not just about the observed point itself.

<sup>14</sup>Because points  $f_{1:9}$  are highly correlated,  $\text{H}[f_{1:9}]$  is already “small”.

$$\begin{aligned}
&\stackrel{(i)}{\leq} \frac{1}{n} \sum_{i=0}^{n-1} \Gamma_i \\
&= \frac{1}{n} \sum_{i=0}^{n-1} \max_{\mathbf{x} \in \mathcal{S}} \sum_{\mathbf{x}' \in \mathcal{A}} \mathbb{I}(f_{\mathbf{x}'}; y_{\mathbf{x}} \mid \mathcal{D}_n) \\
&\stackrel{(ii)}{=} \frac{1}{n} \sum_{i=0}^{n-1} \sum_{\mathbf{x}' \in \mathcal{A}} \mathbb{I}(f_{\mathbf{x}'}; y_{\mathbf{x}_{n+1}} \mid \mathcal{D}_n) \\
&\stackrel{(iii)}{=} \frac{1}{n} \sum_{\mathbf{x}' \in \mathcal{A}} \sum_{i=0}^{n-1} \mathbb{I}(f_{\mathbf{x}'}; y_{\mathbf{x}_{n+1}} \mid \mathbf{y}_{\mathbf{x}_{1:n}}) \\
&\stackrel{(iv)}{=} \frac{1}{n} \sum_{\mathbf{x}' \in \mathcal{A}} \mathbb{I}(f_{\mathbf{x}'}; \mathbf{y}_{\mathbf{x}_{1:n}}) \\
&\leq \frac{1}{n} \sum_{\mathbf{x}' \in \mathcal{A}} \max_{\substack{X \subseteq \mathcal{S} \\ |X|=n}} \mathbb{I}(f_{\mathbf{x}'}; \mathbf{y}_X) \\
&= \frac{1}{n} \sum_{\mathbf{x}' \in \mathcal{A}} \gamma_n(\{\mathbf{x}'\}; \mathcal{S})
\end{aligned}$$

where (i) follows by assumption; (ii) follows from the MM-ITL decision rule; (iii) uses that the posterior variance of Gaussians is independent of the realization and only depends on the *location* of observations; and (iv) uses the chain rule of mutual information. The remainder of the proof is analogous to the proof of Theorem C.12 (cf. Appendix C.5).  $\square$

Noting that

$$\mathbb{I}(f_{\mathbf{x}'}; y_{\mathbf{x}} \mid \mathcal{D}_{n-1}) \leq \sum_{\mathbf{x}' \in \mathcal{A}} \mathbb{I}(f_{\mathbf{x}'}; y_{\mathbf{x}} \mid \mathcal{D}_{n-1})$$

for any  $n \geq 1$ ,  $\mathbf{x} \in \mathcal{X}$ , and  $\mathbf{x}' \in \mathcal{A}$ , Theorem 3.2 can be readily rederived for MM-ITL (cf. Lemmas C.14 and C.18). Hence, the posterior marginal variances of MM-ITL can be bounded uniformly in terms of  $\Gamma_n$  analogously to ITL.

#### D.4 Correlation-based Transductive Learning

We will briefly look at the CTL (*Correlation-based TL*) decision rule

$$\mathbf{x}_n = \arg \max_{\mathbf{x} \in \mathcal{S}} \sum_{\mathbf{x}' \in \mathcal{A}} \text{Cor}[f_{\mathbf{x}}, f_{\mathbf{x}'} \mid \mathcal{D}_{n-1}] \quad (41)$$

which permits no interpretation under principle  $(\dagger)$ . However, if all correlations are non-negative (such as for the standard Gaussian and Matérn kernels), CTL is closely related to ITL, VTL, and MM-ITL (cf. Equations (35) and (39)). In this case, if  $|\mathcal{A}| = 1$ , then all decision rules coincide.

If, on the other hand, correlations may be negative then there is a crucial difference between CTL and the decision rules motivated from principle  $(\dagger)$ . Namely, decision rules following  $(\dagger)$  exhibit a preference for points with high *absolute* correlation to prediction targets as opposed to CTL which prefers points with high *positive* correlation. This stems from the intuitive fact that points with a strong negative correlation are equally informative as points with a strong positive correlation. Nevertheless, we observe in our experiments that (even for a linear kernel which does not ensure non-negative correlations) points selected by ITL and VTL are typically positively correlated with prediction targets.

#### D.5 Summary

We have seen that ITL, VTL, and MM-ITL can be seen as different interpretations of the same fundamental principle  $(\dagger)$ , with the approximations CTL. If  $|\mathcal{A}| = 1$  and correlations are non-negative, then all four decision rules are equivalent. CTL prefers points with high positive



correlation whereas the other decision rules prefer points with high absolute correlation. ITL is the only decision rule that takes into account the mutual dependence between points in  $\mathcal{A}$ , and VTL and MM-ITL differ only in their weighting of the posterior marginal variances of points in  $\mathcal{A}$ .

## E Stochastic Target Spaces

When the target space  $\mathcal{A}$  is large, it may be computationally infeasible to compute the exact objective. A natural approach to address this issue is to approximate the target space by a smaller set of size  $K$ .

**Discretizing the target space** One possibility is to discretize the target space  $\mathcal{A}$ . Compact target spaces can be addressed, e.g., via discretization arguments which are common in the Bayesian optimization literature (see, e.g., appendix C.1 of [Srinivas et al. \(2009\)](#)). That is, if the target space can be covered approximately using a finite (possibly large) set of points, the guarantees of Theorem 3.2 extend directly. This, however, can be impractical when the required size of discretization for sufficiently small approximation error is large. In the following, we briefly discuss a natural alternative approach based on sampling points from  $\mathcal{A}$ .

**Target distributions** Let  $\mathcal{A} \subseteq \mathcal{X}$  be a (possibly continuous) target space, and let  $\mathcal{P}_{\mathcal{A}}$  be a probability distribution supported on  $\mathcal{A}$ . In iteration  $n$ , a subset  $A_n$  of  $K$  points is sampled independently from  $\mathcal{A}$  according to the distribution  $\mathcal{P}_{\mathcal{A}}$  and the objective is computed on this subset. Formally, this amounts to a single-sample Monte Carlo approximation of

$$\mathbf{x}_n \in \arg \max_{\mathbf{x} \in \mathcal{S}} \mathbb{E}_{\mathcal{A} \stackrel{\text{iid}}{\sim} \mathcal{P}_{\mathcal{A}}} [\mathbb{I}(\mathbf{f}_{\mathcal{A}}; y_{\mathbf{x}} \mid \mathcal{D}_{n-1})]. \quad (42)$$

The convergence guarantees from Appendix C can be generalized to the setting of stochastic target spaces by estimating how often points “near” a specified prediction target  $\mathbf{x} \in \mathcal{A}$  are sampled.

**Definition E.1** ( $\gamma$ -ball at  $\mathbf{x}$ ). Given  $\mathbf{x} \in \mathcal{A}$  and any  $\gamma \geq 0$ , we call the set

$$B_{\gamma}(\mathbf{x}) \stackrel{\text{def}}{=} \{\mathbf{x}' \in \mathcal{X} \mid \|\mathbf{x} - \mathbf{x}'\| \leq \gamma\}$$

the  $\gamma$ -ball at  $\mathbf{x}$ . Further, we call  $\mathcal{P}_{\mathcal{A}}(B_{\gamma}(\mathbf{x}))$  the weight of that ball.

**Proposition E.2** (sketch). *Given any  $n \geq 1, K \geq 1, \gamma > 0$ , and  $\mathbf{x} \in \mathcal{A}$ , suppose that  $B_{\gamma}(\mathbf{x})$  has weight  $p > 0$ . Assume that the ITL objective is  $L_I$ -Lipschitz continuous. Then, with probability at least  $1 - \exp(-(1-p)n/(8K))$ ,*

$$\sigma_n^2(\mathbf{x}) \lesssim \eta_S^2(\mathbf{x}) + CL_I \gamma \frac{\gamma^{k(n)}}{\sqrt{k(n)}}$$

where  $k(n) \stackrel{\text{def}}{=} Kpn/2$ .

*Proof sketch.* Let  $Y_i \sim \text{Binom}(K, p)$  denote the random variable counting the number of occurrences of a point from  $B_{\gamma}(\mathbf{x})$  in  $A_i$ . Moreover, we write  $X_i \stackrel{\text{def}}{=} \mathbb{1}\{B_{\gamma}(\mathbf{x}) \cap A_i \neq \emptyset\}$ . Note that

$$\nu \stackrel{\text{def}}{=} \mathbb{E}X_i = \mathbb{P}(B_{\gamma}(\mathbf{x}) \cap A_i \neq \emptyset) = 1 - \mathbb{P}(Y_i = 0) = 1 - (1-p)^K \approx Kp$$

where the approximation stems from a first-order truncation of the Bernoulli series. Let  $X \stackrel{\text{def}}{=} \sum_{i=1}^n X_i$  with  $\mathbb{E}X = n\nu \approx Kpn$ .

Using the assumed Lipschitz-continuity of the objective, we know that  $\mathbb{I}(\mathbf{f}_{A'}; y_{\mathbf{x}} \mid \mathcal{D}_{n-1}) \leq L_I \gamma \mathbb{I}(\mathbf{f}_{\mathcal{A}}; y_{\mathbf{x}} \mid \mathcal{D}_{n-1})$  where  $A' \stackrel{\text{def}}{=} (A \setminus \{\mathbf{x}_{\gamma}\}) \cup \{\mathbf{x}\}$  and  $\mathbf{x}_{\gamma}$  is the point from the  $\gamma$ -ball at  $\mathbf{x}$ . The bound then follows analogously to Theorem 3.2.

Finally, by Chernoff’s bound, at least  $Kpn/2$  iterations contain a point from  $B_{\gamma}(\mathbf{x})$  with probability at least  $1 - \exp(-Kpn/8)$ .  $\square$

This strategy can also be used to generalize the VTL, CTL, and MM-ITL objectives to stochastic target spaces.

## F Closed-form Decision Rules

Below, we list the closed-form expressions for the ITL and VTL objectives. In the following,  $k_n$  denotes the kernel conditional on  $\mathcal{D}_n$ .

## ITL

$$\begin{aligned} I(\mathbf{f}_{\mathcal{A}}; y_{\mathbf{x}} \mid \mathcal{D}_{n-1}) &= \frac{1}{2} \log \left( \frac{\text{Var}[y_{\mathbf{x}} \mid \mathcal{D}_{n-1}]}{\text{Var}[y_{\mathbf{x}} \mid \mathbf{f}_{\mathcal{A}}, \mathcal{D}_{n-1}]} \right) \\ &= \frac{1}{2} \log \left( \frac{k_{n-1}(\mathbf{x}, \mathbf{x}) + \rho^2}{\hat{k}_{n-1}(\mathbf{x}, \mathbf{x}) + \rho^2} \right) \end{aligned} \quad (43)$$

where  $\hat{k}_n(\mathbf{x}, \mathbf{x}) = k_n(\mathbf{x}, \mathbf{x}) - \mathbf{k}_n(\mathbf{x}, \mathcal{A}) \mathbf{K}_n(\mathcal{A}, \mathcal{A})^{-1} \mathbf{k}_n(\mathcal{A}, \mathbf{x})$ .

## VTL

$$\text{tr Var}[\mathbf{f}_{\mathcal{A}} \mid \mathcal{D}_{n-1}, y_{\mathbf{x}}] = \sum_{\mathbf{x}' \in \mathcal{A}} \left( k_{n-1}(\mathbf{x}', \mathbf{x}') - \frac{k_{n-1}(\mathbf{x}, \mathbf{x}')^2}{k_{n-1}(\mathbf{x}, \mathbf{x}) + \rho^2} \right).$$

## G Computational Complexity

Evaluating the acquisition function of ITL in round  $n$  requires computing for each  $\mathbf{x} \in \mathcal{S}$ ,

$$\begin{aligned} I(\mathbf{f}_{\mathcal{A}}; y_{\mathbf{x}} \mid \mathcal{D}_n) &= \frac{1}{2} \log \left( \frac{|\text{Var}[\mathbf{f}_{\mathcal{A}} \mid \mathcal{D}_n]|}{|\text{Var}[\mathbf{f}_{\mathcal{A}} \mid y_{\mathbf{x}}, \mathcal{D}_n]|} \right) && \text{(forward)} \\ &= \frac{1}{2} \log \left( \frac{\text{Var}[y_{\mathbf{x}} \mid \mathcal{D}_n]}{\text{Var}[y_{\mathbf{x}} \mid \mathbf{f}_{\mathcal{A}}, \mathcal{D}_n]} \right) && \text{(backward)}. \end{aligned}$$

Let  $|\mathcal{S}| = m$  and  $|\mathcal{A}| = k$ . Then, the forward method has complexity  $O(m \cdot k^3)$ . For the backward method, observe that the variances are scalar and the covariance matrix  $\text{Var}[\mathbf{f}_{\mathcal{A}} \mid \mathcal{D}_n]$  only has to be inverted once for all points  $\mathbf{x}$ . Thus, the backward method has complexity  $O(k^3 + m)$ .

When the size  $m$  of  $\mathcal{S}$  is relatively small (and hence, all points in  $\mathcal{S}$  can be considered during each iteration of the algorithm), GP inference corresponds simply to computing conditional distributions of a multivariate Gaussian. The performance can therefore be improved by keeping track of the full posterior distribution over  $\mathbf{f}_{\mathcal{S}}$  of size  $O(m^2)$  and conditioning on the latest observation during each iteration of the algorithm. In this case, after each observation the posterior can be updated at a cost of  $O(m^2)$  which does not grow with the time  $n$ , unlike classical GP inference.

