Gaussian Process Bandit Optimisation with Multi-fidelity Evaluations

Kirthevasan Kandasamy [♮], **Gautam Dasarathy** [♢], **Junier Oliva** [♮],
Jeff Schneider [♮], **Barnabás Póczos** [♮]
[♮] Carnegie Mellon University, [♢] Rice University
{kandasamy, joliva, schneide, bapoczos}@cs.cmu.edu, gautamd@rice.edu

Abstract

In many scientific and engineering applications, we are tasked with the optimisation of an expensive to evaluate black box function f. Traditional methods for this problem assume just the availability of this single function. However, in many cases, cheap approximations to f may be obtainable. For example, the expensive real world behaviour of a robot can be approximated by a cheap computer simulation. We can use these approximations to eliminate low function value regions cheaply and use the expensive evaluations of f in a small but promising region and speedily identify the optimum. We formalise this task as a *multi-fidelity* bandit problem where the target function and its approximations are sampled from a Gaussian process. We develop MF-GP-UCB, a novel method based on upper confidence bound techniques. In our theoretical analysis we demonstrate that it exhibits precisely the above behaviour, and achieves better regret than strategies which ignore multi-fidelity information. MF-GP-UCB outperforms such naive strategies and other multi-fidelity methods on several synthetic and real experiments.

1 Introduction

In stochastic bandit optimisation, we wish to optimise a *payoff* function $f : \mathcal{X} \to \mathbb{R}$ by sequentially querying it and obtaining *bandit feedback*, i.e. when we query at any $x \in \mathcal{X}$, we observe a possibly noisy evaluation of $f(x)$. f is typically expensive and the goal is to identify its maximum while keeping the number of queries as low as possible. Some applications are hyper-parameter tuning in expensive machine learning algorithms, optimal policy search in complex systems, and scientific experiments [20, 23, 27]. Historically, bandit problems were studied in settings where the goal is to maximise the cumulative reward of all queries to the payoff instead of just finding the maximum. Applications in this setting include clinical trials and online advertising.

Conventional methods in these settings assume access to only this single expensive function of interest f. We will collectively refer to them as *single fidelity* methods. In many practical problems however, cheap approximations to f might be available. For instance, when tuning hyper-parameters of learning algorithms, the goal is to maximise a cross validation (CV) score on a training set, which can be expensive if the training set is large. However CV curves tend to vary smoothly with training set size; therefore, we can train and cross validate on small subsets to approximate the CV accuracies of the entire dataset. For a concrete example, consider kernel density estimation (KDE), where we need to tune the bandwidth h of a kernel. Figure 1 shows the CV likelihood against h for a dataset of size $n = 3000$ and a smaller subset of size $n = 300$. The two maximisers are different, which is to be expected since optimal hyper-parameters are functions of the training set size. That said, the curve for $n = 300$ approximates the $n = 3000$ curve quite well. Since training/CV on small n is cheap, we can use it to eliminate bad values of the hyper-parameters and reserve the expensive experiments with the entire dataset for the promising candidates (e.g. boxed region in Fig. 1).

In online advertising, the goal is to maximise the cumulative number of clicks over a given period. In the conventional bandit treatment, each query to f is the display of an ad for a specific time, say one

30th Conference on Neural Information Processing Systems (NIPS 2016), Barcelona, Spain.

hour. However, we may display ads for shorter intervals, say a few minutes, to approximate its hourly performance. The estimate is biased, as displaying an ad for a longer interval changes user behaviour, but will nonetheless be useful in gauging its long run click through rate. In optimal policy search in robotics and automated driving vastly cheaper computer simulations are used to approximate the expensive real world performance of the system. Scientific experiments can be approximated to varying degrees using less expensive data collection, analysis, and computational techniques.

In this paper, we cast these tasks as *multi-fidelity bandit optimisation* problems assuming the availability of cheap approximate functions (fidelities) to the payoff f. **Our contributions** are:

1. We present a formalism for multi-fidelity bandit optimisation using Gaussian Process (GP) assumptions on f and its approximations. We develop a novel algorithm, Multi-Fidelity Gaussian Process Upper Confidence Bound (MF-GP-UCB) for this setting.
2. Our theoretical analysis proves that MF-GP-UCB explores the space at lower fidelities and uses the high fidelities in successively smaller regions to zero in on the optimum. As lower fidelity queries are cheaper, MF-GP-UCB has better regret than single fidelity strategies.
3. Empirically, we demonstrate that MF-GP-UCB outperforms single fidelity methods on a series of synthetic examples, three hyper-parameter tuning tasks and one inference problem in Astrophysics. Our matlab implementation and experiments are available at `github.com/kirthevasank/mf-gp-ucb`.

Related Work: Since the seminal work by Robbins [25], the multi-armed bandit problem has been studied extensively in the K-armed setting. Recently, there has been a surge of interest in the optimism under uncertainty principle for K armed bandits, typified by upper confidence bound (UCB) methods [2, 4]. UCB strategies have also been used in bandit tasks with linear [6] and GP [28] payoffs. There is a plethora of work on single fidelity methods for global optimisation both with noisy and noiseless evaluations. Some examples are branch and bound techniques such as dividing rectangles (DiRect) [12], simulated annealing, genetic algorithms and more [17, 18, 22]. A suite of single fidelity methods in the GP framework closely related to our work is Bayesian Optimisation (BO). While there are several techniques for BO [13, 21, 30], of particular interest to us is the Gaussian process upper confidence bound (GP-UCB) algorithm of Srinivas et al. [28].

Many applied domains of research such as aerodynamics, industrial design and hyper-parameter tuning have studied multi-fidelity methods [9, 11, 19, 29]; a plurality of them use BO techniques. However none of these treatments neither formalise nor analyse any notion of *regret* in the multi-fidelity setting. In contrast, MF-GP-UCB is an intuitive UCB idea with good theoretical properties. Some literature have analysed multi-fidelity methods in specific contexts such as hyper-parameter tuning, active learning and reinforcement learning [1, 5, 26, 33]. Their settings and assumptions are substantially different from ours. Critically, none of them are in the more difficult bandit setting where there is a price for exploration. Due to space constraints we discuss them in detail in Appendix A.3.

The multi-fidelity poses substantially new theoretical and algorithmic challenges. We build on GP-UCB and our recent work on multi-fidelity bandits in the K-armed setting [16]. Section 2 presents our formalism including a notion of regret for multi-fidelity GP bandits. Section 3 presents our algorithm. The theoretical analysis is in Appendix C with a synopsis for the 2-fidelity case in Section 4. Section 6 presents our experiments. Appendix A.1 tabulates the notation used in the manuscript.

2 Preliminaries

We wish to maximise a payoff function $f : \mathcal{X} \to \mathbb{R}$ where $\mathcal{X} \equiv [0, r]^d$. We can interact with f only by querying at some $x \in \mathcal{X}$ and obtaining a noisy observation $y = f(x) + \epsilon$. Let $x_\star \in \operatorname{argmax}_{x \in \mathcal{X}} f(x)$ and $f_\star = f(x_\star)$. Let $\mathbf{x}_t \in \mathcal{X}$ be the point queried at time t. The goal of a bandit strategy is to maximise the sum of rewards $\sum_{t=1}^n f(\mathbf{x}_t)$ or equivalently minimise the *cumulative regret* $\sum_{t=1}^n f_\star - f(\mathbf{x}_t)$ after n queries; i.e. we compete against an oracle which queries at $x_\star$ at all t.

Our primary distinction from the classical setting is that we have access to $M-1$ successively accurate approximations $f^{(1)}, f^{(2)}, \ldots, f^{(M-1)}$ to the payoff $f = f^{(M)}$. We refer to these approximations as fidelities. We encode the fact that fidelity m approximates fidelity M via the assumption, $\|f^{(M)} - f^{(m)}\|_\infty \leq \zeta^{(m)}$, where $\zeta^{(1)} > \zeta^{(2)} > \cdots > \zeta^{(M)} = 0$. Each query at fidelity m expends a cost $\lambda^{(m)}$ of a resource, e.g. computational effort or advertising time, where $\lambda^{(1)} < \lambda^{(2)} < \cdots < \lambda^{(M)}$. A strategy for multi-fidelity bandits is a sequence of query-fidelity pairs $\{(\mathbf{x}_t, \mathbf{m}_t)\}_{t \geq 0}$, where

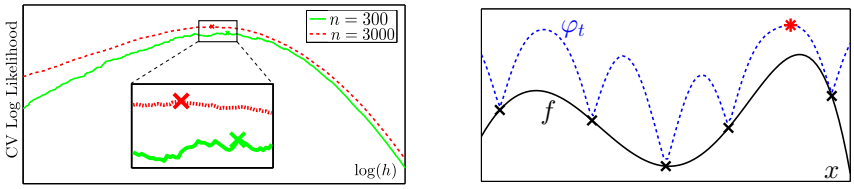

Figure 1: *Left:* Average CV log likelihood on datasets of size $300, 3000$ on a synthetic KDE task. The crosses are the maxima. *Right:* Illustration of GP-UCB at time t. The figure shows $f(x)$ (solid black line), the UCB $\varphi_t(x)$ (dashed blue line) and queries until $t-1$ (black crosses). We query at $\mathbf{x}_t = \text{argmax}_{x\in\mathcal{X}} \varphi_t(x)$ (red star).

$(\mathbf{x}_n, \mathbf{m}_n)$ could depend on the previous query-observation-fidelity tuples $\{(\mathbf{x}_t, \mathbf{y}_t, \mathbf{m}_t)\}_{t=1}^{n-1}$. Here $\mathbf{y}_t = f^{(\mathbf{m}_t)}(\mathbf{x}_t) + \epsilon$. After n steps we will have queried any of the M fidelities multiple times.

Some smoothness assumptions on $f^{(m)}$'s are needed to make the problem tractable. A standard in the Bayesian nonparametric literature is to use a Gaussian process (GP) prior [24] with covariance kernel κ. In this work we focus on the squared exponential (SE) $\kappa_{\sigma,h}$ and the Matérn $\kappa_{\nu,h}$ kernels as they are popularly used in practice and their theoretical properties are well studied. Writing $z = \|x - x'\|_2$, they are defined as $\kappa_{\sigma,h}(x,x') = \sigma \exp\left(-z^2/(2h^2)\right)$, $\kappa_{\nu,h}(x,x') = \frac{2^{1-\nu}}{\Gamma(\nu)}\left(\frac{\sqrt{2\nu}z}{h}\right)^{\nu} B_\nu\left(\frac{\sqrt{2\nu}z}{h}\right)$, where Γ, B_ν are the Gamma and modified Bessel functions. A convenience the GP framework offers is that posterior distributions are analytically tractable. If $f \sim \mathcal{GP}(0,\kappa)$, and we have observations $\mathcal{D}_n = \{(x_i, y_i)\}_{i=1}^n$, where $y_i = f(x_i) + \epsilon$ and $\epsilon \sim \mathcal{N}(0,\eta^2)$ is Gaussian noise, the posterior distribution for $f(x)|\mathcal{D}_n$ is also Gaussian $\mathcal{N}(\mu_n(x), \sigma_n^2(x))$ with

$$\mu_n(x) = \mathbf{k}^\top \Delta^{-1} Y, \qquad \sigma_n^2(x) = \kappa(x,x) - \mathbf{k}^\top \Delta^{-1} \mathbf{k}. \tag{1}$$

Here, $Y \in \mathbb{R}^n$ with $Y_i = y_i$, $\mathbf{k} \in \mathbb{R}^n$ with $\mathbf{k}_i = \kappa(x, x_i)$ and $\Delta = \mathbf{K} + \eta^2 I \in \mathbb{R}^{n\times n}$ where $\mathbf{K}_{i,j} = \kappa(x_i, x_j)$. In keeping with the above, we make the following assumptions on our problem.

Assumption 1. **A1**: *The functions at all fidelities are sampled from GPs,* $f^{(m)} \sim \mathcal{GP}(\mathbf{0}, \kappa)$ *for all* $m = 1, \ldots, M$. **A2**: $\|f^{(M)} - f^{(m)}\|_\infty \leq \zeta^{(m)}$ *for all* $m = 1, \ldots, M$. **A3**: $\|f^{(M)}\|_\infty \leq B$.

The purpose of **A3** is primarily to define the regret. In Remark 7, Appendix A.4 we argue that these assumptions are probabilistically valid, i.e. the latter two events occur with nontrivial probability when we sample the $f^{(m)}$'s from a GP. So a generative mechanism would keep sampling the functions and deliver them when the conditions hold true. A point $x \in \mathcal{X}$ can be queried at any of the M fidelities. When we query at fidelity m, we observe $y = f^{(m)}(x) + \epsilon$ where $\epsilon \sim \mathcal{N}(0, \eta^2)$.

We now present our notion of cumulative regret $R(\Lambda)$ after spending capital Λ of a resource in the multi-fidelity setting. $R(\Lambda)$ should reduce to the conventional definition of regret for any single fidelity strategy that queries only at M^{th} fidelity. As only the optimum of $f = f^{(M)}$ is of interest to us, queries at fidelities less than M should yield the lowest possible reward, $(-B)$ according to **A3**. Accordingly, we set the instantaneous reward q_t at time to be $-B$ if $\mathbf{m}_t \neq M$ and $f^{(M)}(\mathbf{x}_t)$ if $\mathbf{m}_t = M$. If we let $r_t = f_\star - q_t$ denote the instantaneous regret, we have $r_t = f_\star + B$ if $\mathbf{m}_t \neq M$ and $f_\star - f(\mathbf{x}_t)$ if $\mathbf{m}_t = M$. $R(\Lambda)$ should also factor in the costs of the fidelity of each query. Finally, we should also receive $(-B)$ reward for any unused capital. Accordingly, we define $R(\Lambda)$ as,

$$R(\Lambda) \;=\; \Lambda f_\star - \left[\sum_{t=1}^{N} \lambda^{(m_t)} q_t + \left(\Lambda - \sum_{t=1}^{N} \lambda^{(m_t)}\right)(-B)\right] \;\leq\; 2B\Lambda_{res} \;+\; \sum_{t=1}^{N} \lambda^{(m_t)} r_t, \tag{2}$$

where $\Lambda_{res} = \Lambda - \sum_{t=1}^{N} \lambda^{(m_t)}$. Here, N is the (random) number of queries at all fidelities within capital Λ, i.e. the largest n such that $\sum_{t=1}^{n} \lambda^{(m_t)} \leq \Lambda$. According to (2) above, we wish to compete against an oracle that uses all its capital Λ to query $x_\star$ at the M^{th} fidelity. $R(\Lambda)$ is at best 0 when we follow the oracle and at most $2\Lambda B$. Our goal is a strategy that has small regret for all values of (sufficiently large) Λ, i.e. the equivalent of an anytime strategy, as opposed to a fixed time horizon strategy in the usual bandit setting. For the purpose of optimisation, we also define the *simple regret* as $S(\Lambda) = \min_t r_t = f_\star - \max_t q_t$. $S(\Lambda)$ is the difference between $f_\star$ and the best highest fidelity query (and $f_\star + B$ if we have never queried at fidelity M). Since $S(\Lambda) \leq \frac{1}{\Lambda} R(\Lambda)$, any strategy with asymptotic sublinear regret $\lim_{\Lambda\to\infty} \frac{1}{\Lambda} R(\Lambda) = 0$, also has vanishing simple regret.

Since, to our knowledge, this is the first attempt to formalise regret for multi-fidelity problems, the definition for $R(\Lambda)$ (2) necessitates justification. Consider a two fidelity robot gold mining problem

where the second fidelity is a real world robot trial, costing $\lambda^{(2)}$ dollars and the first fidelity is a computer simulation costing $\lambda^{(1)}$. A multi-fidelity algorithm queries the simulator to learn about the real world. But it does not collect any actual gold during a simulation; hence no reward, which according to our assumptions is $-B$. Meantime the oracle is investing this capital on the best experiment and collecting $\sim f_\star$ gold. Therefore, the regret at this time instant is $f_\star + B$. However we weight this by the cost to account for the fact that the simulation costs only $\lambda^{(1)}$. Note that lower fidelities use up capital but yield the lowest reward. The goal however, is to leverage information from these cheap queries to query prudently at the highest fidelity and obtain better regret.