Overall, when  $m$  is small, the computational complexity of ITL is  $O(k^3 + m^2)$ . When  $m$  is large (or possibly infinite) and a subset of  $\tilde{m}$  points is considered in a given iteration, the computational complexity of ITL is  $O(k^3 + \tilde{m} \cdot n^3)$ , neglecting the complexity of selecting the  $\tilde{m}$  candidate points. In the latter case, the computational cost of ITL is dominated by the cost of GP inference.

[Khanna et al. \(2017\)](#) discuss distributed and stochastic approximations of greedy algorithms to (weakly) submodular problems that are also applicable to ITL.

## H Additional GP Experiments & Details

We use homoscedastic Gaussian noise with standard deviation  $\rho = 0.1$  and a discretization of  $\mathcal{X} = [-3, 3]^2$  of size 2500. Uncertainty bands correspond to one standard error over 10 random seeds.

**Additional experiments** Figure 6 includes the following additional experiments:

1. *Extrapolation Setting* ( $\mathcal{A} \cap \mathcal{S} = \emptyset$ ): Right experiment from Figure 2 under the Gaussian kernel. ITL has a similar advantage as in the setting shown in Figure 3.
2. *Heteroscedastic Noise*: Left experiment from Figure 2 under the Gaussian kernel with heteroscedastic Gaussian noise

$$\rho(\mathbf{x}) = \begin{cases} 1 & \text{if } \mathbf{x} \in [-\frac{1}{2}, \frac{1}{2}]^2 \\ 0.1 & \text{otherwise} \end{cases}.$$

If observation noise is heteroscedastic, in considering *posterior* rather than *prior* uncertainty, ITL avoids points with high aleatoric uncertainty, which accelerates learning.

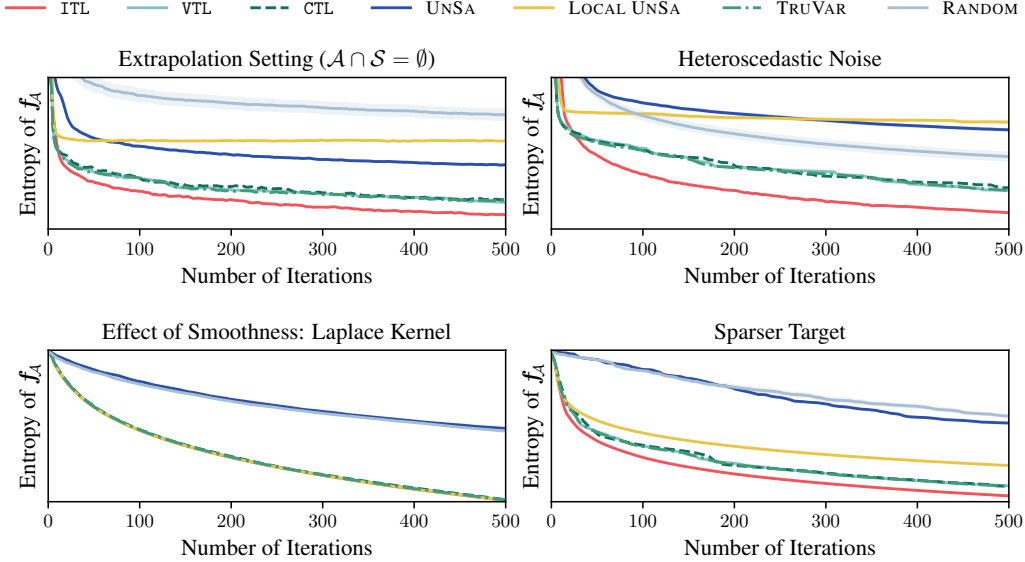


Figure 6: Additional GP experiments

3. *Effect of Smoothness:* Experiment from Figure 3 under the Laplace kernel. All algorithms except for US and RANDOM perform equally well. This validates our claims from Section 3.3: in the extreme non-smooth case of a Laplace kernel and  $\mathcal{A} \subseteq \mathcal{S}$ , points outside  $\mathcal{A}$  do not provide any additional information, and ITL and “local” UNSA coincide.
4. *Sparser Target:* Experiment from Figure 3 under the Gaussian kernel, but with domain extended to  $\mathcal{X} = [-10, 10]^2$ .

**Hyperparameters of TRUVAR** As suggested by Bogunovic et al. (2016), we use  $\tilde{\eta}_{(1)}^2 = 1$ ,  $r = 0.1$ , and  $\delta = 0$  (even though the theory only holds for  $\delta > 0$ ). The TRUVAR baseline only applies when  $\mathcal{A} \subseteq \mathcal{S}$  (cf. Section 6).

**Smoothing to reduce numerical noise** Applied running average with window 5 to entropy curves of Figures 2 and 6 to smoothen out numerical noise.

## I Alternative Settings for Active Fine-Tuning

In our main experiments, we consider the setting  $\mathcal{A} \cap \mathcal{S} = \emptyset$ , i.e., the prediction targets cannot be used for fine-tuning since their labels are not known. This setting is particularly relevant for practical applications where the model is fine-tuned dynamically at test time to each prediction target (or a small set of prediction target). Put differently, in this “transductive” setting, extrapolation to new prediction targets happens at *test-time* with knowledge of the prediction target(s). This is in contrast to a more traditional “inductive” setting, where extrapolation happens at *train-time* without knowledge of the concrete prediction targets, but under the assumption of samples from (or knowledge of) the target distribution. In the following, we briefly survey two settings motivated from an “inductive” perspective.

### I.1 Prediction Targets are Contained in Sample Space: $\mathcal{A} \subseteq \mathcal{S}$

If labels can be obtained cheaply, one can also fine-tune on the prediction targets directly, i.e.,  $\mathcal{A} \subseteq \mathcal{S}$ . Note, however, that the set  $\mathcal{A}$  is still assumed to be small (e.g.,  $|\mathcal{A}| = 100$  in the CIFAR-100 experiment). We perform an experiment in this setting and report the results in Figure 7. The experiment shows that — similarly to the GP experiment from Figure 2 — there can be *additional value* in fine-tuning the model on relevant data selected from  $\mathcal{S}$  beyond simply fine-tuning the model on  $\mathcal{A}$ .

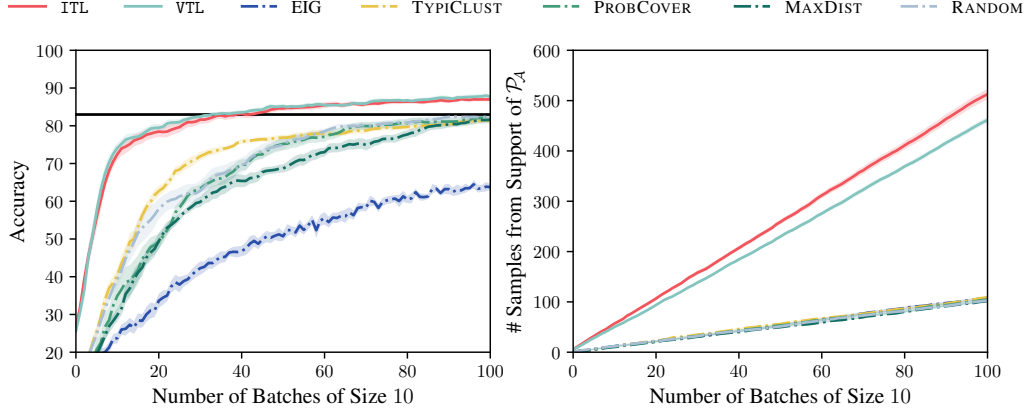


Figure 7: Evaluation of CIFAR-100 experiment in the setting  $\mathcal{A} \subseteq \mathcal{S}$ , i.e., one can also sample from the 100 prediction targets  $\mathcal{A}$ . The solid black line denotes the performance of the model fine-tuned on all of  $\mathcal{A}$ . This experiment shows that there is *additional value* in fine-tuning the model on relevant data from  $\mathcal{S}$  beyond simply fine-tuning the model on  $\mathcal{A}$ . The baselines are summarized in Appendix J.5

## I.2 Active Domain Adaptation

Active DA (Rai et al., 2010; Saha et al., 2011; Berling & Urner, 2015) studies the problem of selecting the most informative samples from a (large) target domain  $\mathcal{A}$ , given a model trained on a source domain  $\mathcal{S}$ . This problem can be cast as an instance of transductive active learning with target space  $\mathcal{A}$  and sample space  $\mathcal{S}' = \mathcal{S} \cup \mathcal{A}$  where the model is already conditioned on all of  $\mathcal{S}$ . This is slightly different from the setting considered in Section 4 where  $\mathcal{A}$  is small and not necessarily part of the sample space. We hypothesize that ITL behaves similarly to recent work on active DA (Su et al., 2020; Prabhu et al., 2021; Fu et al., 2021): querying informative and diverse samples from  $\mathcal{A}$  that are dissimilar to  $\mathcal{S}$ . Evaluating ITL and VTL empirically in this setting is a promising direction for future work.

## J Additional NN Experiments & Details

We outline the active fine-tuning of NNs in Algorithm 1.

---

### Algorithm 1 Active Fine-Tuning of NNs

---

**Given:** initialized or pre-trained model  $f$ , *small* sample  $A \sim \mathcal{P}_A$   
initialize dataset  $\mathcal{D} = \emptyset$   
**repeat**  
  sample  $S \sim \mathcal{P}_S$   
  subsample target space  $A' \stackrel{\text{u.a.r.}}{\sim} A$   
  initialize batch  $B = \emptyset$   
  compute kernel matrix  $\mathbf{K}$  over domain  $[S, A']$   
  **repeat**  $b$  **times**  
    compute acquisition function w.r.t.  $A'$ , based on  $\mathbf{K}$   
    add maximizer  $x \in S$  of acquisition function to  $B$   
    update conditional kernel matrix  $\mathbf{K}$   
  obtain labels for  $B$  and add to dataset  $\mathcal{D}$   
  update  $f$  using data  $\mathcal{D}$

---

In Appendix J.1, we detail metrics and hyperparameters. We describe in Appendices J.2 and J.3 how to compute the (initial) conditional kernel matrix  $\mathbf{K}$ , and in Appendix J.4 how to update this matrix  $\mathbf{K}$  to obtain conditional embeddings for batch selection.

In Appendix J.5, we show that ITL and CTL significantly outperform a wide selection of commonly used heuristics. In Appendices J.6 and J.7, we conduct additional experiments and ablations.

Table 1: Hyperparameter summary of NN experiments. (\*) we train until convergence on oracle validation accuracy.

	MNIST	CIFAR-100
$\rho$	0.01	1
$M$	30	100
$m$	3	10
$k$	1 000	1 000
batch size $b$	1	10
# of epochs	(*)	5
learning rate	0.001	0.001

Hübötter et al. (2024) discusses additional motivation and related work that has previously studied active fine-tuning, but which has largely focused on the training algorithm rather than data selection.

### J.1 Experiment Details

We evaluate the accuracy with respect to  $\mathcal{P}_A$  using a Monte Carlo approximation with out-of-sample data:

$$\text{accuracy}(\hat{\theta}) \approx \mathbb{E}_{(x,y) \sim \mathcal{P}_A} \mathbb{1}\{y = \arg \max_i f_i(x; \hat{\theta})\}.$$

We provide an overview of the hyperparameters used in our NN experiments in Table 1. The effect of noise standard deviation  $\rho$  is small for all tested  $\rho \in [1, 100]$  (cf. ablation study in Table 2).<sup>15</sup>  $M$  denotes the size of the sample  $A \sim \mathcal{P}_A$ . In each iteration, we select the target space  $\mathcal{A} \leftarrow A'$  as a random subset of  $m$  points from  $A$ .<sup>16</sup> We provide an ablation over  $m$  in Appendix J.6.

During each iteration, we select the batch  $B$  according to the decision rule from a random sample from  $\mathcal{P}_S$  of size  $k$ .<sup>17</sup>

Since we train the MNIST model from scratch, we train from random initialization until convergence on oracle validation accuracy.<sup>18</sup> We do this to stabilize the learning curves, and provide the least biased (due to the training algorithm) results. For CIFAR-100, we train for 5 epochs (starting from the previous iterations' model) which we found to be sufficient to obtain good performance.

We use the ADAM optimizer (Kingma & Ba, 2014). In our CIFAR-100 experiments, we use a pre-trained EfficientNet-B0 (Tan & Le, 2019), and fine-tune the final and penultimate layers. We freeze earlier layers to prevent overfitting to the “few-shot” training data.

To prevent numerical inaccuracies when computing the ITL objective, we optimize

$$I(\mathbf{y}_A; y_x \mid \mathcal{D}_{n-1}) = \frac{1}{2} \log \left( \frac{\text{Var}[y_x \mid \mathcal{D}_{n-1}]}{\text{Var}[y_x \mid \mathbf{y}_A, \mathcal{D}_{n-1}]} \right) \quad (44)$$

instead of Equation (43), which amounts to adding  $\rho^2$  to the diagonal of the covariance matrix before inversion. This appears to improve numerical stability, especially when using gradient embeddings.<sup>19</sup>

<sup>15</sup>We use a larger noise standard deviation  $\rho$  in CIFAR-100 to stabilize the numerics of batch selection via conditional embeddings (cf. Table 2).

<sup>16</sup>This appears to improve the training, likely because it prevents overfitting to peculiarities in the finite sample  $A$  (cf. Figure 16).

<sup>17</sup>In large-scale problems, the work of Coleman et al. (2022) suggests to use an (approximate) nearest neighbor search to select the (large) candidate set rather than sampling u.a.r. from  $\mathcal{P}_S$ . This can be a viable alternative to simply increasing  $k$  and suggests future work.

<sup>18</sup>That is, to stop training as soon as accuracy on a validation set from  $\mathcal{P}_A$  decreases in an epoch.