That said, other multi-fidelity settings might require different definitions for $R(\Lambda)$. In online advertising, the lower fidelities (displaying ads for shorter periods) would still yield rewards. In clinical trials, the regret at the highest fidelity due to a bad treatment would be, say, a dead patient. However, a bad treatment on a simulation may not warrant large penalty. We use the definition in (2) because it is more aligned with our optimisation experiments: lower fidelities are useful to the extent that they guide search on the expensive $f^{(M)}$, but there is no reward to finding the optimum of a cheap $f^{(m)}$.

A crucial challenge for a multi-fidelity method is to not get stuck at the optimum of a lower fidelity, which is typically suboptimal for $f^{(M)}$. While exploiting information from the lower fidelities, it is also important to *explore* sufficiently at $f^{(M)}$. In our experiments we demonstrate that naive strategies which do not do so would get stuck at the optimum of a lower fidelity.

A note on GP-UCB: Sequential optimisation methods adopting UCB principles maintain a high probability upper bound $\varphi_t : \mathcal{X} \to \mathbb{R}$ for $f(x)$ for all $x \in \mathcal{X}$ [2]. For GP-UCB, φ_t takes the form $\varphi_t(x) = \mu_{t-1}(x) + \beta_t^{1/2}\sigma_{t-1}(x)$ where μ_{t-1}, σ_{t-1} are the posterior mean and standard deviation of the GP conditioned on the previous $t-1$ queries. The key intuition is that the mean μ_{t-1} encourages an exploitative strategy – in that we want to query where we know the function is high – and the confidence band $\beta_t^{1/2}\sigma_{t-1}$ encourages an explorative strategy – in that we want to query at regions we are uncertain about f lest we miss out on high valued regions. We have illustrated GP-UCB in Fig 1 and reviewed the algorithm and its theoretical properties in Appendix A.2.

3 MF-GP-UCB

The proposed algorithm, MF-GP-UCB, will also maintain a UCB for $f^{(M)}$ obtained via the previous queries at *all* fidelities. Denote the posterior GP mean and standard deviation of $f^{(m)}$ conditioned *only* on the previous queries at fidelity m by $\mu_t^{(m)}, \sigma_t^{(m)}$ respectively (See (1)). Then define,

$$\varphi_t^{(m)}(x) = \mu_{t-1}^{(m)}(x) + \beta_t^{1/2}\sigma_{t-1}^{(m)}(x) + \zeta^{(m)}, \quad \forall\, m, \qquad\qquad \varphi_t(x) = \min_{m=1,\ldots,M} \varphi_t^{(m)}(x). \tag{3}$$

For appropriately chosen β_t, $\mu_{t-1}^{(m)}(x) + \beta_t^{1/2}\sigma_{t-1}^{(m)}(x)$ will upper bound $f^{(m)}(x)$ with high probability. By **A2**, $\varphi_t^{(m)}(x)$ upper bounds $f^{(M)}(x)$ for all m. We have M such upper bounds, and their minimum $\varphi_t(x)$ gives the best bound. Our next query is at the maximiser of this UCB, $\mathbf{x}_t = \text{argmax}_{x\in\mathcal{X}}\, \varphi_t(x)$.

Next we need to decide which fidelity to query at. Consider any $m < M$. The $\zeta^{(m)}$ conditions on $f^{(m)}$ constrain the value of $f^{(M)}$ – the confidence band $\beta_t^{1/2}\sigma_{t-1}^{(m)}$ for $f^{(m)}$ is lengthened by $\zeta^{(m)}$ to obtain confidence on $f^{(M)}$. If $\beta_t^{1/2}\sigma_{t-1}^{(m)}(\mathbf{x}_t)$ for $f^{(m)}$ is large, it means that we have not constrained $f^{(m)}$ sufficiently well at $\mathbf{x}_t$ and should query at the m^{th} fidelity. On the other hand, querying indefinitely in the same region to reduce $\beta_t^{1/2}\sigma_{t-1}^{(m)}$ in that region will not help us much as the $\zeta^{(m)}$ elongation caps off how much we can learn about $f^{(M)}$ from $f^{(m)}$; i.e. even if we knew $f^{(m)}$ perfectly, we will only have constrained $f^{(M)}$ to within a $\pm\zeta^{(m)}$ band. Our algorithm captures this simple intuition. Having selected $\mathbf{x}_t$, we begin by checking at the first fidelity. If $\beta_t^{1/2}\sigma_{t-1}^{(1)}(\mathbf{x}_t)$ is smaller than a threshold $\gamma^{(1)}$, we proceed to the second fidelity. If at any stage $\beta_t^{1/2}\sigma_{t-1}^{(m)}(\mathbf{x}_t) \geq \gamma^{(m)}$ we query at fidelity $\mathbf{m}_t = m$. If we proceed all the way to fidelity M, we query at $\mathbf{m}_t = M$. We will discuss choices for $\gamma^{(m)}$ shortly. We summarise the resulting procedure in Algorithm 1.

Fig 2 illustrates MF-GP-UCB on a 2–fidelity problem. Initially, MF-GP-UCB is mostly exploring $\mathcal{X}$ in the first fidelity. $\beta_t^{1/2}\sigma_{t-1}^{(1)}$ is large and we are yet to constrain $f^{(1)}$ well to proceed to $f^{(2)}$. By $t = 14$, we have constrained $f^{(1)}$ around the optimum and have started querying at $f^{(2)}$ in this region.

Algorithm 1 MF-GP-UCB **Inputs:** kernel κ, bounds $\{\zeta^{(m)}\}_{m=1}^{M}$, thresholds $\{\gamma^{(m)}\}_{m=1}^{M}$.

- For $m = 1, \ldots, M$: $\mathcal{D}_0^{(m)} \leftarrow \varnothing$, $(\mu_0^{(m)}, \sigma_0^{(m)}) \leftarrow (\mathbf{0}, \kappa^{1/2})$.
- for $t = 1, 2, \ldots$
  1. $\mathbf{x}_t \leftarrow \text{argmax}_{x \in \mathcal{X}} \ \varphi_t(x)$. (See Equation (3))
  2. $\mathbf{m}_t = \min_m \{ m \,|\beta_t^{1/2} \sigma_{t-1}^{(m)}(\mathbf{x}_t) \geq \gamma^{(m)} \ \text{ or } \ m = M\}$. (See Appendix B, C for β_t)
  3. $\mathbf{y}_t \leftarrow$ Query $f^{(\mathbf{m}_t)}$ at $\mathbf{x}_t$.
  4. Update $\mathcal{D}_t^{(\mathbf{m}_t)} \leftarrow \mathcal{D}_{t-1}^{(\mathbf{m}_t)} \cup \{(\mathbf{x}_t, \mathbf{y}_t)\}$. Obtain $\mu_t^{(\mathbf{m}_t)}, \sigma_t^{(\mathbf{m}_t)}$ conditioned on $\mathcal{D}_t^{(\mathbf{m}_t)}$ (See (1)).

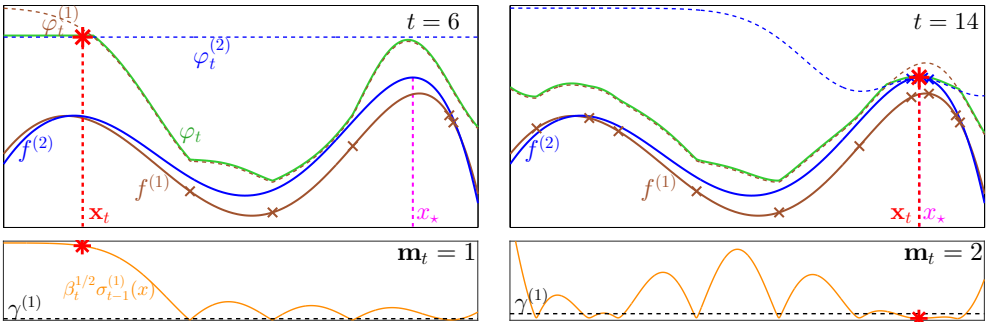

Figure 2: Illustration of MF-GP-UCB for a 2-fidelity problem initialised with 5 random points at the first fidelity. In the top figures, the solid lines in brown and blue are $f^{(1)}, f^{(2)}$ respectively, and the dashed lines are $\varphi_t^{(1)}, \varphi_t^{(2)}$. The solid green line is $\varphi_t = \min(\varphi_t^{(1)}, \varphi_t^{(2)})$. The small crosses are queries from 1 to $t-1$ and the red star is the maximiser of φ_t, i.e. the next query $\mathbf{x}_t$. $x_\star$, the optimum of $f^{(2)}$ is shown in magenta. In the bottom figures, the solid orange line is $\beta_t^{1/2}\sigma_{t-1}^{(1)}$ and the dashed black line is $\gamma^{(1)}$. When $\beta_t^{1/2}\sigma_{t-1}^{(1)}(\mathbf{x}_t) \leq \gamma^{(1)}$ we play at fidelity $\mathbf{m}_t = 2$ and otherwise at $\mathbf{m}_t = 1$. See Fig. 6 in Appendix B for an extended simulation.

Notice how $\varphi_t^{(2)}$ dips to change φ_t in this region. MF-GP-UCB has identified the maximum with just 3 queries to $f^{(2)}$. In Appendix B we provide an extended simulation and discuss further insights.

Finally, we make an essential observation. The posterior for any $f^{(m)}(x)$ conditioned on previous queries at *all* fidelities is not Gaussian due to the $\zeta^{(m)}$ constraints (**A2**). However, $|f^{(m)}(x) - \mu_{t-1}^{(m)}(x)| < \beta_t^{1/2}\sigma_{t-1}^{(m)}(x)$ holds with high probability, since, by conditioning only on queries at the m^{th} fidelity we have Gaussianity for $f^{(m)}(x)$. Next we summarise our main theoretical contributions.

4 Summary of Theoretical Results

For pedagogical reasons we present our results for the $M = 2$ case. Appendix C contains statements and proofs for general M. We also ignore constants and polylog terms when they are dominated by other terms. $\lesssim, \asymp$ denote inequality and equality ignoring constants. We begin by defining the *Maximum Information Gain* (MIG) which characterises the statistical difficulty of GP bandits.

Definition 2. *(Maximum Information Gain) Let $f \sim \mathcal{GP}(\mathbf{0}, \kappa)$. Consider any $A \subset \mathbb{R}^d$ and let $\widetilde{A} = \{x_1, \ldots, x_n\} \subset A$ be a finite subset. Let $f_{\widetilde{A}}, \epsilon_{\widetilde{A}} \in \mathbb{R}^n$ be such that $(f_{\widetilde{A}})_i = f(x_i)$, $(\epsilon_{\widetilde{A}})_i \sim \mathcal{N}(0, \eta^2)$, and $y_{\widetilde{A}} = f_{\widetilde{A}} + \epsilon_{\widetilde{A}}$. Let I denote the Shannon Mutual Information. The Maximum Information Gain of A is $\Psi_n(A) = \max_{\widetilde{A} \subset A, |\widetilde{A}|=n} I(y_{\widetilde{A}}; f_{\widetilde{A}})$.*

The MIG, which depends on the kernel κ and the set A, is an important quantity in our analysis. For a given κ, it typically scales with the volume of A; i.e. if $A = [0, r]^d$ then $\Psi_n(A) \in \mathcal{O}(r^d \Psi_n([0,1]^d))$. For the SE kernel, $\Psi_n([0,1]^d) \in \mathcal{O}((\log(n))^{d+1})$ and for Matérn, $\Psi_n([0,1]^d) \in \mathcal{O}(n^{\frac{d(d+1)}{2\nu + d(d+1)}})$ [28].

Recall, N is the (random) number of queries by a multi-fidelity strategy within capital Λ at either fidelity. Let $n_\Lambda = \lfloor \Lambda / \lambda^{(2)} \rfloor$ be the (non-random) number of queries by a single fidelity method operating only at the second fidelity. As $\lambda^{(1)} < \lambda^{(2)}$, N could be large for an arbitrary multi-fidelity method. However, our analysis reveals that for MF-GP-UCB, N is on the order of n_Λ.

Fundamental to the 2-fidelity problem is the set $\mathcal{X}_g = \{x \in \mathcal{X}; f_\star - f^{(1)}(x) \leq \zeta^{(1)}\}$. $\mathcal{X}_g$ is a high valued region for $f^{(2)}(x)$: for all $x \in \mathcal{X}_g$, $f^{(2)}(x)$ is at most $2\zeta^{(1)}$ away from the optimum. More interestingly, when $\zeta^{(1)}$ is small, i.e. when $f^{(1)}$ is a good approximation to $f^{(2)}$, $\mathcal{X}_g$ will be much smaller than $\mathcal{X}$. This is precisely the target domain for this research. For instance, in the robot gold mining example, a cheap computer simulator can be used to eliminate several bad policies and we could reserve the real world trials for the promising candidates. If a multi-fidelity strategy were to use the second fidelity queries only in $\mathcal{X}_g$, then the regret will only have $\Psi_n(\mathcal{X}_g)$ dependence after n high fidelity queries. In contrast, a strategy that only operates at the highest fidelity (e.g. GP-UCB) will have $\Psi_n(\mathcal{X})$ dependence. In the scenario described above $\Psi_n(\mathcal{X}_g) \ll \Psi_n(\mathcal{X})$, and the multi-fidelity strategy will have significantly better regret than a single fidelity strategy. MF-GP-UCB roughly achieves this goal. In particular, we consider a slightly inflated set $\widetilde{\mathcal{X}}_{g,\rho} = \{x \in \mathcal{X}; f_\star - f^{(1)}(x) \leq \zeta^{(1)} + \rho\gamma^{(1)}\}$, of $\mathcal{X}_g$ where $\rho > 0$. The following result which characterises the regret of MF-GP-UCB in terms of $\widetilde{\mathcal{X}}_{g,\rho}$ is the main theorem of this paper.

Theorem 3 (Regret of MF-GP-UCB – Informal)**.** *Let $\mathcal{X} = [0, r]^d$ and $f^{(1)}, f^{(2)} \sim \mathcal{GP}(\mathbf{0}, \kappa)$ satisfy Assumption 1. Pick $\delta \in (0, 1)$ and run* MF-GP-UCB *with $\beta_t \asymp d \log(t/\delta)$. Then, with probability $> 1 - \delta$, for sufficiently large Λ and for all $\alpha \in (0, 1)$, there exists ρ depending on α such that,*

$$R(\Lambda) \lesssim \lambda^{(2)}\sqrt{n_\Lambda \beta_{n_\Lambda} \Psi_{n_\Lambda}(\widetilde{\mathcal{X}}_{g,\rho})} + \lambda^{(1)}\sqrt{n_\Lambda \beta_{n_\Lambda} \Psi_{n_\Lambda}(\mathcal{X})} + \lambda^{(2)}\sqrt{n_\Lambda^\alpha \beta_{n_\Lambda} \Psi_{n_\Lambda^\alpha}(\mathcal{X})} + \lambda^{(1)}\xi_{n,\widetilde{\mathcal{X}}_{g,\rho},\gamma^{(1)}}$$

As we will explain shortly, the latter two terms are of lower order. It is instructive to compare the above rates against that for GP-UCB (see Theorem 4, Appendix A.2). By dropping the common and subdominant terms, the rate for MF-GP-UCB is $\lambda^{(2)}\Psi^{1/2}_{n_\Lambda}(\widetilde{\mathcal{X}}_{g,\rho}) + \lambda^{(1)}\Psi^{1/2}_{n_\Lambda}(\mathcal{X})$ whereas for GP-UCB it is $\lambda^{(2)}\Psi^{1/2}_{n_\Lambda}(\mathcal{X})$. When $\lambda^{(1)} \ll \lambda^{(2)}$ and $\mathrm{vol}(\widetilde{\mathcal{X}}_{g,\rho}) \ll \mathrm{vol}(\mathcal{X})$ the rates for MF-GP-UCB are very appealing. When the approximation worsens ($\mathcal{X}_g, \widetilde{\mathcal{X}}_{g,\rho}$ become larger) and the costs $\lambda^{(1)}, \lambda^{(2)}$ become comparable, the bound for MF-GP-UCB decays gracefully. In the worst case, MF-GP-UCB is never worse than GP-UCB up to constant terms. Intuitively, the above result states that MF-GP-UCB explores the entire $\mathcal{X}$ using $f^{(1)}$ but uses "most" of its queries to $f^{(2)}$ inside $\widetilde{\mathcal{X}}_{g,\rho}$.

Now let us turn to the latter two terms in the bound. The third term is the regret due to the second fidelity queries outside $\widetilde{\mathcal{X}}_{g,\rho}$. We are able to show that the number of such queries is $\mathcal{O}(n_\Lambda^\alpha)$ for all $\alpha > 0$ for an appropriate ρ. This *strong* result is only possible in the multi-fidelity setting. For example, in GP-UCB the best bound you can achieve on the number of plays on a suboptimal set is $\mathcal{O}(n_\Lambda^{1/2})$ for the SE kernel and worse for the Matérn kernel. The last term is due to the first fidelity plays inside $\widetilde{\mathcal{X}}_{g,\rho}$ and it scales with $\mathrm{vol}(\widetilde{\mathcal{X}}_{g,\rho})$ and polylogarithmically with n, both of which are small. However, it has a $1/\mathrm{poly}(\gamma^{(1)})$ dependence which could be bad if $\gamma^{(1)}$ is too small: intuitively, if $\gamma^{(1)}$ is too small then you will wait for a long time in step 2 of Algorithm 1 for $\beta_t^{1/2}\sigma^{(1)}_{t-1}$ to decrease without proceeding to $f^{(2)}$, incurring large regret $(f_\star + B)$ in the process. Our analysis reveals that an optimal choice for the SE kernel scales $\gamma^{(1)} \asymp (\lambda^{(1)}\zeta^{(1)}/(t\lambda^{(2)}))^{1/(d+2)}$ at time t. However this is of little practical use as the leading constant depends on several problem dependent quantities such as $\Psi_n(\mathcal{X}_g)$. In Section 5 we describe a heuristic to set $\gamma^{(m)}$ which worked well in our experiments.

Theorem 3 can be generalised to cases where the kernels $\kappa^{(m)}$ and observation noises $\eta^{(m)}$ are different at each fidelity. The changes to the proofs are minimal. In fact, our practical implementation uses different kernels. As with any nonparametric method, our algorithm has exponential dependence on dimension. This can be alleviated by assuming additional structure in the problem [8, 15]. Finally, we note that the above rates translate to bounds on the simple regret $S(\Lambda)$ for optimisation.

5 Implementation Details

Our implementation uses some standard techniques in Bayesian optimisation to learn the kernel such as initialisation with random queries and periodic marginal likelihood maximisation. The above techniques might be already known to a reader familiar with the BO literature. We have elaborated these in Appendix B but now focus on the $\gamma^{(m)}, \zeta^{(m)}$ parameters of our method.

Algorithm 1 assumes that the $\zeta^{(m)}$'s are given with the problem description, which is hardly the case in practice. In our implementation, instead of having to deal with $M-1$, $\zeta^{(m)}$ values we set $(\zeta^{(1)}, \zeta^{(2)}, \ldots, \zeta^{(M-1)}) = ((M-1)\zeta, (M-2)\zeta, \ldots, \zeta)$ so we only have one value ζ. This for

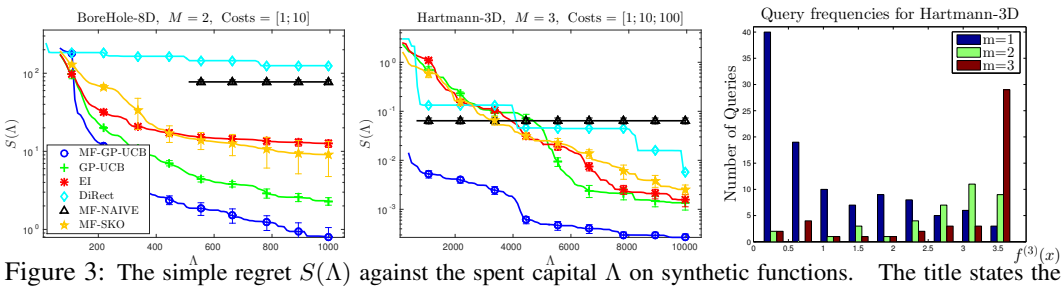

Figure 3: The simple regret $S(\Lambda)$ against the spent capital Λ on synthetic functions. The title states the function, its dimensionality, the number of fidelities and the costs we used for each fidelity in the experiment. All curves barring DiRect (which is a deterministic), were produced by averaging over 20 experiments. The error bars indicate one standard error. See Figures 8, 9 10 in Appendix D for more synthetic results. The last panel shows the number of queries at different function values at each fidelity for the Hartmann-3D example.

instance, is satisfied if $\|f^{(m)} - f^{(m-1)}\|_\infty \leq \zeta$ which is stronger than Assumption **A2**. Initially, we start with small ζ. Whenever we query at any fidelity $m > 1$ we also check the posterior mean of the $(m-1)^{\text{th}}$ fidelity. If $|f^{(m)}(\mathbf{x}_t) - \mu_{t-1}^{(m-1)}(\mathbf{x}_t)| > \zeta$, we query again at $\mathbf{x}_t$, but at the $(m-1)^{\text{th}}$ fidelity. If $|f^{(m)}(\mathbf{x}_t) - f^{(m-1)}(\mathbf{x}_t)| > \zeta$, we update ζ to twice the violation. To set $\gamma^{(m)}$'s we use the following intuition: if the algorithm, is stuck at fidelity m for too long then $\gamma^{(m)}$ is probably too small. We start with small values for $\gamma^{(m)}$. If the algorithm does not query above the m^{th} fidelity for more than $\lambda^{(m+1)}/\lambda^{(m)}$ iterations, we double $\gamma^{(m)}$. We found our implementation to be fairly robust even recovering from fairly bad approximations at the lower fidelities (see Appendix D.3).

6 Experiments

We compare MF-GP-UCB to the following methods. **Single fidelity methods:** GP-UCB; EI: the expected improvement criterion for BO [13]; DiRect: the dividing rectangles method [12]. **Multi-fidelity methods:** MF-NAIVE: a naive baseline where we use GP-UCB to query at the *first* fidelity a large number of times and then query at the last fidelity at the points queried at $f^{(1)}$ in decreasing order of $f^{(1)}$-value; MF-SKO: the multi-fidelity sequential kriging method from [11]. Previous works on multi-fidelity methods (including MF-SKO) had not made their code available and were not straightforward to implement. Hence, we could not compare to all of them. We discuss this more in Appendix D along with some other single and multi-fidelity baselines we tried but excluded in the comparison to avoid clutter in the figures. In addition, we also detail the design choices and hyper-parameters for all methods in Appendix D.

Synthetic Examples: We use the Currin exponential ($d = 2$), Park ($d = 4$) and Borehole ($d = 8$) functions in $M = 2$ fidelity experiments and the Hartmann functions in $d = 3$ and 6 with $M = 3$ and 4 fidelities respectively. The first three are taken from previous multi-fidelity literature [32] while we tweaked the Hartmann functions to obtain the lower fidelities for the latter two cases. We show the simple regret $S(\Lambda)$ against capital Λ for the Borehole and Hartmann-3D functions in Fig. 3 with the rest deferred to Appendix D due to space constraints. MF-GP-UCB outperforms other methods. Appendix D also contains results for the cumulative regret $R(\Lambda)$ and the formulae for these functions.

A common occurrence with MF-NAIVE was that once we started querying at fidelity M, the regret barely decreased. The diagnosis in all cases was the same: it was stuck around the maximum of $f^{(1)}$ which is suboptimal for $f^{(M)}$. This suggests that while we have cheap approximations, the problem is by no means trivial. As explained previously, it is also important to "explore" at the higher fidelities to achieve good regret. The efficacy of MF-GP-UCB when compared to single fidelity methods is that it confines this exploration to a small set containing the optimum. In our experiments we found that MF-SKO did not consistently beat other single fidelity methods. Despite our best efforts to reproduce this (and another) multi-fidelity method, we found them to be quite brittle (Appendix D.1).

The third panel of Fig. 3 shows a histogram of the number of queries at each fidelity after 184 queries of MF-GP-UCB, for different ranges of $f^{(3)}(x)$ for the Hartmann-3D function. Many of the queries at the low $f^{(3)}$ values are at fidelity 1, but as we progress they decrease and the second fidelity queries increase. The third fidelity dominates very close to the optimum but is used sparingly elsewhere. This corroborates the prediction in our analysis that MF-GP-UCB uses low fidelities to explore and successively higher fidelities at promising regions to zero in on $x_\star$. (Also see Fig. 6, Appendix B.)

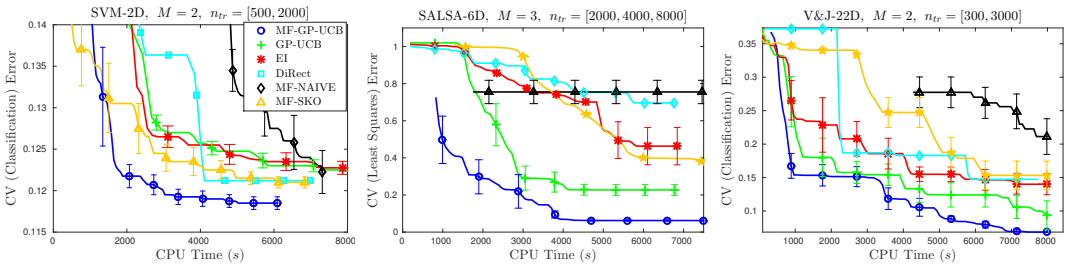

Figure 4: Results on the hyper-parameter tuning experiments. The title states the experiment, dimensionality (number of hyperparameters) and training set size at each fidelity. All curves were produced by averaging over 10 experiments. The error bars indicate one standard error. The lengths of the curves are different in time as we ran each method for a pre-specified number of iterations and they concluded at different times.

Real Experiments: We present results on three hyper-parameter tuning tasks (results in Fig. 4), and a maximum likelihood inference task in Astrophysics (Fig. 5). We compare methods on computation time since that is the "cost" in all experiments. We include the processing time for each method in the comparison (i.e. the cost of determining the next query).

Classification using SVMs (SVM): We trained an SVM on the magic gamma dataset using the SMO algorithm to an accuracy of 10^{-12}. The goal is to tune the kernel bandwidth and the soft margin coefficient in the ranges $(10^{-3}, 10^{1})$ and $(10^{-1}, 10^{5})$ respectively on a dataset of size 2000. We set this up as a $M = 2$ fidelity experiment with the entire training set at the second fidelity and 500 points at the first. Each query was 5-fold cross validation on these training sets.

Regression using Additive Kernels (SALSA): We used the regression method from [14] on the 4-dimensional coal power plant dataset. We tuned the 6 hyper-parameters –the regularisation penalty, the kernel scale and the kernel bandwidth for each dimension– each in the range $(10^{-3}, 10^{4})$ using 5-fold cross validation. This experiment used $M = 3$ and $2000, 4000, 8000$ points at each fidelity.

Viola & Jones face detection (V&J): The V&J classifier [31], which uses a cascade of weak classifiers, is a popular method for face detection. To classify an image, we pass it through each classifier. If at any point the classifier score falls below a threshold, the image is classified negative. If it passes through the cascade, then it is classified positive. One of the more popular implementations comes with OpenCV and uses a cascade of 22 weak classifiers. The threshold values in OpenCV are pre-set based on some heuristics and there is no reason to think they are optimal for a given face detection task. The goal is to tune these 22 thresholds by optimising for them over a training set. We modified the OpenCV implementation to take in the thresholds as parameters. As our domain $\mathcal{X}$ we chose a neighbourhood around the configuration used in OpenCV. We set this up as a $M = 2$ fidelity experiment where the second fidelity used 3000 images from the V&J face database and the first used 300. Interestingly, on an independent test set, the configurations found by MF-GP-UCB consistently achieved over 90% accuracy while the OpenCV configuration achieved only 87.4% accuracy.

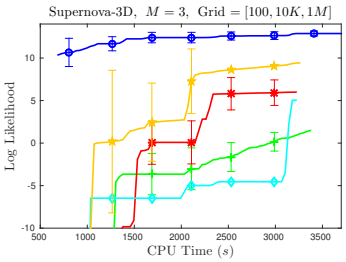

Figure 5: Results on the supernova inference problem. The y-axis is the log likelihood so higher is better. MF-NAIVE is not visible as it performed very poorly.

Type Ia Supernovae: We use Type Ia supernovae data [7] for maximum likelihood inference on 3 cosmological parameters, the Hubble constant $H_0 \in (60, 80)$, the dark matter and dark energy fractions $\Omega_M, \Omega_\Lambda \in (0, 1)$. Unlike typical parametric maximum likelihood problems, the likelihood is only available as a black-box. It is computed using the Robertson–Walker metric which requires a one dimensional numerical integration for each sample in the dataset. We set this up as a $M = 3$ fidelity task. The goal is to maximise the likelihood at the third fidelity where the integration was performed using the trapezoidal rule on a grid of size 10^6. For the first and second fidelities, we used grids of size $10^2, 10^4$ respectively. The results are given in Fig. 5.

Conclusion: We introduced and studied the multi-fidelity bandit under Gaussian Process assumptions. We present, to our knowledge, the first formalism of regret and the first theoretical results in this setting. They demonstrate that MF-GP-UCB explores the space via cheap lower fidelities, and leverages the higher fidelities on successively smaller regions hence achieving better regret than single fidelity strategies. Experimental results demonstrate the efficacy of our method.

References

[1] Alekh Agarwal, John C Duchi, Peter L Bartlett, and Clement Levrard. Oracle inequalities for computationally budgeted model selection. In *COLT*, 2011.

[2] Peter Auer. Using Confidence Bounds for Exploitation-exploration Trade-offs. *J. Mach. Learn. Res.*, 2003.

[3] E. Brochu, V. M. Cora, and N. de Freitas. A Tutorial on Bayesian Optimization of Expensive Cost Functions, with Application to Active User Modeling and Hierarchical RL. *CoRR*, 2010.

[4] Sébastien Bubeck and Nicolò Cesa-Bianchi. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends in Machine Learning*, 2012.

[5] Mark Cutler, Thomas J. Walsh, and Jonathan P. How. Reinforcement Learning with Multi-Fidelity Simulators. In *ICRA*, 2014.

[6] V. Dani, T. P. P. Hayes, and S. M Kakade. Stochastic Linear Optimization under Bandit Feedback. In *COLT*, 2008.

[7] T. M. Davis et al. Scrutinizing Exotic Cosmological Models Using ESSENCE Supernova Data Combined with Other Cosmological Probes. *Astrophysical Journal*, 2007.

[8] J Djolonga, A Krause, and V Cevher. High-Dimensional Gaussian Process Bandits. In *NIPS*, 2013.

[9] Alexander I. J. Forrester, András Sóbester, and Andy J. Keane. Multi-fidelity optimization via surrogate modelling. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Science*, 2007.

[10] Subhashis Ghosal and Anindya Roy. Posterior consistency of Gaussian process prior for nonparametric binary regression". *Annals of Statistics*, 2006.

[11] D. Huang, T.T. Allen, W.I. Notz, and R.A. Miller. Sequential kriging optimization using multiple-fidelity evaluations. *Structural and Multidisciplinary Optimization*, 2006.

[12] D. R. Jones, C. D. Perttunen, and B. E. Stuckman. Lipschitzian Optimization Without the Lipschitz Constant. *J. Optim. Theory Appl.*, 1993.

[13] Donald R. Jones, Matthias Schonlau, and William J. Welch. Efficient global optimization of expensive black-box functions. *J. of Global Optimization*, 1998.

[14] Kirthevasan Kandasamy and Yaoliang Yu. Additive Approximations in High Dimensional Nonparametric Regression via the SALSA. In *ICML*, 2016.

[15] Kirthevasan Kandasamy, Jeff Schenider, and Barnabás Póczos. High Dimensional Bayesian Optimisation and Bandits via Additive Models. In *International Conference on Machine Learning*, 2015.

[16] Kirthevasan Kandasamy, Gautam Dasarathy, Jeff Schneider, and Barnabas Poczos. The Multi-fidelity Multi-armed Bandit. In *NIPS*, 2016.

[17] K. Kawaguchi, L. P. Kaelbling, and T. Lozano-Pérez. Bayesian Optimization with Exponential Convergence. In *NIPS*, 2015.

[18] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi. Optimization by simulated annealing. *SCIENCE*, 1983.