<sup>19</sup>In our experiments, we observe that the effect of various choices of  $\rho$  on this slight adaptation of the ITL decision rule has negligible impact on performance. The more prominent effect of  $\rho$  appears to arise from the batch selection via conditional embeddings (cf. Table 2).

In our experiments, we use last-layer neural tangent embeddings<sup>20</sup> and  $\Sigma = I$  to evaluate ITL and VTL, and select inputs for labeling and training  $f$ . Notably, we use this linear Gaussian approximation of  $f$  only to guide the active data selection and not for inference.

## J.2 Embeddings and Kernels

Using a neural network to parameterize  $f$ , we evaluate the canonical approximations of  $f$  by a stochastic process in the following.

An embedding  $\phi(x)$  is a latent representation of an input  $x$ . Collecting the embeddings as rows in the design matrix  $\Phi$  of a set of inputs  $X$ , one can approximate the network by the linear function  $f_X = \Phi\beta$  with weights  $\beta$ . Approximating the weights by  $\beta \sim \mathcal{N}(\mu, \Sigma)$  implies that  $f_X \sim \mathcal{N}(\Phi\mu, \Phi\Sigma\Phi^\top)$ . The covariance matrix  $K_{XX} = \Phi\Sigma\Phi^\top$  can be succinctly represented in terms of its associated kernel  $k(x, x') = \phi(x)^\top \Sigma \phi(x')$ . Here,

- $\phi(x)$  is the latent representation of  $x$ , and
- $\Sigma$  captures the dependencies in the latent space.

While any choice of embedding  $\phi$  is possible, the following are common choices:

1. *Last-Layer*: A common choice for  $\phi(x)$  is the representation of  $x$  from the penultimate layer of the neural network (Holzmüller et al., 2023). Interpreting the early layers as a feature encoder, this uses the low-dimensional feature map akin to random feature methods (Rahimi & Recht, 2007).
2. *Output Gradients (eNTK)*: Another common choice is  $\phi(x) = \nabla_{\theta} f(x; \theta)$  where  $\theta$  are the network parameters (Holzmüller et al., 2023). Its associated kernel is known as the *empirical neural tangent kernel* (eNTK) and the posterior mean of this GP approximates ultra-wide NNs trained with gradient descent (Jacot et al., 2018; Arora et al., 2019; Lee et al., 2019; Khan et al., 2019; He et al., 2020; Malladi et al., 2023). Kassraie & Krause (2022) derive bounds of  $\gamma_n$  under this kernel. If  $\theta$  is restricted to the weights of the final linear layer, then this embedding is simply the last-layer embedding.
3. *Loss Gradients*: Another possible choice is

$$\phi(x) = \nabla_{\theta} \ell(f(x; \theta), \hat{y}(x))|_{\theta=\hat{\theta}}$$

where  $\ell$  is a loss function,  $\hat{y}(x)$  is the predicted label, and  $\hat{\theta}$  are the current parameter estimates (Ash et al., 2020).

4. *Outputs (eNNGP)*: Another possible choice is  $\phi(x) = f(x)$ , i.e., the output of the network. Its associated kernel is known as the *empirical neural network Gaussian process* (eNNGP) kernel (Lee et al., 2018).
5. *Predictive (Kirsch, 2023)*: Given a Bayesian neural network (Blundell et al., 2015) or probabilistic (deep) ensemble (Lakshminarayanan et al., 2017), which induce samples  $\theta_1, \dots, \theta_K \sim p(\theta)$  from the distribution over network parameters, one can approximate the predictive covariance  $k(x, x') = \text{Cov}_{\theta}[f(x; \theta), f(x'; \theta)]$ . This kernel measures proximity in the prediction space rather than parameter space and as such does not require gradient information. The corresponding feature map is  $\phi(x) = \frac{1}{\sqrt{K}}[f(x; \theta_1) \cdots f(x; \theta_K)]^\top$  where  $\bar{f}(x; \theta_k) \stackrel{\text{def}}{=} f(x; \theta_k) - \frac{1}{K} \sum_{l=1}^K f(x; \theta_l)$ .

In the additional experiments from this appendix we use last-layer embeddings unless noted otherwise. We compare the performance of last-layer and the loss gradient embedding

$$\phi(x) = \nabla_{\theta'} \ell_{\text{CE}}(f(x; \theta), \hat{y}(x))|_{\theta=\hat{\theta}} \quad (45)$$

where  $\theta'$  are the parameters of the final output layer,  $\hat{\theta}$  are the current parameter estimates,  $\hat{y}(x) = \arg \max_i f_i(x; \hat{\theta})$  are the associated predicted labels, and  $\ell_{\text{CE}}$  denotes the cross-entropy loss. This gradient embedding captures the potential update direction upon observing a new point (Ash et al., 2020). Moreover, Ash et al. (2020) show that for most neural networks, the norm of these gradient embeddings are a conservative lower bound to the norm assumed by taking any other proxy label  $\hat{y}(x)$ . In Figure 8, we observe only negligible differences in performance between this and the last-layer embedding.

<sup>20</sup>We observe essentially the same performance with loss gradient embeddings, cf. Appendix J.2.

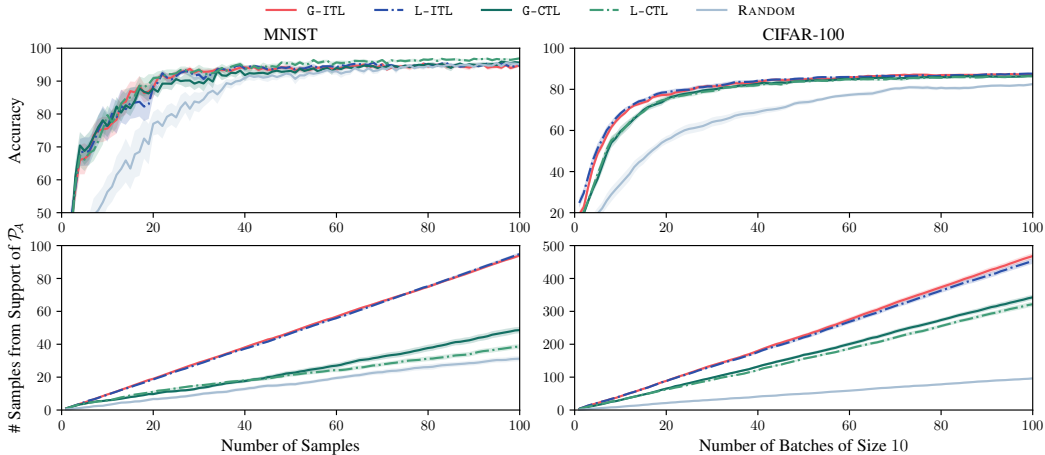


Figure 8: Comparison of loss gradient (“G-”) and last-layer embeddings (“L-”).

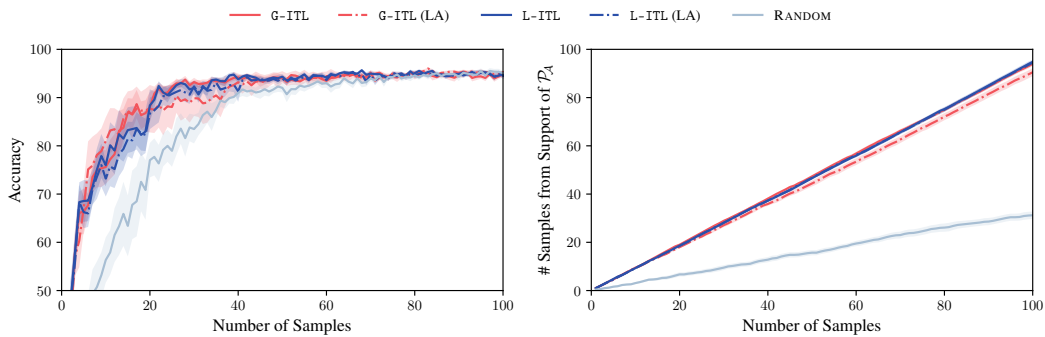


Figure 9: Uncertainty quantification (i.e., estimation of  $\Sigma$ ) via a Laplace approximation (LA, Daxberger et al. (2021)) over last-layer weights using a Kronecker factored log-likelihood Hessian approximation (Martens & Grosse, 2015) and the loss gradient embeddings from Equation (45). The results are shown for the MNIST experiment. We do not observe a performance improvement beyond the trivial approximation  $\Sigma = I$ .

### J.3 Towards Uncertainty Quantification in Latent Space

A straightforward and common approximation of the uncertainty about NN weights is given by  $\Sigma = I$ , and we use this approximation throughout our experiments.

The poor performance of UNSA (cf. Appendix J.5) with this approximation suggests that with more sophisticated approximations, the performance of ITL, VTL, and CTL can be further improved. Further research is needed to study the effect of more sophisticated approximations of “uncertainty” in the latent space. For example, with parameter gradient embeddings, the latent space is the network parameter space where various approximations of  $\Sigma$  based on Laplace approximation (Daxberger et al., 2021; Antorán et al., 2022), variational inference (Blundell et al., 2015), or Markov chain Monte Carlo (Maddox et al., 2019) have been studied. We also evaluate Laplace approximation (LA, Daxberger et al. (2021)) for estimating  $\Sigma$  but see no improvement (cf. Figure 9). Nevertheless, we believe that uncertainty quantification is a promising direction for future work, with the potential to improve performance of ITL and its variations substantially.

### J.4 Batch Selection via Conditional Embeddings

We will refer to the greedy decision rule from Equation (3) as BACE, short for *Batch selection via Conditional Embeddings*. BACE can be implemented efficiently using the Gaussian approximation of  $f_X$  from Appendix J.2 by iteratively conditioning on the previously selected points  $x_{n,1:i-1}$ , and updating the kernel matrix  $K_{XX}$  using the closed-form formula for the variance of conditional

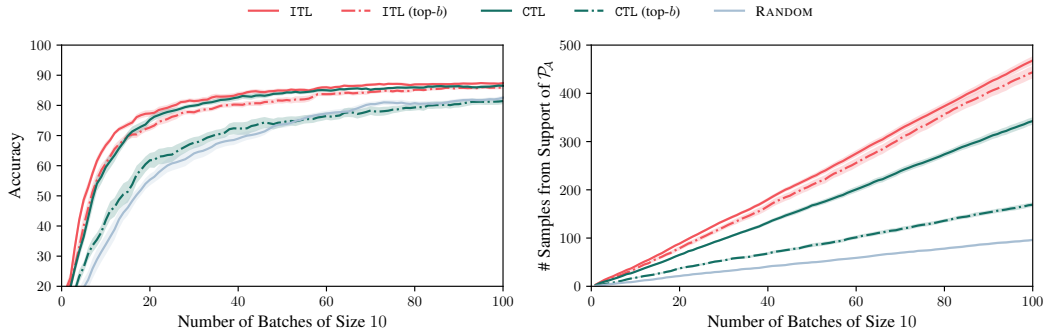


Figure 10: Advantage of batch selection via conditional embeddings over top- $b$  selection in the CIFAR-100 experiment.

Gaussians:

$$\mathbf{K}_{XX} \leftarrow \mathbf{K}_{XX} - \frac{1}{\mathbf{K}_{\mathbf{x}_j \mathbf{x}_j} + \rho^2} \mathbf{K}_{X \mathbf{x}_j} \mathbf{K}_{\mathbf{x}_j X} \quad (46)$$

where  $j$  denotes the index of the selected  $\mathbf{x}_{n,i}$  within  $X$  and  $\rho^2$  is the noise variance. Note that  $\mathbf{K}_{\mathbf{x}_j \mathbf{x}_j}$  is a scalar and  $\mathbf{K}_{X \mathbf{x}_j}$  is a row vector, and hence, this iterative update can be implemented efficiently.

We remark that Equation (3) is a natural extension of previous non-adaptive active learning methods, which typically maximize some notion of “distance” between points in the batch, to the “directed” setting (Ash et al., 2020; Zanette et al., 2021; Holzmüller et al., 2023; Pacchiano et al., 2024). BACE simultaneously maximizes “distance” between points in a batch and minimizes “distance” to points in  $\mathcal{A}$ .

**The sample efficiency of BACE**  $B_n$ , and therefore also the greedily constructed  $B'_n$  (which gives a constant-factor approximation with respect to the objective), yields diverse batches by design. In Figure 10, we compare BACE to selecting the top- $b$  points according to the decision rule (which does *not* yield diverse batches). We observe a significant improvement in accuracy and data retrieval when using BACE. We expect the gap between both approaches to widen further with larger batch sizes.

**Computational complexity of BACE** As derived in Appendix G, a single batch selection step of BACE has complexity  $O(b(k^3 + m^2))$  where  $b$  is the size of the batch,  $k = |\mathcal{A}|$  is the size of the target space, and  $m = |\mathcal{S}|$  is the size of the candidate set. In the case of large  $m$ , an alternative implementation whose runtime does not depend on  $m$  is described in Appendix G.

## J.5 Baselines

In Figure 11, we compare against additional baselines:

- Both TYPICLUST (Hacohen et al., 2022) and PROBCOVER (Yehuda et al., 2022) are recent methods to select points that “cover” the data distribution well. To maintain comparability between algorithms, we use the same embeddings as for ITL which are re-computed before every new batch selection. ITL significantly outperforms TYPICLUST & PROBCOVER, which only attempt to cover  $\mathcal{S}$  well without taking  $\mathcal{A}$  into account (i.e., are “undirected”).
- Mehta et al. (2022) introduced EIG for training neural classification models, which uses the same decision rule as ITL, but approximates the conditional entropy based on the networks’ softmax output rather than using a GP approximation. We approximate the conditional entropy using a single gradient step of the hallucinated updates on the parameters of the final layer, as mentioned by Mehta et al. (2022). We observe that EIG is not competitive for batch-wise selection (CIFAR-100) since it does not encourage batch diversity. Moreover, we observe that EIG is orders of magnitude slower than ITL (since it has to compute  $|\mathcal{S}| \cdot C$  individual gradient steps where  $C$  is the number of classes). We note that since our datasets are balanced, the AEIG algorithm from Mehta et al. (2022) coincides with EIG.