[19] A. Klein, S. Bartels, S. Falkner, P. Hennig, and F. Hutter. Towards efficient Bayesian Optimization for Big Data. In *BayesOpt*, 2015.

[20] R. Martinez-Cantin, N. de Freitas, A. Doucet, and J. Castellanos. Active Policy Learning for Robot Planning and Exploration under Uncertainty. In *Proceedings of Robotics: Science and Systems*, 2007.

[21] Jonas Mockus. Application of Bayesian approach to numerical methods of global and stochastic optimization. *Journal of Global Optimization*, 1994.

[22] R. Munos. Optimistic Optimization of Deterministic Functions without the Knowledge of its Smoothness. In *NIPS*, 2011.

[23] D. Parkinson, P. Mukherjee, and A.. R Liddle. A Bayesian model selection analysis of WMAP3. *Physical Review*, 2006.

[24] C.E. Rasmussen and C.K.I. Williams. *Gaussian Processes for Machine Learning*. UPG Ltd, 2006.

[25] Herbert Robbins. Some aspects of the sequential design of experiments. *Bulletin of the American Mathematical Society*, 1952.

[26] A Sabharwal, H Samulowitz, and G Tesauro. Selecting near-optimal learners via incremental data allocation. In *AAAI*, 2015.

[27] J. Snoek, H. Larochelle, and R. P Adams. Practical Bayesian Optimization of Machine Learning Algorithms. In *NIPS*, 2012.

[28] Niranjan Srinivas, Andreas Krause, Sham Kakade, and Matthias Seeger. Gaussian Process Optimization in the Bandit Setting: No Regret and Experimental Design. In *ICML*, 2010.

[29] Kevin Swersky, Jasper Snoek, and Ryan P Adams. Multi-task bayesian optimization. In *NIPS*, 2013.

[30] W. R. Thompson. On the Likelihood that one Unknown Probability Exceeds Another in View of the Evidence of Two Samples. *Biometrika*, 1933.

[31] Paul A. Viola and Michael J. Jones. Rapid Object Detection using a Boosted Cascade of Simple Features. In *Computer Vision and Pattern Recognition*, 2001.

[32] Shifeng Xiong, Peter Z. G. Qian, and C. F. Jeff Wu. Sequential design and analysis of high-accuracy and low-accuracy computer codes. *Technometrics*, 2013.

[33] C. Zhang and K. Chaudhuri. Active Learning from Weak and Strong Labelers. In *NIPS*, 2015.

Appendix

A Some Ancillary Material

A.1 Table of Notations

M	The number of fidelities.
$f, f^{(m)}$	The payoff function and its m^{th} fidelity approximation. $f^{(M)} = f$.
$\lambda^{(m)}$	The cost for querying at fidelity m.
$\mathcal{X}$	The domain over which we are optimising f.
$x_\star, f_\star$	The optimum point and value of the M^{th} fidelity function.
$\overline{A}$	The complement of a set $A \subset \mathcal{X}$. $\overline{A} = \mathcal{X} \backslash A$.
$\lvert A \rvert$	The cardinality of a set $A \subset \mathcal{X}$ if it is countable.
$\vee, \wedge$	Logical *Or* and *And* respectively.
$\lesssim, \gtrsim, \asymp$	Inequalities and equality ignoring constant terms.
q_t, r_t	The instantaneous reward and regret respectively. $q_t = f^{(M)}(\mathbf{x}_t)$ if $\mathbf{m}_t = M$ and $-B$ if $\mathbf{m}_t \neq M$. $r_t = f_\star - q_t$.
$R(\Lambda)$	The cumulative regret after spending capital Λ. See equation (2).
$S(\Lambda)$	The simple regret after spending capital Λ. See second paragraph under equation (2).
$\zeta^{(m)}$	A bound on the maximum difference between $f^{(m)}$ and $f^{(M)}$, $\lVert f^{(M)} - f^{(m)} \rVert_\infty \leq \zeta^{(m)}$.
$\mu_t^{(m)}$	The mean of the m^{th} fidelity GP $f^{(m)}$ conditioned on $\mathcal{D}_t^{(m)}$ at time t.
$\kappa_t^{(m)}$	The covariance of the m^{th} fidelity GP $f^{(m)}$ conditioned on $\mathcal{D}_t^{(m)}$ at time t.
$\sigma_t^{(m)}$	The standard deviatiation of the m^{th} fidelity GP $f^{(m)}$ conditioned on $\mathcal{D}_t^{(m)}$ at time t.
$\mathbf{x}_t, \mathbf{y}_t$	The queried point and observation at time t.
$\mathbf{m}_t$	The queried fidelity at time t.
$\mathcal{D}_n^{(m)}$	The set of queries at the m^{th} fidelity until time n $\{(\mathbf{x}_t, \mathbf{y}_t)\}_{t:\mathbf{m}_t=m}$.
β_t	The coefficient trading off exploration and exploitation in the UCB. See Theorem 10.
$\varphi_t^{(m)}(x)$	The upper confidence bound (UCB) provided by the m^{th} fidelity on $f^{(M)}(x)$. $\varphi_t^{(m)}(x) = \mu_{t-1}^{(m)}(x) + \beta_t^{1/2}\sigma_{t-1}^{(m)}(x) + \zeta^{(m)}$.
$\varphi_t(x)$	The combined UCB provided by all fidelities on $f^{(M)}(x)$. $\varphi_t(x) = \min_m \varphi_t^{(m)}(x)$.
$\gamma^{(m)}$	The parameter in MF-GP-UCB for switching from the m^{th} fidelity to the $(m+1)^{\text{th}}$.
$\tilde{R}_n$	The cumulative regret for the queries after n rounds, $\tilde{R}_n = \sum_{t=1}^n \lambda^{(m_t)} r_t$.
$T_n^{(m)}(A)$	The number of queries at fidelity m in subset $A \subset \mathcal{X}$ until time n.
$T_n^{(>m)}(A)$	The number of queries at fidelities greater than m in any subset $A \subset \mathcal{X}$ until time n.
n_Λ	Number of plays by a strategy querying only at fidelity M within capital Λ. $n_\Lambda = \lfloor \Lambda / \lambda^{(M)} \rfloor$.
$\Psi_n(A)$	The maximum information gain of a set $A \subset \mathcal{X}$ after n queries in A. See Definition 2.
$\mathcal{X}^{(m)}$	$(\mathcal{X}^{(m)})_{m=1}^M$ is an entirely problem dependent partitioning of $\mathcal{X}$. See Equation (5).
$\mathcal{H}_\tau^{(m)}$	$(\mathcal{H}_\tau^{(m)})_{m=1}^M$ are partitionings of $\mathcal{X}$. See Equation (5). The analysis of MF-GP-UCB hinges on these partitionings.
$\mathcal{H}_{\tau,n}^{(m)}$	An additional n-dependent inflation of $\mathcal{H}_\tau^{(m)}$. See paragraph under equation (5).
$\widehat{\mathcal{H}}_\tau^{(m)}, \widecheck{\mathcal{H}}_\tau{}^{(m)}$	The arms "above"/"below" $\mathcal{H}_\tau^{(m)}$. $\widehat{\mathcal{H}}_\tau^{(m)} = \bigcup_{\ell=m+1}^M \mathcal{H}_\tau^{(\ell)}$, $\widecheck{\mathcal{H}}_\tau{}^{(m)} = \bigcup_{\ell=1}^{m-1} \mathcal{H}_\tau^{(\ell)}$.
$\mathcal{X}_g, \mathcal{X}_b$	The good set and bad sets for $M = 2$ fidelity problems. $\mathcal{X}_g = \mathcal{X}^{(2)}$ and $\mathcal{X}_b = \mathcal{X}^{(1)}$.
$\widetilde{\mathcal{X}}_{g,\rho}, \widetilde{\mathcal{X}}_{b,\rho}$	The inflations of $\mathcal{X}_g, \mathcal{X}_b$ for MF-GP-UCB. $\widetilde{\mathcal{X}}_{g,\rho} = \{x; f_\star - f^{(1)}(x) \leq \zeta^{(1)} + \rho\gamma\}$, and $\ddot{\mathcal{X}}_{b,\tau} = \mathcal{X} \backslash \ddot{\mathcal{X}}_{g,\tau}$.
$\Omega_\varepsilon(A)$	The ε–covering number of a subset $A \subset \mathcal{X}$ in the $\lVert \cdot \rVert_2$ metric.

A.2 Review of GP-UCB

The following bounds the regret R_n for the GP-UCB algorithm of Srinivas et al. [28] after n time steps. The algorithm is given in Algorithm 2.

Theorem 4. *(Theorems 2 in [28]) Let $f \sim \mathcal{GP}(\mathbf{0}, \kappa)$, $f : \mathcal{X} \to \mathbb{R}$ and κ satisfy Assumption 8. At each query, we have noisy observations $y = f(x) + \epsilon$ where $\epsilon \sim \mathcal{N}(0, \eta^2)$. Denote $C_1 = 8/\log(1+\eta^{-2})$. Pick $\delta \in (0,1)$. If $\mathcal{X} = [0,r]^d$, run* GP-UCB *with $\beta_t = 2\log\left(\frac{2\pi^2 t^2}{3\delta}\right) + 2d\log\left(t^2 bdr\sqrt{\frac{4ad}{\delta}}\right)$. Then,*

$$\mathbb{P}\left(\forall n \geq 1,\ R_n \leq \sqrt{C_1 n \beta_n \Psi_n(\mathcal{X})} + 2\right) \geq 1 - \delta$$

Here $\Psi_n(\mathcal{X})$ is the Maximum Information Gain of $\mathcal{X}$ after n queries (see Definition 2).

Algorithm 2 GP-UCB

Input: kernel κ.
For $t = 1, 2 \ldots$

- $\mathcal{D}_0 \leftarrow \varnothing$, $(\mu_0, \sigma_0^2) \leftarrow (\mathbf{0}, \kappa)$.
- $(\mu_0, \kappa_0) \leftarrow (\mathbf{0}, \kappa)$
- **for** $t = 1, 2, \ldots$
  1. $\mathbf{x}_t \leftarrow \text{argmax}_{x \in \mathcal{X}}\, \mu_{t-1}(x) + \beta_t^{1/2}\sigma_{t-1}(x)$
  2. $\mathbf{y}_t \leftarrow$ Query f at $\mathbf{x}_t$.
  3. $\mathcal{D}_t = \mathcal{D}_{t-1} \cup \{(\mathbf{x}_t, \mathbf{y}_t)\}$.
  4. Perform Bayesian posterior updates to obtain μ_t, σ_t (See Equation (1)).

A.3 More Related Work

Agarwal et al. [1] derive oracle inequalities for hyper-parameter tuning with ERM under computational budgets. Our setting is more general as it applies to any bandit optimisation task. Sabharwal et al. [26] present a UCB based idea for tuning hyper-parameters with incremental data allocation. However, their theoretical results are for an idealised non-realisable algorithm. Cutler et al. [5] study reinforcement learning with multi-fidelity simulators by treating each fidelity as a Markov Decision Process. Finally, Zhang and Chaudhuri [33] study active learning when there is access to a cheap weak labeler and an expensive strong labeler. All the work above study problems different to optimisation. Further, none of them are in the bandit setting where there is a price for exploration.

A.4 Some Ancillary Results

We will use the following results in our analysis. The first is a standard Gaussian concentration result and the second is an expression for the Information Gain in a GP from Srinivas et al. [28].

Lemma 5 (Gaussian Concentration)**.** *Let $Z \sim \mathcal{N}(0,1)$. Then* $\mathbb{P}(Z > \epsilon) \leq \frac{1}{2}\exp(-\epsilon^2/2)$.

Lemma 6 (Mutual Information in GP, [28] Lemma 5.3)**.** *Let $f \sim \mathcal{GP}(\mathbf{0}, \kappa)$, $f : \mathcal{X} \to \mathbb{R}$ and we observe $y = f(x) + \epsilon$ where $\epsilon \sim \mathcal{N}(0, \eta^2)$. Let A be a finite subset of $\mathcal{X}$ and f_A, y_A be the function values and observations on this set respectively. Using the basic Gaussian properties they show that the mutual information $I(y_A; f_A)$ is,*

$$I(y_A; f_A) = \frac{1}{2}\sum_{t=1}^{n} \log(1 + \eta^{-2}\sigma_{t-1}^2(x_t)).$$

where σ_{t-1}^2 is the posterior variance after observing the first $t-1$ points.

We conclude this section with the following comment on our assumptions in Section 2.

Remark 7 (Validity of the Assumptions A1, A2, A3). It is sufficient to show that when the functions $f^{(m)}$ are sampled from $\mathcal{GP}(\mathbf{0}, \kappa)$, the latter constraints, i.e. $\|f^{(M)}\|_\infty \leq B$ and $\|f^{(M)} - f^{(m)}\|_\infty \leq \zeta^{(m)}\ \forall m$, occur with positive probability. Then, a generative mechanism would repeatedly sample the $f^{(m)}$'s from the GP and output them when the constraints are satisfied. The claim is true for well behaved kernels. For instance, using Assumption 8 (Appendix C) we can establish a high probability bound on the Lipschitz constant of the GP sample $f^{(M)}$. Since for

a given $x \in \mathcal{X}$, $\mathbb{P}(-B < f^{(M)}(x) < 0)$ is positive we just need to make sure that the Lipschitz constant is not larger than $B/\text{diam}(\mathcal{X})$. This bounds $\|f^{(M)}\|_\infty < B$. For the latter constraint, since $f^{(M)} - f^{(m)} \sim \mathcal{GP}(\mathbf{0}, 2\kappa)$ is also a GP, the argument follows in an essentially similar fashion.

B Some Details on MF-GP-UCB

An Extended Simulation

In Figure 6 we provide an extended version of the simulation of Fig. 2 for a 2 fidelity example. Read the caption under the simulation for more details.

More Implementation Details

Data dependent prior: In our experiments, following recommendations in Brochu et al. [3] all GP methods were initialised with uniform random queries using an initialisation capital Λ_0. For single fidelity methods, we used it at the M^{th} fidelity, whereas for MF-GP-UCB we used $\Lambda_0/2$ at fidelity 1 and $\Lambda_0/2$ at fidelity 2. After initialising the kernel in this manner, we update the kernel every 25 iterations of the method by maximising the GP marginal likelihood.

Choice of β_t: β_t, as specified in Theorems 4, 10 has unknown constants and tends to be too conservative in practice. Following Kandasamy et al. [15] we use $\beta_t = 0.2d\log(2t)$ which captures the dominant dependencies on d and t.

Initial ζ, γ: We set both ζ, γ to 1% of the range of initial queries and update them as explained in the main text.

Maximising φ_t: To determine $\mathbf{x}_t$ we maximised φ_t using DiRect [12]. For other GP methods, the EI, PI, GP-UCB acquisition functions were also maximised using DiRect.

MF-GP-UCB was fairly robust to the above choices except when Λ_0 was set too low in which case, all GP methods performed poorly on some experiments.

C Theoretical Analysis

In this section we present our main theoretical results. While it is self contained, the reader will benefit from first reading the more intuitive discussion in Section 4. The goal in this section is to bound $R(\Lambda)$ for MF-GP-UCB . Recall,

$$\begin{aligned} R(\Lambda) = \;& \Lambda f_\star - \sum_{t=1}^{N} \lambda^{(m_t)} q_t - \left(\Lambda - \sum_{t=1}^{N} \lambda^{(m_t)}\right)(-B) \\ = \;& \underbrace{\left(\Lambda - \sum_{t=1}^{N} \lambda^{(m_t)}\right)(f_\star + B)}_{\tilde{r}(\Lambda)} + \underbrace{\sum_{t=1}^{N} \lambda^{(m_t)} r_t}_{\tilde{R}(\Lambda)}, \end{aligned}$$

where N is the random number of plays within capital Λ and q_t, r_t are the instantaneous reward and regret as defined in Section 2. The first term $\tilde{r}(\Lambda)$ is the residual quantity. It is an artefact of the fact that after the $(N+1)^{\text{th}}$ query, the spent capital would have exceeded Λ. It can be bounded by $\tilde{r}(\Lambda) \leq 2B\lambda^{(M)}$ which is typically small. Our analysis will mostly be dealing with the latter term $\tilde{R}(\Lambda)$ for which we will first bound the quantity $\tilde{R}_n = \sum_{t=1}^{n} \lambda^{(m_t)} r_t$ after n time steps in terms of n. Then, we will bound the random number of plays N within principal Λ. While $N \leq \lfloor \Lambda/\lambda^{(1)} \rfloor$ is a trivial bound, this will be too loose for our purpose. In fact, we will show that after a sufficiently large number of time steps n, with high probability the number of plays at fidelities lower than M will be sub-linear in n. Hence $N \in \mathcal{O}(n_\Lambda)$ where $n_\Lambda = \lfloor \Lambda/\lambda^{(M)} \rfloor$ is the number of plays by any algorithm that operates only at the highest fidelity.

Our strategy to bound $\tilde{R}_n$ will be to identify a (possibly disconnected) measurable region of the space $\mathcal{Z}$ which contains $x_\star$ and has high value for the payoff function $f^{(M)}(x)$. $\mathcal{Z}$ will be determined by

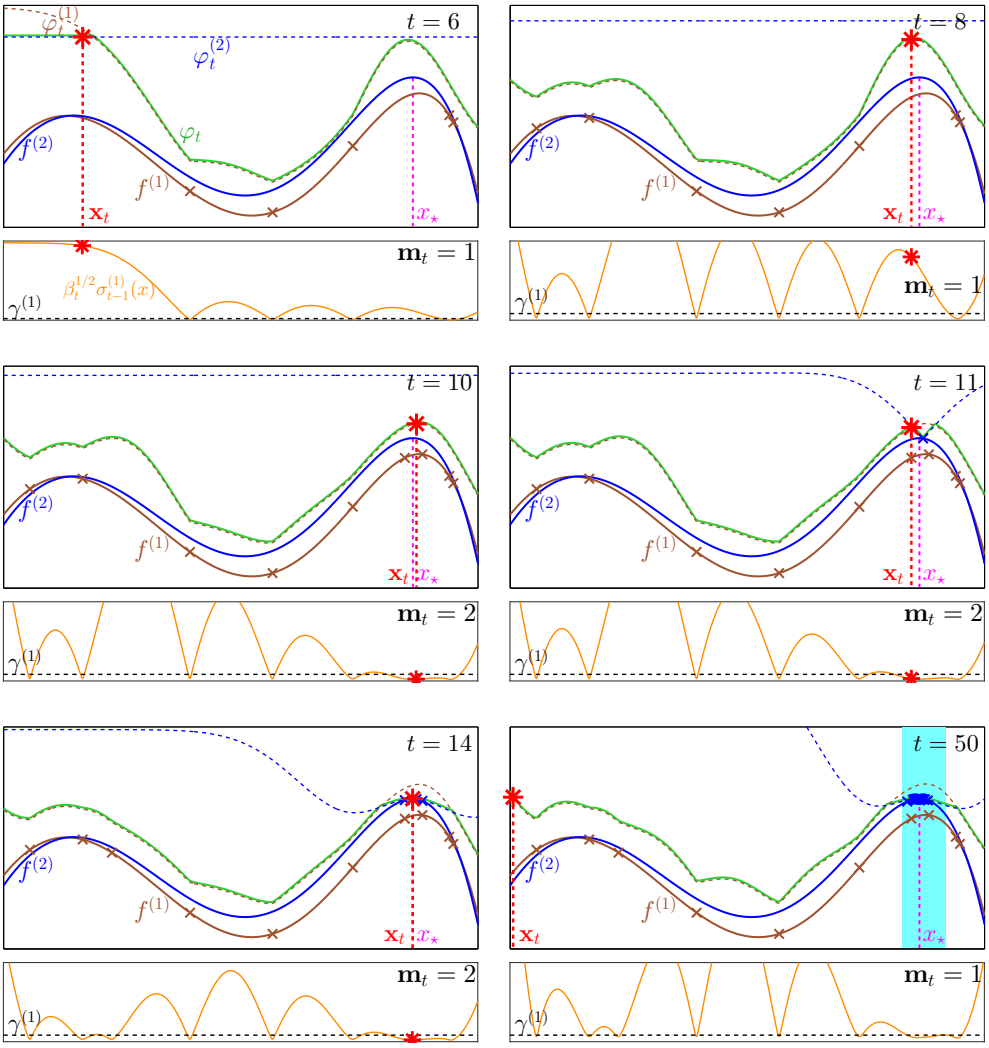

Figure 6: Illustration of MF-GP-UCB for a 2-fidelity problem initialised with 5 random points at the first fidelity. In the top figures, the solid lines in brown and blue are $f^{(1)}, f^{(2)}$ respectively, and the dashed lines are $\varphi_t^{(1)}, \varphi_t^{(2)}$. The solid green line is $\varphi_t = \min(\varphi_t^{(1)}, \varphi_t^{(2)})$. The small crosses are queries from 1 to $t-1$ and the red star is the maximiser of φ_t, i.e. the next query $\mathbf{x}_t$. $x_\star$, the optimum of $f^{(2)}$ is shown in magenta. In the bottom figures, the solid orange line is $\beta_t^{1/2}\sigma_{t-1}^{(1)}$ and the dashed black line is $\gamma^{(1)}$. When $\beta_t^{1/2}\sigma_{t-1}^{(1)}(\mathbf{x}_t) \leq \gamma^{(1)}$ we play at fidelity $\mathbf{m}_t = 2$ and otherwise at $\mathbf{m}_t = 1$.

At the initial stages, MF-GP-UCB is mostly exploring $\mathcal{X}$ in the first fidelity. $\beta_t^{1/2}\sigma_{t-1}^{(1)}$ is large and we are yet to constrain $f^{(1)}$ well to proceed to $m = 2$. At $t = 10$, we have constrainted $f^{(1)}$ sufficiently well at a region around the optimum. $\beta_t^{1/2}\sigma_{t-1}^{(1)}(\mathbf{x}_t)$ falls below $\gamma^{(1)}$ and we query at $\mathbf{m}_t = 2$. Notice that once we do this (at $t = 11$), $\varphi_t^{(2)}$ dips to change φ_t in that region. At $t = 14$, MF-GP-UCB has identified the maximum $x_\star$ with just 4 queries to $f^{(2)}$. In the last figure, at $t = 50$, the algorithm decides to explore at a point far away from the optimum. However, this query occurs in the first fidelity since we have not sufficiently constrained $f^{(1)}(\mathbf{x}_t)$ in this region. The key idea is that it is *not necessary* to query such regions at the second fidelity as the first fidelity alone is enough to conclude that it is suboptimal. Herein lies the crux of our method. The region shaded in cyan in the last figure is the good set $\mathcal{X}_g = \{x; f^{(2)}(x_\star) - f^{(1)}(x) \leq \zeta^{(1)}\}$ discussed in Section 4. Our analysis predicts that most second fidelity queries in MF-GP-UCB will be confined to this set with high probability and the simulation corroborates this claim. In addition, observe that in a large portion of $\mathcal{X}$, φ_t is given by $\varphi_t^{(1)}$ except in a small neighborhood around $x_\star$, where it is given by $\varphi_t^{(2)}$.

the approximations provided via the lower fidelity evaluations. Denoting $\overline{\mathcal{Z}} = \mathcal{X} \backslash \mathcal{Z}$, we decompose $\tilde{R}_n$ as follows,

$$\tilde{R}_n \leq \underbrace{2B \sum_{m=1}^{M-1} \lambda^{(m)} T_n^{(m)}(\mathcal{X})}_{\tilde{R}_{n,1}} + \underbrace{\lambda^{(M)} \sum_{\substack{t:\mathbf{m}_t=M \\ \mathbf{x}_t \in \mathcal{Z}}} \left(f_\star - f^{(M)}(\mathbf{x}_t)\right)}_{\tilde{R}_{n,2}} + \underbrace{\lambda^{(M)} \sum_{\substack{t:\mathbf{m}_t=M \\ \mathbf{x}_t \in \overline{\mathcal{Z}}}} \left(f_\star - f^{(M)}(\mathbf{x}_t)\right)}_{\tilde{R}_{n,3}}. \tag{4}$$

$\tilde{R}_{n,1}$ is the capital spent on the lower fidelity queries for which we receive no reward. $\tilde{R}_{n,2}$ is the regret due to fidelity M queries in $\mathcal{Z}$ and $\tilde{R}_{n,3}$ is due to fidelity M queries outside $\mathcal{Z}$. To control $\tilde{R}_{n,1}$ we will first bound $T_n^{(m)}(\mathcal{X})$ for $m < M$. This will typically be small containing only $\text{polylog}(n)/\text{poly}(\gamma)$ and $o(n)$ terms. The last two terms can be controlled using the MIGs Ψ_n of $\mathcal{Z}, \overline{\mathcal{Z}}$ respectively (Definition 2). As we will see, $\tilde{R}_{n,2}$ will be the dominant term in n in our final expression since most of the fidelity M queries will be confined to $\mathcal{Z}$. $T_M^{(n)}(\overline{\mathcal{Z}})$ will be sublinear in n and hence $\tilde{R}_{n,3}$ will be of low order. When the lower fidelities allow us to eliminate a large region of the space, $\text{vol}(\mathcal{Z}) \ll \text{vol}(\overline{\mathcal{Z}})$ and consequently the maximum information gain of $\mathcal{Z}$ will be much smaller than that of $\overline{\mathcal{Z}}$, $\Psi_n(\mathcal{Z}) \ll \Psi_n(\overline{\mathcal{Z}})$. As we will see, this results in much better regret for MF-GP-UCB in comparison to GP-UCB.

For the analysis, we will need the following regularity conditions on the kernel. It is satisfied for four times differentiable kernels such as the SE and Matérn kernels with smoothness parameter $\nu > 2$ [10].

Assumption 8. *Let $f \sim \mathcal{GP}(\mathbf{0}, \kappa)$, where $\kappa : [0, r]^d \times [0, r]^d \to \mathbb{R}$ is a stationary kernel. The partial derivatives of f satisfies the following high probability bound. There exists constants $a, b > 0$ such that, for all $J > 0$,*

$$\forall\, i \in \{1, \ldots, d\}, \quad \mathbb{P}\left(\sup_x \left|\frac{\partial f(x)}{\partial x_i}\right| > J\right) \leq a e^{-(J/b)^2}.$$

For our proofs we will need to control the conditional variances for queries within a subset $A \subset \mathcal{X}$. To that end, we provide the lemma below.

Lemma 9. *Let $f \sim \mathcal{GP}(0, \kappa)$, $f : \mathcal{X} \to \mathbb{R}$ and each time we query at any $x \in \mathcal{X}$ we observe $y = f(x) + \epsilon$, where $\epsilon \sim \mathcal{N}(0, \eta^2)$. Let $A \subset \mathcal{X}$. Assume that we have queried f at n points, $(x_t)_{t=1}^n$ of which s points are in A. Let σ_{t-1}^2 denote the posterior variance at time t, i.e. after $t-1$ queries. Then, $\sum_{x_t \in A} \sigma_{t-1}^2(x_t) \leq \frac{2}{\log(1+\eta^{-2})} \Psi_s(A)$.*

Proof Let $A_s = \{z_1, z_2, \ldots, z_s\}$ be the queries inside A in the order they were queried. Now, assuming that we have only queried inside A at A_s, denote by $\tilde{\sigma}_{t-1}(\cdot)$, the posterior standard deviation after $t-1$ such queries. Then,

$$\begin{aligned}\sum_{t:x_t \in A} \sigma_{t-1}^2(x_t) &\leq \sum_{t=1}^s \tilde{\sigma}_{t-1}^2(z_t) \leq \sum_{t=1}^s \eta^2 \frac{\tilde{\sigma}_{t-1}^2(z_t)}{\eta^2} \leq \sum_{t=1}^s \frac{\log(1+\eta^{-2}\tilde{\sigma}_{t-1}^2(z_t))}{\log(1+\eta^{-2})} \\ &\leq \frac{2}{\log(1+\eta^{-2})} I(y_{A_s}; f_{A_s})\end{aligned}$$

Queries outside A will only decrease the variance of the GP so we can upper bound the first sum by the posterior variances of the GP with only the queries in A. The third step uses the inequality $u^2/v^2 \leq \log(1+u^2)/\log(1+v^2)$ with $u = \tilde{\sigma}_{t-1}(z_t)/\eta$ and $v = 1/\eta$ and the last step uses Lemma 6. The result follows from the fact that $\Psi_s(A)$ maximises the mutual information among all subsets of size s. ■

We now proceed to the analysis. To avoid clutter in the notation we will use $\gamma = \gamma^{(m)}$ for all m. Generalising this to different $\gamma^{(m)}$'s is straightforward.

Denote $\Delta^{(m)}(x) = f_\star - f^{(m)}(x) - \zeta^{(m)}$ and $\mathcal{J}_\eta^{(m)} = \{x \in \mathcal{X}; \Delta^{(m)}(x) \leq \eta\}$. Let $\tau > 0, \rho > 1$ be given. Central to our analysis will be two partitionings $(\mathcal{X}^{(m)})_{m=1}^M$ and $(\mathcal{H}_\tau^{(m)})_{m=1}^M$ of $\mathcal{X}$. The latter depends on the parameter γ and the given τ, ρ. Let $\mathcal{X}^{(1)} = \overline{\mathcal{J}}_0^{(1)}$, $\mathcal{H}_\tau^{(1)} = \overline{\mathcal{J}}_{\max(\tau,\rho\gamma)}^{(1)}$. Then define,

$$\mathcal{X}^{(m)} = \overline{\mathcal{J}}_0^{(m)} \cap \left(\bigcap_{\ell=1}^{m-1} \mathcal{J}_0^{(\ell)} \right) \quad \text{for } 2 \leq m \leq M-1, \qquad \mathcal{X}^{(M)} = \bigcap_{\ell=1}^{M-1} \mathcal{J}_0^{(\ell)}. \tag{5}$$

$$\mathcal{H}_\tau^{(m)} = \overline{\mathcal{J}}_{\max(\tau,\rho\gamma)}^{(m)} \cap \left(\bigcap_{\ell=1}^{m-1} \mathcal{J}_{\max(\tau,\rho\gamma)}^{(\ell)} \right) \quad \text{for } 2 \leq m \leq M-1, \qquad \mathcal{H}_\tau^{(M)} = \bigcap_{\ell=1}^{M-1} \mathcal{J}_{\max(\tau,\rho\gamma)}^{(\ell)}.$$

In addition to the above, we will also find it useful to define the sets "above" $\mathcal{H}_\tau^{(m)}$ as $\widehat{\mathcal{H}}_\tau^{(m)} = \bigcup_{\ell=m+1}^M \mathcal{H}_\tau^{(\ell)}$ and the sets "below" $\mathcal{H}_\tau^{(m)}$ as $\breve{\mathcal{H}}_\tau{}^{(m)} = \bigcup_{\ell=1}^{m-1} \mathcal{H}_\tau^{(\ell)}$. Intuitively, $\mathcal{H}_\tau^{(m)}$ is the set of points that MF-GP-UCB will query at the m^{th} fidelity but exclude from higher fidelities due to information from fidelity m. $\breve{\mathcal{H}}_\tau{}^{(m)}$ is the set of points that can be excluded from queries at fidelities m and beyond due to information from lower fidelities. $\widehat{\mathcal{H}}_\tau^{(m)}$ are points that need to be queried at fidelities higher than m. In the 2 fidelity setting described in Section 4, the set $\mathcal{X}_g$ is $\mathcal{X}^{(2)}$ and $\breve{\mathcal{X}}_{g,\rho}$ is $\mathcal{H}_\tau^{(2)}$. Finally, for any given $\alpha > 0$ we will also define $\mathcal{H}_{\tau,n}^{(m)} = \{x \in \mathcal{X} : \mathrm{B}_2(x, r\sqrt{d}/n^{\frac{\alpha}{2d}}) \cap \mathcal{H}_\tau^{(m)} \neq \varnothing \wedge x \notin \widehat{\mathcal{H}}^{(m)}\}$ to be an n-dependence inflation of $\mathcal{H}_{\tau,n}^{(m)}$. Here, $\mathrm{B}_2(x, \epsilon)$ is an L_2 ball of radius ϵ centred at x. The sets $\{\mathcal{H}_{\tau,n}^{(m)}\}_{m=1}^M$ depend on ρ, γ, τ, n and α. Notice that for any $\alpha > 0$, as $n \to \infty$, $\mathcal{H}_{\tau,n}^{(m)} \to \mathcal{H}_\tau^{(m)}$. In addition to the above, denote the ε covering number of a set $A \subset \mathcal{X}$ in the $\|\cdot\|_2$ metric by $\Omega_\varepsilon(A)$. Let $T_n^{(m)}(A)$ denote the number of queries in a subset $A \subset \mathcal{X}$ at fidelity m. $\mathcal{D}_n^{(m)} = \{(\mathbf{x}_t, \mathbf{y}_t)\}_{t:\mathbf{m}_t=m}$ denotes the set of query-value pairs at the m^{th} fidelity until time n. Our main theorem is as follows.