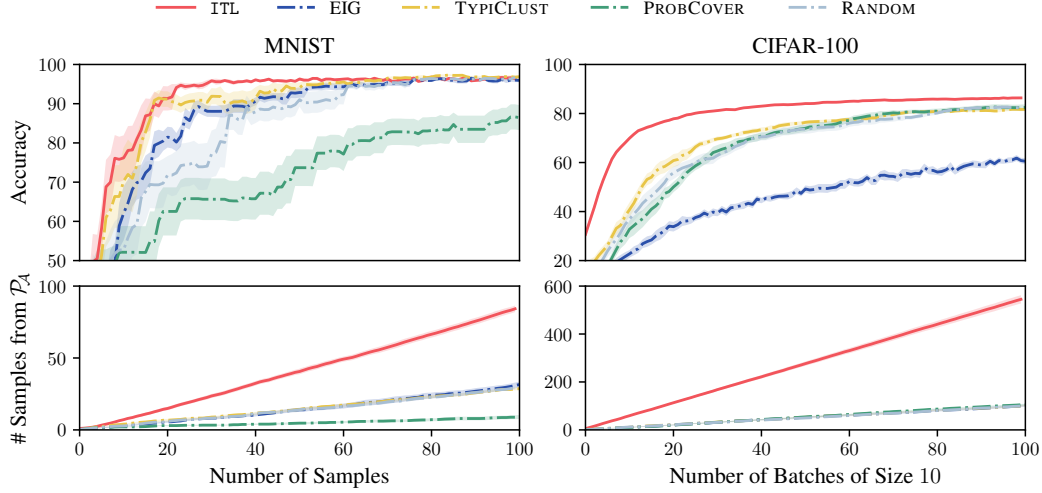


Figure 11: Comparison to baselines for the experiment of Figure 4.

Since, EIG does not have an open-source implementation, we implemented it ourselves following [Mehta et al. \(2022\)](#). For TYPICLUST & PROBCOVER, we use the author’s implementation. In the figure, we show that ITL & VTL substantially outperform all baselines.

In the following, we briefly describe other commonly used “undirected” decision rules.

Denote the softmax distribution over labels  $i$  at inputs  $\mathbf{x}$  by

$$p_i(\mathbf{x}; \hat{\boldsymbol{\theta}}) \propto \exp(f_i(\mathbf{x}; \hat{\boldsymbol{\theta}})).$$

The following heuristics computed based on the softmax distribution aim to quantify the “uncertainty” about a particular input  $\mathbf{x}$ :

- MAXENTROPY ([Settles & Craven, 2008](#)):

$$\mathbf{x}_n = \arg \max_{\mathbf{x} \in \mathcal{S}} H[p(\mathbf{x}; \hat{\boldsymbol{\theta}}_{n-1})].$$

- MAXMARGIN ([Scheffer et al., 2001](#); [Settles & Craven, 2008](#)):

$$\mathbf{x}_n = \arg \min_{\mathbf{x} \in \mathcal{S}} p_1(\mathbf{x}; \hat{\boldsymbol{\theta}}_{n-1}) - p_2(\mathbf{x}; \hat{\boldsymbol{\theta}}_{n-1})$$

where  $p_1$  and  $p_2$  are the two largest class probabilities.

- LEASTCONFIDENCE ([Lewis & Gale, 1994](#); [Settles & Craven, 2008](#); [Hendrycks & Gimpel, 2017](#); [Tamkin et al., 2022](#)):

$$\mathbf{x}_n = \arg \min_{\mathbf{x} \in \mathcal{S}} p_1(\mathbf{x}; \hat{\boldsymbol{\theta}}_{n-1})$$

where  $p_1$  is the largest class probability.

An alternative class of decision rules aims to select diverse batches by maximizing the distances between points. Embeddings  $\phi(\mathbf{x})$  induce the (Euclidean) embedding distance

$$d_\phi(\mathbf{x}, \mathbf{x}') \stackrel{\text{def}}{=} \|\phi(\mathbf{x}) - \phi(\mathbf{x}')\|_2.$$

Similarly, a kernel  $k$  induces the kernel distance

$$d_k(\mathbf{x}, \mathbf{x}') \stackrel{\text{def}}{=} \sqrt{k(\mathbf{x}, \mathbf{x}) + k(\mathbf{x}', \mathbf{x}') - 2k(\mathbf{x}, \mathbf{x}')}.$$

It is straightforward to see that if  $k(\mathbf{x}, \mathbf{x}') = \phi(\mathbf{x})^\top \phi(\mathbf{x}')$ , then embedding and kernel distances coincide, i.e.,  $d_\phi(\mathbf{x}, \mathbf{x}') = d_k(\mathbf{x}, \mathbf{x}')$ .

- MAXDIST (Holzmüller et al., 2023; Yu & Kim, 2010; Sener & Savarese, 2017; Geifman & El-Yaniv, 2017) constructs the batch by choosing the point with the maximum distance to the nearest previously selected point:

$$\mathbf{x}_n = \arg \max_{\mathbf{x} \in \mathcal{S}} \min_{i < n} d(\mathbf{x}, \mathbf{x}_i)$$

- Similarly, K-MEANS++ (Holzmüller et al., 2023) selects the batch via K-MEANS++ seeding (Arthur et al., 2007; Ostrovsky et al., 2013). That is, the first centroid  $\mathbf{x}_1$  is chosen uniformly at random and the subsequent centroids are chosen with a probability proportional to the square of the distance to the nearest previously selected centroid:

$$\mathbb{P}(\mathbf{x}_n = \mathbf{x}) \propto \min_{i < n} d(\mathbf{x}, \mathbf{x}_i)^2.$$

When using the loss gradient embeddings from Equation (45), this decision rule is known as BADGE (Ash et al., 2020).

Finally, we summarize common kernel-based decision rules.

- UNDIRECTED ITL chooses

$$\begin{aligned} \mathbf{x}_n &= \arg \max_{\mathbf{x} \in \mathcal{S}} I(\mathbf{f}_{\mathcal{S}}; y_{\mathbf{x}} \mid \mathcal{D}_{n-1}) \\ &= \arg \max_{\mathbf{x} \in \mathcal{S}} I(f_{\mathbf{x}}; y_{\mathbf{x}} \mid \mathcal{D}_{n-1}). \end{aligned}$$

This can be shown to be equivalent to MAXDET (Holzmüller et al., 2023) which selects

$$\mathbf{x}_n = \arg \max_{\mathbf{x} \in \mathcal{S}} |\mathbf{K}_{\mathbf{x}} + \sigma^2 \mathbf{I}|$$

where  $\mathbf{K}_{\mathbf{x}}$  denotes the kernel matrix over  $\mathbf{x}_{1:n-1} \cup \{\mathbf{x}\}$ , conditioned on the prior observations  $\mathcal{D}_{n-1}$ .

- UNSA (Lewis & Catlett, 1994) which with embeddings  $\phi_{n-1}$  after round  $n - 1$  corresponds to:

$$\mathbf{x}_n = \arg \max_{\mathbf{x} \in \mathcal{S}} \sigma_{n-1}^2(\mathbf{x}) = \arg \max_{\mathbf{x} \in \mathcal{S}} \|\phi_{n-1}(\mathbf{x})\|_2^2.$$

With batch size  $b = 1$ , UNSA coincides with UNDIRECTED ITL. When evaluated with gradient embeddings, this acquisition function is similar to previously used “embedding length” or “gradient length” heuristics (Settles & Craven, 2008).

- UNDIRECTED VTL (Cohn, 1993) is the special case of VTL without specified prediction targets (i.e.,  $\mathcal{A} = \mathcal{S}$ ). In the literature, this decision rule is also known as BAIT (Holzmüller et al., 2023; Ash et al., 2021).

We compare to the abovementioned decision rules and summarize the results in Figure 12. We observe that most “undirected” decision rules perform worse (and often significantly so) than RANDOM. This is likely due to frequently selecting points from the support of  $\mathcal{P}_{\mathcal{S}}$  which are not in the support of  $\mathcal{P}_{\mathcal{A}}$  since the points are “adversarial examples” that the model  $\hat{\theta}$  is not trained to perform well on. In the case of MNIST, the poor performance can also partially be attributed to the well-known “cold-start problem” (Gao et al., 2020).

In Figure 4, we also compare to the following “directed” decision rules:

- COSINESIMILARITY (Settles & Craven, 2008) selects  $\mathbf{x}_n = \arg \max_{\mathbf{x} \in \mathcal{S}} \angle_{\phi_{n-1}}(\mathbf{x}, \mathcal{A})$  where

$$\angle_{\phi}(\mathbf{x}, \mathcal{A}) \stackrel{\text{def}}{=} \frac{1}{|\mathcal{A}|} \sum_{\mathbf{x}' \in \mathcal{A}} \frac{\phi(\mathbf{x})^\top \phi(\mathbf{x}')}{\|\phi(\mathbf{x})\|_2 \|\phi(\mathbf{x}')\|_2}.$$

- INFORMATIONDENSITY (Settles & Craven, 2008) is defined as the multiplicative combination of MAXENTROPY and COSINESIMILARITY:

$$\mathbf{x}_n = \arg \max_{\mathbf{x} \in \mathcal{S}} H[p(\mathbf{x}; \hat{\theta}_{n-1})] \cdot \left( \angle_{\phi_{n-1}}(\mathbf{x}, \mathcal{A}) \right)^\beta$$

where  $\beta > 0$  controls the relative importance of both terms. We set  $\beta = 1$  in our experiments.

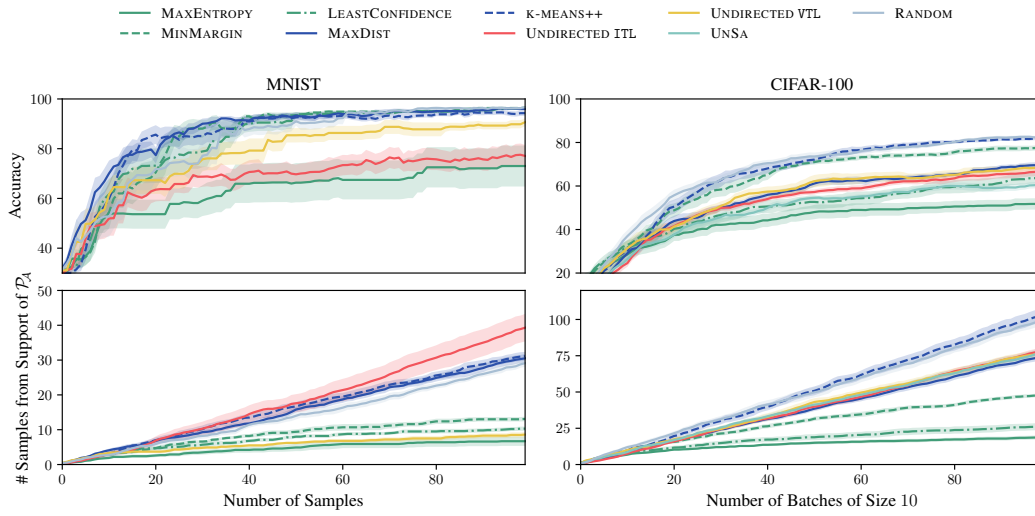


Figure 12: Comparison of “undirected” baselines for the experiment of Figure 4. In the MNIST experiment, UNSA and UNDIRECTED ITL coincide, and we therefore only plot the latter.

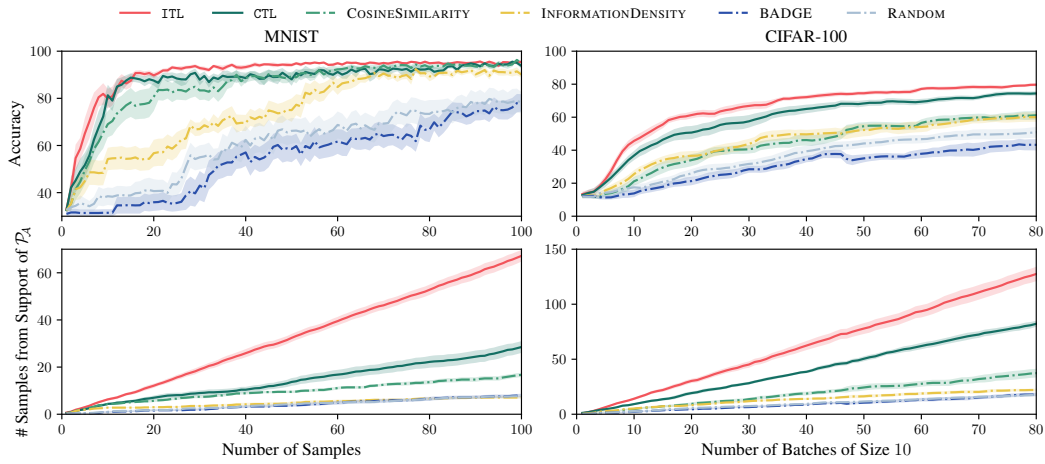


Figure 13: Imbalanced  $\mathcal{P}_S$  experiment.

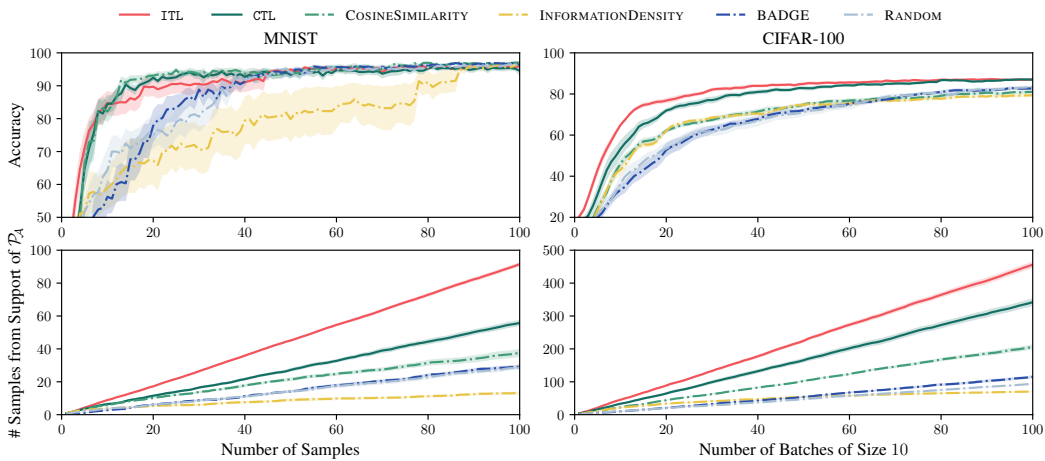


Figure 14: Imbalanced  $A \sim \mathcal{P}_A$  experiment.

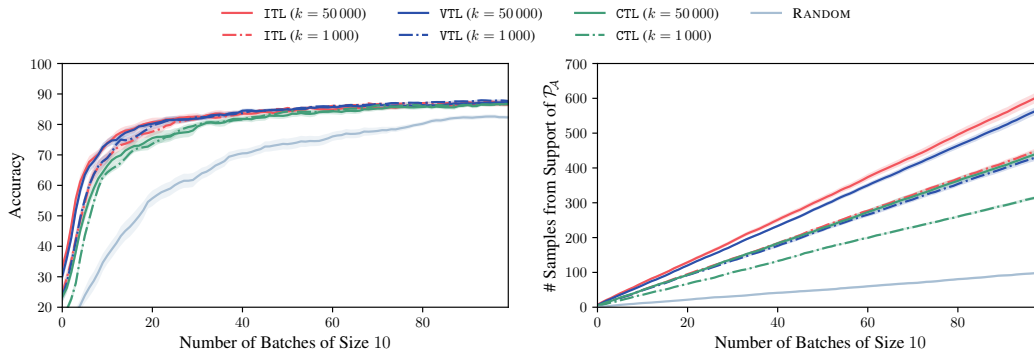


Figure 15: Performance of VTL & choice of  $k$  in the CIFAR-100 experiment.

## J.6 Additional experiments

We conduct the following additional experiments:

1. *Imbalanced  $\mathcal{P}_S$*  (Figure 13): We artificially remove 80% of the support of  $\mathcal{P}_A$  from  $\mathcal{P}_S$ . For example, in case of MNIST, we remove 80% of the images with labels 3, 6, and 9 from  $\mathcal{P}_S$ . This makes the learning task more difficult, as  $\mathcal{P}_A$  is less represented in  $\mathcal{P}_S$ , meaning that the “targets” are more sparse. The trend of ITL outperforming CTL which outperforms RANDOM is even more pronounced in this setting.
2. *Imbalanced  $A \sim \mathcal{P}_A$*  (Figure 14): We artificially remove 50% of part of the support of  $\mathcal{P}_A$  while generating  $A \sim \mathcal{P}_A$  to evaluate the robustness of ITL and CTL in presence of an imbalanced target space  $\mathcal{A}$ . Concretely, in case of MNIST, we remove 50% of the images with labels 3 and 6 from  $A$ . In case of CIFAR-100, we remove 50% of the images with labels  $\{0, \dots, 4\}$  from  $A$ . We still observe the same trends as in the other experiments.
3. *VTL & choice of  $k$*  (Figure 15): We observe that VTL performs almost as well as ITL. Additionally, we evaluate the effect of the number of points  $k$  at which the decision rule is evaluated. Not surprisingly, we observe that the performance of ITL, VTL, and CTL improves with larger  $k$ .
4. *Choice of  $m$*  (Figure 16): Next, we evaluate the choice of  $m$ , i.e., the size of the target space  $\mathcal{A}$  relative to the number  $M$  of candidate points  $A \sim \mathcal{P}_A$ . We write  $p = m/M$ . We generally observe that a larger  $p$  leads to better performance (with  $p = 1$  being the best choice). However, it appears that a smaller  $p$  can be beneficial with respect to accuracy when a large number of batches are selected. We believe that this may be because a smaller  $p$  improves the diversity between selected batches.
5. *Choice of  $M$*  (Figure 17): Finally, we evaluate the choice of  $M$ , i.e., the size of  $A \sim \mathcal{P}_A$ . Not surprisingly, we observe that the performance of ITL improves with larger  $M$ .

## J.7 Ablation study of noise standard deviation $\rho$

In Table 2, we evaluate the CIFAR-100 experiment with different noise standard deviations  $\rho$ . We observe that the performance of batch selection via conditional embeddings drops (mostly for the less numerically stable gradient embeddings) if  $\rho$  is too small, since this leads to numerical inaccuracies when computing the conditional embeddings. Apart from this, the effect of  $\rho$  is negligible.

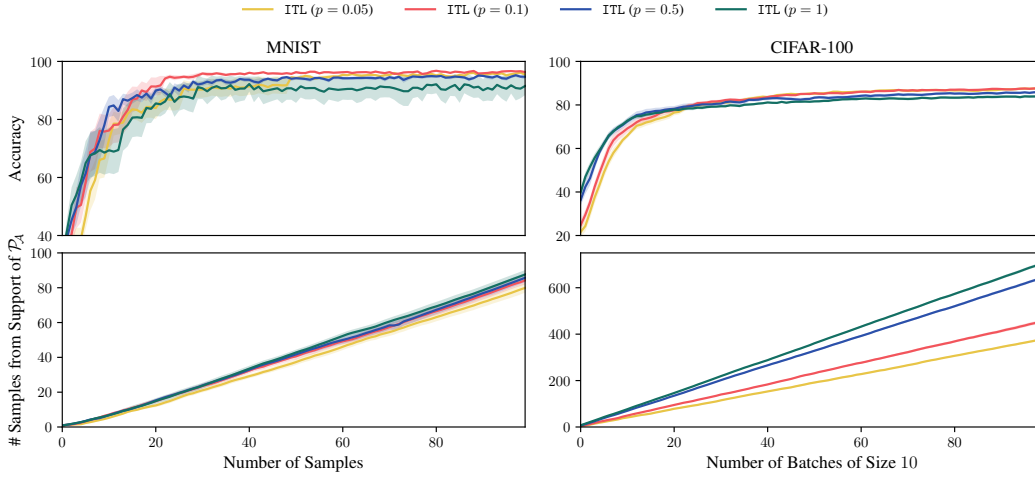


Figure 16: Evaluation of the choice of  $m$  relative to the size  $M$  of  $A \sim \mathcal{P}_A$ . Here,  $p = m/M$ .

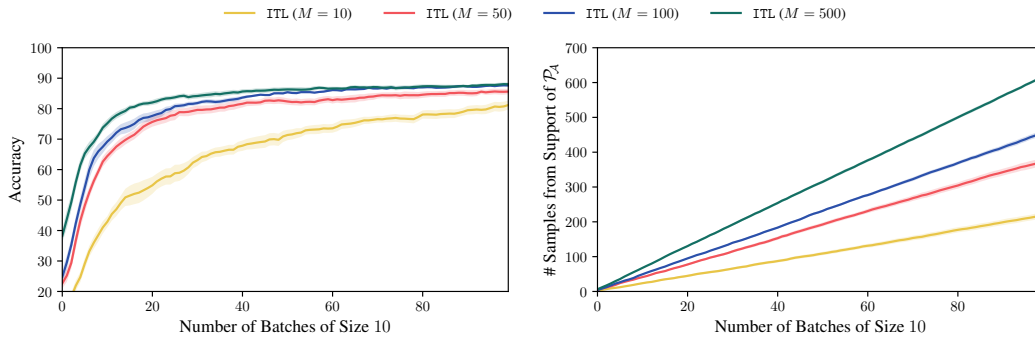


Figure 17: Evaluation of the choice of  $M$ , i.e., the size of  $A \sim \mathcal{P}_A$ , in the CIFAR-100 experiment.

Table 2: Ablation study of noise standard deviation  $\rho$  in the CIFAR-100 experiment. We list the accuracy after 100 rounds per decision rule, with its standard error over 10 random seeds. “(top- $b$ )” denotes variants where batches are selected by taking the top- $b$  points according to the decision rule rather than using batch selection via conditional embeddings. Shown in **bold** are the best performing decision rules, and shown in *italics* are results due to numerical instability.

$\rho$	0.0001	0.01	1	100
G-ITL	<i>78.26 ± 1.40</i>	<i>79.12 ± 1.19</i>	<b>87.16 ± 0.29</b>	<b>87.18 ± 0.28</b>
L-ITL	<b>87.52 ± 0.48</b>	<b>87.52 ± 0.41</b>	<b>87.53 ± 0.35</b>	86.47 ± 0.22
G-CTL	<i>58.68 ± 2.11</i>	<i>81.44 ± 1.04</i>	86.52 ± 0.44	<b>86.92 ± 0.56</b>
L-CTL	<b>86.40 ± 0.71</b>	<b>86.38 ± 0.75</b>	86.00 ± 0.69	84.78 ± 0.39
G-ITL (top- $b$ )	85.84 ± 0.54	85.92 ± 0.52	85.84 ± 0.54	85.55 ± 0.46
L-ITL (top- $b$ )	85.44 ± 0.58	85.46 ± 0.54	85.44 ± 0.59	85.29 ± 0.36
G-CTL (top- $b$ )	82.27 ± 0.67	82.27 ± 0.67	82.27 ± 0.67	82.27 ± 0.67
L-CTL (top- $b$ )	80.73 ± 0.68	80.73 ± 0.68	80.73 ± 0.68	80.73 ± 0.68
BADGE	83.24 ± 0.60	83.24 ± 0.60	83.24 ± 0.60	83.24 ± 0.60
INFORMATIONDENSITY	79.24 ± 0.51	79.24 ± 0.51	79.24 ± 0.51	79.24 ± 0.51
RANDOM	82.49 ± 0.66	82.49 ± 0.66	82.49 ± 0.66	82.49 ± 0.66

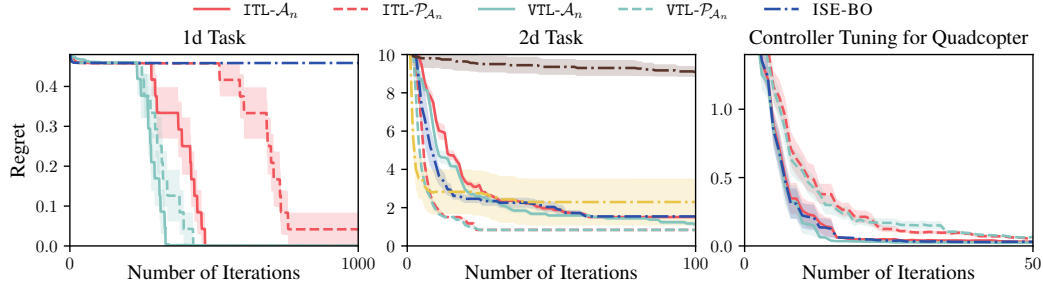


Figure 18: We perform the tasks of Figure 5 using Thompson sampling to evaluate the stochastic target space  $\mathcal{P}_{\mathcal{A}_n}$ . We additionally compare to GOOSE (cf. Appendix K.2.3) and ISE-BO (cf. Appendix K.2.4).

## K Additional Safe BO Experiments & Details

In Appendix K.1, we discuss the use of stochastic target spaces in the safe BO setting. We provide a comprehensive overview of prior works in Appendix K.2 and an additional experiment highlighting that ITL, unlike SAFEOPT, is able to “jump past local barriers” in Appendix K.3. In Appendix K.4, we provide details on the experiments from Figure 5.

### K.1 A More Exploitative Stochastic Target Space

Alternatively to the target space  $\mathcal{A}_n$  which comprises all potentially optimal points, we evaluate the stochastic target space

$$\mathcal{P}_{\mathcal{A}_n}(\cdot) = \mathbb{P}(\arg \max_{\mathbf{x} \in \mathcal{X}: g(\mathbf{x}) \geq 0} f(\mathbf{x}) = \cdot \mid \mathcal{D}_n) \quad (47)$$

which effectively weights points in  $\mathcal{A}_n$  according to how likely they are to be the safe optimum, and is therefore more exploitative than the uniformly-weighted target space discussed so far. Samples from  $\mathcal{P}_{\mathcal{A}_n}$  can be obtained efficiently via Thompson sampling (Thompson, 1933; Russo et al., 2018). Observe that  $\mathcal{P}_{\mathcal{A}_n}$  is supported precisely on the set of potential maximizers  $\mathcal{A}_n$ . We provide a formal analysis of stochastic target spaces in Appendix E. Whether transductive active learning with  $\mathcal{A}_n$  or  $\mathcal{P}_{\mathcal{A}_n}$  performs better is task-dependent, as we will see in the following.

Note that performing ITL with this target space is analogous to output-space entropy search (Wang & Jegelka, 2017). Samples from  $\mathcal{P}_{\mathcal{A}_n}$  can be obtained via Thompson sampling (Thompson, 1933; Russo et al., 2018). That is, in iteration  $n + 1$ , we sample  $K \in \mathbb{N}$  independent functions  $f^{(j)} \sim f \mid \mathcal{D}_n$  from the posterior distribution and select  $K$  points  $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(K)}$  which are a safe maximum of  $f^{(1)}, \dots, f^{(K)}$ , respectively.