Theorem 10. *Let $\mathcal{X} \subset [0,r]^d$ be compact and convex. Let $f^{(m)} \sim \mathcal{GP}(\mathbf{0}, \kappa)$ $\forall\, m$, and satisfy assumptions **A2**, **A3**. Let κ satisfy Assumption 8 with some constants a, b. Pick $\delta \in (0,1)$ and run* MF-GP-UCB *with*

$$\beta_t = 2\log\left(\frac{M\pi^2 t^2}{2\delta}\right) + 4d\log(t) + \max\left\{0\,,\, 2d\log\left(brd\log\left(\frac{6Mad}{\delta}\right)\right)\right\}.$$

For all $\alpha \in (0,1)$, $\tau > 0$, $\rho > \rho_0 = \max\{2, 1+\sqrt{(1+2/\alpha)/(1+d)}\}$ and sufficiently large Λ, we have $R(\Lambda) \in \mathcal{O}\left(\sum_{m=1}^M \lambda^{(m)} \sqrt{n_\Lambda \beta_{n_\Lambda} \Psi_{n_\Lambda}(\mathcal{H}_{\tau,n_\Lambda}^{(m)})} + \frac{\mathrm{diam}(\widehat{\mathcal{H}}_\tau^{(m)})^d \mathrm{polylog}(n_\Lambda)}{\mathrm{poly}(\gamma)}\right)$. Here, $n_\Lambda = \lfloor \Lambda/\lambda^{(M)} \rfloor$ as before.
Precisely, there exists Λ_0 such that for all $\Lambda \geq \Lambda_0$, with probability $> 1-\delta$ we have,

$$\begin{aligned} R(\Lambda) \;\leq\; 2B\lambda^{(M)} \;+\; \lambda^{(M)} \left[\sqrt{2C_1 M n_\Lambda^\alpha \Psi_{2Mn_\Lambda^\alpha}(\breve{\mathcal{H}}^{(M)})} \;+\; \sqrt{2C_1 n_\Lambda \Psi_{2n_\Lambda}(\mathcal{H}_{\tau,n_\Lambda}^{(M)})} \;+\; \frac{\pi^2}{6} \right] \\ +\; 2B \sum_{m=1}^{M-1} \lambda^{(m)} \Bigg[(m-1)(2n_\Lambda^\alpha) \;+\; \frac{1}{\tau}\left(\sqrt{2C_1 n_\Lambda \beta_{2n_\Lambda} \Psi_{2n_\Lambda}(\mathcal{H}_{\tau,n_\Lambda}^{(m)})} + \frac{\pi^2}{6} \right) + \\ \Omega_{\varepsilon_n}(\widehat{\mathcal{H}}_\tau^{(m)}) \left(\frac{2\eta^2}{\gamma^2}\beta_n + 1 \right) \Bigg], \end{aligned}$$

where $C_1 = 8/\log(1+\eta^2)$. For the SE kernel $\varepsilon_n = \frac{\gamma}{\sqrt{8C_{SE}\beta_n}}$, and therefore $\Omega_{\varepsilon_n}(\widehat{\mathcal{H}}^{(m)}) \in \mathcal{O}\left(\frac{\mathrm{diam}(\widehat{\mathcal{H}}^{(m)})^d (\log(n))^{d/2}}{\gamma^d}\right)$. For the Matérn kernel $\varepsilon_n = \frac{\gamma^2}{8C_{Mat}\beta_n}$ and therefore $\Omega_{\varepsilon_n}(\widehat{\mathcal{H}}^{(m)}) \in \mathcal{O}\left(\frac{\mathrm{diam}(\widehat{\mathcal{H}}^{(m)})^d (\log(n))^d}{\gamma^{2d}}\right)$. C_{SE}, C_{Mat} are kernel dependent constants. As $\Lambda \to \infty$, $n_\Lambda \to \infty$ and hence $\mathcal{H}_{\tau,n_\Lambda}^{(m)} \to \mathcal{H}_\tau^{(m)}$ for all $m \in \{1,\ldots,M\}$ and $\alpha \in (0,1)$.

Synopsis: Ignoring the common terms, constants and n_Λ^α terms, the regret for GP-UCB is $\lambda^{(M)}\sqrt{n_\Lambda \Psi_{n_\Lambda}(\mathcal{X})}$ whereas for MF-GP-UCB it is $\sum_m \lambda^{(m)} \sqrt{n_\Lambda \Psi_{n_\Lambda}(\mathcal{H}_{\tau,n}^{(m)})}$. In problems where

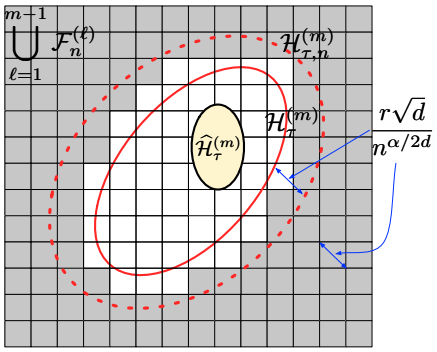

Figure 7: Illustration of the sets $\{\mathcal{F}_n^{(\ell)}\}_{\ell=1}^{m-1}$ with respect to $\mathcal{H}_\tau^{(m)}$. The grid represents a $r\sqrt{d}/n^{\alpha/2d}$ covering of $\mathcal{X}$. The yellow region is $\widehat{\mathcal{H}}_\tau^{(m)}$. The area enclosed by the solid red line (excluding $\widehat{\mathcal{H}}_\tau^{(m)}$) is $\mathcal{H}_\tau^{(m)}$. $\mathcal{H}_{\tau,n}^{(m)}$, shown by a dashed red line, is obtained by inflating $\mathcal{H}_\tau^{(m)}$ by $r\sqrt{d}/n^{\alpha/2d}$. The grey shaded region represents $\bigcup_{\ell=1}^{m-1}\mathcal{F}_n^{(\ell)}$. By our definition, $\bigcup_{\ell=1}^{m-1}\mathcal{F}_n^{(\ell)}$ contains the cells which are entirely outside $\mathcal{H}_\tau^{(m)}$. However, the inflation $\mathcal{H}_{\tau,n}^{(m)}$ is such that $\widehat{\mathcal{H}}_\tau^{(m)} \cup \mathcal{H}_{\tau,n}^{(m)} \cup \bigcup_{\ell=1}^{m-1}\mathcal{F}_n^{(\ell)} = \mathcal{X}$. As $n \to \infty$, $\mathcal{H}_{\tau,n}^{(m)} \to \mathcal{H}_\tau^{(m)}$.

$\mathrm{vol}(\mathcal{H}_{\tau,n}^{(m)}) \ll \mathrm{vol}(\mathcal{H}_{\tau,n}^{(m)})$, and $\lambda^{(m)} \ll \lambda^{(m+1)}$ MF-GP-UCB achieves signficantly better regret than GP-UCB. When the sets become larger (the approximation becomes worse) and the costs become comparable the bound decays gracefully. The $\lambda^{(m)}\sqrt{n_\Lambda^\alpha \Psi_{n_\Lambda^\alpha}(\mathcal{H}_{\tau,n}^{(m)})}$ terms can be made arbitrarily small by picking large enough ρ, provided $\mathcal{H}_{\tau,n}^{(m)}$ is still small relative to $\mathcal{X}$. On the other hand the $\mathrm{diam}(\widehat{\mathcal{H}}_\tau^{(m)})\mathrm{polylog}(n_\Lambda)/\mathrm{poly}(\gamma)$ terms could be big if γ is too small. MF-GP-UCB requires that γ will be chosen large enough so that the above term remains small relative to $\sqrt{n_\Lambda \beta_{n_\Lambda} \Psi_{n_\Lambda}(\mathcal{H}_\tau^{(m)})}$ which is not too restrictive since we expect $\widehat{\mathcal{H}}_\tau^{(m)}$ to be much smaller than $\mathcal{H}_\tau^{(m)}$. Our analysis reveals that an optimal choice for the SE kernel scales $\gamma^{(m)} \asymp (\lambda^{(m)}\zeta^{(m)}/(t\lambda^{(m+1)}))^{1/(d+2)}$ at time step t. However this observation is of little practical consequence as the leading constant depends on several problem dependent quantities such as $\Psi_n(\mathcal{X}_g)$. Our heuristics for setting γ seemed to work well in practice (see Section 5).

Proof of Theorem 10. We will study MF-GP-UCB after n time steps regardless of the queried fidelities and bound $\tilde{R}_n$. Then we will bound the number of plays N within capital Λ. For the analysis, at time n we will consider a $\frac{r\sqrt{d}}{2n^{\frac{\alpha}{2d}}}$-covering of the space $\mathcal{X}$ of size $n^{\frac{\alpha}{2}}$. For instance, if $\mathcal{X} = [0, r]^d$ a sufficient discretisation would be an equally spaced grid having $n^{\alpha/2d}$ points per side. Let $\{a_{i,n}\}_{i=1}^{n^{\frac{\alpha}{2}}}$ be the points in the covering, $F_n = \{A_{i,n}\}_{i=1}^{n^{\frac{\alpha}{2}}}$ be the cells in the covering, i.e. $A_{i,n}$ is the set of points which are closest to $a_{i,n}$ in the covering. Next we define another partitioning of the space similar in spirit to (5) using this partitioning. First let $F_n^{(1)} = \{A_{i,n} \in F_n : \ A_{i,n} \subset \mathcal{J}^{(1)}_{\max(\tau,\rho\gamma)}\}$. Next,

$$F_n^{(m)} = \left\{ A_{i,n} \in F_n : \ A_{i,n} \subset \overline{\mathcal{J}}^{(m)}_{\max(\tau,\rho\gamma)} \quad \wedge \quad A_{i,n} \notin \bigcup_{\ell=1}^{m-1} F_n^{(\ell)} \right\} \quad \text{for } \ 2 \le m \le M-1. \tag{6}$$

Note that $F_n^{(m)} \subset F_n$. We define the following *disjoint* subsets $\{\mathcal{F}_n^{(m)}\}_{m=1}^{M-1}$ of $\mathcal{X}$ via $\mathcal{F}_n^{(m)} = \bigcup_{A_{i,n}\in F_n^{(m)}} A_{i,n}$. We have illustrated $\bigcup_{\ell=1}^{m-1}\mathcal{F}_n^{(\ell)}$ with respect to $\mathcal{H}_\tau^{(m)}$ in Figure 7. By noting that $\mathcal{H}_{\tau,n}^{(1)} = \mathcal{H}^{(1)}$ we make the following observation,

$$T_n^{(m)}(\mathcal{X}) \ \le \ \sum_{\ell=1}^{m-1} T_n^{(m)}(\mathcal{F}_n^{(\ell)}) \ + \ T_n^{(m)}(\mathcal{H}_{\tau,n}^{(m)}) \ + \ T_n^{(m)}(\widehat{\mathcal{H}}^{(m)}). \tag{7}$$

This follows by noting that $\overline{\mathcal{H}_{\tau,n}^{(m)} \cup \widehat{\mathcal{H}}^{(m)}} \subset \bigcup_{\ell=1}^{m-1}\mathcal{F}_n^{(\ell)}$ (See Fig. 7). To control $\tilde{R}_n$ we will bound control each of these terms individually. First we focus on $\widehat{\mathcal{H}}^{(m)}$ for which we use the following lemma. The proof is given in Section C.0.1.

Lemma 11. *Let $f \sim \mathcal{GP}(\mathbf{0}, \kappa)$, $f : \mathcal{X} \to \mathbb{R}$ and we observe $y = f(x) + \epsilon$ where $\epsilon \sim \mathcal{N}(0, \eta^2)$. Let $A \subset \mathcal{X}$ such that its L_2 diameter $\mathrm{diam}(A) \le D$. Say we have n queries $(\mathbf{x}_t)_{t=1}^n$ of which s points*

are in A. Then the posterior variance of the GP, $\kappa'(x,x)$ at any $x \in A$ satisfies

$$\kappa'(x,x) \leq \begin{cases} C_{SE}D^2 + \frac{\eta^2}{s} & \text{if } \kappa \text{ is the SE kernel,} \\ C_{Mat}D + \frac{\eta^2}{s} & \text{if } \kappa \text{ is the Matérn kernel,} \end{cases}$$

for appropriate constants C_{SE}, C_{Mat}.

First consider the SE kernel. At time t consider any $\varepsilon_n = \frac{\gamma}{\sqrt{8C_{SE}\beta_n}}$ covering $(B_i)_{i=1}^{\varepsilon_n}$ of $\widehat{\mathcal{H}}^{(m)}$. The number of queries inside any B_i of this covering at time n will be at most $\frac{2\eta^2}{\gamma^2}\beta_n + 1$. To see this, assume we have already queried $2\eta^2/\gamma^2 + 1$ times inside B_i at time $t \leq n$. By Lemma 11 the maximum variance in A_i can be bounded by

$$\max_{x \in A_i} \kappa_{t-1}^{(m)}(x,x) \leq C_{SE}(2\varepsilon_n)^2 \ + \ \frac{\eta^2}{T_t^{(m)}(A_i)} < \frac{\gamma^2}{\beta_n}.$$

Therefore, $\beta_t^{1/2}\sigma_{t-1}^{(m)}(x) \leq \beta_n^{1/2}\sigma_{t-1}^{(m)}(x) < \gamma$ and we will not query inside A_i until time n. Therefore, the number of m^{th} fidelity queries is bounded by $\Omega_{\varepsilon_n}(\widehat{\mathcal{H}}^{(m)})\left(\frac{2\eta^2}{\gamma^2}\beta_n + 1\right)$. The proof for the Matérn kernel follows similarly using $\varepsilon_n = \frac{\gamma^2}{8C_{Mat}\beta_n}$. Next, we bound $T_n^{(m)}(\mathcal{H}_{\tau,n}^{(m)})$ for which we will use the following Lemma. The proof is given in Section C.0.2.

Lemma 12. *For β_t as given in Theorem 10, we have the following with probability $> 1 - 5\delta/6$.*

$$\forall\, m \in \{1,\ldots,M\}, \ \ \forall\, t \geq 1, \qquad \Delta^{(m)}(\mathbf{x}_t) = f_\star - f^{(m)}(\mathbf{x}_t) \ \leq \ 2\beta_t\sigma_{t-1}^{(m)}(\mathbf{x}_t) \ + \ 1/t^2.$$

First, we will analyse the quantity $\tilde{R}_n^{(m)} = \sum_{\substack{t:\mathbf{m}_t = m \\ \mathbf{x}_t \in \mathcal{H}_{\tau,n}^{(m)}}} \Delta^{(m)}(\mathbf{x}_t)$ for $m < M$. Lemma 12 gives us $\tilde{R}_n^{(m)} \leq 2\beta_n^{1/2}\sum \sigma_{t-1}^{(m)}(\mathbf{x}_t) + \pi^2/6$. Then, using Lemma 9 and Jensen's inequality we have,

$$\left(\tilde{R}_n^{(m)} - \frac{\pi^2}{6}\right)^2 \ \leq \ 4\beta_t\, T_n^{(m)}(\mathcal{H}_{\tau,n}^{(m)}) \sum_{\substack{t:\mathbf{m}_t = m \\ \mathbf{x}_t \in \mathcal{H}_{\tau,n}^{(m)}}} \left(\sigma_{t-1}^{(m)}\right)^2(\mathbf{x}_t) \ \leq \ C_1\beta_t\, T_n^{(m)}(\mathcal{H}_{\tau,n}^{(m)})\Psi_{T_n^{(m)}(\mathcal{H}_{\tau,n}^{(m)})}(\mathcal{H}_{\tau,n}^{(m)}). \tag{8}$$

We therefore have, $\tilde{R}_n^{(m)} \leq \sqrt{C_1 n\beta_n\Psi_n(\mathcal{H}_{\tau,n}^{(m)})} + \pi^2/6$ since trivially $T_n^{(m)}(\mathcal{H}_{\tau,n}^{(m)}) < n$. However, since $\Delta^{(m)}(x) > \tau$ for $x \in \mathcal{H}_{\tau,n}^{(m)}$ we have $T_n^{(m)}(\mathcal{H}_{\tau,n}^{(m)}) < \frac{1}{\tau}\left(\sqrt{C_1 n\beta_n\Psi_n(\mathcal{H}_{\tau,n}^{(m)})} + \pi^2/6\right)$.

Remark 13. *Since $\Psi_n(\cdot)$ is typically sublinear in n, it is natural to ask if we can recursively apply this to obtain a tighter bound on $T_n^{(m)}(\mathcal{H}_{\tau,n}^{(m)})$. For instance, since $\Psi_n(\cdot)$ is $\text{polylog}(n)$ for the SE kernel (Srinivas et al. [28], Theorem 5) by repeating the argument above once we get, $T_n^{(m)}(\mathcal{H}_{\tau,n}^{(m)}) \in \mathcal{O}\left(\frac{1}{\tau^{3/2}}\sqrt{C_1 n^{1/2}\text{polylog}(n)\beta_n\Psi_{\tau^{-3/2}n^{1/2}\text{polylog}(n)}(\mathcal{H}_{\tau,n}^{(m)})}\right)$. However, while this improves the dependence on n it worsens the dependence on τ. In fact, using a discretisation argument similar to that in Lemma 14 and the variance bound in Lemma 11, a $\text{polylog}(n)/\text{poly}(\tau)$ bound can be shown, with the $\text{poly}(\tau)$ term being τ^{d+2} for the SE kernel and τ^{2d+2} for the Matérn kernel. In fact, the same argument can be applied to GP-UCB to show that the number of plays on a τ-suboptimal set is $\text{polylog}(n)/\text{poly}(\tau)$. If we are to avoid this $1/\text{poly}(\tau)$ dependence for GP-UCB the best you can achieve for GP-UCB is a $\mathcal{O}(n^{1/2})$ rate for the SE kernel and $\mathcal{O}(n^{\frac{1}{2}+\frac{d(d+1)}{2\nu+d(d+1)}})$ for the Matérn kernel.*

Finally, to control the first term in (7), we will bound $T_n^{(>m)}(\mathcal{F}_n^{(m)})$. To that end we provide the following Lemma. The proof is given in Section C.0.3.

Lemma 14. *Consider any $A_{i,n} \in F_n^{(m)}$ where $F_n^{(m)}$ is as defined in (6) for any $\alpha \in (0,1)$. Let ρ, β_t be as given in Theorem 10, Then for all $u \geq \max\{3,\ (2(\rho - \rho_0)\eta)^{-2/3}\}$ we have,*

$$\mathbb{P}(T_n^{(>m)}(A_{i,n}) > u) \ \leq \ \frac{\delta}{\pi^2} \cdot \frac{1}{u^{1+4/\alpha}}$$

We will use the above result with $u = n^{\alpha/2}$. Applying the union bound we have,

$$\mathbb{P}\left(\forall\, m \in \{1,\ldots,M\},\ T_n^{(>m)}(\mathcal{F}_n^{(m)}) > |F_n^{(m)}|n^{\alpha/2}\right) \ \leq\ \sum_{m=1}^{M} \mathbb{P}\left(T_n^{(>m)}(\mathcal{F}_n^{(m)}) > |F_n^{(m)}|n^{\alpha/2}\right)$$
$$\leq\ \sum_{m=1}^{M} \sum_{A_{i,n}\in F_n^{(m)}} \mathbb{P}\left(T_n^{(>m)}(A_{i,n}) > n^{\alpha/2}\right) \ \leq\ \sum_{m=1}^{M} |F_n^{(m)}|\frac{\delta}{\pi^2}\frac{1}{n^{2+\alpha/2}} \ \leq\ |F_n|\frac{\delta}{\pi^2}\frac{1}{n^{2+\alpha/2}} \ =\ \frac{\delta}{\pi^2}\frac{1}{n^2}$$

Applying the union bound once again, we have $T_n^{(>m)}(\mathcal{F}_n^{(m)}) \leq n^\alpha$ for all m and all $n \geq \max\{3,\ (2(\rho-\rho_0)\eta)^{2/3}\}^{2/\alpha}$ with probability $> 1-\delta/6$. Henceforth, all statements we make will make use of the results in Lemmas 11, 12 and 14 and will hold with probability $> 1-\delta$.

First using equation (7) and noting $T_n^{(m)}(\mathcal{F}_n^{(\ell)}) \leq T_n^{(>\ell)}(\mathcal{F}_n^{(\ell)})$ for $\ell < m$ we bound $T_n^{(m)}(\mathcal{X})$ for $m < M$.

$$T_n^{(m)}(\mathcal{X}) \ \leq\ (m-1)n^\alpha \ +\ \frac{1}{\tau}\left(\sqrt{C_1 n\beta_n \Psi_n(\mathcal{H}_{\tau,n}^{(m)})} + \frac{\pi^2}{6}\right) \ +\ \Omega_{\varepsilon_n}(\widehat{\mathcal{H}}^{(m)})\left(\frac{2\eta^2}{\gamma^2}\beta_n + 1\right).$$

Using this bound we can control $\tilde{R}_{n,1}$ in (4). To bound $\tilde{R}_{n,2}$ and $\tilde{R}_{n,3}$ we set $\mathcal{Z} = \mathcal{H}_{\tau,n_\Lambda}^{(M)}$ and use Lemma 12 noting that when $\mathbf{m}_t = M$, $r_t = \Delta^{(M)}(\mathbf{x}_t)$. Using similar calculations to (8) and as $T_n^{(M)}(\mathcal{H}_{\tau,n}^{(m)}) \leq n$, we have $\tilde{R}_{n,2} \leq \sqrt{C_1 n\beta_n \Psi_n(\mathcal{H}_{\tau,n}^{(m)})} + \sum_{\mathbf{x}_t\in\mathcal{Z}} 1/t^2$. Next, using Lemma 14 and observing $\overline{\mathcal{Z}} = \overline{\mathcal{H}_{\tau,n}^{(M)}} \subset \bigcup_{\ell=1}^{M-1} \mathcal{F}_n^{(m)} \subset \breve{\mathcal{H}}^{(M)}$, we have,

$$\tilde{R}_{n,3} \ =\ \sum_{\substack{t:\mathbf{m}_t=M\\ \mathbf{x}_t\in\overline{\mathcal{Z}}}} \left(f_\star - f^{(M)}(\mathbf{x}_t)\right) \ \leq\ \sum_{\substack{t:\mathbf{m}_t=M\\ \mathbf{x}_t\in\bigcup_{\ell=1}^{M-1}\mathcal{F}_n^{(m)}}} 2\beta_t^{1/2}\sigma_{t-1}^{(m)}(\mathbf{x}_t) \ +\ \sum_{\mathbf{x}_t\in\overline{\mathcal{Z}}}\frac{1}{t^2}$$
$$\leq\ \sqrt{C_1 M n^\alpha \beta_n \Psi_{Mn^\alpha}(\breve{\mathcal{H}}^{(M)})} \ +\ \sum_{\mathbf{x}_t\in\overline{\mathcal{Z}}}\frac{1}{t^2}.$$

Plugging these bounds back into (4), we obtain a bound on the regret similar to the one given in the theorem except with n replaced by $2n_\Lambda$. The last step in the proof will be to show that for sufficiently large Λ, $N \leq 2n_\Lambda$ which will complete the proof. For this we turn back to our bounds for $T_n^{(m)}(\mathcal{X})$, $m < M$. Next, we can show that the following term upper bounds the number of queries at fidelities less than M,

$$(M-1)n^\alpha \ +\ \sum_{m=1}^{M-1}\frac{1}{\tau}\left(\sqrt{2C_1 n_\Lambda \beta_{2n_\Lambda}\Psi_{2n_\Lambda}(\mathcal{H}_{\tau,n_\Lambda}^{(m)})} + \frac{\pi^2}{6}\right) \ +\ \sum_{m=1}^{M-1}\Omega_{\varepsilon_n}(\widehat{\mathcal{H}}^{(m)})\left(\frac{2\eta^2}{\gamma^2}\beta_n + 1\right).$$

Assume n_0 is large enough so that $n_0 \geq \max\{3,\ (2(\rho-\rho_0)\eta)^{-2/3}\}^{2/\alpha}$ and for all $n \geq n_0$, $n/2$ is larger than the above upper bound. We can find such an n_0 since the bound is $o(n)$. Therefore, for all $n \geq n_0$, $T_n^{(M)}(\mathcal{X}) > n/2$. Since our bounds hold with probability $> 1-\delta$ uniformly over n we can invert the above inequality to bound the number of plays N after capital Λ: $N \leq 2\Lambda/\lambda^{(M)}$ with probability $> 1-\delta$ if $\Lambda \geq \Lambda_0 = \lambda^{(M)}(n_0+1)$. The theorem follows with the observation $N \geq n_\Lambda \implies \mathcal{H}_{\tau,N}^{(m)} \subset \mathcal{H}_{\tau,n_\Lambda}^{(m)} \implies \Psi_N(\mathcal{H}_{\tau,N}^{(m)}) \leq \Psi_N(\mathcal{H}_{\tau,n_\Lambda}^{(m)}) \leq \Psi_{2n_\Lambda}(\mathcal{H}_{\tau,n_\Lambda}^{(m)})$. ■

C.0.1 Proof of Lemma 11

Since the posterior variance only decreases with more observations, we can upper bound $\kappa'(x,x)$ for any $x \in A$ by considering its posterior variance with only the s observations in A. Next the maximum variance within A occurs if we pick 2 points x_1, x_2 that are distance D apart and have all observations at x_1; then x_2 has the highest posterior variance. Therefore, we will bound $\kappa'(x,x)$ for any $x \in A$ with $\kappa(x_2,x_2)$ in the above scenario. Let $\kappa_0 = \kappa(x,x)$ and $\kappa(x,x') = \kappa_0\phi(\|x-x'\|_2)$, where $\phi(\cdot) \leq 1$ depends on the kernel. Denote the gram matrix in the scenario described above by $\Delta = \kappa_0 \mathbf{1}\mathbf{1}^\top + \eta^2 I$. Then using the Sherman-Morrison formula, the posterior variance (1) can be bounded via,

$$\kappa'(x,x) \ \leq\ \kappa'(x_2,x_2) \ =\ \kappa(x_2,x_2) - \left[\kappa(x_1,x_2)\mathbf{1}\right]^\top \Delta^{-1}\left[\kappa(x_1,x_2)\mathbf{1}\right]$$

$$
\begin{aligned}
&= \kappa_0 - \kappa_0\phi^2(D)\mathbf{1}^\top \left[\frac{\kappa_0}{\eta^2}I - \frac{\left(\frac{\kappa_0}{\eta^2}\right)^2 \mathbf{1}\mathbf{1}^\top}{1+\frac{\kappa_0}{\eta^2}s}\right]\mathbf{1} \;=\; \kappa_0 - \kappa_0\phi^2(D)\left(\frac{\kappa_0}{\eta^2}s - \frac{\left(\frac{\kappa_0}{\eta^2}\right)^2 s^2}{1+\frac{\kappa_0}{\eta^2}s}\right)\\
&= \kappa_0 - \kappa_0\phi^2(D)\frac{s}{\frac{\eta^2}{\kappa_0}+s} \;=\; \frac{1}{1+\frac{\eta^2}{\kappa_0 s}}\left(\kappa_0 - \kappa_0\phi^2(D) + \frac{\eta^2}{s}\right)\\
&\leq \kappa_0(1-\phi^2(D)) + \frac{\eta^2}{s}.
\end{aligned}
$$

For the SE kernel $\phi^2(D) = \exp\left(\frac{-D^2}{2h^2}\right)^2 = \exp\left(\frac{-D^2}{h^2}\right) \leq 1 - \frac{D^2}{h^2}$. Plugging this into the bound above retrieves the first result with $C_{SE} = \kappa_0/h^2$. For the Matérn kernel we use a Lipschtiz constant L_{Mat} of ϕ. Then $1-\phi^2(D) = (1-\phi(D))(1+\phi(D)) \leq 2(\phi(0)-\phi(D)) \leq 2L_{Mat}D$. We get the second result with $C_{Mat} = 2\kappa_0 L_{Mat}$. Since the SE kernel decays fast, we get a stronger result on its posterior variance which translates to a better bound in our theorems. ■

C.0.2 Proof of Lemma 12

The first part of the proof mimics the arguments in Lemmas 5.6, 5.7 of Srinivas et al. [28]. By assumption 8 and the union bound we can show,

$$
\mathbb{P}\left(\forall\, m \in \{1,\dots,M\},\ \forall\, i \in \{1,\dots,d\},\ \forall x \in \mathcal{X},\quad \left|\frac{\partial f^{(m)}(x)}{\partial x_i}\right| \;<\; b\log\left(\frac{6Mad}{\delta}\right)\right) \;\geq\; 1 - \frac{\delta}{6}.
$$

Now we construct a discretisation F_t of $\mathcal{X}$ of size $(\nu_t)^d$ such that we have for all $x \in \mathcal{X}$, $\|x - [x]_t\|_1 \leq rd/\nu_t$. Here $[x]_t$ is the closest point to x in the discretisation. (Note that this is different from the discretisation appearing in Theorem 10 even though we have used the same notation). By choosing $\nu_t = t^2 brd\sqrt{6Mad/\delta}$ and using the above we have

$$
\forall\, x \in \mathcal{X},\quad |f^{(m)}(x) - f^{(m)}([x]_t)| \;\leq\; b\log(6Mad/\delta)\|x - [x]_t\|_1 \;\leq\; 1/t^2 \tag{9}
$$

for all $f^{(m)}$'s with probability $> 1 - \delta/6$.