**Experiments** In Figure 18, we contrast the performance of ITL with  $\mathcal{P}_{\mathcal{A}_n}$  to the performance of ITL with the exact target space  $\mathcal{A}_n$ . We observe that their relative performance is instance dependent: in tasks that require more difficult expansion, ITL with  $\mathcal{A}_n$  converges faster, whereas in simpler tasks (such as the 2d experiment), ITL with  $\mathcal{P}_{\mathcal{A}_n}$  converges faster. We compare against the GOOSE algorithm (Turchetta et al., 2019) which is a heuristic extension of SAFEOPT that explores more greedily in directions of (assumed) high reward (cf. Appendix K.2.3). GOOSE suffers from the same limitations as SAFEOPT, which were highlighted in Section 5, and additionally is limited by its heuristic approach to expansion which fails in the 1d task and safe controller tuning task. Analogously to our experiments with SAFEOPT, we also compare against ORACLE GOOSE which has oracle knowledge of the true Lipschitz constants.

The different behaviors of ITL with  $\mathcal{A}_n$  and  $\mathcal{P}_{\mathcal{A}_n}$ , respectively, as well as SAFEOPT and GOOSE are illustrated in Figure 19. We observe that ITL with  $\mathcal{A}_n$  and SAFEOPT expand the safe set more “uniformly” since the set of potential maximizers encircles the true safe set.<sup>21</sup> Intuitively, this is because the set of potential maximizers *conservatively* captures migh points might be safe and

<sup>21</sup>This is because typically, there will always remain points in  $\widehat{\mathcal{S}}_n \setminus \mathcal{S}_n$  of which the safety cannot be fully determined, and since, they cannot be observed, it can also not be ruled out that they have high objective value.

optimal. In contrast, ITL with  $\mathcal{P}_{\mathcal{A}_n}$  and GOOSE focus exploration and expansion in those regions where the objective is likely to be high.




Figure 19: The first 100 samples of (A) ITL with  $\mathcal{A}_n$ , (B) SAFEOPT, (C) ORACLE SAFEOPT, (D) ITL with  $\mathcal{P}_{\mathcal{A}_n}$ , (E) GOOSE, (F) ORACLE GOOSE. The white region denotes the pessimistic safe set  $\mathcal{S}_{100}$ , the light gray region denotes the true safe set  $\mathcal{S}^*$  (i.e., the “island”), and the darker gray regions denotes unsafe points (i.e., the “ocean”).

## K.2 Detailed Comparison with Prior Works

The most widely used method for Safe BO is SAFEOPT (Sui et al., 2015; Berkenkamp et al., 2021) which keeps track of separate candidate sets for expansion and exploration and uses UNSA to pick one of the candidates in each round. Treating expansion and exploration separately, sampling is directed towards expansion in *all* directions — even those that are known to be suboptimal. The safe set is expanded based on a Lipschitz constant of  $g^*$ , which is assumed to be known. In most real-world settings, this constant is unknown and has to be estimated using the GP. This estimate is generally conservative and results in suboptimal performance. To this end, Berkenkamp et al. (2016) proposed HEURISTIC SAFEOPT which relies solely on the confidence intervals of  $g$  to expand the safe set, but lacks convergence guarantees. More recently, Bottero et al. (2022) proposed ISE which queries parameters from  $\mathcal{S}_n$  that yield the most “information” about the safety of another parameter in  $\mathcal{X}$ . Hence, ISE focuses solely on the expansion of the safe set  $\mathcal{S}_n$  and does not take into account the objective  $f$ . In practice, this can lead to significantly worse performance on the simplest of problems (cf. Figure 5). In contrast, ITL balances expansion of and exploration within the safe set. Furthermore, ISE does not have known convergence guarantees of the kind of Theorem 5.1. In parallel independent work, Bottero et al. (2024) proposed a combination of ISE and max-value entropy search (Wang & Jegelka, 2017) for which they derive a similar guarantee to Theorem 5.1.<sup>22</sup> Similar to SAFEOPT, their method aims to expand the safe set in all directions including those that are known to be suboptimal. In contrast, ITL directs expansion only towards potentially optimal regions.

In the 1d task and quadcopter experiment (cf. Figure 5), we observe that SAFEOPT and even ORACLE SAFEOPT converge significantly slower than ITL to the safe optima. We believe this is due to their conservative Lipschitz-continuity/global smoothness-based expansion, as opposed to ITL’s expansion,

<sup>22</sup>We provide an empirical evaluation in Appendix K.2.4.

which adapts to the local smoothness of the constraints. HEURISTIC SAFEOPT, which does not rely on the Lipschitz constant for expansion, does not efficiently expand the safe set due to its heuristic that only considers single-step expansion. This is especially the case for the 1d task. Furthermore, in the 2d task, we notice the suboptimality of ISE since it does not take into account the objective, and purely aims to expand the safe set. ITL, on the other hand, balances expansion and exploration.

### K.2.1 SAFEOPT

SAFEOPT (Sui et al., 2015; Berkenkamp et al., 2021) is a well-known algorithm for Safe BO.

**Lipschitz-based expansion** SAFEOPT expands the set of known-to-be safe points by assuming knowledge of an upper bound  $L_i$  to the Lipschitz constant of the unknown constraints  $g_i^*$ .<sup>23</sup> In each iteration, the (pessimistic) safe set  $\mathcal{S}_n$  is updated to include all points which can be reached safely (with respect to the Lipschitz continuity) from a known-to-be-safe point  $\mathbf{x} \in \mathcal{S}_n$ . Formally,

$$\mathcal{S}_n^{\text{SAFEOPT}} \stackrel{\text{def}}{=} \bigcup_{\mathbf{x} \in \mathcal{S}_{n-1}^{\text{SAFEOPT}}} \{\mathbf{x}' \in \mathcal{X} \mid l_{n,i}(\mathbf{x}) - L_i \|\mathbf{x} - \mathbf{x}'\|_2 \geq 0 \text{ for all } i \in \mathcal{I}_s\}. \quad (48)$$

The expansion of the safe set is illustrated in Figure 20.

We remark two main limitations of this approach. First, the Lipschitz constant is an additional safety critical hyperparameter of the algorithm, which is typically not known. The RKHS assumption (cf. Assumption C.25) induces an assumption on the Lipschitz continuity, however, the worst-case a-priori Lipschitz constant is typically very large, and prohibitive for expansion. Second, the Lipschitz constant is global property of the unknown function, meaning that it does not adapt to the local smoothness. For example, a constraint may be “flat” in one direction (permitting straightforward expansion) and “steep” in another direction (requiring slow expansion). Furthermore, the Lipschitz constant is constant over time, whereas ITL is able to adapt to the local smoothness and reduce the (induced) Lipschitz constant over time.

**Undirected expansion** SAFEOPT addresses the trade-off between expansion and exploration by focusing learning on two different sets. First, the set of *maximizers*

$$\mathcal{M}_n^{\text{SAFEOPT}} \stackrel{\text{def}}{=} \{\mathbf{x} \in \mathcal{S}_n^{\text{SAFEOPT}} \mid u_{n,f}(\mathbf{x}) \geq \max_{\mathbf{x}' \in \mathcal{S}_n^{\text{SAFEOPT}}} l_{n,f}(\mathbf{x}')\}$$

which contains all *known-to-be-safe* points which are potentially optimal. Note that if  $\mathcal{S}_n^{\text{SAFEOPT}} = \mathcal{S}_n$  then  $\mathcal{M}_n^{\text{SAFEOPT}} \subseteq \mathcal{A}_n$  since  $\mathcal{A}_n$  contains points which are potentially optimal and potentially safe *but possibly unsafe*.

To facilitate expansion, for each point  $\mathbf{x} \in \mathcal{S}_n$ , the algorithm considers a set of *expanding points*

$$\mathcal{F}_n^{\text{SAFEOPT}}(\mathbf{x}) \stackrel{\text{def}}{=} \{\mathbf{x}' \in \mathcal{X} \setminus \mathcal{S}_n^{\text{SAFEOPT}} \mid u_{n,i}(\mathbf{x}) - L_i \|\mathbf{x} - \mathbf{x}'\|_2 \geq 0 \text{ for all } i \in \mathcal{I}_s\}$$

A point is expanding if it is unsafe initially and can be (optimistically) deduced as safe by observing  $\mathbf{x}$ . The set of *expanders* corresponds to all known-to-be-safe points which optimistically lead to expansion of the safe set:

$$\mathcal{G}_n^{\text{SAFEOPT}} \stackrel{\text{def}}{=} \{\mathbf{x} \in \mathcal{S}_n \mid |\mathcal{F}_n(\mathbf{x})| > 0\}.$$

That is, an expander is a safe point  $\mathbf{x}$  which is “close” to at least one expanding point  $\mathbf{x}'$ . Observe that here, we start with a safe  $\mathbf{x}$  and then find a close and potentially safe  $\mathbf{x}'$  using the Lipschitz-property of the constraint function. Thus, the set of expanding points is inherently limited by the assumed Lipschitzness (cf. Figure 20), and generally a subset of the potential expanders  $\mathcal{E}_n$  (cf. Equation (27)):

**Lemma K.1.** *For any  $n \geq 0$ , if  $\mathcal{S}_n^{\text{SAFEOPT}} = \mathcal{S}_n$  then*

$$\bigcup_{\mathbf{x} \in \mathcal{S}_n} \mathcal{F}_n^{\text{SAFEOPT}}(\mathbf{x}) \subseteq \mathcal{E}_n.$$

<sup>23</sup>Recall that due to the assumption that  $\|g_i^*\|_k < \infty$ ,  $g_i^*$  is indeed Lipschitz continuous.



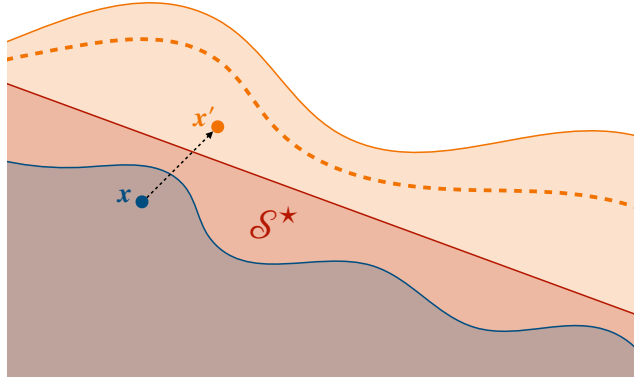


Figure 20: Illustration of the expansion of the safe set à la SAFEOPT. Here, the blue region denotes the pessimistic safe set  $\mathcal{S}$ , the red region denotes the true safe set  $\mathcal{S}^*$ , and the orange region denotes the optimistic safe set  $\widehat{\mathcal{S}}$ . Whereas ITL learns about the point  $\mathbf{x}'$  *directly*, SAFEOPT expands the safe set using the reduction of uncertainty at  $\mathbf{x}$ , and then extrapolating using the Lipschitz constant (cf. Equation (48)). The dashed orange line denotes the expanding points of SAFEOPT which under-approximate the optimistic safe set of ITL (cf. Lemma K.1). Thus, ITL may even learn about points in  $\widehat{\mathcal{S}}$  which are “out of reach” for SAFEOPT.

*Proof.* Without loss of generality, we consider the case where  $\mathcal{I}_s = \{i\}$ . We have

$$\mathcal{E}_n = \widehat{\mathcal{S}}_n \setminus \mathcal{S}_n = \{\mathbf{x} \in \mathcal{X} \setminus \mathcal{S}_n \mid u_{n,i}(\mathbf{x}) \geq 0\}.$$

The result follows directly by observing that  $L_i \|\mathbf{x} - \mathbf{x}'\|_2 \geq 0$ .  $\square$

SAFEOPT then selects  $\mathbf{x}_{n+1}$  according to uncertainty sampling *within* the maximizers and expanders:  $\mathcal{M}_n^{\text{SAFEOPT}} \cup \mathcal{G}_n^{\text{SAFEOPT}}$ . We remark that due to the separate handling of expansion and exploration, SAFEOPT expands the safe set in *all* directions — even those that are known to be suboptimal. In contrast, ITL only expands the safe set in directions that are potentially optimal by balancing expansion and exploration through the single set of potential maximizers  $\mathcal{A}_n$ .

**Based on uncertainty sampling** As mentioned in the previous paragraph, SAFEOPT selects as next point the maximizer/expander with the largest prior uncertainty.<sup>24</sup> In contrast, ITL selects the point within  $\mathcal{S}_n$  which minimizes the posterior uncertainty within  $\mathcal{A}_n$ . Note that the two approaches are not identical as typically  $\mathcal{M}_n^{\text{SAFEOPT}} \cup \mathcal{G}_n^{\text{SAFEOPT}} \subset \mathcal{S}_n^{\text{SAFEOPT}}$  and  $\mathcal{A}_n \not\subseteq \mathcal{S}_n$ .

We show empirically in Section 3.3 that depending on the kernel choice (i.e., the smoothness assumptions), uncertainty sampling within a given target space neglects higher-order information that can be attained by sampling outside the set. This can be seen even more clearly when considering linear functions, in which case points outside the maximizers and expanders can be equally informative as points inside.

Finally, note that the set of expanders is constructed “greedily”, i.e., only considering *single-step* expansion. This is necessitated as the inference of safety is based on single reference points. Instead, ITL directly quantifies the information gained towards the points of interest without considering intermediate reference points.

**Requires homoscedastic noise** SAFEOPT imposes a homoscedasticity assumption on the noise which is an artifact of the analysis of uncertainty sampling. It is well known that in the presence of heteroscedastic noise, one has to distinguish epistemic and aleatoric uncertainty. Uncertainty sampling fails because it may continuously sample a high variance point where the variance is dominated by aleatoric uncertainty, potentially missing out on reducing epistemic uncertainty at points with small aleatoric uncertainty. In contrast, maximizing mutual information naturally takes

<sup>24</sup>The use of uncertainty sampling for safe sequential decision-making goes back to Schreiter et al. (2015) and Sui et al. (2015).

into account the two sources of uncertainty, preferring those points where epistemic uncertainty is large and aleatoric uncertainty is small (cf. Appendix C.1).

**Suboptimal reachable safe set** Sui et al. (2015) and Berkenkamp et al. (2021) show that SAFEOPT converges to the optimum within the closure  $\bar{\mathcal{R}}_\epsilon^{\text{SAFEOPT}}(\mathcal{S}_0)$  of

$$\mathcal{R}_\epsilon^{\text{SAFEOPT}}(\mathcal{S}) \stackrel{\text{def}}{=} \mathcal{S} \cup \{\mathbf{x} \in \mathcal{X} \mid \exists \mathbf{x}' \in \mathcal{S} \text{ such that} \\ f_i^*(\mathbf{x}') - (L_i \|\mathbf{x} - \mathbf{x}'\|_2 + \epsilon) \geq 0 \text{ for all } i \in \mathcal{I}_s\}.$$

Note that analogously to the expansion of the safe set, the “expansion” of the reachable safe set is based on “inferring safety” through a reference point in  $\mathcal{S}$  and using Lipschitz continuity. This is opposed to the reachable safe set of ITL (cf. Definition C.29).

We remark that under the additional assumption that a Lipschitz constant is known, ITL can easily be extended to expand its safe set based on the kernel *and* the Lipschitz constant, resulting in a strictly larger reachable safe set than SAFEOPT. We leave the concrete formalization of this extension to future work. Moreover, we do not evaluate this extension in our experiments, as we observe that even without the additional assumption of a Lipschitz constant, ITL outperforms SAFEOPT in practice.

## K.2.2 HEURISTIC SAFEOPT

Berkenkamp et al. (2016) also implement a heuristic variant of SAFEOPT which does not assume a known Lipschitz constant. This heuristic variant uses the same (pessimistic) safe sets  $\mathcal{S}_n$  as ITL. The set of maximizers is identical to SAFEOPT. As expanders, the heuristic variant considers all safe points  $\mathbf{x} \in \mathcal{S}_n$  that if  $\mathbf{x}$  were to be observed next with value  $\mathbf{u}_n(\mathbf{x})$  lead to  $|\mathcal{S}_{n+1}| > |\mathcal{S}_n|$ . We refer to this set as  $\mathcal{G}_n^{\text{H-SAFEOPT}}$ . The next point is then selected by uncertainty sampling within  $\mathcal{M}_n^{\text{SAFEOPT}} \cup \mathcal{G}_n^{\text{H-SAFEOPT}}$ .

The heuristic variant shares some properties with SAFEOPT, such that it is based on uncertainty sampling, not adapting to heteroscedastic noise, and separate notions of maximizers and expanders (leading to an “undirected” expansion of the safe set). Note that there are no known convergence guarantees for heuristic SAFEOPT. Importantly, note that similar to SAFEOPT the set of expanders is constructed “greedily”, and in particular, does only take into account *single-step* expansion. In contrast, an objective such as ITL which quantifies the “information gained towards expansion” also actively seeks out *multi-step* expansion.

## K.2.3 GOOSE

To address the “undirected” expansion of SAFEOPT discussed in the previous section, Turchetta et al. (2019) proposed *goal-oriented safe exploration* (GOOSE). GOOSE extends any unsafe BO algorithm (which we subsequently call an oracle) to the safe setting. In our experiments, we evaluate GOOSE-UCB which uses UCB as oracle and which is also the variant studied by Turchetta et al. (2019). In the following, we assume for ease of notation that  $\mathcal{I}_s = \{c\}$ .

Given the oracle proposal  $\mathbf{x}^*$ , GOOSE first determines whether  $\mathbf{x}^*$  is safe. If  $\mathbf{x}^*$  is safe,  $\mathbf{x}^*$  is queried next. Otherwise, GOOSE first learns about the safety of  $\mathbf{x}^*$  by querying “expansionist” points until the oracle’s proposal is determined to be either safe or unsafe.

GOOSE expands the safe set identically to SAFEOPT according to Equation (48). In the context of GOOSE,  $\mathcal{S}_n^{\text{SAFEOPT}}$  is called the *pessimistic safe set*. To determine that a point cannot be deduced as safe, GOOSE also keeps track of a Lipschitz-based *optimistic safe set*:

$$\hat{\mathcal{S}}_{n,\epsilon}^{\text{GoOSE}} \stackrel{\text{def}}{=} \bigcup_{\mathbf{x} \in \mathcal{S}_{n-1}^{\text{SAFEOPT}}} \{\mathbf{x}' \in \mathcal{X} \mid \\ u_{n,c}(\mathbf{x}) - L_c \|\mathbf{x} - \mathbf{x}'\|_2 - \epsilon \geq 0\}.$$

We summarize the algorithm in Algorithm 2 where we denote by  $\mathcal{O}(\mathcal{X})$  the oracle proposal over the domain  $\mathcal{X}$ .

It remains to discuss the heuristic used to select the “expansionist” points. GOOSE considers all points  $\mathbf{x} \in \mathcal{S}_n^{\text{SAFEOPT}}$  with confidence bands of size larger than the accuracy  $\epsilon$ , i.e.,

$$\mathcal{W}_{n,\epsilon}^{\text{GoOSE}} \stackrel{\text{def}}{=} \{\mathbf{x} \in \mathcal{S}_n^{\text{SAFEOPT}} \mid u_{n,c}(\mathbf{x}) - l_{n,c}(\mathbf{x}) > \epsilon\}.$$

---

**Algorithm 2** GOOSE
 

---

**Given:** Lipschitz constant  $L_c$ , prior model  $\{f, g_c\}$ , oracle  $\mathcal{O}$ , and precision  $\epsilon$   
 Set initial safe set  $\mathcal{S}_0^{\text{SAFEOPT}}$  based on prior  
 $\widehat{\mathcal{S}}_{n,\epsilon}^{\text{GOOSE}} \leftarrow \mathcal{X}$   
 $n \leftarrow 0$   
**for**  $k$  from 1 to  $\infty$  **do**  
    $\mathbf{x}_k^* \leftarrow \mathcal{O}(\widehat{\mathcal{S}}_{n,\epsilon}^{\text{GOOSE}})$   
   **while**  $\mathbf{x}_k^* \notin \mathcal{S}_n^{\text{SAFEOPT}}$  **do**  
     Observe “expansionist” point  $\mathbf{x}_{n+1}$ , set  $n \leftarrow n + 1$ , and update model and safe sets  
   **end while**  
   Observe  $\mathbf{x}_k^*$ , set  $n \leftarrow n + 1$ , and update model and safe sets  
**end for**

---

Which of the points in this set is evaluated depends on a set of learning targets  $\mathcal{A}_{n,\epsilon}^{\text{GOOSE}} \stackrel{\text{def}}{=} \widehat{\mathcal{S}}_{n,\epsilon}^{\text{GOOSE}} \setminus \mathcal{S}_n^{\text{SAFEOPT}}$  akin to the “potential expanders”  $\mathcal{E}_n$  (cf. Equation (27)), to each of which we assign a priority  $h(\mathbf{x})$ . When  $h(\mathbf{x})$  is large, this indicates that the algorithm is prioritizing to determine whether  $\mathbf{x}$  is safe. We use as heuristic the negative  $\ell_1$ -distance between  $\mathbf{x}$  and  $\mathbf{x}^*$ . GOOSE then considers the set of *potential immediate expanders*

$$\mathcal{G}_{n,\epsilon}^{\text{GOOSE}}(\alpha) \stackrel{\text{def}}{=} \{\mathbf{x} \in \mathcal{W}_{n,\epsilon}^{\text{GOOSE}} \mid \exists \mathbf{x}' \in \mathcal{A}_{n,\epsilon}^{\text{GOOSE}} \text{ with priority } \alpha \text{ such that } u_{n,c}(\mathbf{x}) - L_c \|\mathbf{x} - \mathbf{x}'\|_2 \geq 0\}.$$

The “expansionist” point selected by GOOSE is then any point in  $\mathcal{G}_{n,\epsilon}^{\text{GOOSE}}(\alpha^*)$  where  $\alpha^*$  denotes the largest priority such that  $|\mathcal{G}_{n,\epsilon}^{\text{GOOSE}}(\alpha^*)| > 0$ .

We observe empirically that the sample complexity of GOOSE is not always better than that of SAFEOPT. Notably, the expansion of the safe set is based on a “greedy” heuristic. Moreover, determining whether a single oracle proposal  $\mathbf{x}^*$  is safe may take significant time. Consider the (realistic) example where the prior is uniform, and UCB proposes a point which is far away from the safe set and suboptimal. GOOSE will typically attempt to derive the safety of the proposed point until the uncertainty at *all* points within  $\mathcal{S}_0^{\text{SAFEOPT}}$  is reduced to  $\epsilon$ .<sup>25</sup> Thus, GOOSE can “waste” a significant number of samples, aiming to expand the safe set towards a known-to-be suboptimal point. In larger state spaces, due to the greedy nature of the expansion strategy, this can lead to GOOSE being effectively stuck at a suboptimal point for a significant number of rounds.

### K.2.4 ISE and ISE-BO

Recently, [Bottero et al. \(2022\)](#) proposed an information-theoretic approach to efficiently expand the safe set which they call *information-theoretic safe exploration* (ISE). Specifically, they choose the next action  $\mathbf{x}_n$  by approximating

$$\arg \max_{\mathbf{x} \in \mathcal{S}_{n-1}} \underbrace{\max_{\mathbf{x}' \in \mathcal{X}} \mathbb{I}(\mathbb{1}\{g_{\mathbf{x}'} \geq 0\}; y_{\mathbf{x}} \mid \mathcal{D}_{n-1})}_{\alpha^{\text{ISE}}(\mathbf{x})}. \quad (\text{ISE})$$

In a parallel independent work, [Bottero et al. \(2024\)](#) extended ISE to the Safe BO problem where they propose to choose  $\mathbf{x}_n$  according to

$$\arg \max_{\mathbf{x} \in \mathcal{S}_{n-1}} \max\{\alpha^{\text{ISE}}(\mathbf{x}), \alpha^{\text{MES}}(\mathbf{x})\} \quad (\text{ISE-BO})$$

where  $\alpha^{\text{MES}}$  denotes the acquisition function of max-value entropy search ([Wang & Jegelka, 2017](#)). Similarly to SAFEOPT, ISE-BO treats expansion and exploration separately, which leads to “undirected” expansion of the safe set. That is, the safe set is expanded in all directions, even those that are known to be suboptimal. In contrast, ITL balances expansion and exploration through the single set of potential maximizers  $\mathcal{A}_n$ . With a stochastic target space, ITL generalizes max-value entropy search (cf. Appendix K.1).

We evaluate ISE-BO in Figure 18 and observe that it does not outperform ITL and VTL in any of the tasks, while performing poorly in the 1d task and suboptimally in the 2d task.

---

<sup>25</sup>This is because the proposed point typically remains in the optimistic safe set when it is sufficiently far away from the pessimistic safe set.

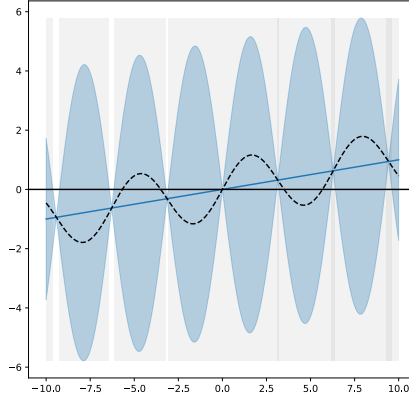


Figure 21: The ground truth  $f^*$  is shown as the dashed black line. The solid black line denotes the constraint boundary. The GP prior is given by a linear kernel with sin-transform and mean  $0.1x$ . The light gray region denotes the initial optimistic safe set  $\hat{\mathcal{S}}_0$  and the dark gray region denotes the initial pessimistic safe set  $\mathcal{S}_0$ .

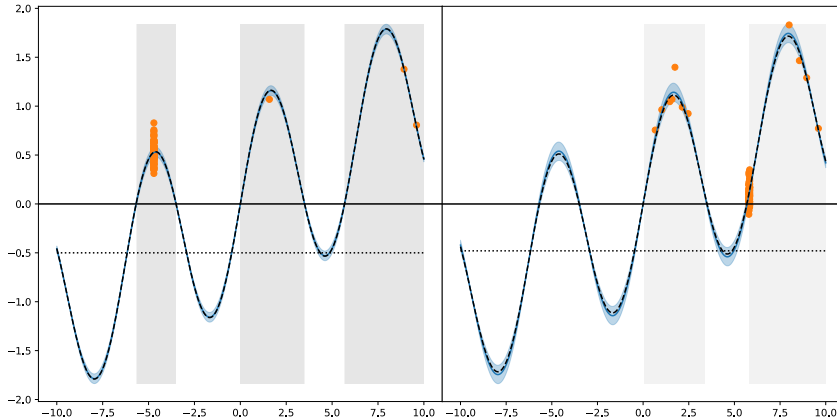


Figure 22: First 100 samples of ITL using the potential expanders  $\mathcal{E}_n$  (cf. Equation (27)) as target space (left) and SAFEOPT sampling only from the set of expanders  $\mathcal{G}_n^{\text{SAFEOPT}}$  (right).