Noting that $\beta_t \geq 2\log(M|F_t|\pi^2 t^2/2\delta)$ for the given choice of ν_t we have the following with probability $> 1 - \delta/3$.

$$
\forall\, t \geq 1,\ \forall\, m \in \{1,\dots,M\},\ \forall\, a \in F_t,\quad |f^{(m)}(a) - \mu_{t-1}^{(m)}(a)| \;\leq\; \beta_t^{1/2}\sigma_{t-1}^{(m)}(a). \tag{10}
$$

The proof uses Gaussian concentration by only conditioning on $\mathcal{D}_t^{(m)}$. Note that instead of a fixed set over all t, we change the set at which we have confidence based on the discretisation. Similarly we can show that with probability $> 1 - \delta/3$ we also have confidence on the decisions $\mathbf{x}_t$ at all time steps. Precisely,

$$
\forall\, t \geq 1,\ \forall\, m \in \{1,\dots,M\},\quad |f^{(m)}(\mathbf{x}_t) - \mu_{t-1}^{(m)}(\mathbf{x}_t)| \;\leq\; \beta_t^{1/2}\sigma_{t-1}^{(m)}(\mathbf{x}_t). \tag{11}
$$

Using (9),(10) and (11) the following statements hold with probability $> 1 - 5\delta/6$. First, using assumption **A2** we can upper bound $f_\star$ by,

$$
f_\star \;\leq\; f^{(m)}(x_\star) + \zeta^{(m)} \;\leq\; f^{(m)}([x_\star]_t) + \zeta^{(m)} + \frac{1}{t^2} \;\leq\; \varphi_t^{(m)}([x_\star]_t) + \frac{1}{t^2}. \tag{12}
$$

Since the above holds for all m, we have $f_\star \leq \varphi_t([x_\star]_t) + 1/t^2$. Now, we bound $\Delta^{(m)}(\mathbf{x}_t)$.

$$
\begin{aligned}
\Delta^{(m)}(\mathbf{x}_t) = \;& f_\star - f^{(m)}(\mathbf{x}_t) - \zeta^{(m)} \;\leq\; \varphi_t([x_\star]_t) + \frac{1}{t^2} - f^{(m)}(\mathbf{x}_t) - \zeta^{(m)}\\
&\leq\; \varphi_t(\mathbf{x}_t) - f^{(m)}(\mathbf{x}_t) - \zeta^{(m)} + \frac{1}{t^2} \leq\; \varphi_t^{(m)}(\mathbf{x}_t) - \mu_{t-1}^{(M)}(\mathbf{x}_t) + \beta_t^{1/2}\sigma_{t-1}^{(M)}(\mathbf{x}_t) - \zeta^{(m)} + \frac{1}{t^2}\\
&\leq\; 2\beta_t^{1/2}\sigma_{t-1}^{(M)}(\mathbf{x}_t) + \frac{1}{t^2}.
\end{aligned}
$$

■

C.0.3 Proof of Lemma 14

First, we will invoke the same discretisation used in the proof of Lemma 12 via which we have $\varphi_t([x_\star]_t) \geq f_\star - 1/t^2$ (12). (Therefore, Lemma 14 holds only with probability $> 1 - \delta/6$, but this event has already been accounted for in Lemma 12.) Let $b_{i,n,t} = \operatorname{argmax}_{x \in A_{i,n}} \varphi_t(x)$ be the maximiser of the upper confidence bound in $A_{i,n}$ at time t. Now using the relaxation $\mathbf{x}_t \in A_{i,n} \implies \varphi_t(b_{i,n,t}) > \varphi_t([x_\star]_t) \implies \varphi_t^{(m)}(b_{i,n,t}) > f_\star - 1/t^2$ and proceeding,

$$
\begin{aligned}
\mathbb{P}(T_n^{(>m)}(A_{i,n}) > u) \;&\leq\; \mathbb{P}\big(\exists t : u+1 \leq t \leq n,\;\; \varphi_t^{(m)}(b_{i,n,t}) > f_\star - 1/t^2 \quad \wedge \quad \beta_t^{1/2}\sigma_{t-1}^{(m)}(b_{i,n,t}) < \gamma\big) \\
&\leq \sum_{t=u+1}^{n} \mathbb{P}\big(\mu_{t-1}^{(m)}(b_{i,n,t}) - f^{(m)}(b_{i,n,t}) > \Delta^{(m)}(b_{i,n,t}) - \beta_t^{1/2}\sigma_{t-1}^{(m)}(b_{i,n,t}) - 1/t^2 \quad \wedge \\
&\qquad\qquad \beta_t^{1/2}\sigma_{t-1}^{(m)}(b_{i,n,t}) < \gamma\big) \\
&\leq \sum_{t=u+1}^{n} \mathbb{P}\big(\mu_{t-1}^{(m)}(b_{i,n,t}) - f^{(m)}(b_{i,n,t}) > (\rho-1)\beta_t^{1/2}\sigma_{t-1}^{(m)}(b_{i,n,t}) - 1/t^2\big) \\
&\leq \sum_{t=u+1}^{n} \mathbb{P}_{Z\sim\mathcal{N}(0,1)}\left(Z > (\rho_0-1)\beta_t^{1/2}\right) \;\leq\; \sum_{t=u+1}^{n} \frac{1}{2}\exp\left(\frac{(\rho_0-1)^2}{2}\beta_t\right) \qquad (13) \\
&\leq \frac{1}{2}\left(\frac{\delta}{M\pi^2}\right)^{(\rho_0-1)^2} \sum_{t=u+1}^{n} t^{-(\rho_0-1)^2(2+2d)} \;\leq\; \frac{\delta}{M\pi^2}u^{-(\rho_0-1)^2(2+2d)+1} \;\leq\; \frac{\delta}{\pi^2}\frac{1}{u^{1+4/\alpha}}.
\end{aligned}
$$

In the second step we have rearranged the terms and used the definition of $\Delta^{(m)}(x)$. In the third step, as $A_{i,n} \subset \overline{\mathcal{J}}^{(m)}_{\max(\tau,\rho\gamma)}$, $\Delta^{(m)}(b_{i,n,t}) > \rho\gamma > \rho\beta_t^{1/2}\sigma_{t-1}^{(m)}(b_{i,n,t})$. In the fourth step we have used the following facts, $t > u \geq \max\{3,\, (2(\rho-\rho_0)\eta)^{-2/3}\}$, $M\pi^2/2\delta > 1$ and $\sigma_{t-1}^{(m)}(b_{i,n,t}) > \eta/\sqrt{t}$ to conclude,

$$
\begin{aligned}
(\rho-\rho_0)\frac{\eta\sqrt{4\log(t)}}{\sqrt{t}} > \frac{1}{t^2} \;&\implies\; (\rho-\rho_0)\cdot\sqrt{2\log\left(\frac{M\pi^2t^2}{2\delta}\right)}\cdot\frac{\eta}{\sqrt{t}} > \frac{1}{t^2} \\
&\implies\; (\rho-\rho_0)\beta_t^{1/2}\sigma_{t-1}^{(m)}(b_{i,n,t}) > \frac{1}{t^2}.
\end{aligned}
$$

In the seventh step of (13) we have bound the sum by an integral and used $\rho_0 \geq 2$ twice. Finally, the last step follows by $\rho_0 \geq 1 + \sqrt{(1+2/\alpha)/(1+d)}$ and noting $M \geq 1$. ■

D Addendum to Experiments

D.1 Other Baselines

For MF-NAIVE we limited the number of first fidelity evalutions to $\max\left(\frac{1}{2}\frac{\Lambda}{\lambda^{(1)}}, 500\right)$ where Λ was the total budget used in the experiment. The 500 limit was set to avoid unnecessary computation – for all of these problems, 500 queries are not required to find the maximum. While there are other methods for multi-fidelity optimisation (discussed under Related Work) none of them had made their code available nor were their methods straightforward to implement - this includes MF-SKO.

In addition to the baselines presented in the figures, we also compared our method to the following methods. The first two are single fidelity and the last two are mutlti-fidelity methods.

- The probability of improvement (PI) criterion for BO. We found that in general either GP-UCB or EI performed better.
- Querying uniformly at random at the highest fidelity and taking the maximum. On all problems this performed worse than other methods.
- A variant of MF-NAIVE where instead of GP-UCB we queried at the first fidelity uniformly at random. On some problems this did better than querying with GP-UCB, probably since unlike GP-UCB it wasn't stuck at the maximum of $f^{(1)}$. However, generally it performed worse.

- The multi-fidelity method from Forrester et al. [9] also based on GPs. We found that this method didn't perform as desired: in particular, it barely queried beyond the first fidelity.

A straightforward way to incorporate lower fidelity information to GP-UCB and EI is to query at lower fidelities and use them in learning the kernel κ by jointly maximising the marginal likelihood. While the idea seems natural, we got mixed results in practice. On some problems this improved the performance of all GP methods (including MF-GP-UCB), but on others all performed poorly. One explanation is that while lower fidelities approximate function values, they are not always best described by the same kernel. The results presented do not use lower fidelities to learn κ as it was more robust. For MF-GP-UCB, each $\kappa^{(m)}$ was learned independently using only the queries at fidelity m.

D.2 Description of Synthetic Experiments

The following are the descriptions of the synthetic functions used. The first three functions and their approximations were taken from [32].

Currin exponential function: The domain is $\mathcal{X} = [0,1]^2$. The second and first fidelity functions are,

$$f^{(2)}(x) = \left(1 - \exp\left(\frac{-1}{2x_2}\right)\right)\left(\frac{2300x_1^3 + 1900x_1^2 + 2092x_1 + 60}{100x_1^3 + 500x_1^2 + 4x_1 + 20}\right),$$
$$f^{(1)}(x) = \frac{1}{4}f^{(2)}(x_1 + 0.05, x_2 + 0.05) + \frac{1}{4}f^{(2)}(x_1 + 0.05, \max(0, x_2 - 0.05)) +$$
$$\frac{1}{4}f^{(2)}(x_1 - 0.05, x_2 + 0.05) + \frac{1}{4}f^{(2)}(x_1 - 0.05, \max(0, x_2 - 0.05)).$$

Park function: The domain is $\mathcal{X} = [0,1]^4$. The second and first fidelity functions are,

$$f^{(2)}(x) = \frac{x_1}{2}\left(\sqrt{1 + (x_2 + x_3^2)\frac{x_4}{x_1^2}} - 1\right) + (x_1 + 3x_4)\exp(1 + \sin(x_3)),$$
$$f^{(1)}(x) = \left(1 + \frac{\sin(x_1)}{10}\right)f^{(2)}(x) - 2x_1^2 + x_2^2 + x_3^2 + 0.5.$$

Borehole function: The second and first fidelity functions are,

$$f^{(2)}(x) = \frac{2\pi x_3(x_4 - x_6)}{\log(x_2/x_1)\left(1 + \frac{2x_7x_3}{\log(x_2/x_1)x_1^2x_8} + \frac{x_3}{x_5}\right)}, \qquad f^{(1)}(x) = \frac{5x_3(x_4 - x_6)}{\log(x_2/x_1)\left(1.5 + \frac{2x_7x_3}{\log(x_2/x_1)x_1^2x_8} + \frac{x_3}{x_5}\right)}.$$

The domain of the function is $[0.05, 0.15; 100, 50K; 63.07K, 115.6K; 990, 1110; 63.1, 116; 700, 820; 1120, 1680; 9855, 12045]$ but we first linear transform the variables to lie in $[0,1]^8$.

Hartmann-3D function: The M^{th} fidelity function is $f^{(M)}(x) = \sum_{i=1}^4 \alpha_i \exp\left(-\sum_{j=1}^3 A_{ij}(x_j - P_{ij})^2\right)$ where $A, P \in \mathbb{R}^{4\times 3}$ are fixed matrices given below and $\alpha = [1.0, 1.2, 3.0, 3.2]$. For the lower fidelities we use the same form except change α to $\alpha^{(m)} = \alpha + (M - m)\delta$ where $\delta = [0.01, -0.01, -0.1, 0.1]$ and $M = 3$. The domain is $\mathcal{X} = [0,1]^3$.

$$A = \begin{bmatrix} 3 & 10 & 30 \\ 0.1 & 10 & 35 \\ 3 & 10 & 30 \\ 0.1 & 10 & 35 \end{bmatrix}, \quad P = 10^{-4} \times \begin{bmatrix} 3689 & 1170 & 2673 \\ 4699 & 4387 & 7470 \\ 1091 & 8732 & 5547 \\ 381 & 5743 & 8828 \end{bmatrix}$$

Hartmann-6D function: The 6-D Hartmann takes the same form as above except $A, P \in \mathbb{R}^{4\times 6}$ are as given below. We use the same modification to obtain the lower fidelities using $M = 4$.

$$A = \begin{bmatrix} 10 & 3 & 17 & 3.5 & 1.7 & 8 \\ 0.05 & 10 & 17 & 0.1 & 8 & 14 \\ 3 & 3.5 & 1.7 & 10 & 17 & 8 \\ 17 & 8 & 0.05 & 10 & 0.1 & 14 \end{bmatrix}, \quad P = 10^{-4} \times \begin{bmatrix} 1312 & 1696 & 5569 & 124 & 8283 & 5886 \\ 2329 & 4135 & 8307 & 3736 & 1004 & 9991 \\ 2348 & 1451 & 3522 & 2883 & 3047 & 6650 \\ 4047 & 8828 & 8732 & 5743 & 1091 & 381 \end{bmatrix}$$

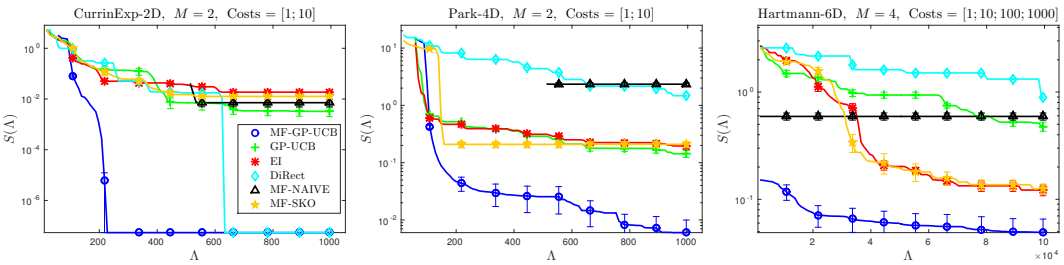

Figure 8: The simple regret $S(\Lambda)$ against the spent capitcal Λ on the synthetic functions. The title states the function, its dimensionality, the number of fidelities and the costs we used for each fidelity in the experiment. All curves barring DiRect (which is a deterministic), were produced by averaging over 20 experiments. The error bars indicate one standard error.

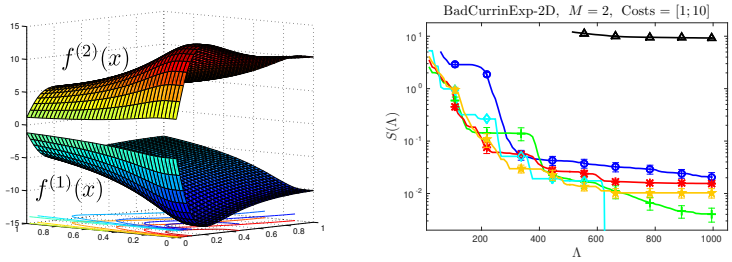

(a) (b)

Figure 9: (a) illustrates the functions used in the Bad Currin Exponential experiment where we took $f^{(1)} = -f^{(2)}$ and (b) shows the simple regret for this experiment. See caption under Fig. 8 for more details.

D.3 More Results on Synthetic Experiments

Figure 8 shows the simple regret $S(\Lambda)$ for the synthetic functions not presented in the main text.

It is natural to ask how MF-GP-UCB performs with bad approximations at lower fidelities. We found that our implementation with the heuristics suggested in Section 5 to be quite robust. We demonstrate this using the Currin exponential function, but using the negative of $f^{(2)}$ as the first fidelity approximation, i.e. $f^{(1)}(x) = -f^{(2)}(x)$. Figure 9 illustrates $f^{(1)}, f^{(2)}$ and gives the simple regret $S(\Lambda)$. Understandably, it loses to the single fidelity methods since the first fidelity queries are wasted and it spends some time at the second fidelity recovering from the bad approximation. However, it eventually is able to achieve low regret.

Finally, we present results on the cumulative regret for the synthetic functions in Figure 10.

Acknowledgements

We wish to thank Bharath Sriperumbudur for the helpful email discussions. This research is partly funded by DOE grant DESC0011114.

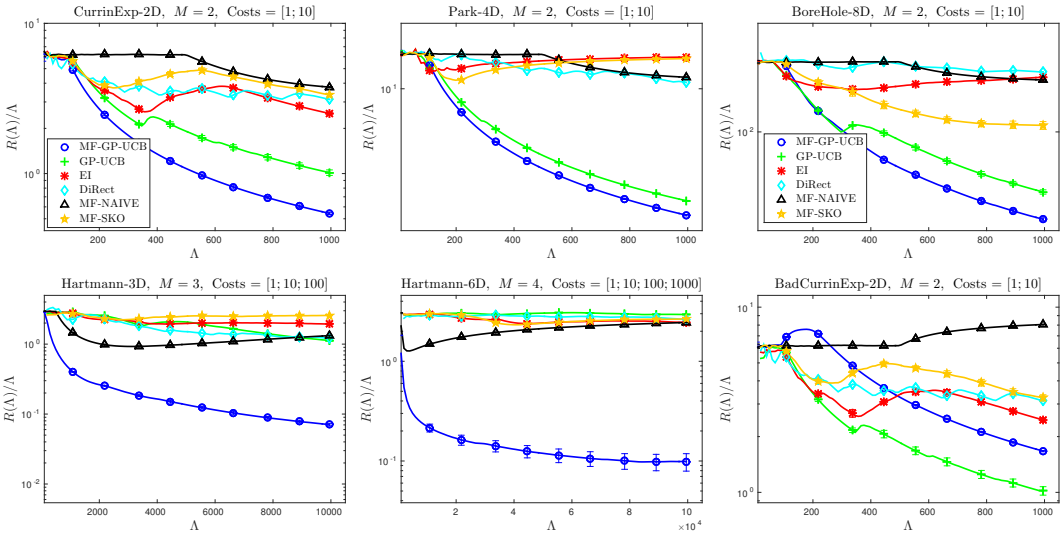

Figure 10: The cumulative regret $R(\Lambda)$ against the spent capitcal Λ on the synthetic functions. The title states the function, its dimensionality, the number of fidelities and the costs we used for each fidelity in the experiment. All curves barring DiRect (which is a deterministic), were produced by averaging over 20 experiments. The error bars indicate one standard error.