### K.3 Jumping Past Local Barriers

In this additional experiment we demonstrate that ITL is able to extrapolate safety beyond local unsafe “barriers”, which is a fundamental limitation of Lipschitz-based methods such as SAFEOPT. We consider the ground truth function and prior statistical model shown in Figure 21. Note that initially, there are three disjoint safe “regions” known to the algorithm corresponding to two of the three safe “bumps” of the ground truth function. In this experiment, the main challenge is to “jump past” the local barrier separating the leftmost and initially unknown safe “bump”.

Figure 22 shows the sampled points during the first 100 iterations of SAFEOPT and ITL. Clearly, SAFEOPT does not discover the third safe “bump” while ITL does. Indeed, it is a fundamental limitation of Lipschitz-based methods that they can never “jump past local barriers”, even if the oracle Lipschitz constant were to be known and tight (i.e., locally accurate) around the barrier. This is because Lipschitz-based methods expand to the point  $x$  based on a reference point  $x'$ , and by definition, if  $x$  is added to the safe set so are all points on the line segment between  $x$  and  $x'$ . Hence, if there is a single point on this line segment which is unsafe (i.e., a “barrier”), the algorithm will *never* expand past it. This limitation does not exist for kernel-based algorithms as expansion occurs in function space.

Moreover, note that for a non-stationary kernel such as in this example, ITL samples the “closest points” in function space rather than Euclidean space. We observe that SAFEOPT still samples “locally at the boundary” whereas ITL samples the most informative point which in this case is

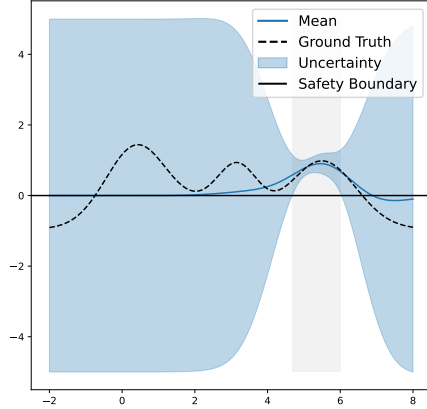


Figure 23: Ground truth and prior well-calibrated model in 1d synthetic experiment. The function serves simultaneously as objective and as constraint. The light gray region denotes the initial safe set  $\mathcal{S}_0$ .

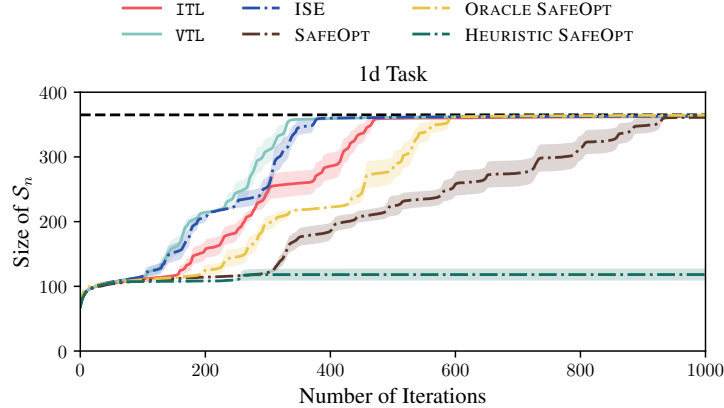


Figure 24: Size of  $\mathcal{S}_n$  in 1d synthetic experiment. The dashed black line denotes the size of  $\mathcal{S}^*$ . In this task, “discovering” the optimum is closely linked to expansion of the safe set, and HEURISTIC SAFEOPT fails since it does not expand the safe set sufficiently.

the local maximum of the sinusoidal function. In other words, ITL adapts to the geometry of the function. This generally leads us to believe that ITL is more capable to exploit (non-stationary) prior knowledge than distance-based methods such as SAFEOPT.

## K.4 Experiment Details

### K.4.1 Synthetic Experiments

**1d task** Figure 23 shows the objective and constraint function, as well as the prior. We discretize using 500 points. The main difficulty in this experiment lies in sufficiently expanding the safe set to discover the global maximum. Figure 24 plots the size of the safe set  $\mathcal{S}_n$  for the compared algorithms, which in this experiment matches the achieved regret closely.

**2d task** We model our constraint in the form of a spherical “island” where the goal is to get a good view of the coral reef located to the north-east of the island while staying in the interior of the island during exploration (cf. Figure 25). The precise objective and constraint functions are unknown to the agent. Hence, the agent has to gradually and safely update its belief about boundaries of the “island” and the location of the coral reef. The prior is obtained by a single observation within the center of the island  $[-0.5, 0.5]^2$ . We discretize using 2 500 points.

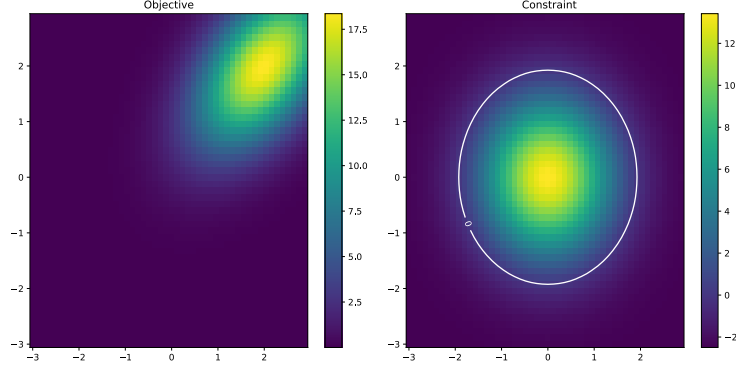


Figure 25: Ground truth in 2d synthetic experiment.

#### K.4.2 Safe Controller Tuning for Quadcopter

**Modeling the real-world dynamics** We learn a feedback policy (i.e., “control gains”) to compensate for inaccuracies in the initial controller. In our experiment, we model the real world dynamics and the adjusted model using the PD control feedback (Widmer et al., 2023),

$$\delta_t(\mathbf{x}) \stackrel{\text{def}}{=} (\mathbf{x}^* - \mathbf{x})[(\mathbf{s}^* - \mathbf{s}_t) (\dot{\mathbf{s}}^* - \dot{\mathbf{s}}_t)], \quad (49)$$

where  $\mathbf{x}^*$  are the *unknown* ground truth disturbance parameters, and  $\mathbf{s}^*$  and  $\dot{\mathbf{s}}^*$  are the desired state and state derivative, respectively. This yields the following ground truth dynamics:

$$\mathbf{s}_{t+1}(\mathbf{x}) = \mathbf{T}(\mathbf{s}_t, \mathbf{u}_t + \delta_t(\mathbf{x})). \quad (50)$$

The feedback parameters  $\mathbf{x} = [\mathbf{x}_p \ \mathbf{x}_d]^\top$  can be split into  $\mathbf{x}_p$  tuning the state difference which are called *proportional parameters* and  $\mathbf{x}_d$  tuning the state derivative difference which are called *derivative parameters*. We use the “critical damping” heuristic to relate the proportional and derivative parameters:  $\mathbf{x}_d = 2\sqrt{\mathbf{x}_p}$ . We thus consider the restricted domain  $\mathcal{X} = [0, 20]^4$  where each dimension corresponds to the proportional feedback to one of the four rotors.

Ground truth disturbance parameters are sampled from a chi-squared distribution with one degree of freedom (i.e., the square of a standard normal distribution),  $\mathbf{x}_p^* \sim \chi_1^2$ , and  $\mathbf{x}_d^*$  is determined according to the critical damping heuristic.

**The learning problem** The goal of our learning problem is to move the quadcopter from its initial position  $\mathbf{s}(0) = [1 \ 1 \ 1]^\top$  (in Euclidean space with meter as unit) to position  $\mathbf{s}^* = [0 \ 0 \ 2]^\top$ . Moreover, we aim to stabilize the quadcopter at the goal position, and therefore regularize the control signal towards an action  $\mathbf{u}^*$  which results in hovering (approximately) without any disturbances. We formalize these goals with the following objective function:

$$f^*(\mathbf{x}) \stackrel{\text{def}}{=} -\sigma \left( \sum_{t=0}^T \|\mathbf{s}^* - \mathbf{s}_t(\mathbf{x})\|_Q^2 + \|\mathbf{u}^* - \mathbf{u}_t(\mathbf{x})\|_R^2 \right) \quad (51)$$

where  $\sigma(v) \stackrel{\text{def}}{=} \tanh((v - 100)/100)$  is used to smoothen the objective function and ensure that its range is  $[-1, 1]$ . The non-smoothed control objective in Equation (51) is known as a *linear-quadratic regulator* (LQR) which we solve exactly for the undisturbed system using ILQR (Tu et al., 2023). Finally, we want to ensure at all times that the quadcopter is at least 0.5 meter above the ground, that is,

$$g^*(\mathbf{x}) \stackrel{\text{def}}{=} \min_{t \in [T]} \mathbf{s}_t^z(\mathbf{x}) - 0.5 \quad (52)$$

where we denote by  $\mathbf{s}_t^z$  the z-coordinate of state  $\mathbf{s}_t$ .

We use a time horizon of  $T = 3$  seconds which we discretize using 100 steps. The objective is modeled by a zero-mean GP with a Matérn( $\nu = 5/2$ ) kernel with lengthscale 0.1, and the constraint is modeled by a GP with mean  $-0.5$  and a Matérn( $\nu = 5/2$ ) kernel with lengthscale 0.1. The prior is obtained by a single observation of the “safe seed”  $[0 \ 0 \ 0 \ 10]^\top$ .

**Adaptive discretization** We discretize the domain  $\mathcal{X}$  adaptively using coordinate LINEBO (Kirschner et al., 2019). That is, in each iteration, one of the four control dimensions is selected uniformly at random, and the active learning oracle is executed on the corresponding one-dimensional subspace.

**Safety** Using the (unsafe) constrained BO algorithm EIC (Gardner et al., 2014) leads constraint violation,<sup>26</sup> while ITL and VTL do not violate the constraints during learning for any of the random seeds.

**Hyperparameters** The observation noise is Gaussian with standard deviation  $\rho = 0.1$ . We let  $\beta = 10$ . The control target is  $\mathbf{u}^* = [1.766 \ 0 \ 0 \ 0]^\top$ .

The state space is 12-dimensional where the first three states correspond to the velocity of the quadcopter, the next three states correspond to its acceleration, the following three states correspond to its angular velocity, and the last three states correspond to its angular velocity in local frame. The LQR parameters are given by

$$\begin{aligned} \mathbf{Q} &= \text{diag} \{1, 1, 1, 1, 1, 1, 0.1, 0.1, 0.1, 0.1, 0.1, 0.1\} \quad \text{and} \\ \mathbf{R} &= 0.01 \cdot \text{diag} \{5, 0.8, 0.8, 0.3\}. \end{aligned}$$

The quadcopter simulation was adapted from Chandra (2023).

Each one-dimensional subspace is discretized using 2 000 points.

**Random seeds** We repeat the experiment for 25 different seeds where the randomness is over the ground truth disturbance, observation noise, and the randomness in the algorithm.

---

<sup>26</sup>On average, 1.6 iterations of the first 50 violate the constraints.

Table 3: Magnitudes of  $\gamma_n$  for common kernels. The magnitudes hold under the assumption that  $\mathcal{X}$  is compact. Here,  $B_\nu$  is the modified Bessel function. We take the magnitudes from Theorem 5 of [Srinivas et al. \(2009\)](#) and Remark 2 of [Vakili et al. \(2021\)](#). The notation  $\tilde{O}(\cdot)$  subsumes log-factors. For  $\nu = 1/2$ , the Matérn kernel is equivalent to the Laplace kernel. For  $\nu \rightarrow \infty$ , the Matérn kernel is equivalent to the Gaussian kernel. The functions sampled from a Matérn kernel are  $\lceil \nu \rceil - 1$  mean square differentiable. The kernel-agnostic bound follows by simple reduction to a linear kernel in  $|\mathcal{X}|$  dimensions.

Kernel	$k(\mathbf{x}, \mathbf{x}')$	$\gamma_n$
Linear	$\mathbf{x}^\top \mathbf{x}'$	$O(d \log(n))$
Gaussian	$\exp\left(-\frac{\ \mathbf{x}-\mathbf{x}'\ _2^2}{2h^2}\right)$	$\tilde{O}\left(\log^{d+1}(n)\right)$
Laplace	$\exp\left(-\frac{\ \mathbf{x}-\mathbf{x}'\ _1}{h}\right)$	$\tilde{O}\left(n^{\frac{d}{1+d}} \log^{\frac{1}{1+d}}(n)\right)$
Matérn	$\frac{2^{1-\nu}}{\Gamma(\nu)} \left(\frac{\sqrt{2\nu}\ \mathbf{x}-\mathbf{x}'\ _2}{h}\right)^\nu B_\nu\left(\frac{\sqrt{2\nu}\ \mathbf{x}-\mathbf{x}'\ _2}{h}\right)$	$\tilde{O}\left(n^{\frac{d}{2\nu+d}} \log^{\frac{2\nu}{2\nu+d}}(n)\right)$
any		$O( \mathcal{X}  \log(n))$



## NeurIPS Paper Checklist

### 1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification:

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

### 2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification:

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

### 3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification:

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

#### 4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: The code is publicly available.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
  - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

#### 5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification:

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

## 6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: See the extensive discussions in the appendices.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

## 7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification:

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.

- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

## 8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: All individual experiments require only small compute resources.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

## 9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines?>

Answer: [Yes]

Justification:

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

## 10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: This paper presents work whose goal is to advance the field of Machine Learning. There are many potential societal consequences of our work, none which we feel must be specifically highlighted here.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

## 11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

## 12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification:

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, [paperswithcode.com/datasets](https://paperswithcode.com/datasets) has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

- If this information is not available online, the authors are encouraged to reach out to the asset’s creators.

### 13. **New Assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [Yes]

Justification:

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

### 14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

### 15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.